# Mixed Structural Models for Decision Making Under Uncertainty Using Stochastic System Simulation and Experimental Economic Methods: Application to Information Security Control Choice

HENLEY BUSINESS SCHOOL
THE UNIVERSITY OF READING

Doctor of Business Administration

**J. Jeffrey Curtis**

July 2016

# Declaration

I confirm that this is my own work and the use of material from other sources
has been properly and fully acknowledged.

________________________________________ Date: ________________________
J. Jeffrey Curtis

# Certificate of readiness to be included in library

I grant powers of discretion to the University Librarian to allow this thesis to be copied in whole or in part without further reference to me. This permission covers only single copies made for study purposes, subject to normal conditions of acknowledgement.

# Acknowledgements

I greatly acknowledge the contribution of my supervisor Professor Walid Hejazi. His guidance, knowledge and support made this thesis possible and he has been a constant source of encouragement throughout this work. I thank Dr. Hejazi for making available his funding resources to support the decision making experiments that are the foundation of this thesis. He remains my greatest academic supporter.

I sincerely acknowledge the contribution of Professor Glenn Harrison and his colleagues whose scholarship in the field of risk management and experimental economic methods has been my preeminent motivation throughout my research. I am very grateful for his guidance provided to me through personal correspondence at key junctures of this work. It has been an honor to attempt to reflect his contributions in this thesis.

I would like to thank Gary Mckenzie of Integrity Assurance and Keith Jonah of Trusted by Design who provided expert review and guidance regarding the logical information security model.

I would like to thank Dr. Nancy Gallini who first inspired me to think economically about business problems 35 years ago – had I paid more attention at the time, I might have been finished much earlier.

Finally, I would like to dedicate this thesis to wife, Fanzeela Ali, who has been very patient with this ambitious personal undertaking over many years. She always ensured that I consumed more than just coffee.

# Abstract

This research is concerned with whether and to what extent information security managers may be biased in their evaluation of and decision making over the quantifiable risks posed by information management systems where the circumstances may be characterized by uncertainty in both the risk inputs (e.g. system threat and vulnerability factors) and outcomes (actual efficacy of the selected security controls and the resulting system performance and associated business impacts). Although 'quantified security' and any associated risk management remains problematic from both a theoretical and empirical perspective (Anderson 2001; Verendel 2009; Appari 2010), professional practitioners in the field of information security continue to advocate the consideration of quantitative models for risk analysis and management wherever possible because those models permit a reliable *economic* determination of optimal operational control decisions (Littlewood, Brocklehurst et al. 1993; Nicol, Sanders et al. 2004; Anderson and Moore 2006; Beautement, Coles et al. 2009; Anderson 2010; Beresnevichiene, Pym et al. 2010; Wolter and Reinecke 2010; Li, Parker et al. 2011)[1]. The main contribution of this thesis is to bring current quantitative economic methods and experimental choice models to the field of information security risk management to examine the potential for biased decision making by security practitioners, under conditions where information may be relatively objective or subjective and to demonstrate the potential for informing decision makers about these biases when making control decisions in a security context.

No single quantitative security approach appears to have formally incorporated three key features of the security risk management problem addressed in this research: 1) the inherently stochastic nature of the information system inputs and outputs which contribute directly to decisional uncertainty (Conrad 2005; Wang, Chaudhury et al. 2008; Winkelvos, Rudolph et al. 2011); 2) the *endogenous* estimation of a decision maker's risk attitude using models which otherwise typically assume risk neutrality or an inherent degree of risk aversion (Danielsson 2002; Harrison, Johnson et al. 2003); and 3) the application of structural modelling which allows for the possible combination and weighting between multiple latent models of choice (Harrison and Rutström 2009). The identification, decomposition and tractability of these decisional factors is of crucial importance to understanding the economic trade-offs inherent in security control choice under conditions of both risk and uncertainty, particularly where established psychological decisional biases such as ambiguity aversion (Ellsberg 1961) or loss aversion (Kahneman and Tversky 1984) may be assumed to be endemic to, if not magnified by, the institutional setting in which these decisions take place. Minimally, risk averse managers may simply be overspending on controls, over-compensating for anticipated losses that do not actually occur with the frequency or impact they imagine. On the other hand, risk-seeking managers, where they may exist (practitioners call them 'cowboys' – they

[1] For practitioner's guidance, see the "Octave" methodology as developed by Carnegie Mellon's Software Engineering Institute: http://www.cert.org/octave/ and the U.S. National Institute of Standards and Technology Special Publication 800-30 "Risk Management Guide for Information Technology Systems" http://csrc.nist.gov/publications/nistpubs/800-30/sp800-30.pdf . Although these and other approaches allow for both qualitative and quantitative risk assessment, most methods advocate the use of risk assessment and control selection based on quantitative measurement wherever possible as a means of objectively evaluating, communicating and managing information security risks.

are a familiar player in equally risky financial markets) may be simply gambling against ultimately losing odds, putting the entire firm at risk of potentially catastrophic security losses. Identifying and correcting for these scenarios would seem to be increasingly important for now universally networked business computing infrastructures.

From a research design perspective, the field of *behavioural economics* has made significant and recent contributions to the empirical evaluation of psychological theories of decision making under uncertainty (Andersen, Harrison et al. 2007) and provides salient examples of *lab experiments* which can be used to elicit and isolate a range of latent decision-making behaviours for choice under risk and uncertainty within relatively controlled conditions versus those which might be obtainable in the field (Harrison and Rutström 2008). My research builds on recent work in the domain of information security control choice by 1) undertaking a series of lab experiments incorporating a stochastic model of a simulated information management system at risk which supports the generation of observational data derived from a range of security control choice decisions under both risk and uncertainty (Baldwin, Beres et al. 2011); and 2) modeling the resulting decisional biases using structural models of choice under risk and uncertainty (El-Gamal and Grether 1995; Harrison and Rutström 2009; Keane 2010). The research contribution consists of the novel integration of a model of stochastic system risk and domain relevant structural utility modeling using a mixed model specification for estimation of the latent decision making behaviour. It is anticipated that the research results can be applied to the real world problem of ‘tuning’ quantitative information security risk management models to the decisional biases and characteristics of the decision maker (Abdellaoui and Munier 1998).

# Table of Contents

# List of Figures

# List of Tables

# List of Abbreviations

| | |
|---|---|
| **ALE** | Annual Loss Expectancy |
| **CARA** | Constant Absolute Risk Aversion |
| **CIA** | Confidentiality, Integrity and Availability |
| **CPT** | Cumulative Prospect Theory |
| | Conditional Probability Table |
| **CRRA** | Constant Relative Risk Aversion |
| **CVaR** | Conditional Value at Risk |
| **CVE** | Common Vulnerabilities and Exposures |
| **CySeMol** | Cyber Security Model |
| **DES** | Discrete Event Simulation |
| **EU** | Expected Utility |
| **EUT** | Expected Utility Theory |
| **GPD** | Generalized Pareto Distribution |
| **LRT** | Likelihood-Ratio test |
| **MPL** | Multiple Price List |
| **P2CySeMol** | Predictive, Probabilistic Cyber Security Model |
| **PRA** | Probability Risk Assessment |
| **PRM** | Probabilistic Relational Model |
| **PT** | Prospect Theory |
| **QSR** | Quadratic Scoring Rule |
| **RAPSA** | Risk Analysis and Probabilistic Survivability Assessment |
| **RDU** | Rank Dependent Utility |
| **RRA** | Relative Risk Aversion |
| **SCADA** | Supervisory Control and Data Acquisition |
| **SEU** | Subjective Expected Utility |
| **SEUT** | Subjective Expected Utility Theory |
| **SQL** | Structured Query Language |

| | |
|---|---|
| **SSA** | Survivable Systems Analysis |
| **VaR** | Value at Risk |
| **XSS** | Cross Site Scripting |

*Uncertainty must be taken in a sense radically distinct from the familiar notion of risk, from which it has never been properly separated.... The essential fact is that 'risk' means in some cases a quantity susceptible of measurement, while at other times it is something distinctly not of this character; and there are far-reaching and crucial differences in the bearings of the phenomena depending on which of the two is really present and operating... It will appear that a measurable uncertainty, or 'risk' proper, as we shall use the term, is so far different from an unmeasurable one that it is not in effect an uncertainty at all.*

**Frank H. Knight, *Risk, Uncertainty, and Profit (1921)***

*In real life you do not know the odds; you need to discover them, and the sources of uncertainty are not defined. Economists, who do not consider what was found by non-economists worthwhile, draw an artificial distinction between Knightian risk (which you can compute) and Knightian uncertainty (which you cannot compute), after one Frank Knight, who rediscovered the notion of unknown uncertainty and did a lot of thinking but perhaps never took risks, or perhaps lived in the vicinity of a casino. Had he taken financial or economic risk he would have realized that these "computable" risks are largely absent from real life! They are laboratory contraptions.*

**“The uncertainty of the nerd” Nassim Nicholas Taleb, *The Black Swan (2007)***

*Even apart from the instability due to speculation, there is the instability due to the characteristic of human nature that a large proportion of our positive activities depend on spontaneous optimism rather than mathematical expectations, whether moral or hedonistic or economic. Most, probably, of our decisions to do something positive, the full consequences of which will be drawn out over many days to come, can only be taken as the result of animal spirits - a spontaneous urge to action rather than inaction, and not as the outcome of a weighted average of quantitative benefits multiplied by quantitative probabilities.*

***Keynes J. M., 1936,***
***The General Theory of Employment, Interest and Money, McMillan University Press.***

*Noise makes financial markets possible, but also makes them imperfect.'*

***Black F., 1986,”Noise,” Journal of Finance, 41, 529-543.***

*If I don't know I don't know, I think I know. If I don't know I know I know, I think I don't know.*

***R.D. Laing, 1927 - 1989***

*Chance discovers direction, impulse discovers intention, design discovers structure.*

***Seamus Heaney***

*There is no general principle that prevents the creation of an economic theory based on other hypotheses than that of rationality*

***Kenneth Arrow (1986)***

*It is quite surprising and disappointing to me that almost 40 years after the establishment of the concept of risk aversion by Pratt and Arrow, our profession has not yet been able to attain a consensus about the measurement of risk aversion. Without such a consensus, there is no hope to quantify optimal portfolios, efficient public risk prevention policies, optimal insurance deductibles, and so on. It is vital that we put more effort on research aimed at refining our knowledge about risk aversion. For unclear reasons, this line of research is not in fashion these days, and it is a shame.*

***Gollier, C. (2001). The Economics of Risk and Time. Cambridge, MA: MIT Press.***

*Essentially, all models are wrong, but some are useful*

***George Box and Norman Draper (1987)***

# 1 – Introduction: Research Problem, Hypotheses and Contribution

Institutions and their information technology service providers are increasingly faced with a wide range of ethical, legal and contractual obligations to protect the *security* of the information management systems and the information which they collect and then use and disclose to customers, other service providers, governments and other 3rd parties during the provision of services. This research builds on the academic and practitioner literature concerning the 'economics of information security' and investigates security control selection in the context of information security risk management objectives under specified conditions of risk and uncertainty. Specifically, I consider that any 'optimal' selection of policy, procedural and technical security controls intended to prevent or compensate for loss of the confidentiality, integrity or availability of systems or information requires decisions by information managers under circumstances which are characterized by the inherent uncertainty of the *likelihood* and *impact* of prospective business outcome scenarios based, in part, on the *perceived* ex ante and actual ex post efficacy of the selected controls. These circumstances may be generalized as complex, 'multi-stakeholder, multi-objective, multi-attribute' decision problems involving control decisions resulting in endogenous risks to the information assets involved and can be modeled using stochastic system simulations coupled with multi-criteria decisional utility functions (Keeney 1982; Keeney 1993; Baldwin, Mont et al. 2009; Sen 2010) from which decisional biases can be estimated using experimental economics laboratory approaches for eliciting risk attitude and preference under uncertainty (Harrison and Rutström 2008). It is hypothesized that control decisions will, for example, depend on the *subjective* perception of risk involving the potential loss of system performance and versus the expected lifecycle costs of control. Decisions made in conditions of both known risk factors and those involving uncertain risk have been demonstrated, within both the psychological and economics literature, to be subject to decisional biases including economic agency effects, bounded rationality, self-interest, heuristics and risk framing. I propose that structural modeling of these decisional biases can improve our understanding of the *descriptive* and *normative* theories of information security control decisions where observational data is based on 'behavioural economics' experiments which generate quantitative data under conditions which permit a broad range of testing for decisional biases under uncertainty. This research is therefore intended to better describe and potentially allow 'correction' for bias in management control decisions and contributes to the integration of experimental design and empirical research into quantitative institutional risk management approaches for information privacy and security control.

### Privacy vs. Security Control

As a primary dichotomy for this research, I make a fundamental distinction between information *privacy* versus *security* control selection, both of which motivate this research and are important dimensions of modern institutional information management goals and objectives. The information privacy literature is essentially concerned with the legal, contractual and practical obligations and capabilities of an individual (or their agent) to control personal information, whereas information security literature is more generally concerned with institutional control over all of the institutional entity's information systems and data often including, but not limited to, the personal information of its clients and employees which are considered to be held in 'trust' by the individuals to whom the data pertains (Gordon and Loeb 2006). Privacy, as an individual human right with corresponding legal requirements, is recognized in an increasing number of regional national and international jurisdictions and is, in this respect, antecedent to the resulting contractual *confidentiality* obligations between the data owner and the institution which collects, uses or discloses the data, and forms but a subset of the overall institutional motivation for the selection and management of security controls generally. Information security management, in this respect, is fundamentally concerned with maintaining not only the *confidentiality* (appropriate access to and use of systems and information), but also the *integrity* (accuracy, reliability, verifiability, completeness, computability) and *availability* (continuous accessibility, recoverability) of data and data management systems using policy, procedural and technical controls which may be considered appropriate in the circumstances under which systems and information operate (Anderson 2001; Gordon and Loeb 2002; Gordon and Richardson 2004; Appari 2010)[2].

Consistent with this view, information security research is primarily concerned about the resulting vector of control choices made by data and system *custodians* who address confidentiality, integrity and availability goals. The perspective shift from individual privacy (or even personal security) control selection to institutional security assurance allows us to take into account the goals, objectives and decision making activities of the institutional data custodian or manager whose risk attitude regarding prospective system or data loss and the associated control choices to prevent or recover from loss may be aligned with, but is also inherently distinct from, that of the individual seeking data protection over their own personal information and which may be further biased by other business objectives within the institution (Berg 2005). Furthermore, where the control decision may be based on multiple individual decision makers, we might expect latent decisional biases to be compounded, offset or perhaps subject to strategic behaviour between institutional decision makers or between institutions.

**Information Security and Security Risk Management**

Practitioner selection of information security controls in the field is typically undertaken as a response to the perceived risk to business value anticipated resulting from the continuous operation of information management systems that are intended to create value for the business. Networked computing systems may be considered inherently risky from a security perspective since the absolute assured operation of the systems can be generally expected to decline from 'error free' operation as business dependency and

[2] For an introduction to information security theory, see: Deborah Russell and G. T. Gangemi, Sr., *Computer Security Basics* (New York: Thunder Mountain Press,1994) or John D. Howard, *An Analysis of Security Incidents on the Internet 1989–1995,* Ph.D. thesis, Department of Engineering and Public Policy, Carnegie Mellon University, April 7, 1997.

system complexity increases, errors which directly or indirectly impact the security attributes of the system (Avižienis, Laprie et al. 2004; Nicol, Sanders et al. 2004). Security risk in terms of declining prospective confidentiality, integrity or availability can therefore be characterized as the divergence from an expected level of operation (even if not perfect operation) and the associated degradation of business value attributed to the otherwise successful operation the systems. Notionally, management is expecting some overall level of successful systems operation in order to support business needs. The divergence from expectation is what I will consider to be understood as *risk*, regardless of whether that risk can be measured or not. Stated another way, if there is no anticipated divergence in the actual performance or business outcome, there is no risk to manage. Operationally, practitioners accept that that the system will, within a defined period of time, likely operate either better than or worse than 'expected' and it is this inherent variability that I will characterize as 'security risk' (Rockafellar and Uryasev 2013). This research attempts to both model this risk, which is of growing interest in itself for security practitioners, and to consider how managers react to this risk in practice.

'Risk management' for information security, is therefore a growing field of interest for business operators, system providers and customers since it represents a fiduciary duty within the institution to address the prospective liability of operating inherently risky systems to generate business value. Security risk management is typically based on some form of 'threat and risk assessment' approach which indicates that, for an assumed *risk neutral* decision maker, controls anticipated to be effective prospective data or system losses should be employed up to the value of the expected loss being protected against, whether that value is calculated quantitatively or qualitatively, and whether the loss is even stated explicitly since most operators accept that the resulting business operations will not turn out exactly as 'expected'.

It is axiomatic that the cost of a control should not exceed the expected value of the loss it is intended to protect. In practice, the prospective nature of both the control effectiveness and any associated attributable resulting loss is what essentially underpins the commonly understood notion of an 'expected' value and, for the purposes of this research, is what makes the decisional considerations interesting from both a psychological and economic perspective[3]. I will refer to the inputs to the loss value calculation as 'risk factors' (Fama and French 1993) consisting of a vector of business environmental and usage conditions and system *states* affecting the resulting probability and magnitude of a prospective business loss. In the prospective case, the value of an *expected* loss is understood as the probability of a *threat* acting on a system, process or control *vulnerability likelihood*, multiplied by the anticipated *impact* of the resulting business loss translated into a monetary measure should the threat actually be successful. It should be noted

[3] In practice, determining ex post, 'root cause', attribution between failed control(s) and business impacts is itself problematic in most cases where multiple controls are involved and where alignment between system monitoring and associated business operations may be weak. This will be discussed in the [next section] as a component of the decision maker's inherent uncertainty over ex ante control selection and also presents challenges for lab experiments based on models which, while inherently causal in construction (even if stochastic or non-deterministic) may not be entirely accepted by lab participants as valid representations of connections between security controls and modelled losses. I will also return to this issue in Section [XX} in a discussion of Smith's valuable cautions regarding divergence between theory and experiment (Harrison 2010; Smith 2010)

that, in this common information security risk model[4], the absolute absence of one or more of these factors means that there is no meaningful qualitative or quantitative prospective loss to be anticipated. This is important from the perspective of risk modelling and the testing of decisions using risk models since particular controls involving absolute avoidance of certain system or business activity risk factors may be effectively employed in some circumstances which can effectively eliminate the risk. Disabling certain types of system networking, media use, or not engaging with third party suppliers, for example, are examples which essentially eliminate the associated risks (since the specific threat is not expected to occur if the activity generating the threat is not undertaken), although in practice this is recognized as either reducing business value (since some value generating activity may not be undertaken as a result) or leading to substitute activities which involve alternate risks.

It is therefore fundamental to this research to recognize that the security *risk management* process – the identification of risk factors, qualitative or quantitative risk measurement, risk treatment via the *ex-ante* selection of controls, and ultimately risk acceptance by stakeholders - is fundamentally subject to inherent uncertainty, meaning that the actual stream of operational events and the performance of the selected controls should be expected to vary from the 'average' or typical outcome at the point of control selection due to a range of exogenous and endogenous factors which may occur over a specified future time period, and for which the *probability distribution* of occurrence may be essentially uncertain. For example, the threat to the data may be greater or less than expected (e.g. from a threat perspective, no one was actually interested in stealing the data, or did not attempt to do so with the frequency or intensity anticipated; or from a physical protections perspective, no fire actually broke out in the computer server room; etc.), or because the vulnerability of the data management system turns out to be greater or lesser than expected (i.e. the incumbent system or process to which the controls were applied is actually more or less vulnerable to misuse than estimated due to the actual efficacy of the control; etc.), or because the impact of any actual loss (as measured in, for example, dollars or operational time lost, embarrassment to the institution, inconvenience suffered by the customers, etc.) turns out to be greater or less than what was expected at the time the control was selected (Conrad 2005).

As a matter of pure definition and also the implication for management's consideration of risk factors, I will discuss the distinction between probability and likelihood at several junctions in the paper - for now I consider the term 'probability' to be descriptive for purposes of defining risk in the management context I am primarily concerned with. From management's perspective, I also note that risk calculations are typically undertaken as, and therefore only yield, point estimates of 'expected loss' and, in my experience, risk is not well understood by practitioners to be represented using the 2nd or higher moments of probability distributions, or as a range of random values generated from the joint probabilities of the threats, vulnerabilities and impacts or, even more generally, as a series of prospective, iterative outcomes. The

[4] http://www.nist.gov/cyberframework/

immediate direct analog between daily system performance and attributable business losses to a time series of financial market returns, for example, where both are generated by combinations of known and unknown stochastic processes (even if the processes can be modelled) was a key insight and motivator for this research. From that perspective, it is my observation that few dominant practitioner security frameworks[5] take into consideration the inherent *stochastic* nature of the threat, vulnerability or impact risk factors or the resulting documented potential for bias in a practitioner decision maker attempting to prospectively select controls under conditions characterized by either known or uncertain risk factor probabilities. My overall hypothesis is therefore that due to the inherent prospective uncertainty of the risk factors in practice, information security controls may in fact be selected in a psychologically biased manner measured from some normative comparator and so careful consideration should therefore be given to the behavioural aspects of control selection under these conditions[6].

The uncertainty of the risk factors, within what is otherwise a valid information security risk management approach (Appari 2010), is what makes the task of control selection a process of decision making under conditions of uncertainty. Misperception of the moments of the frequency distribution of the incidence of any or all of the risk factors, or particularly of the joint probability distribution of the factors is likely to result in outcomes which make the investment in controls a risky bet. The analogous 'asset allocation' problem, here applied to security control 'portfolio selection' has strong analogs in behavioural finance experiments in connection with equity premia and portfolio optimization problems (Benartzi and Thaler 1995) and the solution has characteristics similar to 'dynamic asset allocation' models which similarly incorporate 'investor' risk attitudes under conditions of market uncertainty (Dantzig and Infanger 1993; Infanger 2006). I am therefore interested in the extent to which management 'bets' on controls over time (where there are, practically, many of these bets involved, often in sequence, often informed sequentially as new controls become required and old controls perform or fail in the overall information risk management process) can be modeled as a series of economic trade-offs made under conditions of uncertainty, and whether the resulting descriptive theory of decision making can be modeled and the models estimated based on choice data derived from experimental lab procedures which reveal both the type and degree of decisional bias described in both the psychological and emergent behavioural economics literature. The overall research question addressed and the resulting contribution of this research is: *Are information security managers biased and to what extent when making security control decisions under risk and uncertainty?*

[5] See, for example, http://www.cert.org/octave/

[6] Some alternative economic approaches to information security control selection are notable. See, for example, (Gordon and Loeb 2002) who argue that there is a natural threshold level of optimal control selection regardless of the system risk profile. While they do not take into consideration a 'behavioural' approach to control selection, they do take an 'economic approach' to security outcome value maximization and argue that managers should generally concentrate on data and systems with midrange vulnerabilities, concluding that "…to maximize the expected benefit from investment to protect information, a firm should spend only a small fraction of the expected loss due to a security breach."

**Considered Hypotheses**

Following the approaches of Antoniou, Andersen and Sen (Andersen, Harrison et al. 2006; Andersen, Harrison et al. 2007; Antoniou, Harrison et al. 2010; Sen 2010), several relevant hypotheses regarding security manager's decision making biases have been considered in order to develop the associated lab experiments proposed in this research:

1) Risk attitude and subjective estimates of security risk are equal under conditions of exogenous and endogenous risk (Sen 2010)
2) Risk aversion is not equal over prospective operational gains versus security losses. (Sen 2010)
3) Managers exhibit loss aversion and probability weighting over low vs. high probability security risks. (Andersen, Harrison et al. 2006)
4) The strength and weight of quantitative estimates of security risk influence a manager's subjective expected utility over control decisions. (Antoniou, Harrison et al. 2010)
5) Managers are 'ambiguity averse' over quantitative estimates of security risk. (Andersen, Fountain et al. 2009)
6) Prior outcomes affect a manager's current security decisions. (Andersen, Harrison et al. 2006)
7) A manager's experience of significant security losses increases their risk aversion. (Sen 2010)
8) A manager's subjective prior estimate of risk declines over time. (Sen 2010)
9) Managers integrate control choices over accumulated security gains or losses. (Andersen, Harrison et al. 2006)
10) A manager's reference point for control choices is not the marginal value of the current choice. (Andersen, Harrison et al. 2006)
11) Loss averse managers tend to be more successful over time than non-loss averse managers. (Andersen, Harrison et al. 2006)
12) Unitary groups of security decision makers exhibit less risk aversion than single decision makers. (Sen 2010)
13) (A): Unitary groups of security decision makers operating without coordination over time (free-riding) conditions exhibit more risk aversion than single decision makers.
13) (B): Quasi-unitary groups of security decision makers operating without coordination over time (free-riding possible) conditions exhibit less risk aversion than consensus decision makers.
14) Symplectic groups of security decision makers (i.e. with competing goals) exhibit more risk aversion than single decision makers.(Gordon, Loeb et al. 2008)
15) Security managers behave according to RDU power utility for reported vs. recovered beliefs over continuous security loss distributions.
16) Individuals who prefer lower risk will choose the more efficient risk management method (precaution vs. insurance) to accomplish this goal.

Although certain hypotheses concerning risk attitudes can be potentially tested across different experiments, certain hypotheses are better tested using specific experiments which control for particular aspects of the risk and uncertainty associated with stochastic systems. The goal for this business management research was to test a reasonable range of behavioural decision making hypotheses while also replicating a range of experiments that would be able to directly incorporate IT system simulation and the practitioner-relevant context of business losses attributable to security incidents. The resulting range of hypotheses, experiments and treatments have achieved a balance across these objectives which showcase the potential for this experimental approach to decision making under risk and uncertainty in the information security domain. The hypotheses specifically undertaken for this research are reviewed in detail in **Section 7** and have been assigned to individual lab experiments detailed in **Section 8.** The results of all experiments are reported in **Section 9**.

# 2 – Literature Review

This research involves a microeconomic approach to evaluating decision making over security controls which permits the modeling and estimation of decisional biases under conditions of uncertainty (Kunreuther and Heal 2003; Grossklags, Christin et al. 2008; Grossklags, Christin et al. 2008; Camp 2009; Grossklags, Johnson et al. 2010; Grossklags, Johnson et al. 2010; Johnson, Grossklags et al. 2010). Current legal and contractual information protection regimes require institutional administrators to interpret privacy and security requirements and to thereafter deploy controls over information assets which would be considered 'appropriate' from the normative perspective of these regimes within the specific operational circumstances of the firm controlling the information. The 'uncertain' circumstances include, but are not limited to: uncertainty with respect to the normative intent of the law; subjective interpretation of legal and contractual control requirements and compliance obligations; and an uncertain set of future system characteristics, internal and external system threats, and control performance results, most if not all of which are inherently stochastic. This uncertainty in operational results, combined with inherent decision maker risk attitudes and preferences is hypothesized to result in 'biased' decisions over controls otherwise designed to reduce attributed business losses resulting from uncertain system performance. My focus for investigation in this research is to determine how administrators can make better information security control decisions in the presence of these uncertainties and given the inherent decisional biases that may be present. The following literature review informs the construction of a formal, researchable research question and the resulting research methodology, data collection, compilation and analysis approach.

### Individual vs. Institutional Perspectives on Information Privacy and Security Assurance

The modern legal rights of a person to control information about themselves arose out of the liberal democratic recognition of the liberty and dignity of the individual and were specifically interpreted within an American legal context at the end of the 19th century to include a conception of the 'inviolate personality' (Warren and Brandeis 1890), including the privacy of the physical person, sanctity of the home, and the right to both personal and informational concealment and secrecy (Westin 1967; Richards and Solove 2007). Comparatively, English and European law has traditionally recognized (and continues to emphasize) a distinct individual right to control information about one's private affairs or circumstances under conditions of disclosure to others which have been 'undertaken in confidence', thereby leading to expectations of 'confidential relations' and which therefore establish a corresponding *fiduciary* obligation or 'trust within relationships' between the individual to whom the information pertains and the receiving party in respect of the information exchanged (Schwartz 1994; Richards 2006; Richards and Solove 2007). Since the beginning of the 20th century, Anglo-American jurisprudence and the resulting regulatory and contractual frameworks for privacy protections have diverged substantially over the inherent differences therefore arising between *decisional* privacy rights (the right to control or make decisions affecting the self) and *informational* privacy rights (the right to control personal information). The intersection of these approaches continues to inform and shape scholarly debate regarding the objective and subjective valuations of personal information and the extent to which personal information ought to and can be practically controlled by the individual to whom the information pertains and, correspondingly, by those who are expected to use and disclose it under conditions of trust or confidence (Bennett 2000; Schwartz 2000; Richards and Solove 2007).

Driven by large scale commercial and non-commercial collection, use and disclosure of personal information by individuals, businesses and governments using increasingly networked electronic information management systems, the issue of how to enforce *personal* information rights from an ethical, legal, contractual and technological control perspective has resulted in a large body of information privacy and security literature developed by legal scholars, economists, sociologists, psychologists, institutional theorists, and information systems and security theorists, architects and technologists. My literature review references, as a starting point, the 'privacy economics' research being undertaken, for example, by Allesandro Acquisti at Carnegie Mellon[7] and the 'security economics' research undertaken Ross Anderson at the University of Cambridge[8] beginning in the late 1990s. Seminal papers referenced by these authors have been reviewed for connections across the fields of law, economics, economics and law, privacy, security, decision theory, and information and technology management. Given the diversity of both the social and technical sciences represented across my topic of interest, I have purposefully explored both the theoretical and empirical literature for all relevant topics since the emergent *microeconomic* models under examination explicitly attempt to account for institutional, psychological and decisional biases in the context of legal, contractual and operational information management risks.

### Opportunities for Information Privacy and Security Research in Healthcare Institutions

My professional background is as the chief privacy officer in a large academic health sciences (research and teaching) hospital in Canada. Healthcare delivery markets are of increasing interest to privacy and security stakeholders since they manifest, and often magnify, many of the information risk management considerations which may be expected to motivate privacy concerns in the first place (Schwartz 1997): 1) the presence of subjectively sensitive, personally identifying information which is inherently required for the provision of healthcare services and the control of which impacts the informational privacy rights of individuals; 2) the increasing scale and scope of demand for healthcare services and associated service delivery infrastructures which are dependent on aggregated personal health information data banks involving millions of records; 3) the increasing demand for *secondary use* of personal information by governments and private institutions for research and planning purposes which, while often motivated from a public health or system planning perspective, is typically unrelated to the original purposes or contractual conditions under which the information was originally supplied by the patient to the healthcare provider (Bregman-Eschet 2006; Bellamy 2010); 4) increasingly decentralized but electronically networked health

[7] http://www.heinz.cmu.edu/~acquisti/economics-privacy.htm
[8] http://www.cl.cam.ac.uk/~rja14/econsec.html

information management systems and databanks which increase the inherent risk of the information management system as a whole, and therefore of privacy breaches resulting from loss of 'security' (loss of data or system confidentiality, integrity, accuracy or computability, or availability) (Appari 2006; Johnson 2009; Appari 2010); and 5) the advent of specific statutory, regulatory, contractual and institutional policy regimes which seek to balance the privacy individuals with the flow of information necessary to support an optimally functioning healthcare delivery system and the realization of secondary social benefits such as public health tracking and chronic disease management (Anderson 1996; Schwartz 1997; Sage 1999; Sage 2008; Miller and Tucker 2009; Bansal, Zahedi et al. 2010; Hall 2010).

In terms of the amount and sensitivity of personal information involved, and the inherent security risks presented to the information by increasingly networked electronic medical records, healthcare information systems therefore present a valuable opportunity for research into how both individuals and institutions make decisions over the control of the information. Indeed, as a specific area for information security research, healthcare appears to be relatively under-examined to date versus traditional security research domains. As Appari notes: "Despite [this] growing stream of research on information security, very limited research has focused on studying information security risks in the healthcare sector, which is heavily regulated and calls upon business models different from other industries…Surprisingly, very little attention has been given to the economics of information security risks…" (Appari 2010).

### The Economics of Privacy

From both a positive and normative perspective, the economics literature has been particularly instrumental in framing information privacy and security issues from the standpoint of a *resource allocation problem*, and offers valuable insights and methodologies for analyses of the problem of personal information protection and the associated *prescriptive* legal, policy, procedural and technological controls which would be appropriate under certain market conditions, all under the heading of the **'economics of privacy'** (Posner 1978; Stigler 1980; Posner 1981; Laudon 1993; Varian 1996; Bennett 2000; Acquisti 2004; Hermalin and Katz 2004; Hui and Png 2005; Zhan and Rajamani 2008; Acquisti 2010; Appari 2010).

According to Allesandro Acquisti:

> *"Privacy economics deals with informational trade-offs: it tries to understand and sometimes quantify the costs and benefits that data subjects (as well as potential data holders) bear or enjoy when their personal information is either protected or shared. Privacy economics also tries to understand how to use market mechanisms, technology or policy to achieve a desirable balance between information revelation and protection, with satisfaction shared among the individual, organization and society as a whole".* (Acquisti 2009)

In seeking a 'balance between information revelation and protection', the recognition here is that that data holders play a perhaps equally important part in both making decisions over and experiencing the results of controls, and that 'satisfaction' (possibly understood as 'utility' vs. 'value') is somehow divisible between the decision makers. The claim is that microeconomic approaches to privacy assurance (and many existing legal and regulatory regimes which recognize individual information rights) essentially consider the *effectiveness and efficiency* of the market in internalizing Coasean-style ‘privacy externalities’ among the participants and must therefore proceed to specify optimal remedies from both the information *supplier’s* point of view (for example, a commercially transacting consumer using the Internet) and the person or organization obtaining access to the information for business purposes but whose custody and control potentially present an inherent information *privacy risk* to the owner of the information (for example, a bank or airline which shares information with marketing companies) (Coase 1960; Dahlman 1979; Acquisti 2010). I will briefly review the essential aspects of this market mechanism before turning to the comparative economics of security.

The microeconomic analysis of economic externalities has been undertaken in many markets including environmental protection and telecommunications regulation where unintended costs resulting from the underlying economic activity are incurred by a party who did not agree to the action, leading to several types of market inefficiencies (Pigou 1932; Coase 1960; Buchanan 1962; Baumol 1972; Dahlman 1979; Haddock 2003; Eagle 2004; Haddock 2007; Barnett 2009). From this economic analysis perspective, *loss of control* of personal information can be viewed as an *unintended cost to the individual* resulting from the underlying economic activity involving the use of the information and, in situations where losses affect many individuals, might be considered analogous to, for example, pollution production during manufacturing. From a privacy perspective, loss of control over information is therefore an externalized, non-priced cost imposed on the individual to whom the information pertains and who, in the absence of a legal or contractual means to assign the costs of the loss to the data custodian (i.e. without a corresponding ability to either charge the custodian for the loss or pay the custodian to prevent the loss), suffers the entire and often unforeseen consequences of the loss. Market ‘inefficiencies’ may also result from the *expectation of loss* on the part of the individual whom, in the absence of an *ex ante* ability to transfer liability for potential costs or an *ex post* ability to charge the data custodian for actual costs incurred may, for example, *ex ante* restrict the amount of socially optimal data provided to the market (Shapiro 1997). Healthcare regulators typically warn that patients who mistrust healthcare providers’ ability to confidentially manage their personal health information may restrict the supply of information or obfuscate key facts about themselves which may be perceived to be at greater risk creating unintended healthcare or other consequences for both the care provider and the patient.

On the other hand, it has been argued that the achievement of optimal privacy protection does not require the establishment of legal or contractual property rights over the information, as long as the parties are able to successfully contract over the avoidance of loss and the costs to contract are (sufficiently) low or non-existent (Hermalin and Katz 2004; Chellappa and Shivendu 2006). In reality, the ability of an individual or groups of persons to effectively contract for *ex ante* control over or *ex post* compensation for loss of their confidential information is not costless (and may be entirely impracticable) in many markets and may therefore result in suboptimal welfare for either the individual or society in general, if not outright market failure (Kahn, McAndrews et al. 2000; Noam 2002; Schwartz 2004). Where markets for personal information subsequently fail due to inordinate transactions costs, lack of information symmetry or other factors[9], and where the cost of failure is considered to be too large for one or more of the transacting parties to bear or is considered sub-optimal from a social welfare perspective, regulators have seen fit to rebalance interests by assigning specific regulatory *property rights* to individuals in their personal information with, at minimum, associated privacy assurance control obligations assigned to data custodians. The leading example of this has been the evolution of *fair information handling practices* (FIPs) which originated from OECD concerns over the privacy assurance of transnational data flows and has resulted in a set of principles, legislation and regulations governing the collection, use and disclosure of personal information within and between most Western jurisdictions since the early 1980's. FIPs continue to strongly influence jurisdictional control over the institutional handling of personal information in both commercial and non-commercial markets (Bennett 1992; Schwartz 1994; Rotenberg 2001; Prins 2006; Anderson, Böhme et al. 2009; Mendez and Mendez 2010; Schwartz 2010).

Several authors have concluded that some assignment of property rights over personal information, including legal rights of *alienability* (the ability to sell the rights to the information), and the associated incentives to protect personal information through the *individual* employment of controls (e.g. personal data encryption technologies or explicit contracting for the imposition of custodial controls) are necessary for the optimal management of personal information (Lessig 2000; Litman 2000; Samuelson 2000; Lessig 2002; Schwartz 2004; Hall 2010). Others have argued that absolute assignment of property rights over personal information is impracticable and would result in *pareto*-inferior results for both individuals and society within the current context of interconnected electronic information sources and uses especially where prospective losses are difficult to estimate and liabilities cannot therefore be reliably assigned to individual custodial actors or institutions. A number of prominent authors argue that some restricted combination of property rights *and* industry regulation requiring the establishment of controls by custodians is pareto-superior to a purely market-based approach to data protection involving personal information (Schwartz 1997; Schwartz 2000; Rotenberg 2001; Rodwin 2010), and personal health information specifically (Schwartz 1997), although the effectiveness of *ex ante* regulation vs. *ex post* liability in achieving desired outcomes is unclear (Romanosky, Telang et al. 2008; Romanosky and Acquisti 2009; Romanosky 2010). Significantly, for the purposes of establishing my central hypotheses, it

[9] Noam (2002) indicates several market requirements for 'privacy transactions' to materialize: 1) sufficiently low transaction costs; 2) a legal environment that permits transactions to be carried out; 3) an industry structure that permits transactions to occur; 4) symmetry of information among the transacting parties; and 5) no "market failure." Examples of market failure include the presence of externalities, free riding and moral hazards. For example, in a market where customers could pay to avoid telemarketing calls, it is possible that some unscrupulous marketers might call customers just to get them to pay to avoid the calls. The market 'fails' because the number of socially optimal unwanted calls would actually increase.

should be noted that most, if not all, of these studies do not consider whether the actors involved may be biased in their risk preferences and attitudes and therefore over prospective decisions over privacy or security outcomes and associated controls. This is important to this research since Coasean predictions of preference and behaviour have been tested and found wanting (Kahneman, Knetsch et al. 1990) and should therefore be taken into consideration, if not explicitly modelled, when predicting behaviour under circumstances where existing privacy and security regulations may affect transactions costs, informational rights endowments and opportunity costs of control for both individuals and information custodians. Jolls et al comment on the failure of Coase in making predictions in situations where there may be real differences in the way people react to perceived wealth effects and opportunity costs: "Many economists and economically oriented lawyers think of the Coase theorem as a tautology; if there were really no transaction costs (and no wealth effects), and if an alternative allocation of resources would make some agents better off and none worse off, then of course the agents would move to that allocation. Careful empirical study, however, shows that the Coase theorem is not a tautology; indeed, it can lead to inaccurate predictions. That is, even when transaction costs and wealth effects are known to be zero, initial entitlements alter the final allocation of resources. These results are predicted by behavioural economics, which emphasizes the difference between opportunity and out-of-pocket costs. (Jolls, Sunstein et al. 1998)

### The Economics of Security

The potential failure to protect individual privacy can also be viewed uniquely from the perspective of the *data custodian* with the associated change in emphasis from that of *individual* privacy assurance to one of *institutional* data *security* assurance, associated risk management approaches and managerial decision making over data management controls. The selection of perspective is instrumental in understanding emerging opportunities for information privacy and security research where the control decision criteria and control selection methods used by data custodians, as well as the actual resulting controls and control performance, may be expected to be inherently different than that of an individual (or groups of individuals) who have a claim to only a portion of the data under institutional custody at any one time. The information security literature is also based on several considerations which permit an economic analysis of the relevant information control issues based on: 1) the managerial decision makers' relative valuations of the information in use; 2) the implied or expressed ethical, legal and contractual 'property rights' over the information within the context of a confidential relationship; 3) the managerial decision makers' interpretations of the resulting business outcome objectives and normative information control obligations and; 4) the managerial decision makers' subjective expectations over *prospective* information threats, vulnerabilities and the likelihood and impact of experiencing an information loss scenario; and 5) the respective the managerial decisions regarding the means to control and thereby optimize the expected net benefits derived from the use and sharing of confidential information under secure conditions.

The shift in perspective from individual privacy to institutional security assurance therefore allows for an analysis of managerial, firm and industry implications from the broader (and for my purposes, more practical) perspective of the 'economics of security' (Soo Hoo 2000; Anderson 2001; Gordon and Loeb 2002; Cavusoglu, Mishra et al. 2004; Anderson and Moore 2006; Gordon and Loeb 2006; Cavusoglu, Raghunathan et al. 2008; Acquisti, John et al. 2009; Appari 2010). Ross Anderson and others define the economics of security in terms of the incentives to secure data and systems from loss (Anderson and Moore 2006; Camp 2006; Anderson and Moore 2007):

> According to one common view, information security comes down to technical measures. Given better access control policy models, formal proofs of crypto-graphic protocols, approved firewalls, better ways of detecting intrusions and malicious code, and better tools for system evaluation and assurance, the problems can be solved. … I put forward a contrary view: information insecurity is at least as much due to perverse incentives. Many of the problems can be explained more clearly and convincingly using the language of microeconomics: network externalities, asymmetric information, moral hazard, adverse selection, liability dumping and the tragedy of the commons. (Anderson 2001)

Approaching this research from the perspective of the economics of security further differentiates this work from studies of privacy assurance by the fact that businesses and government employ managers who are responsible and accountable on behalf of the institution for making decisions over the selection, specification, deployment, management and performance of information controls and who might therefore be expected to introduce subjective decisional perspectives including risk preferences and attitudes regarding the control of the information (Hirshleifer 1971; Hirshleifer 1980; Stigler 1980; Eisenhardt 1989; Robey and Boudreau 1999; D'Arcy 2009). In this research, I am therefore interested in examining managers' decisional perspectives, attitudes and decision making methods which may be subject to institutional effects including agency bias and other behavioural economic factors which have been demonstrated in other institutional management settings to impact portfolio allocation decisions (here notionally represented by a prospective portfolio of security controls) by affecting risk attitudes and associated risk behaviours. A number of behavioural and functional models can be employed which have direct application to the topic, including institutional and agency theory (Meyer and Rowan 1977; DiMaggio 1983; Zucker 1987; Eisenhardt 1989; Robey and Boudreau 1999; Appari 2009); actor-network theory (Bonner and Chiasson 2005; Latour 2005; Bonner, Chiasson et al. 2009; Law 2009); behavioural economics, including behavioural law and economics (Jolls, Sunstein et al. 1998; Mullainathan 2000; Acquisti and Grossklags 2006) and decision making under uncertainty (Ellsberg 1961; Kahneman 1979; Thaler 1981). Furthermore, I propose that many of the emergent theoretical and experimental approaches to analyzing the economics of privacy which incorporate behavioural influences on decision making under conditions of uncertainty undertaken by Acquisti and others (Acquisti 2004; Acquisti 2005; Acquisti and Grossklags 2005; Acquisti and Grossklags 2006; Acquisti, John et al. 2009; Bansal, Zahedi et al. 2010) might be applicable to the study of the control selection behaviour of security managers as they attempt to balance informational management priorities with other institutional objectives.

Apart from purely theoretical economic models of security management, from an *operational* standpoint there are numerous emerging models for both qualitative and quantitative decision making regarding security controls. Immediate issues arise in practice when attempting to evaluate either inputs or outcomes from security alternatives, some of which are inherent in the complexity of the systems that are under control. On the other hand, it would seem that, regardless of the explicit approach taken, there is an increasing need for quantification of the problem, and a need to express degrees of 'satisfaction' across alternatives in the context of what are inherently uncertain circumstances. Quantification remains stubbornly elusive in many non-trivial institutional settings. "[Actual] decision makers need information about the utility, such as the reliability or estimated costs and impacts, of different security options to make appropriate decisions. Quantification has been suggested as a solution to such needs" (Soo Hoo 2000; Verendel 2009). Methodologically, quantitative approaches to system security *evaluation* represent a relatively recent area of study and the relevant survey literature indicates that many quantitative approaches have grown out of comparatively static *system dependability/fault models* which, for example, do not generally incorporate core features of the 'security problem' such as stochastic, external threat models or control attributes (Nicol, Sanders et al. 2004; Conrad 2005) or 'strategic' approaches to managerial decision making within real life security contexts (Kunreuther and Heal 2003; Verendel 2009). In addition, there are basic challenges in establishing meaningful *security metrics* which would support robust quantification of security *risks*. Verendel's survey work reviewed 90 quantitative security papers between 1981 and 2008 to evaluate whether security can be represented using quantitative information and concluded that "…quantified security is a *weak hypothesis*: for most cases it is unknown if the methods are valid or not in representing operational security." (Verendel 2009). There are also relatively few recent examples of quantitative approaches to security risk management involving conceptions of economic trade-offs under conditions of loss uncertainty (Grossklags, Johnson et al. 2010; Johnson, Grossklags et al. 2010), or which explicitly include managerial decision bias in the perception of security risk (Blakley, McDermott et al. 2001; Schechter 2004; Conrad 2005; Schechter 2005; Schroeder 2005; Huang, Hu et al. 2006; Sklavos 2006; Grunske and Joyce 2008; Verendel 2008; Beres, Casassa Mont et al. 2009; Verendel 2009).

There are, nonetheless, recent and specific approaches in the literature to qualitative and quantitative security risk management models which incorporate some of the behavioural and institutional bias effects noted above (Cavusoglu, Mishra et al. 2004; Cavusoglu 2006; Poindexter, Earp et al. 2006; Beautement, Sasse et al. 2008; Cavusoglu, Raghunathan et al. 2008; Beautement, Coles et al. 2009; Beautement and Pym 2010). In addition, certain qualitative approaches to examining individual decisional bias from the economics of privacy literature can be modified to specifically address security decision making by institutional managers. Acquisti, for example, has been instrumental in developing novel experimental designs demonstrating decisional bias for individuals faced with making privacy control decisions, although his focus is almost exclusively on the person to whom the information pertains (Acquisti 2004; Acquisti 2005). I propose that there are also seminal examples of economic lab experiments undertaken to determine general deviations from expected utility theory and rational choice assumptions which could be modified and used as a basis for determining risk preference and decisional bias in security managers (Kagel and Roth 1995; Jolls, Sunstein et al. 1998).

There are several representative studies of the qualitative analysis and quantitative *modeling* of information security decision bias in the recent literature which have been reviewed in depth for their applicability to this research.  Soo Hoo presents one of the earliest security decision analyses which questions the traditional use of established "annual loss expectancy" form quantitative estimates of risk (Soo Hoo 2000). Schroeder conducted a qualitative survey for security control selection using non-professional subjects (Schroeder 2005) utilizing virtually the identical questionnaire panel originally utilized in Kahneman and Tyversky's seminal work on *prospect theory* (Kahneman 1979), modified for an information security context. Verendel directly addresses a prospect theory approach to security decision making under conditions of uncertainty which reframes the decisions explicitly in terms of the perceived security risk. (Verendel 2008). This approach, acknowledging the limitations of obtaining accurate security metrics, indicates that, even with accurate metrics, there is value in understanding the potential for bias in the decision maker in order to make accurate predictions of utility maximization. This seems to be an appropriate approach within this research for describing how managers do and ought to make privacy and security decisions under conditions of uncertainty.  As noted by Jolls et al, an emphasis on the *predictive* value of the chosen model without the practical ability to generate accurate security risk measures is an important aspect of taking a behavioural approach to positive economics: "The difference between conventional and behavioural [law and] economics is not just a difference in the validity of the assumptions about human behaviour. While the [traditional] assumptions of unbounded rationality, willpower, and self-interest are unrealistic, the force of behavioural economics comes from the difference in its predictions … In this sense [this approach] is consistent with the precept originally proposed by Milton Friedman: Economics should not be judged on whether the assumptions are realistic or valid, but rather on the quality of its predictions." (Jolls, Sunstein et al. 1998).

Antoniou's (Antoniou 2010; Antoniou, Harrison et al. 2010) doctoral dissertation undertook considerable research within the seminal and recent decision making literature which establishes the perspective that behavioural decision research challenges traditional neoclassical economic assumptions of expected utility maximization (von Neumann 1947): specifically, the assumption that decision makers are able to consistently apply Bayes Rule to estimate unobservable probability distributions and that they update their expectations of, in this case, future financial asset prices by accurately incorporating new market information over time, leading to alternative models of utility and risk attitude formation (Simon 1955; Savage 1971).  Antoniou, Harris and (and others) appropriately note that 'experimental economics' approaches specifically seek to qualify many of the psychological experimental methods (e.g. by providing monetary incentives to subjects to avoid hypothetical decisions, etc.) and, even if sound, are inherently revisionist compared to much of the established decision making research literature. For the core of the empirical studies undertaken, Antoniou correctly cites three biases from the seminal literature which may, in fact, be observed to operate in financial and here, I propose, with direct analogs in security 'markets' involving decision making under uncertainty (Hirshleifer 2001) and provides relevant theoretical and empirical research references for each bias discussed and directly addresses each of these in turn: 1) heuristic simplification (Tversky and Kahneman 1974), specifically involving subjective estimates of the strength and weight of information (Griffin and Tversky 1992) and leading to over or under reaction and associated pricing changes (Barberis, Shleifer et al. 1998; Sorescu and Subrahmanyam 2006); 2) self-deception, manifesting as investor ambiguity aversion to noisy information (Gilboa and Schmeidler 1989) leading to 'market' pessimism and correction (Camerer and Weber 1992); and 3) investor sentiment, resulting from decision maker optimism and overconfidence (Johnson and Tversky 1983) and resulting in serially correlated price changes and price momentum (Daniel, Hirshleifer et al. 1998).

The most basic of the biases is risk attitude or preference, generally expressed as the degree to which a decision maker is either relatively 'risk averse' or 'risk seeking' over either risky or uncertain prospects. Bernoulli proposed and Von Neumann formalized the axioms under known probabilities in which a decision maker would reveal the 'certainty equivalent' of the expected value of a risky lottery which defined the 'expected utility' of the risky decision for the decision maker (von Neumann 1947). Correspondingly, a person, when asked to choose between either a sure money equivalent or a lottery based on probabilities of outcomes in which the two choices have the same expected value, who subsequently chooses the sure money equivalent is deemed 'risk averse', whereas the one who chooses the lottery is deemed 'risk seeking' (at least for the single choice presented). What is revealed is the person's subjective valuation of the expected value of the lottery which, if the decision maker correctly understands the game, will never actually occur in practice: rather, one of the discrete risky outcomes will occur and the decision maker understands it will never be equivalent to the sure money option (except as an average measure over many repeated games or unless the lottery choice is trivial where one of the outcomes will occur with a 100% probability and is technically equal to the sure money outcome). Savage proposed that this subjective calculation will occur even in the absence of certain probabilities i.e. under conditions of uncertainty, where the decision maker will behave 'as if' he or she held consistent beliefs regarding the underlying probabilities and will always hold a subjective expected utility over the prospect. If the game were repeated, Savage also proposed that the decision maker would update their beliefs regarding the underlying probabilities in a manner consistent with Bayes theorem. This normative perspective on decision making under uncertainty was tested by Allais and Ellsberg who showed that individual violations of one or more of the axioms indicated that subjective expected utility did not in fact hold in all decision cases and that other decisional biases (Ambiguity aversion, probability weighting, Loss aversion, etc.) were in play, leading to alternative theories of choice (e.g. Ranked Dependent Utility).

In the context of ambiguity aversion for example, Antoniou cites seminal work by Fischer Black which first identified the problem investors face attempting to separate market signal from noise (Black 1986) and which leads to decision making using information they believe to be true using heuristics identified in both the psychological and behavioural economics literature. Antoniou also references the seminal theoretical psychological and economic literature on ambiguity aversion (Ellsberg 1961; Savage 1971), provides substantive references for more recent theoretical and empirical literature surveys (Camerer and Weber 1992) and outlines the psychological and economic experiment foundations for explaining the bias phenomenon (Keren and Gerritsen 1999). He also cites similar studies which incorporate ambiguity aversion within asset markets and which inform several market anomalies such as the equity premium puzzle (Mehra and Prescott 1985). If ambiguity bias is indeed present, then decision makers should be expected to predictably respond pessimistically to noisy signals and 'market' investors (here taken as control decision makers in a security setting) can therefore expect a period of price declines (under investment in control) followed by a correction as 'prices' (the value of control investment) adjust to actual information received over time (DeBondt and Thaler 1985).

From a bias quantification perspective, Glenn Harrison has been particularly prolific in developing experimental methods for testing and modelling risk attitudes and biases using lab experiments (Harrison and McKee 1985; Harrison 1994; Harrison and Rutström 2008). His contributions to my approach are many and form a significant basis for both methodology and analysis, here directly applied within the context of information security control selection. First, he emphasizes that experimental methods in behavioural economics must be carefully matched to the *theoretical model of decision making* being hypothesized since the subject's experience of the experiment may not be the same as that intended by the researcher (Harrison 2010; Smith 2010). Second, rather than assuming risk attitude, he argues that risk attitude should (almost always) be estimated within a *structural*[10] utility model, since other model parameters will be biased if the subject's risk attitude is incorrectly assumed (Harrison 1986; Wilcox 2008; Keane 2010; Quinn 2011). Third, he advocates the use of *mixed* structural models of latent choice, for example combining expected utility theory (EUT) and prospect theory (PT) specifications within the same structural model, with a weighting function, to allow for heterogeneity *between* theories of choice within a sample (Harrison and Rutström 2009). Fourth, he advocates the *joint* estimation of the parameters of a complete structural model rather than using non-nested approaches for model selection and hypothesis testing (Harrison and Rutström 2008). He also provides important guidance on the treatment of model errors which, in an experimental context concerning latent choice modeling, likely reflect a much wider range of effects aside from stochastic sampling errors (Harrison 2010). Finally, Harrison has fully

[10] Quinn: "There is no single, simple definition of what constitutes a structural model…we will consider a structural model to be any economic model that specifies a theory about how agents optimize and then use the implications of that theory literally as a basis for empirical estimation. In doing so, structural models often try to estimate 'deep parameters' governing human behaviour (for example, the return to capital, or preferences over risk) rather than merely identifying particular causal relationships in a particular context (for example, the 'average treatment effect' of a given policy on some outcome)". (Quinn 2011)

documented many of his experiments, including detailed specification of ‘linking functions’ in the context of binary choice experiments which permit the maximum likelihood estimation of the proposed structural choice models and form the crucial bridge between theory and econometrics in his experiments. Through personal correspondence he has also provided source code and data from his experiments for estimation of the models using STATA (Harrison 2008), and has provided me with guidance regarding appropriate customization of the code in practice and advice on sampling and statistical testing. Details of the Harrison approach and the applicability within this research are described in more depth in Section 3.

### Control Choice Under Risk and Uncertainty Using System Simulation

As noted above, the drive for quantification of security risk and the associated economic evaluation of control choice alternatives is valuable for management in terms of providing a better means to compare control alternatives, even if perception of the risks and the resulting control choices may be subject to individual psychological bias or institutional decision making effects. A key problem with the approach is therefore to obtain sufficient and robust enough operational data to establish a risk profile of the security operation in question. Several authors have recently argued for the use of system simulation to solve the problem, both in the field and in the lab. A collection of researchers connected with HP labs in the UK have contributed to the use of both simulation models in the context of ‘complex, multi-attribute” decision theory for security control research. Following on the multi-attribute trade-off modeling of Keeney (Keeney 1982; Keeney 1993), and the formalized information systems modeling work of Collinson (Collinson, Monahan et al. 2009), Beautement undertook several studies involving the specification of a system at risk incorporating a utility framework for decision over control selection where the trade-off between system confidentiality and availability given that a level of investment in security control is analogous to the trade-off between inflation and unemployment based on setting interest rates (Beautement, Coles et al. 2009; Beautement and Pym 2010). Ioannidis has also explored security control selection as a utility maximization problem under uncertainty incorporating loss aversion (Ioannidis, Pym et al. 2011) and specifically analyses the trade-offs in a security CIA context using quadratic loss functions  (Zellner 1986; Ioannidis, Pym et al. 2009). Beres also specifies the control selection problem in terms of ‘loss functions’ and conducts a qualitative survey of risk preferences and compares these to ‘optimal’ solutions obtained by a Monte Carlo simulation of the HP “Gnosis” security simulation model (Beres, Casassa Mont et al. 2009; Beresnevichiene, Pym et al. 2010). Baldwin also conducted a qualitative decision experiment which divided 12 decision makers between a control group and an ‘intervention’ group that undertook decisions after using a decision support tool driven by an HP simulation model of the relevant system risks (Baldwin, Beres et al. 2011). The intervention group gave markedly higher weight to the ‘economic’ trade-offs incorporated into the simulation between control cost, productivity and security outcomes than the control group.

The use of simulations as the setting for control decisions has also been explored in areas of public policy, notably environmental policy in which the underlying system at risk may be relatively complex and where endogeneity of risk is a factor. In his doctoral dissertation supervised by Harrison, Sen (Sen 2010) has replicated a study by Fiore (Fiore, Harrison et al. 2009) which provides a lab experiment methodology incorporating binary choice over "endogenous risk" and is directly applicable for my purposes. Risk is endogenous when an individual is able to undertake mitigation or self-protection actions, including self-protection and self-insurance, that reduce the risk that he faces and is therefore a valid frame for economic decisions over security controls which mitigate risk (Ehrlich and Becker 1972; Shogren and Crocker 1991; Quiggin 2002). Indeed, the relatively recent increased attention being paid to 'cyber-insurance' is itself indicative of a maturing market for security control selection in which decision makers may be expected to balance a portfolio of controls between traditional self-protection (e.g. detective or preventative security controls) and insurance (including self-insurance) ranging from compensating and recovery controls to business or operational loss cyber insurance. I will return in depth to this aspect of the control portfolio in Game 5 concerning the cyber insurance experiment.

Controlling for endogeneity of risk is therefore important within this management domain and includes psychological aspects of decisional bias such as moral hazard effects and overconfidence (Antoniou 2010). The use of system simulation to generate a wide range of risk scenarios which supports specific risk inputs and outputs and which can incorporate specific panels of controls presents a valuable platform for security control experimentation. The analog to established experiments in other domains also supports the ability to benchmark against similar decisional bias hypotheses, here applied to the domain of information security. Sen, for example (Sen 2010), utilizes a computer based, 3D 'virtual reality' simulation of a forest 'system' at risk of wildfire within which the decision maker owns a house, and for which the subject must decide whether to purchase a fire prevention control (in this case a 'controlled burn' to decrease the proximate fire fuel from the area surrounding the house) which lowers but does not eliminate the likelihood that the house will burn down if there is a wildfire. As a matter of overall approach, Sen's hypotheses, experimental setup, data modeling and analyses are highly appropriate as a methodological guide for my research model and design, with the noted modifications I detail in Section 3. Sen's approach also concerns both individual and group decision making under conditions involving known and unknown (subjective) risk estimation and employs Harrison's generalized approach to modeling and estimation of risk attitude.
Sen's inclusion of group decision making treatments and allowance for the possibility of decisional biases is also motivated by relatively recent research interest in the area of behavioural economics. In his survey of the area, Kugler notes that

> Traditional game theory, the science of rational behaviour in interactive [group] settings, makes a few assumptions—mostly based on the concept of *homo-economicus*…If one accepts these assumptions, comparing the behaviour of individual decision makers and the behaviour of unitary

> groups[11] seems almost dull. When there is a unique game-theoretic equilibrium or optimal choice, both individuals and groups should follow the normative prediction, and [aside from the relatively small number of multiple equilibria scenarios] their choices should not differ at all…Considering the enormous recent body of literature on individual behaviour in interactive contexts, it becomes clear that while traditional game theory is still very useful as a normative theory, it fares less well as a descriptive tool. If game theory is expected to provide a realistic account of human behaviour, its assumptions have to be adjusted. One should take into account heterogeneity in levels of rationality, different extents of other-regarding preferences, and different forms of uncertainty attitudes among decision makers. Once these assumptions are integrated into classical game theory, the analysis of group decision making becomes interesting and important. Therefore, investigating group decisions in games has slowly picked up in the late 1990s and after the turn of the century, leading Camerer in his widely-used textbook *Behavioural Game Theory* to conclude that the study of group decisions making is among the top ten research programs in behavioural and experimental economics. (Kugler, Kausel et al. 2010)

Finally, an interesting extension of the 'unitary group' model for decision making is that premised on principal-agent theory (Myerson 1982). Gordon suggests a principal-agent model based on CFO monitoring of security risk profiles and associated compensation incentives for security officers as a check against 'empire building' i.e. over spending on controls (Gordon, Loeb et al. 2008). This approach recognizes that group members may have different objectives in terms of utility maximization and that unitary decision experiments may not be the best model to represent actual security decision tasks Gordon has also proposed several other conditioning factors on the control allocation task, including real options considerations (Gordon, Loeb et al. 2003) and their controversial finding that for certain classes of vulnerabilities, the optimal control investment does not exceed a 1/e fraction (about 36%) of the total value at risk (Gordon and Loeb 2002). Willemson and others have commented on the sensitivity of the Gordon and Loeb model to underlying assumptions regarding the specific functional form of the effect of controls on the vulnerability (i.e. typically, increasing at a decreasing rate) and show that the conclusion may not be robust for all plausible vulnerability models.

[11] A unitary group is a group that has to come up with a joint decision and does not face any internal conflicts of interests in terms of payoffs.

# 3 – Methodology

## Empirical Approaches to Decision Research

The framework for understanding and choosing an appropriate research model applicable to my research topic starts with a specific perspective on the philosophy of science and the philosophy of research. The philosophy of science involves both an *epistemology* (What is knowledge? How is knowledge acquired? How do we know what we know?) and a *methodology* for the actual acquisition of knowledge. According to Karl Popper, a scientific approach to knowledge requires a 'logically consistent portrait of the world' with validation based on evidence and empirical falsification of theory, a view originally proposed by David Hume in the 1740s. For Hume, everything else is natural philosophy or metaphysics: non-observational and, therefore, to Hume, largely 'nonsense'.

In my field of interest, the paradigm of the *decision maker* – in this case, of an information manager with both a fiduciary and (we assume) a personal interest in the prospective success of the security controls to be decided upon - is normatively based on the established economic theory of the rational utility maximizer (von Neumann 1947). This theory has encountered substantial challenges since the early 1950s (Simon 1955) and has, arguably, undergone a substantial paradigm shift (Kuhn 1962) with the increased interest in descriptive vs. normative theories of decision making.[12] That is, based on the actually observed evidence of decisions being made 'in the field', we start by questioning whether the presumed decision theory is sound and are therefore open to its testing and possible revision with specific implications for our view of the security administrator and his/her 'administrative' world. The starting point for consideration of the security decision problem is to essentially recognize the difference between decisions made under known probabilities of outcomes (risk) and that class of problems involving unknown probabilities (uncertainty) [13]. What I am examining, here in the context of security control choice, is therefore alternate theories of decision making which are required to explain a particular class of decisional problems including, but not restricted to, choice under *uncertainty* where, for example, the expected utility of a discrete decision (as in Bernoulli's St. Petersburg Paradox) is observed to *measurably* deviate from its expected value (as proposed in Pascal's Wager)[14]. The overall question is whether and to what degree these 'deviations' may be present

[12] According to Simon: "Recent developments in economics, and particularly in the theory of the business firm, have raised great doubts as to whether this schematized model of economic man provides a suitable foundation on which to erect a theory - whether it be a theory of how firms do behave, or of how they "should" rationally behave. It is not the purpose of this paper to discuss these doubts, or to determine whether they are justified. Rather, I shall assume that the concept of "economic man" (and, I might add, of his brother "administrative man") is in need of fairly drastic revision, and shall put forth some suggestions as to the direction the revision might take…Broadly stated, the task is to replace the global rationality of economic man *with a kind of rational behaviour that is compatible with the access to information and the computational capacities that are actually possessed by organisms* [emphasis added], including man, in the kinds of environments in which such organisms exist." (Simon 1955)

[13] According to the economist Frank Knight, *uncertainty* is different from *risk*, in which there is a specific probability assigned to each outcome (as when flipping a fair coin). Uncertainty involves a situation that involves unknown probabilities, and where the estimated probabilities of a closed set of possible outcomes need not add to unity from the perspective of the decision maker. See Knight, F.H. (1921) Risk, Uncertainty, and Profit. Boston, MA: Hart, Schaffner & Marx; Houghton Mifflin Company

[14] Generally, this class of decisions include: single *choice under uncertainty* ((Brooke 2010); *intertemporal choice* (Acquisti and Grossklags 2003); *strategic decision making*, generally with two or more decision makers either cooperatively or in competition

in the domain of security control choice. The pertinent scientific questions for this research are therefore: 1) Which philosophy of research and which research methodology, including the choice of quantitative and qualitative methods for data collection and analysis, is appropriate for describing these decision making processes?; and 2) within the context of organizations who need to make security choices, what are the normative implications of these research findings for managers and institutions who, presumably, want to make better informational control decisions?

The dominant research approaches for assessing theoretical claims are positivism and, variously, interpretivism[15] and its philosophical and practical middle ground, relativism. *Positivist* approaches assume that the researcher can objectively collect data, form hypotheses, test hypotheses and either accept or reject hypotheses on the basis of the data and that the results are generalizable outside of the research frame. Interpretivism assumes that human behaviour cannot be observed (or described) objectively and so the researcher should not presuppose a *causal structure* (or therefore propose a hypothesis based on an assumed causal structure), but rather seek to reveal the 'hidden generative structures' of what are (at least within business research) essentially social phenomena. Results of interpretivist research may therefore not be generalizable outside of the observed behaviour, but are also not expected to be since there are no presupposed generalized phenomena to discover. Extensions of the interpretivist approach essentially reject the empirical approach to measuring social phenomena entirely, and are exemplified by Jurgen Habermas' view that social science and 'described' history (and so business research) can only have a "…situation-specific understanding of meaning that can be explicated only hermeneutically... access to a symbolically pre-structured reality cannot be gained by observation alone" [16] (Habermas 1967). Interpretivist research into economics, business and, specifically, information technologies is, however, well represented in the research literature and include agency theory (Eisenhardt 1989), institutionalism and neo-institutionalism (Meyer and Rowan 1977; DiMaggio 1983) and actor-network theory (Latour 2005; Bonner, Chiasson et al. 2009; Law 2009). *Relativism* appears as a compromise between positivism and interpretivism for business research and assumes that structure is, to a certain extent, discoverable, objectively describable, somewhat causal and, possibly, generalizable, but requires the researcher to carefully balance both positivist and interpretivist approaches within the study.

As will be noted below, the observed gaps between the normative theory and the descriptive realities of decision making led early researchers to the development of lab experiments in both psychology and economics which directly supported *empirical decision research* and which forms the basis for quantitative

involving Nash or non-Nash equilibria  (Rosenthal 1989); and 'complex decisions' where the optimal behaviour or outcome is dependent on 'multi-stakeholder, multi-objective, multi-attribute' models (Keeney 1993).

[15] This branch is also referenced, variously, under the topics of Phenomenalism, Subjectivism, Subjective Idealism, Anti-Positivism, Post-Positivism, Subjective Idealism, Sociological Interpretivism, Interactionism, and others.

[16] *Hermeneutics* refers to the interpretation of texts or events which can only be understood as a product of the culture, symbols and signs which produce them. It is also one of the bases of *critical theory* postmodern critical research in which the rejects the idea that a researcher's work is considered an "objective depiction of a stable other" (Lindlof and Taylor 2002)

experimental economics. This research will therefore take a positivist (quantitative) approach using an experimental strategy following Yin indicated in Table 1:

**Table 1 - Research Strategies**

| | Type of Research Question | Control Over Event | Contemporary Event | Used in this thesis |
|---|---|---|---|---|
| **Experimental** | How, why | Yes | Yes | Yes |
| **Archival** | Who, what, where, how many/much | No | Yes | No |
| **History** | How, why | No | No | No |
| **Case Study** | How, why | No | Yes | Yes |

(Yin 2013) from (Johansson 2005)

The following sections outline this evolution, identify seminal experimental approaches relevant to my topic and provide a basis for my proposed experimental approach based on integrating current work in both experimental economics and the economics of information security.

### Psychological and Economic Considerations for Decision Making Under Uncertainty

This research addresses the potential for, modeling, parameter estimation and implications of the presence of cognitive biases[17] in decision making, generally under conditions of uncertainty, in the context of information security control selection. My literature review to date documents that bias can be expected to be present in decisions made by individuals over control of their own personal information (Acquisti 2010; Baddeley 2010; Rivenbark 2011), however my focus is exclusively on institutional agents, specifically privacy, security or chief information/technology officers. These persons are tasked with the appropriate selection of institutional information security controls which are intended to protect the information entrusted to them by other individuals, and who may encounter both endogenous and exogenous constraints (whether perceived or otherwise) within the applicable decision making environments and processes. One of the most significant endogenous constraints identified in the literature and successfully experimentally *elicited* in both lab and field settings is the potential bias of the decision maker, for which the literature identifies several dimensions (Jolls, Sunstein et al. 1998): 1) bounded rationality, as first introduced by Herbert Simon, involving natural and situational limitations of human computation skill or flawed memory (Simon 1955), and particularly the experimentally observed use of *heuristics* – short cuts or 'rules of thumb' to solve non-trivial problems and to make decisions using subjective assessments of event probabilities (Tversky and Kahneman 1974; Griffin and Tversky 1992); 2) bounded willpower, including limitations to undertake decisions which may be in the decision makers self-interest; and 3) bounded self-interest, including aspects of both altruism and spitefulness. This research focuses exclusively on an identified class of bounded rationality biases by quantitatively modeling and testing for decisional bias

[17] https://en.wikipedia.org/wiki/List_of_cognitive_biases

under conditions of uncertainty within a lab setting and will thereby attempt to identify the implications of these biases specifically within the domain of information security control selection supporting institutional risk management.

The testing of individual probability estimates concerning the risk of an information system, for example, is of core interest to this research since the business context in which security control decisions are undertaken inevitable requires some subjective estimation of probabilities over security risks which are inherently uncertain but which may be revealed in operations to some degree over time and therefore may reasonably be understood from the perspective of the decision maker as a type of 'repeated game with learning' (involving a series of individual repeated choices which are possibly, although not necessarily, interdependent). This problem setting is analogous to established experimental economics approaches for eliciting attitudes over uncertain bets based on subjective risk modeling and favors a Bayesian updating[18] approach to decision making as the baseline model of individual behaviour for the personal estimation of the actual underlying probability of the system at risk (Antoniou, Harrison et al. 2010). In an institutional security operational setting, a range of subjective probability estimates updated over time are required for managing the risk of the information management system, and are therefore directly analogous to the documented experimental approaches noted below, whereby the decision maker is assumed to: 1) form some prior estimate of the probability of risk; 2) makes control choices (i.e. 'purchases' controls with uncertain efficacy over uncertain risks producing uncertain outcomes – essentially a 'bet' that the control will pay off as expected) based on their subjective probability estimate(s) of the risk factors; 3) observes some operational data over time (e.g. the degree to which the controls are actually effective); and 4) repeats the 'betting' process iteratively (i.e. control selection) to attempt to better her chances of avoiding further information loss over time.

### Quantitative Approaches to Information Security Simulation and Control

While many authors have addressed the theoretical, institutional and operational aspects of information security risk management (Anderson 2001; Camp 2006; Anderson and Moore 2007; Verendel 2008; Beautement and Pym 2010), comparatively few have attempted to quantitatively examine the behavioural economics of decision making specifically within the managerial information security control selection process (Soo Hoo 2000; Schechter 2004; Schroeder 2005; Ioannidis, Pym et al. 2009; Appari 2010; Beresnevichiene, Pym et al. 2010; Baldwin, Beres et al. 2011). Verendel's literature review of over 90 quantitative security studies suggests that, while 'quantified' – or quantifiable – security measurement itself may be a 'weak hypothesis', a key determinant of an appropriate modeling, policy and operational approach to control selection, *even assuming accurate security performance metrics*, would inevitably need to take into consideration  the potential for bias in the security decision maker (Verendel 2009).

[18] This theorem is named for Thomas Bayes and often called **Bayes' law** or **Bayes' rule**. Bayes' theorem expresses the conditional probability, or "posterior probability", of a hypothesis H (that is, its probability after evidence E is observed) in terms of the "prior probability" of H, the prior probability of E, and the conditional probability of E given H. It implies that evidence has a confirming effect if it is more likely given H than given not-H. http://en.wikipedia.org/wiki/Bayes'_theorem

Furthermore, while the problem of bias may be fundamental to any decision making setting, it is particularly salient to decisions made under conditions where the prospective risk factors may be uncertain. This applies directly to the information security setting where the underlying probability of information system threats, vulnerabilities and impacts contributing to information security risks may all be uncertain to some degree and whose 'likelihood' is only revealed to the decision maker over time[19]. Quantitative economic approaches for testing this type of subjective learning on the part of a decision maker typically start from the assumption of the rational use of Bayes Rule on the part of the decision maker and proceed to test the degree to which this assumption is in fact violated in practice under experimental conditions. As a positivist approach undertaken within a lab setting designed to control for identifiable confounding factors such as risk framing or event ordering effects, methodologies for quantitatively examining decisions made under these conditions require the researcher to have an understanding of the theoretical bases of behavioural decision research, the ability to formulate and integrate microeconomic models of expected utility and revealed risk attitude and the capability to appropriately undertake empirical lab experiments using valid models of information systems at risk which are designed to elicit quantitative estimates of bias under conditions which can be expected to produce reliable, consistent and generalizable results outside of the lab setting. The combination of these methods for risk attitude elicitation within information security settings is a new area of research which represents opportunity for quantitative models representing both system simulation and decision theory.

### Specification of Structural Models for Choices Under Uncertainty

Traditional security assurance models seek to optimize nominal measures of the system's "CIA" outcomes over a defined period for a given level of control investment: *Confidentiality* (unauthorized access and data loss prevention), *Integrity* (data and system accuracy and completeness) and *Availability* (accessibility when required for purpose). Assuming some ability to attribute outcome metrics to these states and given the stochastic nature of the system states and transitions, the objective control decision is therefore a multi-attribute trade-off problem under uncertainty recognizing that, for example, increased availability can be expected to lower confidentiality, but not always or to the same extent, and depending on the level of control investment affecting one or both attributes. Baldwin's theoretical model assumes that the decision maker has preferences over certain *prospective* states of the system, which can be evaluated based on the decision maker's *expected utility* of the quantifiable business and operational outcomes generated by the system in an observed future period.

[19] 'Likelihood' addresses unknown facts, whereas probability addresses unknown prospects. Likelihood is the hypothetical probability that an event *that has already occurred* would yield a specific outcome. The concept differs from that of a "probability" in that a probability refers to the occurrence of future events (which obviously have not yet occurred and are therefore inherently probabilistic), while a likelihood refers to past events with definite but unknown outcomes. http://mathworld.wolfram.com/Likelihood.html . For example, if a coin is flipped 10 times and it is known to be a fair coin, what is the probability of it landing heads-up every time? Likelihood is used when describing a function of a parameter given an outcome. For example, if a coin is flipped 10 times and it has landed heads-up 10 times, what is the likelihood that the coin is fair? https://en.wikipedia.org/wiki/Likelihood_function#Historical_remarks

The choice of the specific utility function used to model the decision maker's latent preferences over these trade-offs is important since a number of theoretical assumptions must be made prior to the chosen structural specification and this will have implications for both the resulting descriptive and normative models of choice. The most significant assumption is that of the decision maker's form and degree of *risk aversion*. A common *structural utility function* incorporating risk aversion over gains is the *Constant Relative Risk Aversion* (CRRA) *power function* which generally takes the form

$$U(x) = x^{1-r} / (1-r) \tag{3.1}$$

where "r" is the 'coefficient of risk aversion' and "x" is the monetary value of the outcome (>0), where r=0 corresponds to risk neutrality, r<0 to risk loving, and r>1 to risk aversion.  The coefficient of risk aversion is constant for the power function family. 'Expected utility' is then specified as the *probability conditioned* utility,

$$EU = [\ p \times U(x)\ ] \tag{3.2}$$

where p is the probability of the outcome and U(x) is the CRRA specified utility function.

This specification allows the *expected utility* of the outcome to be greater or less than its *expected value* based on the degree to which the person is averse to the uncertainty of the outcome (von Neumann 1947). This is particularly relevant to the information security context where the underlying risk of the system may be unknown and trade-offs between *prospective* outcomes of the system involve *subjective* estimates of the *probabilities* of the resulting system and outcome states. In this context, the utility model may also be extended to allow for *subjective expected utility* (Friedman, Savage 1948) in which the decision maker's estimate of the likelihood of the prospective outcome becomes salient (Antoniou, Harrison et al. 2010).

Variations on the power function permit the specifications of utilities which vary in the degree of *relative* (versus constant) risk aversion being assumed. Harrison and others have extensively reviewed the appropriate specification of utility functions under different circumstances. For example, the *Expo-Power* function (Saha 1993) incorporates an *income* effect as well as risk aversion:

$$U(x) = [1-\exp(-\alpha x^{1-r})]/\alpha \tag{3.3}$$

where $\alpha$ and r are parameters to be estimated (Pratt 1976; Saha 1993). *Relative Risk Aversion* (RRA) is then

$$r + \alpha(1-r)x^{1-r} \tag{3.4}$$

so RRA varies with income if $\alpha \neq 0$. This function ‘nests’ CRRA (as $\alpha \rightarrow 0$) and CARA (Constant Absolute Risk Aversion) as $r \rightarrow 0$ (Harrison, Lau et al. 2010).

Domain appropriate selection of the objective utility function is a major concern in the recent economics of security literature since the security trade-offs are proposed to be multi-attribute CIA preferences on prospective system states and outcomes conditioned by investment in a vector of controls. Further, rather than being a single attribute utility function of the system state as noted above, several recent authors have proposed a number of utility variations appropriate for a security context, but which generally assume risk aversion instead of structurally estimating it. Beautement (Beautement, Coles et al. 2009) proposes the security analog of the ‘Central Bank Problem’ where policy makers attempt to set interest rates to control trade-offs between inflation and unemployment: the Bank is able to accurately set interest rates, but the resulting inflation and unemployment levels are stochastic and cannot be accurately predicted in a future period based on the control choice. The analog expected utility over outcomes in terms of security is therefore, for example, a stochastic *loss function* of system or information confidentiality (C) and availability (A) (which move in opposite directions for a given level of control) and where Availability is a function of the level of control investment ***I*** :

$$C = -\lambda A + \varepsilon_C \qquad (3.5)$$

$$A = -\psi \boldsymbol{I} + \varepsilon_A \qquad (3.6)$$

$$E(U(C,A)) = E(((\exp[\alpha A] - \alpha A - 1)/ \alpha^2 + \frac{\varphi}{2} C^2 \qquad (3.7)$$

In this specification, the assumed objective is to maximize the expected utility of availability and confidentiality which trade-off over a feasible set of outcomes. Beautement later simplifies the specification to

$$U(C,A) = \alpha(A - \beta C) \qquad (3.8)$$

where $\alpha$ and $\beta$ are parameters to be estimated and represent a simple ratio between confidentiality and availability.

As noted above, modifications of the *structural form* of the utility function should also be considered since violations of EUT model assumptions have been well documented in both the psychological and economics literature (Allais 1953; Ellsberg 1961) in which the EUT specification fails to be descriptive of actual decisions made under uncertainty, notably the Independence Axiom underlying the linearity in probabilities. Deviations from the behaviour prescribed by EU Theory have become known as “decision theoretic" paradoxes (Ben-Tal and Ben-Israel 1991), in particular: the Allais paradox or 'common consequence effect' (Allais 1953); the common ratio effect (Cubitt, Starmer et al. 1998), over sensitivity to small probabilities (Kahneman 1979); and the utility evaluation effect (Machina 1983). In context, Beres and other security researchers typically assume that security decision makers are risk averse utility maximizers (Verendel 2008; Beresnevichiene, Pym et al. 2010; Ioannidis, Pym et al. 2011), and the specification of the resulting utility (or loss) functions used to model preferences follow from that assumption without requiring the *parameterization of risk aversion*. Although plausible in an institutional security control context, this is nonetheless a strong assumption within the context of performing experiments that are meant to elicit decisional biases which are hypothesized to be dependent on risk attitudes and are intended to be descriptive and not normative. Andersen indicates that the assumption of a certain risk attitude (perhaps even risk neutrality) may be expositionally useful,  but in general risk attitude should be jointly estimated within the specified structural model (Andersen, Harrison et al. 2006). An exception appears to be Ioannidis who specifies a modified prospect theory model where institutional *loss aversion* (not necessarily risk aversion), for example, may be reasonably assumed but which explicitly parameterizes risk attitudes over gains versus losses. Other theoretical aspects of the PT specification (reference dependence and probability weighting) may also be reasonably salient within the security context and so the associated structural model should ideally be constructed and tested for variations in these factors, not their a priori assumption or omission. This approach is consistent with Harrison's view to the experimental elicitation and modeling of risk attitudes under uncertainty:

> Two broad methods of estimating risk attitudes have been used. One involves the calculation of bounds implied by observed choices, typically using utility functions which only have a single parameter to be inferred. A major limitation of this approach is that it restricts the analyst to utility functions that can characterize risk attitudes using one parameter. This is because one must infer the bounds that make the subject indifferent between the switch points, and such inferences become virtually incoherent statistically when there are two or more parameters. Of course, for popular [single parameter] functions such as CRRA or Constant Absolute Risk Aversion (CARA) this is not an issue, but if one wants to move beyond those functions then there are problems. It is possible to devise one-parameter functional forms with more flexibility than CRRA or CARA in some dimension, as illustrated nicely by the one-parameter Expo-Power function developed by Abdellaoui, Barrios and Wakker [2007; §4]. But in general we need to move to structural modeling with maximum likelihood to accommodate richer models. (Harrison, Lau et al. 2010)

In the context of Beautement, this would involve, for example, utility functions which infer some type of discount rate which is hypothesized to be separate from risk aversion. Harrison emphasizes the connection between theory and method within lab experiments which illustrates the endogenous nature of the parameter estimation, here in the context of a choice over expected income $M_1$ in time *t*, versus $M_2$ in time *t+n*, where an n-period discount rate is assumed:

> As one relaxes the assumption that the decision maker has a linear [risk neutral] utility function, it is apparent… that the implied discount rate decreases if U(M) is concave in M. Thus one cannot infer the level of the individual's discount rate without knowing or assuming something about their utility function. This identification problem implies that discount rates cannot be estimated based on discount rate experiments with choices defined solely over time-dated money flows, and that separate tasks to identify the extent of diminishing marginal utility must also be implemented.

> …Thus there is a clear implication from theory to experimental design: you need to know the non-linearity of the utility function before you can *conceptually* define the discount rate. There is also a clear implication for econometric method: you need to jointly estimate the parameters of the utility function and the discount rate, to ensure that sampling errors in one propagate correctly to sampling errors of the other. In other words, if we know the parameters of the utility function less precisely, due to small samples or poor parametric specifications, we have to use methods that reflect the effect of that imprecision on our estimates of discount rates. (Harrison, Lau et al. 2010)

Following Harrison, I therefore propose that the risk factors within the security control selection context should be estimated using a structural model of the latent choice processes in which the core parameters defining risk attitudes are *jointly* estimated.

**Use of Binary Lotteries for Estimating Decisional Bias**

As detailed in Section 4 below, I am proposing to generate data using a series of lab experiments requiring participants to make individual choices over what are essentially *binary lotteries*, representing two probabilistic outcome states which are characterized by known, uncertain or mixed outcome probabilities. The core hypothesis is that the choices made by the respondent consistently reflect and can therefore be used to model a) the respondent's degree of risk aversion; and 2) other decisional biases, based on the degree to which they are statistically observed to select one lottery of a pair over the other, in the context of the information available to the participant in each experiment. Respondents consistently (but perhaps not perfectly consistently) choosing a lower risk, lower *value* lottery over another higher expected value but risky lottery is, for example, hypothesized to prefer the 'certainty equivalent' of the lower value lottery versus the higher expected value (but less certain outcome) of the alternative lottery (Hershey and Schoemaker 1985). The following example lottery pair in Table 1 indicates the prototypical choice problem with known probabilities, or risk, facing the participant. The participant must pick either the 'left' or 'right' lottery. In this case, the Left lottery pays $50 for sure, while the Right Lottery has a 50% chance of paying out $40, and a 50% chance of paying out $80:

**Table 2 - Binary Lottery Choice Under Risk (Known Probabilities)**

| | **Left Lottery** | **Right Lottery** | |
|---|---|---|---|
| **Probability** | 100% | 50% | 50% |
| **Payout** | $50 | $80 | $40 |
| **Expected Payout** | $50 | $40 | $20 |
| | | $60 | |

In this simplified example under single play ('one shot game') conditions, a hypothesized risk neutral person would choose the Right Lottery since the expected value of the payout is higher than that of the Left Lottery. However a risk averse person is hypothesized to prefer the Left (L) lottery, where the guarantee of winning $50 represents a 'certainty equivalent' compared to the expected, but risky, payout of the Right (R) Lottery. Under Expected Utility Theory (EUT), the lower the observed choice of a certainty equivalent for a given expected value outcome, the greater is the hypothesized risk aversion of the participant. For the purposes of estimating risk aversion, it is however necessary to repeat the game in order to generate enough observations to make statistical inferences about the degree of risk aversion. Theory demands that we make assumptions (or, specifically, relax assumptions) regarding either or both the participant's understanding of the particular game and of the nature of their risk attitude and preference over one shot vs. 'compound lotteries'. For example, we assume that the decision maker understands the mechanics of the lottery payouts enough to be consistent in their decisions given their latent risk preference. That is, that the participant realizes that, in a one shot game under the Right lottery, the expected value ($60) is never actually realized and that they are essentially choosing between 3 discrete outcomes: a sure $50, or an unsure $80 vs. $40. Conversely, in a repeated game context, we assume that the participant understands that the long run average (expected) Right Hand payout is absolutely larger than that of the Left lottery and, all else considered equal, their risk aversion should diminish towards risk neutrality with the increasing length of the game. It is also generally assumed that, in an experiment consisting of a series of one shot games in which one of the chosen lottery pairs is actually chosen at random to be played out for money, the decision maker effectively 'reduces' the series of potential lottery outcomes to a single lottery and that their preference across the sub-lotteries is consistent with their overall risk attitude. A discussion of the assumptions behind this 'reduction of compound lotteries' and its potential impact on the design and outcome of the experiments will be presented below (Harrison, Martínez-Correa et al. 2015).

Under the assumption of expected utility maximization therefore, the elicitation of attitudes over repeated lottery choices using various levels of certainty equivalents permits the estimation of the respondent's 'coefficient of risk version' That is, given enough observed choices, we should be able to estimate the *likelihood of observing the actual choices made* given a) utility dependent upon a value of risk aversion *r* and b) other parameters of an *assumed latent structural model of choice* specified using a given utility model. Generically, following Harrison, this is done econometrically by creating an 'index' (denoted here by '∇') of the difference in the expected utilities of each lottery pair, assuming an initial level of risk aversion:

$$\nabla EU = (EU_R - EU_L) \tag{3.9}$$

where EU is a functional specification of expected utility dependent on a risk aversion parameter *r*. For example,

$$U(x) = x^{1-r} / 1\text{-}r \tag{3.10}$$

where "*r*" is the 'coefficient of risk aversion' and "x" is the monetary value of the outcome, where *r*=0 corresponds to risk neutrality, *r*<0 to risk loving, and *r*>1 to risk aversion. The estimated value for *r* can be benchmarked against known values for *r* in general populations (Hey and Orme 1994; Antoniou, Harrison et al. 2010) and used in subsequent experiments either as a constant to estimate subjective probabilities, or compared to the value of *r* when, for example, jointly estimated with subjective probabilities to test whether *r* changes under different utility preference assumptions.

The resulting index (which is scaled from –∞ to +∞) is then used to define the *cumulative probability* of the observed choice (right or left lottery) using the cumulative normal distribution function Φ(.). The agent chooses lottery 'R' if

$$\nabla \text{ EU} + \varepsilon > 0 \tag{3.11}$$

Where ε is a normally distributed error term with mean zero and variance $\sigma^2$. Thus we can establish a 'probit link function', showing the probability that lottery R is chosen given a value of the index, as

$$\text{prob(choose R)} = \text{P} (\nabla \text{ EU} + \varepsilon) > 0 = \text{P}(\varepsilon / \sigma > -\nabla \text{ EU} / \sigma) = \Phi(\nabla\text{EU} / \sigma) \tag{3.12}$$

The figure below illustrates the 'linking' of the index value y* (the difference in expected value between the right and left lotteries) to the probability of observing the index value using both a probit (cumulative normal) and a logit (logistic) specification:

**Figure 1 - Normal and Logistic Cumulative Density Functions**

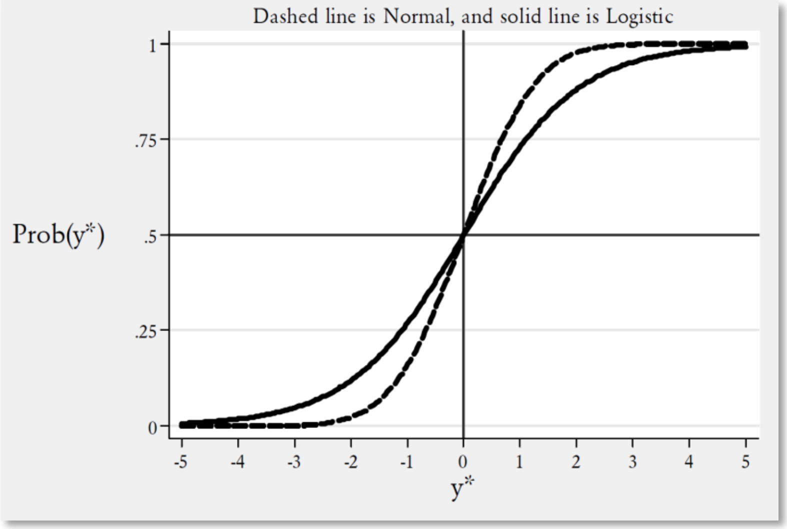

(Andersen, Harrison et al. 2007)

The Probit functional specification with x representing the index is

$$\text{F (x; } \mu, \sigma^2) = \Phi\left(\frac{x-\mu}{\sigma}\right) = \frac{1}{2}\left[1 + \text{erf}\left(\frac{x-\mu}{\sigma\sqrt{2}}\right)\right], x \in \mathbb{R} \tag{3.12}$$

where “erf” is an error function $\frac{x-\mu}{\sigma\sqrt{2}}$ for positive *x* values. The corresponding *logit* functional specification would be

$$G(\nabla\,EU) = \Lambda(\nabla\,EU) = \exp(\nabla\,EU)\,/\,[1 + \exp(\nabla\,EU)] \qquad (3.13)$$

By assuming different values of risk aversion, we can then iteratively estimate the linking function to determine statistically that value of *r* which would maximize the likelihood of observing the actual choices made (Pratt 1976).[20]

The estimation of *r* is assumed to *structurally* depend on the latent characteristics of the decision maker which are hypothesized to influence her degree of risk aversion. This also permits pooling of the data across subjects where otherwise *r* would only be estimable per subject. This approach is important to this research since we are fundamentally interested in whether and how an individual decision maker’s risk attitude may affect their control choices even in the presence of objective or ‘good’ information about prospective security outcomes. As noted by Andersen,

> It is [a] simple matter to generalize this ML [maximum likelihood] analysis to allow the core parameter r to be a linear function of observable characteristics of the individual or task. For example, assume that we collected information on the sex of the subject, and coded this as a binary dummy variable called Female. In this case we extend the model to be $r = r_0 + r_1 \times$ Female, where $r_0$ and $r_1$ are now the parameters to be estimated. In effect the prior model was to assume $r = r_0$ and just estimate $r_0$. This extension significantly enhances the attraction of structural ML estimation, particularly for responses pooled over different subjects, since one can condition estimates on observable characteristics of the task or subject.(Andersen, Harrison et al. 2007)

The *conditional* log-likelihood to be estimated using maximum likelihood methods, assuming a vector of demographic characteristics X would then be

$$\ln L(r\,; y, \mathbf{X}) = \sum_i \left[ (\ln \Phi(\nabla EU) \times \mathbf{I}(y_i = 1)) + \left(\ln\left(1 - \Phi(\nabla EU)\right) \times \mathbf{I}(y_i = -1)\right) \right] \qquad (3.14)$$

[20] “In non-technical parlance, "likelihood" is usually a synonym for "probability" but in statistical usage, a clear technical distinction is made. One may ask "If I were to flip a fair coin 100 times, what is the *probability* of it landing heads-up every time?" or "Given that I have flipped a coin 100 times and it has landed heads-up 100 times, what is the *likelihood* that the coin is fair?" but it would be improper to switch "likelihood" and "probability" in the two sentences. If a probability distribution depends on a parameter, one may on the one hand consider—for a given value of the parameter—the probability (density) of the different outcomes, and on the other hand consider—for a given outcome—the probability (density) this outcome has occurred for different values of the parameter. The first approach interprets the probability distribution as a function of the outcome, given a fixed parameter value, while the second interprets it as a function of the parameter, given a fixed outcome. In the latter case the function is called the "likelihood function" of the parameter, and indicates how likely a parameter value is in light of the observed outcome.” http://en.wikipedia.org/wiki/Likelihood_function

where **I**(.) is the indicator function (i.e. indicating choice of Right versus Left, $y_i$ =1(-1) denotes the choice of the Option R (L) lottery in choice task *i*. In this specification, **X** is a vector of individual characteristics reflecting age, sex, race, and so on. The parameter *r* is thus defined as a linear function of the characteristics in vector **X**. (Harrison and Rutström 2009)

The figure below illustrates the resulting simple estimate of *r* as a constant (i.e. not a function of X) graphically for Harrison's replication of the seminal Hey and Orme binary choice study (Hey and Orme 1994), using 158 subjects making 60 choices each for over 9,000 observations. The maximum likelihood estimate of *r* assuming a power utility function is .777 with a std. error = .025, p<0.001 and is illustrated by the resulting 'kernel density function' for *r* across all participants:

**Figure 2 - Example Kernel Density Function of Risk Attitudes**

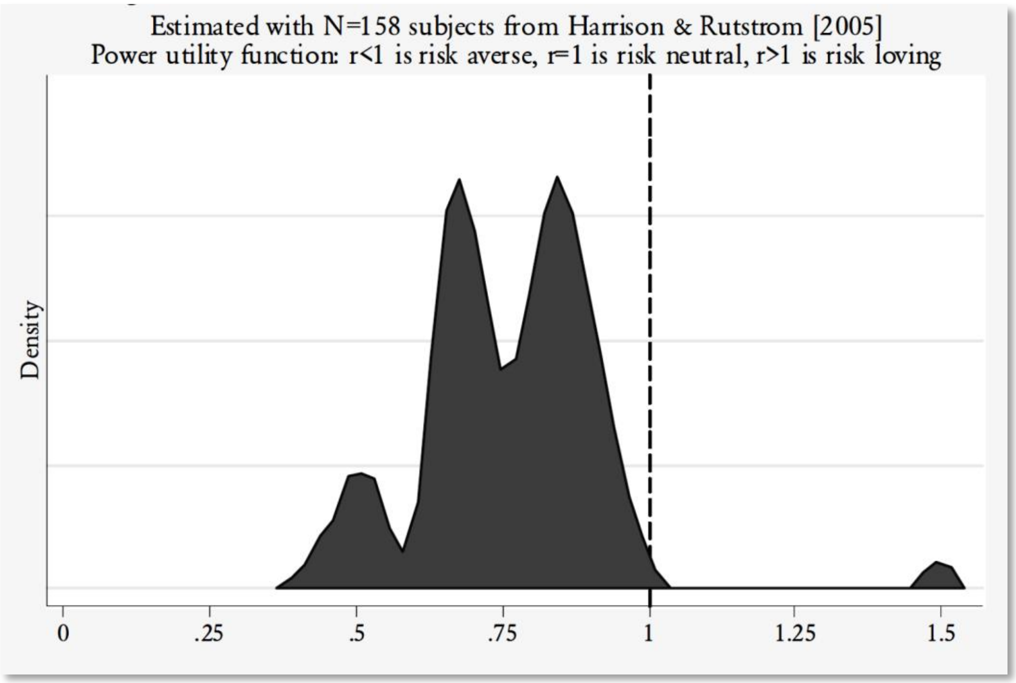

(Andersen, Harrison et al. 2007)

In the context of the power function estimate for *r*, the experiment concludes that participants are relatively risk averse in lotteries involving prospective gains.  It should be noted that there are some observations in which the inferred risk attitude would be 'risk seeking', but not many, and for lotteries defined in terms of prospective gains this finding is consistent with most research outcomes (Harrison and Rutström 2008). In this example the prospects were gains, but the experiment can be replicated using losses, or mixed outcomes that allow the ability to test for risk aversion over both gains and losses. In a real world security context typically defined in terms of prospective 'business losses' (whether actual transactional losses or attributed losses due to the inability to transact due to lack of system availability or integrity, or attributed loss of information confidentiality), we will need to take care in making assumptions about how the decision maker 'frames' the choice and whether, when they do so, the choice is actually perceived as a marginal loss. Kahneman's simple example of two symmetric choice lotteries "A": zero stake with a prospective gain of $100 (either through a $100 sure win or a 50/50 prospect of $200 or nothing) and "B": a stake of $200 with a sure loss of $100 or a 50/50 loss of $200 or nothing), each producing the same expected 'gain' or 'loss' depending on the wealth reference point and whether the observer frames the outcome as a gain or a loss, illustrates how the perception of marginal loss or gain can change based on the reference point of the decision maker (Tversky and Kahneman 1992). I will explicitly test for differences in this 'framing effect' in Game #1, and particularly whether the decision maker is subject to 'asset integration' i.e. decision from a wealth perspective of cumulative gains or losses over time.

### Handling Model Errors

Harrison emphasizes the need for careful handling of the linking function error term noted above. This is because the structural choice model explicitly makes the statistical assumption that the probability of choosing a lottery is not 1 when the EU of that lottery exceeds the EU of the other lottery. Indeed, if there were no errors from the perspective of EUT (i.e. if the subject was in fact behaving perfectly according to EUT), this function would be a step function, as in Figure 3: zero for all values of y*<0, anywhere between 0 and 1 (i.e. indifference) for y*=0, and 1 for all values of y*>0:

**Figure 3 - Hardnose Theorist Cumulative Density Function**

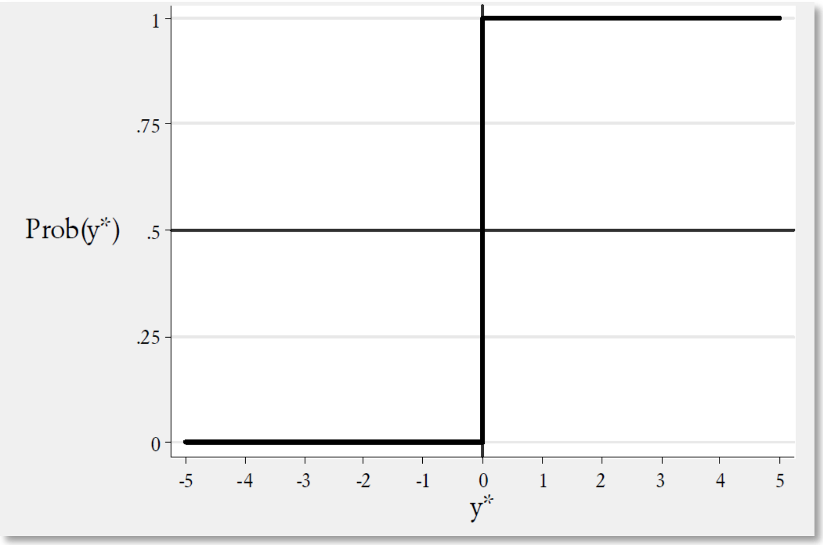

(Fréchette and Schotter 2015)

A popular form of incorporating structural error into the model according to Harrison is the *Fechner* error term:

$$\nabla EU = (EU_R - EU_L)/\mu \tag{3.15}$$

where μ is a structural "noise parameter" used to allow some errors from the perspective of the deterministic EUT model (Harrison and Rutström 2008). If we assume no error in a subject's ability to distinguish the higher EU lottery choice, then μ would be assumed to have a value of 1. If this is not assumed, then we must allow for this error in the structural model of choice and jointly estimate the error along with the other parameters of the model. Conceptually, the error term in this specification takes into account a number of potential 'errors' associated with not only the choices but the choice task itself and has been discussed at length by Smith (Smith 2010) and commented on by Harrison (Harrison 2010). This is a crucial caution from the perspective that, in moving from the theory to the empirical testing of theory, we are making inferences about an appropriate model which will in fact represent the theory under consideration. The insight here is that we cannot infer risk aversion, or other bias, using tests of binary choice, with a model that assumes away the possibility of error in the test methodology. The specification of the structural error term is therefore important for making any subsequent statements about decision bias which are clearly conditional on the model assumed to elicit the risk attitude of the subject. According to Harrison,

> "…some specifications place the error at the final choice between one lottery or after the subject has decided which one has the higher expected utility; some place the error earlier, on the comparison of preferences leading to the choice; and some place the error even earlier, on the determination of the expected utility of each lottery …In short, *one cannot divorce the job of the theorist from the job of the econometrician*, and some assumption about the process linking latent preferences and observed choices is needed. That assumption might be about the mathematical form of the link, but it cannot be avoided. Even the very definition of risk aversion needs to be specified using stochastic terms unless we are to impose absurd economic properties on estimates (Wilcox 2008; Wilcox 2010)" (Harrison, Lau et al. 2010).

The error term is also potentially the repository for deeper theoretical issues concerning the nature of risk preference itself as represented by the concept of risk 'aversion' and its implication of subjective lower utility over risky prospects i.e. concave Bernoulli utility functions (Harrison 2010). In a contrarian view, Friedman and Isaac have recently undertaken and extensive review of the lack of resulting prediction power associated with the assumption of non-linear individual utility curves and the inference of what is supposed to be latent in the specification of a utility curve itself: risk aversion (Friedman, Isaac et al. 2014).

**Mixed Structural Model Specifications**

Harrison advocates the specification of *mixed structural utility models* involving both EUT and PT latent choice models, particularly in situations where there is theoretical reason to believe that more than one latent decision making process may be present in the generation of choices within and across decision makers (Harrison and Rutström 2009). In the domain of information security, it may be plausible, for example, to assume that professional security managers are predominantly *loss averse* (the strong preference for avoiding losses as opposed to acquiring gains[21]) and that the specified utility model should therefore structurally allow for *loss* aversion and not simply risk aversion; or that subjects may be allowing for probability weighting when perceiving risky or uncertain prospects i.e. the overweighting of low probabilities and the underweighting of high probabilities (Kahneman 1979). Incorporating loss aversion would then involve using a utility model based on prospect theory. Following the example above, a power utility function can be defined separately over gains and losses:

$$U(x) = x^{\alpha} \text{ if } x \geq 0, \text{ and} \tag{3.15}$$

$$U(x) = -\lambda(-x)^{\beta} \text{ for } x < 0 \tag{3.16}$$

In the above equation $\alpha$ and $\beta$ are the risk aversion parameters, and $\lambda$ is the coefficient of loss aversion. We can see immediately that the assumption of this model either invalidates the assumption of the EUT specification (since *r* in EUT actually depends on whether the choice is a gain or a loss in utility) or, per Harrison, requires us to allow the structural model to incorporate both latent choice theories. "Cumulative PT" (CPT) also incorporates *probability weighting* as noted above. The form of the weighting function proposed by Tversky and Kahneman (Tversky and Kahneman 1992) assumes weights

$$w(p) = p^{\gamma}/[\, p^{\gamma} + (1-p)^{\gamma}\, ]^{1/\gamma} \tag{3.17}$$

The PT utility function can be used instead of the EUT utility function, and w(p) is used instead of p, but the steps are otherwise identical to that described above for EUT. The same 'error' process is assumed to apply when the subject forms a PT-based preference for one lottery over the other. The difference in prospective utilities is defined similarly as

$$\nabla PU = PU_R - PU_L \tag{3.18}$$

where $\nabla$PU is an analogous index of the difference in the expected values of the choices. Thus the likelihood, conditional on the PT model being true, depends on the estimates of $\alpha$, $\beta$, $\lambda$ and $\gamma$ given the above specification and observed choices. Similarly, the conditional log-likelihood is

[21] We recognize that a security control decision can be framed as a gain – or at least a 'net gain' – problem instead of a loss avoidance scenario, where it is recognized that security controls, like IT systems generally, are implemented to support business goals which fundamentally involve generating income for the business (where no business is in business to simply 'avoid losses'). This framing is more obvious with certain types of controls (e.g. authentication swipe cards) that have some direct user benefits such as improving speed to log onto systems for which positive business value can be clearly attributed but, we argue, is not the typical frame of reference for a security practitioner. Admittedly, the focus on 'business value' is an increasingly common frame of reference for all IT systems design and implementation and for 'IT governance' and risk management generally. A separate study would be needed to compare control decisions made in terms of gains vs. those framed here as pure loss and is discussed in the section on Future Research.

$$\ln L^{PT}(\alpha, \beta, \lambda, \gamma\,; y, X) = \sum_i l_i^{PT} = \sum_i [(\ln \mathrm{G}(\nabla \mathrm{PU}) \mid y_i = 1) + (\ln(1 - \mathrm{G}(\nabla \mathrm{PU})) \mid y_i = 0)] \qquad (3.19)$$

If we let $\pi^{\mathrm{EUT}}$ denote the probability that the EUT model is correct, and $\pi^{\mathrm{PT}} = (1\text{-}\pi^{\mathrm{EUT}})$ denote the probability that the PT model is correct, a 'grand likelihood' function can then be written as the probability weighted average of the conditional likelihoods. Thus the likelihood for the overall model estimated is defined by

$$\ln L(r, \alpha, \beta, \lambda, \gamma, \pi^{\mathrm{EUT}}; y, X) = \sum_i ln[(\pi^{\mathrm{EUT}} \times l_i^{\mathrm{EUT}}) + (\pi^{\mathrm{PT}} \times l_i^{\mathrm{PT}})] \qquad (3.20)$$

This log-likelihood can be econometrically maximized using observed choice data to find estimates of the specified parameters.

**Model and Hypothesis Testing**

We are able to test the estimated results of each structural model in three ways and apply the results to our hypotheses. First, standard sign and confidence interval tests can be applied to most parameter estimates. Second, we consider the parameter estimates of the various specifications in the context of the hypotheses being tested. This would include (variously):

- For EUT specifications, the presence and strength of risk aversion in both income and non-income utility specifications.
- For PT specifications, that the risk aversion over gains versus losses is different, and that the participants exhibit loss aversion and probability weighting. For PT results generally, the results can also be compared to benchmarks from established studies (Tversky and Kahneman 1992)
- For mixed EUT/PT models, that the weighting between the models is different from .5 (i.e. that one data generating process dominates the other) and whether and to what extent the other parameter estimates change from the pure EUT/PT specifications. The distribution of parameter estimates can also be examined graphically using *kernel density* diagrams (essentially smoothed histograms) to illustrate multi-modal distributions of parameter effects across and within subjects.
- Whether and to what extent professional and demographic factors affect the various parameter estimates. This allows us to test for the heterogeneity of risk attitudes within the sample and therefore help to characterize classes of decision makers and their biases based on those factors (e.g. male, math degree, more than 10 years in the business, etc.).
- For subjective utility specifications, the difference between estimation of subjective probabilities using assumed versus jointly estimated risk attitudes, and the effect of strength and weight on the subjective probabilities.

Third, we can employ a range of ‘goodness of fit’ tests for the non-nested (separately specified and estimated) EUT and PT models and ‘likelihood-ratio’ tests of one model against another. STATA produces a ‘Wald Chi-squared statistic’ which approximates an F-statistic type of test on the overall model fit based on the optimized maximum likelihood value derived for the model. ‘Likelihood ratio’ tests may then be used to compare the likelihood values from two estimated models (Quinn 2011). Harrison reviews the seminal test literature but emphasizes that mixed model specification inherently embeds the testing of one model against another by virtue of the specification of model weighting parameters and is considered superior to non-nested specifications since non-nested models implicitly assume only one data generating process:

> An issue that comes up when estimating choice models is how to discriminate between them. The answer we prefer is “not to,” in the sense that one should allow the data-generating process to admit more than one choice model and estimate a mixture model…But how does [the mixed model specification] response compare to more classical responses, using formal tests to discriminate between models?
>
> The first [non-nested test] is due to Cox (Cox 1961; Cox 1962), and assumes that one of two models is the true model for all data. His test compares the difference between the actual likelihood ratio of the two models with the expected likelihood ratio, suitably normalized by the variance of that difference, under the hypothesis that one of the models is the true data generating process. The statistic is applied symmetrically to both models, in the sense that each takes a turn at being the true model, and leads to one of four conclusions: one model is the true model, the other model is the true model, neither model is true, or both models are true…The ambiguity in the resulting inference has bothered many who might otherwise use the test.
>
> The next step in the statistical literature was the development by Atkinson (Atkinson 1970)…The main problem with this exposition, noted by virtually every commentator in the ensuing discussion, was the interpretation of [a] mixing parameter. Atkinson focused on testing the hypothesis that this parameter equaled ½, “which implies that both models fit the data equally well, or equally badly.” There is a colloquial sense in which this is a correct interpretation, but it can easily lead to confusion if one maintains the hypothesis that there is only one true data generating process, as the commentators do. **In that case one is indeed confusing model specification tests with model selection tests** [emphasis added]. If instead the possibility that there are two data generating processes is allowed, then natural interpretations of tests of this kind arise.
>
> Perhaps the most popular modern variant of the generalized LRT approach of Cox is due to Vuong (Vuong 1989). He proposes the null hypothesis that both models are the true models, and then allows two one-sided alternative hypotheses. The statistic he derives takes observation-specific ratios of the likelihoods under each model, so that …the ratio for observation *i* is the likelihood of observation *i* under EUT divided by the likelihood of observation *i* under PT. It then calculates the log of these ratios, and tests whether the expected value of these log-ratios over the sample is zero. Under reasonably general conditions a normalized version of this statistic is distributed according to the standard normal, allowing test criteria to be developed. Thus the resulting statistic typically provides evidence in favor of one of the models which may or may not be statistically significant.
>
> The third test is due to Clarke (Clarke 2007), and is similar in motivation to the Vuong test but ends up using a non-parametric sign test to discriminate between the models. The Vuong and Clarke tests can be compared in terms of their asymptotic efficiency as tests, and it turns out that when the distribution of the log of the likelihood ratios is normally distributed that the Vuong test is better. But if this distribution exhibits sharp peaks, in the sense that it is mesokurtic, then the Clarke test is better... Many of the likelihood ratios we deal with have the latter shape, as we will see using the EUT and PT example from above, so it is useful to be armed with both tests. (Harrison 2008)

I undertake these various non-nested model tests in the context of the respectively applicable model scenarios in Game #2.

# 4 – Research Design

### Overview

Using recent information system simulation work (Baldwin, Mont et al. 2009; Fenz, Tjoa et al. 2009; Collinson, Monahan et al. 2010; Tjoa, Jakoubi et al. 2011; Sommestad, Mathias et al. 2013; Holm, Shahzad et al. 2015), together with guidance on utility models appropriate to the modeling of preferences over security control choices (Beautement, Coles et al. 2009; Beresnevichiene, Pym et al. 2010), and using a modified version of both Antoniou and Harrison's (Antoniou, Harrison et al. 2010)[22] replication of Griffin and Tversky's multiple price list lab experiment (Griffin and Tversky 1992) with additional multi-period decision features which incorporate asset integration (Andersen, Harrison et al. 2006) and endogeneity into both the experiment and the mixed EUT/PT model framework (Shogren and Crocker 1994; Fiore, Harrison et al. 2009; Sen 2010), I propose to implement a multi-part experiment to elicit risk attitudes where the perceived risk is stochastic but may also be endogenously controllable by the subject over multiple decision periods by making successive risk control 'wagers' to affect either or both likelihood (using 'self-protection') or impact (using insurance, including 'self-insurance' or internal resources reserves to be used in case of actual breach). Individual decisions are hypothesized to be based on individual subjective perceptions of whether and to what extent the simulated system appears to be at risk and the perceived prospective efficacy of the proposed control option(s). In the experiment, 'successful' control wagers and outcomes inherently depend on the respondent's ability to update their subjective prior estimates of the apparent system risks over a one shot or series of selection rounds which vary by both the strength (i.e. the perceived influence of 'extreme' risk factor observations,) and weight (i.e. the number of data points making up a factor observation)[23] and where the latent decision making model is allowed to vary between EUT and PT utility specifications (Harrison and Rutström 2009) permitting probability weighting and loss aversion (Tversky and Kahneman 1992). The proposed experiments generate robust data on both each respondent's risk attitude and their subjective beliefs of the actual risk profile of the simulated system at risk which then enables the modeling and testing of the presence of various decisional biases within a relevant context of professional interest to the subjects.

### Experimental Research Methodologies for Behavioural Economics Studies

'Behavioural decision research' involving lab experiments is a well-established branch of psychological research and its methods have, more recently, been increasingly represented in theoretical, and empirical economic studies of decision making under conditions of both risk and uncertainty (Savage 1954; Machina 1992; Abdellaoui, Baillon et al. 2008; Cox and Sadiraj 2008; Harrison and Rutström 2008; Wilcox 2008; Andersen, Fountain et al. 2009). Early psychological studies focused on qualitative approaches including

[22] See pp. 138-191. Antoniou's doctoral dissertation experiment is used to incorporate subjective decisional bias into asset pricing theories in order to inform the behaviour of financial asset prices. The balance of Antoniou's thesis addresses other issues within behavioural finance which also look specifically at non-Bayesian updating behaviours, but my focus will be on his experiment which specifically replicates the Griffin and Tversky study (Griffin and Tversky 1992).

[23] Ibid. pp. 144-145

normative analysis, descriptive studies, and prescriptive interventions (Edwards 1954; Simon 1955; Edwards 1961). Early contributions by economists (Allais 1953; Ellsberg 1961; Savage 1971) took the established economic theories of rational decision making into the lab and their findings formed the basis of the modern economic considerations of risk and uncertainty in decision making. Herbert Simon, who won the Nobel prize in 1978 for his work on decision making within organizations, pioneered the theory of 'bounded rationality' which changed the dominant economic paradigm from *homo economicus* rational utility maximization to that of individual 'satisfaction seeking' through the use of heuristics (Simon 1955). This was also the basis for Kahneman and Tyversky's work (earning these psychologists the Nobel Prize in Economics in 2002) which identified, through experimental means, the presence of asymmetric risk attitudes in decisions made under conditions of outcome uncertainty (Kahneman and Tversky 1972; Kahneman 1979; Kahneman, Knetsch et al. 1990). Griffin and Tversky's seminal psychological experiments also demonstrated the possibility of non-Bayesian updating by subjects (Griffin and Tversky 1992; Holt 2002) and, as noted below, continues to be the basis for experimental economic studies which both support and refute their findings depending on the experimental controls and models employed by the researcher (Antoniou 2010).

The concept of bounded rationality and the possibility of alternatives to unbiased expected utility maximization has therefore formed the basis of much subsequent theoretical and empirical work in economic decision making and continues to inform an emergent body of experimental research under the banners of both 'empirical decision research' and 'experimental economics'. Cox (Cox and Sadiraj 2008) surveyed five essential decisional theories and the associated utility models and testing approaches: expected value theory (Bernoulli 1738 (1954)); expected utility theory (von Neumann 1947); cumulative prospect theory (Tversky and Kahneman 1992); rank dependent utility theory (Quiggin 1982) and dual theory of expected utility (Yaari 1987). The motivation for Cox's work is to attempt to "develop a distinction between "normative" and "descriptive" theories of choice…and their common problems when viewed as 'positive (that is, testable) theories" – i.e. their inability to predict the actual behaviour of experimentally tested subjects without accounting for risk aversion. The most recent experimental economic approaches are therefore explicitly designed to assist in the modeling, parameter estimation and testing of various decisional utility models and involve lab experiments which generate robust data derived from subject choices concerning systems at risk. (Holt and Smith 2007; Fiore, Harrison et al. 2009; Antoniou, Harrison et al. 2010; Sen 2010)).

Seminal experimental methods for measuring subjective economic utility and bias include Becker, DeGroot, and Marschak (Becker, Degroot et al. 1964) and Savage (Savage 1971) who established experimental wagering schemes to infer personal (subjective) probability estimates over uncertain prospects. Experimental economists continue to use variations on these approaches employing a range of devices including paired choices (including binary lotteries), multiple price lists, auctions, multi-party bargaining or similar schemas to elicit subjective risk attitudes and to test for decisional bias in situations involving uncertain probabilities and payoffs (Machina and Schmeidler 1992; Hey and Orme 1994; Holt 2002). Vernon Smith originally developed the concept of ‘induced value’ – the elicitation of revealed utilities or values within an experimental ‘game’ - in which motivated subjects could be made to reveal their preferences over outcomes under conditions which replicated aspects of real world markets, but in a laboratory setting (Smith 1976; Smith 1982; Smith 2010). Smith’s experimental work built on Leonid Hurwicz’ earlier work in ‘mechanism design’ (for which Hurwicz eventually won a Nobel prize in 1997 as the oldest recipient of the award at the time) whereby economic systems were modeled as games within a laboratory setting. Smith also established a formalized empirical approach to laboratory experiments directed at economic decision makers which could be used to observe behaviour under controlled settings and employ methods designed to quantify and test traditional models of rational utility maximization. Grether and Plott (Grether and Plott 1979) and Grether (Grether 1980; Grether 1992) are credited with developing seminal experimental economic methods for validating the psychological findings of Griffin and Tversky, including the degree to which the subject may be violating Bayes rules by employing heuristics within the decisional process.

Variants of these foundational experimental methods continue to be employed by economic researchers with distinct changes to the experimental design and the employed utility maximization models to be estimated depending on the context of the decision problem, the nature of the latent decision making process(es) assumed and the specific decisional bias which is under examination and are described at length in Harrison and Rutström’s 2008 survey of lab experiments for risk aversion (Harrison and Rutström 2008). While earlier seminal surveys noted by Harrison reviewed methods employed for the elicitation of expected utility (Fishburn 1967; Farquhar 1984), Harrison and Antoniou (Antoniou 2010) also note that early experimental methods often made inappropriate risk aversion assumptions (i.e. risk neutrality) or failed to take into account the need for appropriate experimental controls which would ensure accurate measurement of preferences, for example failing to provide monetary incentives to properly induce subjects to make bets consistent with their presumed latent risk attitudes and preferences and their ‘true’ subjective probability estimates  (i.e. participants have to be betting with and for real money outcomes, or generally incented to make non-arbitrary bets which reliably reveal personal probability estimates (Holt 2002)). Criticism of the risk attitude elicitation procedures undertaken taken by Griffin and Tversky are well documented in the economics literature and include: the lack of appropriate financial incentives for the decision maker; the likely presence of ‘hypothetical bias’[24]; and the study’s assumption of ‘risk neutrality’ (Antoniou 2010). The appropriate design and employment of experimental controls to compensate for these issues is an important consideration for researchers undertaking these types of experimental studies and are well

[24] For example, the explicit testing of ‘hypothetical bias’ in lab conditions where probabilistic event outcomes are only described but not demonstrated to the subject, say by describing a biased coin flip or dice roll instead of actually flipping a biased coin or rolling actual dice.

documented in the relevant experimental economics literature (Antoniou, Harrison et al. 2010; Harrison, Lau et al. 2010).

The data gathering experiments typically feature subjects making monetary bets over risky (known probability) or uncertain (unknown probability) prospective outcomes in the context of, variously, *binary lotteries* (Grether 1992), *multiple price lists* (Holt 2002; Harrison, Johnson et al. 2005) and, in two person experiments testing for Nash equilibria, *first price sealed bid auctions* (Vickrey 1961; Harrison 1989; Harrison 1990). The resulting data from the experiments may be used to quantitatively model and test theories of decision making both within the context of the experiment and, possibly, to predict 'out-of-context' behaviour under certain specified conditions (i.e. outside of the experimental frame, choice types or tasks and range of values over which decisions are made) (Rabin and Thaler 2001; Cox and Sadiraj 2008; Wilcox 2010).

In summary, recent experimental economics research has employed experiments consisting of 'binary 'choices', under conditions characterized by both risk and uncertainty, to generate choice data and to then undertake the econometric estimation of stochastic, mixed structural models of latent decision-making processes which allow for the specification, joint estimation and weighting over n>1 decision making processes and of the associated parameters of those processes (Lopes and Oden 1999; Andersen, Harrison et al. 2006). At the same time, recent advances in the mathematical modeling of information systems supports the construction of increasingly realistic *simulations* of those systems under conditions which can be characterized in terms of risk and uncertainty which can be employed to provide the quantitative inputs and simulated outcomes over which experiments involving decisions, binary choice or otherwise, regarding system control can be designed. The combination of these two fundamental constructs – quantitative decision scenarios based on stochastic system simulation and established formats for decisional experiments based on simulated quantitative choices - permits the testing of both established and innovative hypotheses within the decision research domain to proceed within the specific context of information security decision making under uncertainty.

The following two sections describe 1) the development of the system simulator which operates on inputs and outcomes which are readily recognizable by security practitioners; and 2) the development of experiments which use are based on well documented choice experiments but which have been translated into the domain specific context of information security control choice.

### Overview of the Use of Simulation in the Context of Binary Choice Experiments

The format of lab experiments for decision research have been typically based on the use of very simple physical systems including pre-calculated pie-chart lotteries, the use of dice, coloured balls in transparent urns, bingo cages with numbered or coloured balls or other stylized 'systems' at risk, aspects of which the participant can directly observe and able to represent both risky and uncertain probabilistic prospects. These systems are used to represent or dynamically generate decision scenarios within which participants are asked to then make decisions regarding preferences over binary choices involving the prospective outcomes based on the choices they make. In some cases the participant may be paid to play and the outcome of the played out choice is paid to the participant in some form of either points or money. In the example given above, the participant is asked to choose between two pre-determined 'binary' lotteries in which the participant indicates their preference for one lottery over the other and the chosen lottery is then 'played out' to determine the payoff for that choice. In order to generate sufficient choice data to estimate the decisional parameter of interest (say risk aversion), the participant may make a series of independent choices over several rounds and then one of the choice pairs is randomly selected to be 'played out' for a payoff to the participant. Variations on this format may, alternatively, involve a 'multiple price list' format indicating at what 'price' the player would choose one prospect over the other. In both of these examples, the 'crossover point' between choosing one lottery over the other is used to econometrically determine degree to which the participant prefers a 'sure thing' to a relatively risky prospect.

The 'lotteries' can involve both known and unknown probabilities depending on what aspect of decisional bias are being tested. For example, seminal experiments regarding risk aversion typically used a series of paired binary lotteries, with one lottery being 'played out' at random for money at the end of the series. The premise is that, if properly incentivized, the participant will make consistent (even if biased) choices throughout the series on the prospect of being paid to maximize their potential earnings and that the resulting choice data is inherently reflective of the underlying risk preferences, decisional biases and latent decision making model of the participant[25]. Scenarios can also involve gains, losses or a mixture of gains and losses in order to test preference variance in the context of gains and losses. In loss scenarios, participants may be provided with a 'stake' so that a reference point is established in the context of prospective loss. This also allows gain and loss scenarios to be compared such that scenarios can be constructed in which the expected value of a prospective gain (without stake) is equivalent to the expected value of a prospective loss (with a stake). In their seminal research, Kahneman and Tversky showed that under these conditions, persons are generally risk averse in gains, but risk seeking in losses (Tversky and Kahneman 1992). The construction of similar 'equivalence' tests for decisional biases using carefully constructed outcome scenarios is a fundamental construct in designing these experiments and its application in my experimental context will be explained in the next section.

Since I am proposing to test for decisional bias in the context of the real world problem of choosing information security controls, I propose that the 'system at risk' should be appropriately sophisticated to 1) reflect an adequate contextual validity for both the inputs and outcomes of the system to an experienced

[25] The selection of one lottery pair at random to pay out after a series of lottery pair choices creates what is referred to as a 'compound lottery' and has implications for the way in which decision makers process final decision probabilities and prospective outcomes. Harrison et al review the theoretical and experimental implications of this 'reduction of compound lotteries' problem and its potential impacts on decision making biases (Harrison, Martínez-Correa et al. 2015).

practitioner; and 2) permit the construction of a sufficient range of decision scenarios which support the desired range of decisional biases to be tested. This means that the system generating the choice scenarios should ideally be one which the practitioner acknowledges to be an acceptable representation of an information system at risk based on both their subjective experience of information systems in the real world and its appropriateness in the context of the decision task being presented in the experiment. On the other hand, the main objective of this research is to test for decisional bias, not to develop a model of an information system at risk. This means that I have had to limit the scope of scale of the system model to that which is manageable for the purposes of supporting the experiments. In addition, while the system generating the prospective outcomes should be sufficiently sophisticated in terms of generating plausible stochastic outcomes, the experiments themselves should likely employ choice scenarios that can be benchmarked directly to established experiments involving simpler physical analogs in order to ensure methodological comparability. This requires that the simulation model be able to generate 'constructed' scenarios that permit the incremental testing of choices in which (possibly) only specific inputs are changed between scenarios.

For example, as an analog, Antoniou (Antoniou, Harrison et al. 2010) performed a choice experiment consisting of two bins containing an equal number of 10-sided dice whose individual faces are either 60% 'blue' or 60% 'white' and where each bin contains predominantly blue-faced or white-faced dice respectively. A set number of dice are then selected at random and in secret by the experimenter from one of the two bins and the dice are rolled in secret. The numbers of resulting blue and white faces are then announced to the player who then places 'bets' with a fixed stake amount according to a 'bookie table' offering varying odds that the dice were *in fact* selected from the predominantly blue-faced or white-faced die bin respectively. By observing the betting 'crossover point' at which participants expect that predominantly blue dice results are subjectively expected to have in fact come from the white bin (and, symmetrically, vice versa), and by varying the number of dice selected per roll, the experimenter is able to econometrically estimate, variously: the participants' subjective probability estimation of the likelihood of dice coming from either box; the participants subjective expected utility of a given bet; and the participants' subjective weighting between the 'strength' (percentage of blue vs. white faces on any roll) and 'weight' (number of dice rolled or evidence 'points') of evidence presented in each experiment.

To test this 'strength/weight hypothesis' in the context of information security, I propose an analogous experiment in which participants are required to subjectively estimate whether a sample of daily business transaction losses due to system degradation are generated from a 'yellow' (relatively high risk) or an 'orange' (relatively low risk) information system, where the (known) probability of observations being above/below a certain dollar threshold are 60%/40% for the yellow system and 40%/60% for the orange system respectively. To my knowledge this is the first experiment of this type to be made specifically in the context of information security using security professionals instead of lay participants (or MBA students).

Using the analogous simulation of an information system at risk, the experiment permits the testing for decisional biases based on the *strength* of the evidence (i.e. what percentage of observations are above the threshold and by how much) and by the *weight* of the evidence (i.e. the number of observations sampled across in a series of experiments involving 3, 5, 9, 12 and 17 observations respectively). Assuming the construction and presentation of an appropriately realistic and flexible system at risk, similar experiments based on established decision research studies can be constructed to test an appropriate range of decisional behaviours of interest in the context of both psychological decisional biases and information security.

Important methodological progress was achieved from my continued in-depth review of Harrison's work on choice under uncertainty in lab settings and the referenced seminal publications in the experimental economic field generally. As noted previously, significant conceptual analogs for the system at risk and the resulting choice experiments was initially motivated by Antoniou (Antoniou 2010) and then by both Fiore (Fiore, Harrison et al. 2009) and Sen (Sen 2010). Fiore (Fiore, Harrison et al. 2009), conducted a lab experiment in which participants are presented with a computer simulation of a forest fire and are asked to make a series of decisions about how to prevent a house within the forest from burning down. The participants are given ownership over the 'virtual' house of nominal value located within the virtual forest which is stochastically subject to forest fire and are asked whether they prefer to purchase a government provided 'controlled burn' around the house from a staked cash fund that would reduce (by an unknown amount), but not eliminate, the chance that the house would burn if a forest fire broke out during the simulation. The forest fire 'system' (including 3D visuals of the fire that could be played forward and backward by the participants) and the determination of whether the house burns is simulated in computer software and participants are paid out at the end of the simulation based on the net value of the house (if it still exists) and their individual net cash left over after paying (or not paying) for the control. The associated econometric *joint estimation* of subjective risk attitude and subjective beliefs over the likely behaviour and outcome of the system specifically incorporating the endogenous efficacy of the (single) control provided an excellent example of a relatively complex simulation of a system at risk over which subjects are asked to make binary choice control bets under both risk and uncertainty. Sen extended Fiore's work using the same simulation system (Sen 2010) but performed a series of experiments that have informed my hypothesis development.

I will now turn to the consideration of what constitutes an 'appropriate' simulation of an information system at risk within the context of this research and describe my specific choice of an approach and a software platform to generate the required simulations.

### Considerations for Use of System Simulation in a Lab Setting

Although lab experiments can be and are increasingly delivered via computer interface, Harrison stresses the need to conduct computerized experiments in person within a formal lab setting in order to establish instructional certainty of the choice task for each participant, eliminate or minimize cheating or collusion and, for this particular class of experiments, to ensure transparency of the stochastic outcome generation process for payoffs as a key element of satisfying participants of the fairness of the experimental methodology (Harrison almost always uses dice rolling in his experiments) (Harrison and Rutström 2008). All of Harrison's doctoral candidates whose dissertations I have been able to review in depth have implemented experiments in person (Antoniou 2010; Antoniou, Harrison et al. 2010; Rivenbark 2010; Sen 2010) although both Rivenbark's and Sen's experiments were still comparatively complex for the genre. For the purposes of the proposed research, use of non-simplistic computer simulations, while increasingly popular in experimental settings because they provide a factor-rich analog to the real world decision problem under study, remain challenging to implement from a number of perspectives and will have to be carefully specified prior to undertaking any actual programming or implementation (Andersen, Harrison et al. 2006; Fiore, Harrison et al. 2009; Crookall 2010; Lawson and Lawson 2010). At the same time, while Baldwin's system model (and the proposed simulation model for my experiment) is complex, he was not running binary choice experiments and, aside from the simulation model, the data generation methods were comparatively simple survey tasks. While the simulation is a core part of my proposed research, it is nonetheless 'background' to the choice tasks and subsequent econometric modeling and these aspects deserve due consideration, possibly at the expense of a more robust simulation. The research can always utilize a simpler system at risk, but it cannot compromise on the data generation and modeling.

### Simulation of an Information System at Risk

Significant research was undertaken on my part between October 2012 and October 2013 to determine the appropriate logical model and (physical) software platform for a simulated information 'system at risk' that would 1) incorporate enough real-world factors to be considered valid from the perspective of an information security practitioner, 2) incorporate endogenous risk factors (Shogren and Crocker 1991) which the participant would potentially i.e. interactively be able to influence, although not deterministically, to affect the outcome of a risk scenario simulation; and 3) was sufficiently manageable from a programming perspective to permit me to create a system and associated controls that represented the stochastic outcomes I was interested in using as the basis for choice under uncertainty. The following section reviews my research regarding three potential candidate platforms, each of which approach the problem of system modelling from a different perspective and whose review provided perspective on the 'best' approach for my purposes.

# 5 - Modelling Overview: System Simulation and Choice Modelling

My proposed research model for security control choice builds on theoretical work undertaken by Collinson (Collinson and Pym 2009) for discrete process simulation in the context of system security modelling, and associated theoretical work on utility maximization under uncertainty within the information security domain (Beresnevichiene, Pym et al. 2010). My own simulation model forms the basis for subsequent behavioural economics lab experiments which present participants with an information system at risk and permits participants to make domain relevant choices over specified controls under conditions involving both risk and uncertainty. The experiment generates individual participant choice data which permits consideration of multi-criteria utility functions, specifically involving loss functions (Zellner 1986; Beautement, Coles et al. 2009; Ioannidis, Pym et al. 2009) and the specification and econometric estimation of *choice models* which allow for a possible combination and weighting between multiple latent models of choice ('mixed models') (Harrison and Rutström 2009). The research contribution is to specify structural models of latent decision making behaviour as generally proposed by Harrison which permit the joint estimation of a range of relevant biases within the context of information security including, variously: the risk aversion of the subject and the weighting between EUT and PT decision models (Harrison and Rutström 2009) in the context of multi-period decisions involving asset accumulation (Game 1) (Thaler and Johnson 1990; Cubitt and Sugden 2001; Andersen, Harrison et al. 2006); subjective Bayesian updating based on the strength and weight of the quantitative evidence presented system simulation (Game 2) (Antoniou, Harrison et al. 2010); decisions over exogenous versus endogenous simulated risk (Game 3) (Fiore, Harrison et al. 2009; Sen 2010); the recovery of subjective probability estimates for continuous distributions of simulated security losses (Game 4) (Harrison and Ulm 2015); decisions over security precaution versus insurance using system simulation (Game 5) (Bajtelsmit, Coats et al. 2015).

### 1 - HP Labs GNOSIS System

Collinson (Collinson, Monahan et al. 2009) outlines the essential theoretical modeling approach (Figure 4) undertaken in the context of system simulation. Beres and then Baldwin correspondingly specify a *logical* model (Figure 5) which incorporates both stochastic system simulation and utility modeling:

**Figure 4 – Theoretical System Simulation Modeling Approach**

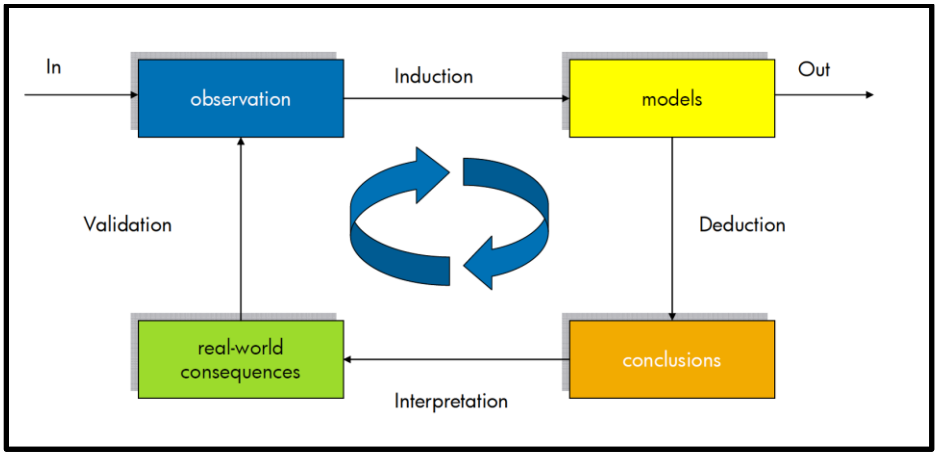

(Collinson and Pym 2009)

In this approach, a theoretical *discrete process system model* is proposed based on induction from real world observations of actual systems. The model permits the simulation, manipulation and analysis of stochastic discrete system *states* and *transitions* and the subsequent *deduction* of the system's causal interactions and outcomes which may be complex and not obvious depending on the model's specification and underlying event joint probability distributions. In the security context, the system model includes security risk factors (assets, threats, vulnerabilities, and economic and operational impacts of discrete operational states) and component control interventions over vulnerabilities based on decision maker economic preferences. Decision maker behaviour and the resulting system states (attributed outcomes) are then interpreted as having real world *implications* (transactional losses) which can be reasonably compared to observed outcomes by professional participants for actual systems and decision makers in real work contexts. The degree of alignment between the simulation with real world behaviours and attributed outcomes permits continuous adjustment of the system and the desired utility models to capture additional components, interactions or hypothesized choice models as required.

Baldwin (Baldwin, Beres et al. 2011) specifies a *descriptive* research model for security control choice based on Collinson's system modeling work which was appropriate for my purposes as a starting platform for the analysis of the control selection problem under both risk and uncertainty. The proposed mathematical model permits the *stochastic* representation of the objective system at risk which permits generation of non-deterministic and therefore uncertain outcomes for the decision maker and supports the elicitation of economic preferences and utilities representative of 'multi-objective, multi attribute' decision problems under uncertainty. The establishment of appropriate *binary discrete choice experiments* under scenario-based the prospective outcomes under both risk and uncertainty can then be used to evaluate specific risk attitude and preference models.

**Figure 5 - Economic Framing and System Model**

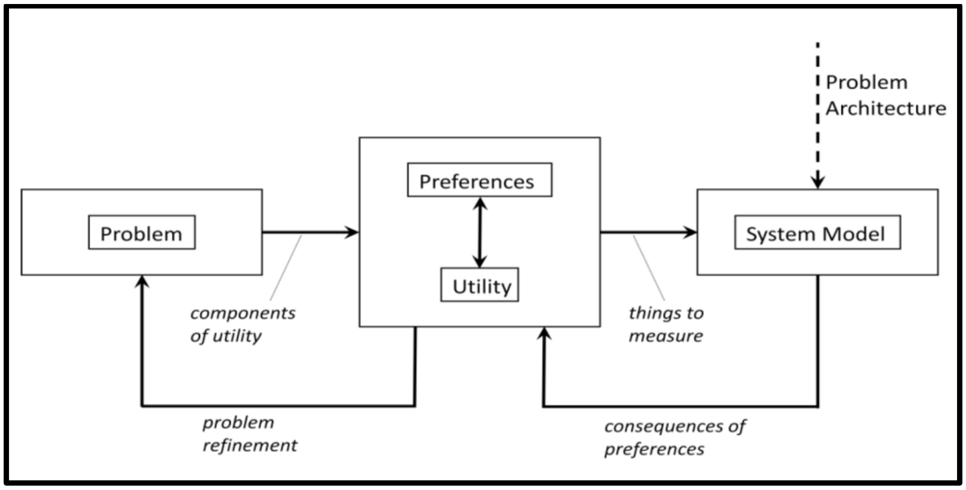

(Baldwin, Beres et al. 2011)

Choosing a *specific* system model to represent the security aspects of the system at risk was a key consideration for this research since the achievement of the model itself was only an input to the development of the experiments and not the ultimate goal of this research. An early insight was based on prior work which analyzed system *component* risks as opposed to overall system risk (Beres, Griffin et al. 2008; Baldwin, Mont et al. 2009; Beautement, Coles et al. 2009; Beresnevichiene, Pym et al. 2010), but which was applied to a 'whole system' view of a specific set of decision problems. Beres (Beresnevichiene, Pym et al. 2010) describe the key conceptual components of the system as defined by Collinson and their mathematical instantiation in a proprietary modeling language based on prior work undertaken in conjunction with HP Labs in the UK (Collinson and Pym 2009). The *Gnosis* modeling language and toolset developed by HP labs (Collinson, Monahan et al. 2010) is a discrete process simulation language which enables stochastic, Monte Carlo simulation of the system at risk and, although not designed specifically for security system modeling, has been successfully applied in the security domain and is freely available for research use[26]. Figure 6 illustrates the essential system model deployed by Baldwin for use in his work which is representative of a typical discrete process simulation involving external security threats, in this case a component vulnerability and an associated malware threat:

[26] http://www.hpl.hp.com/research/systems_security/gnosis.html

**Figure 6 - Representative schematic of a discrete process system model (Baldwin)**

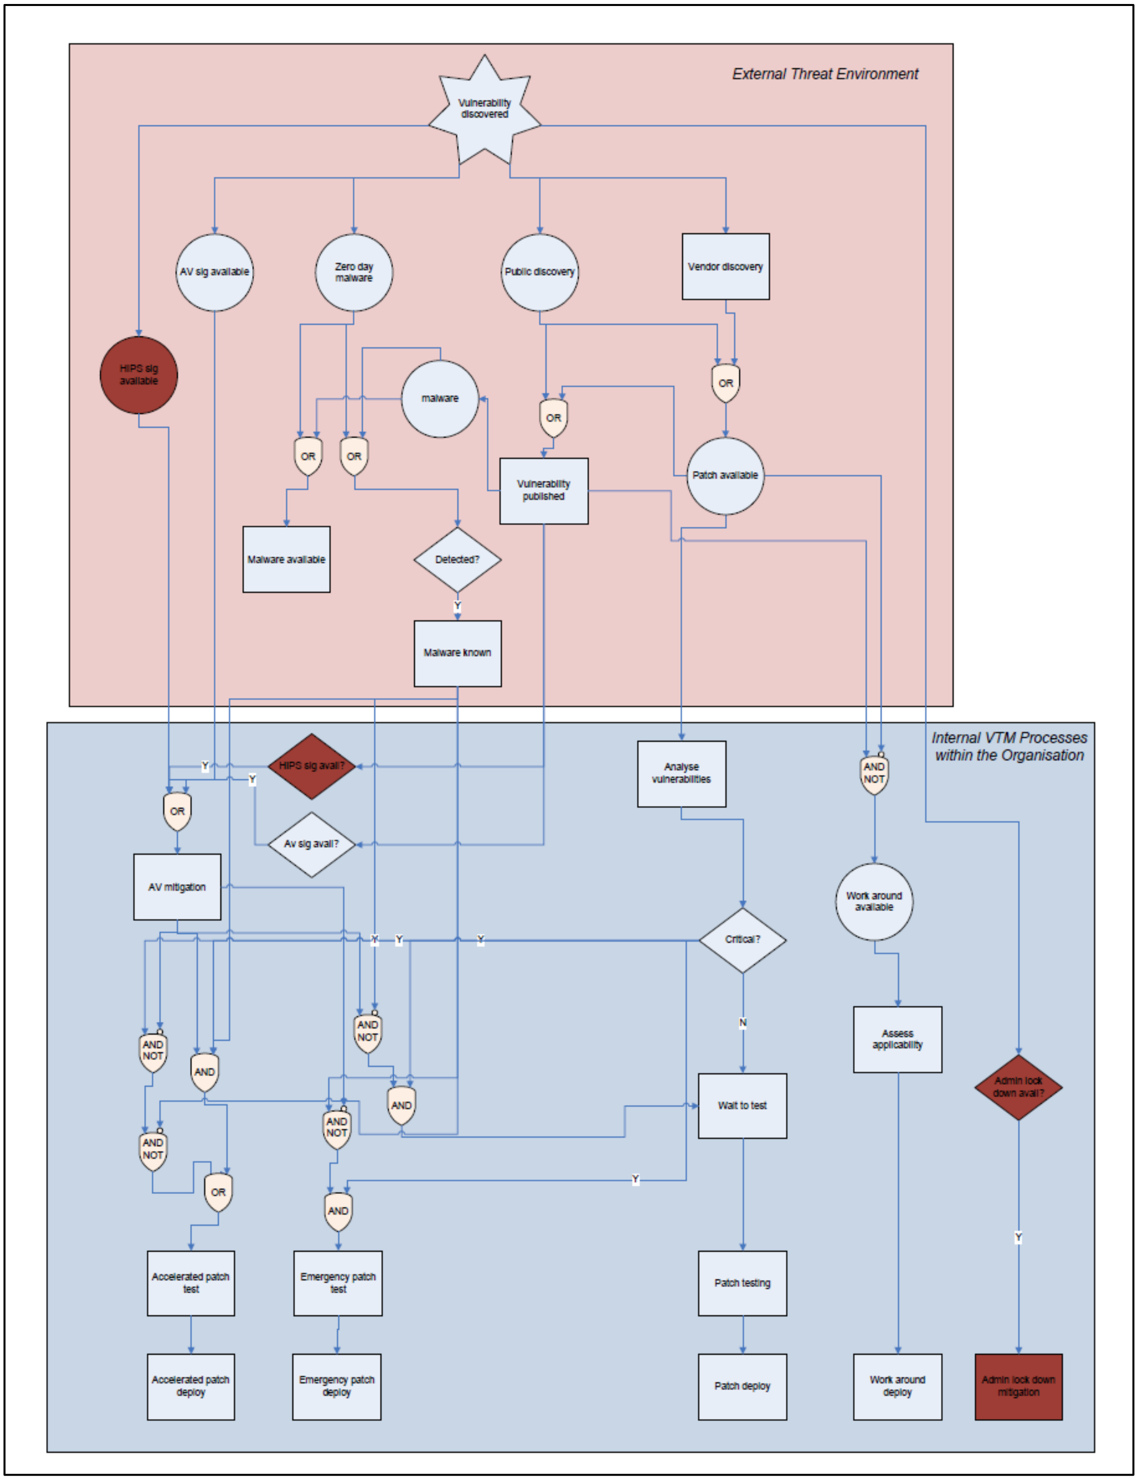

(Baldwin, Beres et al. 2011)

In this model, the boxes represent process 'event' generators and decision points for controls with connectors representing transition paths between located system resources and processes. After the specification of an initial system state and an attack vector, the system can then be simulated, allowing the paths to run concurrently and to generate quantitative metrics representing the states of the various resources and processes over time. Additional dynamic or stochastic parameters can then be specified, for example with respect to how long control activities take to activate or the activation gradient. In Baldwin's example, nine discrete measurements were generated representing various aspects of system risk based on prior qualitative surveying of a panel of respondents' economic and operational concerns:

- **Machine-days exposed to know malware** - # of machines x # days each of them is exposed from the time that malware is known in the wild
- **Exposure window** – # of day's workstation environment is exposed (vulnerability not mitigated) from vulnerability discovery time
- **Malware infection** - # of machine infections
- **Helpdesk calls** – number of helpdesk calls per each vulnerability.
- **Productivity of operational staff** – number of hours spent by operational staff doing security tasks
- **User satisfaction** – level of user satisfaction between 0-4.
- **Policy violations** - % of vulnerability cases where mitigations took longer to be deployed than policy dictated timeline
- **Capital cost** – one off dollar value of controls
- **Operational cost** – dollars spent yearly on a control.

Using the Gnosis system, Baldwin was able to display simulated outcomes to respondents with comparison to one of four *control interventions* to permit the respondent to interactively consider potential tradeoffs based on control choice and expected results.

### Model Pros and Cons

After reviewing the conceptual and logical modelling examples from the associated body of work based on the Gnosis system, I gave strong consideration to the apparent feasibility of my learning and then manipulating the Gnosis system to create an appropriate system at risk for my own purposes. The Gnosis platform, while conceptually and logically modelling discrete events of an information system, is physically based on an object oriented modelling language. While freely available for research use, use of this platform would have involved the need to learn the Gnosis language and then develop a model which, at that point, potentially required a wide range of unknown system features, inputs and outputs appropriate for my emerging set of hypotheses and choice experiments. The system also apparently lacked any kind of graphical user interface which would make implementing the code prospectively more feasible since I was not familiar with the modelling language or the coding environment. This prompted me to explore alternative software platforms that, while involving discrete system simulation, might provide a similarly robust but more flexible and ‘user friendly’ modelling environment.

The HP research essentially established a formal mathematical model of an information system at risk consisting of a series of temporally parallel and dependent processes, the steps of which could be specified in terms of discrete events. Importantly, the discrete event model can be simulated in order to evaluate alternative operational input and output modes of the system (Collinson, Monahan et al. 2009; Collinson and Pym 2009; Collinson, Monahan et al. 2010). While the modelling platform developed for this was not immediately suitable for my purposes, the HP work indicated that I should likely explore alternative discrete event modelling environments to identify a more feasible platform for my use. Additional literature search for “discrete event simulation” led to a review of several candidate platforms[27] that were superior in terms of apparent learning curve and ease of use. Platforms that did not directly employ a graphical user interface were rejected as were those which did not provide a readily deployable demonstration version that I could try before buying. Schriber (Schriber, Brunner et al. 2013) provides a short guide to DES models and reviews several popular platforms. His review includes an overview of the *ExtendSim* platform discussed below.

At the same time as I was exploring platforms generally, I continued to search for research combining ‘security’ and ‘discrete event simulation” generally. Felde (Felde 2010) wrote a Master’s thesis which, of particular interest in my specific research context, modelled and simulated a hospital information system process using the *ExtendSim* platform, a commercial platform indicated by my other search[28]. Felde was able to demonstrate a reasonably complex model of a discrete event process in the context of information security decision making concerning the deployment of system access smartcards vs. passwords for hospital physicians. Felde’s work contributed to my key understanding of modeling for discrete event simulation and platform choice. Felde indicates that there are specific steps in the overall modelling process according to Banks and as represented by the platform vendor themselves (Banks 1984) :

[27] https://en.wikipedia.org/wiki/List_of_discrete_event_simulation_software
[28] http://www.extendsim.com/

**Table 3 - Comparing modelling and simulation methodologies**

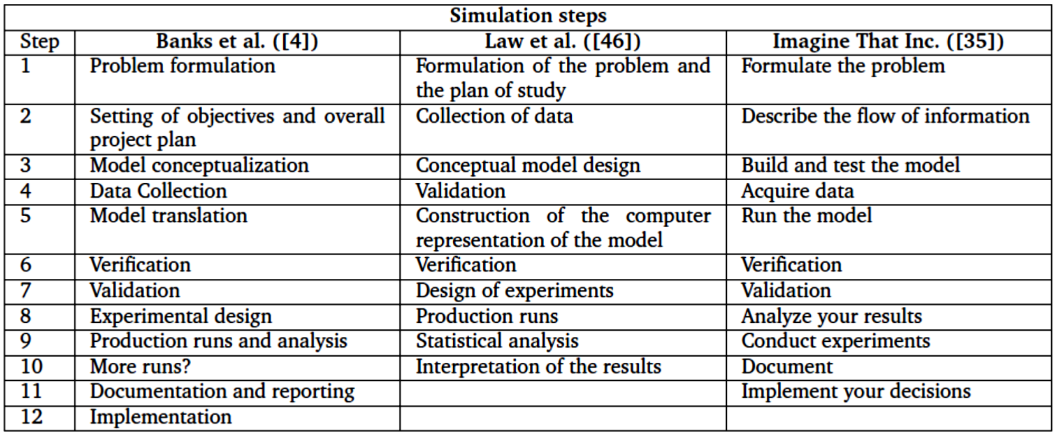

| Simulation steps | | | |
|---|---|---|---|
| Step | Banks et al. ([4]) | Law et al. ([46]) | Imagine That Inc. ([35]) |
| 1 | Problem formulation | Formulation of the problem and the plan of study | Formulate the problem |
| 2 | Setting of objectives and overall project plan | Collection of data | Describe the flow of information |
| 3 | Model conceptualization | Conceptual model design | Build and test the model |
| 4 | Data Collection | Validation | Acquire data |
| 5 | Model translation | Construction of the computer representation of the model | Run the model |
| 6 | Verification | Verification | Verification |
| 7 | Validation | Design of experiments | Validation |
| 8 | Experimental design | Production runs | Analyze your results |
| 9 | Production runs and analysis | Statistical analysis | Conduct experiments |
| 10 | More runs? | Interpretation of the results | Document |
| 11 | Documentation and reporting | | Implement your decisions |
| 12 | Implementation | | |

Source: Felde (2010)

Felde represents the modelling and simulation steps indicated above by a flow diagram indicating several stages of model verification and validation:

**Figure 7 - Categorization of steps found in modelling and simulation methodology**

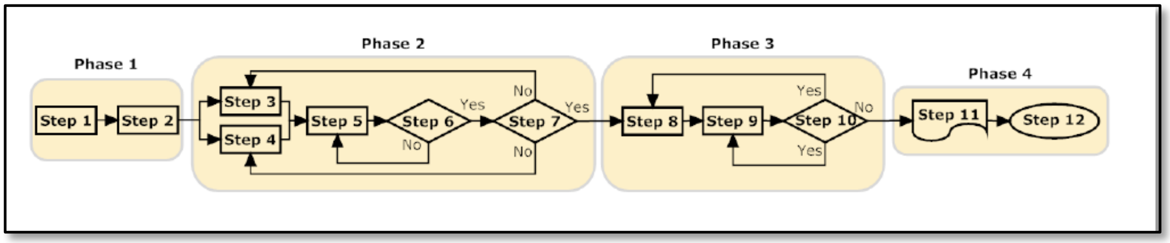

Source: Felde (2010)

Felde concludes by indicating that his work will only cover the first 3 Phases above, translated into 5 distinct steps:

1. Preparation (Phase 1)
2. Design and data collection (Phase 2)
3. Model building (Phase 2)
4. Simulation parameter estimation (Phase 3)
5. Analyzing the simulation output (Phase 3)

This outline guided my preparation for model development and continued through the actual model development indicated below.

Felde also presented a decision matrix that assisted in my consideration for the various available DES platforms:

**Table 4 - Comparing simulation environments**

| Simulation Environment | Price | Ease of Use | Applicability | Combining drag and drop customization | Support | Area Specific |
|---|---|---|---|---|---|---|
| AnyLogic | High | Yes | Yes | Yes | Yes | No |
| ExtendSim | Free | Yes | Yes | Yes | Yes | No |
| FlexSim | High | Yes | No | Yes | Yes | Yes |
| Micro Saint | High | Yes | Yes | Yes | Yes | Yes |
| OMNet + + | High | Yes | Yes | Yes | Yes | Yes |

Source: Felde (2010)

Felde indicated that the ExtendSim platform was offered to him for free during his research which presented a possible support path for my own research.

On further examination of Felde's work, I elected to review the platform and Felde's implementation in greater depth as a reference point for my simulation approach. Felde's work was initially attractive because it dealt specifically with security control decisions within a hospital context - a security issue and business setting which I was familiar with professionally: whether to deploy smart cards or passwords to physicians and to what extent each mechanism affected patient service times overall, and specifically in the event of a lost card or password reset support process. This was also the first model I had encountered in the literature which expressly set out to model the 'business impact' of a security related IT process 'at risk'. A key motivation emerged while pursuing ExtendSim and his modelling approach: the prospect of being able to directly benchmark my sufficient understanding of the platform through a direct replication of his model which was available from the ExtendSim website as a sample of academic work accomplished using the platform. Proceeding with my own modelling without a core understanding of a platform's specific features and limitations was considered a significant risk to this research and I retained this perspective throughout the selection of my ultimate modelling approach. Sufficient mastery of the modelling environment was also particularly important since, as noted, my achievement of a simulation of a system at risk was only an input to the ultimate goal of the research. I thus made the assumption that I would likely discover elements of the simulation over time that would need to be added or modified after initial modelling and simulation results were incorporated into one or more lab experiment interfaces. The uncertainty of that necessarily iterative design approach indicated that I would need to be absolutely sure of the platform before proceeding. At this time, I had also anticipated that I would possibly need to run the simulation in real time (i.e. on demand, as was apparently the case in Fiore noted above) even if in the 'background' via some type of participant user interface and so understanding how this could be achieved with any specific tool was considered an early key to overall success.

While I was able to obtain and run the model in ExtendSim, unfortunately his flowchart descriptions of the model in his thesis were not available in English and the prospect of decoding this underlying guide to the realized electronic model did not appear realistic. Significantly, Felde also describes several rounds of model simplifications of the control provisioning, use and reset model, eventually indicating that "…we will only focus on how the [provisioning and reset control] security measures affect the time use on a business activity…the security measure which yields the lowest time related KPI value is considered the 'best' security measure." (pg. 40) This was unfortunately overly restrictive for my purposes since I was interested in determining whether the subjective selection of controls (even if not considered optimal or 'best') is biased taking into account both the cost of the control and the prospective business losses under varying control scenarios. Felde's sole focus on provisioning and support effects of the 'controls' mirrored some aspects of the HP approach, and were similarly not wide enough in their applicability to enterprise level security decision making.

Given the lack of direct documentation for his implemented model and the incompleteness of the scope of the 'security model', I elected to pursue Felde's reference to Neubauer's work, particularly given the indication of the incorporation of costing factors into the system modelling. Neubauer introduced a model in which business downtime costs are explicitly related to the level of security control investment and sets out to model the trade-offs between security costs and the attributed business losses from security related incidents:

**Figure 8 - IT Dependent Business Process Model**

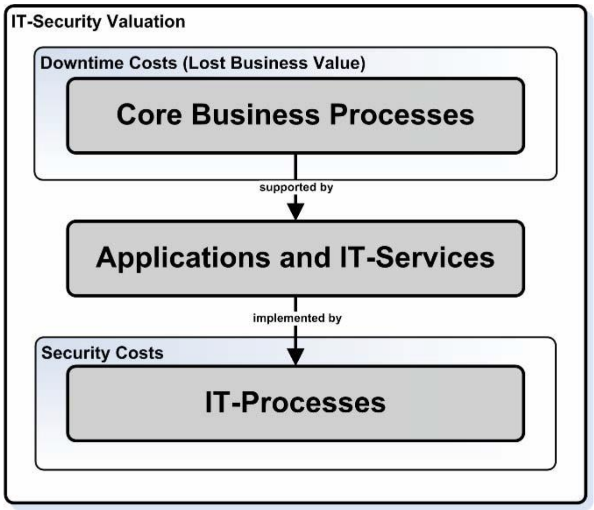

(Neubauer, Klemen et al. 2005)

In terms of a normative decision regarding control selection, the 'optimal' selection of controls can then be represented by the point at which the sum of business losses and control costs are minimized, where losses are assumed to be inversely proportional to control investments:

**Figure 9 - Security Cost vs. Benefit Trade-off**

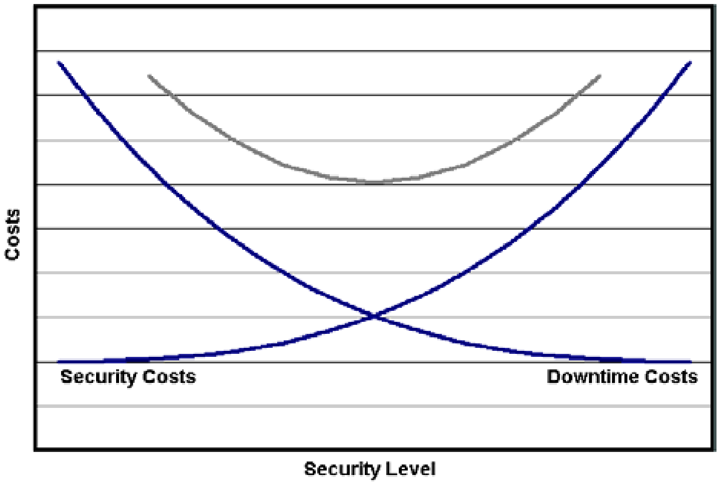

(Neubauer, Klemen et al. 2005)

A literature search for related Neubauer papers took my research in several complementary directions. First, Neubauer has collaborated with colleagues at both the SBA Research in Austria[29] and the Vienna University of Technology (TUWEIN)[30], focusing variously on multi-objective security control selection (Neubauer, Stummer et al. 2006; Neubauer, Ekelhart et al. 2008), business process modelling for IT investment decision making (Neubauer and Stummer 2007), and a survey of methods for the evaluation of IT security investments (Neubauer and Hartl 2009). Notably, Neubauer also collaborated on a 'roadmap' for risk aware business process modelling (Jakoubi, Neubauer et al. 2009) and a survey of methods for the integration of security risk into business process modelling (Jakoubi, Tjoa et al. 2009). This body of work supported my review of business process modelling and the integration of security controls within process models which has subsequently informed the ontological and structural basis for my understanding of security risk modelling.

The resulting 'roadmap' indicates the incremental components of risk aware business process modelling incorporating security risks:

[29] http://www.sba-research.org/

[30] http://www.ifs.tuwien.ac.at/

**Figure 10 - Risk-Aware Business Process Management**

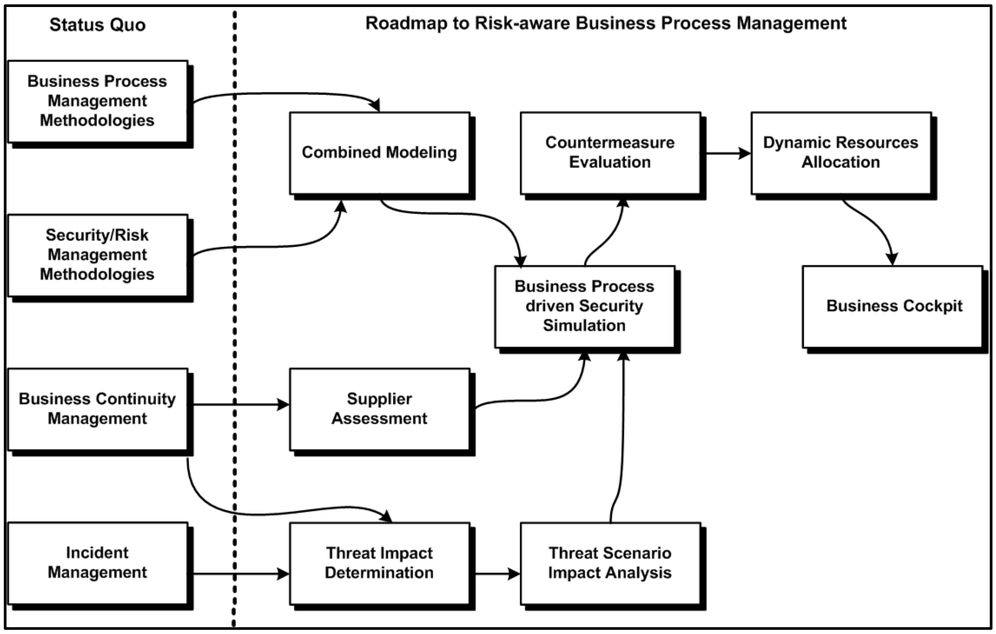
Status Quo
Roadmap to Risk-aware Business Process Management
Business Process Management Methodologies
Security/Risk Management Methodologies
Business Continuity Management
Incident Management
Combined Modeling
Countermeasure Evaluation
Dynamic Resources Allocation
Business Process driven Security Simulation
Business Cockpit
Supplier Assessment
Threat Impact Determination
Threat Scenario Impact Analysis

(Jakoubi, Tjoa et al. 2009)

Their ‘survey’ indicates which method supports simulation and economic evaluation , which they term the ‘ROPE’ (Risk Oriented Process Evaluation) approach:

**Figure 11 - Survey of Methods for Security Simulation and Economic Evaluation**

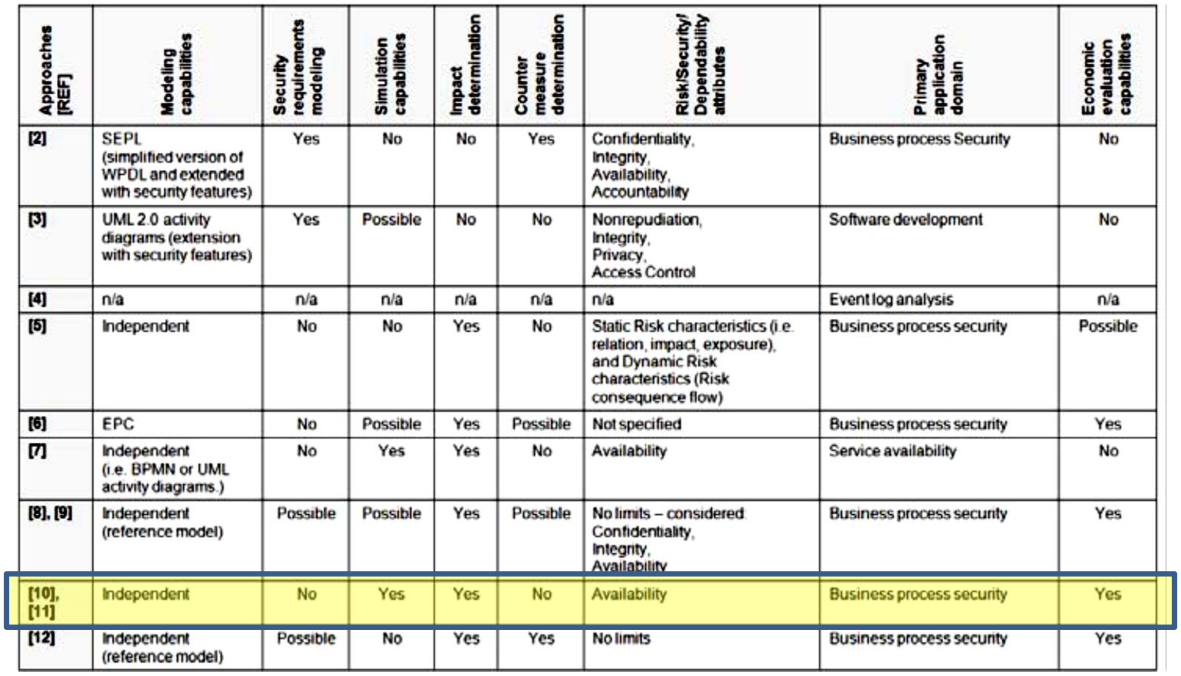

| Approaches [REF] | Modeling capabilities | Security requirements modeling | Simulation capabilities | Impact determination | Counter measure determination | Risk/Security/ Dependability attributes | Primary application domain | Economic evaluation capabilities |
|---|---|---|---|---|---|---|---|---|
| [2] | SEPL (simplified version of WPDL and extended with security features) | Yes | No | No | Yes | Confidentiality,<br>Integrity,<br>Availability,<br>Accountability | Business process Security | No |
| [3] | UML 2.0 activity diagrams (extension with security features) | Yes | Possible | No | No | Nonrepudiation,<br>Integrity,<br>Privacy,<br>Access Control | Software development | No |
| [4] | n/a | n/a | n/a | n/a | n/a | n/a | Event log analysis | n/a |
| [5] | Independent | No | No | Yes | No | Static Risk characteristics (i.e. relation, impact, exposure), and Dynamic Risk characteristics (Risk consequence flow) | Business process security | Possible |
| [6] | EPC | No | Possible | Yes | Possible | Not specified | Business process security | Yes |
| [7] | Independent (i.e. BPMN or UML activity diagrams.) | No | Yes | Yes | No | Availability | Service availability | No |
| [8], [9] | Independent (reference model) | Possible | Possible | Yes | Possible | No limits – considered Confidentiality,<br>Integrity,<br>Availability | Business process security | Yes |
| [10], [11] | Independent | No | Yes | Yes | No | Availability | Business process security | Yes |
| [12] | Independent (reference model) | Possible | No | Yes | Yes | No limits | Business process security | Yes |

(Jakoubi, Tjoa et al. 2009)

Similar to aspects of the HP business process modelling methods described above, Neubauer et al contributed to work on methodologies for ‘risk aware business process’ modelling (Jakoubi, Tjoa et al. 2007; Jakoubi, Goluch et al. 2008; Tjoa, Jakoubi et al. 2008) and specifically the formulation of the ROPE methodology and the use of discrete event simulation based on the ROPE approach (Jakoubi, Tjoa et al. 2010; Tjoa, Jakoubi et al. 2010; Tjoa, Jakoubi et al. 2011). Separately, other Neubauer colleagues have pursued the development of *security ontologies* (Fenz and Weippl 2006), and the application of security ontologies in the generation of Bayesian networks for security threat risk determination (Fenz, Tjoa et al. 2009; Fenz, Ekelhart et al. 2011; Fenz 2012). The use of both discrete event simulation and Bayesian networks, while related in terms of their incorporation of risk elements and specifically security risks into a business process model, approach the modelling solution from separate perspectives. A significant insight for this research has been to combine elements of each approach, together with Monte Carlo techniques as is explained below, to achieve a tractable, flexible and valid simulation of an information system at risk for my lab purposes. The following three sections describe these core modelling approaches, my insights from each and the fourth describes my synthesis of the key aspects of these approaches which resulted in my own simulation model.

### 2 - ROPE Methodology Using *Matlab* and *SimuLink* for Discrete Event Simulation

Research conducted at the Vienna University of Technology and SBA Research in Austria[31] undertakes discrete event simulation modelling of a ‘risk aware business process’ with a specific focus on system availability and integrity (Tjoa, Jakoubi et al. 2011; Fenz 2012; Fenz, Ekelhart et al. 2012). In contrast to HP labs, their approach utilizes the commercially available MatLab[32] platform, which incorporates a robust graphical user interface (GUI) to build a complex, multi-layered mathematical model expressing the relations between threats, detection mechanisms, safeguards, recovery measures and their effects on business processes, the simulation of which can directly generate any number of system outputs that would be of interest to the decision maker looking at control alternatives.

The risk aware business process modelling approach recognizes that business processes generally, and business processes dependent on information systems specifically, are at risk from disruptions caused by ‘threats’ to one or more of the *process actions* (e.g. call centre call processing) performed by one or more *resources* (e.g. information systems, call operators) in a specified environment. The following figure indicates the conceptual model of a business process at risk from threat disruptions and the corresponding detection, counter measure and recovery controls that mitigate potential and realized threats to the business process:

[31] Ibid

[32] http://www.mathworks.com/discrete-event-simulation/

**Figure 12 - Detection, Counter and Recovery Model of Security Control**

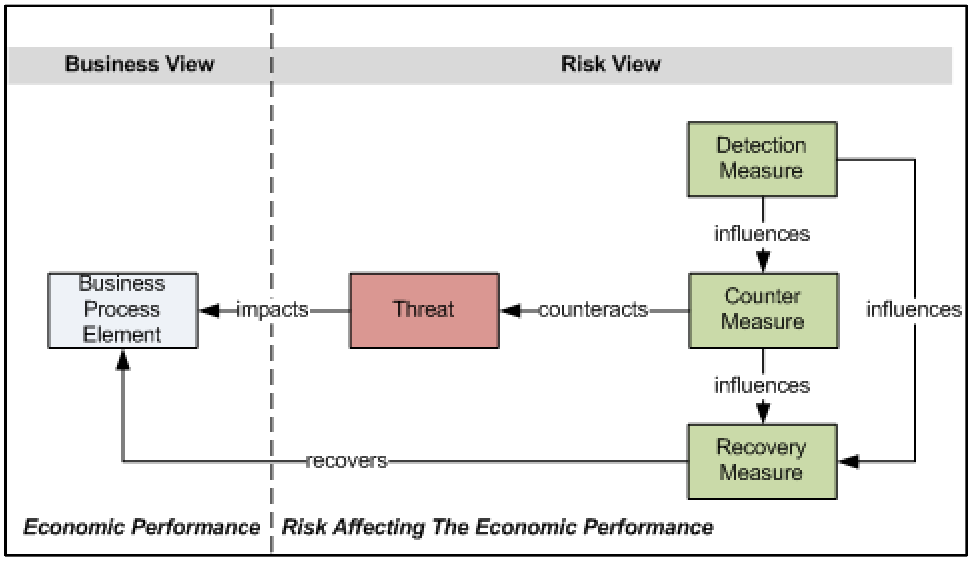

Source: Tjoa, Jakoubi 2010

Jakoubi et al (2008) introduce their "ROPE" (Risk Oriented Process Evaluation) methodology to explicitly combine business process optimization with risk management and present a methodology to simulate the effects of business process security risks over time. The ROPE Methodology consists of three modelling-layers:

- **Business process (BP) modelling layer:** Representation of the company's business processes – business process activities serve as linkage to the next modelling-level.
- **CARE** (Condition, Action, Resource and Environment) modelling layer: Identification of a business process activity's elements (Action, Resource and Environment) and their coherences (Condition).
- **TIP** (Threat Impact Process) modelling layer: Identification of potential threats as well as their management- and recovery strategy:.

The ROPE methodology extends the conceptual 'process at risk' model by adding in process *resource dependencies*:

**Figure 13 - Risk Aware Business Process Model: It Resource Dependencies**

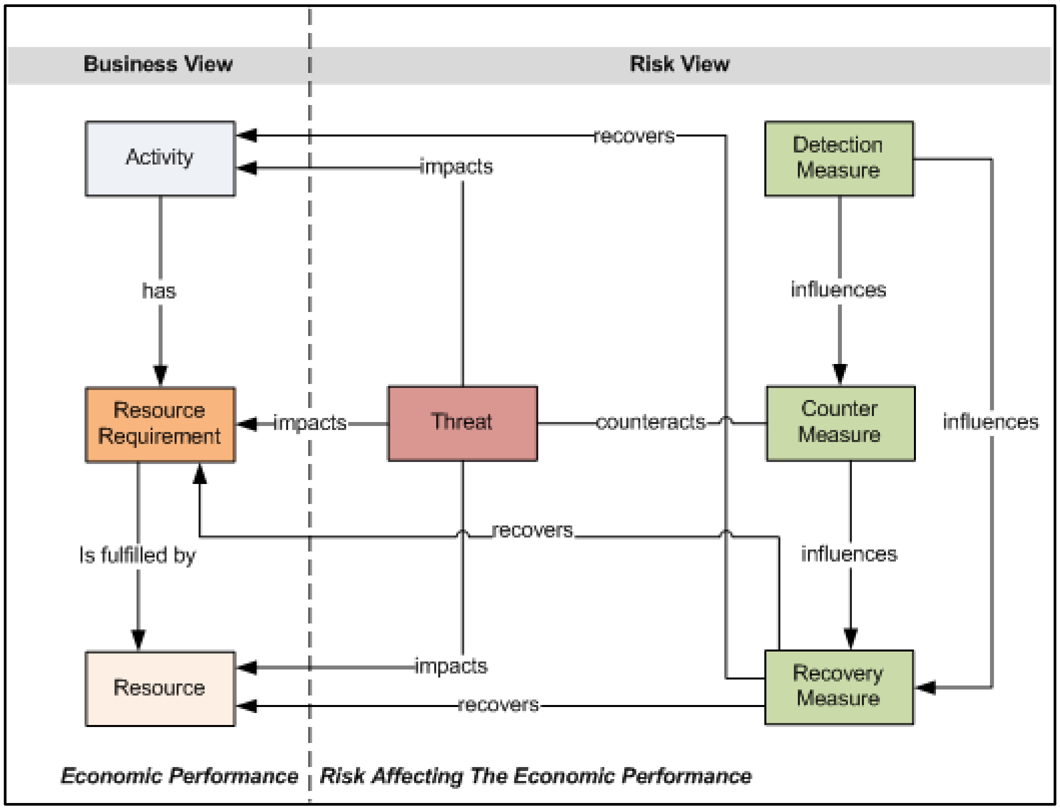

Source: Tjoa, Jakoubi 2010

A *logical* representation of the extended ROPE approach is indicated in the following figure:

**Figure 14 - Logical model of a risk aware business process**

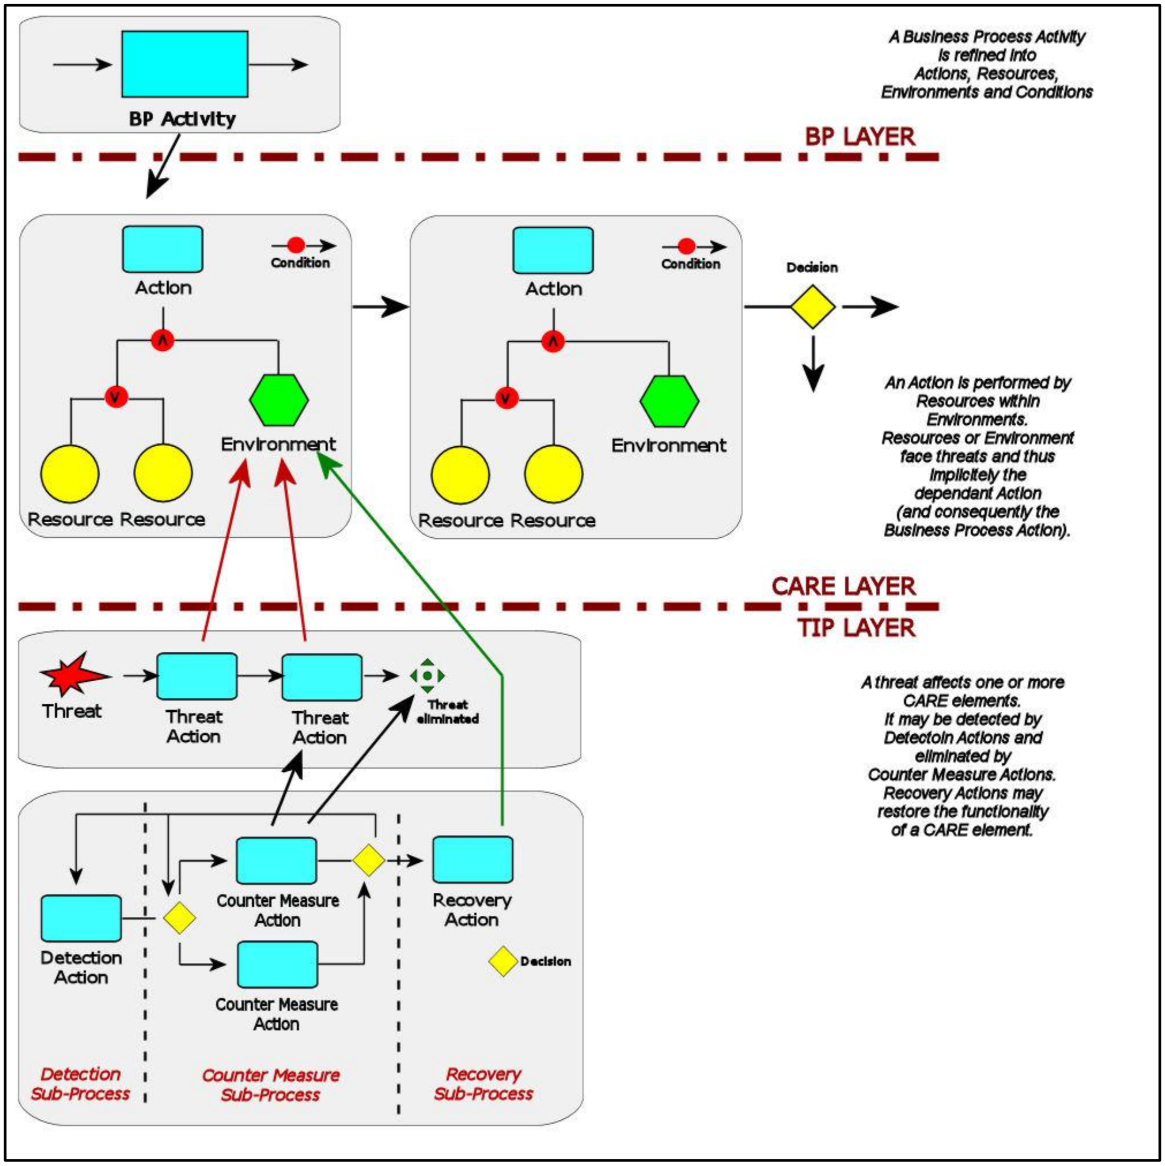
BP Activity
A Business Process Activity is refined into Actions, Resources, Environments and Conditions
BP LAYER
Action
Condition
Action
Condition
Decision
Environment
Environment
Resource Resource
Resource Resource
An Action is performed by Resources within Environments. Resources or Environment face threats and thus implicitely the dependant Action (and consequently the Business Process Action).
CARE LAYER
TIP LAYER
Threat
Threat Action
Threat Action
Threat eliminated
A threat affects one or more CARE elements. It may be detected by Detectoin Actions and eliminated by Counter Measure Actions. Recovery Actions may restore the functionality of a CARE element.
Counter Measure Action
Recovery Action
Detection Action
Counter Measure Action
Decision
Detection Sub-Process
Counter Measure Sub-Process
Recovery Sub-Process

*(Jakoubi et al. 2007)*

Since the discrete event model is intended to support simulation over time, the unavailability of impacted resources, modelled as a percentage degradation of absolute (100%) availability affects the temporal performance of the business process which continues to degrade overtime depending on 1) the efficacy of the threat detection, which impacts 2) how long the threat is active in the environment before detection; 3) once detected, how long it takes to counteract the threat; and 4) how long it takes to recover the degraded resource to full performance.

The following figure indicates a repeated business process consisting of three sequential sub-processes (A, B, C) where, in the third iteration, a threat becomes effective on process B shortly after process B begins, delaying its completion until recovered, and subsequently delaying process C in the third iteration and the subsequent start of process B in the fourth iteration:

**Figure 15 - Time dependent business activity model with security incident**

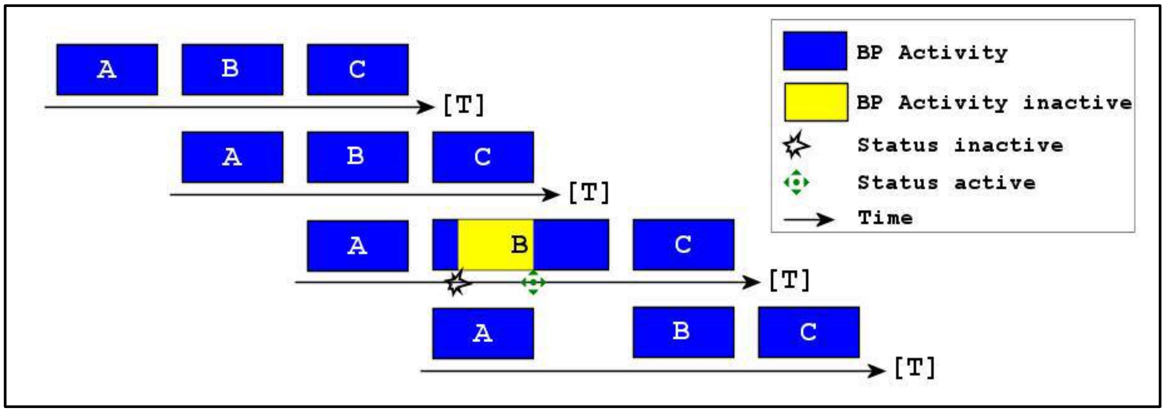

(Jakoubi, Goluch et al. 2008)

In the above example, while the B sub-process recovers in the third iteration, the interruption in the timing of the process is permanent (i.e. even though A starts on time, it is subsequently always waiting extra time for B to complete its prior execution before it can start B) and the system would have to somehow 'speed up' in a single cycle in order to return to an equivalent temporal state in which no interruption had occurred. In business process systems which are dealing with exogenous demand (e.g. calls arriving at a call centre) the process would effectively create an inventory or backlog of calls that could not be eliminated without conceptually either speeding up or adding some type of efficiency in the A/B/C sub processes. Dynamically modelling the business process backlog from period-to-period improves the ability of the model to account for the both the presence and effectiveness of the controls (here the incident manager) over time and permits the differentiation between prevention, detection, counteraction and recovery controls:

**Figure 16 - Graphical representation of a business process backlog**

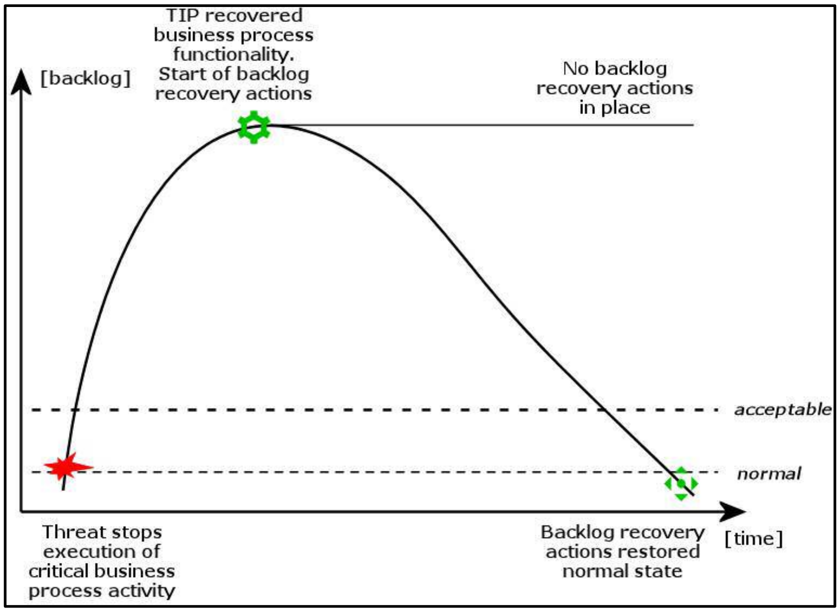

(Jakoubi et al. 2008)

The authors go on to present a contextual scenario of a travel agency dealing with customer booking orders where agency's internal booking system depends on a single information system server providing essential order processing services for the employees. A second ongoing business process of the agency is first level IT support, which is performed by the agency's incident manager. The model simulates the introduction of a computer virus degrading the computing resources and the impact on both call backlog and resource requirements to recover the process.  The following figure indicates the logical impact of a computer virus affecting the agency's server resource and the subsequent response process of the first level support resource to detect, counteract and recover the server:

**Figure 17 - Logical model of a successful virus attack on a travel agency (ROPE)**

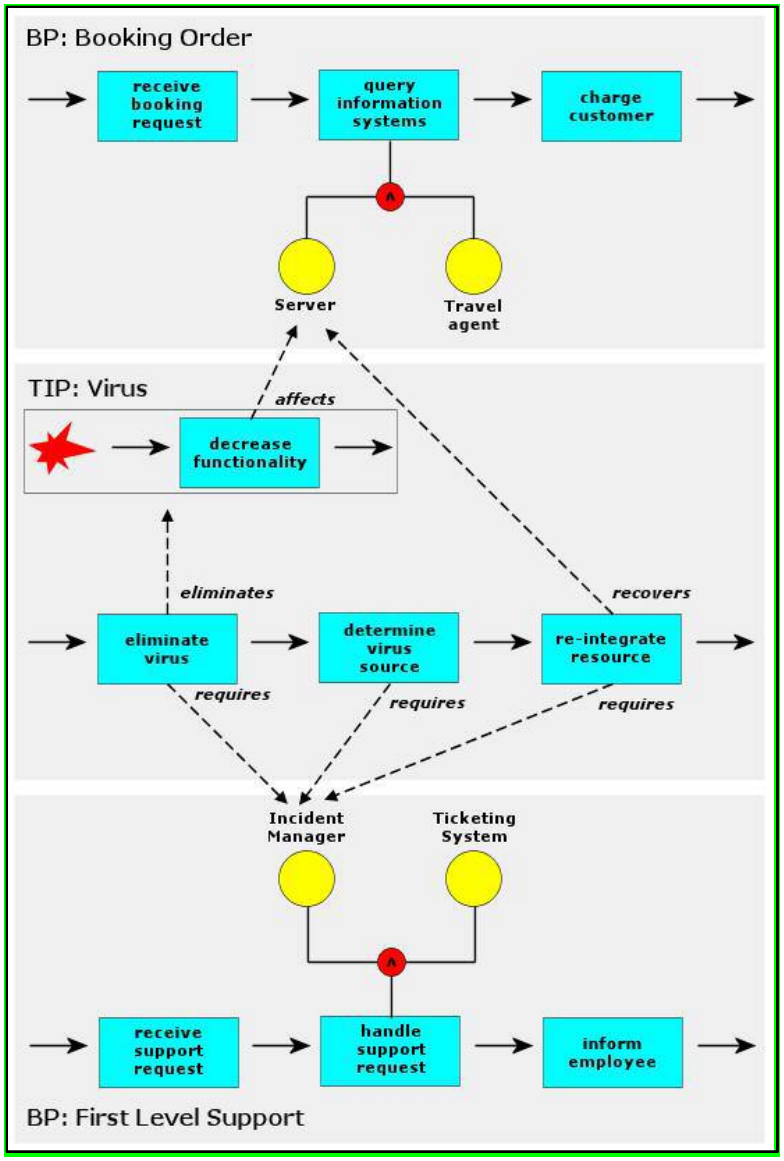

(Jakoubi et al. 2008)

Although not undertaken by Jakoubi et al in the above example, the approach is significant since it permits the determination of the (temporary) backlog in call processing the resulting increased resource requirements (activities) to return the process to a normal state, both of which represent additional costs to the business over time.

Tjoa et al proceed to illustrate three additional detailed scenarios that demonstrate the effectiveness of this approach using the SimuLink discrete event simulation tool within MatLab to both model the business process at risk and to simulate the dynamic action of the system to prevent, detect, counter and recover from disruption. The following reviews these three models in some depth in order to demonstrate the applicability and flexibility of the modelling approach and to illustrate key aspects of specific platform reviewed in the course of this research. The first example is not security centric but illustrates the basic activity and resource-dependency approach of this method, and introduces two key concepts of risk aware business process activity modelling which have carried through to my model: 1) the concept of an activity's *degree of completion (DC)* and its *integrity loss (IL)*, and 2) the concept that the system has memory such that any incomplete prevention, detection, counteraction or recovery of a sub-process or resource attribute in the previous time period affects the starting attributes of the sub-process or resource in the current time period. Degree of completion is modelled indirectly via the 'availability' of the resource', whereas integrity is modelled directly, both measured between 0 and 100%. In environments characterized by stochastic demand for resources, this modelling approach presents an opportunity to reflect not only average process activity over several time periods but the variability of the process' ability to meet demand across the time period. If business costs can be attributed to unfulfilled process demand (whether due to lack of availability or integrity or both), the system can more accurately reflect the impact of availability and integrity threats.

In terms of overall usability or throughput, both availability and integrity loss ultimately impact system output, although a distinction between availability and integrity at the process level needs to be maintained[33] depending on the system purpose, the desired type of risks (threats) to be modelled and the resulting effects to be estimated. For example, systems that are less than 100% available continue to process every request accurately, although at a lower rate (or 'degree') of completion and can result, variously, in increased system user wait times, request abandonment and overall lower productivity. Integrity compromised systems (as especially assuming full availability) potentially produce erroneous outputs (and where each unit of output may in fact be erroneous) at the given availability or capacity. Depending on the type of business purpose and computing environment, integrity may in fact be more important than availability and result in, for example, rework or output accuracy liabilities where service level agreements may be applicable (e.g. the return of imperfect goods to the manufacturer). In other cases such as healthcare or industrial control (e.g. SCADA) systems where the quality or safety of the output is carefully controlled, even small degrees of integrity loss may be absolutely unacceptable, resulting in the shutdown of the resource or sub-system despite the otherwise availability of the resource itself (in which case, availability itself goes to zero even though not directly threatened). These security distinctions based in the CIA triad permits the risk aware system to dynamically reflect 'backlogs' of process and dependent resource demand and reflect real world conditions in which the impacts of security incidents may persist within a system for several time periods before being fully resolved.

[33] See https://www.tofinosecurity.com/blog/scada-security-basics-why-industrial-networks-are-different-it-networks for examples of the consideration between availability and integrity in industrial control applications versus other IT systems. The subsequent discussion regarding the work by Ekstedt et all elaborates on their considerations for integrity over availability in their application to SCADA systems.

**Case #1: Dynamic Resource Allocation:**

Tjoa et al modelled a call centre using discrete event simulation with dynamic resource allocation properties to determine the optimal configuration of human resources between call queues (Tjoa, Jakoubi et al. 2010). The conceptual model consists of three increasing levels of call support ‘actions’ performed by 3, 2 and 2 human resources respectively, where the second and third level support human resources can be re-allocated to the prior level under conditions where the backlog of calls at the prior level exceeds a certain threshold and the resources are returned to their original position if the backlog of their original position exceeds a certain threshold.

The following diagram represents the ‘conceptual’ model of the process. The incoming queue of calls, and the call resolution determination are stochastic based on historical data and indicating that one-third of the Level One calls are typically escalated to Level 2, and 1/2 of Level 2 calls are escalated to Level 3, with 100% of Level 3 calls being resolved at Level 3:

**Figure 18 – Discrete Event Simulation Example #1: Call Centre Resourcing**

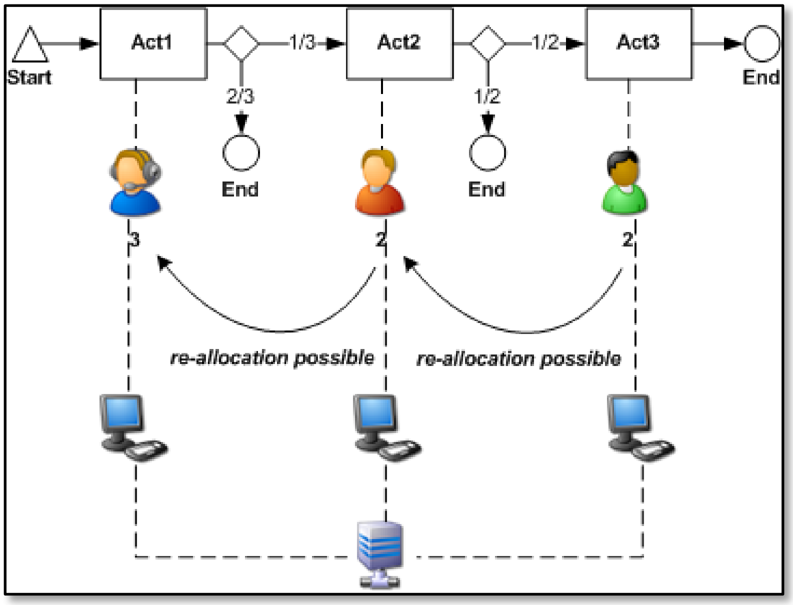

(Tjoa, Jakoubi et al. 2010)

The corresponding logical model, expressed in *MathWorks* using *Simulink* illustrates the modelling of the process, actions and resource inputs and outputs at various levels of abstraction. The following diagram indicates the “Main Layer” of the model which represents the logical instantiation of the conceptual model, with the corresponding single server, 7 human resources and their computer workstation resources and the three levels of call processing activity:

**Figure 19 – Example #1 Instantiated Model in MathWorks/SimuLink (Main Layer)**

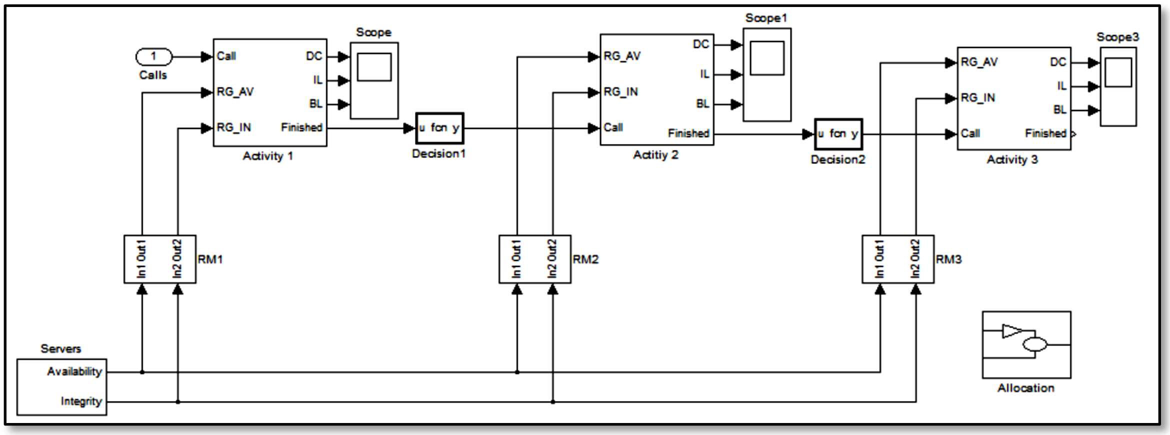

(Tjoa, Jakoubi et al. 2010)

The model monitors the inputs and outputs of the process and displays results of the simulated process overtime using display 'Scopes' connected to the activities. The model can be simulated under risk scenarios in which a generic security threats affect the availability or integrity of the PC and /or Server resources and can be used to determine the effectiveness of both controls to prevent/detect/counter/recover from the threats and, specifically as is the goal of this model, to determine whether dynamic human resource reallocation between call levels is a valuable control. The process has one financial attribute which is 'proceeds' (the cumulative attributed business value of completed calls); each activity has three output attributes: degree of completion (DC), integrity loss (IL) and call backlog (BL), and three inputs: PC resource availability (RG_AV), PC resource integrity (RG_IN) and the presence of a call (Call). The DC attribute depends on the availability ("AV") attribute of the PC resources and the IL attribute depends on the integrity attribute ("IL") for the server resource. Lower AV and IL lowers DC and increases BL. Conceptually, decreases in AV cause direct decreases in DC due to a fewer number of (here a percentage lower level of) resources available per period, while decreases in IL causes decreases in DC due to a lower level of activity cycles (speed of completion) per available resource .

The embedded Activity blocks are identical other than their respective per period input and output levels, including any re-allocated human resource instances which are dependent on queue lengths. The activities take three inputs:

**Figure 20 - Example #1 Instantiated Model (Activity Layer)**

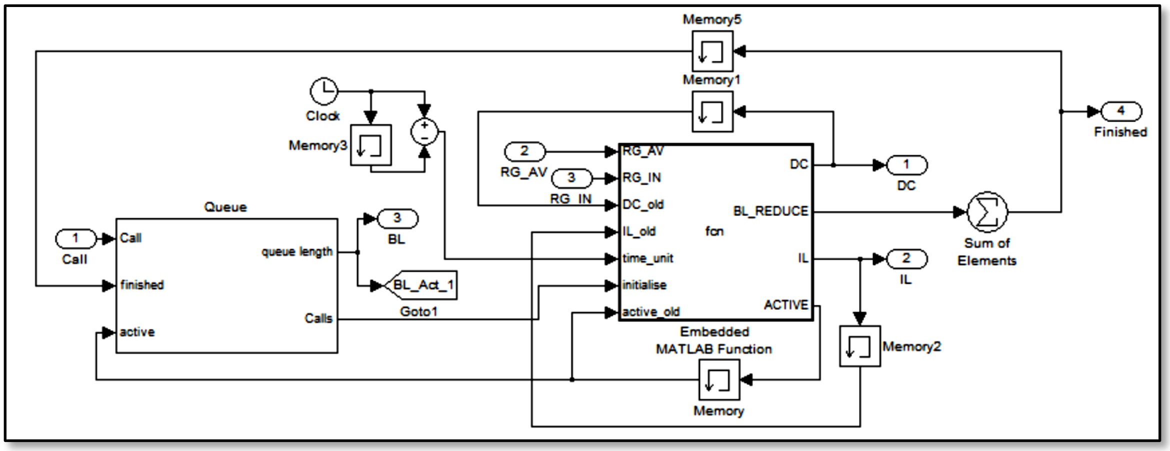

(Tjoa, Jakoubi et al. 2010)

The model also includes an “Allocation” process which monitors the call backlog of each activity and dynamically re-allocates a single human resource plus PC to a lower service activity level (increasing the number of “Active” resources at that activity level and decreasing the number of “Active” resources at the source activity level) as required:

**Figure 21 - Example #1 Instantiated Model (Resource Allocation Layer)**

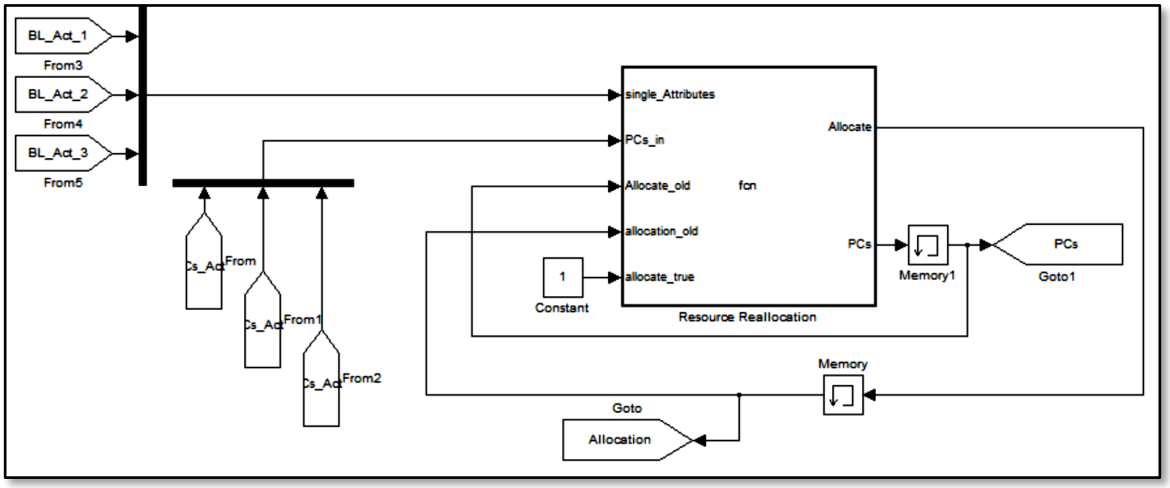

(Tjoa, Jakoubi et al. 2010)

The system is then simulated to reflect the impact on, respectively, activity backlogs and overall proceeds. In the following output, a threat impacts an Activity 1 PC resource at time step 400 and persists until the threat begins to be countered at step 730 and is fully recovered around step 980. In this case, no dynamic allocation is permitted and results in essentially stagnant business proceeds between initial impact and recovery:

**Figure 22 - Example #1 system simulation output: No dynamic reallocation of resources**

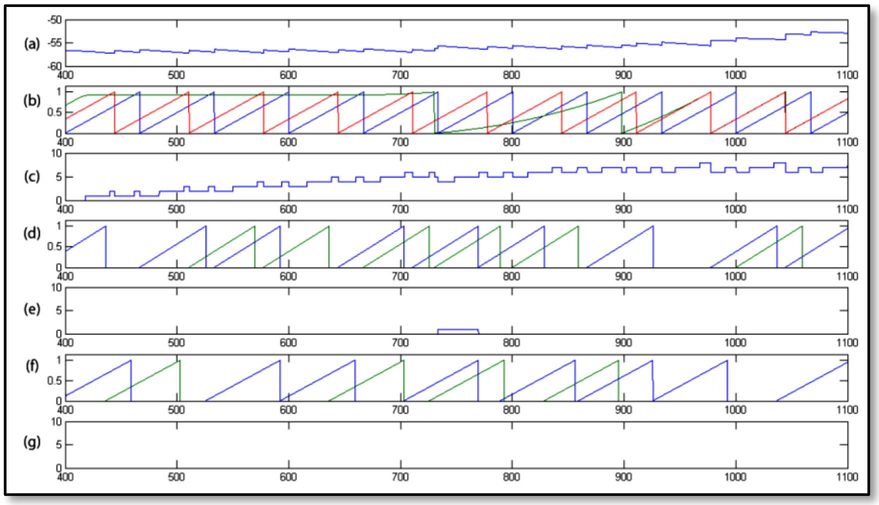

← Business Revenues

← Activity 1 (3 resources)

← Activity 1 Backlog

← Activity 2 (2 resources)

← Activity 2 Backlog

← Activity 3 (2 resources)

← Activity 3 Backlog

(Tjoa, Jakoubi et al. 2010)

In the alternative scenario permitting dynamic reallocation of resources between Activity levels, a Level 2 resource is reallocated to level 1 around step 510 as Activity 1's backlog reaches a predefined threshold. The reallocated resource stays in Activity 1 (increasing Activity 1's throughput) until Activity 2's backlog increases to its threshold at which point the Activity 2 resource returns to Activity 2. After reducing Activity 2's backlog the Activity 2 resource again returns to Activity 1 but cannot effectively reduce the backlog, and a resource from Activity 3 is reallocated to Activity 1 around time 750. This increases Activity 3's backlog but decreases Activity 2's backlog (because more calls are completed at Activity 1):

**Figure 23 - Example #1 system simulation output: No dynamic reallocation of resources**

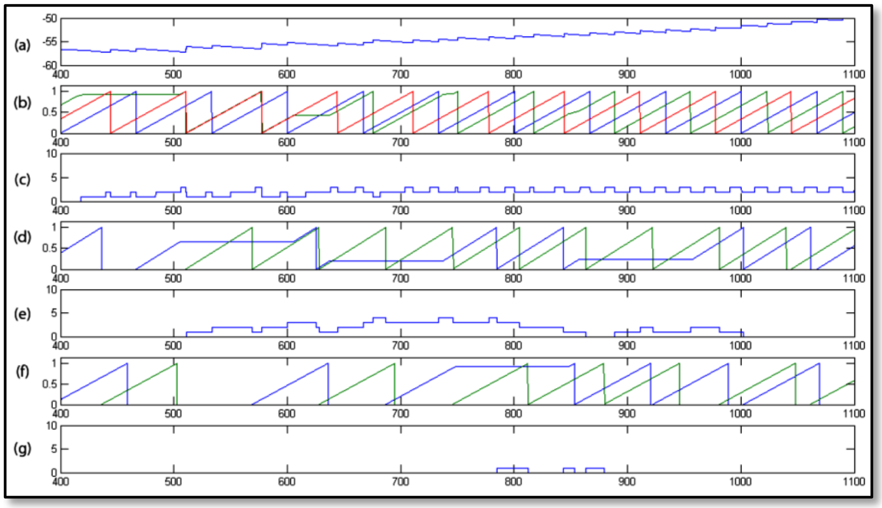

← Business Revenues

← Activity 1 (3 +1 resources)

← Activity 1 Backlog

← Activity 2 (2 -1 resources)

← Activity 2 Backlog

← Activity 3 (2 resources)

← Activity 3 Backlog

(Tjoa, Jakoubi et al. 2010)

In this scenario, Tjoa et al successfully demonstrate that proceeds are largely unaffected under simulation conditions characterized by stochastic security threat risk to computing resources using dynamic resource allocation.

**Case #2: System Availability Simulation:**

In the case presented above, Tjoa et al were able to demonstrate the impact of generic threats on computing resources and the computer dependent human resources, introducing the concept of *degree of completion* of a process activity based on the availability and integrity of the underlying resources. The case does not however directly illustrate the modelling of the core security control factors that they advocate as part of their ROPE conceptual and logical model: prevention, detection, counter, and recovery. In the following case, they fully implement these control sub-processes within the process model and demonstrate the mitigation of impacts on resource availability and integrity. These concepts support the modelling of real world security 'defense in depth' where unprevented but detected attacks must first be countered to eliminate the threat after which the impacted systems can then recovered over time, gradually restoring business processing to a normal, pre-attack state. The business cost of the impact on system availability from realized threats must necessarily also take into account these counteraction and recovery costs which may be significant depending on the severity and persistence of the attack and the effectiveness of the controls. Their contribution in the following case consists of the ability to model and then simulate the required system process and security control components and directly informed my subsequent modelling approach using these conceptual modelling and simulation components:

1) **Modelling Concepts:**
    - Business Process Activities
    - Required Resources
    - Threats endangering these resources
    - Detection, counter and recovery measures
    - Relations between these components
2) **Simulation-based determination of risk impacts** (e.g. time, control costs, backlogs, etc.):

In the following model, Tjoa et al also formally specify two further 'proactive' controls: *preventive* measures and *blocking* measures. Together, these controls further extend the defense in depth concept by reducing, respectively, the probability of a threat occurring and the probability of fully stopping a threat before it impacts the system. These additional control factors correspond to established threat and risk modelling approaches wherein risk can be measured a product of the likelihood of a threat affecting an asset's (resource) control vulnerabilities, multiplied by the impact on the asset should the threat be effective within the system. Under this model, likelihood itself can be decomposed into the probability of the threat occurring multiplied by the likelihood of the vulnerability being present[34]. It should be noted that in this model, the lack of a threat or of a corresponding vulnerability (or of both) effectively reduces the expected risk qualitatively or quantitatively to zero. The following diagrams indicate the impact of a threat on the availability of an activity, which is degraded for a period of time due to the availability impact on the underlying resource:

**Figure 24 - Resource Availability and Impact on System**

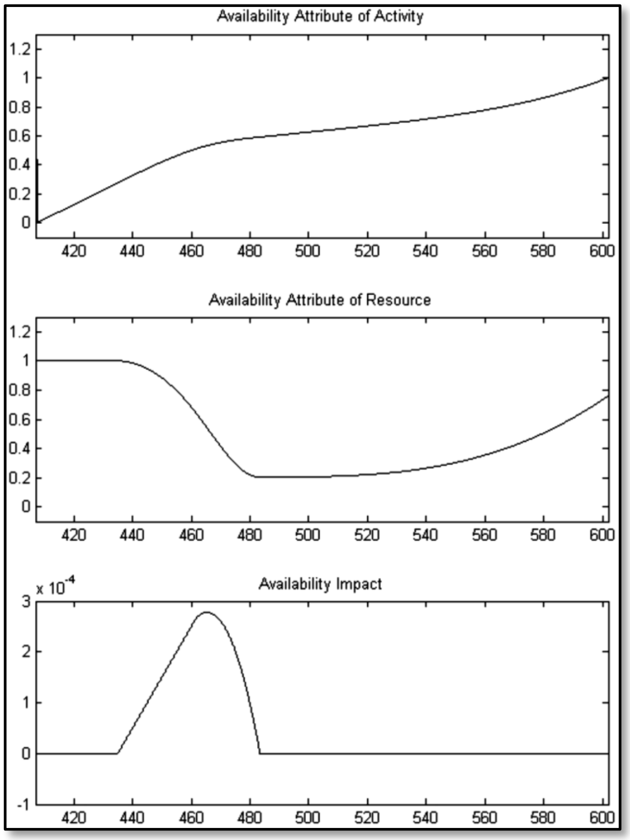

(Tjoa, Jakoubi et al. 2011)

In the above example, an activity begins processing at time 410 until a threat becomes effective around time 435 when the availability of the underlying resource begins to degrade. The resource degradation continues until the threat is countered around time 480 at which point the resource begins recovering. Here activity availability is degraded resulting in fewer executions (completions) per cycle. Similar effects are modelled for integrity threats which, as noted above, degrade the cycle time per activity execution. Both

[34] As noted above, we distinguish between probability of a prospective event and the likelihood or degree of belief that a state pre-exists. In this model, we are required to estimate the probability that a threat will occur and the degree of belief that a vulnerability actually exists i.e. the threat will occur with some expected propensity, whereas the vulnerability is in fact either present or absent with some degree of uncertainty. Both estimates are expressed as a percentage between 0 and 100, however the concepts are distinct. As noted, probability is typically estimated based on the historical frequency of the threat action whereas degree of belief in the presence of a vulnerability may be better estimated using Bayes rule based on both prior expectations of the presence of the vulnerability and available evidence indicating that a vulnerability may be present.

impacts result in attributed business losses where proceeds are measured in terms of total transactions completed per period of time.

The demonstrated case study again involves a call centre with two levels of support, however this time the model incorporates two computing resources (a shred server and individual PCs) and proactive controls for threat prevention and blocking, and reactive controls for detection, counteraction and recovery. The objective of the case was to demonstrate decision support capability for two prospective control improvement scenarios: the selection of an anti-virus toolkit where each kit varied in acquisition costs ($12,000 vs. $22,000) and availability threat blocking detection rates (75% vs. 95%) respectively. In addition, each kit had a similar integrity threat blocking capability of 84%. The Main layer of the model is illustrated below:

**Figure 25 - Example #2 Instantiated Model in MathWorks/SimuLink (Main Layer)**

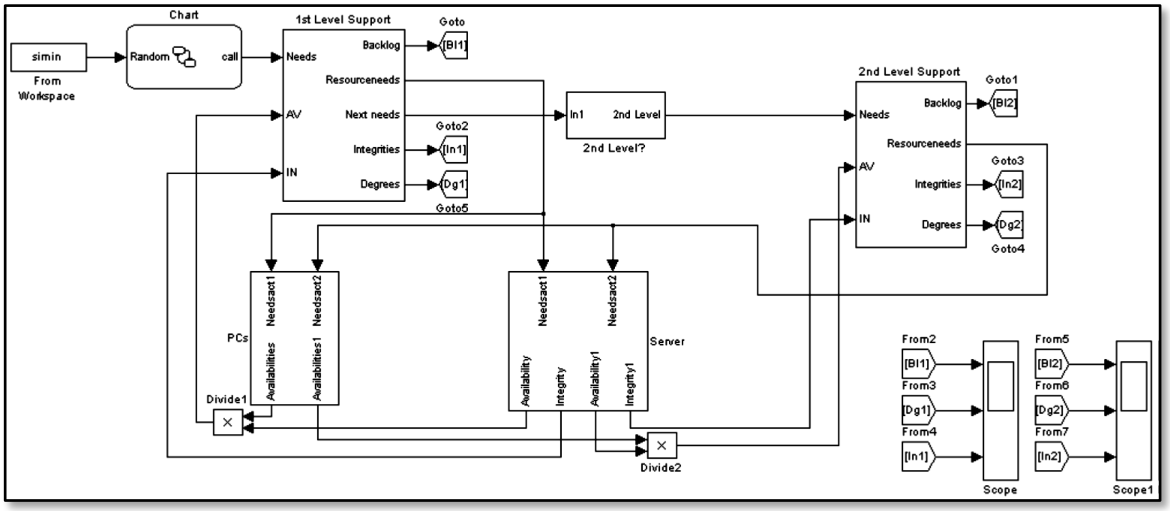

(Tjoa, Jakoubi et al. 2011)

Each Activity layer is identical and takes as inputs the server integrity and the blended availability level of both the server and the Activity's PC resource. The Activity's degree of completion of a call in queue is monitored from start (DC = 0) until it completes the call (DC = 1). A decision block "2nd level?" stochastically determines whether a level one call is escalated to Level 2:

**Figure 26 - Example #2 Instantiated Model (Activity Layer)**

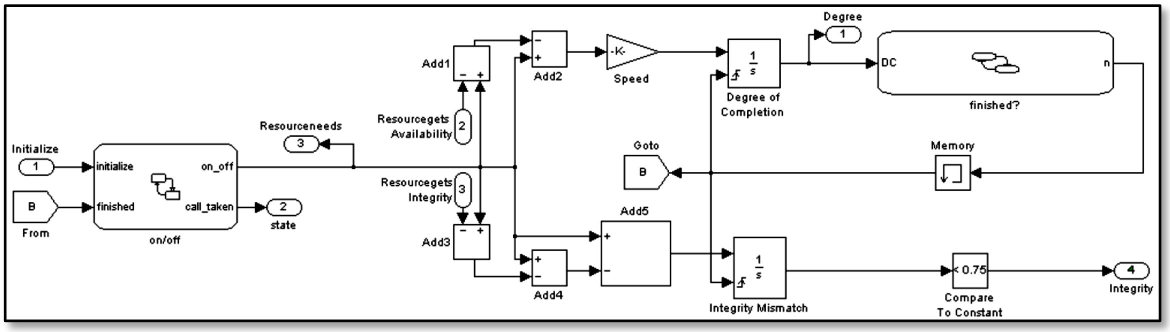

(Tjoa, Jakoubi et al. 2011)

The resources contain sub-processes for 1) determination of unprevented and unblocked threats, which proceed to impact resource either or both of availability and integrity; and 2) detection, counter and recovery processes that relieve the system of the threat and recover the resource to full availability and/or integrity overtime. The impact on the resource depends on 1) the current threat impact level, which is based on whether the threat has been detected and, once detected, the level of threat counteraction acting on the threat; and 2) the level of resource recovery that is being effected after counteraction has been completed:

**Figure 27 - Example #2 Instantiated Model (Resource Layer)**

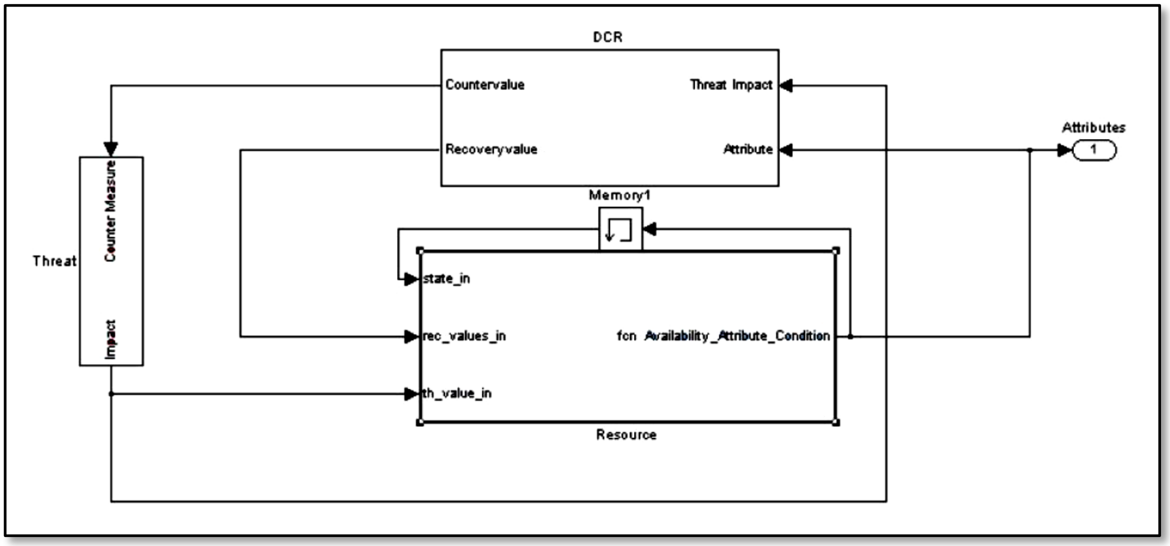

(Tjoa, Jakoubi et al. 2011)

**Figure 28 - Example #2 Instantiated Model (Threat Layer)**

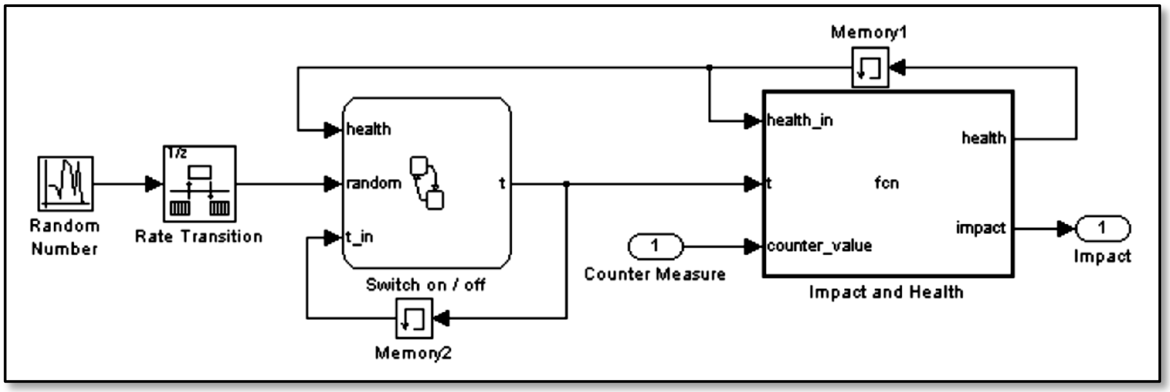

(Tjoa, Jakoubi et al. 2011)

**Figure 29 - Example #2 Instantiated Model (Detection, Counter and Recovery Layer)**

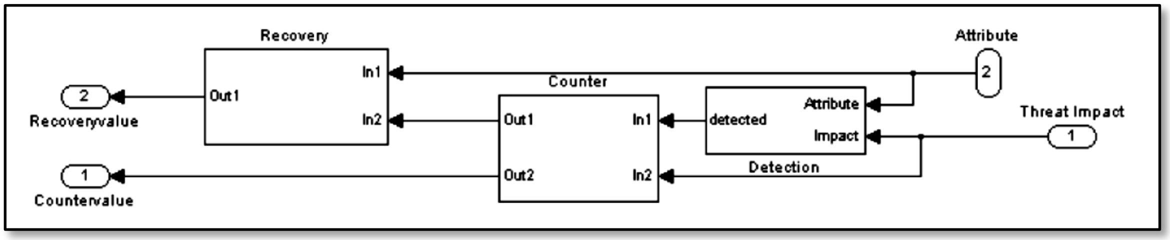

(Tjoa, Jakoubi et al. 2011)

Conceptually, the per period effectiveness of the detection, counter and recovery activities and the resulting reduction in threat impact and resource availability are modelled essentially as deterministic parametric functions, although effectiveness could be modelled as stochastic probability distributions as required.

The resulting simulations indicate the effects of the two different proposed blocking controls. The following simulation outputs illustrate the differences on outcomes:

**Figure 30 - Example #2 Control Scenario Comparison**

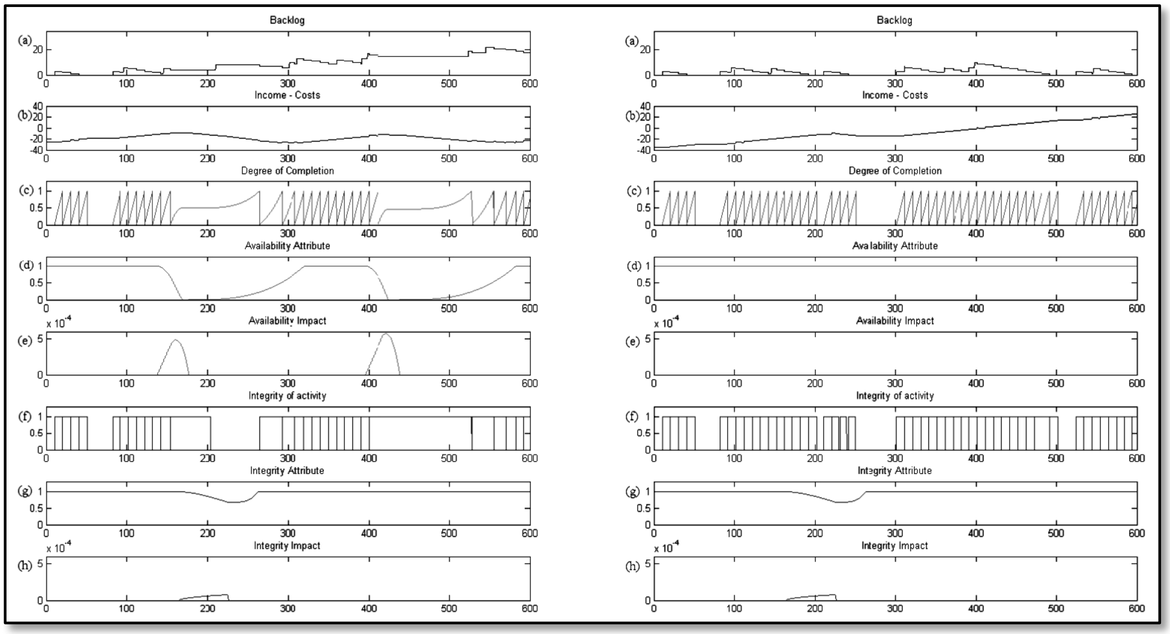

**Simulation #1: Availability Blocking Measure: 75%** **Simulation #2: Availability Blocking Measure: 95%**

(Tjoa, Jakoubi et al. 2011)

In the first simulation, two successful (unblocked) threats enter the system, first at approximately time 150 and again at time 380, and the Availability attribute declines to zero within about 50 time units until the counter measure is completely effective. At that point the Recovery measure initiates and takes nearly 150 time units to fully recover the availability of the resource. The Backlog and Income plots indicate that the DCR controls take some time to counteract the threats leading to an increasing (i.e. unrecoverable at current resources) call backlog and fluctuating income due to sustained business losses during the periods when the activities are compromised. Comparatively, the second simulation indicates that all threats were successfully blocked resulting in no decrease in resource availability. Correspondingly, call backlog fluctuates but does not increase over time and income rises over the simulation period. Since there are no Availability impacts in simulation 2, the impact of the single integrity threat can be seen clearly beginning at time 180 which, due to the relative shorter counteraction and recovery times, temporarily slows degree of completion but does not have significant resulting impact on either call backlog or income. The authors conclude by indicating that the model has demonstrated the ability to support decisions regarding security controls in a way that is informed by the incorporation of security-related control elements into the model which affect the availability and integrity of the underlying resources on which business processes depend.

### Discussion and Implications for Application within this Research

This modelling approach was significant in my research for several reasons: 1) Contrary to the procedural modelling approaches suggested by the HP research, it provided an intuitive, GUI-based approach to discrete event simulation (DES) for a business process at risk of security threats; 2) the approach incorporated relatively realistic threat impact and security control sub-processes for the computing resources which could be easily replicated within other prospective models involving business process resource and risk dependencies. Specifically, I liked their incorporation of the prevent/block/detect/counter/recover schema which directly reflected practitioner experience in designing ‘defense in depth’ security control systems; 3) the models , although only minimally dependent upon stochastic inputs in these cases (i.e. for the call queues and the threat incidence generation) clearly supported the stochastic determination of any or all inputs (e.g. control efficacy) as might be required to increase the realism of the control sub-processes; 4) the outputs of the model were time based and could be set up to produce any length of simulation that might be required; 5) the platform on which the models were built (MatLab and its DES sub-platform, Simulink) is commercially available at a reasonable cost for exploration at the decision stage in which I was considering its effectiveness for my purposes. This was particularly important since I had at that point not decided on a platform since the implications for ‘locking-in’ were significant at that point in time and required careful consideration; 6) it was reasonable to assume that the modelled examples (i.e. as modelled within MatLab by the authors themselves) might be readily available from the academic authors for academic review and further research use, unlike the HP models which were essentially proprietary and commercial in nature.

On this basis, I decided to pursue the platform further to determine whether I could modify the existing models for my purposes. Through personal correspondence with the authors, I was able to obtain the MatLab model for Case #2 above to understand the benefits and limitations of working with the platform and the prospect for modifying it for my own purposes. While the model presented a valid simulation of an information system at risk from a practitioner perspective, as implemented there were several significant considerations for improvements that would have to be overcome before it would be considered an appropriate modelling approach and platform for my use: 1) the model was essentially deterministic (i.e. other than the stochastic queues and threat instances it always produced the same output results) since it did not yet incorporate any type of realistic sub-component threat attack model or otherwise stochastic attributes for security threats, vulnerabilities or controls. This meant it would likely require significant modification to achieve any degree of input and output ‘uncertainty’ profiles appropriate for my lab purposes and it was not immediately clear how those modifications would be accomplished in a reasonable period of development time, despite the GUI nature of the tool, since I was not familiar with the Matlab environment; 2) It was not clear at that stage of my review whether or how the model could be reasonably run in ‘real time’ within a lab setting that would facilitate an adequate level of user interaction between the simulation (i.e. interactive control over control types or levels of control) and the individual decision scenarios [35]; 3) it was not clear how individual instances of the simulation could be run on independent participant workstations (and if performed in the field, possibly using their own corporate workstations) at a reasonable software licensing cost; 4) even if multiple versions of the software were available to me, it was not clear how the software would be conveniently installable in the field on workstations I did not own or control prior to the lab sessions; 5) it was not clear whether or how the model could be modified to run with a suitable 'front end' overlay onto the MatLab modelling environment itself that users could thereafter manipulate themselves during the labs; 6) It was not clear whether or how input selection parameters could be easily changed by a user using some type of user friendly/user-controllable 'widget' (e.g. such as a slider or press buttons) or whether the outputs which (as illustrated above) could be abstracted from the monitoring 'scopes' used within the MatLab environment so far. From my perspective, there was likely a need to be able to create input tools (sliders, buttons, etc.) and graph outputs that mode closely matched something available in Visual Basic in Excel – the comparative devices in MatLab were at that point unknown to me.

For these reasons, after reviewing and manipulating the available model for some weeks, I decided that I should continue to explore alternative platforms to address the potential issues noted. On the plus side, the model did, however, provide invaluable perspective on the suitable conceptual framework for security control grouping that I would eventually incorporate consisting of 'preventive', 'blocking' and 'reactive' (detection, counter-measures and recovery) controls (Tjoa, Jakoubi et al. 2011).

**3 – Use of Security Ontologies and Bayesian Networks for Threat Probability Determination**

The research by Tjoa et al (Tjoa, Jakoubi et al. 2011) led to my review of related but different approaches by other members of the SBA Research group. Fenz et al have undertaken extensive research in threat determination (Ekelhart, Fenz et al. 2006), security ontologies (Ekelhart, Fenz et al. 2006) , the use of security ontologies for threat determination (Fenz and Neubauer 2009), the use of security ontologies within Bayesian networks ((Fenz, Tjoa et al. 2009; Fenz 2012) and the construction of Bayesian networks based on security ontologies using a prototype modelling tool (Ekelhart, Fenz et al. 2009; Fenz, Ekelhart et al. 2011; Fenz 2012). Their approach is essentially characterized by the integration of two relatively recent security modelling approaches, each addressing a specific aspect of the security risk modelling problem: 1) the specification of a *security ontology* which formally defines the relationship between threats, assets, vulnerabilities, and controls (Fenz and Ekelhart 2009); and 2) threat probability (success) determination based on the calculation of the conditional probability of the system attributes using Bayesian networks (Fenz and Tjoa 2008; Fenz and Neubauer 2009; Fenz, Tjoa et al. 2009). In addition, the authors implement a prototype of a software construction of the Bayesian network using the defined security ontology based

[35] As it turned out, this would not have been an issue for any of the simulation platforms reviewed since the lab interfaces themselves were eventually built using pre-calculated simulation outputs. This was done to accommodate models which ultimately could not be simulated in a short enough time to be realistically accommodated on a standard desktop computer within a lab setting and to eliminate the need for the acquisition and deployment of a proprietary simulation platform, and likely requiring multiple copies to accommodate group trials setting, within the lab interface itself.

on the *Aurum* platform[36]. The conceptual ontological approach and the Bayesian determination of threat probabilities are reviewed below. I then discuss my consideration for using this modelling approach and the key insights that led to further research and ultimately the adoption of conditional probabilities directly within my chosen model and software platform.

In the first instance, the authors advocate for the use of security ontologies to enable the definition, formalization, reuse and machine computability of security relationships. An ontology "…defines the basic terms and relations comprising the vocabulary of a topic area, as well as the rules for combining terms and relations to define extensions to the vocabulary" (Neches, Fikes et al. 1991) or is "…a hierarchically structured set of terms for describing a domain that can be used as a skeletal foundation for a knowledge base." (Swartout, Patil et al. 1997). Most notably, from the perspective of developing a security ontology that can be incorporated into a software model of system risk, an ontology is "…a formal, explicit specification of a shared conceptualization." (Studer, Benjamins et al. 1998), where

- "formal" means machine-readable
- "explicit specification" means entity concepts, properties, relations, functions, constraints, and axioms  are explicitly defined
- "shared" means consensually based knowledge of the explicit specification; and
- "conceptualization" means an abstract model and simplified view of some phenomenon in the world that we want to represent

On reflection, these definitional criteria are important for my consideration of a security model which will need to 1) demonstrate real world validity (i.e. it is both 'explicit', defining structured elements and their logical relationships over which relations and therefore calculations can be made; and shared, such that it is representative of a recognizable standard within a professional practitioner or academic community on which members rely for a common understanding of events and causal inference); 2) be both conceptualized (simplified) enough to sufficiently define and be implementable within a constrained time or resource modelling approach; and 3) be convertible into a machine readable and calculable format – i.e. in some type of modelling software - which supports scaled and scoped simulation of the model in some available (consumer grade) computing system.

The logical security relationships on which the ontology is built are illustrated below:

[36] AURUM is a proprietary information security risk management platform originally developed by the SBA group and now offered as part of the consulting services of XYLEM Group - http://www.securityontology.com/

**Figure 31 - Logical Security Concept Relationships**

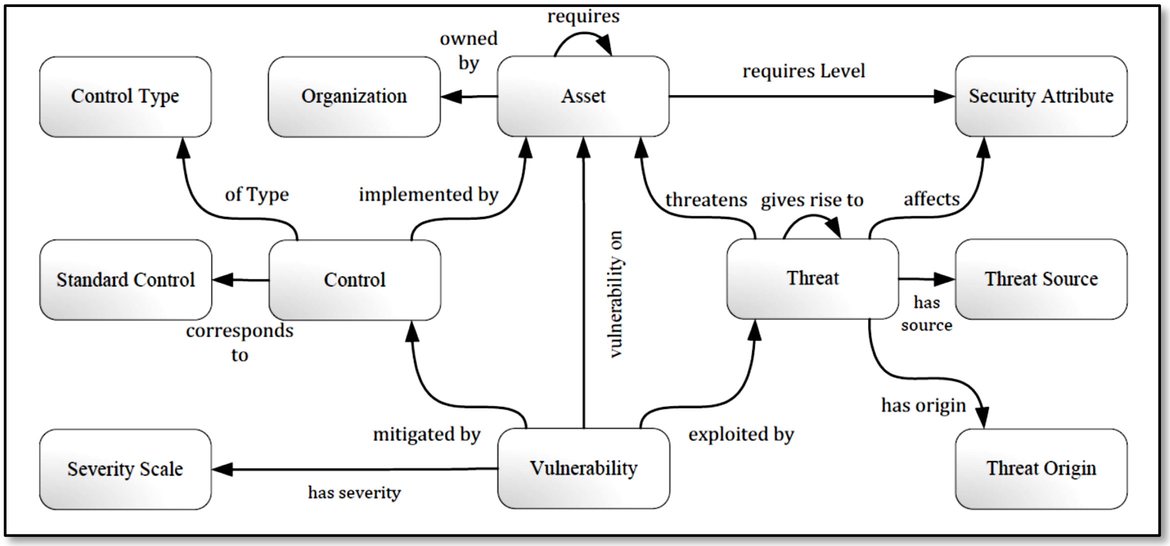

(Fenz, Tjoa et al. 2009)

In the above diagram, the logical relationships between the security components correspond to established security practitioner standards (NIST October 1995): system and process assets, which are subject to threats which affect the assets security attributes (confidentiality, integrity and availability, etc.) due to inherent or control effectiveness vulnerabilities, which can be mitigated using security controls affecting threat prevention, blocking, detection, counteraction and asset recovery (Fenz, Tjoa et al. 2009). An example of the logical specification of the risk of  a 'fire threat' would be as follows where the practitioner guidance for the threat, control and impact relationships can be traced to both the GSHB[37] and ISO/IEC 27001[38] standards (Fenz and Ekelhart 2009):

[37] IT Grundschutz Manual, 2004 - a German standard for IT security based on the British Standards Institute https://www.bsi.bund.de/EN/Topics/ITGrundschutz/itgrundschutz_node.html
[38] ISO/IEC. ISO/IEC 27001:2005, Information technology - Security techniques – Information security management systems - Requirements, 2005 http://www.iso.org/iso/home/standards/management-standards/iso27001.htm

**Figure 32 - Fire Threat Example**

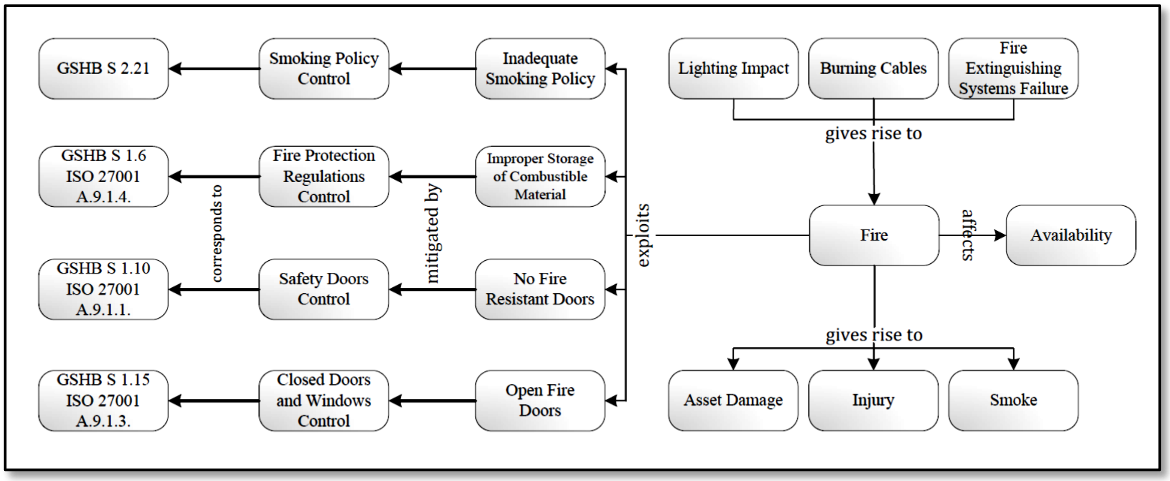
GSHB S 2.21
Smoking Policy Control
Inadequate Smoking Policy
Lighting Impact
Burning Cables
Fire Extinguishing Systems Failure
gives rise to
GSHB S 1.6 ISO 27001 A.9.1.4.
Fire Protection Regulations Control
Improper Storage of Combustible Material
corresponds to
mitigated by
exploits
Fire
affects
Availability
GSHB S 1.10 ISO 27001 A.9.1.1.
Safety Doors Control
No Fire Resistant Doors
gives rise to
GSHB S 1.15 ISO 27001 A.9.1.3.
Closed Doors and Windows Control
Open Fire Doors
Asset Damage
Injury
Smoke

(Fenz and Ekelhart 2009)

The logical expression of the ontology allows the calculation of the threat probability for any associated asset:

1) Attacker motivation and capability affect attacker effectiveness.
2) Attacker effectiveness combines with the a prior probability of the threat occurring and affects the probability of exploiting a vulnerability.
3) The probability of exploiting a vulnerability is also a function of combined control effectiveness across the relevant asset controls.
4) The resulting vulnerability exploitation probability is summed over all vulnerabilities and is combined with the additional probability of threats resulting from predecessor threats to determine the overall threat probability,
5) and any ensuing successor threat probabilities

The order of operation of the logical relationships are illustrated below:

**Figure 33 - Logical Security Model Order of Operation for Threat Determination**

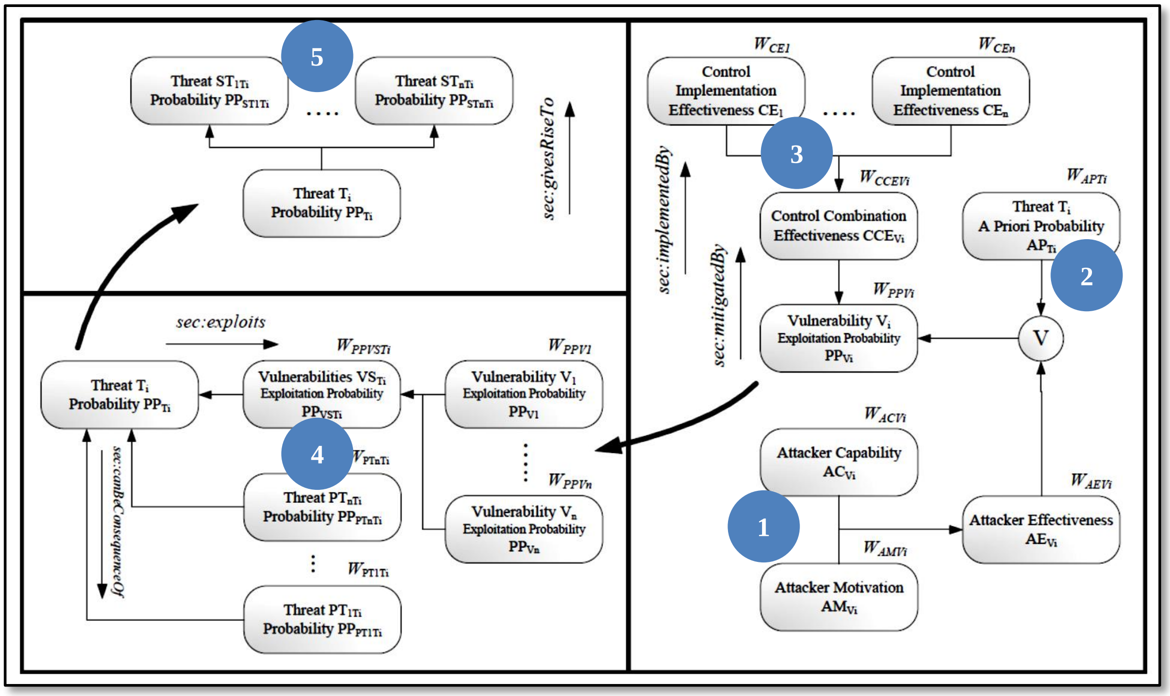
Threat $ST_{1T_i}$ Probability $PP_{ST1T_i}$
5
....
Threat $ST_{nT_i}$ Probability $PP_{STnT_i}$
sec:givesRiseTo
Threat $T_i$ Probability $PP_{T_i}$
$W_{CE1}$
$W_{CEn}$
Control Implementation Effectiveness $CE_1$
....
Control Implementation Effectiveness $CE_n$
3
$W_{CCEVi}$
$W_{APTi}$
sec:implementedBy
Control Combination Effectiveness $CCE_{V_i}$
Threat $T_i$ A Priori Probability $AP_{T_i}$
2
sec:mitigatedBy
$W_{PPVi}$
Vulnerability $V_i$ Exploitation Probability $PP_{V_i}$
V
sec:exploits
$W_{PPVSTi}$
$W_{PPV1}$
Threat $T_i$ Probability $PP_{T_i}$
Vulnerabilities $VS_{T_i}$ Exploitation Probability $PP_{VST_i}$
Vulnerability $V_1$ Exploitation Probability $PP_{V1}$
$W_{ACVi}$
Attacker Capability $AC_{V_i}$
4
$W_{PTnTi}$
$W_{PPVn}$
sec:canBeConsequenceOf
Threat $PT_{nT_i}$ Probability $PP_{PTnT_i}$
Vulnerability $V_n$ Exploitation Probability $PP_{Vn}$
1
$W_{AEVi}$
Attacker Effectiveness $AE_{V_i}$
$W_{AMVi}$
$W_{PT1Ti}$
Attacker Motivation $AM_{V_i}$
Threat $PT_{1T_i}$ Probability $PP_{PT1T_i}$

(Fenz and Ekelhart 2009)

The fully *encoded* ontology consists of the concepts ('component nodes'), relations ('linking functions' expressing relationships between the nodes) and axioms ('node scales' – potential states of the node, and 'node weights' – parent/child weights, equal unless specifically weighted for certain vulnerabilities) representing the full security schema which will be used to construct the Bayesian network itself:

**Figure 34 - Concepts Nodes for Security Ontology**

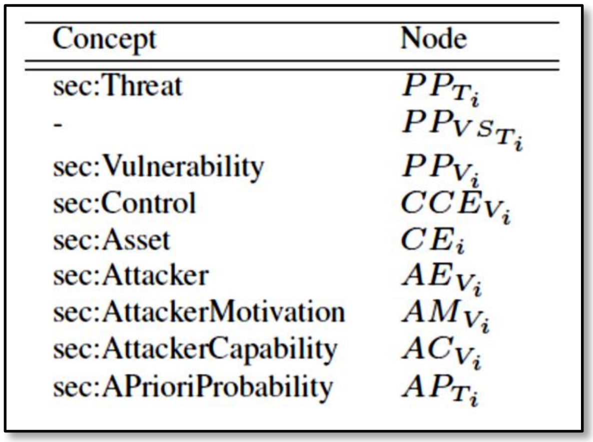

| Concept | Node |
|---|---|
| sec:Threat | $PP_{T_i}$ |
| - | $PP_{VS_{T_i}}$ |
| sec:Vulnerability | $PP_{V_i}$ |
| sec:Control | $CCE_{V_i}$ |
| sec:Asset | $CE_i$ |
| sec:Attacker | $AE_{V_i}$ |
| sec:AttackerMotivation | $AM_{V_i}$ |
| sec:AttackerCapability | $AC_{V_i}$ |
| sec:APrioriProbability | $AP_{T_i}$ |

(Fenz and Ekelhart 2009)

**Figure 35 - Security Ontology Concept Relations**

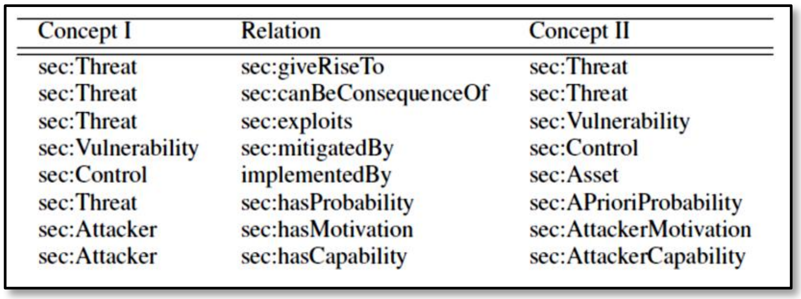

| Concept I | Relation | Concept II |
|---|---|---|
| sec:Threat | sec:giveRiseTo | sec:Threat |
| sec:Threat | sec:canBeConsequenceOf | sec:Threat |
| sec:Threat | sec:exploits | sec:Vulnerability |
| sec:Vulnerability | sec:mitigatedBy | sec:Control |
| sec:Control | implementedBy | sec:Asset |
| sec:Threat | sec:hasProbability | sec:APrioriProbability |
| sec:Attacker | sec:hasMotivation | sec:AttackerMotivation |
| sec:Attacker | sec:hasCapability | sec:AttackerCapability |

(Fenz and Ekelhart 2009)

**Figure 36 - Security Ontology Child-Parent Relationships**

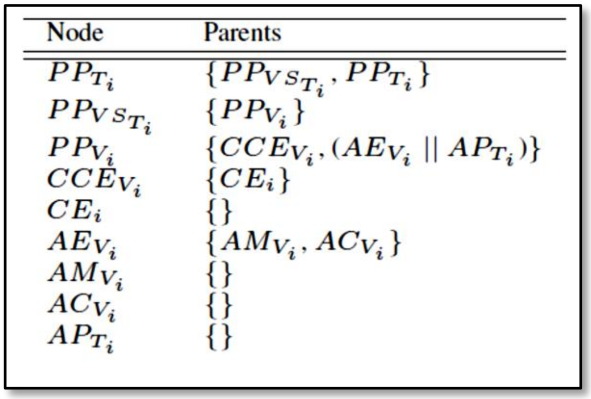

| Node | Parents |
|---|---|
| $PP_{T_i}$ | $\{PP_{VS_{T_i}}, PP_{T_i}\}$ |
| $PP_{VS_{T_i}}$ | $\{PP_{V_i}\}$ |
| $PP_{V_i}$ | $\{CCE_{V_i}, (AE_{V_i} \mid\mid AP_{T_i})\}$ |
| $CCE_{V_i}$ | $\{CE_i\}$ |
| $CE_i$ | $\{\}$ |
| $AE_{V_i}$ | $\{AM_{V_i}, AC_{V_i}\}$ |
| $AM_{V_i}$ | $\{\}$ |
| $AC_{V_i}$ | $\{\}$ |
| $AP_{T_i}$ | $\{\}$ |

(Fenz and Ekelhart 2009)

The concepts are then ‘instanced’ by indicating relevant known real world evidence for each node using Likert scales or probabilities: presence/absence; high/medium/low; attacker motivation; attacker capability; control implementation effectiveness; a priori threat probability, etc.

### Use of Bayes Theorem for Determination of Security Threat Probabilities

In standard security risk models, the expected value of a security risk to an asset is based on the probability of a threat successfully impacting the asset multiplied by the attributed impact given a successful threat. The resulting expected loss can be plotted as a 2 dimensional ‘heat map’ indicating relative expected risks across all threat and asset combinations:

**Figure 37 - Security Risk Assessment 'Heat Map'**

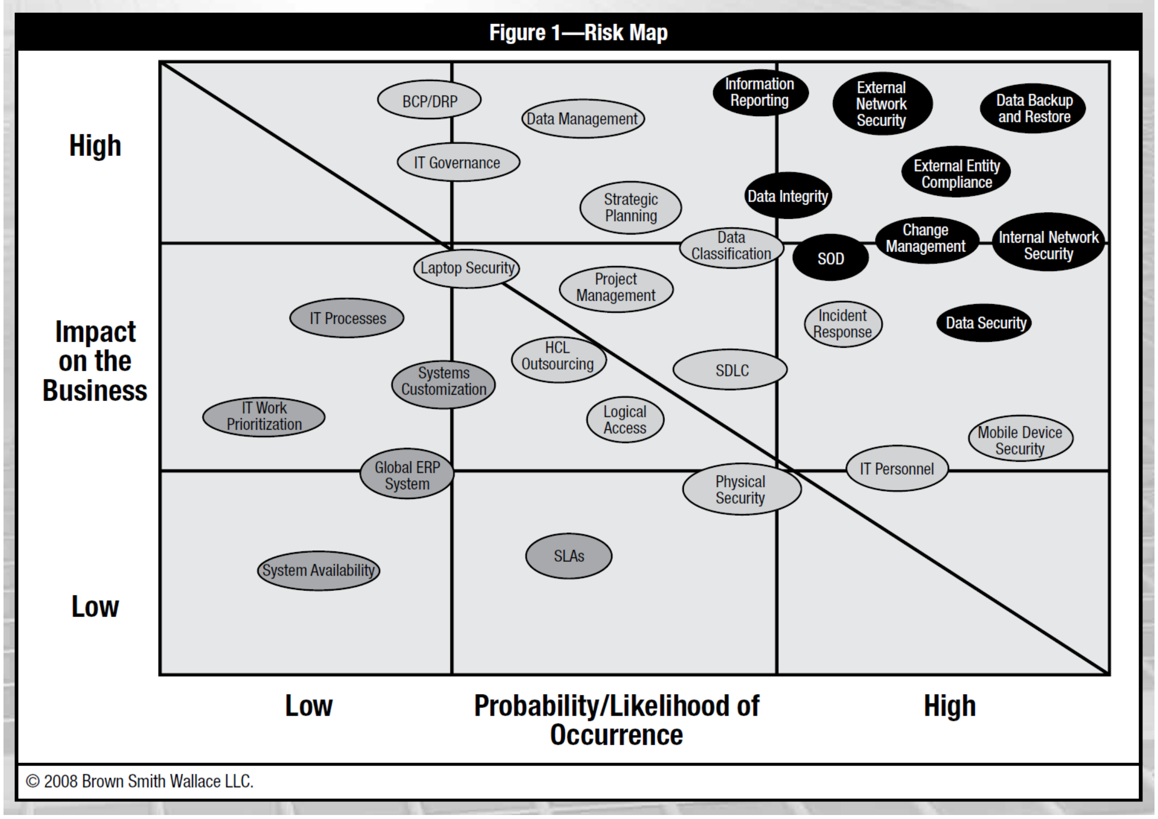

(Schmittling and Munns 2010)

The probability of threat success is assumed to be *conditional* based on the underlying *nature* of the system's risk factors: the presence and interaction of attacker motivations, asset vulnerabilities, control effectiveness, etc. In this model, certain combinations of factors are assumed to generally (but not necessarily) produce more 'successful' threats than others: systems with a given vulnerability and a relatively weaker control (or no control) produce more successful threats than systems with the same vulnerability but with a relatively stronger control. The combination of risk factors can be considered as a vector of factors which combine to create the likelihood of a threat occurring and is essentially a causal model of threat occurrence. If the relationship between the risk factors and successful threats were purely deterministic (i.e. if the presence of a particular vector of risk factors always/never caused a successful threat) then, given information about the presence of those risk factors, we would be sure of the probability of a successful threat (100% or 0% respectively).

In the absence of sufficient information about the underlying risk factors (or the absence of a posited model of threat probability), the best we can do to estimate the probability of a successful threat is to observe the *frequency* of successful threats in an actual system over time. The longer we observe, given the law of large numbers, the better is our estimate of the true probability although we can say nothing about the underlying cause of the threats or how the probability might differ if, or as, the risk factors change. Presuming that the factors do in fact change over time in a real world setting, the probability based on frequency or proportion of outcomes (successful threats over total system activity periods observed) simply reflects the average of all types of possible risk factor effects and their resulting threat probability. The problem of threat probability determination asks: can we do better?

As a useful analogy for what follows, we can think of a 10-sided die as being the 'system at risk', where the die has faces that either say "Successful Threat" or "Unsuccessful Threat" and a roll of the die produces a random binary outcome: Successful or Unsuccessful. Imagine that you do not know what the proportion of 'Successful' to 'Unsuccessful' faces is or how the faces are assigned to the die – how would you estimate the probability of rolling a 'Successful' face on a single throw (i.e. by analogy, experiencing a 'successful threat' to the system)? In the absence of other information, intuitively, the best guess is 50/50. On the one hand, if you could run an *experiment*, you might roll the die a thousand times and compare the proportion of 'Successful' to 'Unsuccessful' faces and conclude that this proportion is a good approximation for the 'underlying' probability of a successful threat. This might be satisfactory, but only if you: 1) had the luxury of running experiments (observing outcomes without intervening) – note this might be costly in a real world setting with an actual information system at risk or, analogously, even with dice if you were betting on the outcomes; or 2) had no chance to examine the die or some portion of the die directly in order to know fully or estimate what the actual proportion of 'Successful' to 'Unsuccessful' faces might be i.e. how risky the system was in fact; or 3) to understand or (better) somehow influence the way in which the values 'Successful' and 'Unsuccessful' are 'assigned' to the faces. In the latter case, we might imagine some malevolent die painter who prefers to paint Success or Unsuccessful, but we don't know which. In a real world security setting, we can certainly imagine both malevolent attackers and lazy security control managers where each would have an effect on the underlying probability of successful threats.

If you could examine the die, purely from the perspective of determining the probability of successful threat, it would not matter by what means the faces had been assigned, only that there was a specific proportion representing the objective probability of 'successful' assuming the proportion did not change over the course of the experiment. In this case, you do not need to know anything more about the creation of the die and, in fact, don't need to roll the die at all to forecast the probability of events related to its possible outcomes: the probability of 'successful' can be known for sure. Comparatively, having to roll the die one roll at a time is equivalent to randomly examining the individual faces on the die, but without ever knowing the true underlying proportions. In this latter case, additional information on the general nature of the die could provide some guidance as to its magnitude e.g. if you knew roughly how many 'successfuls' were on the die (i.e. the presence or absence of some but not all risk factors) or by what mechanism the faces were 'assigned' "Success" over "Unsuccessful". Notice that this knowledge or assumption of the underlying ‘mechanics of assignment’ is especially valuable if we only have a limited number of rolls in which to subjectively forecast the probability. How would knowledge of this mechanic improve the ability to estimate the probability of a successful threat?

In reality, what is generally observed are systems that have certain risk factors present but which experience varying probabilities of successful threats, and for the same system over time. This means that the probability of a successful threat must somehow be *statistically inferred* from the relationship between successful threats and the presence and interaction of the risk factors. The inference of the probability depends on a subjective interpretation of what this probability represents: either as a *proportion of outcomes* as noted above or, alternatively, as a *degree of belief* based on the conditional relationship between threat success and the risk factors. A *frequentist* approach to the calculation of the probability of a successful attack P(A) is simply the long run relative frequency or proportion of outcomes of, “A” out of all possible outcomes assuming some underlying threat generation mechanism of the system, “B” (i.e. without the ability to see or influence the threat generation, or simply ignoring it). On the other hand, from an epistemological perspective, taking evidence of the risk factors into account, a better alternative interpretation of the probability of A can be understood as the probability of A *conditional on the evidence of B.* The probability of A would then be derived as follows:

By the commutative property[39],

$$P(A\ and\ B) = \ P(B\ and\ A) \tag{4.1}$$

If A and B are independent events, then

$$P(A\ and\ B) = \ P(A) * P(B) = P(B) * P(A) \tag{4.2}$$

If B depends on A, we write

$$P(B|A) \tag{4.3}$$

And the probability of A and B is then

$$P(A\ and\ B) = \ P(A) * P(B|A) \tag{4.4}$$

Similarly, if A depends on B then

$$P(B\ and\ A) = \ P(B) * P(A|B) \tag{4.5}$$

Since the right sides of 4.0 and 5.0 above are equal, then

$$P(B) * P(A|B) = \ P(A) * P(B|A) \tag{4.6}$$

[39] http://mathworld.wolfram.com/Commutative.html

Dividing both sides by P(B), we get *Bayes' Theorem*:

$$P(A|B) = \frac{P(A)*P(B|A)}{P(B)} \tag{4.7}$$

Bayes Theorem says that the probability of a threat P(A) is a matter of a *degree of belief* conditional on 1) a *prior* assumption of the probability of the threat P(A); 2) the *likelihood* of the presence of the risk factors *given a successful threat* P(B|A); and 3) the probability of the risk factors occurring P(B). This allows the probability of a successful threat to vary depending on the *likelihood* of the underlying stochastic process of how the risk factors create threats, when a threat actually occurs. This model itself, not just the risk factors, can vary and thus contributes to the probability of a threat occurring. Comparing this approach to the frequentist approach, P(B) can be thought of as *evidence* of risks which are not considered when simply observing the frequency of threats P(A) but when introduced, imply a risk model and change your belief in the underlying probability of threats. The factor $\frac{P(B|A)}{P(B)}$ therefore represents the 'total' impact of B (i.e. both the underlying risk model combined with the existence of the risks factors) on the probability of A.

The Bayesian approach can dramatically change the estimation of probabilities depending on the incidence of the threat and the presence of the risk factors. By another analogy, if a coin is flipped 99 times, and each time a head comes up, what would you expect to be the probability of a head turning up on the 100th flip? The answer is that *it depends on whether you believe the coin is fair* (and therefore that the probability of heads is 50%). If you are a frequentist and you assume the coin is fair, you would bet that the probability is 50% regardless of the evidence to the contrary. You would correspondingly *believe* that the probability of 99 heads is a row is possible, but highly unlikely (also since you assume the coin is fair). If you are a Bayesian, you would incorporate the evidence of the 99 heads to update your *assumption* of a fair coin (i.e. a 50% chance of flipping heads) in estimating the probability of heads on the next flip. In the context of threat determination, the analog is that we have a system (a coin) that produces successful threats (heads) with some probability *based on the underlying threat producing characteristics of the system* (fairness of the coin). The probability of a threat is then analogous to the probability of tossing a head which depends on our belief in the nature of the risk producing system (fairness of the coin).

For example, if the prior expectation of the probability of successful threats P(A) = 1.1% (a successful threat occurring 4 days in 365) but the probability of a vulnerability being present is P(B) = 3.3% (vulnerabilities occur 12 days in 365) and the probability of a vulnerability being present given a successful threat is P(B|A) = 99% (i.e. vulnerabilities nearly always cause threats), then

$$P(successful\ threat|vulnerability) = \frac{P(successful\ threat)P(vulnerability|successful\ threat)}{P(vulnerability)}$$

$$P(successful\ threat|vulnerability) = \frac{.011 * .99}{.033} = .33 = 33\%$$

Assume instead that the vulnerability is only found to occur in 75% of the threat cases. Then

$$P(successful\ threat|vulnerability) = \frac{.011 * .75}{.033} = .25 = 25\%$$

For the purposes of modelling threats which are assumed conditional on the presence of risk factors, this is a much more valuable approach to threat probability determination since it incorporates both a model of threat determination and the evidence of the incidence of risks and threats.

**Use of Bayesian Networks for Modelling Complex Security Threat Probabilities**

In the above examples, threat probability was determined by a single risk factor notionally represented by a 'vulnerability' which is unrealistic for the purposes of modelling even a small scale real world system at risk. In reality, threats are caused by combinations of risk factors (including vulnerabilities) that influence the probability of a threat occurring. To incorporate a more realistic approach, researchers have proposed the use of Bayesian networks to model 'attack graphs' which can be used to compute the *joint conditional probabilities* of threats for entire systems (Holm, Ekstedt et al. 2012; Kiesling, Ekelhart et al. 2014). The approach is based on Bayesian probabilities but supports deeper risk relationships between the risk factors. Bayesian networks are 'graphical structures for representing the probabilistic relationships among a large number of variables and doing probabilistic inference with those variables' (Neapolitan 2004).

For example, suppose that a successful threat is essentially the result of a malicious actor attack (e.g. the introduction of a computer virus) acting on an asset vulnerability (e.g. an unpatched operating system) in an organization that has an associated virus control (e.g. patching procedures) which reduces but does not eliminate all potential viruses. The resulting risk relationships leading to the probability of a successful threat can be expressed as a Bayesian network and a set of conditional probabilities representing the relationship between the risk factors. Here the probability of a successful threat is conditional on the presence of a virus (the 'threat') acting on a present system vulnerability (C), but possibly mitigated by the presence of a vulnerability patching control (B). The *conditional probabilities* between the risk factors can be represented in binary lookup tables indicating the result of the presence or absence of the factors on the dependent network node:

**Figure 38 - Three Node Bayesian Network**

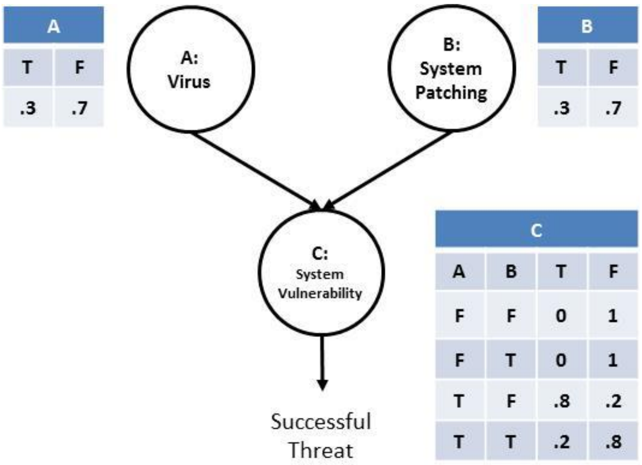

| A | |
|---|---|
| T | F |
| .3 | .7 |

| B | |
|---|---|
| T | F |
| .3 | .7 |

| C | | | |
|---|---|---|---|
| A | B | T | F |
| F | F | 0 | 1 |
| F | T | 0 | 1 |
| T | F | .8 | .2 |
| T | T | .2 | .8 |

Note that the sophistication of the inferred system model has increased substantially: 1) if there is no virus present (A=F), the system vulnerability cannot lead to a successful threat; 2) if there is a virus present but no patching is present, there is an 80% chance that the virus will act on the vulnerability and create a successful threat(C=T=.8); 3) on the other hand, if both a virus and patching is present, there is still a chance that a virus can act on the vulnerability and subsequently cause the threat to be successful, but with less probability (C=T=.2).

The modelling question we are interested in is ‘what is the probability of a successful threat given the presence of one or more risk factors?’. P(C|A,B) can be directly calculated by summing over the exhaustive set of conditions of A, B and the resulting C:

**Table 5 - Conditional Probability Calculation**

| | A,B,C | | | | |
|---|---|---|---|---|---|
| | **T,T,T** | **T,F,T** | **F,T,T** | **F,F,T** | |
| **P(A)** | 0.3 | 0.3 | 0.7 | 0.7 | |
| **P(B)** | 0.3 | 0.7 | 0.3 | 0.7 | |
| **P(C )= T** | 0.4 | 0.4 | 0.4 | 0 | **Sum** |
| **P(A)*P(B)*P(C )** | **0.036** | **0.084** | **0.084** | **0** | **0.204** |

Here the probability of a successful threat <u>over all possible combinations of risk inputs</u> is 20.4%. This percentage (like any average) is interesting but not overly helpful to a decision maker more interested in making specific control choices than understanding what the average probability of success would be over all possible combinations of threats and controls. A more interesting question from an interventional perspective is: “What is the probability of a successful threat given the presence of patching?’ In this case, we use the Bayesian probability calculation noted above:

$$P(C|B) = \frac{P(C)*P(B|C)}{P(B)} \tag{4.8}$$

The question here is how to estimate the *likelihood function* P(B|C) since C is also dependent on A.

P(B|C) can be written as a conditional probability and summing over all nuisance variables (in this case A), where T equals the condition “TRUE”:

$$P(B|C) = P(B = T|C = T) = \frac{P(C=T,B=T)}{P(C=T)} = \frac{\sum_{A\in\{T,F\}} P(C=T,A,B=T)}{\sum_{A,B\in\{T,F\}} P(C=T,A,B)} \tag{4.8}$$

Using the conditional probability tables above, we can then evaluate each of the terms in the sums - in the numerator:

$$\begin{aligned} &P(C = T, A = T, B = T) \\ &\qquad = P(C = T|A = T, B = T) * P(A = T) * P(B = T) \\ &\qquad = .4 * .3 * .3 \\ &\qquad = 0.036 \end{aligned} \tag{4.9}$$

$$\begin{aligned} &P(C = T, A = F, B = T) \\ &\qquad = P(C = T|A = F, B = T) * P(A = F) * P(B = T) \\ &\qquad = .4 * .7 * .3 \\ &\qquad = 0.084 \end{aligned} \tag{4.10}$$

In the denominator:

$$\begin{aligned} &P(C = T, A = T, B = T) \\ &\qquad = P(C = T|A = T, B = T) * P(A = T) * P(B = T) \\ &\qquad = .4 * .3 * .3 \\ &\qquad = 0.036 \end{aligned} \tag{4.11}$$

$$\begin{aligned} &P(C = T, A = T, B = F) \\ &\qquad = P(C = T|A = F, B = T) * P(A = F) * P(B = T) \\ &\qquad = .4 * .3 * .7 \end{aligned} \tag{4.12}$$

$$= 0.084$$

$$\begin{aligned} P(C = T, A = F, B = T) \\ &= P(C = T|A = T, B = T) * P(A = T) * P(B = T) \qquad (4.13) \\ &= .4 * .7 * .3 \\ &= 0.084 \end{aligned}$$

$$\begin{aligned} P(C = T, A = F, B = F) \\ &= P(C = T|A = F, B = T) * P(A = F) * P(B = T) \qquad (4.14) \\ &= .4 * .7 * .3 \\ &= 0.0 \end{aligned}$$

$$P(B|C) = P(B = T|C = T) = \frac{.036_{TTT} + .084_{TFT}}{.036_{TTT} + .084_{TTF} + .084_{TFT} + 0.0_{TFF}} \qquad (4.15)$$

$$P(B|C) = \frac{.12}{.204} = .588$$

We can now calculate the probability of a successful threat, given the actual presence of patching P(B). This allows the model to incorporate available evidence for the various risk factors in place and not just the likelihood of their effect. For example, if we determine that P(B) = .8 (i.e. that patching is highly likely to be present), then

$$P(C|B) = \frac{P(C) * P(B|C)}{P(B)} \qquad (4.16)$$

$$P(C|B) = \frac{1 * .588}{.8} = .73$$

This indicates that if the known probability of patching was 80%, and the prior expectation of successful breach (regardless of whether patching is present) was 100%, the *posterior* or conditional probability of a successful breach given an 80% level of patching would only be 73%. Intuitively, as the actual probability (or presence) of patching P(B) falls, the probability of a successful threat increases. The factors are constrained based on the underlying likelihood function P(B|A). Note on the one hand that the probability that patching is actually present has a lower limit based on the product of P(A) * P(B|A). More intuitively (and in the context of estimating the probability of threat, not the probability of patching which can be observed), we can recognize that the posterior probability of threat (i.e. conditional on patching) has an upper limit (i.e. less than 100%) based on the ratio of P(B|A)/ P(B), where this represents the effect of patching on threat success. This demonstrates a methodology for the determination of the conditional probability of successful threat in a multi risk factor model and illustrates the sensitivity of the *posterior* probability of successful threat to both the underlying *prior* probability of threat regardless of condition, the likelihood of risk factor effects and the presence of the risk factors.

**Model scaling and considerations for use within a risk management context**

The example of a Bayesian network for threat determination provided above has only three nominal risk factors (one attack, one control, and one vulnerability) the attribution of the conditional probabilities between the factors and the resulting threat probability calculation is therefore relatively straightforward. Clearly, in a more complex model representing an integrated business process involving multiple sub processes, multiple information assets (each with multiple attack vectors, vulnerabilities and associated controls) and multiple asset dependencies (i.e. successful threats on one asset possibly influencing successful threats on another) the construction of the Bayesian net itself and the network node relationships is also relatively complex. In their research, Fenz et al have specified the use of multiple tools to support model creation, including an open-source security ontology, a commercially available Bayes net generation platform (Norsys Netica[40]) and SBA group's proprietary risk management platform (AURUM) which is not currently available for open research use. While the ontology itself is potentially valuable for my specifying the relationships between the risk factors, the ontology is only valuable to me to the extent that 1) a representative Bayesian network can be practically specified from it and 2) the resulting threat probabilities can be appropriately used within a subsequent risk model that suits my simulation purposes. It should be noted that the Bayesian network itself does not produce a simulation of the system at risk in the same manner as a discrete event simulation which embeds the conditional relationships of the security factors but which essentially can produce a simulation of the system as the relationships are constructed (i.e. the core result of a DES model is the representation of repeated time period system states based on the underlying conditional probabilities). The consideration for me was whether to spend time constructing a Bayes network only to have to also then construct the simulation model as well as a separate representation of the security risk model. Since the goal of the research was to construct a set of lab experiments based on simulated results and not a superior simulation itself, this seemed to be prospectively not worth the effort. Although SBA has made the resulting Bayes net employed in their work available for research review (Fenz 2012) because of the size of the model it does not run on a free copy of Norsys Netica, I was not prepared to invest additional funds and time into the platform without further consideration of alternatives.

More importantly for my purposes in desiring to express inherent uncertainty in the model *inputs* (i.e. the risk factors) as well as the outputs, it should be noted that the Bayesian network described by Fenz et al effectively calculates an overall probability of threat success per asset node, where each node is a Boolean value indicating the presence or absence of the corresponding risk factor. Conceptually, this limits the model's ability to capture qualitative risk factor attributes and corresponding concepts like control 'effectiveness' where, even if the control is present, it may not be entirely effective and the effectiveness at

[40] http://www.norsys.com

any point in time could range, qualitatively, from “Very Low” to “Very High” and points in between or, quantitatively, at some percentage level across a continuous probability distribution from 0 to 100%. In the model described above, the presence of the control is determined (1/0) and the conditional impact on the dependent risk factor (e.g. vulnerability) is then also determined to be 1 or zero, where the effectiveness and the resulting impact could, conceptually, be any value between 1 and zero. Fenz acknowledges this limitation and indicates a possible solution proposed by Fenton et al involving the use of ‘ranked’ Bayesian network nodes which represent, for example, Likert scale qualitative variables that are abstractions of the underlying continuous quantities (Fenton, Neil et al. 2007). Recent applications of Fenton’s approach include software development fault prediction (Perkusich, Soares et al. 2015) and smartphone security (Herland 2015). In Fenton’s conception, ranked nodes allow both the incorporation of qualitative node values and weighting between nodes. For example, using the Bayesian network described above and weighting within and between nodes where the probability of the virus and the associated patching control offset each other but are now reported on a 5 point Likert scale perhaps representing the qualitative ‘strength’ of the Virus and the associated Patch respectively (instead of either being simply present or absent – ‘TRUE’ or ‘FALSE), we can instead produce a more realistic alternative specification:

**Figure 39 - Weighted Bayesian Network Example: Matched Virus vs. Patching Effectiveness**

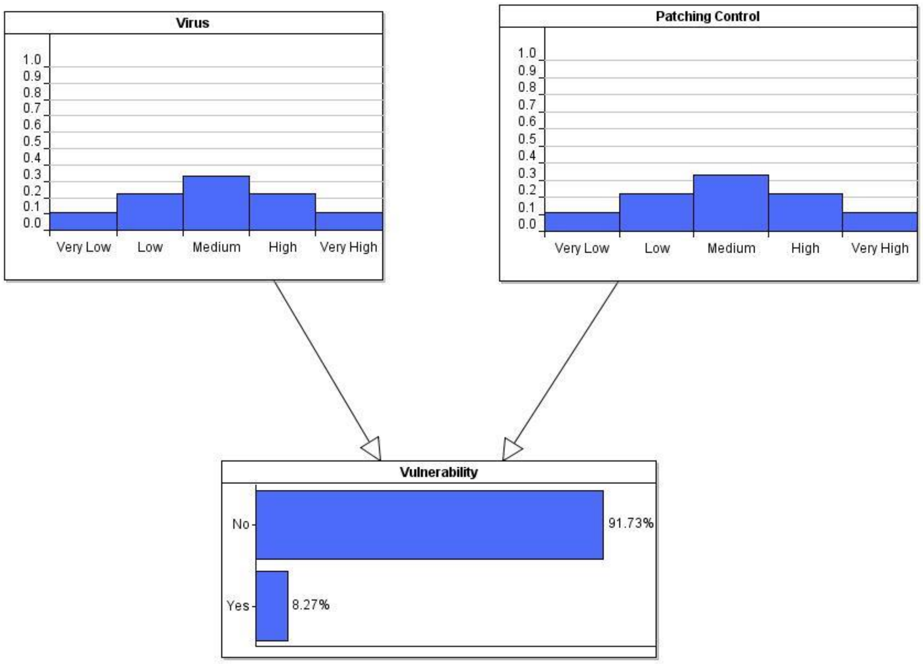

In the above example, in the absence of evidence, we expect  that both viruses and patching controls will be generally well matched in profile with the strength of the factors having a central tendency varying from “Very Low” to Very High” as indicated. If we assume a certain amount of variability (mismatch) in virus strength vs. patch, the patch blocks the virus 91% of the time but does not 8.27% of the time. On the other hand, if we expected the virus profile to be relatively ‘stronger’ (i.e. skewed towards “Very High”), we get the following result, where the probability that the patch does not block the virus rises to 16.5% since it is more likely that a stronger virus is not matched by an equally strong patch:

**Figure 40 - Weighted Bayesian Network Example: Skewed Virus Strength**

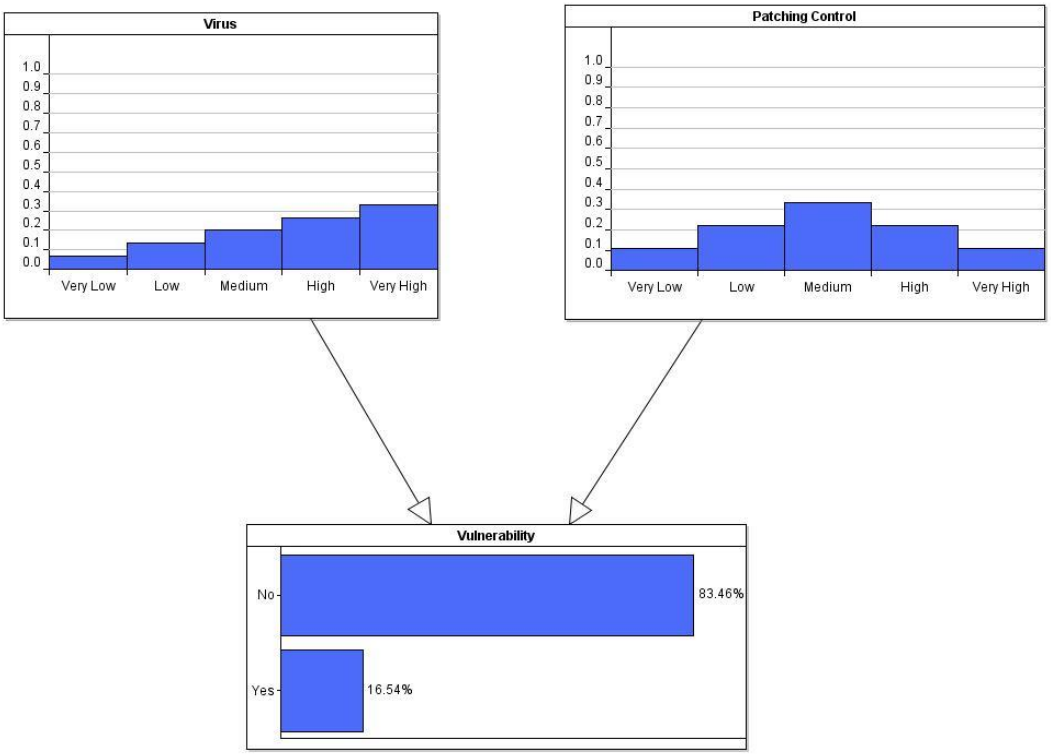

All Bayesian network models are able to incorporate evidence based on the probability that the node exist i.e. P(Virus) or P(Patching Control) regardless of the scale on which the node is represented. If we now assume that a ‘Very High’ level virus is actually present in the system, we can incorporate this evidence into the model which results in the probability of unsuccessful blocking rising to 36%:

**Figure 41 - Weighted Bayesian Network Example: Evidence of 'Very High' Virus Strength**

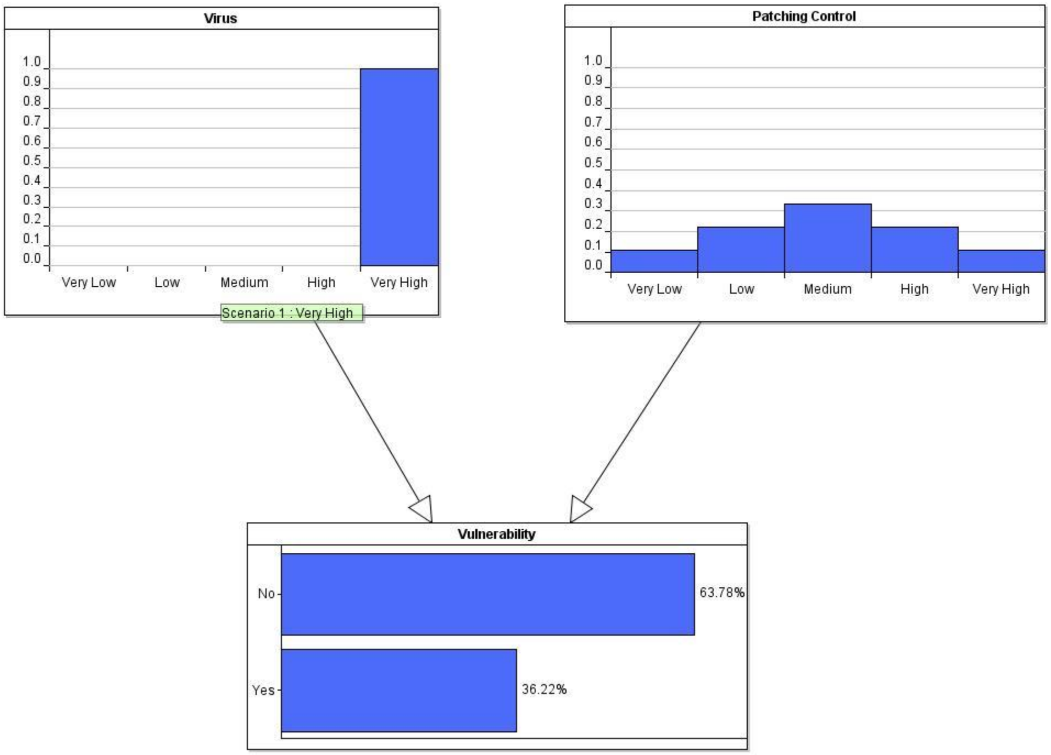

Finally, if we again assume a relatively 'stronger' Virus pattern, but discover evidence of 'Very Low' patching, the probability of a successful virus jumps to 58% since the patch is largely ineffective in all but a handful of 'Very Low' virus cases where most anticipated Viruses are relatively stronger than 'Very Low':

**Figure 42 - Weighted Bayesian Network Example: 'Very High Virus vs. Very Low Patching**

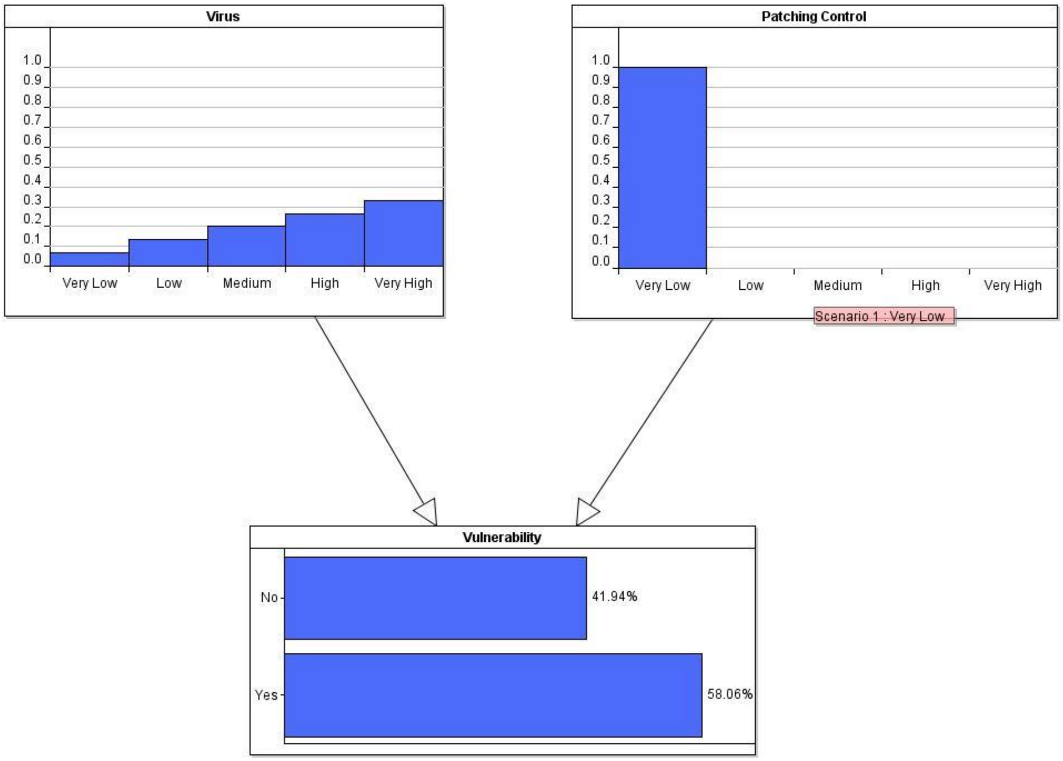

### Discussion and Implications for Application within this Research

This modelling approach was significant in my research for several reasons: 1) The approach incorporates both a security ontology and the use of Bayesian networks based on a security ontology to model threat success probabilities, a key component in the specification of a business risk model based on information security risk; 2) the security ontology was developed based on recognized and standardized security risk practitioner frameworks (Fenz and Ekelhart 2009) and was potentially amenable to alternative control framework mappings with which I am professionally familiar (e.g. ISO/IEC 27002) (Alcazar and Fenz 2012); 3) The developed ontology had been used directly within a commercially available Bayesian network modelling platform to generate a demonstrable threat probability model (Fenz 2012); 4) the Bayesian threat probability calculation had been integrated into an overall system risk model (AURUM) which was demonstrably used to determine overall business risk based on attributed control and business impact costs (Ekelhart, Fenz et al. 2009; Fenz, Ekelhart et al. 2011).

While Bayesian network modelling based on a security ontology provides a superior threat determination approach versus that employed by Tjoa et al, there are several practical considerations in implementation: 1) the example Bayesian network model employed by Fenz et al was freely available for research use, however the network modelling platform employed (Norsys Netica), although free for small scale (i.e. low number of nodes) models, was not freely available for the scale of model used by Fenz and I was reluctant to commit to further software platforms without understanding more about the feasibility of using the described ontology within the indicated platform since there were alternatives present in the literature. I was therefore not able to review the actual Bayesian network described by Fenz; 2) the business risk modelling platform used by Fenz (AURUM) was proprietary and not currently available for either research or commercial use[41] (it now apparently forms the basis for a commercial consulting service operated, in part, by Fenz et al); 3) as with the DES approach, it was not clear at that stage of my review whether or how the threat probability model (a subset of the overall security risk model itself) could be reasonably run in ‘real time’ within a lab setting on consumer grade workstations. Bayesian networks are classed an “NP hard’ computing problems implying potentially long computational times that scale exponentially with the number of conditionally dependent nodes and attributes (Charniak 1991; Kiesling, Strauß et al. 2012; Kiesling, Ekelhart et al. 2014)[42]. Without the ability to test the computation time of models of the size implied by Fenz et al (or even minimally valid from a practitioner perspective), I was reluctant to commit to the Bayesian network approach without some readily determined baseline for both computation time and resulting ease of lab handling.

For these reasons, I proceeded to evaluate analogous threat modelling approaches to ensure that I was pursuing an approach that would support my overall modelling goal without unnecessarily committing to a proprietary software modelling platform. Having reviewed the Bayesian network security modelling literature so far, it became apparent that different academic researchers were pursuing similar approaches (i.e. use of Bayesian networks with security ontologies) but were either not familiar with or at least not citing each other’s work. At this stage it became clear that, at least for the purpose of validating this particular approach, a review of a comparative research effort was required.

[41] Personal correspondence with the authors

[42] NP-hardness (non-deterministic polynomial-time hard), in computational complexity theory, is a class of problems that are, informally, "at least as hard as the hardest problems in NP" https://en.wikipedia.org/wiki/NP-hardness . In computational complexity theory, NP is one of the most fundamental complexity classes. The abbreviation NP refers to "nondeterministic polynomial time." Intuitively, NP is the set of all decision problems for which the instances where the answer is "yes" have efficiently verifiable proofs of the fact that the answer is indeed "yes". More precisely, these proofs have to be verifiable in polynomial time by a deterministic Turing machine. In an equivalent formal definition, NP is the set of decision problems where the "yes"-instances can be accepted in polynomial time by a non-deterministic Turing machine. The equivalence of the two definitions follows from the fact that an algorithm on such a non-deterministic machine consists of two phases, the first of which consists of a guess about the solution, which is generated in a non-deterministic way, while the second consists of a deterministic algorithm that verifies or rejects the guess as a valid solution to the problem. Because of the many important problems in this class, there have been extensive efforts to find polynomial-time algorithms for problems in NP. However, there remain a large number of problems in NP that defy such attempts, seeming to require super-polynomial time. Whether these problems are not decidable in polynomial time is one of the greatest open questions in computer science. However, in practical uses, instead of spending computational resources looking for an optimal solution, a good enough (but potentially suboptimal) solution may often be found in polynomial time. Also, the real life applications of some problems are easier than their theoretical equivalents. https://en.wikipedia.org/wiki/NP_(complexity) . An algorithm is said to be of polynomial time if its running time is upper bounded by a polynomial expression in the size of the input for the algorithm… Problems for which a deterministic polynomial time algorithm exists belong to the complexity class P, which is central in the field of computational complexity theory. Cobham's thesis states that polynomial time is a synonym for "tractable", "feasible", "efficient", or "fast".[9] https://en.wikipedia.org/wiki/Time_complexity#Polynomial_time An algorithm is said to take superpolynomial time if T(n) is not bounded above by any polynomial. It is $\omega(n^c)$ time for all constants c, where n is the input parameter, typically the number of bits in the input. For example, an algorithm that runs for $2^n$ steps on an input of size n requires superpolynomial time (more specifically, exponential time). An algorithm that requires superpolynomial time lies outside the complexity class P. Cobham's thesis posits that these algorithms are impractical, and in many cases they are. https://en.wikipedia.org/wiki/Time_complexity#Superpolynomial_time

**4 – Enterprise Architecture Modelling Approach - KTH Group (Sweden)**

An alternative approach to modelling threat risk determination as incorporated by Sommestad et al (Sommestad, Ekstedt et al. 2008; Sommestad, Ekstedt et al. 2013) involves an information security modelling platform called "CySeMoL" (Cyber Security Model)[43] based on the construction of a 'probabilistic relational model' (PRM) (Friedman, Getoor et al. 1999; Liu and Man 2005; Löf, Stomberg et al. 2010; Buschle, Ullberg et al. 2011; Sommestad 2012; Buschle 2014) which generates an associated Bayesian network from which modelled risk attributes can be inferred. According to Sommestad,

> "A PRM specifies how a Bayesian network should be constructed from a *reference object model* which has been instantiated from a class diagram representing security relationships. Classes have attributes and security relational references. The 'child' attributes in the PRM are associated with a set of 'parents'. The parents of a given attribute are other attributes in the object model that it depends on. Each attribute is associated with a *conditional probability table* that defines the attribute's value given all possible combinations of states in the attribute's parents. The probabilistic model enables the value of attributes in an instantiated object model to be inferred. Such inference can also infer values for attributes with no assigned state. In essence, a PRM defines how a Bayesian network shall be generated using the attributes of an object model. Thus, a PRM constitutes a formal machinery for calculating the probabilities of object properties in various architecture instantiations. For example, a PRM could be used to assess the availability of systems given that certain administrators are assigned to the systems." (Sommestad, Ekstedt et al. 2013).

Motivations for their approach involve the concept of 'enterprise architecture' modelling and its various applications to information technology decision making for IT governance and business alignment (Prado 2009), system quality analysis (Närman, Johnson et al. 2007), system maintainability (Lagerström and Johnson 2008), enterprise software architecture design and integration (Johnson 2002), critical infrastructure design and risk management (Ekstedt 2004), system availability (Närman, Franke et al. 2012; Franke, Johnson et al. 2014), architectural fault tree analysis using Bayesian networks (Franke, Flores et al. 2009) security control selection (Johansson 2005; Johnson, Ekstedt et al. 2007), and the use of Bayesian networks for security defense graphs (Sommestad, Ekstedt et al. 2008; Sommestad, Ekstedt et al. 2009). The authors have focused much of their research to date on critical infrastructure risk modelling for the Swedish electricity generation industry (Sommestad, Ekstedt et al. 2010; Sommestad, Ekstedt et al. 2010; Sommestad, Holm et al. 2011) although the evolved modelling platform is conceptually applicable to any enterprise security risk model.

From my perspective this approach is similar to that of the SBA group (Holm 2014) in that a security 'reference model' consisting of security object 'classes' and relationships between objects (essentially , an ontology, as described) is specified from which a Bayesian risk network can be constructed to represent the probabilistic relationships between the objects. In Sommestad's approach, the construction of an instantiated business process at risk is facilitated by an enterprise architecture modelling platform called EAAT (Johnson, Johansson et al. 2007; Franke, Hook et al. 2009) which incorporates both the system's

[43] https://www.kth.se/en/ees/omskolan/organisation/avdelningar/ics/research/cc/cysemol/description-1.432380

physical component relationships and the conditional probabilities of attributes (dependent on logical relationships between component attributes) and can be used to directly estimate probabilistic ‘states’ of the individual attributes and the system overall. While similar to both the HP and SBA approach in terms of utilizing an underlying security ontology concept, the integrated modelling platform in this case is freely available from the authors as a research platform and an instantiated information system can be directly modelled using predefined assets, threat and control classes using a relatively straightforward graphical user interface. Once again, in correspondence with the authors, I was able to obtain both the software platform and, importantly, documented instantiated model samples from several of their papers and I was subsequently successful in manipulating the model and in running simulations to observe the behaviour of the model inputs and outputs. In addition, the conditional probability tables for the individual system components are described in full detail via the reference documentation also supplied by the authors.

The platform’s architecture template or ‘metamodel’ defines the classes of the possible model assets together with class ‘attributes’ and ‘reference slots’ to other classes and attribute parent references. The classes in the CySeMol metamodel consist of Asset, Owner, Threat, ThreatAgent, AttackStep, and five types of Countermeasure: ContingencyCountermeasure, PreventiveCountermeasure, DetectiveCountermeasure, ReactiveCountermeasure and AccountabilityCountermeasure. Sommestad also describes how the model can be extended to incorporate utility or other decision attributes of interest such as Availability or Business Loss (Hsu and Joehanes 2004). In the following metamodel example, the ‘Availability’ attribute of class ‘System’ is dependent on both the System’s own ‘Reliability’ attribute and the ‘Competence’ attribute of the class ‘System Administrator’. The addition of the conditional probability tables for each attribute completes the definition of a PRM which can now be calculated within a Bayesian network to determine the conditional probability states of the system attributes:

**Figure 43 - Example CySeMol metamodel asset dependencies**

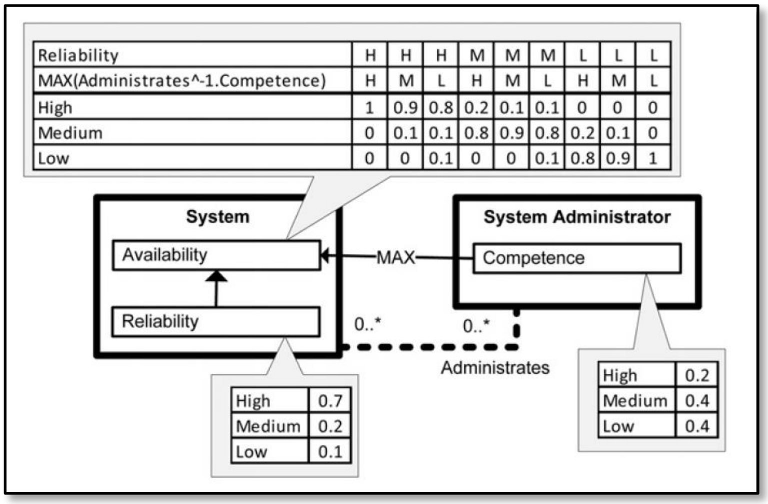

| Reliability | H | H | H | M | M | M | L | L | L |
|---|---|---|---|---|---|---|---|---|---|
| MAX(Administrates^-1.Competence) | H | M | L | H | M | L | H | M | L |
| High | 1 | 0.9 | 0.8 | 0.2 | 0.1 | 0.1 | 0 | 0 | 0 |
| Medium | 0 | 0.1 | 0.1 | 0.8 | 0.9 | 0.8 | 0.2 | 0.1 | 0 |
| Low | 0 | 0 | 0.1 | 0 | 0 | 0.1 | 0.8 | 0.9 | 1 |

| | |
|---|---|
| High | 0.7 |
| Medium | 0.2 |
| Low | 0.1 |

| | |
|---|---|
| High | 0.2 |
| Medium | 0.4 |
| Low | 0.4 |

(Sommestad, Ekstedt et al. 2010)

Once the metamodel is sufficiently complete from an ontological perspective, a specific architecture can then be concretely 'instantiated' by specifying desired objects in a specific system using both the predefined conditional probabilities of the metamodel and, if available, evidence for the attributes of concrete (physical) system assets present in the specific business situation under study. The high level metamodel for CySeMol is depicted below:

**Figure 44 - CySeMol abstract-PRM metamodel for security risk**

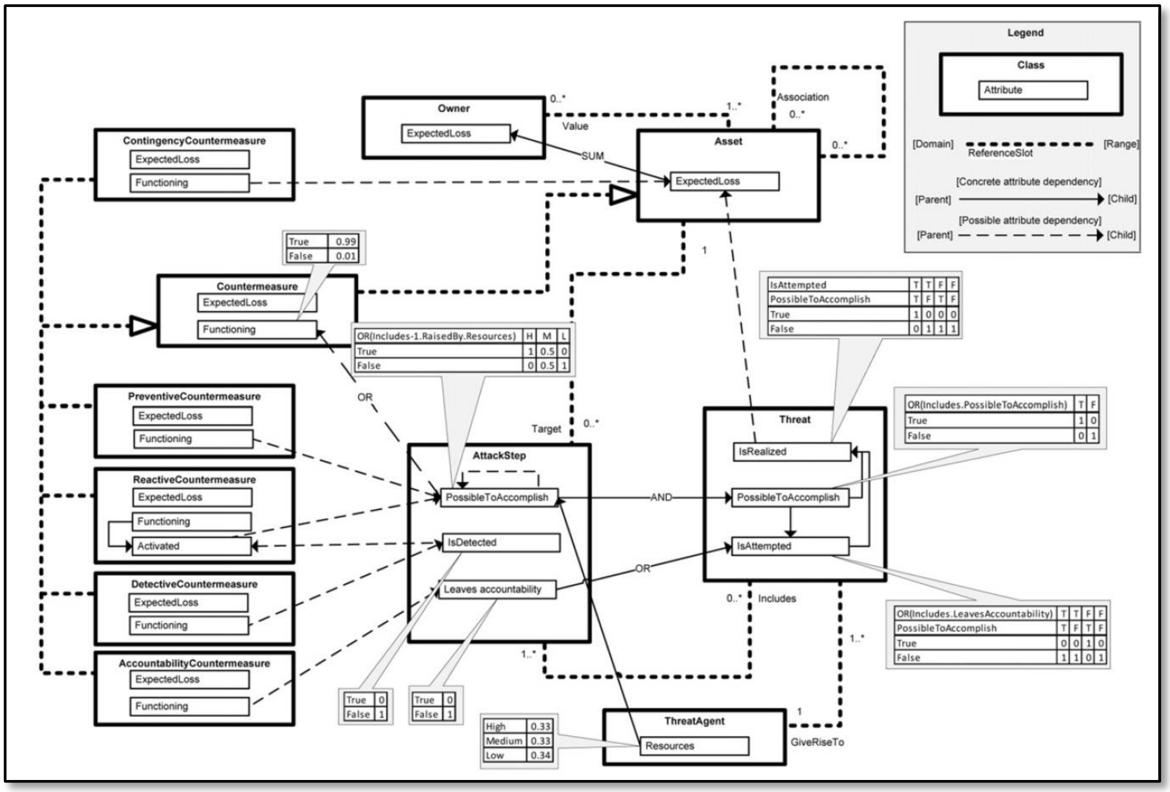

**(Sommestad, Ekstedt et al. 2010)**

The complete CySeMol metamodel comprises 22 classes, 102 attributes, and 32 class relationships (reference slots) and is depicted below:

## Figure 45 - P2CySeMol Class Model*

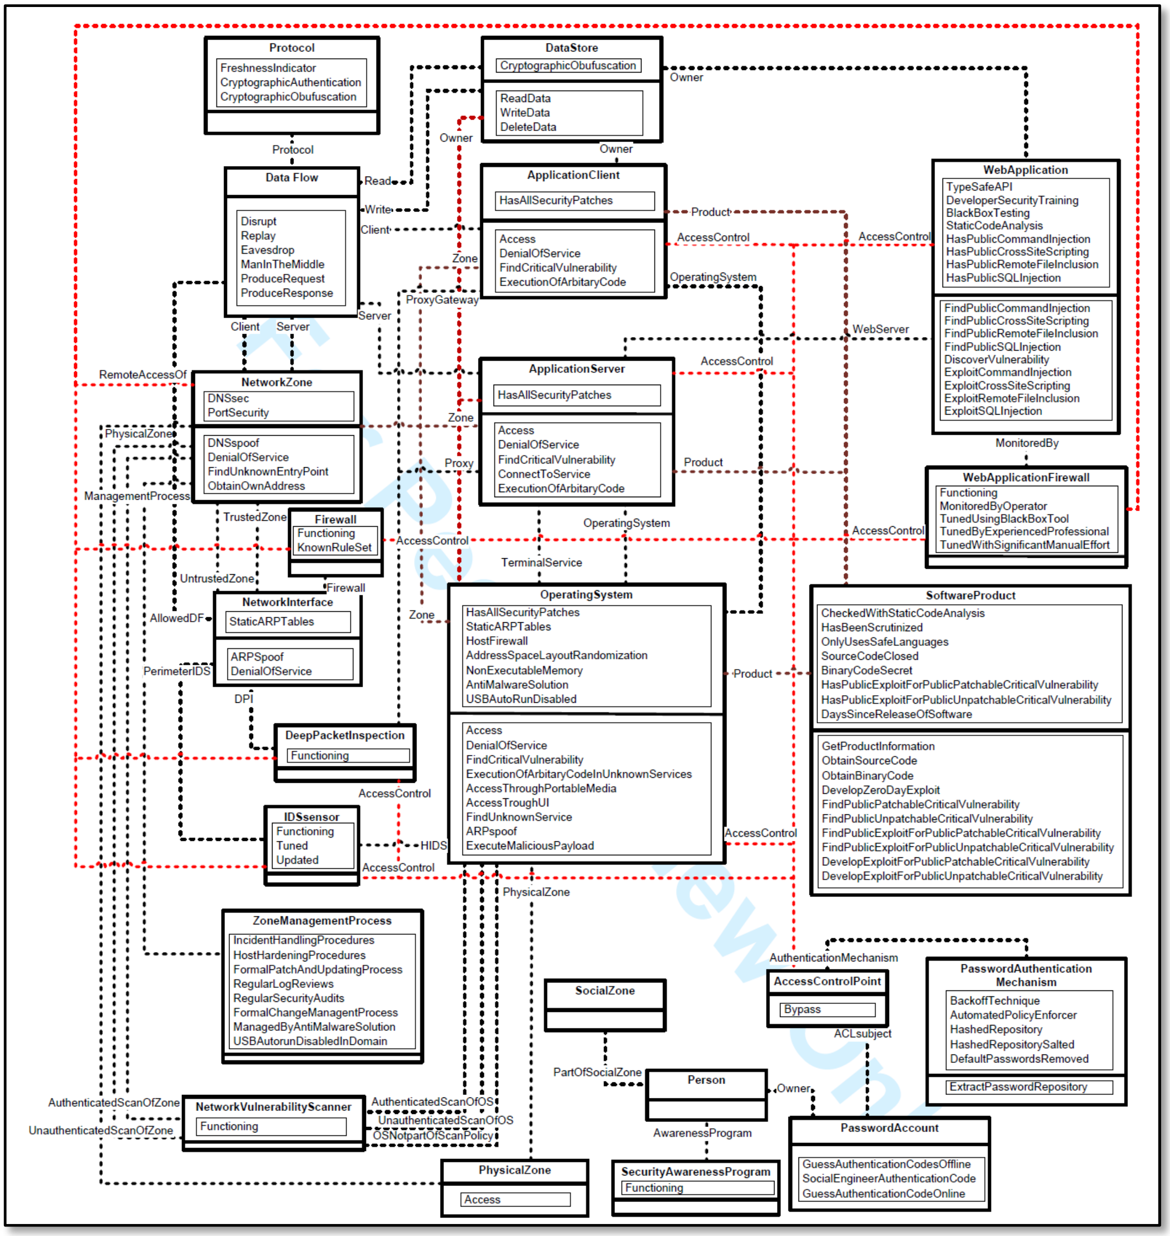

*The upper box of an asset contains the defenses associated with it. The lower box contains the attack steps associated with the asset. Colors are used only to make asset relations more clear.

(Holm, Shahzad et al. 2014)

In this metamodel, for example, if an attacker attempts to log on to a *SoftwareInstance*, the attacker may be required to bypass an *AccessControl- Point*, i.e., if the *SoftwareInstance* has a relationship to an *AccessControlPoint* such as a desktop computer. An *AccessControlPoint* is associated with an *AuthenticationMechanism* and *Account*s that belong to *Person*s who are authorized to access the system and who may have taken security awareness training through a functioning *SecurityAwarenessProgram*. An *Account'*s password may be compromised by being guessed online, offline, or through social engineering (Sommestad, Ekstedt et al. 2013). The model user is required to specify real-world appropriate concrete security object concepts such as network zones, data flows, and software installations and assign values to attributes that determine whether countermeasures such as DNSSEC and non-executable memory are functioning.

For my purposes, CySeMol essentially represents a model for determining the conditional probability of a successful threat on individual systems asset attributed. Threats impact assets through the class 'AttackStep' which is determined at each asset node of the system, the success of which is dependent on preventative/detective and counter/recovery controls. The abstract class 'Countermeasure' allows for the specification and detailed quantification regarding how countermeasures depend on each other and how they influence risk. Five subclasses of Countermeasure are defined in CySeMol which permit me to model a system at a level of granularity reasonably recognizable by practitioners: PreventiveCountermeasure (e.g. firewalls), DetectiveCountermeasure (e.g. an intrusion detection system), ReactiveCountermeasure (e.g. incident handling by administrators), ContingencyCountermeasure (e.g. system or data backups) and Accountability- Countermeasure (e.g. system logging). Indeed, one of the key attractions of the Sommestad model for my purposes is its similar use of the Preventive/Detective and Counter/Reactive control components as utilized by Tjoa et al described above, where this represents a reasonable practitioner grouping of real world control components. The CySeMol platform also ensures that any concretely specified system model describing only the in-scope business system assets and their functional relationships will always produce a valid Bayesian network: from a modelling perspective, the system designer is only required to use the metamodel classes provided when specifying concrete assets which ensures that only valid connections between instantiated objects are created and that the resulting conditional probabilities are therefore always calculable (Sommestad, Ekstedt et al. 2010). Although loss attributes are specifiable as scalar values in CySeMol, my intention was to specify these outside of the model based on further business process risk dependencies. Consideration for how this linking would be accomplished within a system simulator appropriate for use within a lab setting that potentially incorporated both CySeMol and another yet unspecified model which would take the output of CySeMol as its input is described in the next section of this paper.

### Using the CySeMol Modelling Platform in this Research

The software platform for CySeMol and its 2$^{nd}$ generation version, P2CySeMol ("Predictive, Probabilistic Cyber Security Model"), have been described in detail by Sommestad et al. (Johnson, Ullberg et al. 2013; Holm 2014; Holm, Shahzad et al. 2015). P2CySeMol also incorporates a number of improvements supporting the practical use of the platform, including: calculation of all possible attack paths (previously these needed to be specified); eliminating the need to specify attack path depth and designated targets of attacks; and greatly improved computation speeds. P2CyseMol also adds an additional number of practitioner recognizable ‘real world’ logical attacks and defenses specifically regarding web applications, arbitrary code exploits, network vulnerability scanning, and signature-based network intrusion detection. The authors have also carefully documented their extensive validation of the platform at both the system level and at the class component level to verify the conditional probability tables over the course of the model development (Holm, Ekstedt et al. 2012; Holm, Ekstedt et al. 2013; Holm, Shahzad et al. 2015). As indicated earlier, the potential computational speed of a Bayesian network has been a consistent concern in the literature and is also a concern in this research given that the absolute size of my model was not yet determined and the prospect of wanting to use any resulting model within an interactive lab simulation environment. Holm et al indicate the significant speed increases of P2CySeMol over CySeMol although, depending on the model size, prospective calculation times were still a key consideration in my employment of this platform (Holm, Shahzad et al. 2015).

As noted above, the P2CySeMol object modelling platform is freely available and the authors make available the validated instantiated object model described above through the KTH website[44] (Holm, Shahzad et al. 2014). The interface for P2CySeMol permits drag and drop modelling which allows interactive design of a system model based on user knowledge which may be limited to the logical system architectural specifications. The following screenshot indicates the basic workspace with my instantiated model, derived from their validated example model:

[44] https://www.kth.se/en/ees/omskolan/organisation/avdelningar/ics/research/cc/cysemol/downloads-1.432383 (Checked December 2015)

**Figure 46 - Instantiated CySeMol model with end user access nodes highlighted**

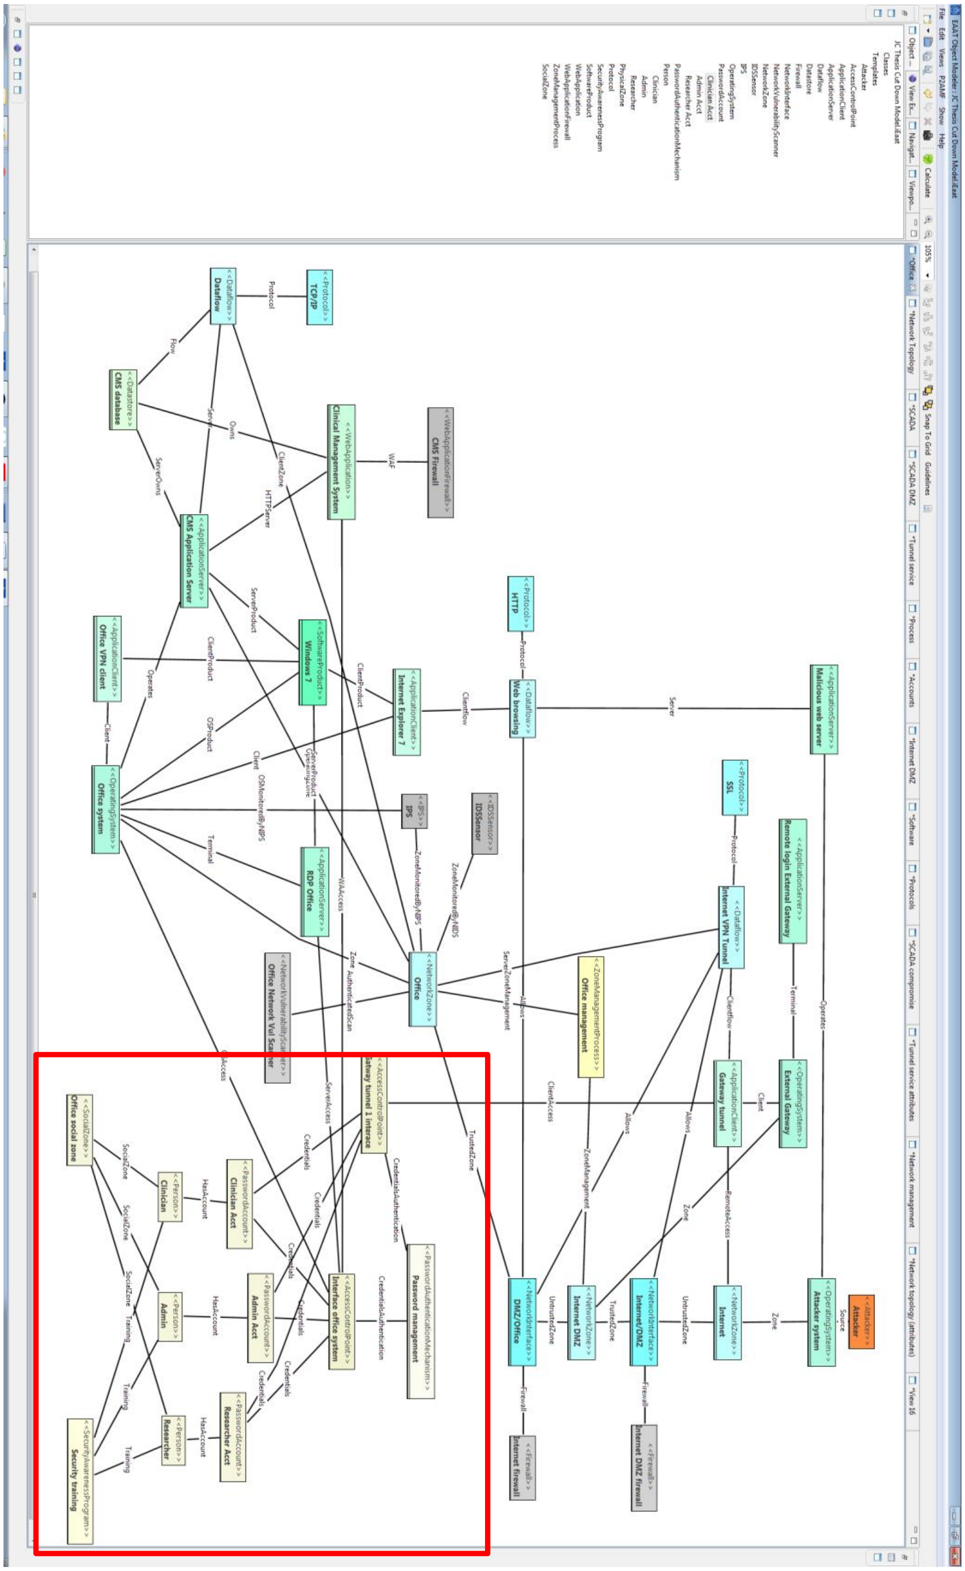

On the left of the screenshot is a list of templated Class objects which allows the user to drag and drop system objects into the open workspace on the right. An object can only be logically connected to another object based on the child object’s attribute parent references. For example, in the above model in the lower right corner, there are three human user groups (“Clinicians”, “Researchers” and “Administrators”) represented by a <<Person>> object, each of whom belong to a <<SocialZone>> (which is susceptible to Social Engineering attacks), each of whom receive “Security Training” via a <<SecurityAwarenessProgram>> , have “Accounts” with individual passwords (<<PasswordAccount>>) that allow them to log onto workstations (“Interface Office System”) requiring password authentication (<<AccessControlPoint>>) which is connected to a web-based “Clinical Management System” (<<WebApplication>>):

**Figure 47 - Instantiated CySeMol model: end user access node detail**

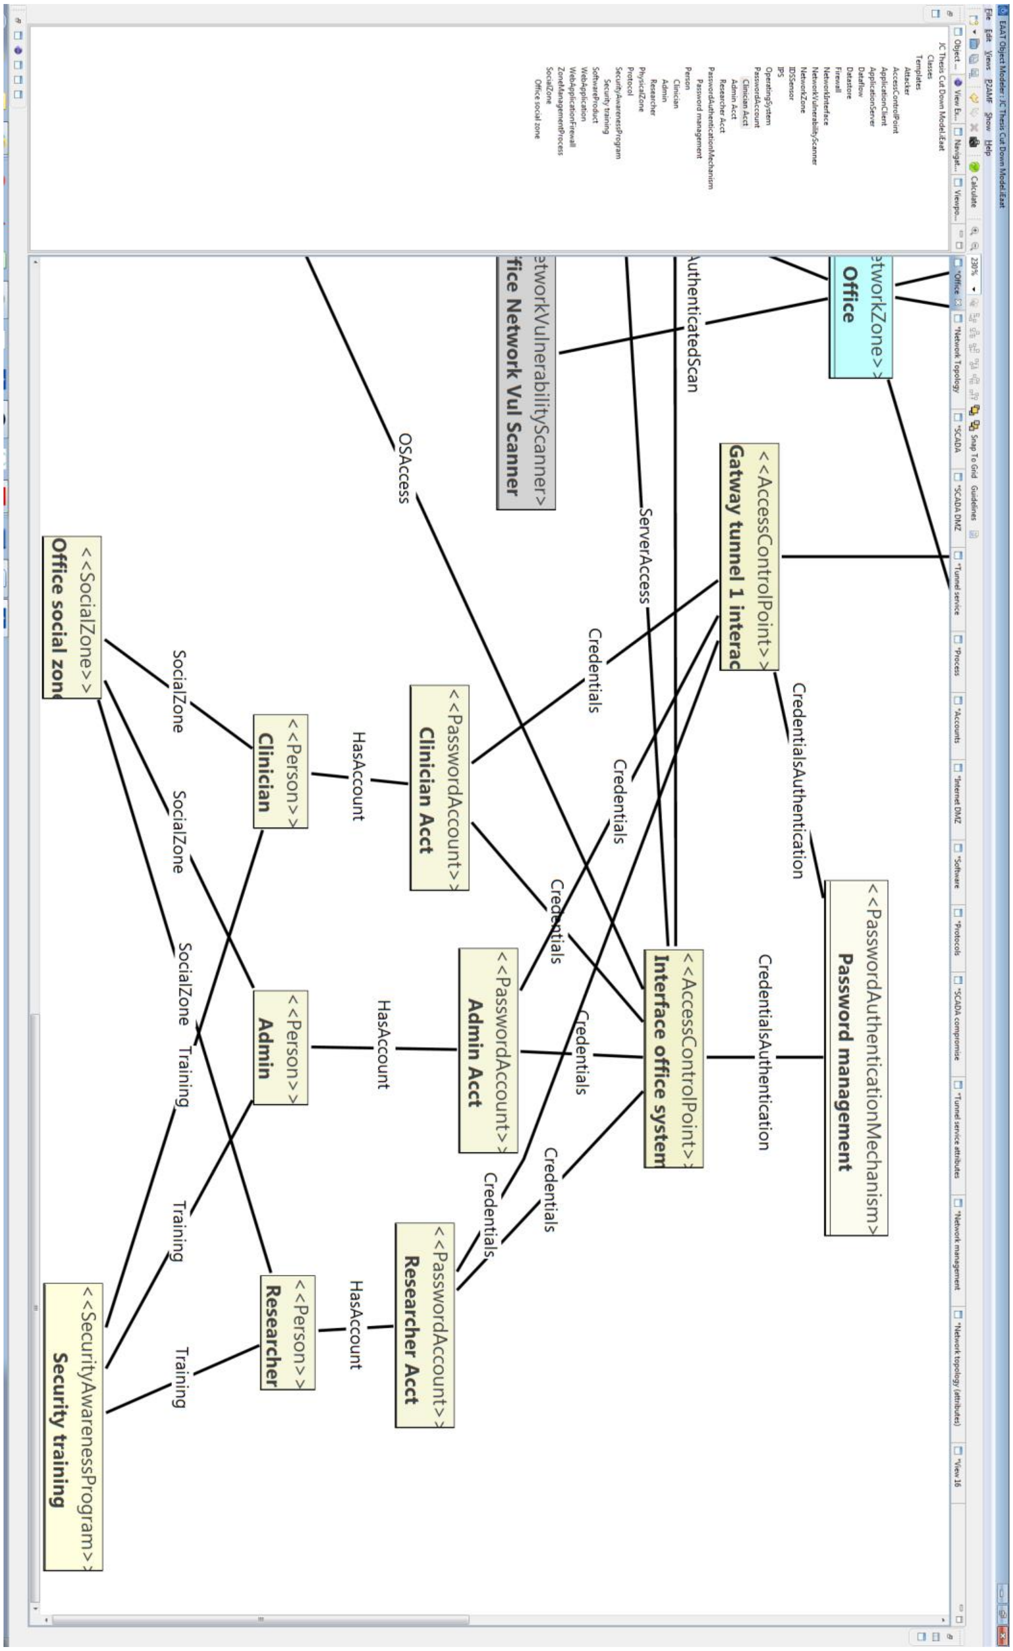

Connections between the model nodes are one-way depending on the relationship (e.g. <<Person>> has a <<PasswordAccount>> account, not the other way around) and from a security perspective may be "Trusted" or "Untrusted" (e.g. between Network Zones) which implies attacker access to a connected Trusted node. Each object has a set of specified Attacks and Controls whose state can be viewed by expanding the Box label to show this detail:

## Figure 48 - Instantiated CySeMol model: node attacks and countermeasures

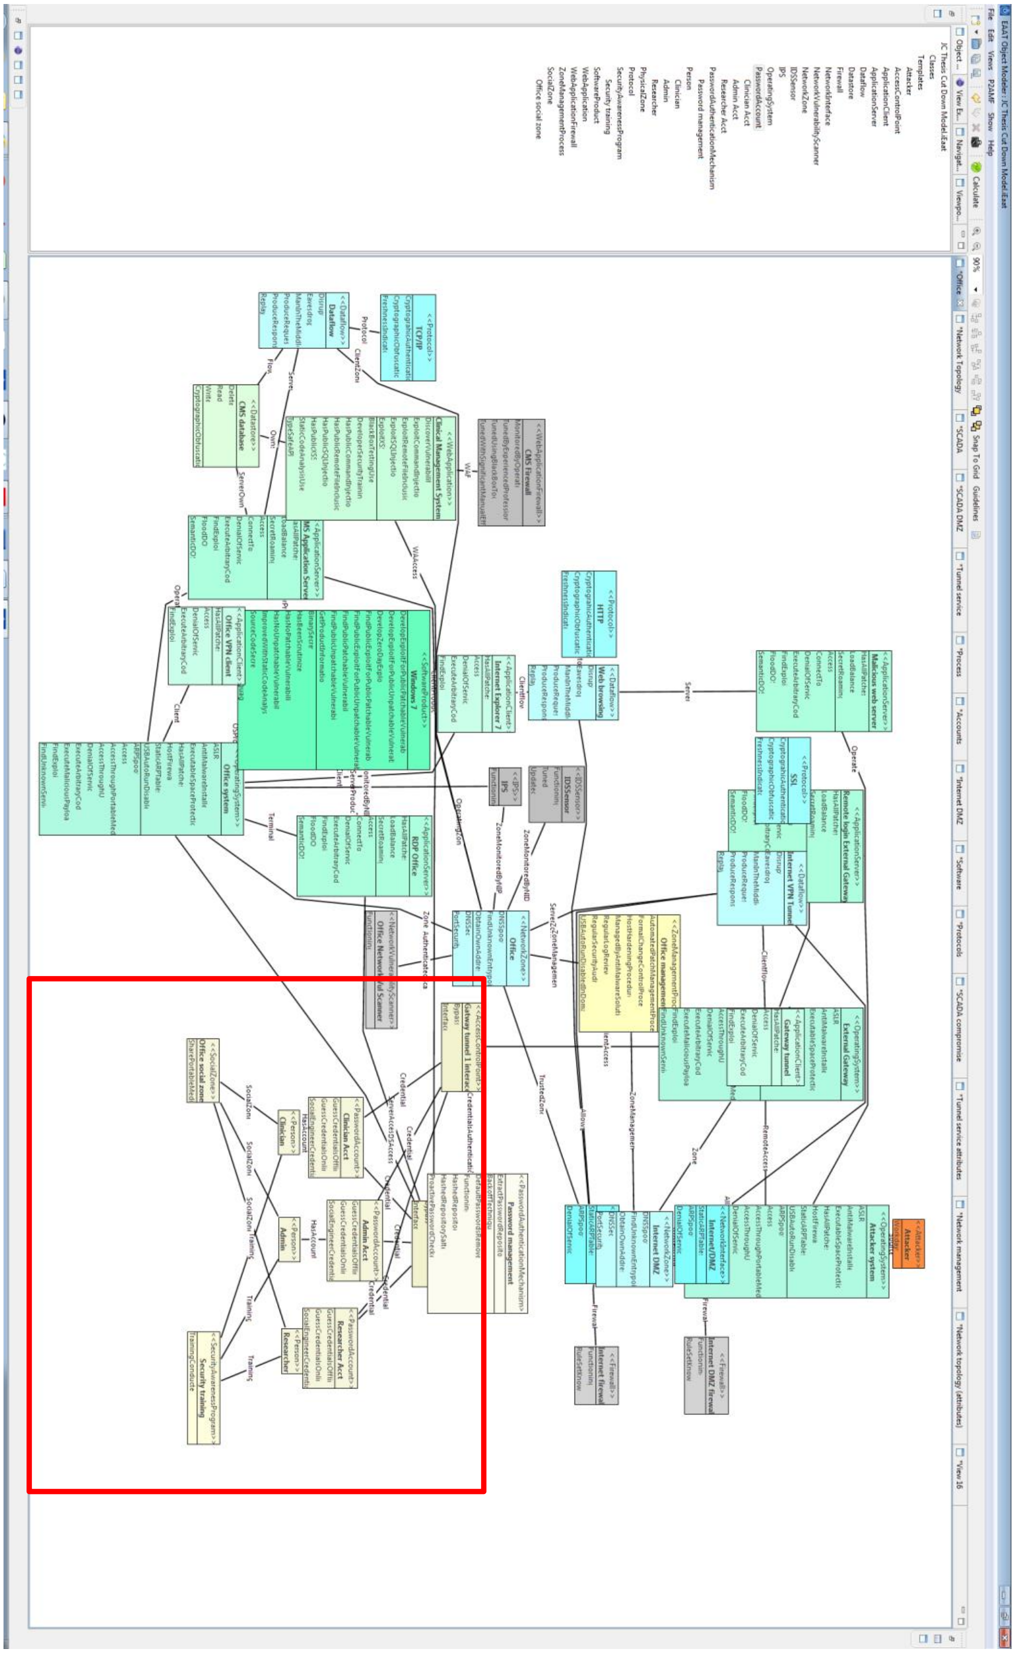

Figure 48 - Instantiated CySeMol model: node attacks and countermeasures

Once again looking at the same overall model but with Attack and Control attribute details shown, we see that <PasswordAccount>> has three potential Attack types: ‘Guessing Credentials Offline’, ‘Guessing Credentials Online’ and ‘Social Engineering Credentials’. The Office <<SocialZone>> is susceptible to ‘Sharing Portable Media’, and ‘Security Training’ may or may not be conducted to prevent ‘Social Engineering’ attacks. The <<PasswordAuthenticationMechanism>> has one attack (“ExtractPasswordRepository”) and 6 corresponding controls (“BackoffTechnique”, “DefaultPasswordRemoved”, “Functioning”, “HashedRepository”, “HashedRepositorySalted” and “ProactivePasswordChecker”:

**Figure 49 - Instantiated CySeMol model: end user access nodes detail**

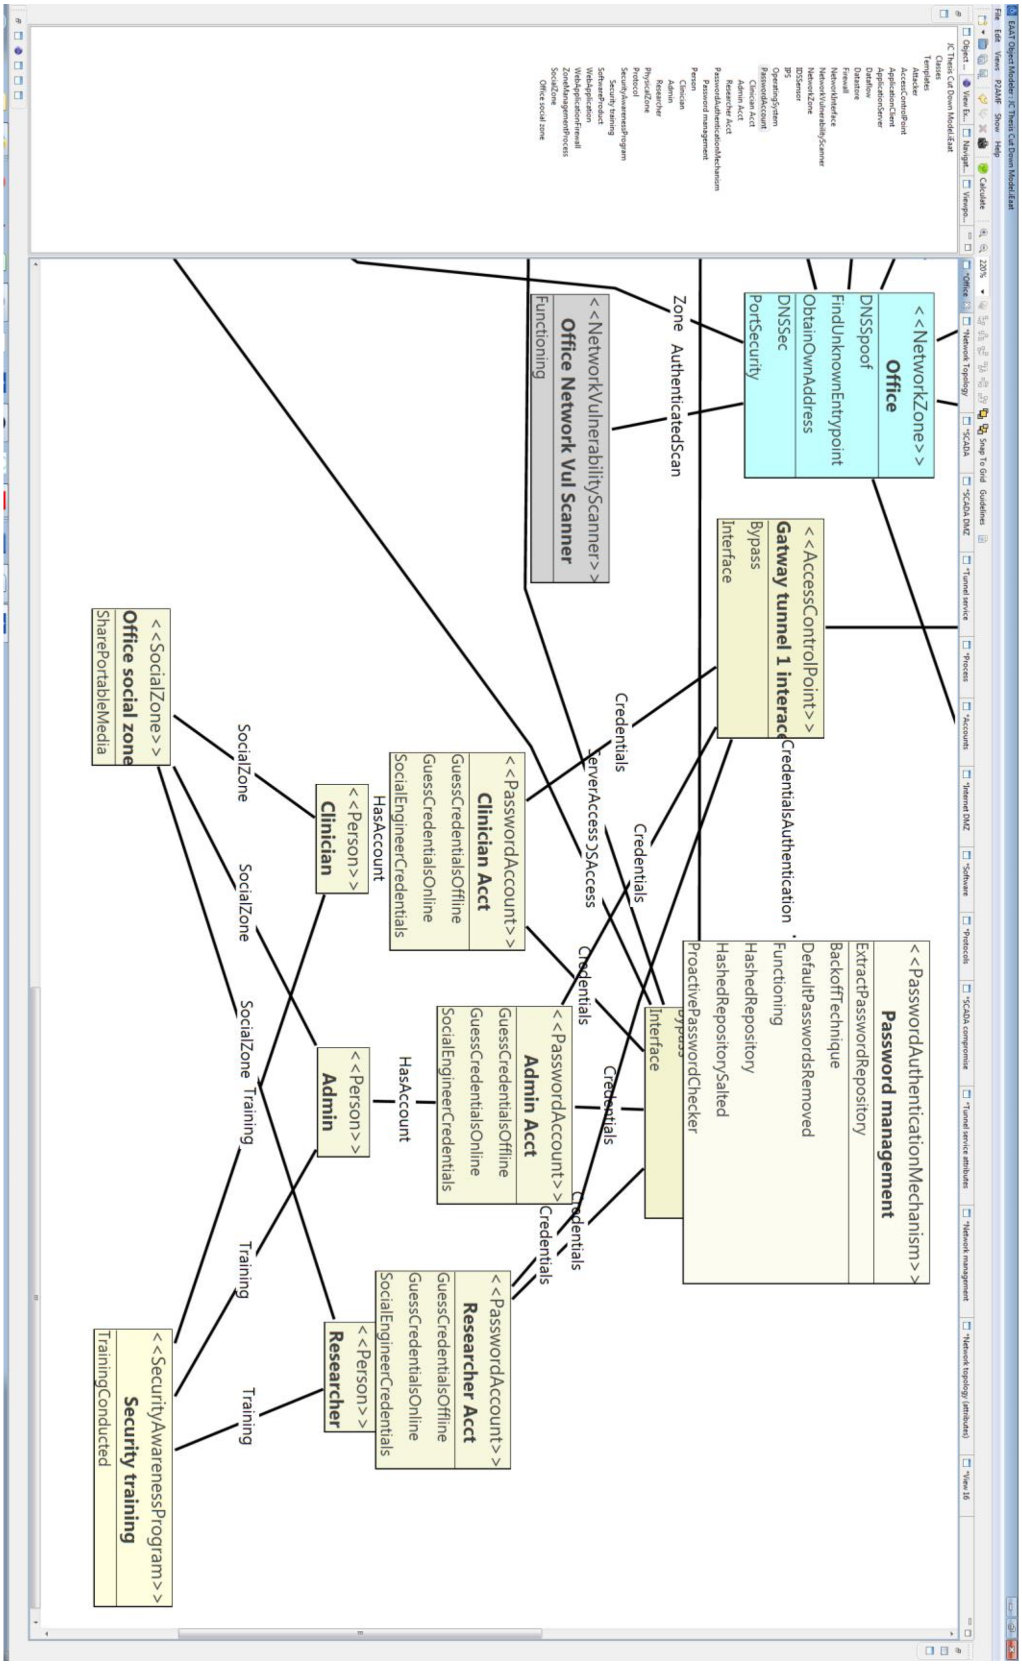

Each of the Objects potential conditional relationships, Attacks and Controls are described in the CySeMol manual, and include descriptions of the associated attribute conditional probability tables for successful attack and references for expert determination of the attack/control dependencies and probabilities of compromise. For example, the Object <<WebApplication>> has three potential connections: <<WebApplicationFirewall>>, <<Datastore>>, and <<ApplicationServer>>:

**Figure 50 - CySeMol Web application class dependencies**

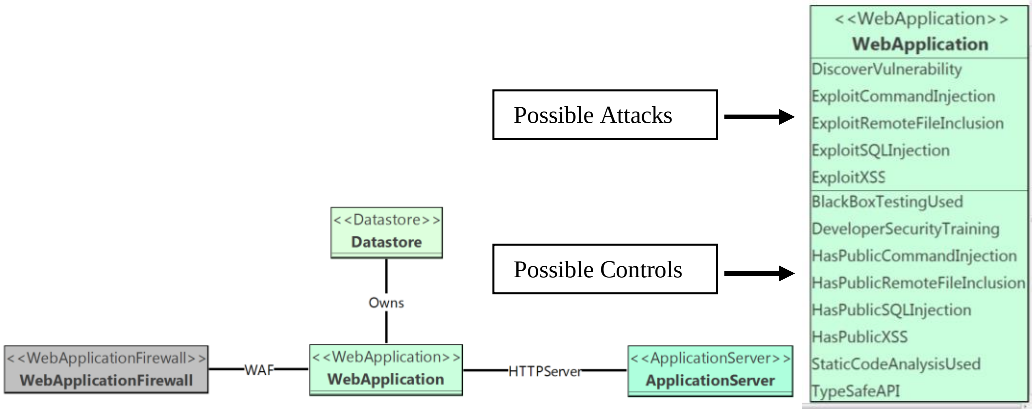

(Holm, Ekstedt et al. 2013)

In CySeMol, the” WebApplication” must have an “ApplicationServer” connected to it in order to run the Application, but may have either or both of a “Datastore” and a “WebApplicationFirewall”, each of which have their own attack vectors and defense controls. Each attack vector for the WebApplication object is then described as being conditionally successful based on the ontologically specified and logically related attack and control attributes both within and external to WebApplication. For example, the attack “Exploit Command Injection” in the above has the following conditional dependencies:

**Figure 51 - CySeMol Web Application class 'Exploit Command Injection' vulnerability**

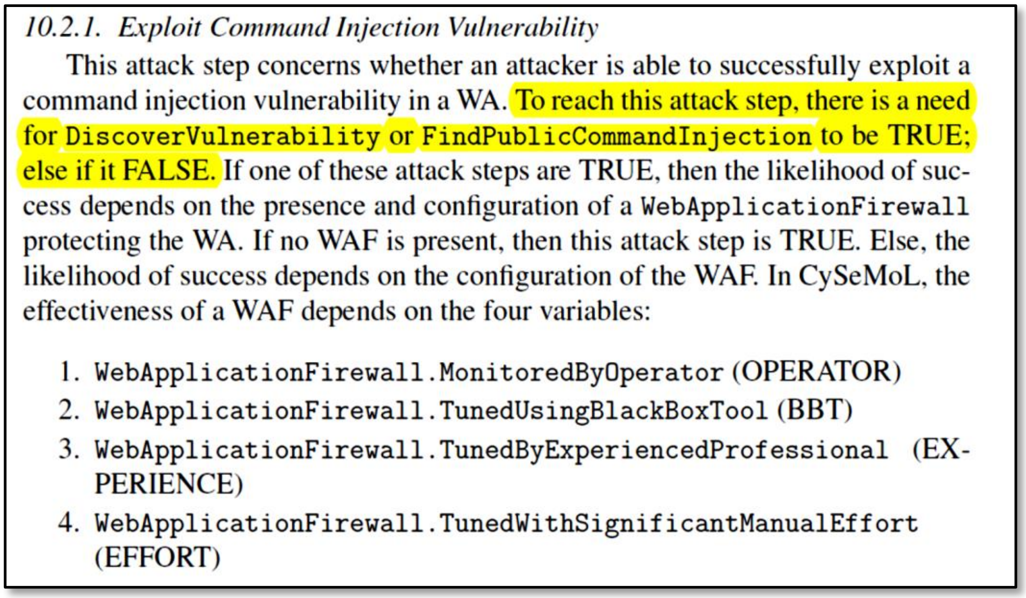

*10.2.1. Exploit Command Injection Vulnerability*

This attack step concerns whether an attacker is able to successfully exploit a command injection vulnerability in a WA. To reach this attack step, there is a need for `DiscoverVulnerability` or `FindPublicCommandInjection` to be TRUE; else if it FALSE. If one of these attack steps are TRUE, then the likelihood of success depends on the presence and configuration of a `WebApplicationFirewall` protecting the WA. If no WAF is present, then this attack step is TRUE. Else, the likelihood of success depends on the configuration of the WAF. In CySeMoL, the effectiveness of a WAF depends on the four variables:

1. `WebApplicationFirewall.MonitoredByOperator` (OPERATOR)
2. `WebApplicationFirewall.TunedUsingBlackBoxTool` (BBT)
3. `WebApplicationFirewall.TunedByExperiencedProfessional` (EXPERIENCE)
4. `WebApplicationFirewall.TunedWithSignificantManualEffort` (EFFORT)

(Holm, Ekstedt et al. 2013)

In this case, if the attacker is able to either manually discover a novel command injection, cross-site scripting (XSS) vulnerability, remote file inclusion (RFI) vulnerability, or SQL injection (SQLi) vulnerability, or find a publicly known Command Injection vulnerability for the WebApplication, the attack has a chance of being successful; otherwise it fails (FALSE). If one or more of these vulnerabilities is present and there is no WebApplicationFirewall present, the attack automatically succeeds (TRUE). The presence and configuration of a "Web Application Firewall" attached to the Web Application otherwise further conditionally determines the probability of successful attack and is itself expressed as a conditional probability table according to the indicated controls present for the firewall: the presence of a watchful human Operator ("Operator"); a 'tuned' configuration ("BBT"); configuration tuning done by a professional technician ("Experience"), and configuration tuning beyond default settings ("Effort"). The specific combinations of the presence or absence of these controls is then associated with the probability of the attack being successful which is evaluated as an exponential distribution based on the amount of Attacker Time:

**Figure 52 - CySeMol Web Application class ‘Exploit Command Injection’ vulnerability: conditional probability of successful attack in the presence of a Web Application Firewall**

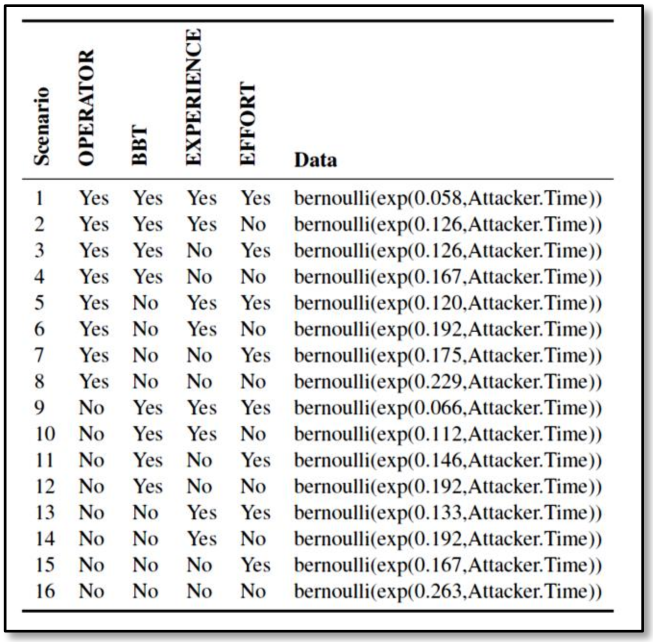

| Scenario | OPERATOR | BBT | EXPERIENCE | EFFORT | Data |
|---|---|---|---|---|---|
| 1 | Yes | Yes | Yes | Yes | bernoulli(exp(0.058,Attacker.Time)) |
| 2 | Yes | Yes | Yes | No | bernoulli(exp(0.126,Attacker.Time)) |
| 3 | Yes | Yes | No | Yes | bernoulli(exp(0.126,Attacker.Time)) |
| 4 | Yes | Yes | No | No | bernoulli(exp(0.167,Attacker.Time)) |
| 5 | Yes | No | Yes | Yes | bernoulli(exp(0.120,Attacker.Time)) |
| 6 | Yes | No | Yes | No | bernoulli(exp(0.192,Attacker.Time)) |
| 7 | Yes | No | No | Yes | bernoulli(exp(0.175,Attacker.Time)) |
| 8 | Yes | No | No | No | bernoulli(exp(0.229,Attacker.Time)) |
| 9 | No | Yes | Yes | Yes | bernoulli(exp(0.066,Attacker.Time)) |
| 10 | No | Yes | Yes | No | bernoulli(exp(0.112,Attacker.Time)) |
| 11 | No | Yes | No | Yes | bernoulli(exp(0.146,Attacker.Time)) |
| 12 | No | Yes | No | No | bernoulli(exp(0.192,Attacker.Time)) |
| 13 | No | No | Yes | Yes | bernoulli(exp(0.133,Attacker.Time)) |
| 14 | No | No | Yes | No | bernoulli(exp(0.192,Attacker.Time)) |
| 15 | No | No | No | Yes | bernoulli(exp(0.167,Attacker.Time)) |
| 16 | No | No | No | No | bernoulli(exp(0.263,Attacker.Time)) |

(Holm, Ekstedt et al. 2013)

The prior probability of the presence of each firewall configuration item (e.g. the presence of a watchful firewall administrator operator) is proposed as a default state (e.g. FALSE) in the documentation based on expert judgement and experience (Holm, Ekstedt et al. 2013). In my application, each modelled control can be turned on or off interactively based on a qualitative selection of the general level of intended control (None, Low, Medium, High, Very High) which is converted to a relative percentage of average effectiveness and then evaluated as a Boolean value (TRUE or FALSE) using a Bernoulli function. For example, a “Low” control profile may not include the presence of a watchful firewall operator at all (0% probability of TRUE = always FALSE), while a “High” or “Very High” control profile (e.g. 80% or 90% effective) may result in this control being effectively present most of the time (TRUE) but with some degree of stochasticity (i.e. the Operator may be present, but not very ‘watchful’). The resulting conditional probability tables are then constructed for each Object attribute based on presence or absence or absence of its parent attribute dependencies. In the example of the “WebApplication” attack “Exploit Command Injection”, we obtain a $2^n$ column by n row CPT, where n equals the number of conditional attributes (so: $2^6$ = 64 columns x 6 rows in this case) which allows us to lookup the indicated conditional probability of a successful attack being “TRUE” or “FALSE” for all possible Boolean combinations of the attribute dependencies. The following table indicates the first 10 columns of that table:

**Figure 53 - CySeMol Web Application class ‘Exploit Command Injection’ vulnerability: conditional probability table**

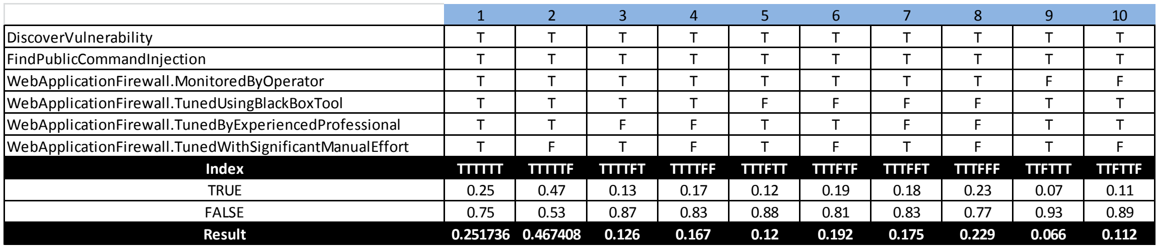

| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|---|---|---|---|---|---|---|---|---|---|---|
| DiscoverVulnerability | T | T | T | T | T | T | T | T | T | T |
| FindPublicCommandInjection | T | T | T | T | T | T | T | T | T | T |
| WebApplicationFirewall.MonitoredByOperator | T | T | T | T | T | T | T | T | F | F |
| WebApplicationFirewall.TunedUsingBlackBoxTool | T | T | T | T | F | F | F | F | T | T |
| WebApplicationFirewall.TunedByExperiencedProfessional | T | T | F | F | T | T | F | F | T | T |
| WebApplicationFirewall.TunedWithSignificantManualEffort | T | F | T | F | T | F | T | F | T | F |
| **Index** | **TTTTTT** | **TTTTTF** | **TTTTFT** | **TTTTFF** | **TTTFTT** | **TTTFTF** | **TTTFFT** | **TTTFFF** | **TTFTTT** | **TTFTTF** |
| TRUE | 0.25 | 0.47 | 0.13 | 0.17 | 0.12 | 0.19 | 0.18 | 0.23 | 0.07 | 0.11 |
| FALSE | 0.75 | 0.53 | 0.87 | 0.83 | 0.88 | 0.81 | 0.83 | 0.77 | 0.93 | 0.89 |
| **Result** | **0.251736** | **0.467408** | **0.126** | **0.167** | **0.12** | **0.192** | **0.175** | **0.229** | **0.066** | **0.112** |

(Holm, Ekstedt et al. 2013)

Since the model objects (WebApplications, PasswordAccounts, etc.) are specifically linked based on their logical security relationships and the objective logical architecture of the modelled system, the conditional probability of any given attribute can be evaluated by determining the associated attribute values in the object’s conditional probability table. For any arbitrarily complex instantiated model, since the associated attribute values themselves may be conditionally complex (that is, any single input attribute may have numerous linked conditional attributes which need to be evaluated first, and where each of these have similar but distinct conditional states based on their inherent object relationships and the architectural complexity of the model), the resulting Bayesian calculations quickly become complex and typically require both computational horsepower and specific algorithmic sampling methods to ensure that a reasonably *approximated* (i.e. not exact) solution can be reached in a reasonable amount of time (Cooper 1990; Dagum and Luby 1993; Welch 1996; Ellis and Wong 2008; Larrañaga, Karshenas et al. 2013)[45]. To address this issue, P2CyseMol has lowered the required computation time greatly from its predecessor CySeMol, but this time remains significant for anything but arbitrarily small models and/or relatively high powered desktop computers:

**Figure 54 - CySeMol vs. P2CySeMol computation times**

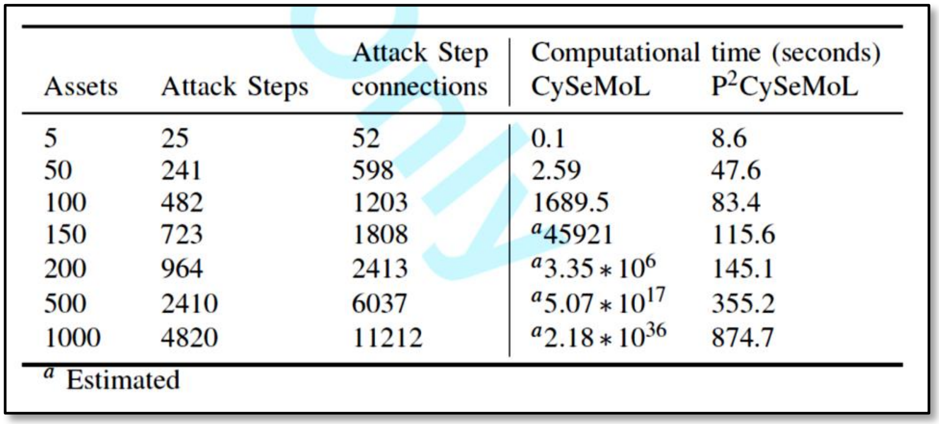

| Assets | Attack Steps | Attack Step connections | Computational time (seconds)<br>CySeMoL | <br>$P^2$CySeMoL |
|---|---|---|---|---|
| 5 | 25 | 52 | 0.1 | 8.6 |
| 50 | 241 | 598 | 2.59 | 47.6 |
| 100 | 482 | 1203 | 1689.5 | 83.4 |
| 150 | 723 | 1808 | $^{a}$45921 | 115.6 |
| 200 | 964 | 2413 | $^{a}3.35 * 10^{6}$ | 145.1 |
| 500 | 2410 | 6037 | $^{a}5.07 * 10^{17}$ | 355.2 |
| 1000 | 4820 | 11212 | $^{a}2.18 * 10^{36}$ | 874.7 |

$^{a}$ Estimated

(Holm, Shahzad et al. 2015)

[45] See note above re: NP-hardness of this problem

In my model with just under 50 objects, the calculation time is just under two minutes. While this is reasonably fast during my iterative development of the overall model, the prospect of having lab participants waiting two minutes or longer (based on the anticipation of using lowered powered laptops) between trial simulations of the model so that they can then select controls - the main point of the choice experiments - remained potentially problematic.

**Discussion and Implications for Application within this Research**

At this point in model development I was reasonably sure that the conditional probability approach was appropriate for the purpose of representing the required logical security model relationships that would be valid from a practitioner perspective and that CySeMol/P2CySeMol was a promising platform for this purpose. In addition to the computation time issue noted above, however, several additional modelling considerations remained that were not directly addressable via CySeMol. Since the objective of the modelling effort was to generate business losses attributable to system CIA states over which participants would make prospective control decisions, the calculation of successful attacks, however valid, was only a precursor to the required risk outputs and substantial model modification would be required beyond simply calculating successful attack probabilities:

1) **Representation and interactive manipulation of stochastic control effectiveness:** A core element of my modelling approach is that the attributes of both the business process model and particularly the logical security model are inherently stochastic, making the conditions involving security control choice subject to both risk (objective and known) and uncertainty (unknown and therefore subjective) regarding the likelihood and probability of most model attributes and outcomes. One of the features of P2CySeMol, and an improvement over CySeMol, is its ability to incorporate uncertainty in the model whereby attributes values are made stochastic using probability distributions instead of fixed percentages and where the existence of nodes and links between nodes may also be uncertain (Johnson, Ullberg et al. 2013). While this was promising within P2CyseMol, its implementation was not obvious from the available documentation and presented a substantial learning curve. More importantly, it was also not clear how variations on control effectiveness (average and variance of effectiveness) would be manipulable within the proposed lab setting where, despite the platform's GUI, the unlikely prospect of lab participants changing individual or groups of attribute probability distributions *directly within P2CySeMol* was a substantial barrier to its use in the lab. This highlighted the need for consideration of the incorporation of the system simulation model within some type of secondary participant end-user interface, similar to that undertaken by Fenz et al using their AURUM tool noted above (Ekelhart, Fenz et al. 2009) although it was not immediately clear how that would be best accomplished. This highlighted the need to begin prototyping the lab experiment interface itself and to consider linkages or the embedding of the simulation within that interface.

2) **Expression of the system risk impacts:** Consistent with all of the logical security models presented above, successful attacks are proposed to affect attribute states (operability, accessibility, etc) and therefore the process performance of system components expressed in terms of standard 'CIA' attributes (Confidentiality, Integrity and Availability) which in turn affect the dependent business processes. As described in the HP and Tjoa models for example, successful attacks are posited to lower the level of system or component availability based on an assumed negative functional relationship between a successful attack and attributes affecting the component's availability attribute. Lower system or component availability then diminishes the ability of the dependent business process. Tjoa proposes that a given business process has a 'degree of completion' in the range (0,1) which is a function of time *t* and where degree of completion is a function of the required computing resource availability and integrity. The amount of decline in the availability of a required computing resource is itself a function of a 'threat impact function' which is mitigated by a factor representing the amount of the threat 'counter' control. (Tjoa, Jakoubi et al. 2011). Conceptually, decreases in component availability 'cause' business processes to slow down per unit of time but the processes nonetheless eventually complete, whereas decreases in component integrity (either data or system computational accuracy) essentially cause process output errors requiring process rework, requiring re-computation once the system recovers sufficient integrity and similarly lowering 'acceptable' output per unit of time. While this model presents a clear causal relationship between attacks, system availability/integrity and business losses, it does not present a granular model of the way that attacks specifically degrade the CIA attributes of affected components or sub-components. CySeMol/P2CySeMol on the other hand, explicitly models the effect of attacks on vulnerabilities at the sub-component level (Holm, Ekstedt et al. 2012; Sommestad, Ekstedt et al. 2013) however it does not directly calculate the resulting CIA impact of the successfully attacked component or of the system overall on a designated business process. From that perspective the Tjoa approach is more useful in terms of generating business losses, but CySeMol presents a much more valid model of a system at risk of attack.

3) **Attribution of business losses:** my proposed lab experiments are based on participants making choices over controls faced with risky or uncertain prospective monetary business losses directly attributed, in the case of attack's affecting component availability, to the inability of system users to undertake transactions which would otherwise have been completed were the system fully available. Tjoa specifies an 'income function' which enables the attribution of business income to the degree of completion of a business process in each time period and where declines in process degree of completion correspondingly diminish income. While the CySeMol documentation conceptually explained how utility or business losses could be incorporated using the EAAT modelling framework generally (Johnson, Iacob et al. 2013; Johnson, Lagerstrom et al. 2013), the examples indicated that substantial additional programming would be required to accomplish this *directly within the tool* which

presented an unknown effort and timing risk to my work. In correspondence with Ekstedt, he indicated that the incorporation of the above elements 1-3, while underway within his own team, “…would constitute a full bingo on his [research] scorecard…”, currently represented by multi-doctoral work at the KTH Institute, and was therefore considered ambitious within my own research agenda. I resolved that this would have to be

4) **Appropriate expression of business losses in the context of decision making over controls:** Simultaneous with my review of the conditional probability and Bayesian network models noted above, it became apparent that I would need to more appropriately represent business losses attributed to security posture in a context which allowed participants to make relatively ‘interactive’ control choices affecting losses. As noted above, the Bayesian network models themselves, while appropriately representing the logical relationship between system components and attack success, do not generate associated business losses directly – these would have to be ‘attributed’ to processes dependent on system availability. More importantly, while the Bayesian network estimation of attack success probability is recognized as a stochastic process which can be (must be) estimated based on Monte Carlo simulation (since an entire network, being NP-hard, is typically not directly estimable but can only be approximated) and produce probability distributions for component attribute states, Bayesian networks are not time series simulation models per se and do not produce time-based or even discrete observations of the network component or states. Simulation in that context is used to produce observations on attribute states, but there is no direct way of producing observations *across all attributes* for individual simulation periods, nor are the observations considered temporal. In essence,  as present above, CySeMol can only tell you what the distribution of attack success is for each node of the network, not what the actual (or simulated) attack success was for a day, or over an entire year:

**Figure 55 – CySeMol Histogram representing the probability of successful node attack**

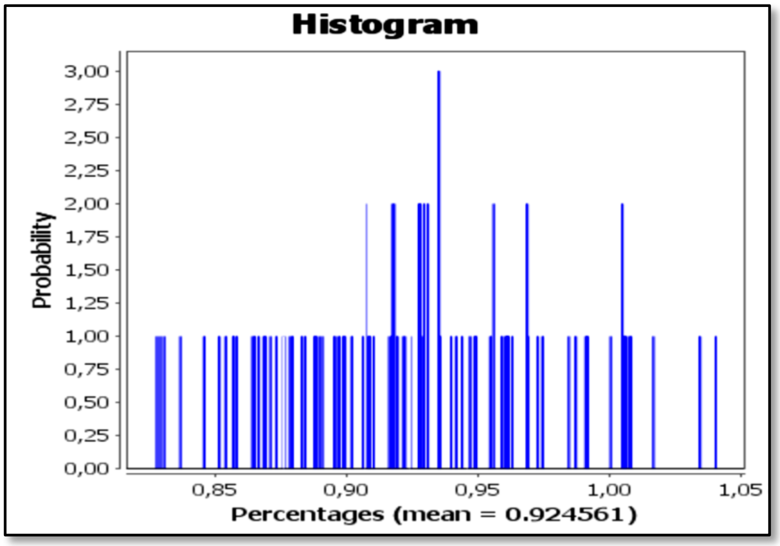

In contrast, the various DES models reviewed are based on temporal input/output business processes and associated control activities operating sequentially and repeatedly over discrete time intervals (minutes, hours, days) where the probability of successful attack is repeatedly simulated at each interval over some chosen time horizon (a year). Importantly, the DES approach can incorporate system 'memory' and can therefore represent inter-period dependencies that more realistically reflect real-world business and security control processes. For example, the Tjoa model explicitly models the current period level of attack strength based on the prior period's attack strength diminished by the level of counter controls present at the start of the current period. This approach permits the decision maker to naturally observe the inter-period (day to day) and overall behaviour of the stochastic system much as a real world operator would if they were running the designated business process.

More importantly, in the context of losses attributed to security CIA postures in which the distribution of attribute states and business losses may not be normal, a decision maker is not (should not be) only concerned with the total or average of a series of attribute states and losses, but the nature of these observations expressed as a *probability density function*. In these cases, we should be particularly concerned with the higher moments of the distribution of losses - variance, skewness and kurtosis – and not simply the mean. This also becomes apparent when we consider emerging alternatives to security 'self-protection' (control investments to lower the probability of loss) and alternatively consider investments in 'self-insurance' to lower the impact of losses *that may actually occur* and which typically exceed some defined upper threshold. Simulating and expressing security related business losses in this context lead to my review of modelling approaches supporting *value-at-risk* (VaR) and *conditional value at risk* (CVaR) methodologies familiar in financial engineering, insurance and information security (Duffie and Pan 1997; Artzner, Delbaen et al. 1999; Rockafellar and Uryasev

2000; Wang, Chaudhury et al. 2008; Herath and Herath 2011). The following section details my review of the literature concerning the nature of coherent risk measures, and the potential for application of VaR and CVaR measures of risk within this research. The subsequent sections then returns to the issue of modelling and simulation of security losses whose probability distributions exhibit characteristics which can be expected to reflect real world results and are appropriate for the control choice experiments within this research.

## 5 - Using VaR, CVaR and Monte Carlo Modelling Approaches for Simulation of Systems at Risk

Alternative models for information security risk simulation have been presented relatively recently in the information security literature using VaR and CVaR concepts to more appropriately express risk measures for stochastic systems (Jaisingh and Rees 2001; Conrad 2005; Ozcelik and Rees 2005; Conrad, Oman et al. 2006; Wang, Chaudhury et al. 2008; Thomas 2009; Gheorghe 2012; Sawik 2013; Thomas, Antkiewicz et al. 2013; WEF 2015). This quantitative approach to risk modelling borrows from both enterprise finance and operations risk management[46], incorporates Monte Carlo simulation for stochastic model inputs and expresses risk in terms of a probability distribution of loss outcomes. This approach, which enables the representation of the potential *distribution* of information security risk outcomes, has appeal for my purposes since it explicitly incorporates uncertainty in both model input and outputs and the model results are demonstrably non-deterministic across time periods making the fact of uncertainty clear to the decision maker. As noted in the above section, representation of historical or simulated losses relevant for stochastic systems which can be expected to exhibit non-normal outcome profiles and for which risk decisions should approached using concepts drawn from *extreme value theory* (Embrechts, Resnick et al. 1999). For this research I undertook significant literature review to understand the potential application of VaR/CVaR concepts to information security risk and associated literature review in the area of 'coherent risk estimators' and decision making under uncertainty for risk management (Artzner, Delbaen et al. 1999; Acerbi, Nordio et al. 2001; Acerbi 2002; Acerbi and Tasche 2002; Klinke and Renn 2002). Importantly, this literature review also prompted the concept of adding 'self-insurance' as an endogenous security control alongside traditional information security 'self-protection' controls (Shogren and Crocker 1991) which I have successfully incorporated into the lab experiments using an analogous experiment from the literature (Bajtelsmit, Coats et al. 2015) but incorporating my simulation of a system at risk.

In terms of simulation and presentation of data, since the system at risk in our context is operating continuously across time, this requires not point estimates but time series or converted probability distributions of losses that capture the nature of the system risk being considered. Capturing these differences in a practitioner security context implies presenting not just the average risk level (i.e. the expected value of loss) but also the distribution (variance) of losses over some designated time period.

The downside of this modelling approach is that the utilized attack/threat and impact models within the existing VaR/CVaR security literature generally lack a level of logical attack/vulnerability model validity from a practitioner perspective and so additional work on that aspect was be required in order to pursue this approach in a way that avoided unnecessary simplicity and contributed to the field. Positively, my selected approach permits relative flexibility and extensibility, ease of implementation and of portability and use within a lab setting though the use of Excel with associated commercial Monte-Carlo plug-in tools which

[46] See for example "A New Approach for Managing Operational Risk - Addressing the Issues Underlying the 2008 Global Financial Crisis, Society of Actuaries 2009 http://www.soa.org/Files/Research/Projects/research-new-approach.pdf

significantly reduced the required simulation programming effort without compromising the potential validity of the information security model being presented.

The following sections briefly review the literature regarding the use of ‘coherent’ measures of risk for stochastic systems which exhibit ‘tail losses’. I then explain my resulting modelling and system simulation approach which borrows from both the Bayesian/conditional probability approaches and recent literature on modelling of security risks using Monte Carlo simulation techniques.

### Coherent Risk Measures for Security Control Choice

The core hypothesis of this research is that information security decision makers may be subjectively biased when choosing controls to reduce risk and that the biases can be tested using experiments involving choices made under controlled presentations of risk information. The latent biases are assumed to manifest when there is uncertainty in the available information requiring subjective interpretation of the risk factors on the part of the decision maker, and that subjective interpretation may be involved even when the information is known or demonstrably certain. In practice, since security risks may contain both known and unknown/uncertain information, security control selection presents a potentially ideal context in which to test for decisional biases. At the same time, the question of what security risk information should be used and how the information should be presented to test for decisional bias arises due to the specific nature of the risk model being presented. As noted above, traditional qualitative and quantitative risk models often present risk as a point estimate measured as the *probability of a loss event* multiplied by the *business impact* should the loss event actually occur. Qualitative measures of risk involve assigning a relative rating scale for the risk factors representing the likelihood of loss and the amount of impact (Low, Medium, High, etc.) and may either represent ranges of quantified factors (i.e. a conversion of known probabilities into qualitative ranges) or management’s relative estimates of the factors without direct quantification. This approach also requires some ordering convention to combine the inputs into the final risk rating (e.g. that Low probability x Low impact = Low Risk, but that Low x Medium = Medium, etc.). Quantitative risk measures are comparatively more robust since risks are automatically ranked on a continuous numerical scale based on quantified probabilities and impacts measured in ‘expected loss’ units (e.g. dollars). Both approaches require that management agree on how to measure impacts that may not necessarily be quantifiable in monetary terms e.g. assigning a value to ‘loss of business reputation’. In both approaches, the measure of risk can be calculated over any time period but in either method are always considered as a point estimate of the average or expected value of loss over that period.

One of the main practitioner criticisms of this approach is that neither method distinguishes effectively between “high probability/low impact” and “low probability/high impact” risks and therefore cannot directly inform appropriate control strategies that differentiate between the frequency and impact of these risks. The distinction is important since the impact of low probability/high impact events *that actually* *occur* can be potentially catastrophic to a business, while the impact of high probability/low impact events may presumably be reasonably absorbed by a business for a short period of time until their incidence can be brought under control. The choice of specific controls that lower the probability of successful attack versus those that lower the impact of successful attacks are inherently different choices from a managerial perspective. Low probability/high impact risks, for example, generally favour some type of *recovery* control (such as cyber insurance) which reduces impact, while high probability/low impact risks favour administrative or technical controls that directly affect the system at risk in order to reduce the probability of occurrence itself. Although one can think of scenarios in which many small impacts may also have catastrophic results taken as a whole, we restrict our view here to a single realized risk event and the associated impact if the risk is experienced once. In qualitative risk models this aversion towards high impact events may be reflected by ranking the low probability/high impact risk ratings relatively higher than the otherwise quantitatively equivalent high probability/low impact risks. From a perspective that generally assigns negative utility to loss, it is clear that low probability/high risks are of a different, if not greater concern, even if taken as an equivalently 'expected' value for the purposes of decision making. The following section explores approaches to quantifying these risks in the context of a stochastic system generating losses in which a decision maker must decide on a portfolio of control inputs that affect the losses non-deterministically.

Quantitative modelling considerations for systems exhibiting low probability/high impact events originated in the context of extreme value theory for financial risk modelling (Embrechts, Resnick et al. 1999; Gilli and Këllezi 2006) and has been more recently extended to operational loss modelling for catastrophic events (Berliner 1985; Kunreuther, Novemsky et al. 2000; Kunreuther and Heal 2003), catastrophe insurance (Kunreuther and Pauly 2004; Kunreuther and Michel-Kerjan 2012; Kunreuther and Pauly 2015) and information security or 'cyber' risk insurance coverage (Biener, Eling et al. 2015; Eling and Wirfs 2015). What is common to these approaches is the decision maker's latent concern for the distributional characteristics of losses which are generally not normal and exhibit long 'tails', what Uryasev refers to as "non-symmetric return-loss distributions" (Uryasev 2000). In operational circumstances characterized by infrequent but large losses, it is the 'downside risk' exceeding some limit that becomes of most concern since individual losses can be catastrophic (Nawrocki 1999). The focus on downside risk was formalized by Roy who, contra to Markowitz, proposed that investors faced with potentially asymmetric and catastrophic losses (i.e. the loss of the entire portfolio) are primarily concerned not with the overall mean/variance return of a portfolio, but rather with the downside risk of their portfolio as compared to some benchmark loss and would therefore adopt a 'safety-first' approach to investing which ensures that losses do not exceed some 'disaster level' (Markowitz 1952; Roy 1952). The safety first principle and various derivations of it have been broadly tested in the context of choice over uncertain prospects (Levy and Sarnat 1972; Levy and Levy 2009; Zeisberger 2013; Zeisberger 2014) and include studies indicating both over – (Lopes and Oden 1999)  and underweighting (Camerer 2000; Hertwig, Barron et al. 2004) of small probabilities and the concept of 'regret' over those losses (Loomes and Sugden 1982).

The following diagram illustrates the typical frequency distribution of business 'portfolio' losses in this context, bounded by zero on the left and, depending on the business context, with potentially no upper limit on the right. Two risk statistics reflecting the low probability right hand 'tail loss' are introduced: 'Value at Risk' (VaR) and 'Conditional Value at Risk" (CVaR):

**Figure 56 - Typical frequency distribution of losses with VaR and CVaR measures**

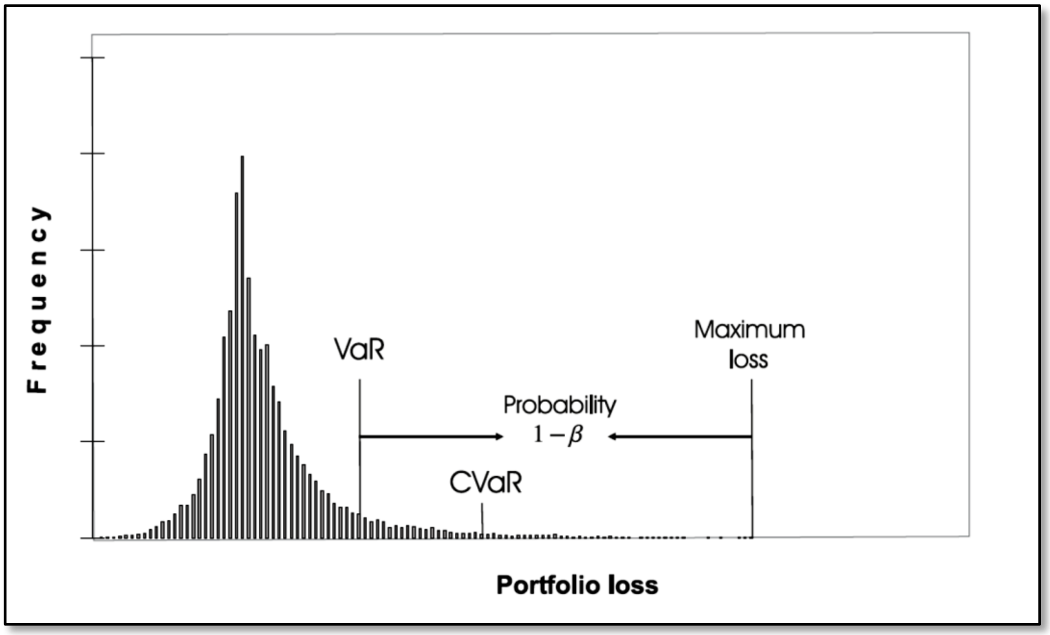

(Uryasev 2000)

According to Uryasev:

> Probabilistic and quantile (percentile) functions are commonly used for the analysis of models with uncertainties or variabilities in parameters. In financial applications, the percentile of the losses is called Value-at-Risk (VaR). VaR, a widely used performance measure, answers the question: what is the maximum loss with a specified confidence level? Percentiles are also used for defining other relevant performance measures, such as Conditional Value-at-Risk (CVaR). CVaR (also called Mean Excess Loss, Mean Shortfall, or Tail VaR) is the average loss for the worst x% scenarios (e.g., 5%). (Uryasev 2000)

VaR is typically expressed as the maximum loss (or other measure of the performance of a system) at a specified confidence level (e.g. at the 95% level of the probability distribution of losses), typically representing some catastrophic or 'disaster' level of concern to management. CVaR is then defined as the average or expected value of losses exceeding the selected disaster threshold where the business operator is concerned with the nature of all possible losses exceeding the chosen threshold (Uryasev 2000). Risk measures that are similar to CVaR and/or may coincide with it, are Expected Shortfall and Tail VaR

(Acerbi and Tasche 2002; Krokhmal 2007). In terms of associated internal controls concerned with mitigating impact in the context of credit losses, for example, CVaR would represent the capital coverage required to meet individual account losses exceeding a management chosen maximum threshold or confidence level:

**Figure 57 - CVaR as the expected value of losses exceeding the VaR threshold**

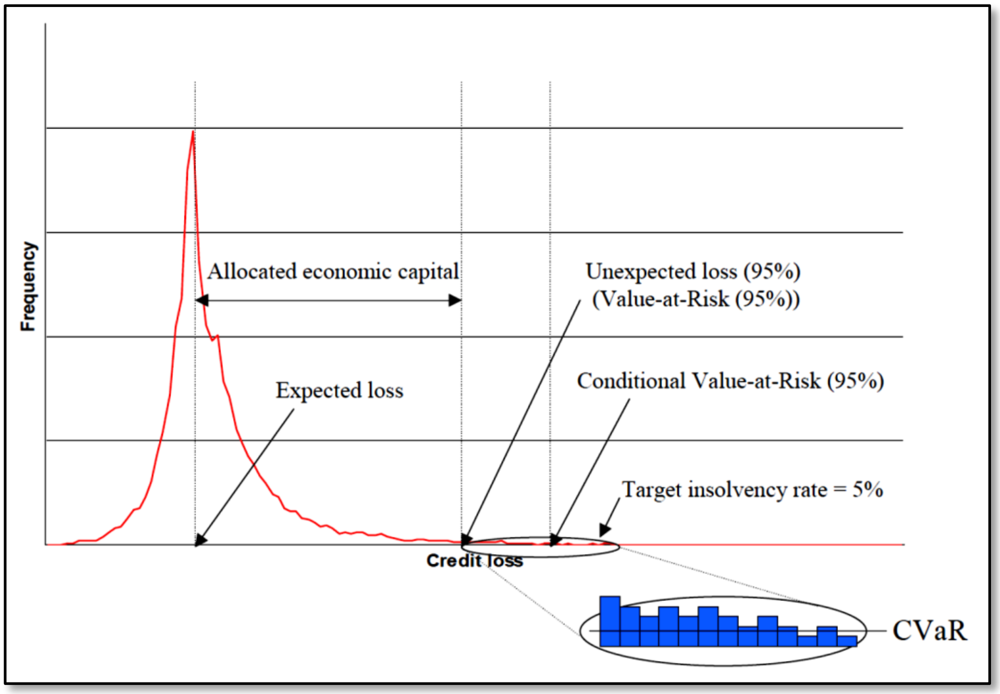

(Uryasev 2000)

From an information security perspective, CVaR could be used to represent the expected value of business losses attributed to successful attacks from security incidents exceeding a chosen VaR loss level (Jaisingh and Rees 2001; Conrad 2005; Ozcelik and Rees 2005; Conrad, Oman et al. 2006; Wang, Chaudhury et al. 2008; Thomas 2009; Gheorghe 2012; Sawik 2013; Thomas, Antkiewicz et al. 2013; WEF 2015). From a security controls perspective, CVaR could also represent the expected value of *insurable* losses for a single firm resulting from security incidents exceeding a chosen VaR. Boehmer has indicated the use of VaR and CVaR in security operations management in the development of security outcome scenarios for SCADA systems critical infrastructure risks based on the stochastic properties of these systems:

**Figure 58 – Downside 'risk corridor' for the time interval (T)**

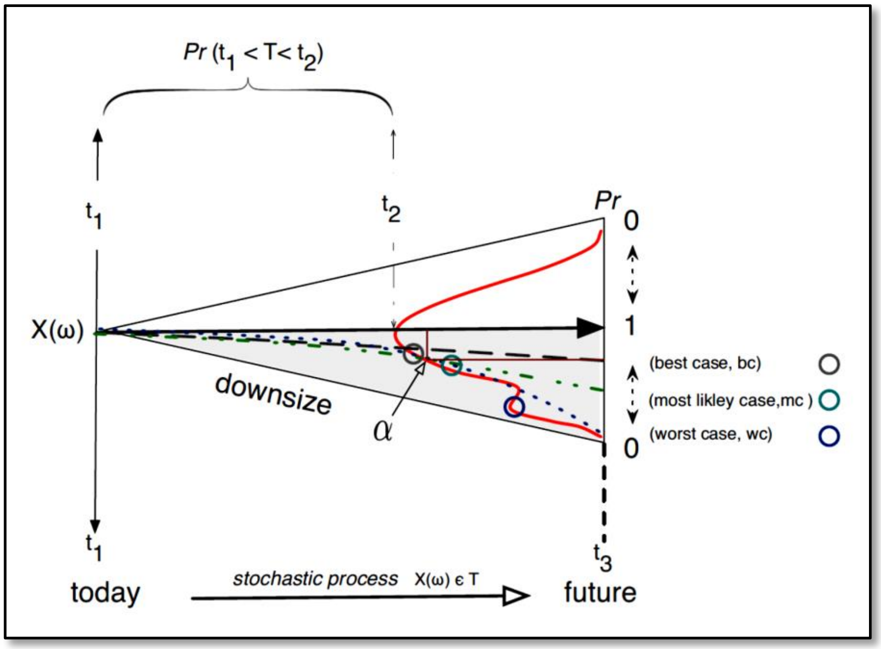

(Boehmer 2011; Boehmer 2012)

The distributional characteristics of losses attributed to security incidents are therefore a key consideration for the decision maker choosing between controls that affect either or both the probability or impact of the logical security risk factors. Biener has recently demonstrated that cyber risk losses share the asymmetric distributional characteristics of operational business losses although the tail risk itself may not be as wide as that of operational losses generally:

**Figure 59 - Boxplots of cyber and non-cyber risk business losses**

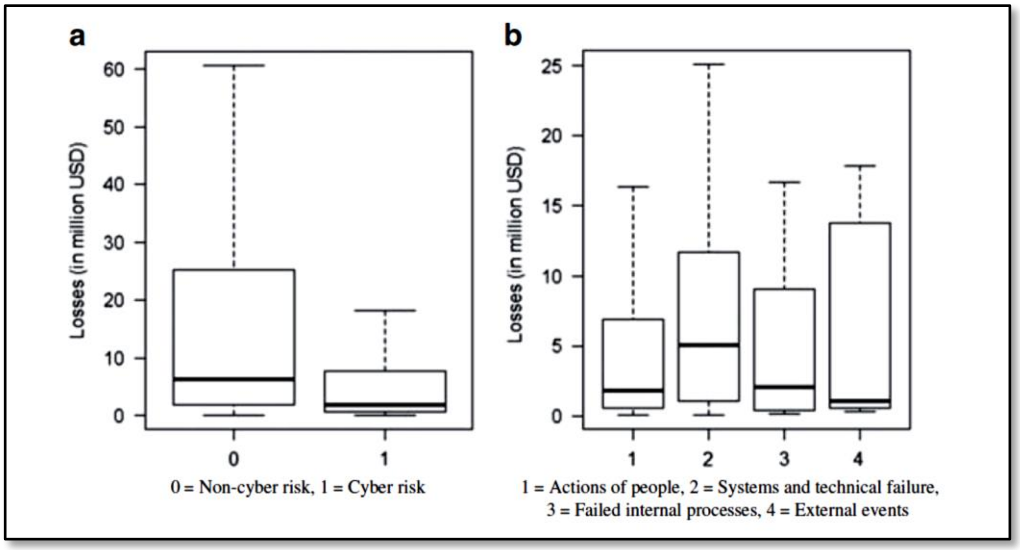

(a) non-cyber risk vs cyber risk; (b) cyber risk by event category

(Biener, Eling et al. 2015)

Although attributed information security losses may not be of the same scale as operational losses generally, the distribution of security losses and the specific risk measures associated with that distribution should be of concern to managers since the average and variance of losses do not capture the potential for large or catastrophic losses (Artzner, Delbaen et al. 1999; Krokhmal 2007; Krokhmal, Zabarankin et al. 2011).

The notion of stochasticity in the business process and security model itself may also be a cause for poor management uptake and use of representative information about the distribution of losses. Krokhmal illustrates the essential conceptual transition required for a decision maker from a deterministic situation without ‘risk’ (and therefore without uncertainty), to a *prospective* situation inherently involving ‘chance’ and which defines the decision problem from a behavioural economic perspective based on, among other biases, risk attitude or preference (Krokhmal, Zabarankin et al. 2011). He first introduces the general deterministic optimization case where we choose a decision or design vector *x* (in our case, a portfolio of security controls) to maximize some objective function (e.g. losses, or perhaps net gains) *f(x)* (>=0), subject to some costs *g(x)* (<=0), and then introduces ‘uncertainty’ described by a random element *ξ*, which leads to situations where, instead of just *f(x)* and *g(x)* we have *f(x, ξ)* and *g(x, ξ)* and where *ξ* can be represented as a probability that is known or can be estimated. This brings us back to the essential decision problem in the context of risk and uncertainty, but informed by the inherent nature of stochastic systems:

> A serious difficulty, however, is that the decision *x* must be chosen before the outcome from this distribution can be observed. One cannot then simply replace *f(x)* by *f(x, ξ)*, because a choice of *x* only produces a random variable *X = f (x, ξ)* whose realization is not yet known, and it is difficult to

> make sense of "minimizing a random variable" as such. Likewise, *g(x)* cannot just be replaced by *g(x, ξ)*…Over the years, a number of approaches have been developed to address these issues; a familiar and commonly used approach is to replace functions *f(x, ξ)* and *g(x, ξ)* with their expected values, e.g., *f(x, ξ)* → $E_\xi$[*f(x, ξ)*].
>
> Being intuitively appealing and numerically efficient, this generic method has its limitations, which have long been recognized in literature. In particular, replacing a random objective function with its expected value implies that (i) the decision obtained as a solution of the stochastic programming problem will be employed *repeatedly* under identical or similar conditions (also known as the "long run" assumption); and (ii) the variability in realizations of the random value *f(x, ξ) is not significant* [emphasis added]. As it poses no difficulty to envisage situations when these two assumptions do not hold, a work-around has to be devised that will allow for coping with models that do not comply with (i) and (ii).
>
> A rather general remedy is to bring the concept of risk into the picture, with "risk" broadly defined as a *quantitative expression of a system of attitudes, or preferences with respect to a set of random outcomes.(Krokhmal, Zabarankin et al. 2011)*

Two key insights can be drawn from this. First, there is the reminder that we are dealing with prospective and (either) risky or uncertain circumstances which essentially defines the problem in terms of the perspective of the decision maker i.e. consideration of a choice in which the chooser seeks 'optimization' of the future state must either assume some type of risk preference or seek to describe risk preference in terms of decisions made in the face of such risky or uncertain prospects. From descriptive point of view, this is the basis for the behavioural economics and lab experiment approach to testing for decisional biases. Second, Krokhmal reminds us that metrics based on the expected valuation of risky or uncertain outcomes implies that we assume either that the decision maker would make the same choice across all potential outcomes (i.e. as if deciding only over a sufficiently large number of repeated choices approaching 'the long run') or that the decision maker effectively ignores the variability in outcomes for any sub-set of choices smaller than that which would produce a near approximation to the long run outcome itself. Since this is clearly not generally the case with human decision makers in experience, what emerges is a definition of 'risk' as not the description of an inherently stochastic process itself, but rather as the "…quantitative expression of a system of attitudes, or preferences with respect to a set of random outcomes." i.e. essentially a description of the decision maker him/herself in terms of their epistemological view of the stochastic system. To my knowledge this is the only such description of risk in the literature and directly informs the contribution of this study in the description of biases in the context of security control choice under risk and uncertainty.

From a practitioner's point of view, this is another way of saying that probability times impact itself is not a very good measure of risk in a security context and may be expected to lead to poor control decisions as a result. A number of questions arise from this insight including whether and to what extent management may be focused on 'expected' or average losses versus the higher moment distributional characteristics of those losses and the bias introduced by risk measures which, in that sense, are not 'coherent' in the sense introduced by Artzner et al from a decision making perspective (Artzner, Delbaen et al. 1999). This also leads to considerations for the generation and representation of loss information that is appropriate for management to use in making control decisions in these circumstances. The following discussion will relate the general notion of the 'coherency' of a risk measure (Artzner, Delbaen et al. 1999; Rockafellar 2007) to the topic of this research which is whether and to what extent uncertainty potentially affects the decisions of security risk managers.

The focus on tail risk in a security context is therefore important since, when we refer to risk, we are clearly referring to variability in business outcome 'losses' and not gains, and generally not some measure of 'net' losses where operational gains from a particular system security posture (e.g. low authorization thresholds and therefore relatively more efficient access to corporate information by insiders versus some benchmark) are seen to offset operational losses from a security incident (e.g. theft of corporate information by an unauthorized insider) based on the same security control posture. Although the 'net' outcome of the security posture could certainly be calculated, it's not clear that this would represent a practical measure against which management would typically make decisions regarding deployed controls intended to primarily prevent losses and not, in the first order, to produce gains. In this sense, we define security losses and the 'disutility' that they create as the opposite' of the utility generated from business gains and conceptually equivalent to a measure of 'regret' over potential uncertain losses (Rockafellar and Uryasev 2002; Rockafellar and Uryasev 2013), where regret *V(X)* of a stochastic loss X in dollars "…can be the compensation deemed appropriate for taking on the burden of the uncertain loss *X*' (Rockafellar and Uryasev 2013). In the context of decisions over prospective losses this is, conceptually, the 'mirror' case of the equity premium required to induce investors to undertake a risk position in gains. 'Regret' has been proposed as an expression of negative utility over losses compared to a benchmark (Dembo and King 1992) and a fundamental element of a 'risk quadrangle' supporting the estimation and optimization of decisions involving inherent loss or hazards (Rockafellar and Uryasev 2013). In this conception, 'risk' of loss can be decomposed into four elements which together define the overall model of prospective uncertain loss over which questions regarding both optimization and estimation/description can be formulated:

> The context is that of random variables [*X*] that can be thought of as standing for uncertain "costs" or "losses" in the broadest sense, not necessarily monetary (with a negative "cost" corresponding perhaps to a "reward"). The language of cost gives the orientation that we would like the outcomes of these random variables to be lower rather than higher, or to be held *below some threshold* [emphasis added]. All sorts of indicators that may provide signals about hazards can be viewed from this perspective. The quadrangle elements provide numerical "quantifications" of them (not only finite numbers but in some cases infinity (∞) which can be employed for various purposes.

$$\begin{array}{lcr} & \text{risk } \mathcal{R} \longleftrightarrow \mathcal{D} \text{ deviation} & \\ \text{optimization} & \uparrow\downarrow \quad \mathcal{S} \quad \downarrow\uparrow & \text{estimation} \\ & \text{regret } \mathcal{V} \longleftrightarrow \mathcal{E} \text{ error} & \end{array}$$

*R(X)* provides a numerical surrogate for the overall hazard in X,
*D(X)* measures the "nonconstancy" in X as its uncertainty,
*ℰ (X)* measures the "nonzeroness" in X,

> *V(X)* measures the "regret" in facing the mix of outcomes of X,
> *S(X)* is the "statistic" associated with X through E and V.
>
> (Rockafellar and Uryasev 2013)

The idea of a 'risk measure', R(X), is then based on 1) the inherent stochasticity of the losses; and 2) the idea that the risk measure must reflect the desire for losses to be 'adequately' (i.e. from a decision makers point of view, ranging from 'some of the time' to 'most of the time' or equivalently 'adequately') below some chosen value:

> In denoting a random cost by *X* and a constant by *C*, a key question is how to give meaning to a statement that *X* is "adequately" $\leq C$ with respect to the preferences of a decision maker who realizes that uncertainty might inescapably generate some outcomes of *X* that are > *C*. The role of a risk measure *R*, in the sense intended here, is to answer this question by aggregating the overall uncertain cost in *X* into a single numerical value *R(X)* in order to
>
> *model "X adequately ≤ C" by the inequality R(X) ≤ C.*
>
> There are familiar ways of doing this. One version could be that *X* is $\leq C$ on average, as symbolized by $\mu(X) \leq C$ with $\mu(X)$ the mean value, or in equivalent notation (both are convenient to maintain), $E(X) \leq C$ with $E(X)$ the expected value. Then $R(X) = \mu(X) = E(X)$. A tighter version could be $\mu(X) + \lambda\sigma(X) \leq C$ with $\lambda$ giving a positive multiple of the standard deviation $\sigma(X)$ so as to provide a safety margin reminiscent of a confidence level in statistics; then $R(X) = \mu(X) + \lambda\sigma(X)$. The alternative idea that the inequality should hold at least with a certain probability $\alpha \in (0, 1)$ corresponds to $q_\alpha(X) \leq C$ with $q_\alpha(X)$ denoting the α-quantile of *X*.
>
> …A choice of R corresponds to an expression of preferences toward risk, but it might not yet be clear why some measures of risk are better motivated or computationally more tractable than others. The key challenge is that most applications require more than just looking at *R(X)* for a single *X*, as far as optimization is concerned. Usually instead, there is a random variable that depends on parameters $x_1, \ldots ,x_n$. We have $X(x_1, \ldots ,x_n)$ and it becomes important to know how the numerical surrogate $R(X(x_1, \ldots ,x_n))$ depends on $x_1, \ldots ,x_n$. *This is where favorable conditions imposed on R, like convexity and monotonicity* [emphasis added], are indispensable.
>
> …Suppose next, though, that these cost-like expressions are *uncertain* through dependence on additional variables - random variables - whose realizations will not be known until later. A decision *x* merely results then in *random variables*…which can only be shaped in their distributions through the choice of x, not pinned down to specific values. Now there is no longer a single, evident answer to how optimization should be viewed, but risk measures can come to the rescue.
>
> (Rockafellar and Uryasev 2013)

Here we have the idea that the expected value of loss *could* be a risk measure but is better informed by incorporating the *variance* of loss. Rockafellar then introduces the idea that losses themselves are normally understood to be a function of some decision variable *x*, representing costs which, in our context, represent some vector of security controls, over which the decision maker chooses in order to 'optimize' losses expressible in *R(X)*:

> A typical situation in optimization that illustrates the compelling need for measures of risk revolves around a family of random "costs" that depend on a decision vector *x* belonging to a subset …

$X_i(x)$ for $i = 0, 1, \ldots, m$; where $x = (x_1, \ldots, x_n)$. (1.1)

> The handicap is that *x can usually do no more than influence the probability distribution of each of the "costs"* [emphasis added]. A potential aim in choosing $x$ from S [a subset of all possible choices of $x$] would be to keep the random variable $X_i(x)$ adequately $\leq c_i$ for $i = 1, \ldots, m$, while achieving the lowest $c_0$ such that $X_0(x)$ is adequately $\leq c_0$. The way "adequately" is modeled could be different for each $i$, and the notion of a risk measure provides the perfect tool. A selection of risk measures $R_i$ that pins down the intended sense of "adequately" in each case leads to an optimization problem having the form
>
> *choose* $x \epsilon S$ *to minimize* $R_0(X_0(x))$ *subject to* $R_i(X_i(x)) \leq c_i$ *for* $i = 1, \ldots, m$ (1.2)
>
> (Rockafellar and Uryasev 2013)

The nature of the specific deviation of $X$, $D(X)$, is then introduced as a generalization of $\sigma(X)$, where such measures include standard deviation as a special case but need not be symmetric with respect to losses versus gains. More importantly, measures of $D(X)$ can be chosen to reflect the risk preference of the decision maker, particularly where the primary interest is in losses, and catastrophic losses in specific contexts (Rockafellar, Uryasev et al. 2006; Grechuk, Molyboha et al. 2009) and have been motivated by experience in financial portfolio selection, here mapped to the idea of a portfolio of security controls instead of the selection of individual financial securities:

> In portfolio theory, the rate of return of the portfolio is a random variable $X(x)$ depending on the vector $x$ that gives the proportions of various securities included in the portfolio. Bounds are placed on $\sigma(X(x))$ or this quantity is minimized subject to side conditions on $x$. Such an approach can be justified when the random variables have normal distributions, but when the heavy tail behaviour of non-normal distributions enters the scene, doubts arise. It may be better, then, to replace standard deviation by a different deviation measure, which perhaps could even act on $X(x)$ asymmetrically.
>
> (Rockafellar and Uryasev 2013)

One of the most interesting components of this approach is then the inclusion of the concept of *regret* which endogenizes the concept of utility theory within the model of risk and reflects the consideration that decision makers will generally exhibit *loss aversion* relative to a prior state of wealth as a *benchmark* (Tversky and Kahneman 1992) and which makes intuitive sense from a security practitioner's risk perspective. This also reinforces the idea that the 'correct' risk statistic in the context of operational gains and losses may be one that primarily, if not exclusively, reflects losses and not gains or even net losses/gains and can be thought of as an 'expectation-bounded risk measure' (Rockafellar, Uryasev et al. 2002). The introduction of 'expectations' within the risk model is an attempt to fundamentally incorporate the decision makers' consideration of prospective choice under risk and uncertainty, where, as noted by Krokhmal above, the decision maker faces an inherently uncertain system at risk where the objective function $f(x, \xi)$ and the corresponding control choices $g(x, \xi)$ are both stochastic and the decision maker must necessarily therefore possess *states of belief*, $\omega$, with some probability P, over the randomness of the objective function and, underlying that, of the controls logically affecting losses. Risk aversity in the context of loss is then defined as a latent condition of the 'regret' of the decision maker who desires to specifically and routinely avoid losses under these circumstances: Rockafellar notes that:

> Recall that [in the simplest case] *R(X) =E(X)* for constant *X ≡ C*. Aversity has the interpretation that the *risk of loss in a nonconstant random variable X cannot be acceptable unless, in particular, X(ω)<0 on average* [emphasis added]. Note that relations to expectation, and consequently to the particular choice of P, have not entered axiomatically until this point (Rockafellar 2007).
>
> In a measure of regret *V*, the value *V(X)* stands for the net displeasure perceived in the potential mix of outcomes of a random "cost" *X* which may sometimes be > 0 (bad) and sometimes ≤0 (OK or better). Regret comes up in penalty approaches to constraints in stochastic optimization and, *in mirror image, corresponds to measures of "utility" U in a context of gains Y instead of losses X* [emphasis added]...Regret obeys *V*(0) = 0, so…we have to focus on utility measures that have *U*(0) = 0; we say then that *U* is a measure of relative utility…Focusing on relative utility in this sense is a positive feature of the quadrangle scheme because it can help to capture the sharp difference in attitude toward outcomes above or below a *benchmark* that is increasingly acknowledged as influencing the preferences of decision makers.
>
> Measures of regret *V*, like measures of deviation *D*, are profoundly related to measures of risk *R*… Especially important will be a one-to-one correspondence between measures of deviation and measures of risk under "aversity," regardless of coherency [of the risk measure]. A powerful property of measures of regret, which soon will be discussed, is their ability to generate measures of risk through trade-off formulas. By means of such formulas, an optimization problem in the form of (1.2) [above] *may be recast in terms of regret instead of risk* [emphasis added], and this can be a great simplification. Furthermore, by revealing a deep connection between risk measures and utilty, regret reconciles the seemingly different approaches to optimization based on those concepts.
>
> (Rockafellar and Uryasev 2013).

Finally, the model introduces a 'measure of error' in *X*, $\mathcal{E}$(X), "…that quantifies the nonzeroness in *X*" and reflects the inherently stochastic nature of *X*, where *X* is a non-deterministic function of some vector of inputs *x*. In the quadrangle model, the error term is fundamental to both the concept of deviation (uncertainty) and the concept of regret over the outcome of a random variable based on *x*:

> Given an error measure $\mathcal{E}$ and a random variable *X*, one can look for a constant *C* nearest to *X* in the sense of minimizing *$\mathcal{E}$(X - C).* The resulting minimum "$\mathcal{E}$-distance," denoted by *D(X)*, turns out to be a deviation measure…The *C value in the minimum, denoted by S(X)* [emphasis added], can be called the "statistic" associated with *X* by $\mathcal{E}$…The emergence of a particular deviation measure *D* and statistic *S* from the choice of an error measure $\mathcal{E}$ has intriguing implications for statistical estimation in the sense of generalized regression…There is furthermore a deep connection between regression and an optimization problem like (1.2) [above]. The *x*-dependent random variables $X_i(x)$ there might be replaced by convenient approximations $\hat{X}_i(x)$ developed through regression…The optimization and estimation sides of the quadrangle are bound together not only through such considerations, but also in a more direct manner.
>
> (Rockafellar and Uryasev 2013).

In total, the quadrangle model specifies the relationships between a measure of risk, loss deviation, loss error and decision maker regret for both estimation (descriptive) and, in the context of choice, optimization (normative) purposes in which, given a measure of decision maker regret and the description of inherent error in the system at risk, the decision maker can select a portfolio of controls *x* in order to 'optimize' the resulting losses in the context of the decision makers' utility over losses with reference to a benchmark:

> The rule that projects from $\mathcal{E}$ onto *D* is echoed by a certainty-uncertainty trade-off formula which projects a regret measure *V* onto a risk measure *R*. This formula, in *which C + V(X - C)* is minimized over *C*, generalizes a rule in (Rockafellar and Uryasev 2000) and (Rockafellar, Uryasev et al. 2002), for VaR-CVaR computations…Under a simple relationship between *V* and $\mathcal{E}$*, the optimal C value in the trade-off is the same statistic S(X) as earlier* [emphasis added], but that

> conceptual bond has been missed. *Nothing has hitherto suggested that "error" in its context of approximation might be inherently related to the very different concept of "regret" and, through that, to "utility"*.[emphasis added]. … The paired arrows on the sides of [the diagram], in contrast to the two-way arrows on the top and bottom, correspond to the fact that the simple formulas…for getting R and D from V and $\mathcal{E}$ are not uniquely invertible. Antecedents *V* and $\mathcal{E}$ for *R* and *D* always exist, even in multiplicity, so the real issue for inversion is the identification of *natural, nontrivial antecedents.*

$$\begin{tikzcd}
\text{'Optimization'} & R(X) = E(X) + D(X) \arrow[r, leftrightarrow] \arrow[d, shift left] & D(X) = R(X) - E(X) \arrow[d, shift left] & \text{'Estimation'} \\
& V(X) = E(X) + \mathcal{E}(X) \arrow[r, leftrightarrow] \arrow[u, shift left] & \mathcal{E}(X) = V(X) - E(X) \arrow[u, shift left] &
\end{tikzcd}$$

where

$$R(X) = \min_{C}\{C + V(X - C)\} \qquad D(X) = \min_{C}\{E(X - C)\}$$

and where

$$\operatorname{argmin}_{C}\{C + V(X - C)\} = S(X) = \operatorname{argmin}_{C}\{C + \mathcal{E}(X - C)\}$$

where "argmin" refers to 'the C value that achieves the "min"

> This theorem integrates, in a new and revealing way, various results or partial results that were separately developed elsewhere…Although the parallel between E → D and V → R, which ties the two sides of the quadrangle fully together, is mathematically elementary, it has not come into focus easily despite its conceptual significance. That, especially, is where the theorem innovates. What was absent in the past was the broad concept of a measure of regret, not limited to an expectation, and the realization it could anchor a fourth corner in the relationships, thereby serving as a conduit for bringing in "utility" beyond expected utility.
>
> (Rockafellar and Uryasev 2013)

Having defined the relationship between a measure of risk in losses and a decision maker's regret as a function of risk, Rockafellar then proceeds to define VaR and CVaR in terms of these statistics:

> The key in this case is provided by the (cumulative) distribution function $F_X(x)$ = prob{$X \leq x$} of a random variable *X* and the quantile values associated with it. If, for a probability level α ϵ (0, 1), there is a unique *x* such that $F_X$(x) = α , that *x* is the α-quantile $q_\alpha(X)$. In general, however, there are two values to consider as extremes i.e. the lower threshold of losses and the upper bound, potentially infinite]:
>
> $$q^+_\alpha(X) = \inf\{x \mid F_X(x) > \alpha\} \qquad q^-_\alpha(X) = \sup\{x \mid F_X(x) < \alpha\} \tag{2:1}$$
>
> It is customary, when these differ, to take the lower value as "the" α-quantile, noting that, because $F_X$ is right-continuous, this is the lowest *x* such that $F_X$(x) = α. Here, instead, we will consider the entire interval between the two competing values as the quantile,
>
> $$q_\alpha(X) = [q^-_\alpha(X), q^+_\alpha(X)], \tag{2:2}$$
>
> bearing in mind that this interval usually collapses to a single value [i.e. the lower value]. That approach will fit better with our way of defining a "statistic" by the argmin notation [noted in the boxed figure above]. Also important to understand, in our context of interpreting *X* as a "cost" or "loss," is that the notion of *value-at-risk* in finance coincides with quantile. There is an upper value-

> at-risk $\mathrm{VaR}^{+}_{\alpha}(X) = q^{+}_{\alpha}(X)$ along with a lower value-at-risk $\mathrm{VaR}^{-}_{\alpha}(X) = q^{-}_{\alpha}(X)$, and, in general, a value-at-risk interval $\mathrm{VaR}_{\alpha}(X) = [\mathrm{VaR}^{+}_{\alpha}(X),\ \mathrm{VaR}^{-}_{\alpha}(X)\ ]$ identical to the quantile interval $q_{\alpha}(X)$.
>
> Besides value-at-risk, the example coming under consideration involves the *conditional value-at-risk* of *X* at level $\alpha \in (0, 1)$ as defined by
>
> $$\mathrm{CVaR}_{\alpha}\ (X) = \text{expectation of } X \text{ in its } \alpha\text{-tail}, \qquad (2:3)$$
>
> which is also expressible by
>
> $$\mathrm{CVaR}_{\alpha}\ (X) = \frac{1}{1-\alpha}\int_{\alpha}^{1} \mathrm{VaR}_{\tau}(X)d\tau \qquad (2:4)$$
>
> The second formula is due to (Acerbi 2002) in different terminology, while the first follows the pattern in (Rockafellar and Uryasev 2002), where "conditional value-at-risk" was coined.
>
> (Rockafellar and Uryasev 2013)

The risk quadrangle can then be re-expressed in terms of VaR and CVaR statistics:

> A Quantile-Based Quadrangle (at any confidence level $\alpha \in (0, 1)$,
>
> $S(X) = \mathrm{VaR}_{\alpha}(X) = q_{\alpha}(X)$ = quantile
> $R(X) = \mathrm{CVaR}_{\alpha}(X) = \bar{q}_{\alpha}(X)$ = superquantile
> $D(X) = \mathrm{CVaR}_{\alpha}(X\text{-}E(X)) = \bar{q}_{\alpha}(X\text{-}E(X))$ = superquantile-deviation
> $V(X) = \frac{1}{1-\alpha}E(X_{+})$ = average absolute loss[47], scaled
> $\mathcal{E}(X) = E[\frac{1}{1-\alpha}X_{+} + X_{-}]$ = normalized Koenker-Bassett[48] error
>
> (Rockafellar and Uryasev 2013)

The choice of $R(X) = \mathrm{CVaR}_{\alpha}(X)$ is a well-known result in finance in the consideration for the choice of 'coherent risk measures' (Artzner, Delbaen et al. 1999) as well as for considerations of optimization, whether formally modelled as described here by Rockafellar in terms of 'stochastic optimization', or considered more practically in the context of day-to-day managerial choice over a portfolio of security controls *x*:

> Consider a stochastic optimization problem in the form of (1.2). It is tempting, and common in many applications, to contemplate taking $\underline{R_i}$ to be a quantile $q_{\alpha_i}$. The constraint $R_i(X_i(x)) \leq c_i$ would require then that *x* be chosen so that the random "cost" $X_i(x)$ *is* $\leq c_i$ with probability at least $\alpha_i$. However, this apparently natural approach suffers from the fact that $q_{\alpha_i}(X_i(x))$ may be poorly behaved as a function of *x* as well as subject to the indeterminacy, or discontinuity, associated with (2.2). That could hamper computation and lead to instability of solutions.
>
> An alternative to a quantile would be to take $R_i$ to be a superquantile $\bar{q}_{\alpha_i}$. The constraint $R_i(X_i(x)) \leq c_i$, as an expression of $X_i(x)$ being "adequately" $\leq c_i$, is then more conservative and has an interpretation in terms of "buffered probability of failure," cf. (Rockafellar and Royset 2010) . Moreover it is better behaved and able to preserve convexity of $X_i(x)$ with respect to *x*, if present. A further advantage in optimization from such an approach is seen from the projection from *V* to *R* on the left side of the quadrangle:
>
> $$\bar{q}_{\alpha_i}(X_i(x)) \leq c_i \iff C_i + \frac{1}{1-\alpha_i} E[\max\{0, X_i(x) - C_i\}] \leq c_i \text{ for a choice of } C_i \in \mathbb{R}$$

[47] (Dembo and King 1992)
[48] (Koenker and Bassett Jr 1978)

> Thus, a superquantile (or CVaR) constraint *can be reformulated as something simpler through the introduction of another decision variable $C_i$ alongside of x* [emphasis added].
>
> (Rockafellar and Uryasev 2013)

'Coherency' is important attribute for risk measures since not all risk measures are coherent and portfolio choices made with non-coherent risk measures may result in sub-optimal and outcomes inconsistent with the decision maker's risk preferences. Rockafellar discusses the nature and need for *coherency* in the risk measure *R* and, in the context if the risk quadrangle, his replacement of 'coherent' with the term 'regular' illustrating that, while $R(X) = \mathrm{E}(X)$ may be coherent in Artzner's sense, in the context of uncertainty it lacks the normative decisional characteristic of loss aversion and it cannot be used to represent a measure of acceptable *loss* in the context of uncertainty, where 'acceptability' means losses typically less than a benchmark:

> Specifically, we suppose there is an underlying space Ω with elements $\omega$ standing for future states, or scenarios, along with a measure which assigns probabilities to various subsets of Ω…Random variables from now on are functions $X : \Omega \rightarrow \mathbb{R}$, but we restrict attention to those for which $E[X^2] < \infty$, indicating this by $X \epsilon L^2(\Omega)$[49]. Here *E* is the expectation with respect to the background probability measure on Ω. Any $X \epsilon L^2(\Omega)$ also has $E|X| < \infty$, so that $E(X)$ is well defined and finite…In many applications Ω may consist of finitely many elements $\omega$, each having a positive probability weight.
>
> ….
>
> The quantifiers *R*, *D*, *V* and $\mathcal{E}$, all of which assign numerical values, possibly including $+\infty$ to random variables X, are said to be "functionals" [denoted as '$\mathcal{F}$'] on $L^2(\Omega)$. Some of the properties that come up may be shared, so it is expedient to state them in terms of a general functional $\mathcal{F} : L^2(\Omega) \rightarrow (-\infty,\infty)$:
>
> $\mathcal{F}$ is *convex* if $\mathcal{F}((1-\tau)X + \tau X') \leq (1-\tau)\mathcal{F}(X) + \tau\mathcal{F}(X')$ for all X, X′, and τ.
>
> $\mathcal{F}$ is *positively homogeneous* if $\mathcal{F}(0) = 0$ and $\mathcal{F}(\lambda X) = \lambda \mathcal{F}(X)$ for all $\lambda \epsilon (0;1)$.
>
> $\mathcal{F}$ is *subadditive* if $\mathcal{F}(X + X') \leq \mathcal{F}(X) + \mathcal{F}(X')$ for all $X \leq X'$.
>
> $\mathcal{F}$ is *monotonic* (nondecreasing, here) if $\mathcal{F}(X) \leq F(X')$ when $X \leq X'$.
>
> $\mathcal{F}$ is *closed* if, for all $C \epsilon \mathbb{R}$, the set $\{ X \mid \mathcal{F}(X) \leq C$ g is closed.

These properties mirror the axioms for coherency proposed by Artzner, where:

> Convexity will be valuable for much of what we undertake. Positive homogeneity is a more special property which, in the study of risk, was emphasized more in the past than now. An elementary fact of convex analysis is that
>
> $$\mathcal{F} \text{ convex + positively homogeneous } \Leftrightarrow \text{ subadditive + positively homogeneous.} \tag{3:2}$$
>
> …Other important consequences of convexity emerge only in combination with closedness. One that will be applied in several ways is the following rule coming out of convex analysis.
>
> …An immediate consequence, for instance, is that

[49] Where, for our understanding, $L^2(\Omega)$ is a 'LaPlace transform' of the expected future states of the system into a probability distribution of losses expressible in *X*. https://en.wikipedia.org/wiki/Laplace_transform

$$\text{for } \mathcal{F} \text{ closed convex: if } \mathcal{F}(X) \leq 0 \text{ whenever } X \leq 0, \text{ then } \mathcal{F} \text{ is monotonic.} \tag{3:4}$$

> To assist with closedness, it may help to note that this property of $\mathcal{F}$ holds when $\mathcal{F}$ is continuous, and moreover, as long as $\mathcal{F}$ does not take on ∞, that stronger property is automatic in broad circumstances of interest to us. Namely,

$$\mathcal{F} \text{ is continuous on } L^2(\Omega) \text{ when } \begin{cases} \mathfrak{F}\ is\ finite, convex, and\ closed, or \\ \mathfrak{F}\ is\ finite, convex, and\ monotonic, or \\ \mathfrak{F}\ is\ finite, convex, and\ \Omega \text{ is finite} \end{cases} \tag{3.5}$$

Rockafellar then concludes by explaining ‘regular’ risk measures which are generally *closed convex and incorporate risk aversity* in place of otherwise ‘coherent’ risk measures, or measures which are coherent but may not incorporate risk aversity:

> …The role of a measure of risk, *R*, is to assign to a random variable X, standing for an uncertain “cost” or “loss,” a numerical value *R(X)* that can serve as a surrogate for overall (net) cost or loss. However, the assignment must meet reasonable standards in order to make sense.
>
> The class of *coherent* measures of risk has attracted wide attention in finance in this regard. A functional *R* belongs to this class, as introduced in {Artzner, 1999 #241}if it is convex and positively homogeneous …as well as monotonic, and, in addition, satisfies
>
> $$R(X + C) = R(X) + C \text{ for all } X \text{ and constants } C. \tag{3.6}$$
>
> Closedness of *R* was not mentioned in (Artzner, Delbaen et al. 1999)but the context there supposed *R* to be finite (and actually finite, too), so that closedness and even continuity of *R* were implied by coherency through (3.5) Subsequent researchers considered dropping the positive homogeneity, and with it the term “coherent,” speaking then of a “convex measure of risk” or a “convex risk function,” cf. Follmer and Schied [2004], Ruszczynski and Shapiro [2006a].
>
> By a *regular measure of risk* we will mean a functional R with values in (-∞,∞] that is *closed convex* with
>
> $$R(C) = C \text{ for constants } C \tag{3:7}$$
>
> and furthermore
>
> $$R(X) > E(X) \text{ for nonconstant } X \tag{3:8}$$
>
> Property (3.8) is *aversity to risk*.
>
> An example of a coherent measure of risk that is not regular is *R(X)* = E*(X)*, which lacks aversity. On the other hand, $R(X) = VaR_{\alpha}^{-}(X)$ fails to be a regular measure of risk by lacking closedness, convexity and the aversity in (3.8), in general, although it does have positive homogeneity, satisfies (3.6) and is monotonic. It fails to be a coherent measure of risk through the absence of convexity.
>
> (Rockafellar and Uryasev 2013)

### VaR and CVaR Risk Measures for Cyber Insurance Control Choice

VaR and CVaR measures of risk also have application in insurance since they can be used to reflect both the buyer’s preferred benchmark loss level and therefore the deductible level on an insurance policy i.e. the maximal level of loss which the firm is routinely willing to absorb and therefore the lowest preferred loss which would be covered were insurance to be purchased to cover losses at this level or higher (Herath and Herath 2011). The corresponding CVaR would then represent, from the individual firm's perspective, the average expected loss exceeding the deductible. An insurance premium reflective of this CVaR, (i.e. ignoring losses from a market of polled firms and reflecting only this firm) would be deemed 'actuarially fair' if the cost of insurance equalled the expected value of the loss within a specified time frame. For example, assume that the VaR of simulated security losses over 30 years at a 95% confidence level for a firm with a relatively "High" level of security control was approximately $22,150 per day and the corresponding CVaR was $32,110, then the actuarially fair premium would be $32,110 *5% = $1,605.51. This means that, over the long run (i.e. of at least 30 years), and assuming the firm was risk neutral, they should be willing to pay $1605.51 per day to insure against any daily loss exceeding $22,150 and the annual premium would be $1,605.51 * 365 = $586,012:

**Figure 60 - High Controls System: Probability Distribution of Simulated Daily Losses over 30 Years (n=10,950)**

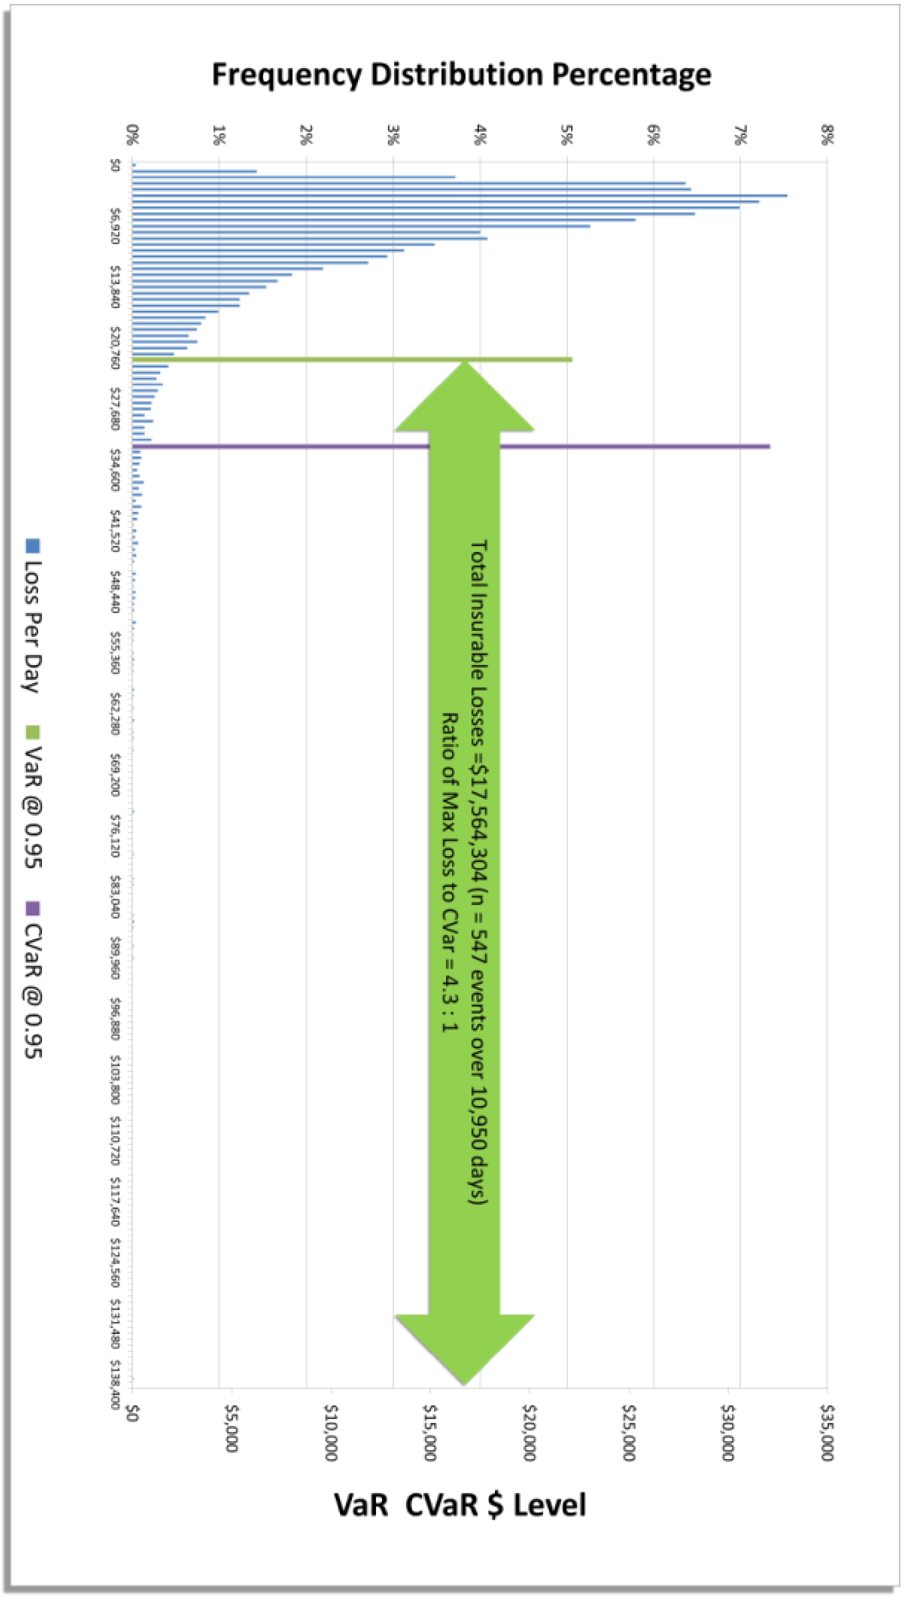

The insurable losses are of course dependent on the underlying security loss profile itself, which in turn is dependent on the relative security level of the controls intended to prevent losses. In the above example we were looking at a firm with relatively ‘high’ strength controls. If the same firm were to deploy relatively ‘medium’ strength controls, the profile of losses is similar but different in two key aspects:

**Figure 61 - Medium Controls System: Probability Distribution of Daily Losses over 30 Years (n=10,950)**

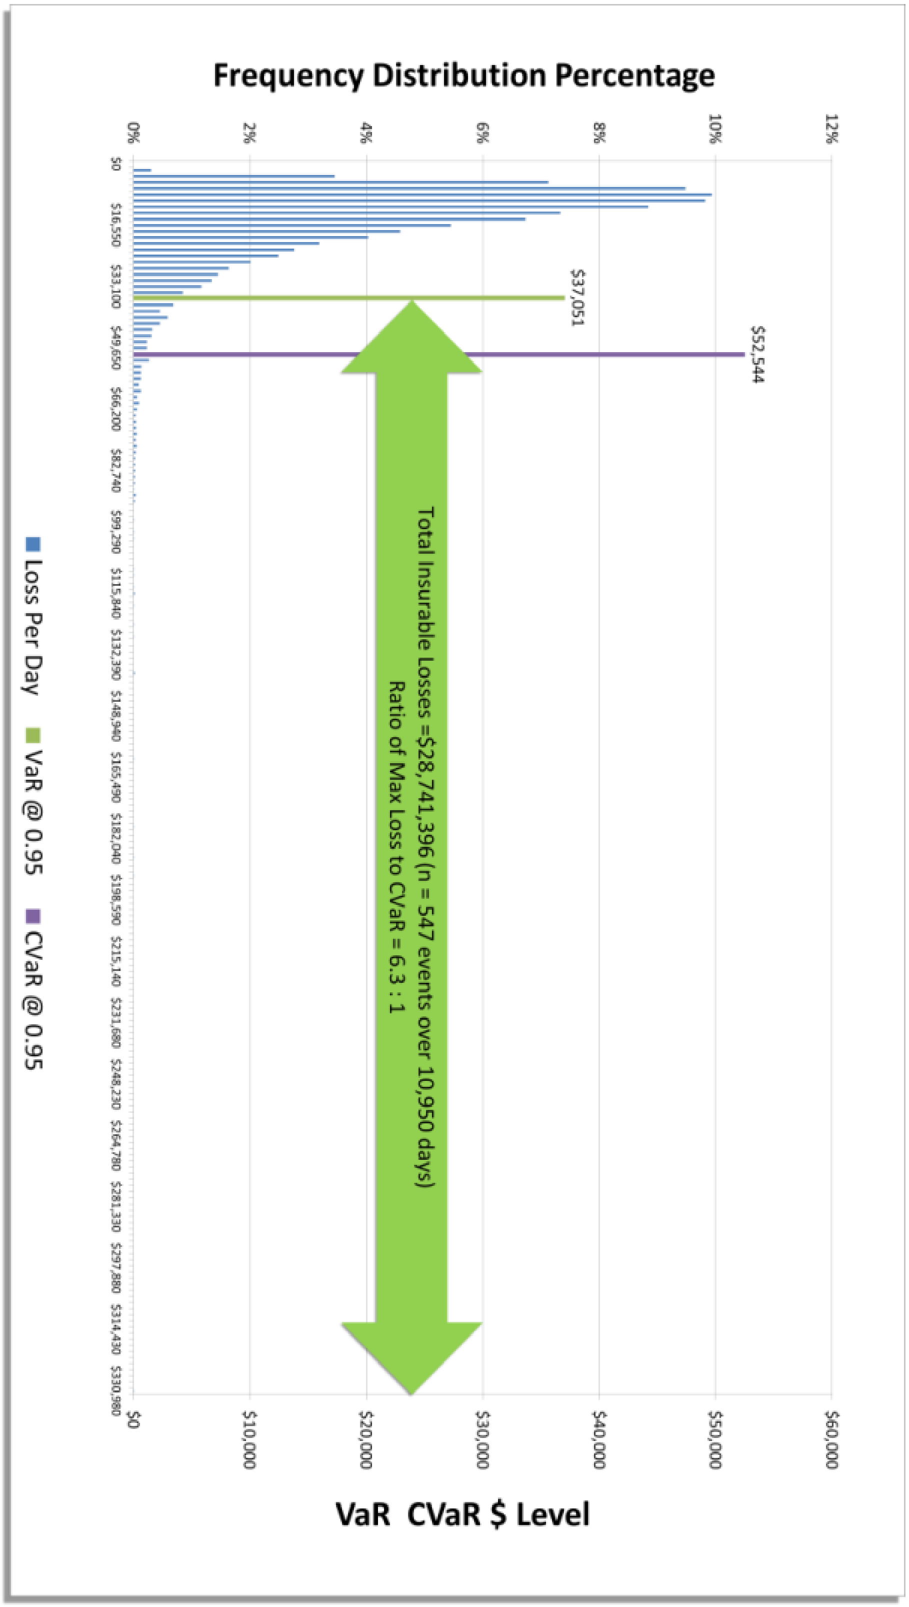

First, since the Medium controls prevent fewer successful attacks, the total dollar amount of losses is almost double in the medium case. Second, the maximum loss in the medium case is also more than double that of the high controls case. This has implications for the trade-off between controls that prevent losses and controls that reduce impacts, like insurance, if losses were to actually occur.

Since the losses in any year are stochastic, the annual difference between the premium paid and the losses claimed will be either positive or negative, but would be expected to roughly net out over the 30 years (since this is the period over which the 'actuarially  fair' premiums are calculated. The following table and chart represent the simulated losses over 30 years under high control conditions and the corresponding annual and cumulative net difference between claims and premiums paid year:

**Table 6 - High Controls Case: Annual vs. Cumulative Excess of Claims Paid Less Premium over 30 Years (VaR @ 95%)**

**(Daily Premium = CVaR Loss amount)**

| Year | Total Loss Claims Exceeding 95% VaR in 365 days (n=~ 18 claimable events per year) | Total Annual Premium: based on 95% VAR Loss + Average Loss exceeding 95% VAR, over 30 years | Annual Excess of Claims Paid Less Premium | Cumulative Excess of Claims Paid Less Premium |
|---|---|---|---|---|
| 1 | $556,966.91 | $586,011.98 | -$29,045.07 | -$29,045.07 |
| 2 | $640,254.38 | $586,011.98 | $54,242.40 | $25,197.32 |
| 3 | $770,744.48 | $586,011.98 | $184,732.50 | $209,929.82 |
| 4 | $558,602.89 | $586,011.98 | -$27,409.09 | $182,520.72 |
| 5 | $680,641.01 | $586,011.98 | $94,629.03 | $277,149.75 |
| 6 | $580,431.88 | $586,011.98 | -$5,580.10 | $271,569.64 |
| 7 | $491,582.91 | $586,011.98 | -$94,429.07 | $177,140.57 |
| 8 | $370,963.63 | $586,011.98 | -$215,048.35 | -$37,907.79 |
| 9 | $572,136.66 | $586,011.98 | -$13,875.32 | -$51,783.11 |
| 10 | $659,259.65 | $586,011.98 | $73,247.67 | $21,464.56 |
| 11 | $449,415.28 | $586,011.98 | -$136,596.70 | -$115,132.15 |
| 12 | $518,657.61 | $586,011.98 | -$67,354.37 | -$182,486.52 |
| 13 | $661,433.67 | $586,011.98 | $75,421.69 | -$107,064.84 |
| 14 | $534,074.04 | $586,011.98 | -$51,937.94 | -$159,002.78 |
| 15 | $705,593.42 | $586,011.98 | $119,581.44 | -$39,421.35 |
| 16 | $479,677.63 | $586,011.98 | -$106,334.35 | -$145,755.70 |
| 17 | $297,912.32 | $586,011.98 | -$288,099.66 | -$433,855.37 |
| 18 | $646,726.18 | $586,011.98 | $60,714.20 | -$373,141.17 |
| 19 | $535,903.64 | $586,011.98 | -$50,108.34 | -$423,249.51 |
| 20 | $713,460.81 | $586,011.98 | $127,448.83 | -$295,800.69 |
| 21 | $601,462.07 | $586,011.98 | $15,450.09 | -$280,350.60 |
| 22 | $600,164.07 | $586,011.98 | $14,152.09 | -$266,198.52 |
| 23 | $710,551.48 | $586,011.98 | $124,539.50 | -$141,659.02 |
| 24 | $779,110.51 | $586,011.98 | $193,098.53 | $51,439.50 |
| 25 | $529,438.30 | $586,011.98 | -$56,573.68 | -$5,134.18 |
| 26 | $595,274.19 | $586,011.98 | $9,262.21 | $4,128.02 |
| 27 | $614,666.64 | $586,011.98 | $28,654.66 | $32,782.68 |

| | | | | |
|---|---|---|---|---|
| 28 | $615,461.09 | $586,011.98 | $29,449.11 | $62,231.79 |
| 29 | $561,230.84 | $586,011.98 | -$24,781.14 | $37,450.64 |
| 30 | $532,506.22 | $586,011.98 | -$53,505.76 | -$16,055.12 |
| | | | | |
| **Sum** | **$17,564,304.41** | **$17,580,359.53** | **-$16,055.12** | |

**Figure 62 - High Controls Case: Annual vs. Cumulative Excess of Claims Paid Less Premium over 30 Years (VaR @ 95%)**

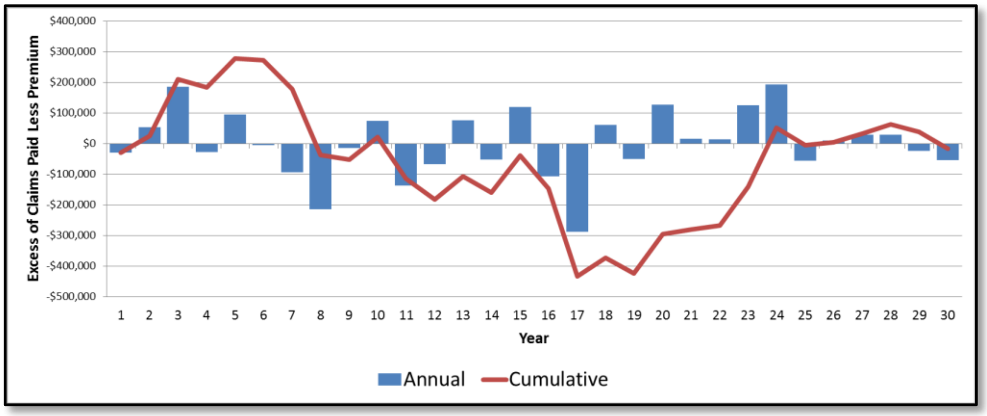

The results should be interesting to a security practitioner making a decision over controls, in this case over cyber insurance covering losses. First, the selection of the VaR level is important since it represents both the level of unacceptable losses which should be chosen carefully depending on the operational needs of the business and particularly since it represents the insurance deductible level – losses below this level will have to be sustained by the business despite paying for insurance. Second, since the resulting CVaR will be more heavily weighted towards the more frequent lower losses at a lower VaR, the net benefits of the insurance over any sub-series of years less than 30 declines as the VaR increases. Depending on the length of the tail losses above the chosen VaR, increasing the VaR will generally cause the CVaR to rise exponentially at increasingly higher levels of VaR. At a 99% level for example, the net benefit of the insurance is negative for longer stretches of time since insurable losses are less frequent (3.6 per year), although over the entire 30 years the benefits still net out to roughly zero. At the 99.9% level (representing insurance coverage for only truly catastrophic loss, approximately once every 3 years) and, in this particular simulation, the insurance actually only pays off once in 30 years:

**Figure 63 - High Control Case: Annual vs. Cumulative Excess of Claims Paid Less Premium over 30 Years (VaR = 99%)**

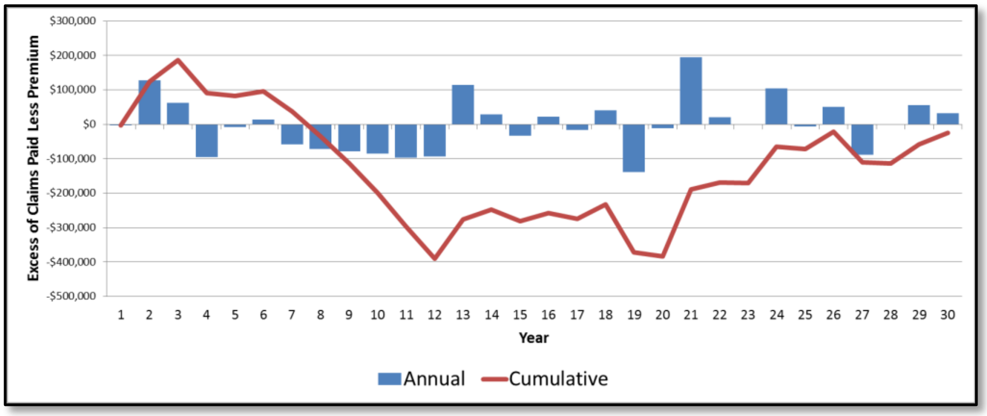

**Figure 64 - High Control Case: Annual vs. Cumulative Excess of Claims Paid Less Premium over 30 Years (VaR @ 99.99%)**

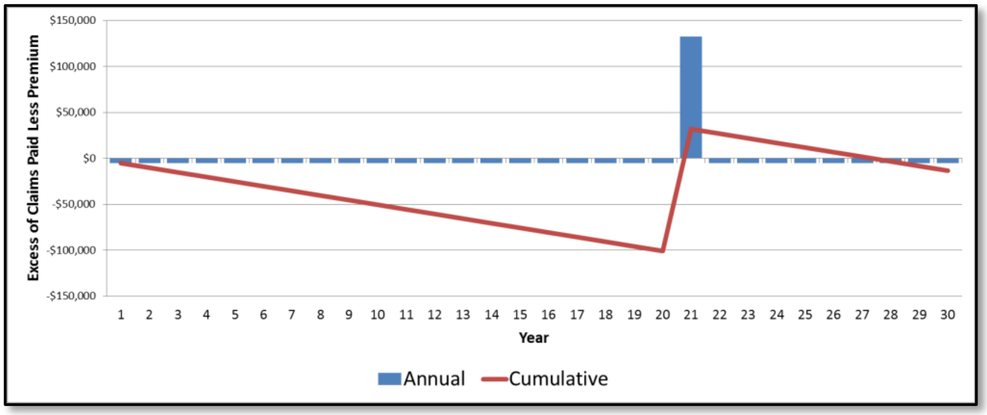

**Figure 65 - High Control Case: VaR vs. CVaR Loss Levels and # of Insurable Claims Per Year**

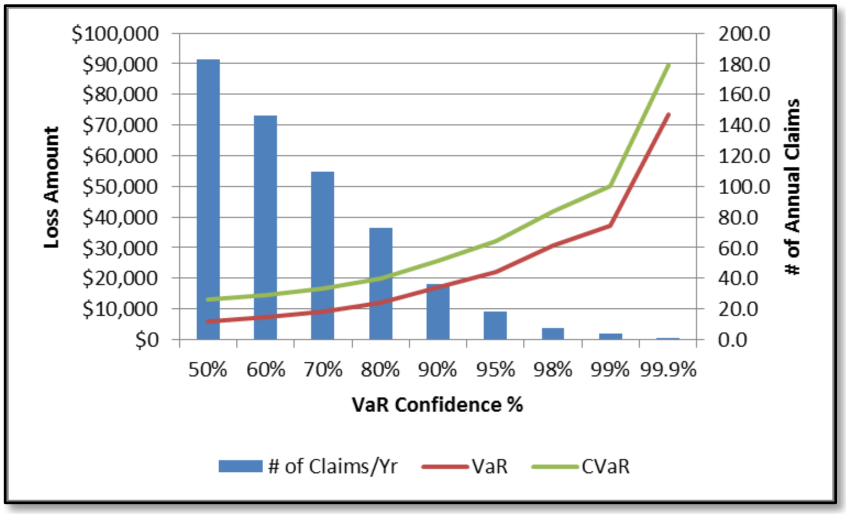

It is apparent that the potential range of losses and the maximum loss above the VaR changes the overall value of the insurance, so the shape of the loss distribution, particularly in the tail which is typically sparse, becomes important. Comparing the High and Medium control cases at the 95% VaR level, we see that the decision maker under the Medium Controls case although covered for the same number of insurable events over 30 years (since by definition the VaR level in both cases represents a level of losses in which only 5% of events will exceed the VaR) both the amount of insurance paid and the volatility of the net insurable losses is substantially greater than in the High Controls case:

**Figure 66 - Medium vs. High Control Case: Annual vs. Cumulative Excess of Claims Paid Less Premium over 30 Years (VaR = 95%)**

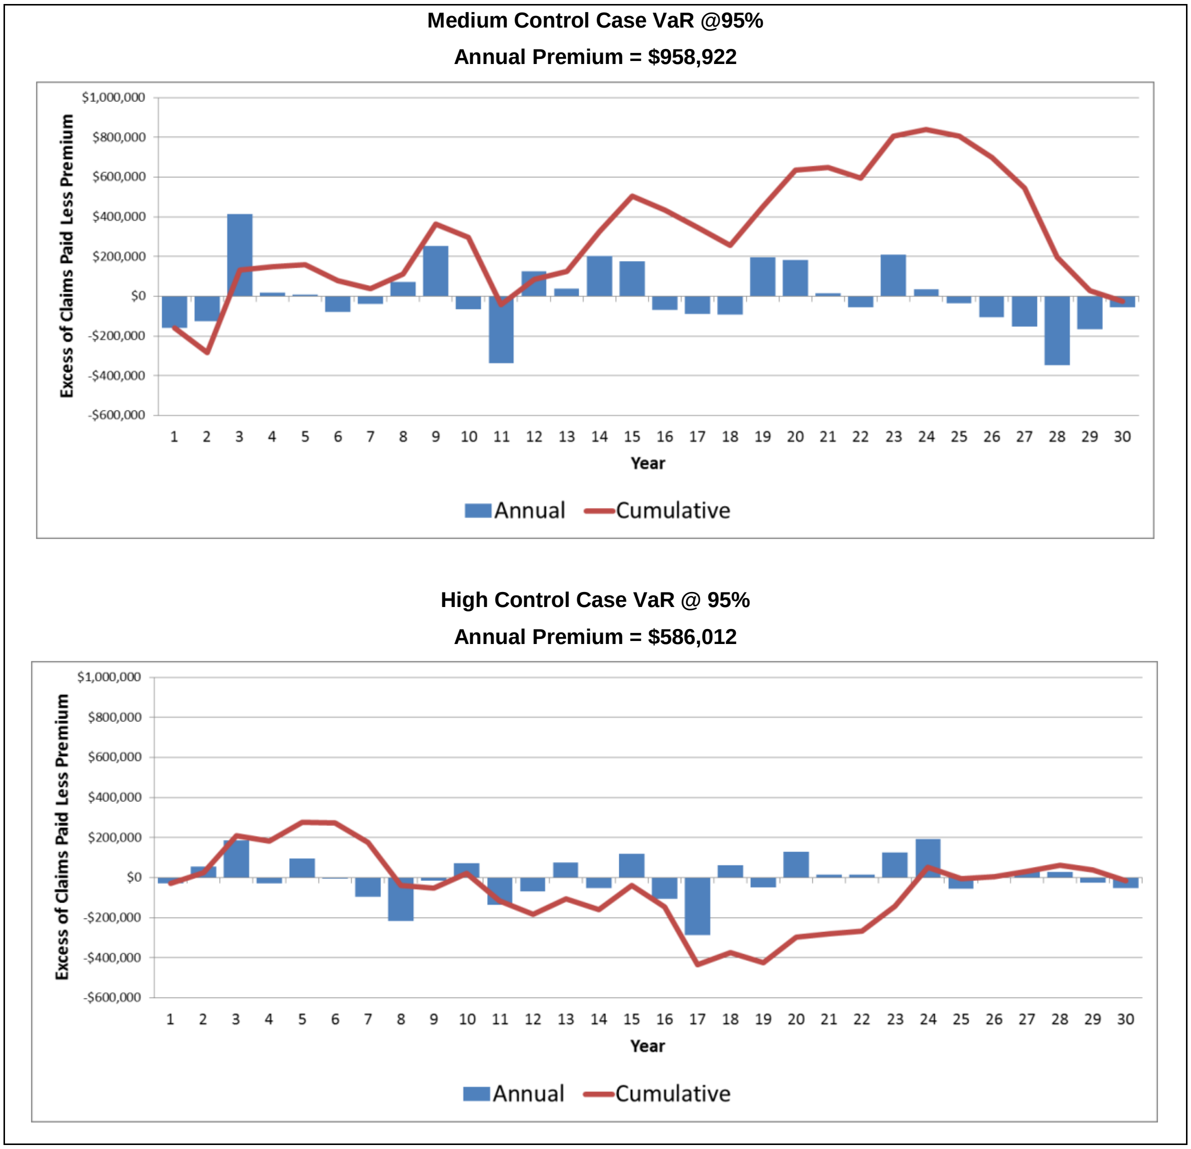

Medium Control Case VaR @95%

Annual Premium = $958,922

High Control Case VaR @ 95%

Annual Premium = $586,012

For the decision maker, a more salient difference arises if, instead of selecting a VaR confidence level in terms of a percentage, she selects a minimum absolute dollar loss level (which may be more realistic for management to consider, particularly in terms of catastrophic loss levels) and then compares the required insurance to accommodate the same level of loss under Medium vs. High controls. For example, if we take the Medium Controls case as representing a ‘base case’, at a 95% VaR the loss limit equals $37,051. For management to insure losses over this dollar level under the High controls case, the premiums are one-fifth of that in the Medium controls case and the volatility of net losses is correspondingly lower (as it would always be under High vs. Medium controls):

**Figure 67 - Medium vs. High Control Case: Annual vs. Cumulative Excess of Claims Paid Less Premium over 30 Years (VaR = 95%)**

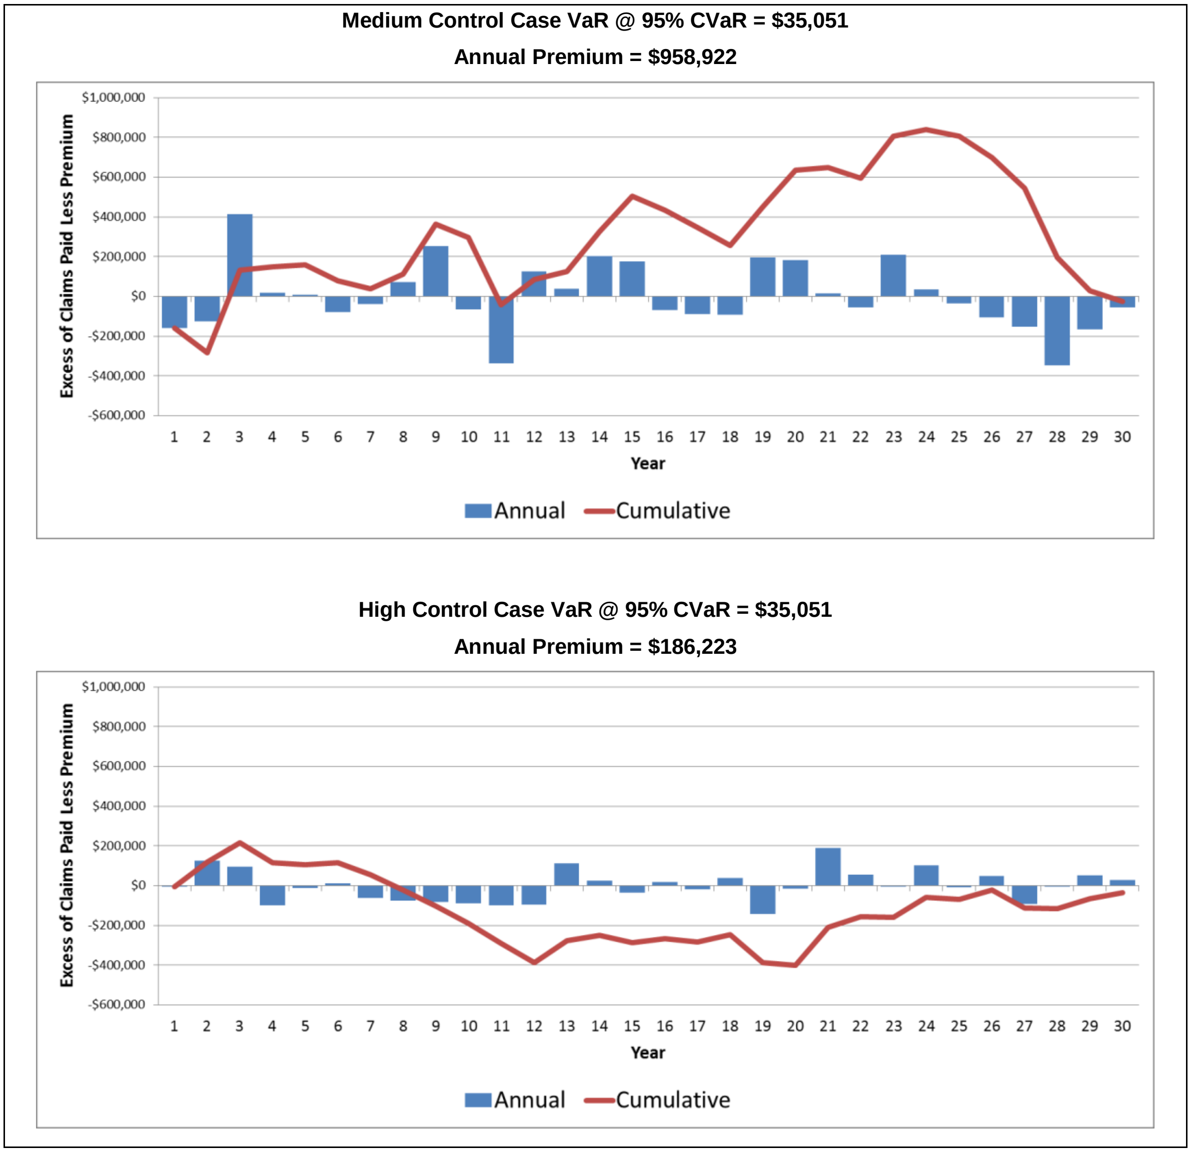

Medium Control Case VaR @ 95% CVaR = $35,051

Annual Premium = $958,922

High Control Case VaR @ 95% CVaR = $35,051

Annual Premium = $186,223

This type of comparison should prompt management to consider two further decisions: 1) that insurance cannot be considered outside of the context of the preventive controls that avoid losses in the first place, and 2) the ‘optimal’ level of control, whether frequency or impact reducing depends, among other ‘biases’, on the decision makers’ preference for avoiding risk and, in the specific context of loss, the degree of ‘regret’ involved (Loomes and Sugden 1982; Gollier 2016).

The use of probability distributions to represent stochastic system outcomes and the selection of risk indicators based on higher moments of the distribution are therefore important tools for the representation of security risk information, although it is not entirely clear to what extent managers are prepared to interpret and act on this type of risk information about stochastic system processes (Lampel, Shamsie et al. 2009). Dominant organizational discourse on risk may also hamper efforts to portray and evaluate risks in a manner that facilitates good decision making (Hardy and Maguire 2016). From a quantitative perspective, the lack of operational information regarding the frequency of security losses means that managers will likely resort to subjective assumptions about the frequency of security events and therefore the underlying probability distribution of all possible events which will be reflected in their resulting control choices (Weston 2014). This is problematic since the assumed shape of the probability distribution (if any is assumed) matters when considering what is a 'rare' event[50]. Even if representative event information is available, organizational context may inhibit the collection of representative samples (Feiler, Tong et al. 2013). With potentially minimal and unrepresentative data on hand, there is a high likelihood that the events are a non-representative sample of the underlying probability distribution, and do not particularly represent the likelihood of low probability events (Lopes 1982).

### Review of VaR-based Security Risk Modelling Approaches

Having reviewed the nature of the probability distributional qualities of stochastic systems which are subject to extreme values, I now return to the specific issue of modelling and simulation of information system security processes and outcomes that are expected to exhibit these characteristics. Limitations of the reviewed models to this point included the primary difficulty of incorporating stochastic inputs within proprietary modelling platforms where alternative approaches using standard desktop tools with Monte Carlo plug ins might support a shorter learning curve and greater flexibility in creating model alternatives and scenarios appropriate for the labs. It had also become clear that a standardized presentation of simulated results would need to include a time series view and a corresponding probability distribution view of the attributed loss results, particularly if the higher moments of the distribution were of interest, and if specific risk measures based on VaR and CVaR statistics were to be incorporated into the experiments. Since the intention was to model a system at risk without basing these on any particular system or actual historical results, I was not constrained by lack of data going into the modelling exercise however any simulated results would still have to represent plausible outcomes with distributional characteristics that would be readily recognizable by practitioners. This required that the stochastic elements of the model be based on plausible expert assumptions for their range of effect. A further literature review was initiated to review security modelling efforts that directly incorporated both Monte Carlo simulation and VaR/CVaR components.

Alternative models for information security risk simulation have been presented relatively recently in the information security literature using VaR and CVaR concepts to more adequately capture the risk profile

[50] Weston uses the example of financial market losses which exhibit 'fat tail' distributions and indicates that the "…individual's choice behaviour reflects a distributional assumption whether or not she makes the assumption consciously and/or with explicit knowledge of the statistical properties of the distribution. For example, if single-day financial-market drops of 10% were normally distributed, they would be considered rare since they would be expected to occur once every 500 years (Buchanan, 2004); in actuality, they are merely unusual since they occur two orders of magnitude more frequently, about once every five years (Mandelbrot & Hudson, 2004)." (Weston 2014)

risk in business information systems subject to security risks (Soo Hoo 2000; Jaisingh and Rees 2001; Conrad 2005; Ozcelik and Rees 2005; Conrad, Oman et al. 2006; Wang, Chaudhury et al. 2008; Thomas 2009; Gheorghe 2012; Sawik 2013; Thomas, Antkiewicz et al. 2013; WEF 2015). These quantitative approaches to risk modelling borrow from both enterprise finance and operations risk management[51], incorporates Monte Carlo simulation for managing stochastic model inputs and typically expresses risk in terms of a probability distribution of loss outcomes. Soo Hoo proposed a quantitative approach based on simulation of annual loss expectancy (ALE) and reviewed the limitations of ALE in its inability (by definition) to distinguish between high/low and low/high probability/outcome risks. To resolve this limitation, he advocated incorporating both sensitivity analysis, which dynamically varied the input parameters of a logical risk model to determine control policy ‘crossover points’ and then stochastic analysis to determine ‘stochastically dominant’ decision alternatives which characterizes the solutions across the cumulative distribution function of an alternative (Soo Hoo 2000). Levy illustrated the use of stochastic dominance in a number of decision contexts including experimental tests for prospect theory (Levy and Levy 2002) and the ‘safety first’ principle for portfolio selection noted above (Levy and Levy 2009). Jaisingh (Jaisingh and Rees 2001) and Ozcelik (Ozcelik and Rees 2005) provide conceptual descriptions of the calculation and use of VaR in an information security context but do not discuss the non-coherency limitations of using VaR as a risk measure. Conrad introduced an explicit Monte Carlo method for simulating security risks (Conrad 2005) into a static ‘efficacy model’ for security risk assessment following Longstaff (Longstaff, Chittister et al. 2000). Conrad also applied a Monte Carlo simulation approach to the Risk Analysis and Probabilistic Survivability Assessment (RAPSA) process (Taylor, Krings et al. 2002) which combines the quantification approach of Probability Risk Assessment (PRA) with Survivable Systems Analysis (SSA) (Ellison, Fisher et al. 1997) in order to focus control decisions and resources on the mission critical services of a system (Conrad, Oman et al. 2006).Thomas presents a conceptual model for stochastic, ‘piecewise’ modelling of the “total cost of security” where extreme events need to be modelled separately from everyday protective controls to focus on VaR-based insurance approaches to absorbing catastrophic risks (Thomas 2009). He also reviews a ‘dynamic branching model’ which incorporates temporal, non-static event trees to simulate security breach costs (Thomas, Antkiewicz et al. 2013). The World Economic Forum also extensively promoted the quantification of security risk using Monte Carlo methods and VaR for in their 2015 annual report on cyber risks (WEF 2015). It appears that many of these models could be executed using Excel with a Monte Carlo plug in and so I pursued further detailed study of these methods to be able to contrast the effort required and results expected from these approaches compared to the models noted in the previous section.

### Simulation using Generalized Probability Distributions with Available Historical Data

Stochastic modelling and simulation approaches are particularly valuable when historical operating data on which to base a model are available but may be limited or poorly understood by management which may

[51] See for example “A New Approach for Managing Operational Risk - Addressing the Issues Underlying the 2008 Global Financial Crisis, Society of Actuaries 2009 http://www.soa.org/Files/Research/Projects/research-new-approach.pdf

often be the case in practice, particularly for organizations whose security or IT risk management programs or expertise are immature (Hardy and Maguire 2016) and where tail observations may be characteristically sparse in any case. In these situations, care must be taken to ensure that models, even if based on actual operating data, and the resulting simulation results are robust and reliable. Wang describes a loss modelling methodology based on an 'exceedance over thresholds extreme value model' (Pickands III 1975; Davison and Smith 1990) in which attributed security losses are stochastically simulated based on actual incident statistics for one year of operations. The objective is to select an extreme loss threshold using maximum likelihood methods to estimate a parameterized Generalized Pareto Distribution (GPD) model for losses from which a VaR level can reliably be calculated based on sparse operating data (Coles, Bawa et al. 2001), and then to employ the resulting model to estimate alternative security policy costs (Wang, Chaudhury et al. 2008). Coles method is instructive since it allows management to statistically determine a relevant threshold loss level from existing operating data as opposed to selecting some arbitrary VaR value above which to model extreme values. Wang's model first estimates daily losses attributed to stochastic incidents resulting from business activities:

**Figure 68 - Actual vs. Simulated daily losses due to activity based security events (Wang)**

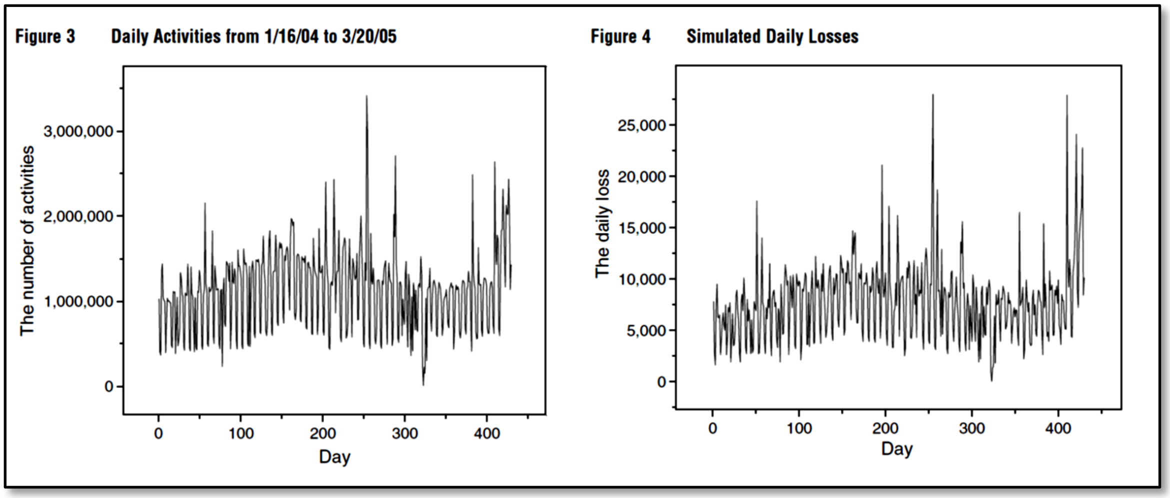

(Wang, Chaudhury et al. 2008)

Since the attributed losses are largely a function of business activity levels, the overall time series of losses expectedly aligns closely to the pattern of business activities. In this case the model assumes 15,000 users generating approximately 1.13 million (user level) activities per day. Wang then identifies two types of activity-based security incidents based on interviews with management, reflecting the high/low low/high dichotomy:

- **Type A: High-Frequency/Low-Impact incidents** expected once every 15,000 activities, with approximately 75 such incidents per day. Attributed losses resulting from incidents of this type include

> the cost of help-desk calls and loss of individual productivity and [are] estimated to be $100 per incident on average.
>
> - **Type B: Low-Frequency/High-Impact incidents**, including successful incidents/exploits that cause disruptions of network services expected once every 50,000,000 activities; on average, there will be one or two such incidents every two months where the cost is approximately $10,000 per incident, which includes the loss of user productivity, business disruptions, and possible damage to business reputation.
>
> (Wang, Chaudhury et al. 2008)

Wang then proceeds to model loss values occurring above a loss threshold using a generalized Pareto distribution (Coles, Bawa et al. 2001)[52]. The residuals from the modelled vs. the simulated losses are then compared at different threshold levels to determine the minimum threshold level above which there are enough observations to ensure reliability of the resulting GPD model of extreme values: the threshold should be high enough to justify the assumptions of the GPD model, but low enough to capture a reasonable number of observations at the tail end of the distribution (Dupuis 1999). Wang's resulting model selects $8,500 as the threshold value (where the average loss is $7,700) leaving 178 out of 430 loss observations on which to model tail losses. The resulting model is indicated in Figure 70 below. The model was checked for sensitivity against two threshold levels and using two 'de-clustering' factors (*r*) to ensure temporal independence of the observations:

[52] *Consider the distribution of X conditionally on exceeding some high threshold u, and let $Y = X - u$, and $Y > 0$. We know*

$$F_u(y) = Pr\{Y \leq y \mid Y>0\}$$

*As $u \rightarrow \omega_F$ sup{x: F(x) <1}, we found a limit distribution,*

$$F_u(y) \approx G(y; \sigma, \xi)$$

*where G is Generalized Pareto Distribution (GPD)*

$$G(y; \sigma, \xi) - 1 - \left(1 + \xi \frac{y}{\sigma}\right)^{-1/\xi}$$

*defined on $\{y: y > 0 \text{ and } (1+ \xi y/ \sigma )>0\}$, where $\sigma$ and $\xi$ are the two parameters of the distribution.* (Coles, Bawa et al. 2001)

**Figure 69 – Mean Residual Life Plot for Daily Losses and Generalized Pareto Distribution Parameter Estimate (Wang)**

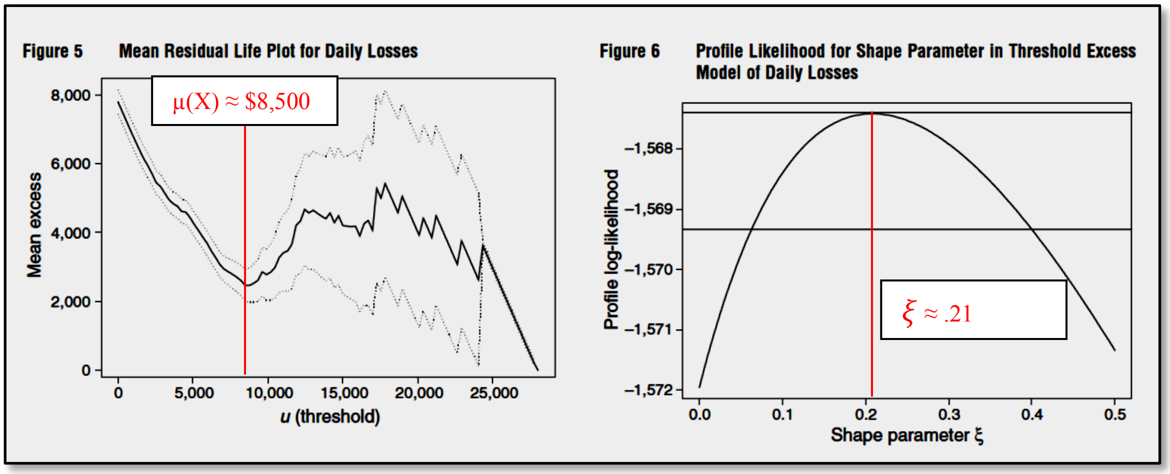

(Wang, Chaudhury et al. 2008)

**Figure 70 – GPD Model Estimation Results (Wang)**

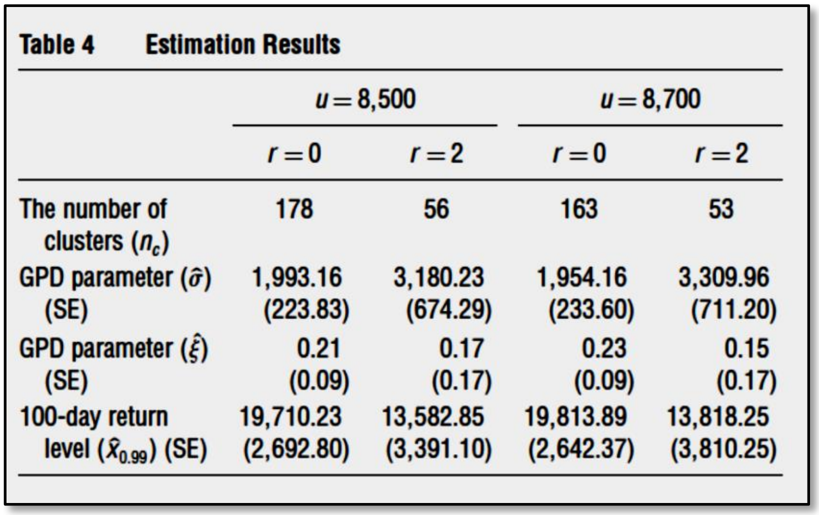

**Table 4 Estimation Results**

| | $u = 8,500$ | | $u = 8,700$ | |
|---|---|---|---|---|
| | $r = 0$ | $r = 2$ | $r = 0$ | $r = 2$ |
| The number of clusters ($n_c$) | 178 | 56 | 163 | 53 |
| GPD parameter ($\hat{\sigma}$) (SE) | 1,993.16<br>(223.83) | 3,180.23<br>(674.29) | 1,954.16<br>(233.60) | 3,309.96<br>(711.20) |
| GPD parameter ($\hat{\xi}$) (SE) | 0.21<br>(0.09) | 0.17<br>(0.17) | 0.23<br>(0.09) | 0.15<br>(0.17) |
| 100-day return level ($\hat{x}_{0.99}$) (SE) | 19,710.23<br>(2,692.80) | 13,582.85<br>(3,391.10) | 19,813.89<br>(2,642.37) | 13,818.25<br>(3,810.25) |

(Wang, Chaudhury et al. 2008)

**Figure 71 – GPD Model Diagnostic Plots (Wang)**

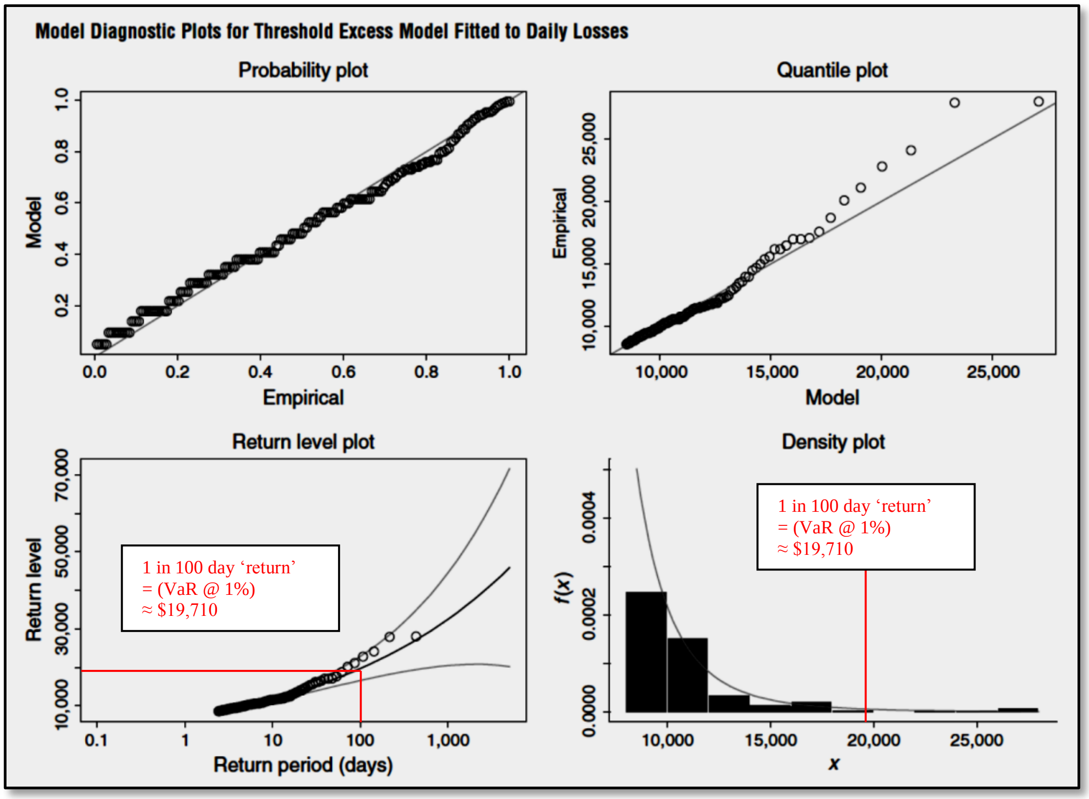

1 in 100 day 'return'
= (VaR @ 1%)
≈ $19,710

1 in 100 day 'return'
= (VaR @ 1%)
≈ $19,710

(Wang, Chaudhury et al. 2008)

Wang then compares two control policy decisions to the modelled base case. In addition, he tests the model's sensitivity to inputs using random variables for both probability and loss generated by probability distributions as opposed to using fixed values, although he does not explain how this was done:

**Control Solution A:** Install a firewall. This solution will decrease the probability of both types of incidents by half at a daily cost of $550.

**Control Solution B:** Increase the frequency of backup and introduce detection systems. This solution will decrease losses from both incident types by half and reduce the frequency of type A intrusions to 1 per 20,000 activities [from 1 per 15,000]. It has a daily cost of $2,000.

**Figure 72 – Sensitivity Analysis Parameters and Scenario Results (Wang)**

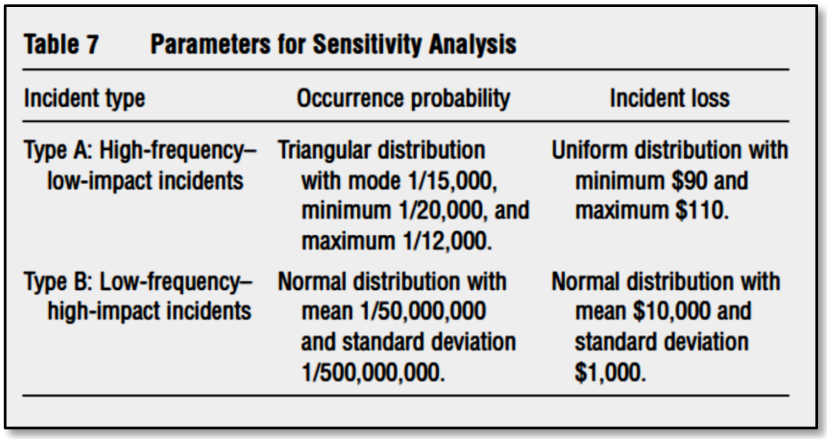

**Table 7 Parameters for Sensitivity Analysis**

| Incident type | Occurrence probability | Incident loss |
|---|---|---|
| Type A: High-frequency–low-impact incidents | Triangular distribution with mode 1/15,000, minimum 1/20,000, and maximum 1/12,000. | Uniform distribution with minimum $90 and maximum $110. |
| Type B: Low-frequency–high-impact incidents | Normal distribution with mean 1/50,000,000 and standard deviation 1/500,000,000. | Normal distribution with mean $10,000 and standard deviation $1,000. |

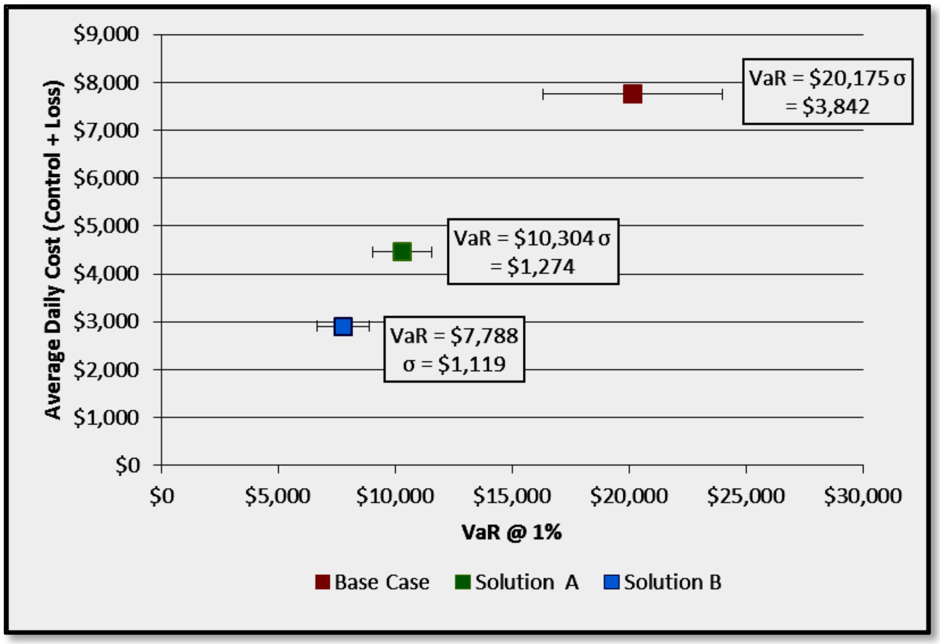

(Wang, Chaudhury et al. 2008)

The advantage of this approach is that it enables modelling of the distribution of losses above a threshold where data may typically be sparse. Here, however, Wang focuses on the resulting VaR for losses and not the expected value of losses above VaR (i.e. CVaR) which, as noted above, may substantially exceed VaR depending on the underlying loss dependencies. As noted in the section above, management decisions based solely on the reduction in VaR vs. the reduction in average losses above VaR may not reflect the correct risk measure where the average loss above VaR represents a better measure of extreme losses (although we might expect in this model that CVaR would likely also decline in each scenario). This model also introduces the use of probability distributions (specifically normal, triangular and uniform distributions based on expert judgment) and Monte Carlo simulation methods to calculate stochastic loss outcomes.

# 6 – System Simulation Platform Selection and Modification

Having therefore completed a review of three state of the art information security modelling approaches and associated software platforms, I made a final decision to utilize the proprietary enterprise architectural modelling approach of the CySeMol tool to model a system at risk which ensures the objective logical interconnectivity of the proposed 'system' components (attack types, assets, vulnerabilities and controls) and which already incorporates independent expert validation of the model components undertaken by Ekstedt et al, but to abstract the predefined system component classes and probability relationship factors underlying the Bayesian relationships between the components into a form that would permit their remodelling and simulation  directly within Excel. Use of Excel thereafter permits me to then incorporate both Monte Carlo model simulation and design the necessary data display and control selection tools required to generate various simulation scenarios. I also decided to use Excel to manage the 'front end' user interface for the resulting choice games which would be directly manipulated by the individual participants in my lab experiments in order to ensure maximum flexibility in designing the interface and to readily manage the necessary data capture and reporting within an already personally familiar platform. On reflection, this approach balances several theoretical objectives for model validity and the pragmatic issues of prospectively designing and running a lab simulation to generate control choice data. In the case that simulations needed to be run interactively (an aspect that was not certain at the time of model selection) being Excel based, an Excel based system would also be able to run a complex Monte Carlo simulation quickly on a typical late model laptop using an academically available (i.e. cost effective) software add in[53] facilitating both lab portability and the anticipated need for individuals to run separate replicated simulations on individual computers within the lab at relatively low cost.

### Logical Model Specification Using P2CySeMol and Excel

I began with a P2CySeMol model based on Holm et al which was readily available and well documented by the authors (Holm, Shahzad et al. 2015). While their model inherently involved SCADA control systems, the model included an extensive "Office Zone" administrative management sub-model which was suitable as the basis for a corporate administrative IT 'system at risk' for my purposes. I simplified and altered their base model to reflect the core components of an enterprise clinical management system that might be found at a medium to large hospital, offering both clinical and administrative users web-based, on premises and remote access to the system. Use of the P2CySeMol platform ensured that the logical architectural connections between the instantiated system components remained appropriately connected in terms of both their system functionality and their logical security relationships. Simplification of the functional and security architecture of the overall model (but not the components themselves which were pre-specified in the CySeMol Class Model) were undertaken in order to ensure that the model size (i.e. number of components) remained small enough to practically debug but large enough to reflect an

[53] I will be using Frontline's Monte Carlo Risk Solver for Excel: http://www.solver.com/risk-solver-platform

acceptable minimum scope of functional and non-functional (security) architectural components necessary to represent a plausible operational scenario.

As noted above, the anticipated computational requirements for simulation of a conditional probability model, as a Bayesian network in CySeMol but to be recreated in Excel, roughly scales exponentially with the number of nodes, and at the time it was not known how this would translate once instantiated in Excel. The approach was to convert only the number of components needed to produce plausible operational loss data that would reflect changes to the control and operational inputs of the model in a way that reflected practitioner experience. Several iterations of the model which incrementally added the vulnerability and control features of the components as noted in the next section were attempted in order to gauge the feasibility of the full target model. Incremental testing indicated that the model simulation time increased substantially as additional nodes and their corresponding vulnerability and control elements, each represented by probability distributions to support Monte Carlo simulation, were added to the model. While supporting the stochastic simulation aspect of the model, these elements added significantly to the computation time for a 365 day simulation, rising from nearly instantaneous calculation to over one minute with the addition of several core stochastic components. Additionally, and significantly as a contribution of this work, the introduction of lag behaviours to the controls meant that each day could not be simulated independently, but rather had to be generated as a time series, where one day was dependent on the outcome of the previous day. Together with the addition of substantial stochastic attributes for breach probability and impact determination, it became clear that the model might not be practically simulated 'in real time' during the labs and that, if simulations were to be computed during the labs, substantial time would have to be allotted for users to compute the model individually. This problem was effectively eliminated later in development during the creation of the user interfaces by precomputing the required simulations which, at that point, needed to represent only a limited number of control variations to support the required lab experiments. Precomputation of the simulations also effectively allowed the creation of precise business loss probability distributions representing 30 years of data in each required simulation scenario which could be randomly selected by the lab interface to 'simulate' business losses over any selected time period within the 30 years instantaneously without the need for interactive Monte Carol computation. This allowed users to experience stochastic system outputs without the interruption of actually computing the scenario results in real time without loss of model fidelity. Although not an 'immersive' simulation environment as utilized by Fiore et al, this approach to the generation of simulated results supported the need to focus the users' attention on the decision making elements of the model rather than the computational aspects of the simulation model (Fiore, Harrison et al. 2009).

Figure 73 – Instantiated P2CySeMol Model (Curtis)

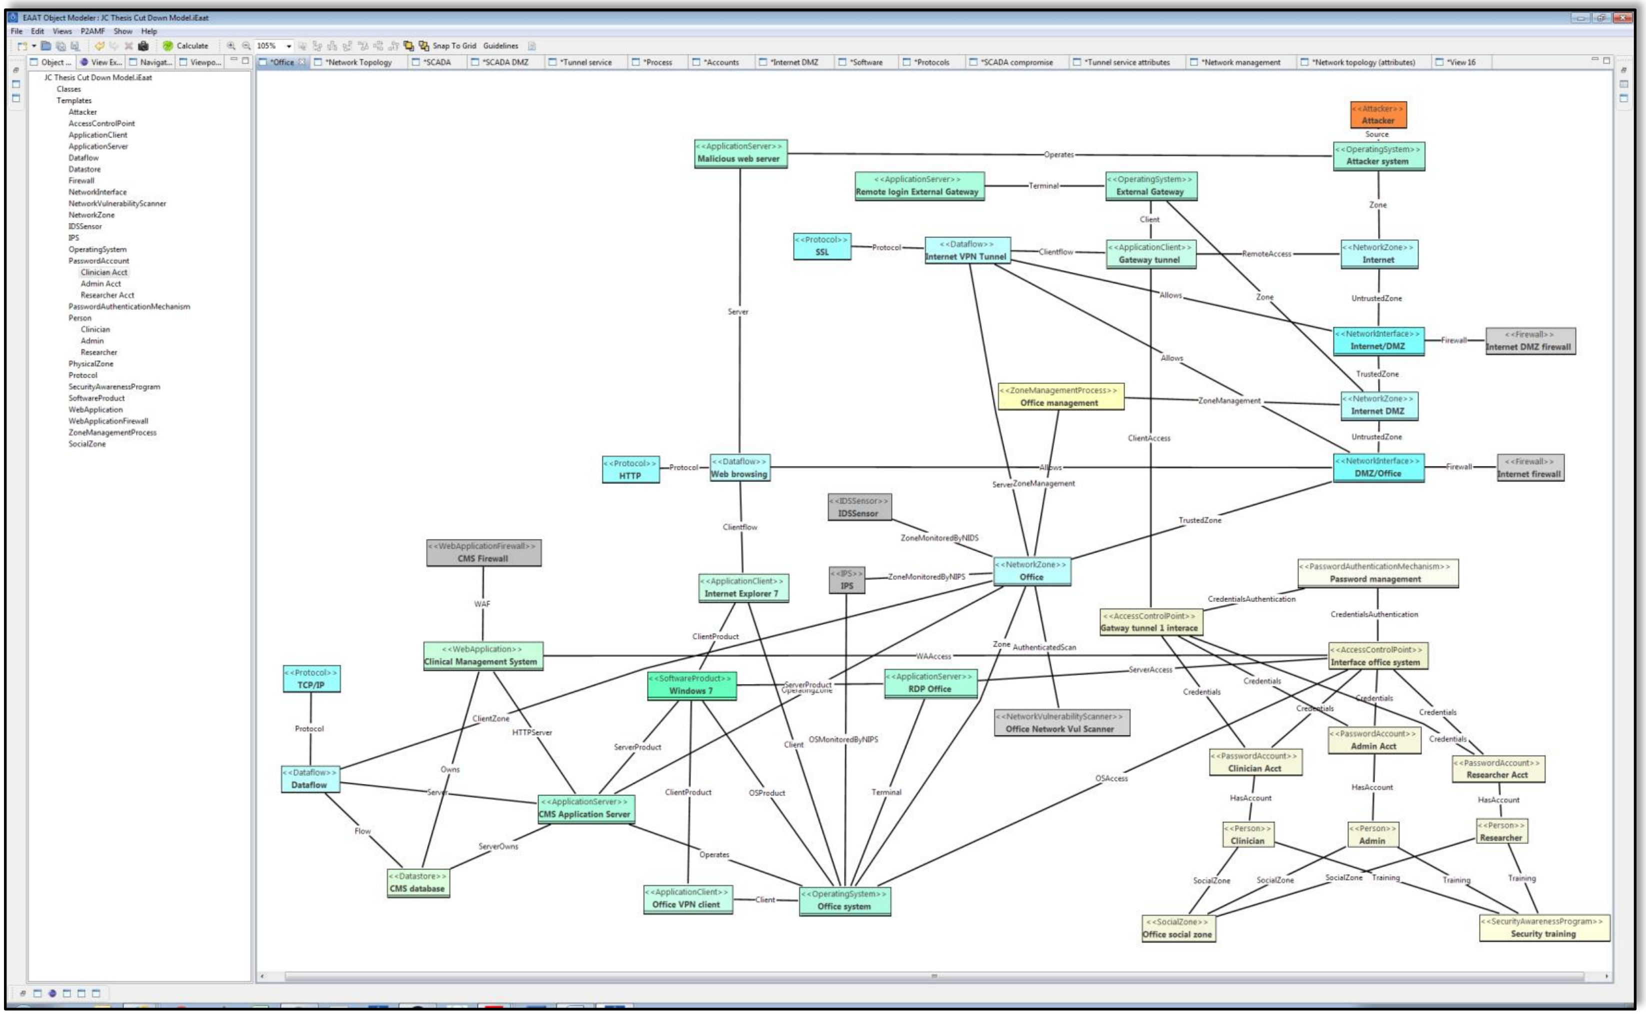
EAAT Object Modeler : JC Thesis Cut Down Model.iEaat
File Edit Views P2AMF Show Help
Calculate
105%
Snap To Grid Guidelines
Object ... View Ex... Navigat... Viewpo...
*Office *Network Topology *SCADA *SCADA DMZ *Tunnel service *Process *Accounts *Internet DMZ *Software *Protocols *SCADA compromise *Tunnel service attributes *Network management *Network topology (attributes) *View 16
JC Thesis Cut Down Model.iEaat
Classes
Templates
Attacker
AccessControlPoint
ApplicationClient
ApplicationServer
Dataflow
Datastore
Firewall
NetworkInterface
NetworkVulnerabilityScanner
NetworkZone
IDSSensor
IPS
OperatingSystem
PasswordAccount
Clinician Acct
Admin Acct
Researcher Acct
PasswordAuthenticationMechanism
Person
Clinician
Admin
Researcher
PhysicalZone
Protocol
SecurityAwarenessProgram
SoftwareProduct
WebApplication
WebApplicationFirewall
ZoneManagementProcess
SocialZone
<<Attacker>> Attacker
Source
<<OperatingSystem>> Attacker system
<<ApplicationServer>> Malicious web server
Operates
<<ApplicationServer>> Remote login External Gateway
Terminal
<<OperatingSystem>> External Gateway
Client
Zone
<<Protocol>> SSL
Protocol
<<Dataflow>> Internet VPN Tunnel
Clientflow
<<ApplicationClient>> Gateway tunnel
RemoteAccess
<<NetworkZone>> Internet
Allows
Zone
UntrustedZone
Server
<<NetworkInterface>> Internet/DMZ
Firewall
<<Firewall>> Internet DMZ firewall
Allows
TrustedZone
<<ZoneManagementProcess>> Office management
ZoneManagement
<<NetworkZone>> Internet DMZ
ClientAccess
UntrustedZone
<<Protocol>> HTTP
Protocol
<<Dataflow>> Web browsing
Allows
<<NetworkInterface>> DMZ/Office
Firewall
<<Firewall>> Internet firewall
ServerZoneManagement
<<IDSSensor>> IDSSensor
Clientflow
TrustedZone
ZoneMonitoredByNIDS
<<WebApplicationFirewall>> CMS Firewall
<<NetworkZone>> Office
<<PasswordAuthenticationMechanism>> Password management
<<ApplicationClient>> Internet Explorer 7
<<IPS>> IPS
ZoneMonitoredByNIPS
CredentialsAuthentication
WAF
<<AccessControlPoint>> Gatway tunnel 1 interace
CredentialsAuthentication
ClientProduct
Zone
AuthenticatedScan
<<WebApplication>> Clinical Management System
WAAccess
<<AccessControlPoint>> Interface office system
<<Protocol>> TCP/IP
<<SoftwareProduct>> Windows 7
ServerProduct
OperatingZone
<<ApplicationServer>> RDP Office
ServerAccess
Credentials
Credentials
Credentials
Credentials
Credentials
Credentials
ClientZone
<<NetworkVulnerabilityScanner>> Office Network Vul Scanner
Protocol
HTTPServer
<<PasswordAccount>> Admin Acct
Owns
ServerProduct
Client
OSMonitoredByNIPS
OSAccess
<<PasswordAccount>> Clinician Acct
<<PasswordAccount>> Researcher Acct
<<Dataflow>> Dataflow
Server
ClientProduct
OSProduct
Terminal
HasAccount
HasAccount
HasAccount
<<ApplicationServer>> CMS Application Server
Flow
ServerOwns
<<Person>> Clinician
<<Person>> Admin
<<Person>> Researcher
Operates
<<Datastore>> CMS database
SocialZone
SocialZone
SocialZone
Training
Training
Training
<<ApplicationClient>> Office VPN client
Client
<<OperatingSystem>> Office system
<<SocialZone>> Office social zone
<<SecurityAwarenessProgram>> Security training

Figure 73 indicates my resulting logical model executed in P2CySeMol. Individual system components are based on CySeMol class model objects which have defined attack vectors accomplished through the documented component security vulnerabilities: “Has Vulnerability”, “Find Vulnerability” and “Exploit Vulnerability”, sub-elements of which can be either ‘True’ or ‘False’ based on ‘evidence’ of their existence, either pre-specified or conditionally calculated within the model. Components also have defense capabilities – security controls - which are also based on evidence either incumbent, calculated within or set by the user using the user interface. The user specified controls are grouped into one of four types according to the conceptual framework described above by Tjoa et al (Tjoa, Jakoubi et al. 2011): Preventive controls (including Blocking controls), Detective controls, Counter controls and Recovery controls. Each type of control can be set to one of 5 levels of effectiveness: None, Minimal, Low, Medium, High or Very High. Higher levels of control generally have higher average and lower variance of ‘effectiveness’ which results in a higher probability score for the effectiveness of the corresponding component control element. Each class of control contains one or more administrative (people or process) or technical control elements which have been assigned to each component’s ‘Control’ elements where it would be reasonably valid to assume that that the level of control would likely be actively set by management. The ability to vary the nominal control levels within and across the control types provides the ability to generate system simulations of daily business losses reflecting different levels of desired security control posture which are used to characterize control choice scenarios for the lab experiments involving different security risk and uncertainty levels.

In the following example, I will first illustrate the incorporation of a typical system component, the ‘Web Application’ component of the system architecture (the web-based clinical management system software application itself), in order to illustrate the logical conversion of the CySeMol components to calculable Excel objects. I then explain the nature of the Preventive and Blocking controls applicable to this component and to other components generally and the associated determination of ‘successful’ attacks that have avoided prevention and are able to affect the confidentiality, integrity and availability attributes of the IT system. I then illustrate how successful attacks are possibly Detected in order to initiate both Counter and Recovery control of the successful attacks to eliminate the effects of the attacks on the system. The section ends with a discussion of the calculation and presentation of business loss impacts attributed to successful attacks despite those attacks having been countered and recovered.

### Determination of Successful Attacks

In the CySeMol class model a web-based application is distinct from its database or other supporting components like its associated Web Application Firewall, etc. This permits realistic granularity of the associated vulnerabilities and controls permitting each component to have its own corresponding attack and defense vectors:

**Figure 74 - CySeMol Class Object "Web Application" Component Elements including simulated element states (T/F) (Curtis)**

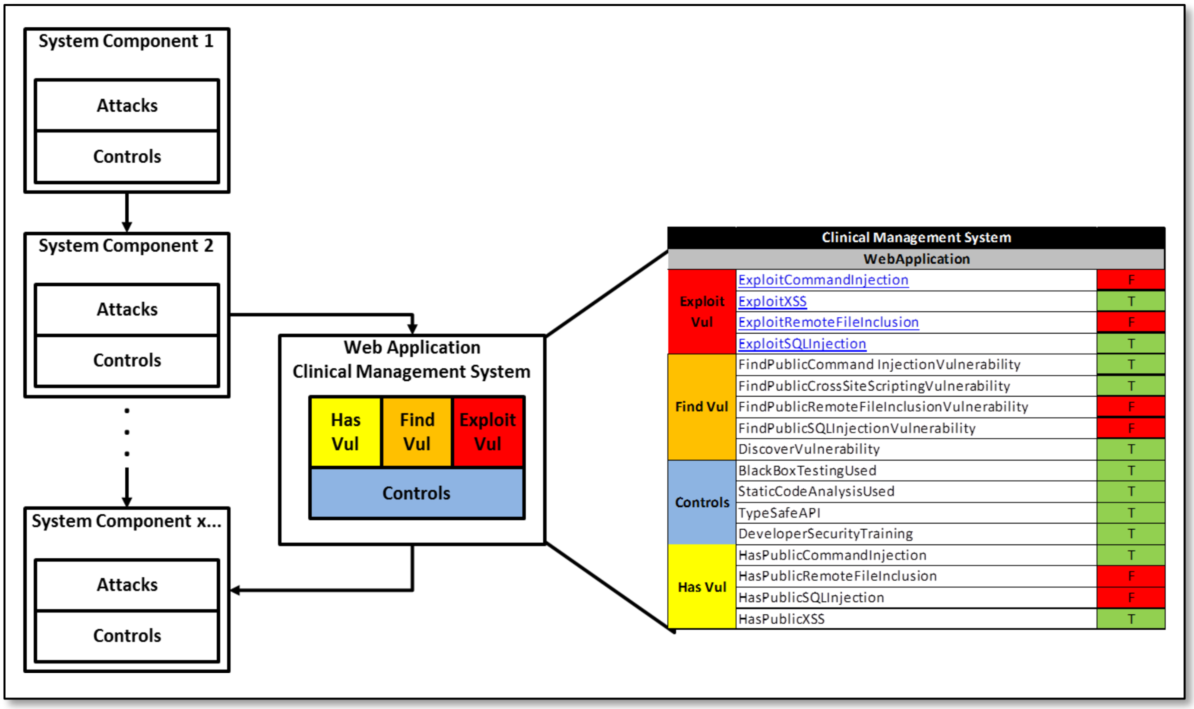

| Clinical Management System | | |
|---|---|---|
| WebApplication | | |
| Exploit Vul | ExploitCommandInjection | F |
| | ExploitXSS | T |
| | ExploitRemoteFileInclusion | F |
| | ExploitSQLInjection | T |
| Find Vul | FindPublicCommand InjectionVulnerability | T |
| | FindPublicCrossSiteScriptingVulnerability | T |
| | FindPublicRemoteFileInclusionVulnerability | F |
| | FindPublicSQLInjectionVulnerability | F |
| | DiscoverVulnerability | T |
| Controls | BlackBoxTestingUsed | T |
| | StaticCodeAnalysisUsed | T |
| | TypeSafeAPI | T |
| | DeveloperSecurityTraining | T |
| Has Vul | HasPublicCommandInjection | T |
| | HasPublicRemoteFileInclusion | F |
| | HasPublicSQLInjection | F |
| | HasPublicXSS | T |

In the above example, the clinical management system Web Application component has four potential security exploits: "Command Injection", "Cross-Site Scripting (XSS)", "Remote File Inclusion", and "SQL Injection". The state of the component at this point in a simulation is represented by the right hand column, indicating whether the component element is present (T = True) or absent (F = False), either by incumbency, based on controls set by the user or as a result calculated by the model. As would be expected, the presence of Vulnerabilities (T=True) generally increases the probability of successful attack while the presence of Controls (T = True) generally lowers the probability of successful attack. The applicable vulnerabilities are commonly understood web application vulnerabilities based on expert practitioner documented profiles contained within a publicly managed vulnerability database[54]. While the vulnerabilities are not exhaustive as instantiated within the CySeMol object class model, they have been extensively validated by Holm et al using expert panels and represent sufficiently valid profiles of vulnerabilities and attack vectors for my modelling purposes (Holm, Ekstedt et al. 2012; Holm, Ekstedt et al. 2013; Sommestad, Ekstedt et al. 2013). 'Successful' (T = True) attacks on individual exploitable component elements ("Exploit Vul") which remain either undetected or not Countered in the current or

[54] CVE stands for "Common Vulnerabilities and Exposures". The CVE database "…is a list of information security vulnerabilities and exposures that aims to provide common names for publicly known cyber security issues. The goal of CVE is to make it easier to share data across separate vulnerability capabilities (tools, repositories, and services) with this "common enumeration."" See http://cve.mitre.org/

subsequent rounds of the simulation will result in some level of attributed security incidents affecting the confidentiality, integrity or availability of system components and business processes resulting in associated business losses.

The state of each component’s Exploit Vulnerability elements (True or False) are based on the individual element’s own corresponding conditional probability table specified in the CySeMol class model. In the following example table, we detail the Web Application component’s “Exploit Command Injection” attack which has 6 direct dependencies based on the state of two vulnerabilities within the Web Application itself (‘Discover Vulnerability’ and ‘Find Public Command Injection’) and four ‘Preventive’ controls based in the associated component ‘Web Application Firewall’ (‘Monitored by Operator’, ‘Tuned Using Black Box Tool’, ‘Tuned by Experienced Professional’, and ‘Tuned With Significant Manual Effort’). The state of ‘Discover Vulnerability’ for the Web Application is based on a conditional probability table and depends on the Controls set on the Web Application itself (TypeSafeAPI, DeveloperSecurityTraining, BlackBoxTesting, StaticCodeAnalysis). ‘Find Public Command Injection’ is based on whether the attacker has logical access to the Web Application (determined elsewhere in the model) and whether the element ‘HasPublicCommandInjection’ is True (i.e. if there is in fact no PublicCommandInjection to discover in the Web Application, then it cannot be discovered regardless of the other Control settings or the state of the Firewall, etc.):

**Figure 75 - CySeMol Class Object “Web Application” ‘ExploitCommandInjection’ Element Conditional Probability Table**

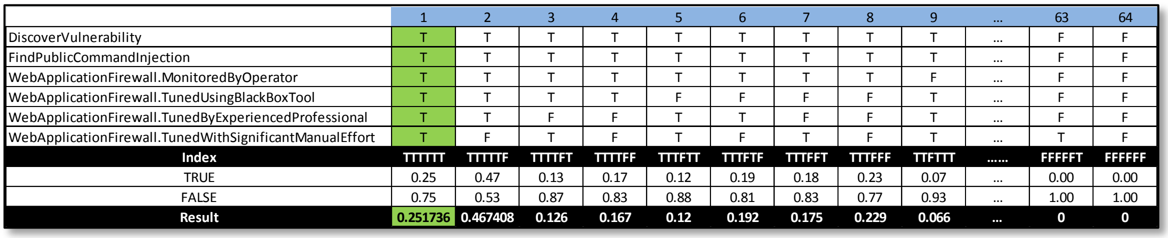

| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | … | 63 | 64 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| DiscoverVulnerability | T | T | T | T | T | T | T | T | T | … | F | F |
| FindPublicCommandInjection | T | T | T | T | T | T | T | T | T | … | F | F |
| WebApplicationFirewall.MonitoredByOperator | T | T | T | T | T | T | T | T | F | … | F | F |
| WebApplicationFirewall.TunedUsingBlackBoxTool | T | T | T | T | F | F | F | F | T | … | F | F |
| WebApplicationFirewall.TunedByExperiencedProfessional | T | T | F | F | T | T | F | F | T | … | F | F |
| WebApplicationFirewall.TunedWithSignificantManualEffort | T | F | T | F | T | F | T | F | T | … | T | F |
| **Index** | **TTTTTT** | **TTTTTF** | **TTTTFT** | **TTTTFF** | **TTTFTT** | **TTTFTF** | **TTTFFT** | **TTTFFF** | **TTFTTT** | **……** | **FFFFFT** | **FFFFFF** |
| TRUE | 0.25 | 0.47 | 0.13 | 0.17 | 0.12 | 0.19 | 0.18 | 0.23 | 0.07 | … | 0.00 | 0.00 |
| FALSE | 0.75 | 0.53 | 0.87 | 0.83 | 0.88 | 0.81 | 0.83 | 0.77 | 0.93 | … | 1.00 | 1.00 |
| **Result** | **0.251736** | **0.467408** | **0.126** | **0.167** | **0.12** | **0.192** | **0.175** | **0.229** | **0.066** | **…** | **0** | **0** |

(Holm, Ekstedt et al. 2013)

In the above example, in the case where each of the Vulnerability and Control conditions is determined to be ‘True’ (T), the resulting conditional probability of successfully exploiting the Web Application via “Command Injection” is approximately 25% as defined by CySeMol. This probability of success is then used as the input probability in a Bernoulli function for the component in order to determine whether the exploit is in fact successful (Bernoulli (.25) = 1 True, or 0 False). The Bernoulli function is calculated for each specific component exploit (and for each component’s exploits in the logical model) for each iteration of the simulation and may result in either success or failure on each iteration based on the condition of the input elements and the result of the Bernoulli trial:

**Figure 76 - CySeMol Class Object “Web Application” Component Elements: Simulated Element States based on Conditional Probability input, Bernoulli function calculation (Curtis)**

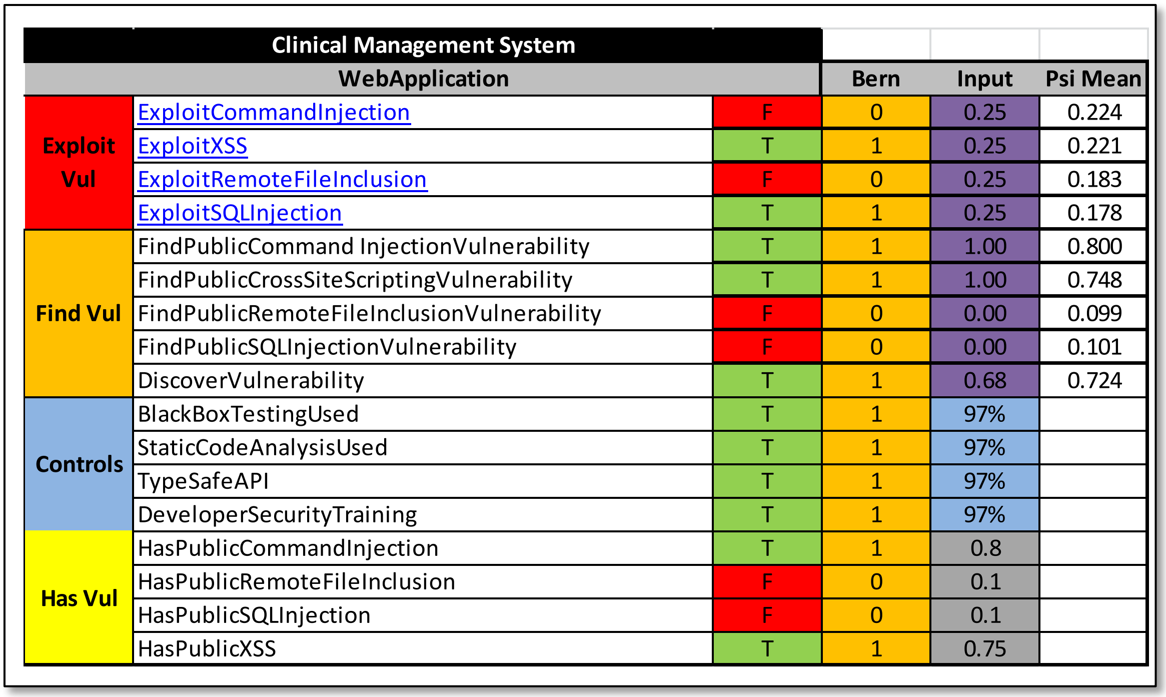

| Clinical Management System | | | | | |
|---|---|---|---|---|---|
| WebApplication | | | Bern | Input | Psi Mean |
| Exploit Vul | ExploitCommandInjection | F | 0 | 0.25 | 0.224 |
| | ExploitXSS | T | 1 | 0.25 | 0.221 |
| | ExploitRemoteFileInclusion | F | 0 | 0.25 | 0.183 |
| | ExploitSQLInjection | T | 1 | 0.25 | 0.178 |
| Find Vul | FindPublicCommand InjectionVulnerability | T | 1 | 1.00 | 0.800 |
| | FindPublicCrossSiteScriptingVulnerability | T | 1 | 1.00 | 0.748 |
| | FindPublicRemoteFileInclusionVulnerability | F | 0 | 0.00 | 0.099 |
| | FindPublicSQLInjectionVulnerability | F | 0 | 0.00 | 0.101 |
| | DiscoverVulnerability | T | 1 | 0.68 | 0.724 |
| Controls | BlackBoxTestingUsed | T | 1 | 97% | |
| | StaticCodeAnalysisUsed | T | 1 | 97% | |
| | TypeSafeAPI | T | 1 | 97% | |
| | DeveloperSecurityTraining | T | 1 | 97% | |
| Has Vul | HasPublicCommandInjection | T | 1 | 0.8 | |
| | HasPublicRemoteFileInclusion | F | 0 | 0.1 | |
| | HasPublicSQLInjection | F | 0 | 0.1 | |
| | HasPublicXSS | T | 1 | 0.75 | |

The probabilities for the “Has Vulnerability” elements indicated in grey are fixed according to the expert determination established by Holm et al within CySeMol, although their resulting True/False condition for each simulation round is still determined based on a Bernoulli function. The model is recursive to the extent that each of the input vulnerability and exploit elements may have its own conditional probability table with similar dependencies on other component elements in the logical model. In my resulting logical model there are 36 logical and physical architectural nodes (e.g. Web Application, Firewalls, Network Zones, Zone Management Policies, Password Management, etc.), 70 individual conditional probability tables associated with corresponding component exploits and vulnerabilities, and 64 individual Preventive or Blocking component controls of which 44 can be set by the user.

### Limitations of the Excel Specification of the CySeMol class model

It should be noted that while my Excel instantiation of the CySeMol class objects simplifies the overall architectural model (while retaining the logical architectural and security relationships of the class objects) it also introduces specific difficulties in terms of the recursive calculability of the conditional probability relationships since my approach is to approximate the logical relationships calculated as a Bayesian network without resorting to Bayesian estimation techniques. The state dependency of each element on a conditional probability table which is further composed of input elements that may also have their own conditional probability tables presents two practical problems for efficiently modelling the system: 1) the number of dependencies rises exponentially with the number of components and 2) the dependencies may become self-referential, requiring (ideally) some form of recursive estimation method in order to find the solution to the state of the system. This former problem is reasonably managed by making the model as small as possible which, in this case was achievable with minimizing the number of required components and their associated conditional probability tables that achieve a valid and plausible ‘system at risk’. The second problem is not directly addressable within the form of the linear model executed in Excel since the model is essentially a spreadsheet based on lookup tables and explicitly not a calculable Bayesian network. This translation into Excel presents problems for the Excel model which manifest in specific instances of circular references since Excel in this case cannot resolve certain individual component probabilities which are recursively dependent on their own (initial or calculated) state.

For example, The “Office Social Zone” component has one exploit ‘Share Portable Media’, meaning that there is a potential security risk to the system based on persons who are in the same office ‘social’ setting sharing portable media such as USB keys. While the sharing of portable media may, in and of itself, lead to a confidentiality impact (i.e. a user may have unintended access to confidential information on the USB key), the primary risk in the CySeMol class model derives from the fact that the media may contain malware or other security viruses, possibly intentionally implemented by an attacker with access to the Social Zone, which could lead to subsequent attack of other system components (Beautement, Coles et al. 2009). According to CySeMol: “This attack step is TRUE if a Person connected to the SocialZone has a PasswordAccount related to an ‘OperatingSystem’ that has been compromised by the attacker [i.e., if the Operating system component element ‘OperatingSystem.Access’ = TRUE there is a risk associated with subsequently sharing portable media, otherwise sharing portable media would have no effect]. It is FALSE in other cases.” The challenge for executing this relationship in Excel is that operating system access is itself dependent on whether portable media has already been shared, so that there is a directed but simultaneous relationship between the states of operating system access by the attacker and the sharing of portable media. This only manifests in the Excel model once the relationship is traced through the ensuing conditional probability table dependencies:

> Office Social Zone: ‘Share Portable Media’ (True) ➔ OS: ‘OS Access’ (True) ➔ OS: ‘Execute Malicious Code’ (True) ➔ OS: ‘Access via Portable Media’ (True) ➔ Office Social Zone: ‘Share Portable Media’ (True)
>
> where ➔ denotes ‘depends on whether’

While it is clear that the intended logical relationship in the Bayes network is that ‘sharing portable media potentially causes malicious OS Access’ and not the other way around, the expression of the risk of ‘sharing portable media’ must reflect both the ensuing OS compromise and the dependency of the sharing risk on the fact of OS access by other means. This type of relationship would normally be resolved within the Bayes network by the logical expression of the conditional probabilities within the calculable model

$$P(A|B) = \frac{P(A) * P(B|A)}{P(B)}$$

$$P(\text{OS Access}|\text{Share Portable Media}) = \frac{P(\text{OS Access}) * P(\text{Share Portable Media}|\text{OS Access})}{P(\text{Share Portable Media})}$$

Doing so in a spreadsheet without the use of a Bayesian network (or as a system of equations) cannot resolve what Excel considers to be a circular reference between the terms

$$P(\text{OS Access}|\text{Share Portable Media}) \text{ and } P(\text{Share Portable Media}|\text{OS Access})$$

Resolving this directly within the Excel model requires that we 'break' the causal link at some point between 'OS Access' and 'Share Portable Media' without substantively changing the necessary potential logical relationship between the two components (Operating System and Office Social Zone) and their corresponding component exploit and vulnerability elements. Where required, minimal changes to these relationships (typically the replacement of conditional probability relationships with fixed or one-way dependencies) were made to a number of component attributes to resolve this issue. It is my perspective that these changes were not material to the resulting simulated outcomes of the model which does not specifically profile, for example, portable media sharing, and is therefore acceptable for the purposes of generating simulations within in this research.

**Control Specification: Prevention and Blocking of Attacks (Web Application Secure Coding Example)**

This section details the incorporation of the various controls into the model components in a way that is intended to capture the inherently stochastic nature of the control effectiveness in a real world operational setting and which contributes directly to both the resulting probability distribution characteristics of the resulting business losses and the diversity of the loss profile scenarios that are required for the lab experiments as being characterized as decisions under both risk and uncertainty. While the CySeMol class model specified specific controls and the conditional probability relationships between controls and vulnerabilities, it does not directly incorporate the inherently stochastic nature of individual control *effectiveness* in a way that sufficiently reflects the real world variability of technical and administrative controls in operation. In CySeMol, controls are generally specified as being either present (True) or absent (False) and the thereafter the stochastic nature of their effectiveness is essentially captured within the conditional probability tables as a total effect together with the other conditional elements of the particular table.

In the context of our running example of the Web Application exploit 'Command Injection' for example, recall that the probability of a successful attack is, in the first instance, dependent on the state of two essential vulnerabilities: 1) whether the Web Application component ‘has’ a publicly discoverable command injection vulnerability, which is either strictly true or false, and whether the attacker can ‘find’ a public vulnerability, or 2) whether the attacker can discover a novel vulnerability in the Web Application which is, form the perspective of specifying controls directly on the Web Application, comparatively more complex. If neither of these conditions is ‘True’, then the Command Injection exploit cannot be successful. Finding a public command injection only requires logical access to the web application (which itself may be complex but is not a function of the Web Application control per se), but in the end is either possible or not. Discovering a vulnerability on the other hand is directly conditional on the control posture of the Web Application itself determined by whether the attacker can logically connect to the Web Application (‘ApplicationServer.ConnectTo’) and, perhaps more importantly, in four different *preventive* controls reflecting the secure coding practices and resulting code of the Web Application software itself:

**Figure 77 - CySeMol Class Object “Web Application” ‘ExplotCommandInjection’ Element Conditional Probability Table and associated ‘Discover Vulnerability’ Element Conditional Probability Table**

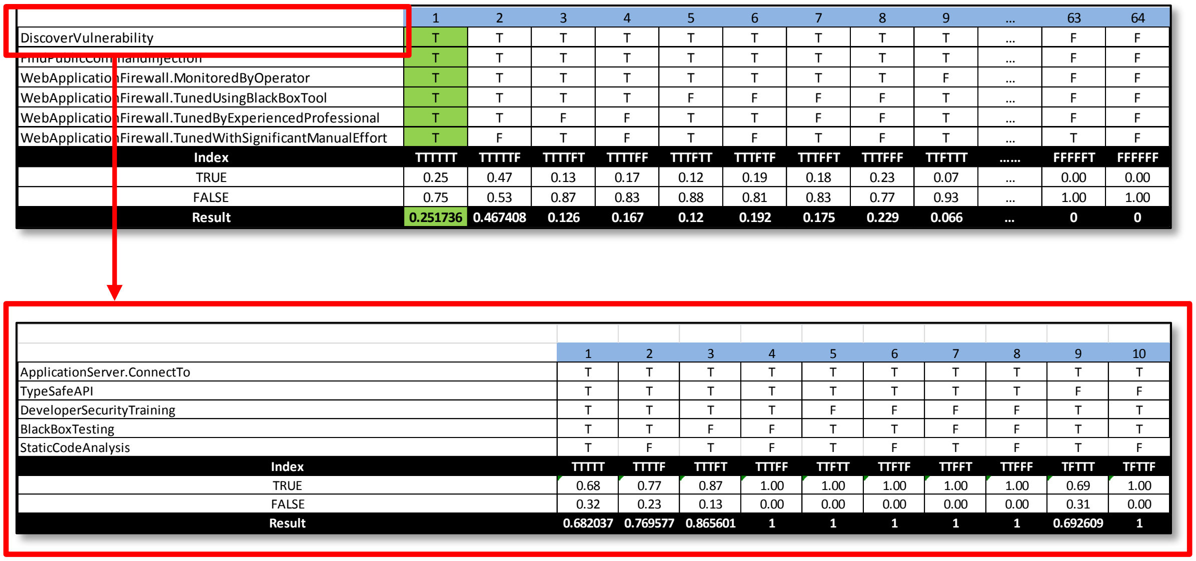

|  | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | ... | 63 | 64 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| DiscoverVulnerability | T | T | T | T | T | T | T | T | T | ... | F | F |
| FindPublicCommandInjection | T | T | T | T | T | T | T | T | T | ... | F | F |
| WebApplicationFirewall.MonitoredByOperator | T | T | T | T | T | T | T | T | F | ... | F | F |
| WebApplicationFirewall.TunedUsingBlackBoxTool | T | T | T | T | F | F | F | F | T | ... | F | F |
| WebApplicationFirewall.TunedByExperiencedProfessional | T | T | F | F | T | T | F | F | T | ... | F | F |
| WebApplicationFirewall.TunedWithSignificantManualEffort | T | F | T | F | T | F | T | F | T | ... | T | F |
| **Index** | **TTTTTT** | **TTTTTF** | **TTTTFT** | **TTTTFF** | **TTTFTT** | **TTTFTF** | **TTTFFT** | **TTTFFF** | **TTFTTT** | **......** | **FFFFFT** | **FFFFFF** |
| TRUE | 0.25 | 0.47 | 0.13 | 0.17 | 0.12 | 0.19 | 0.18 | 0.23 | 0.07 | ... | 0.00 | 0.00 |
| FALSE | 0.75 | 0.53 | 0.87 | 0.83 | 0.88 | 0.81 | 0.83 | 0.77 | 0.93 | ... | 1.00 | 1.00 |
| **Result** | **0.251736** | **0.467408** | **0.126** | **0.167** | **0.12** | **0.192** | **0.175** | **0.229** | **0.066** | **...** | **0** | **0** |

|  | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|---|---|---|---|---|---|---|---|---|---|---|
| ApplicationServer.ConnectTo | T | T | T | T | T | T | T | T | T | T |
| TypeSafeAPI | T | T | T | T | T | T | T | T | F | F |
| DeveloperSecurityTraining | T | T | T | T | F | F | F | F | T | T |
| BlackBoxTesting | T | T | F | F | T | T | F | F | T | T |
| StaticCodeAnalysis | T | F | T | F | T | F | T | F | T | F |
| **Index** | **TTTTT** | **TTTTF** | **TTTFT** | **TTTFF** | **TTFTT** | **TTFTF** | **TTFFT** | **TTFFF** | **TFTTT** | **TFTTF** |
| TRUE | 0.68 | 0.77 | 0.87 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 0.69 | 1.00 |
| FALSE | 0.32 | 0.23 | 0.13 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.31 | 0.00 |
| **Result** | **0.682037** | **0.769577** | **0.865601** | **1** | **1** | **1** | **1** | **1** | **0.692609** | **1** |

(Holm, Ekstedt et al. 2013)

The determination of the probabilities of successfully discovering a vulnerability based on the $2^4 = 16$ possible variations of the four controls are complex within CySeMol and are each based on a Bernoulli function of probability of success given the amount of effort expended by the Attacker in days:

**Figure 78 - CySeMol Web Application class 'Exploit Command Injection' vulnerability: conditional probability of successful attack in the presence of a Web Application Firewall**

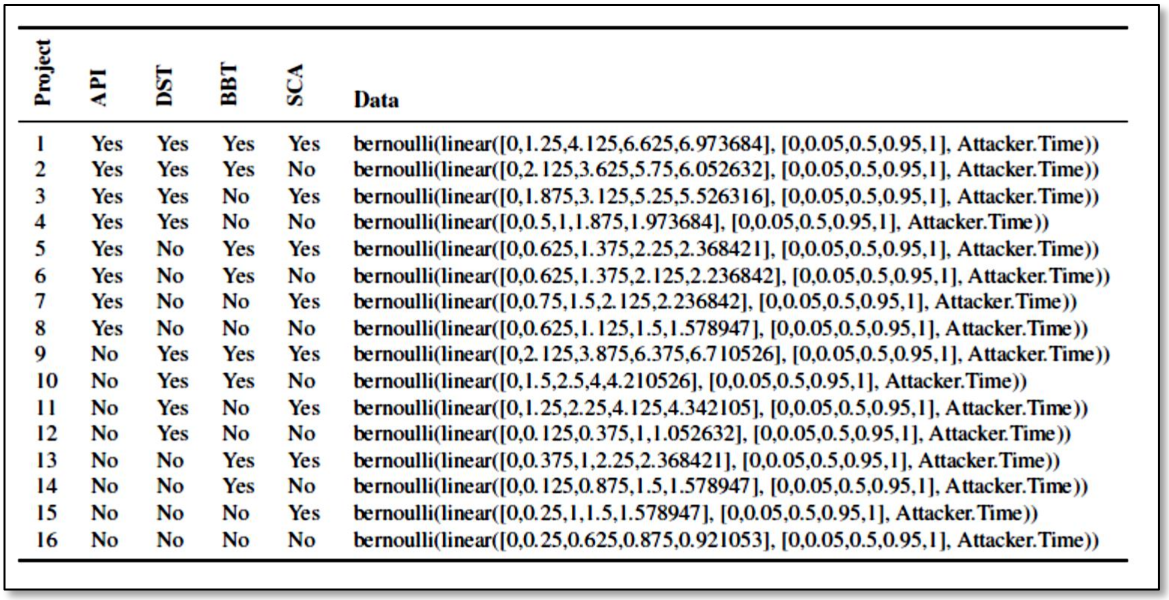

| Project | API | DST | BBT | SCA | Data |
|---|---|---|---|---|---|
| 1 | Yes | Yes | Yes | Yes | bernoulli(linear([0,1.25,4.125,6.625,6.973684], [0,0.05,0.5,0.95,1], Attacker.Time)) |
| 2 | Yes | Yes | Yes | No | bernoulli(linear([0,2.125,3.625,5.75,6.052632], [0,0.05,0.5,0.95,1], Attacker.Time)) |
| 3 | Yes | Yes | No | Yes | bernoulli(linear([0,1.875,3.125,5.25,5.526316], [0,0.05,0.5,0.95,1], Attacker.Time)) |
| 4 | Yes | Yes | No | No | bernoulli(linear([0,0.5,1,1.875,1.973684], [0,0.05,0.5,0.95,1], Attacker.Time)) |
| 5 | Yes | No | Yes | Yes | bernoulli(linear([0,0.625,1.375,2.25,2.368421], [0,0.05,0.5,0.95,1], Attacker.Time)) |
| 6 | Yes | No | Yes | No | bernoulli(linear([0,0.625,1.375,2.125,2.236842], [0,0.05,0.5,0.95,1], Attacker.Time)) |
| 7 | Yes | No | No | Yes | bernoulli(linear([0,0.75,1.5,2.125,2.236842], [0,0.05,0.5,0.95,1], Attacker.Time)) |
| 8 | Yes | No | No | No | bernoulli(linear([0,0.625,1.125,1.5,1.578947], [0,0.05,0.5,0.95,1], Attacker.Time)) |
| 9 | No | Yes | Yes | Yes | bernoulli(linear([0,2.125,3.875,6.375,6.710526], [0,0.05,0.5,0.95,1], Attacker.Time)) |
| 10 | No | Yes | Yes | No | bernoulli(linear([0,1.5,2.5,4,4.210526], [0,0.05,0.5,0.95,1], Attacker.Time)) |
| 11 | No | Yes | No | Yes | bernoulli(linear([0,1.25,2.25,4.125,4.342105], [0,0.05,0.5,0.95,1], Attacker.Time)) |
| 12 | No | Yes | No | No | bernoulli(linear([0,0.125,0.375,1,1.052632], [0,0.05,0.5,0.95,1], Attacker.Time)) |
| 13 | No | No | Yes | Yes | bernoulli(linear([0,0.375,1,2.25,2.368421], [0,0.05,0.5,0.95,1], Attacker.Time)) |
| 14 | No | No | Yes | No | bernoulli(linear([0,0.125,0.875,1.5,1.578947], [0,0.05,0.5,0.95,1], Attacker.Time)) |
| 15 | No | No | No | Yes | bernoulli(linear([0,0.25,1,1.5,1.578947], [0,0.05,0.5,0.95,1], Attacker.Time)) |
| 16 | No | No | No | No | bernoulli(linear([0,0.25,0.625,0.875,0.921053], [0,0.05,0.5,0.95,1], Attacker.Time)) |

(Holm, Ekstedt et al. 2013)

In this approach, the more controls that are present the longer it takes the attacker to compromise the system, ranging from less than a day when none of the four controls are present to approximately 7 days when all four controls are present.

While we accept that this appropriately reflects the probability of discovering a vulnerability given the state of the controls, it does not account for the state of the controls themselves. In other words, from a decision making perspective, whether we assume that a vulnerability can either be found or discovered, from a controls perspective the managerial focus for this component should, appropriately for our purposes, shift to the determination of the state of the four specific Web Application controls: 'TypeSafeAPI', 'DeveloperSecurityTraining', 'BlackBoxTesting', and 'StaticCodeAnalysis'. The default state of these underlying control conditions in CySeMol is, respectively, False, False, True and True, as determined through expert interviews (Holm and Ekstedt 2012), but could be specified as specific evidence according to the operational circumstances of the model (i.e. Developer Security Training could be specified as either True or False depending on what we believed to be the case in our specific simulation circumstances). For our purposes, simply specifying that 'Developer Security Training' is either present or absent is too restrictive for the determination of the presence of the controls since, arguably, a qualitative control like Developer Security Training, even if undertaken, may be more or less relatively effective or not depending on, for example, how well the training was delivered, understood or whether it in fact translated into secure coding of the application itself. While the resulting state of the control will always be necessarily either True or False for the purposes of calculating the conditional probability of the dependent element(s) of the logical security model (here, the 'DiscoverVulnerability' element, which is itself a dependency for the

‘Command Injection’ Exploit), we require that the probability of a control being ‘True’ should be increasing with the a specifiable level of control effectiveness as determined by management i.e. variably selectable from None through to Very High by either the researcher or possibly by the lab participant which in turn increases the probability that Developer Security Training is in fact ‘True’.

This suggests that an additional layer of control effectiveness should be introduced prior to the determination of the state of individual controls, further reflecting the stochastic nature of individual control influence in operations. The following describes how this was accomplished within the Excel model in the context of our running example.

For the purposes of generating a reasonable range of scenarios we propose to limit the choice of controls from ‘Low’ Through ‘High’, eliminating the possibility of ‘Minimal’ or ‘None’ since, through model testing, control scenarios involving ‘Minimal’ or None’ tended to produce unstable (i.e. constantly increasing) and therefore operationally unacceptable results based on the lag effects built into the model as described below. For lab experiment purposes, the differential between Low and Very High is also sufficiently broad enough to produce the range of losses that is reflective of the desired lab choice scenarios.

A console is built into the simulator enabling the user to specify 4 levels of control across the 4 types of control:

**Figure 79 – Simulation Controls Console (Curtis)**

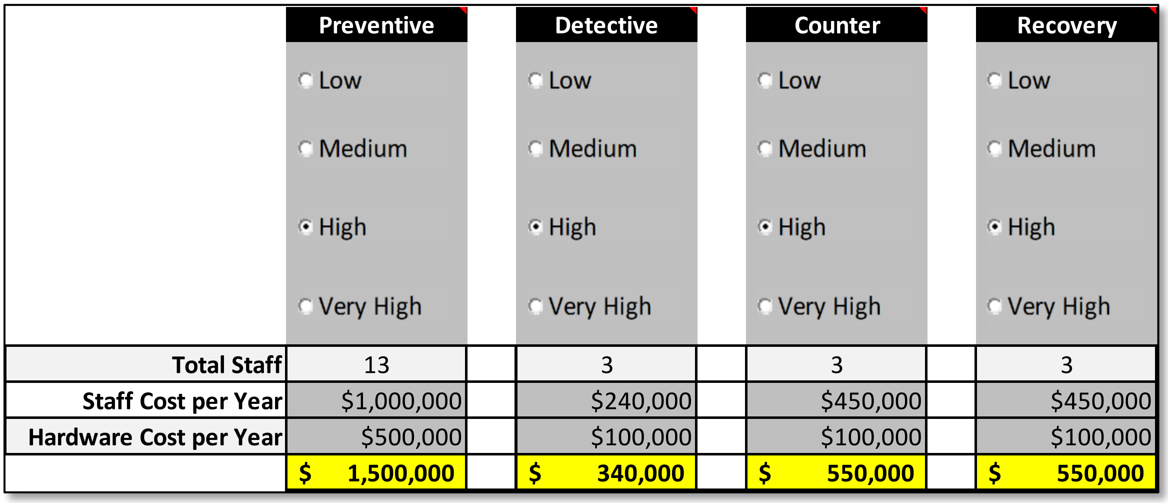

| | Preventive | Detective | Counter | Recovery |
|---|---|---|---|---|
| | ○ Low | ○ Low | ○ Low | ○ Low |
| | ○ Medium | ○ Medium | ○ Medium | ○ Medium |
| | ● High | ● High | ● High | ● High |
| | ○ Very High | ○ Very High | ○ Very High | ○ Very High |
| **Total Staff** | 13 | 3 | 3 | 3 |
| **Staff Cost per Year** | $1,000,000 | $240,000 | $450,000 | $450,000 |
| **Hardware Cost per Year** | $500,000 | $100,000 | $100,000 | $100,000 |
| | **$ 1,500,000** | **$ 340,000** | **$ 550,000** | **$ 550,000** |

Control types include those controls applicable to specific component control elements, here based on those described by Jakoubi, Tjoa et al (Jakoubi, Tjoa et al. 2007; Tjoa, Jakoubi et al. 2011) modified here to be applicable to specific component controls within the CySeMol class model:

**Table 7 – Preventive, Blocking, Detective, Counter and Recovery Control Types**

| Preventive/Blocking | Detective | Counter | Recovery |
|---|---|---|---|
| *Reduces the probability of a successful threat* | *Reduces the Impact of a successful threat* | | |
| • Security Policy<br>• Security Awareness Training<br>• Web Application Secure Coding Techniques<br>• Cryptography<br>• Secure Infrastructure Management<br>• Blocking: Intrusion Prevention System, anti-virus controls, etc. | • Help Desk Monitoring<br>• Network/Security Operations Centre monitoring<br>• Compliance and Auditing monitoring<br>• Sandboxing | • Quarantining/isolation<br>• System/component failover to backup/secondary system<br>• Virus/Malware Removal | • Data or system restoration from backup<br>• Primary system testing and restart<br>• Restoration of service to restored system<br>• Cutover to primary system |

(Jakoubi, Tjoa et al. 2007; Tjoa, Jakoubi et al. 2011)

In anticipation that control costs might be a consideration in some scenarios, costs associated with staffing and hardware were attributed to each set of controls although the resulting experiments ultimately did not directly utilize these factors. The introduction of control costs, together with the losses produced by the system might suggest that the model should be regarded from the decision maker's point of view primarily as an optimization problem similar to the approach of Wang (Wang, Chaudhury et al. 2008) or Sawik (Sawik 2013) and others reviewed above (Neubauer and Hartl 2009). While optimization of the controls against prospective losses is certainly one application of this type of simulation model, optimization of the 'return' on security investment, or otherwise determining what an 'optimal' security posture would be in the context of the available controls is not the focus of this research and therefore significant time was not dedicated to determining highly realistic costing values either for the controls or for the attributed business losses as described in the next section. Some of the lab experiments also require scaling of scenario losses to dollar value ranges that support specific dollar value choice trade-offs based on the benchmarked experiment where overall gains or losses are in the sub-$100 range and marginal gain/loss values are in the -$15 to $15 dollar range. These gain/loss values are clearly not realistic in terms of the operational expectations for a real world enterprise system at risk and any associated controls costs, if included as a factor in the decision, would need to be similarly scaled to relatively low values again not reflective of any security or IT system management budget in the real world.  The overall goal for the model was therefore to produce a plausible annual *profile* of daily business losses characterized by the average, variance, and extreme value distribution of the losses based on the behaviour of the system at different *nominal* control levels without attributing realistic costs to the control themselves or introducing the notion of risk neutral optimization. Indeed, the core purpose of the research is to allow participants to reveal their risk attitude and decisional biases under informational uncertainty by carefully controlling other decisional factors, including examination of the plausibility or optimization of the control costs, that might otherwise confound choices over outcomes that are meant to reveal decisional bias and not assume risk neutrality in the context of an optimization problem.

Each of the Preventive control elements in the table are then converted to a control widget that could be used to interactively set the level of control desired for the simulation scenario. In the following I explain the configuration of a typical control using the example of **‘Prevention – Web Application Secure Coding techniques’**.

**Figure 80 – Setting the Desired Level of Control for ‘Web Application Secure Coding Practices’**

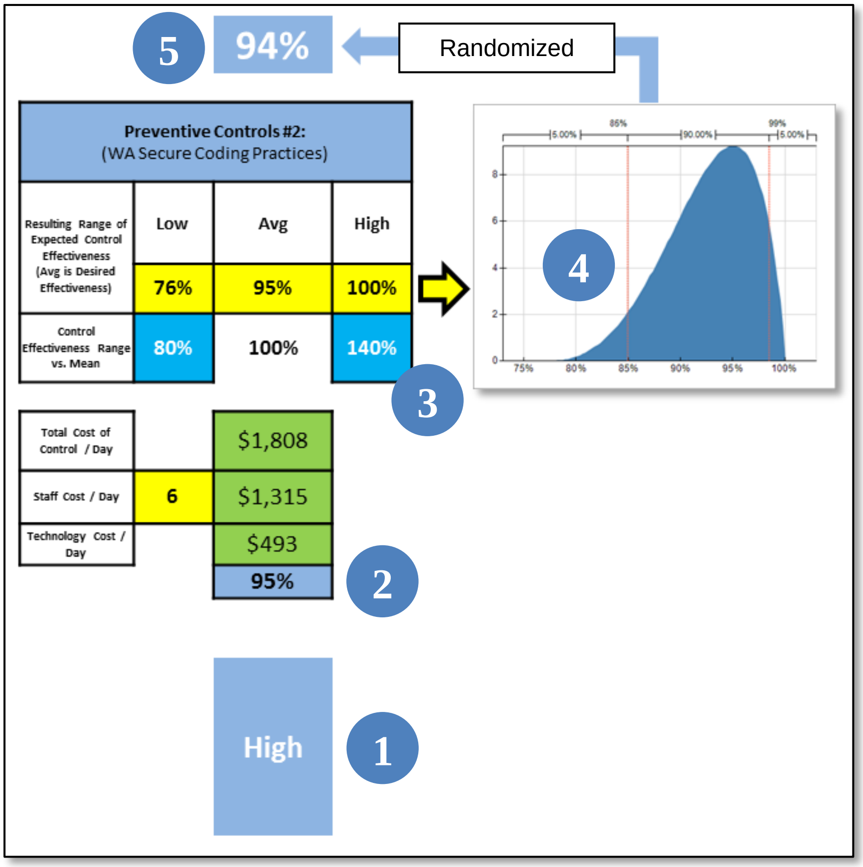

| Preventive Controls #2: (WA Secure Coding Practices) | | | |
|---|---|---|---|
| | Low | Avg | High |
| Resulting Range of Expected Control Effectiveness (Avg is Desired Effectiveness) | 76% | 95% | 100% |
| Control Effectiveness Range vs. Mean | 80% | 100% | 140% |

| | | |
|---|---|---|
| Total Cost of Control / Day | | $1,808 |
| Staff Cost / Day | 6 | $1,315 |
| Technology Cost / Day | | $493 |
| | | 95% |

1) Using this onscreen widget **(Figure 80)**, the user is permitted to specify one of 6 nominal levels of control for a particular simulation: None, Minimal, Low, Medium, High and Very High, although for practical purposes only Low through Very High were used for the final lab simulations as noted above.
2) The widget then looks up the corresponding “Average Effectiveness of Control” from a lookup table (**Figure 81)**:

**Figure 81 – Control Effectiveness Lookup Table**

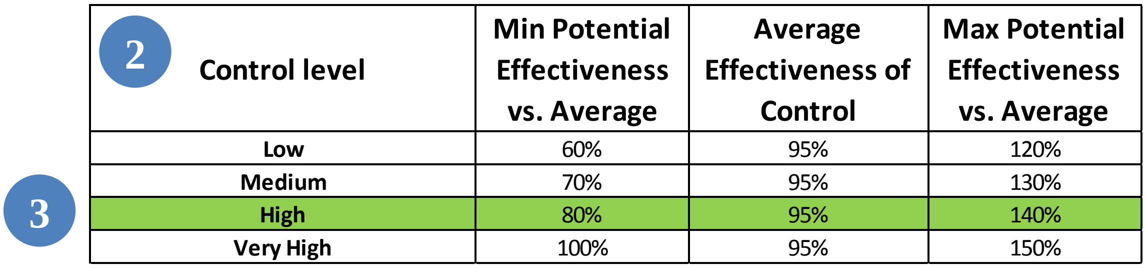

| Control level | Min Potential Effectiveness vs. Average | Average Effectiveness of Control | Max Potential Effectiveness vs. Average |
|---|---|---|---|
| **Low** | 60% | 95% | 120% |
| **Medium** | 70% | 95% | 130% |
| **High** | 80% | 95% | 140% |
| **Very High** | 100% | 95% | 150% |

3) A range of effectiveness for the control is specified, indicating the potential for a standard control to vary in its actual effectiveness during the simulation. This reflects the idea that a control's effectiveness in operation may vary, for example, based on the diligence with which it is deployed, the number of staff resources applied to manage the control, or the actual acceptance of the control by users, etc., although any control has a chance to perform at least as well as was intended by the average or expected effect. In the above example, a 'High' control which is 95% effective on average would be expected to perform no worse than 80% of the average and possibly up to 140% of the average to a maximum of 100%. This permits the model to employ controls with lower or diverse average effects (e.g. inferior grade policies or software controls) but to also capture the idea that a 'higher' or better level or quality of deployment is expected to narrow the variance of the control's actual performance results regardless of the anticipated average effectiveness.
4) The resulting effectiveness range (Min = 76%; Average = 95%; Max = 100%) is then converted to a PERT distribution[55] **(Figure 80)** which is typically used for the estimation of probability distributions with limited information based on expert judgment (Malcolm, Roseboom et al. 1959; Vose 2008).
5) The PERT distribution is then used to randomize the effectiveness value of the control on each Monte Carlo simulation round (i.e. daily) where most of the time the value will reflect performance around the average but may vary as low or as high as the minimum or maximum respectively. In this example, the randomized value on the day is 94% **(Figure 80)**. This permits the actual effectiveness of the control to vary within the anticipated range of effectiveness and reflects the inherent uncertainty of the control's effectiveness in use despite its intended effectiveness.
6) The randomized effectiveness value (76%-100%) is then used as the input probability in a Bernoulli function to calculate whether the controls covered by the widget were in fact effective for the simulation round **(Figure 82)**. In this example, the Web Application Secure Coding Practices consist of 4 sub controls, each with its own individual model dependencies: 'Black Box testing'; Static Code Analysis'; 'Type Safe API'; and 'Developer Security Training'. While as a group these controls are considered 94% effective, only 3 of the 4 controls have Bernoulli values resulting in 1 for this simulation round: the sub-control 'TypeSafeAPI' is considered to be ineffective for that simulation

[55] PERT stands for Program Evaluation and Review Technique (PERT) (Malcolm, Roseboom et al. 1959). See https://reference.wolfram.com/language/ref/PERTDistribution.html for a description of the specification and use of the PERT function.

round (Value = ‘F’ or FALSE) and increases the probability that an attack on a vulnerability associated with that control will be successful for that round.

**Figure 82 - Conversion of Control Effectiveness Level to Bernoulli Values**

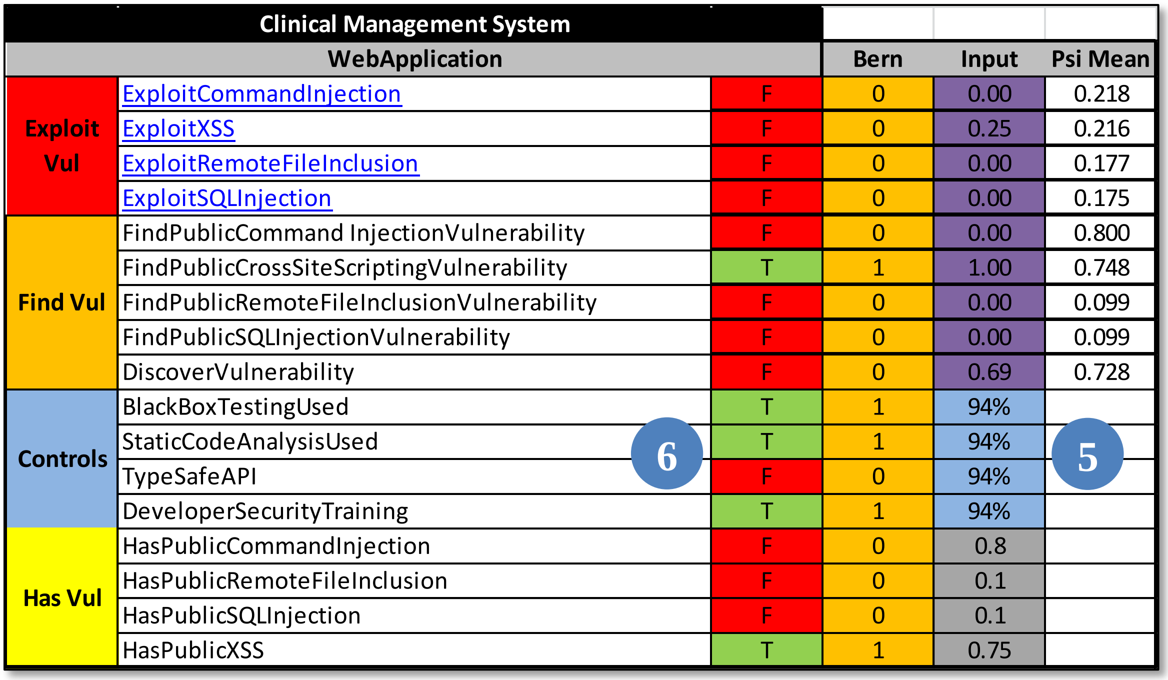

| Clinical Management System | | | | | |
|---|---|---|---|---|---|
| WebApplication | | | Bern | Input | Psi Mean |
| Exploit Vul | ExploitCommandInjection | F | 0 | 0.00 | 0.218 |
| | ExploitXSS | F | 0 | 0.25 | 0.216 |
| | ExploitRemoteFileInclusion | F | 0 | 0.00 | 0.177 |
| | ExploitSQLInjection | F | 0 | 0.00 | 0.175 |
| Find Vul | FindPublicCommand InjectionVulnerability | F | 0 | 0.00 | 0.800 |
| | FindPublicCrossSiteScriptingVulnerability | T | 1 | 1.00 | 0.748 |
| | FindPublicRemoteFileInclusionVulnerability | F | 0 | 0.00 | 0.099 |
| | FindPublicSQLInjectionVulnerability | F | 0 | 0.00 | 0.099 |
| | DiscoverVulnerability | F | 0 | 0.69 | 0.728 |
| Controls | BlackBoxTestingUsed | T | 1 | 94% | |
| | StaticCodeAnalysisUsed | T | 1 | 94% | |
| | TypeSafeAPI | F | 0 | 94% | |
| | DeveloperSecurityTraining | T | 1 | 94% | |
| Has Vul | HasPublicCommandInjection | F | 0 | 0.8 | |
| | HasPublicRemoteFileInclusion | F | 0 | 0.1 | |
| | HasPublicSQLInjection | F | 0 | 0.1 | |
| | HasPublicXSS | T | 1 | 0.75 | |

### Attribution of Confidentiality, Integrity and Availability Events

The system then attributes security incidents to specific successful exploits which are inferred to result in associated confidentiality (C), integrity (I) or availability (A) impacts to the component based on CVE guidance on[56]. The total number of potential ‘breaches’ in C, I or A per unit of simulation is determined by the number of successful exploits multiplied by the chance of exploit. In the following Figure, only the ‘Exploit Command Injection’ was successful out of the four possible exploits, and has resulted in the possibility of both one Integrity breach and one Availability breach:

[56] ibid

**Figure 83 – Calculation of C, I or A breaches from Successful (non-prevented) Attacks**

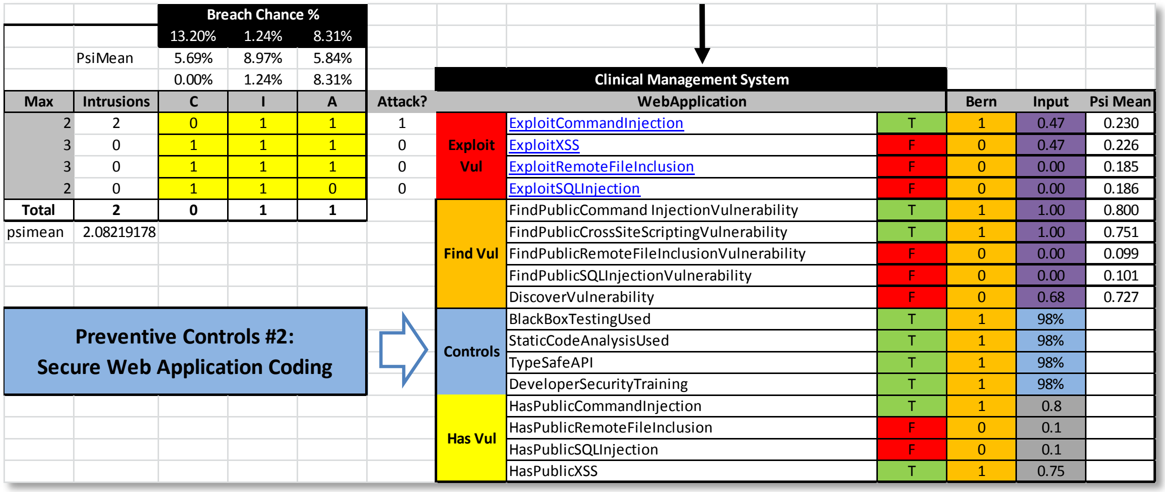

|  |  | Breach Chance % |  |  |
|---|---|---|---|---|
|  |  | 13.20% | 1.24% | 8.31% |
|  | PsiMean | 5.69% | 8.97% | 5.84% |
|  |  | 0.00% | 1.24% | 8.31% |
| **Max** | **Intrusions** | **C** | **I** | **A** |
| 2 | 2 | 0 | 1 | 1 |
| 3 | 0 | 1 | 1 | 1 |
| 3 | 0 | 1 | 1 | 1 |
| 2 | 0 | 1 | 1 | 0 |
| **Total** | **2** | **0** | **1** | **1** |
| psimean | 2.08219178 |  |  |  |

**Preventive Controls #2:**
**Secure Web Application Coding**

| **Attack?** | **Clinical Management System**<br>**WebApplication** |  |  | **Bern** | **Input** | **Psi Mean** |
|---|---|---|---|---|---|---|
| 1 | **Exploit Vul** | ExploitCommandInjection | T | 1 | 0.47 | 0.230 |
| 0 |  | ExploitXSS | F | 0 | 0.47 | 0.226 |
| 0 |  | ExploitRemoteFileInclusion | F | 0 | 0.00 | 0.185 |
| 0 |  | ExploitSQLInjection | F | 0 | 0.00 | 0.186 |
|  | **Find Vul** | FindPublicCommand InjectionVulnerability | T | 1 | 1.00 | 0.800 |
|  |  | FindPublicCrossSiteScriptingVulnerability | T | 1 | 1.00 | 0.751 |
|  |  | FindPublicRemoteFileInclusionVulnerability | F | 0 | 0.00 | 0.099 |
|  |  | FindPublicSQLInjectionVulnerability | F | 0 | 0.00 | 0.101 |
|  |  | DiscoverVulnerability | F | 0 | 0.68 | 0.727 |
|  | **Controls** | BlackBoxTestingUsed | T | 1 | 98% |  |
|  |  | StaticCodeAnalysisUsed | T | 1 | 98% |  |
|  |  | TypeSafeAPI | T | 1 | 98% |  |
|  |  | DeveloperSecurityTraining | T | 1 | 98% |  |
|  | **Has Vul** | HasPublicCommandInjection | T | 1 | 0.8 |  |
|  |  | HasPublicRemoteFileInclusion | F | 0 | 0.1 |  |
|  |  | HasPublicSQLInjection | F | 0 | 0.1 |  |
|  |  | HasPublicXSS | T | 1 | 0.75 |  |

Once the possible types of breaches are determined, the system adds up all of the ‘non-prevented’ breaches and passes these to an algorithm which then determines whether the breaches are detected within the simulation round. Non-detected breaches are added to the following simulation round’s count of non-prevented breaches which introduces a one-period lag or system memory effect at this stage in the time series. This is considered more realistic than a system in which we assume all non-prevented breaches are fully detected in each round and makes the system security performance more sensitive to the level of detection controls.

### Detection of Successful Exploits

Detection of breaches is a function of a range of monitoring controls that affect the organization’s ability to identify breaches as they are occurring and before they have a chance to have substantial impact on the system and the associated dependent business processes. Successful detection lowers the number of breaches that can possibly impact the system within the current simulation round. The number of detected breaches out of the total number of non-prevented breaches is computed based on the Detection control level set by the user and the resulting random control effectiveness percentage calculated in a manner similar to the Prevention Controls described above, where each non-prevented breach has a chance of being detected in proportion to the effectiveness of the Detection control. Successfully detected breaches are passed to the Counter Phase which attempts to eliminate the breach. ‘Non-detected’ breaches are subsequently added to the next day’s newly non-prevented breaches and are subject to detection again in the next simulation round.

Modelling undetected breaches in this manner introduces a realistic one-period lag effect on the resulting time series of losses where undetected breaches continually degrade system performance and must be detected before counter controls can be applied. The system at this stage appropriately simulates as expected: if insufficient resources are spent on Detection, the number of undetected breaches increases over time and eventually overwhelm the Detection resources which cannot catch up. Underspending on detection can therefore cause the number of undetected breaches to increase to the point where the attributed impacts exceed a threshold (e.g. system availability dips below a operationally viable level) and the business experiences catastrophic losses. Figure 84 indicates the conversion of successful C, I, or A breaches into detected and non-detected security incidents respectively:

**Figure 84 - Conversion of Non-Prevented Breaches to Detected and Non-Detected Breaches**

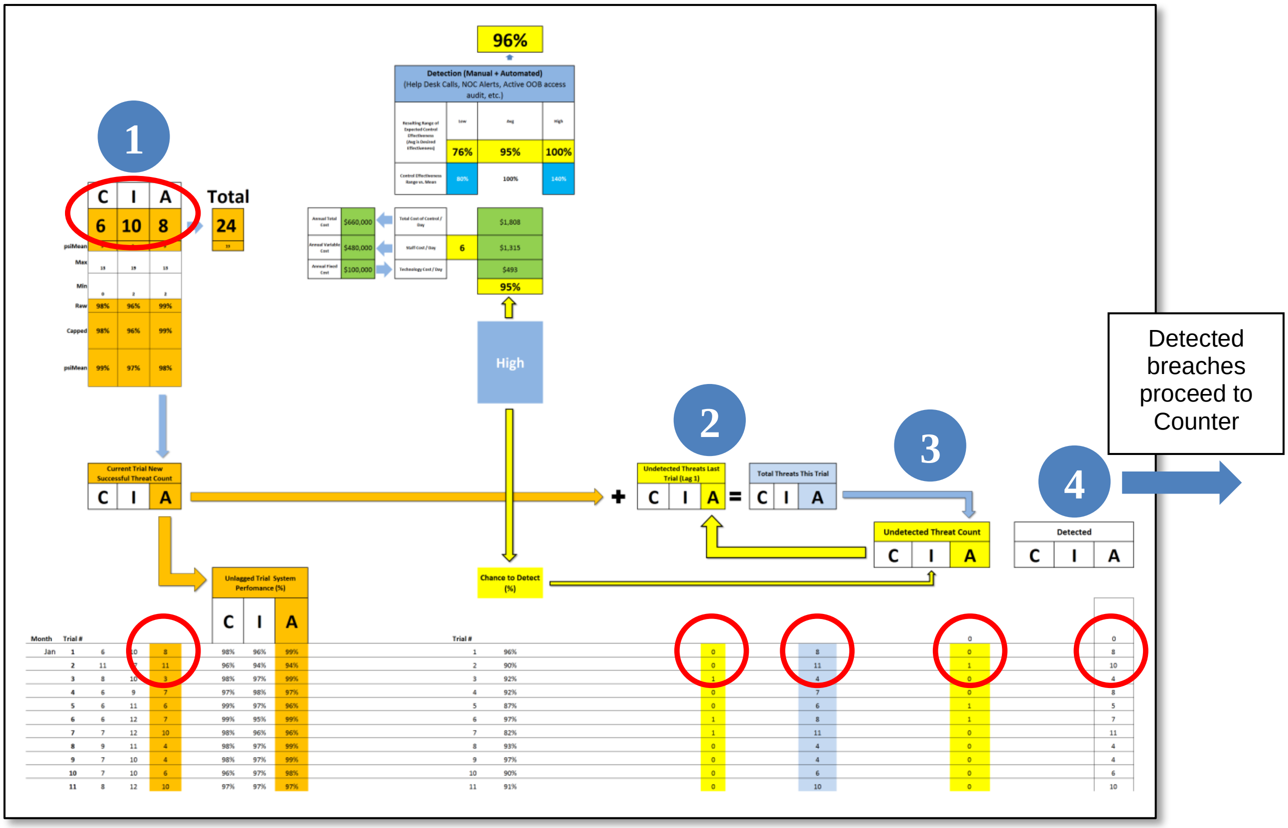
96%
Detection (Manual + Automated) (Help Desk Calls, NOC Alerts, Active OOB access audit, etc.)
76%
95%
100%
1
C I A
6 10 8
Total
24
$660,000
$480,000
$100,000
6
$1,808
$1,315
$493
95%
High
2
3
4
Detected breaches proceed to Counter
Current Trial New Successful Threat Count
C I A
Undetected Threats Last Trial (Lag 1)
C I A
=
Total Threats This Trial
C I A
Undetected Threat Count
C I A
Detected
C I A
Unlagged Trial System Perfomance (%)
C I A
Chance to Detect (%)
Month
Trial #

1. New successful (i.e. ‘non-prevented’) C, I, or A breaches in the current simulation round are summed across all exploited system components (➔ e.g. 8 Availability breaches)
2. Last period’s ‘Undetected breaches’ are added to the current round’s ‘non-prevented’ breaches
3. Total current round breaches are then subject to ‘Detection’ based on the effectiveness level of the ‘Detection’ controls
4. Successfully detected breaches in current round are passed to the ‘Counter’ phase and do not immediately affect system availability in the current round unless subsequently ‘not-countered’. Undetected breaches in the current round impact System availability in the current round (see section 7 below)

### Countering Detected Breaches

In the Counter Phase, detected breaches have a chance of being countered or interrupted by support staff assigned to eliminate detected threats in the environment. Similar to the Detection Control, Counter efforts add a lag effect to the resulting time series of losses since un-countered breaches in one period must be successfully countered in a subsequent period before Recovery controls can be applied and any resulting uncountered breaches have a chance to degrade system performance.

For this control, available counter resources (staff) are assumed to be able to attempt to counter all detected breaches within the day[57]. Similar to the treatment for Detection, the number of ‘detected but not countered’ breaches is then computed based on the Counter control level and the resulting effectiveness percentage set for the simulation. The total number of ‘detected-but-not-countered’ breaches is calculated and then added to the next day’s newly detected breaches; countered breaches are passed to the Recovery Phase. The system at this stage appropriately simulates as expected: if Counter controls are insufficient to address the number of detected breaches, the number of un-countered breaches increases over time and eventually overwhelms the Counter resources which cannot catch up. Underspending on Countering can therefore cause the number of un-countered breaches to increase to the point where the attributed impacts exceed a threshold (e.g. system availability dips below an operationally viable level) and the business experiences catastrophic losses. Figure 85 illustrates the conversion detected breaches into countered and non-countered incidents:

[57] In a real world setting, individual staff would typically be assigned multiple detected breach ‘tickets’ to work on at the same time but may not be able to attempt to counter each detected breach within a given day. In that case, some detected breaches would remain automatically ‘uncountered’ (since countering was not even attempted) and would be directly added to the next day’s detected breaches. In the current model we assume that we can attempt to counter each detected breach within the day regardless of the number of resources assigned to countering. This factor could be easily adjusted to simulate increased or decreased staffing levels, ‘coverage’ or throughput by countering personnel, aside from the overall effectiveness assigned to the Counter control itself. For example, if we assumed that staff could effectively attempt to counter one breach per day, if the number of detected breaches exceeded the number of available staff assigned to counter breaches on the day, the difference between the number of staff and the number of detected breaches needing to be countered could be automatically added to the next day’s newly detected breaches. In sensitivity tests of that version of the model however, even at high control effectiveness, low non-attempted numbers of breaches quickly add up over several rounds and eventually overwhelm the system. Additional tuning of the model would therefore be required to reflect the difference between overall Counter control *effectiveness* vs. resource *coverage* of open breaches.

**Figure 85 - Conversion of Detected Breaches to Countered and Non-Countered Security Breaches**

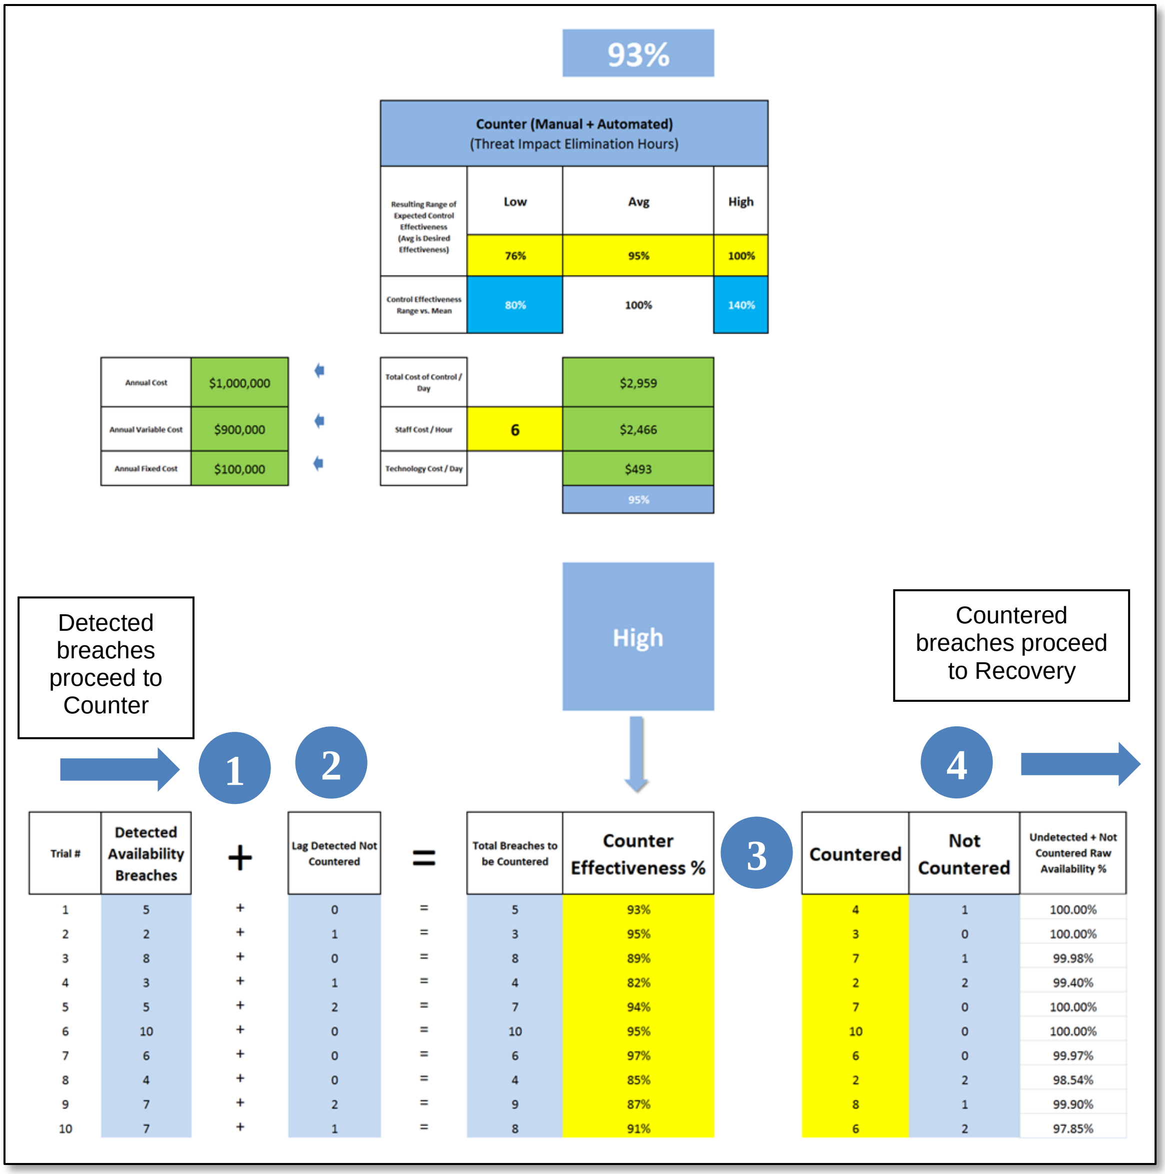

93%
Detected breaches proceed to Counter
Countered breaches proceed to Recovery
High
1
2
3
4

| Counter (Manual + Automated)<br>(Threat Impact Elimination Hours) | | | |
|---|---|---|---|
| Resulting Range of Expected Control Effectiveness (Avg is Desired Effectiveness) | Low | Avg | High |
| | 76% | 95% | 100% |
| Control Effectiveness Range vs. Mean | 80% | 100% | 140% |

| | |
|---|---|
| Annual Cost | $1,000,000 |
| Annual Variable Cost | $900,000 |
| Annual Fixed Cost | $100,000 |

| | | |
|---|---|---|
| Total Cost of Control / Day | | $2,959 |
| Staff Cost / Hour | 6 | $2,466 |
| Technology Cost / Day | | $493 |
| | | 95% |

| Trial # | Detected Availability Breaches | + | Lag Detected Not Countered | = | Total Breaches to be Countered | Counter Effectiveness % | Countered | Not Countered | Undetected + Not Countered Raw Availability % |
|---|---|---|---|---|---|---|---|---|---|
| 1 | 5 | + | 0 | = | 5 | 93% | 4 | 1 | 100.00% |
| 2 | 2 | + | 1 | = | 3 | 95% | 3 | 0 | 100.00% |
| 3 | 8 | + | 0 | = | 8 | 89% | 7 | 1 | 99.98% |
| 4 | 3 | + | 1 | = | 4 | 82% | 2 | 2 | 99.40% |
| 5 | 5 | + | 2 | = | 7 | 94% | 7 | 0 | 100.00% |
| 6 | 10 | + | 0 | = | 10 | 95% | 10 | 0 | 100.00% |
| 7 | 6 | + | 0 | = | 6 | 97% | 6 | 0 | 99.97% |
| 8 | 4 | + | 0 | = | 4 | 85% | 2 | 2 | 98.54% |
| 9 | 7 | + | 2 | = | 9 | 87% | 8 | 1 | 99.90% |
| 10 | 7 | + | 1 | = | 8 | 91% | 6 | 2 | 97.85% |

1. Current round detected breaches proceed to the ‘Counter’ phase.
2. Last round’s ‘detected but not countered’ breaches are added to the current round’s detected breaches to determine total number of potentially counterable breaches.
3. The number of countered and non-countered breaches for this round is determined based on the effectiveness level of the Counter control.
4. Countered breaches proceed to ‘Recover’ phase. Non-countered breaches are added to next round’s potentially counterable breaches and affect system Availability.

### Recovering from Countered Breaches

Finally, in the Recovery Phase, the system has a chance to recover system functionality from detected and countered breaches by support staff assigned to the Recovery control function. Similar to the Detection and Counter Controls, Recovery efforts add a lag effect to the resulting time series of losses since any Countered but un-recovered breaches must be addressed in the next round of Recovery. Similar to the treatment for Counter controls, the number of 'countered but not recovered' breaches is computed based on the Recovery control level and the resulting effectiveness percentage set by the user in a manner similar to the Prevention, Detection and Counter Controls. Each Countered breach addressed by the assigned Recovery staff member has a chance of being not recovered in inverse proportion to the effectiveness of the control. The total number of number of 'countered but not recovered' breaches is then added to the next day's newly countered breaches. Recovered breaches are then subtracted from the total number of breaches in the day that have a chance to affect C, I or A factors:

**Figure 86 - Conversion of Countered Breaches to Recovered and Non-Recovered Security Breaches**

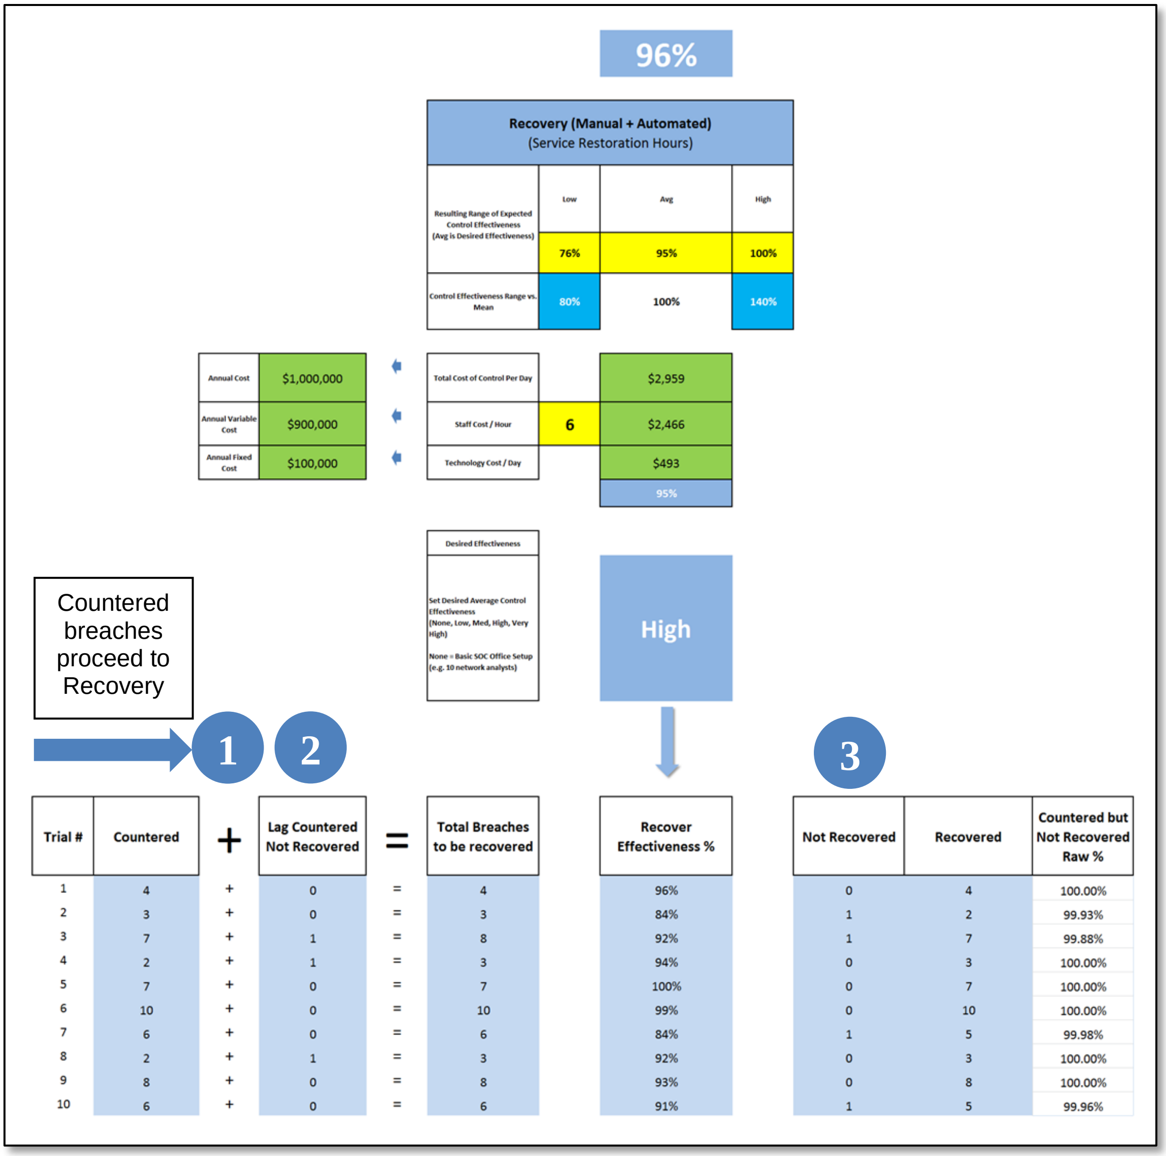

| Recovery (Manual + Automated)<br>(Service Restoration Hours) | Low | Avg | High |
|---|---|---|---|
| Resulting Range of Expected Control Effectiveness<br>(Avg is Desired Effectiveness) | 76% | 95% | 100% |
| Control Effectiveness Range vs. Mean | 80% | 100% | 140% |

| | |
|---|---|
| Annual Cost | $1,000,000 |
| Annual Variable Cost | $900,000 |
| Annual Fixed Cost | $100,000 |

| | | |
|---|---|---|
| Total Cost of Control Per Day | | $2,959 |
| Staff Cost / Hour | 6 | $2,466 |
| Technology Cost / Day | | $493 |
| | | 95% |

Desired Effectiveness

Set Desired Average Control Effectiveness (None, Low, Med, High, Very High)

None = Basic SOC Office Setup (e.g. 10 network analysts)

| Trial # | Countered | + | Lag Countered Not Recovered | = | Total Breaches to be recovered | Recover Effectiveness % | Not Recovered | Recovered | Countered but Not Recovered Raw % |
|---|---|---|---|---|---|---|---|---|---|
| 1 | 4 | + | 0 | = | 4 | 96% | 0 | 4 | 100.00% |
| 2 | 3 | + | 0 | = | 3 | 84% | 1 | 2 | 99.93% |
| 3 | 7 | + | 1 | = | 8 | 92% | 1 | 7 | 99.88% |
| 4 | 2 | + | 1 | = | 3 | 94% | 0 | 3 | 100.00% |
| 5 | 7 | + | 0 | = | 7 | 100% | 0 | 7 | 100.00% |
| 6 | 10 | + | 0 | = | 10 | 99% | 0 | 10 | 100.00% |
| 7 | 6 | + | 0 | = | 6 | 84% | 1 | 5 | 99.98% |
| 8 | 2 | + | 1 | = | 3 | 92% | 0 | 3 | 100.00% |
| 9 | 8 | + | 0 | = | 8 | 93% | 0 | 8 | 100.00% |
| 10 | 6 | + | 0 | = | 6 | 91% | 1 | 5 | 99.96% |

1. Current round countered breaches proceed to the ‘Recover’ phase.
2. Last round’s ‘countered but not recovered’ breaches are added to the current round’s countered breaches to determine total number of potentially recoverable breaches.
3. The number of recovered and non-recovered breaches for this round is determined based on the effectiveness level of the Recover control. Recovered breaches terminate and are not counted as affecting system Availability. Non-recovered breaches are added to next round’s potentially recoverable breaches and affect system Availability.

The following figure illustrates the combination of undetected, detected-not-countered and countered-not-recovered breaches for a typical High Controls system for one year:

**Figure 87 – One Year Time Series of System Availability and Daily Business Losses based on Effective Breach Counts**

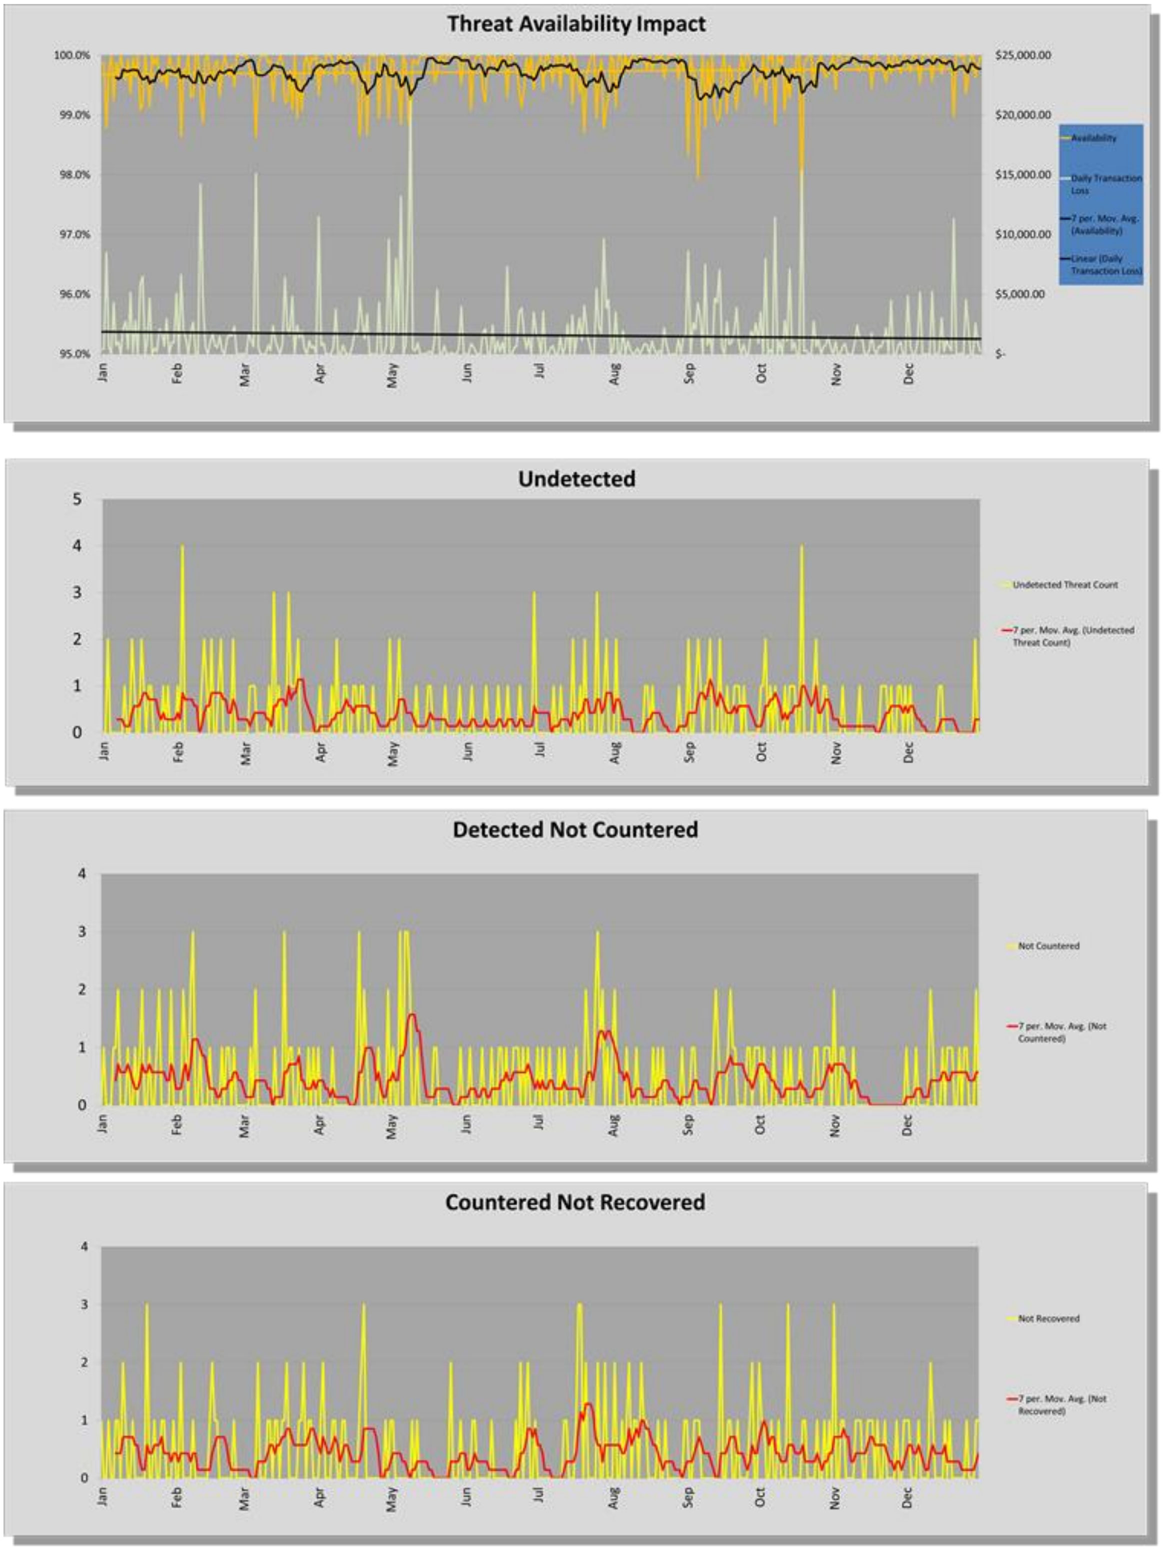

### Calculating Availability (A) Impacts

In the current model although we calculate breach events attributable to Confidentiality (C), Integrity (I) or Availability (A) impacts, we have focused solely on Availability impacts for the purposes of calculating business losses. Availability is commonly measured as a percentage of system computing capacity from zero to 100% where operating standards are typically expected to be above 95% and where ‘high availability’ systems incorporating redundancy are expected to operate in excess of 99.9% or higher. (Gray and Siewiorek 1991)Availability generally affects the throughput of computing transactions where in cases of less than 100% availability, the average user would experience a range of work slowdown or interruption effects (e.g. computational delays leading to the clock cursor, or intermittent lack of connectivity, etc.). In contrast, Integrity events are assumed to typically result in computing errors such that transactions would require rework.

To calculate the Availability effects, we take a weighted average of the individual Availability impacts of the two sets of impactful breaches: 1) Non-Detected and Detected-but-not-Countered breaches (which are assumed to have similar availability impact profiles since both are not yet countered in the system); and 2) Countered-but-not-Recovered breaches which are assumed to have a different ‘impact rate’ than undetected or non-countered breaches where countering controls halt and then diminish the effect of the breach over time (Tjoa, Jakoubi et al. 2011). The following lookup table is used to individually determine the impact on Availability of the two sets of breaches. The marginal impact decrement from 100% is an exponential function of the impact rate, where the impact rate is assumed to be half a large for breaches in Recovery as compared to undetected or detected but not countered breaches:

**Figure 88 – Availability Impact Lookup Table (Undetected and Detected-but-not-Countered Breaches)**

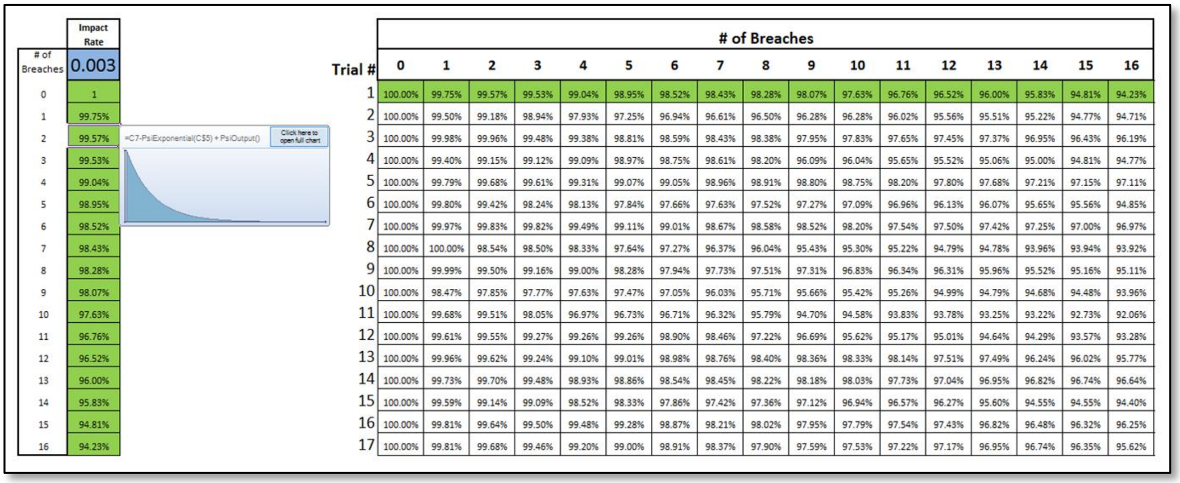

| # of Breaches | Impact Rate |
|---|---|
| | 0.003 |
| 0 | 1 |
| 1 | 99.75% |
| 2 | 99.57% |
| 3 | 99.53% |
| 4 | 99.04% |
| 5 | 98.95% |
| 6 | 98.52% |
| 7 | 98.43% |
| 8 | 98.28% |
| 9 | 98.07% |
| 10 | 97.63% |
| 11 | 96.76% |
| 12 | 96.52% |
| 13 | 96.00% |
| 14 | 95.83% |
| 15 | 94.81% |
| 16 | 94.23% |

=C7-PsiExponential(C$5) + PsiOutput()

| | # of Breaches | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Trial # | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 |
| 1 | 100.00% | 99.75% | 99.57% | 99.53% | 99.04% | 98.95% | 98.52% | 98.43% | 98.28% | 98.07% | 97.63% | 96.76% | 96.52% | 96.00% | 95.83% | 94.81% | 94.23% |
| 2 | 100.00% | 99.50% | 99.18% | 98.94% | 97.93% | 97.25% | 96.94% | 96.61% | 96.50% | 96.28% | 96.28% | 96.02% | 95.56% | 95.51% | 95.22% | 94.77% | 94.71% |
| 3 | 100.00% | 99.98% | 99.96% | 99.48% | 99.38% | 98.81% | 98.59% | 98.43% | 98.38% | 97.95% | 97.83% | 97.65% | 97.45% | 97.37% | 96.95% | 96.43% | 96.19% |
| 4 | 100.00% | 99.40% | 99.15% | 99.12% | 99.09% | 98.97% | 98.75% | 98.61% | 98.20% | 96.09% | 96.04% | 95.65% | 95.52% | 95.06% | 95.00% | 94.81% | 94.77% |
| 5 | 100.00% | 99.79% | 99.68% | 99.61% | 99.31% | 99.07% | 99.05% | 98.96% | 98.91% | 98.80% | 98.75% | 98.20% | 97.80% | 97.68% | 97.21% | 97.15% | 97.11% |
| 6 | 100.00% | 99.80% | 99.42% | 98.24% | 98.13% | 97.84% | 97.66% | 97.63% | 97.52% | 97.27% | 97.09% | 96.96% | 96.13% | 96.07% | 95.65% | 95.56% | 94.85% |
| 7 | 100.00% | 99.97% | 99.83% | 99.82% | 99.49% | 99.11% | 99.01% | 98.67% | 98.58% | 98.52% | 98.20% | 97.54% | 97.50% | 97.42% | 97.25% | 97.00% | 96.97% |
| 8 | 100.00% | 100.00% | 98.54% | 98.50% | 98.33% | 97.64% | 97.27% | 96.37% | 96.04% | 95.43% | 95.30% | 95.22% | 94.79% | 94.78% | 93.96% | 93.94% | 93.92% |
| 9 | 100.00% | 99.99% | 99.50% | 99.16% | 99.00% | 98.28% | 97.94% | 97.73% | 97.51% | 97.31% | 96.83% | 96.34% | 96.31% | 95.96% | 95.52% | 95.16% | 95.11% |
| 10 | 100.00% | 98.47% | 97.85% | 97.77% | 97.63% | 97.47% | 97.05% | 96.03% | 95.71% | 95.66% | 95.42% | 95.26% | 94.99% | 94.79% | 94.68% | 94.48% | 93.96% |
| 11 | 100.00% | 99.68% | 99.51% | 98.05% | 96.97% | 96.73% | 96.71% | 96.32% | 95.79% | 94.70% | 94.58% | 93.83% | 93.78% | 93.25% | 93.22% | 92.73% | 92.06% |
| 12 | 100.00% | 99.61% | 99.55% | 99.27% | 99.26% | 99.26% | 98.90% | 98.46% | 97.22% | 96.69% | 95.62% | 95.17% | 95.01% | 94.64% | 94.29% | 93.57% | 93.28% |
| 13 | 100.00% | 99.96% | 99.62% | 99.24% | 99.10% | 99.01% | 98.98% | 98.76% | 98.40% | 98.36% | 98.33% | 98.14% | 97.51% | 97.49% | 96.24% | 96.02% | 95.77% |
| 14 | 100.00% | 99.73% | 99.70% | 99.48% | 98.93% | 98.86% | 98.54% | 98.45% | 98.22% | 98.18% | 98.03% | 97.73% | 97.04% | 96.95% | 96.82% | 96.74% | 96.64% |
| 15 | 100.00% | 99.59% | 99.14% | 99.09% | 98.52% | 98.33% | 97.86% | 97.42% | 97.36% | 97.12% | 96.94% | 96.57% | 96.27% | 95.60% | 94.55% | 94.55% | 94.40% |
| 16 | 100.00% | 99.81% | 99.64% | 99.50% | 99.48% | 99.28% | 98.87% | 98.21% | 98.02% | 97.95% | 97.79% | 97.54% | 97.43% | 96.82% | 96.48% | 96.32% | 96.25% |
| 17 | 100.00% | 99.81% | 99.68% | 99.46% | 99.20% | 99.00% | 98.91% | 98.37% | 97.90% | 97.59% | 97.53% | 97.22% | 97.17% | 96.95% | 96.74% | 96.35% | 95.62% |

The resulting Availability impacts are uniquely calculated for each type of breach, for each simulation round, based on the number of breaches in that round:

**Figure 89 – Availability Impacts based on # of Undetected + Detected-but-not-Countered Breaches**

| | # of Breaches | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Trial # | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 |
| 1 | 100.00% | 99.75% | 99.57% | 99.53% | 99.04% | 98.95% | 98.52% | 98.43% | 98.28% | 98.07% | 97.63% | 96.76% | 96.52% | 96.00% | 95.83% | 94.81% | 94.23% |
| 2 | 100.00% | 99.50% | 99.18% | 98.94% | 97.93% | 97.25% | 96.94% | 96.61% | 96.50% | 96.28% | 96.28% | 96.02% | 95.56% | 95.51% | 95.22% | 94.77% | 94.71% |
| 3 | 100.00% | 99.98% | 99.96% | 99.48% | 99.38% | 98.81% | 98.59% | 98.43% | 98.38% | 97.95% | 97.83% | 97.65% | 97.45% | 97.37% | 96.95% | 96.43% | 96.19% |
| 4 | 100.00% | 99.40% | 99.15% | 99.12% | 99.09% | 98.97% | 98.75% | 98.61% | 98.20% | 96.09% | 96.04% | 95.65% | 95.52% | 95.06% | 95.00% | 94.81% | 94.77% |
| 5 | 100.00% | 99.79% | 99.68% | 99.61% | 99.31% | 99.07% | 99.05% | 98.96% | 98.91% | 98.80% | 98.75% | 98.20% | 97.80% | 97.68% | 97.21% | 97.15% | 97.11% |
| 6 | 100.00% | 99.80% | 99.42% | 98.24% | 98.13% | 97.84% | 97.66% | 97.63% | 97.52% | 97.27% | 97.09% | 96.96% | 96.13% | 96.07% | 95.65% | 95.56% | 94.85% |
| 7 | 100.00% | 99.97% | 99.83% | 99.82% | 99.49% | 99.11% | 99.01% | 98.67% | 98.58% | 98.52% | 98.20% | 97.54% | 97.50% | 97.42% | 97.25% | 97.00% | 96.97% |
| 8 | 100.00% | 100.00% | 98.54% | 98.50% | 98.33% | 97.64% | 97.27% | 96.37% | 96.04% | 95.43% | 95.30% | 95.22% | 94.79% | 94.78% | 93.96% | 93.94% | 93.92% |
| 9 | 100.00% | 99.99% | 99.50% | 99.16% | 99.00% | 98.28% | 97.94% | 97.73% | 97.51% | 97.31% | 96.83% | 96.34% | 96.31% | 95.96% | 95.52% | 95.16% | 95.11% |
| 10 | 100.00% | 98.47% | 97.85% | 97.77% | 97.63% | 97.47% | 97.05% | 96.03% | 95.71% | 95.66% | 95.42% | 95.26% | 94.99% | 94.79% | 94.68% | 94.48% | 93.96% |

**Figure 90 - Availability Impacts based on # of Countered-but-not-Recovered Breaches**

| | # of Breaches | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Trial # | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 |
| 1 | 100.00% | 99.76% | 99.76% | 99.61% | 99.55% | 99.53% | 99.45% | 99.29% | 99.01% | 98.90% | 98.79% | 98.79% | 98.56% | 98.27% | 98.11% | 97.92% | 97.88% |
| 2 | 100.00% | 99.93% | 99.67% | 99.48% | 99.19% | 99.07% | 99.03% | 98.67% | 98.63% | 98.60% | 98.51% | 98.43% | 97.93% | 97.92% | 97.16% | 97.15% | 97.01% |
| 3 | 100.00% | 99.88% | 99.67% | 99.40% | 98.62% | 98.53% | 98.48% | 98.45% | 98.39% | 98.32% | 98.16% | 98.12% | 98.09% | 97.76% | 97.71% | 97.63% | 97.25% |
| 4 | 100.00% | 99.78% | 99.59% | 99.39% | 99.37% | 99.36% | 98.91% | 98.88% | 98.88% | 98.84% | 98.84% | 98.73% | 98.65% | 98.46% | 98.40% | 98.40% | 98.16% |
| 5 | 100.00% | 99.84% | 99.83% | 99.79% | 99.43% | 99.19% | 98.95% | 98.71% | 98.63% | 98.24% | 98.15% | 97.95% | 97.92% | 97.78% | 97.60% | 97.59% | 97.47% |
| 6 | 100.00% | 99.99% | 99.90% | 99.29% | 99.09% | 99.01% | 98.91% | 98.13% | 97.88% | 97.81% | 97.70% | 97.51% | 97.42% | 97.27% | 97.01% | 96.95% | 96.85% |
| 7 | 100.00% | 99.98% | 99.79% | 99.66% | 99.60% | 99.30% | 99.18% | 99.17% | 99.11% | 98.84% | 98.56% | 98.50% | 98.19% | 98.07% | 97.99% | 97.73% | 97.64% |
| 8 | 100.00% | 99.96% | 99.93% | 99.89% | 99.59% | 99.52% | 99.40% | 99.35% | 99.34% | 99.23% | 99.21% | 98.96% | 98.81% | 98.76% | 98.76% | 98.75% | 98.52% |
| 9 | 100.00% | 99.79% | 99.78% | 99.75% | 99.67% | 99.64% | 99.63% | 99.52% | 98.92% | 98.88% | 98.87% | 98.85% | 98.81% | 98.72% | 98.18% | 98.16% | 98.05% |
| 10 | 100.00% | 99.96% | 99.86% | 99.82% | 99.67% | 99.39% | 99.31% | 98.80% | 98.75% | 98.74% | 98.64% | 98.52% | 98.48% | 98.11% | 97.96% | 97.65% | 97.61% |

The weighted average Availability % is then calculated based on the number of number of 'Undetected + Not Countered' and 'Countered-not-Recovered' breaches respectively:

$$\frac{(\text{Undetected} + \text{Detected but not Countered})}{\text{UnPrevented}+\text{Lag Detected not Countered}+\text{Lag Countered not Recovered}-\text{Recovered}} \times \text{Availability \%}_{\text{(Undetected + Not Countered)}} \tag{6.1}$$

$$+$$

$$\frac{(\text{Not Recovered})}{\text{UnPrevented}+\text{Lag Detected not Countered}+\text{Lag Countered not Recovered}-\text{Recovered}} \times \text{Availability \%}_{\text{(Countered not Recovered)}} \tag{6.2}$$

Here each set of breaches which can potentially impact the system contribute to a decline in the system Availability percentage, with each individual breach in each simulation round and therefore each simulation round contributing a stochastically different impact rating, even in rounds with the same breach activity level. While the cumulative impact of successful non-recovered breaches is continuously increasing, the choice of exponential function is deliberate to reflect the idea that any particular breach has an exponentially declining impact on Availability.

The impact percentage itself is arbitrary and can be scaled through the exponential function to suit the assumed severity of individual breach impacts. What is primarily important for my purposes is not the level of resulting decrement but the stochastic behaviour and distribution of the Availability measure (and corresponding loss distribution) as breaches increase since we are primarily interested in the 'tail loss' of the distribution. Here we see that for a Medium Controls system simulated over 30 individual years, the profile of the cumulative number of breaches experienced per year over the first 255 days ranges substantially from 100% availability (zero breaches, no impact) to a decrement of several percentage points below 100% depending on the number cumulative breaches[58]:

**Figure 91 – Monte Carlo Simulation of Cumulative Breaches over 255 days**

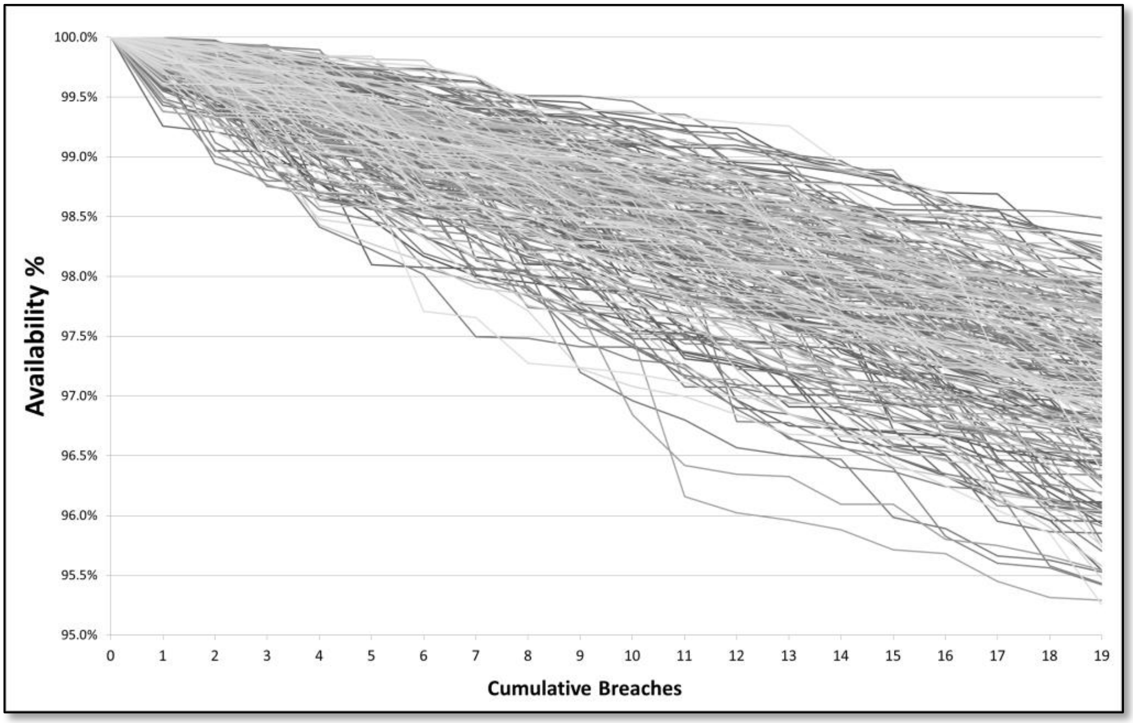

[58] 255 days represents the limit at which all 365 day Medium System simulations operated above 95% Availability. Fore a Medium Controls system, some simulations result in Availability ratings below 95% beyond 19 cumulative breaches reflecting that, because of the one period lag effect introduced to simulate non-detected/countered/recovered breaches, at lower control levels a system can become 'unstable' where availability continuously declines until reaching unacceptable levels of performance. Comparatively, for High and Very High Control systems, the number of cumulative breaches generally remains under 10 per day (i.e. most breaches are fully recovered within a day or two at most) and the resulting system Availability level does not degrade below approximately 98%.

Using the above model, we are able to simulate daily system Availability for any combination of Preventive, Blocking, Detection, Counter and Recovery controls over a single year (365 days), or for multiple years. The following figures indicates a resulting example time series of the Availability of the system and the attributed business Loss per Day based on all controls set to “High”:

**Figure 92 - High Controls Time Series (95% Control Effectiveness, 365 Days)**

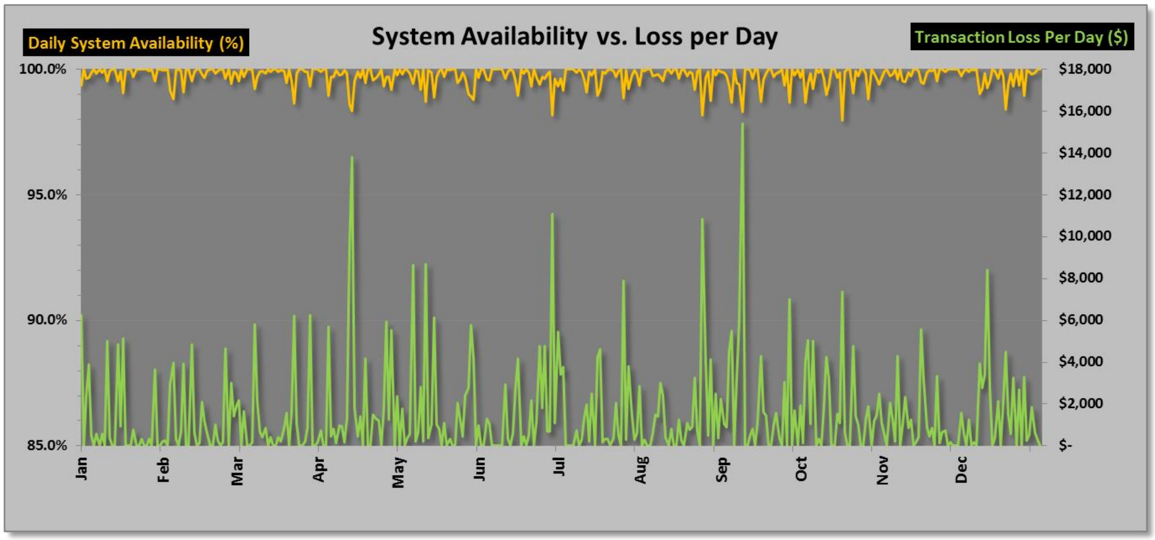

With controls set to ‘High’, system Availability generally stays above 99.8% and is mostly above 99.9%, a level of reliability that might be routinely expected in most enterprise systems. The resulting business losses per day attributed to declines in system Availability can be conveniently displayed as a histogram of the probability distribution of losses which supports further analysis and the economic comparison of control scenarios that involve the moments of the loss distribution including estimates of the average and 95% VaR and CVar:

**Figure 93 - High Controls Probability Distribution (95% Control Effectiveness, 365 Days)**

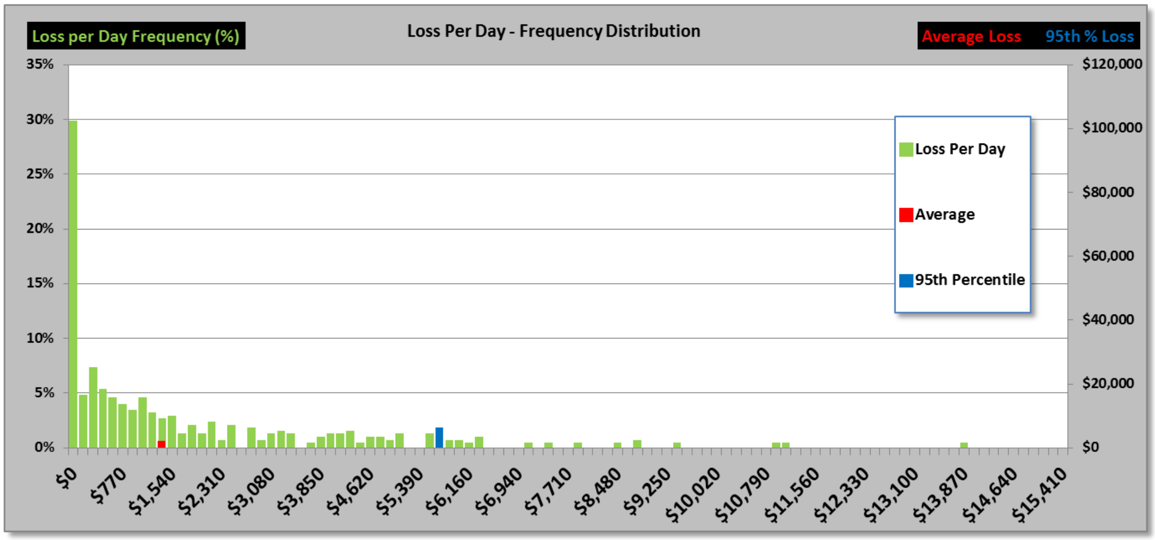

The model permits simulation of daily time series for any period between 1 year and 30 years. Expectedly, simulating the time series over 30 years produces a smoother probability distribution than doing so for a single year and reveals additional information about the underlying nature of the High Control system’s loss distribution.

For example, in the one year simulation above, the largest single daily loss was $15,410; in the 30 year simulation, the largest single loss is more than double at $39,930, although the proportion of ‘zero-loss’ (i.e. Availability = 100%) days stays constant at approximately 30%. In contrast, a ‘Low Controls’ system simulated over 30 years has an higher upper loss of $88,030, double that of the High Controls system, and has only 18% zero-loss days:

**Figure 94 – High vs. Low Controls System Probability Distributions (Equal Average Effectiveness)**

**(95% Control Effectiveness, 10,950 Days)**

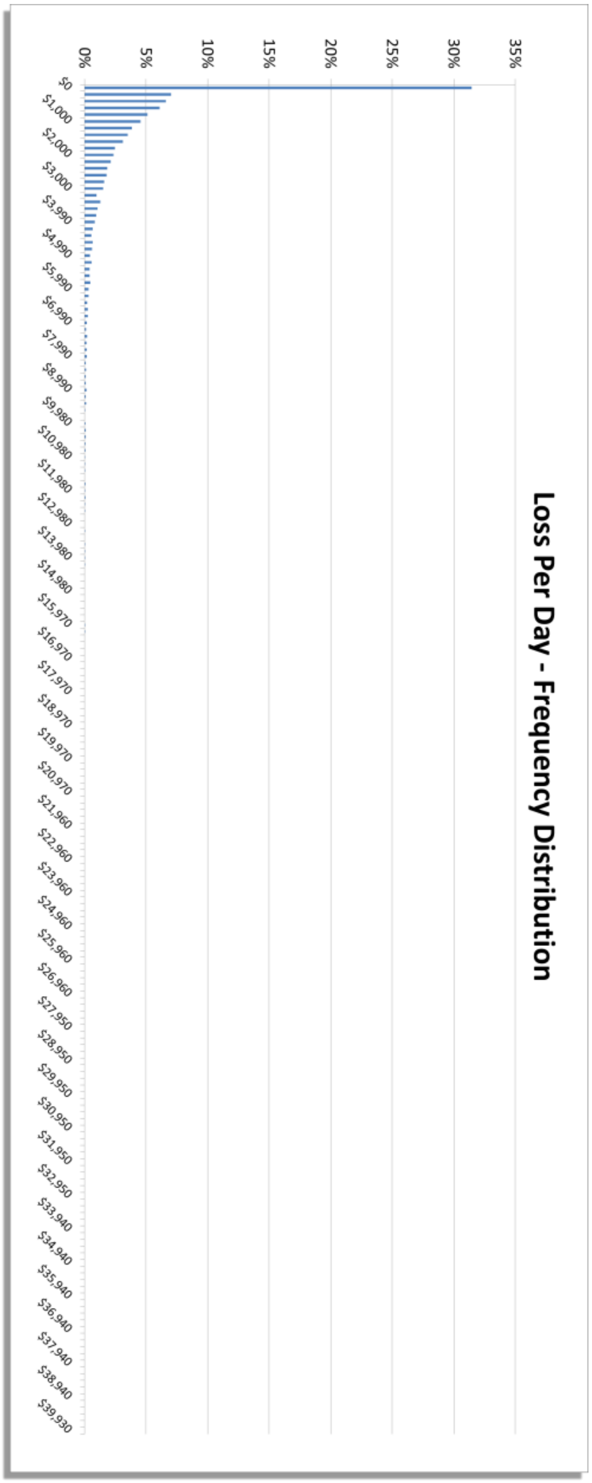

**High Controls System**

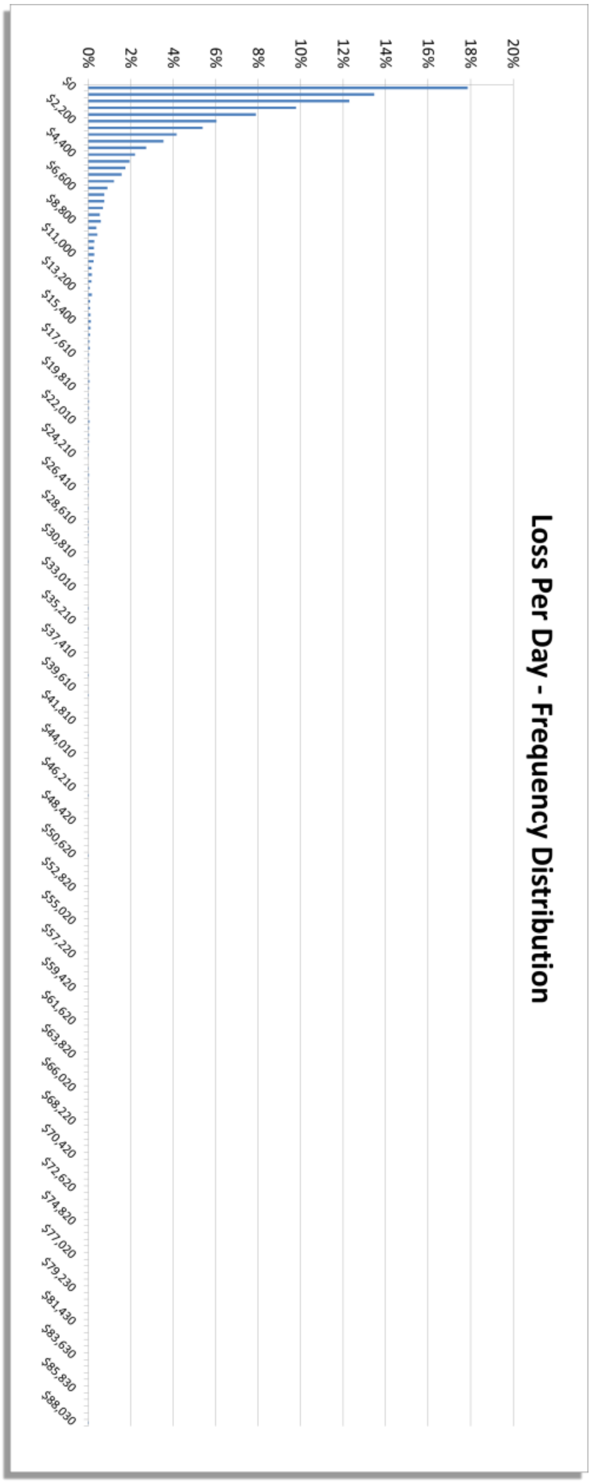

**Low Controls System**

Even more interesting from the perspective of a decision maker facing shorter term control decisions, the one year loss profile of a High Control (Figure 92, above) vs. a Low Control system (Figure 95) indicates that both systems are capable of producing similar maximum losses – around $18K and $16K for the Low and High Controls system respectively – although this level of losses is infrequent for the High Controls system and comparatively common for the Low Controls system:

**Figure 95 - Low Controls Time Series (95% Control Effectiveness, 365 Days)**

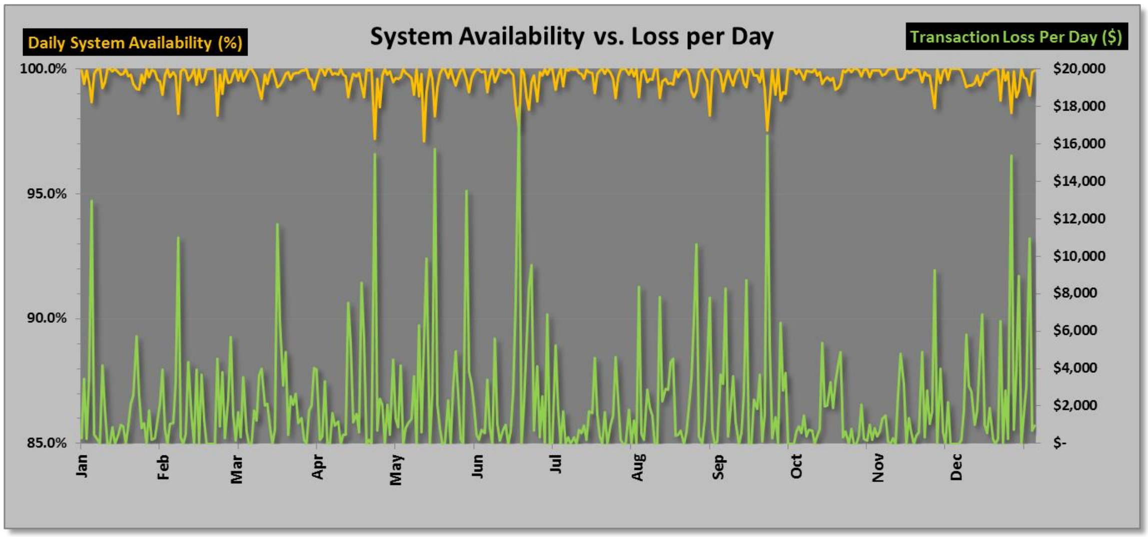

## Calibration of the Model Parameters

As noted above, the system availability and associated business loss model is clearly sensitive to the mean and deviation parameters for the stochastic inputs including control effectiveness factors for the individual controls, the number of users and user transactions per day and the attributed transaction values. In the examples above, the average effectiveness of each control has been fixed at 95% with variability of effectiveness around the average allowed to increase as control levels get weaker. This results in the overall profile of the losses being relatively similar across Control levels with most of the difference in the zero day percentage and the extreme loss end of the distribution. If we allow the average effectiveness to vary, a different pattern emerges which would be of concern to the decision maker, not surprisingly, who is more interested in the average effectiveness of the control (and therefore the average loss):

## Figure 96 - High vs. Low Controls System Probability Distributions (Different Average Effectiveness)

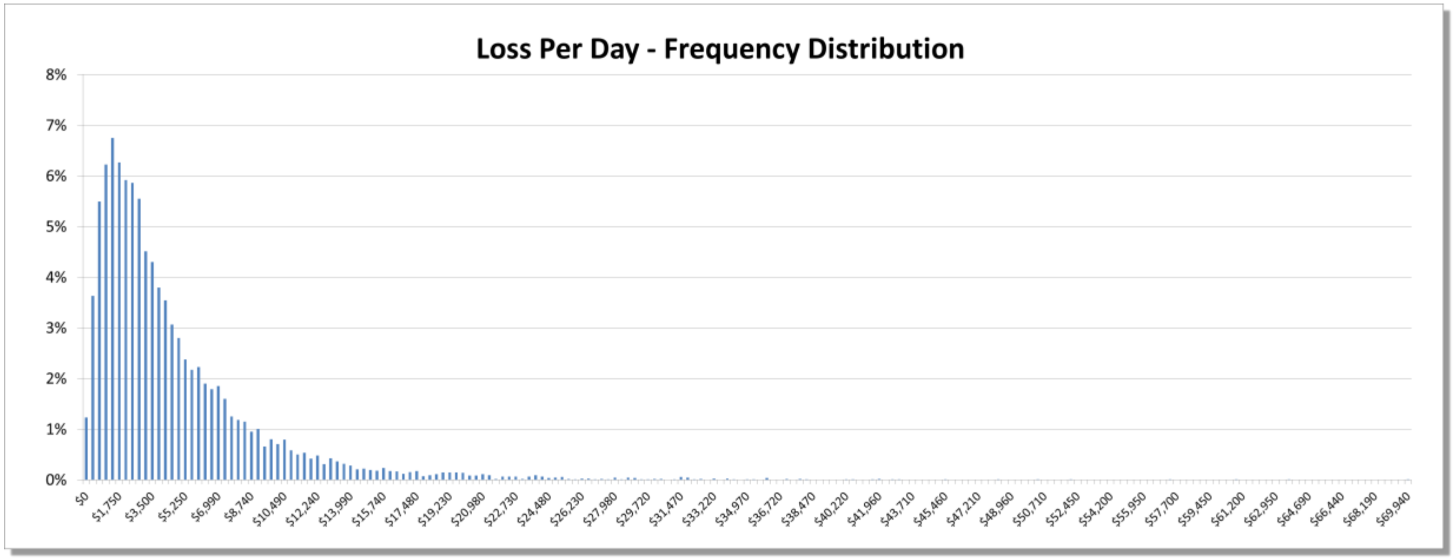

**High Controls System**
(80% Average Effectiveness, 80%-140% Variability, 10,950 days)

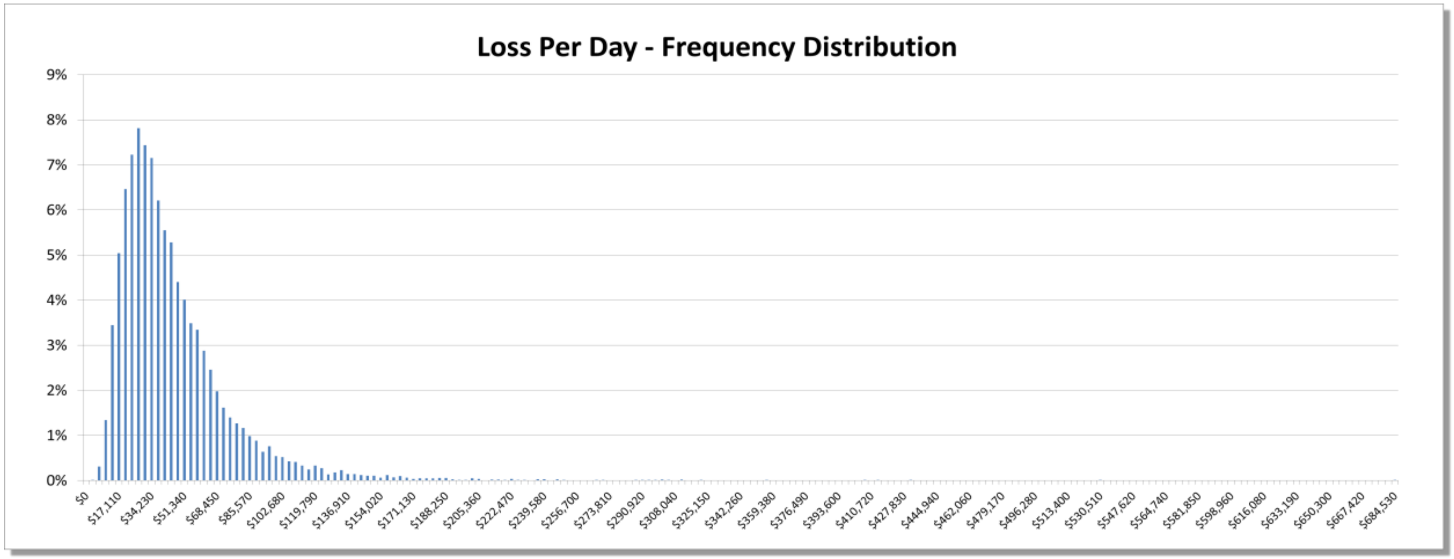

**Low Controls System**
(40% Average Effectiveness, 60%-120% Variability, 10,950 Days)

In this case the Low Controls scenario produces no zero days at all and the maximum daily loss balloons to just under 10 times the 95% effectiveness case. The High Controls case also produces many fewer zero days, although the maximum loss doesn't quite double compared to the 95% effectiveness case.

Overall, as we lower the average control effectiveness, the distribution shifts to the right (higher average losses) and the shape of the distribution, although still heavily skewed right, takes on a more Gaussian shape. The insight here is that we must be sensitive to both the absolute and the relative values of the stochastic input factors in order to characterize plausible probability distributions within and across control levels. On the other hand, it is also apparent that the dollar loss values themselves are essentially arbitrary and what is of greater importance is the relative shape of the loss distribution i.e. the moments of the distribution and particularly the extreme losses. Although more authoritative factor values could be determined, the resulting scale of the losses only materially matters in the context of this research, variously, in either the *differential* between control scenarios i.e. between prospective loss scenarios, or as may be related to some *benchmark* cost of the controls which can be thought of as a *stake* in a gamble over prospective losses. The absolute level of control costs vs. the corresponding loss levels, are otherwise being ignored since we are not proposing to *model* a cost-benefit trade-off or optimization problem for management. Rather we are interested in the way that a decision maker *perceives* the risk of the system and the cost-benefit trade-off under different framing scenarios, in controlled circumstances involving both known and uncertain loss profiles. 'Optimization' in this approach, is subjective and we allow the biases of the decision maker to be estimated rather than indicating what an otherwise risk neutral decision maker would consider to be the 'optimal' choice under risk or uncertainty.

We also note that we are explicitly focusing on attributed dollar losses (which introduces the question of appropriate business transaction value averages and variances) when system availability might actually be a better (and scale free) measure of system performance from management's point of view. System Availability is often of main concern in practice since it can be directly measurable using relatively standardized technologies, although attributing availability degradation to security related events is problematic since non-security related events can also affect system availability and performance (e.g. network congestion or component mechanical malfunction). As a measure of capacity to transact however, availability is often chosen over other performance measures as the basis of service level contracts since very high levels of availability (99.99% or higher) can be converted into both downtime measured in minutes per year and steady state or throughput transaction measurements (e.g. in systems where throughput of high volumes of real time transactions are a primary consideration such as stock exchanges or air traffic control). Notwithstanding the attractiveness of Availability as a measure of system performance, from a modelling and simulation perspective the problem of how to reflect the 'effectiveness' of security controls in maintaining system Availability remains a key issue since its logical specification and the choice of parameter values directly affects the behaviour of the model, and therefore the potential for attributed losses well before any attribution of transaction rates and value factors.

In the baseline models for the simulations used in most of the lab experiments and as represented above, the average effectiveness between controls is fixed at 95% while the variance in effectiveness is allowed to decrease and shift positively as higher levels of control are imposed. Alternatively, allowing the *average* effectiveness to vary substantially (i.e. to be substantially lower across controls than some arbitrarily high value such as 95% effective) is a questionable premise since it does not seem plausible that an enterprise operation would employ lower performing controls *on average* (i.e. if so, we would have to assume that management was either naïve to the average effect or somehow understood that buying 'lower' control levels means actually buying controls that are known to perform worse on average by *design*, and prior to deployment). To illustrate, assuming the mechanics of our logical security model are reasonably valid, the resulting Availability profile under Low Controls would likely be immediately unacceptable to most organizations and certainly unacceptable to an enterprise at the scale on which our model is proposed. For example, the Low Controls Availability profile at 60% average effectiveness looks like this:

**Figure 97 - Low Controls Time Series (60% Average Effectiveness, 365 Days)**

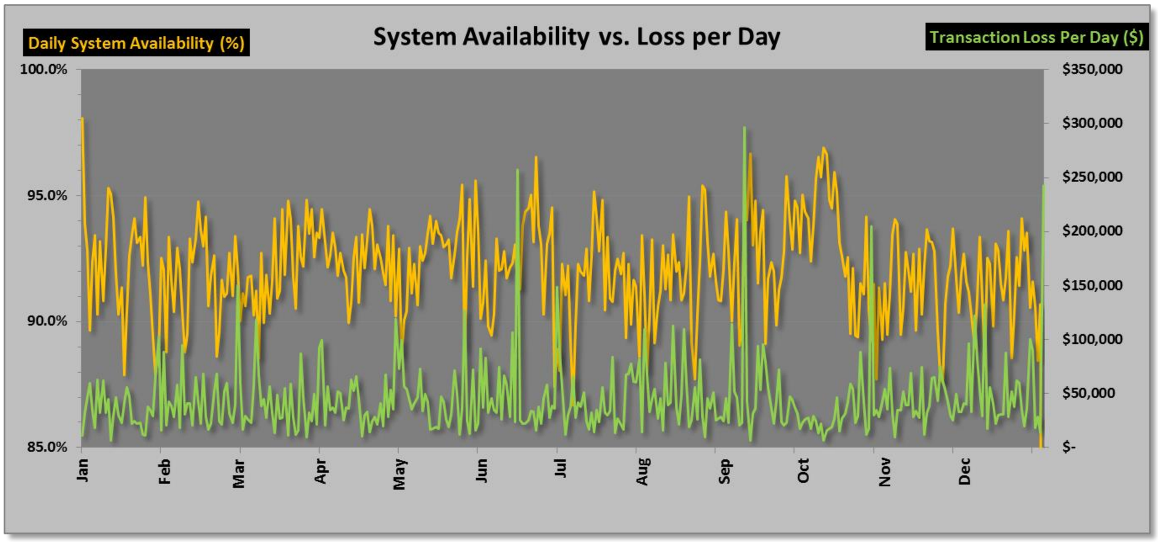

Even before the attributed losses would be of concern, the availability profile of the system itself is likely unacceptable 'on face', operating below 95% most of the time and sub-90% much of the time. Looking at the undetected, uncountered and unrecovered breach components, we see that the security system effectively never 'catches up' and always has non-zero lagged breaches. Perhaps because of this, the resulting time series' appear to be non-stationary with non-detected/countered/recovered values rising and availability decreasing overall throughout the one year period:

**Figure 98 – Low Controls Time Series – (60% Average Effectiveness, 365 days)**

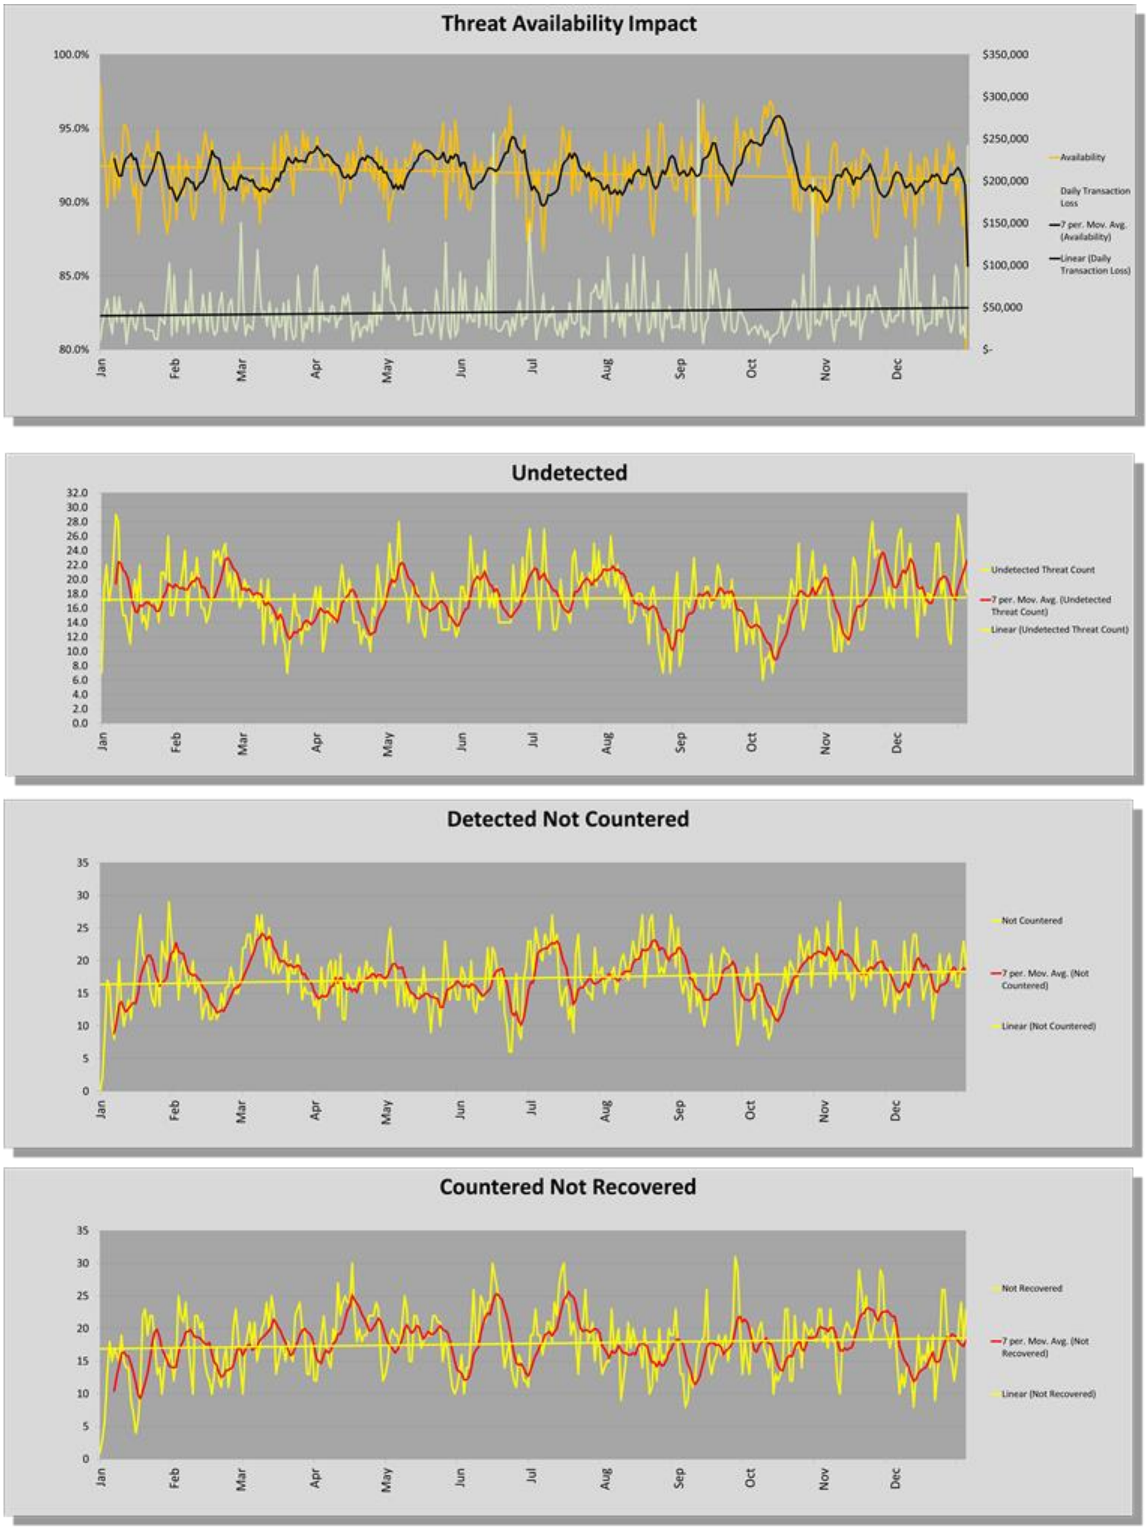

The context of the business operating environment is clearly important in this consideration: we can perhaps imagine a small business office having difficulty setting up and consistently running it's IT systems and, correspondingly, being subject to frequent and large security interruptions that would result in this kind of profile, but this would likely prompt the business (or most businesses) to simply buy 'better' controls per se, as reflected in both the average effectiveness (think advertised or vendor promised performance level) as well as in the 'deployed effectiveness' of the controls (i.e. improvements to the variance of the operational effectiveness of the control in practice).

For this reason, setting the average effectiveness between levels closer together, if not equal in value, while continuing to allow for variance of effectiveness within and across control levels (reflecting 'deployed' effectiveness) seems to make better sense from a face validity perspective and this is the approach that has been taken generally in the simulations for the lab experiments. It should also be noted that certain experiments require that the difference in control scenario losses be calibrated in any case to reflect the specific objectives of the experiment with the result that both the average and variance of control effectiveness have been manipulated to achieve specific loss scenario results. For example, one of the experiments calls for the 'Low' controls system to generate losses in excess of a benchmark 60% of the time, while the comparatively 'High' controls system produces losses above the benchmark only 40% of the time. While this may seem arbitrary or contrived from a practitioner's point of view, it enables a non-trivial decision making scenario that reflects both risk and a sufficient degree of uncertainty to enable the testing of decisional bias. Comparatively, optimization problems such as presented by Wang or Sawik, while informative of 'efficient' portfolio choices in context, either implicitly assume risk neutrality or ignore potential decisional bias altogether. My hypothesis is that, even with full information, if managers are psychologically biased then decisions will not be optimal in any case and so the degree of bias should be well understood before substantial investment is made in modelling or simulation of actual systems for control selection design purposes. The model proposed here achieves a reasonable reflection of a system at risk while enabling the framing and testing of biases across a range of decision scenarios based on risk and uncertainty.

### Attributing Business Operational Losses to Availability Impacts

As has been noted throughout, the primary objective of the simulation model is to enable the attribution of business operational losses to an IT system and the dependent business processes at risk of degraded system Availability. Once the model calculates Availability measures per day, we propose a straightforward stochastic sub-model of system use, dependent on Availability, based on the number of users and an attributed value per user transaction. The following tables illustrate the variability in user types, number of users and transactions per user by type, and the resulting range of transactions per day. Here we propose three user roles: Clinicians, Researchers and Administrators. As was done for the range of control effectiveness, the number of users of the system in each class per day, and the number of transactions per user are stochastic based on PERT distributions specified by the anticipated minimum, average and maximum. The total number of daily transactions spans from about 150,000 to well over 1 million per day with implications for daily losses driven by the scale of system use:

**Figure 99 – Daily Transactions Calculator by User Type**

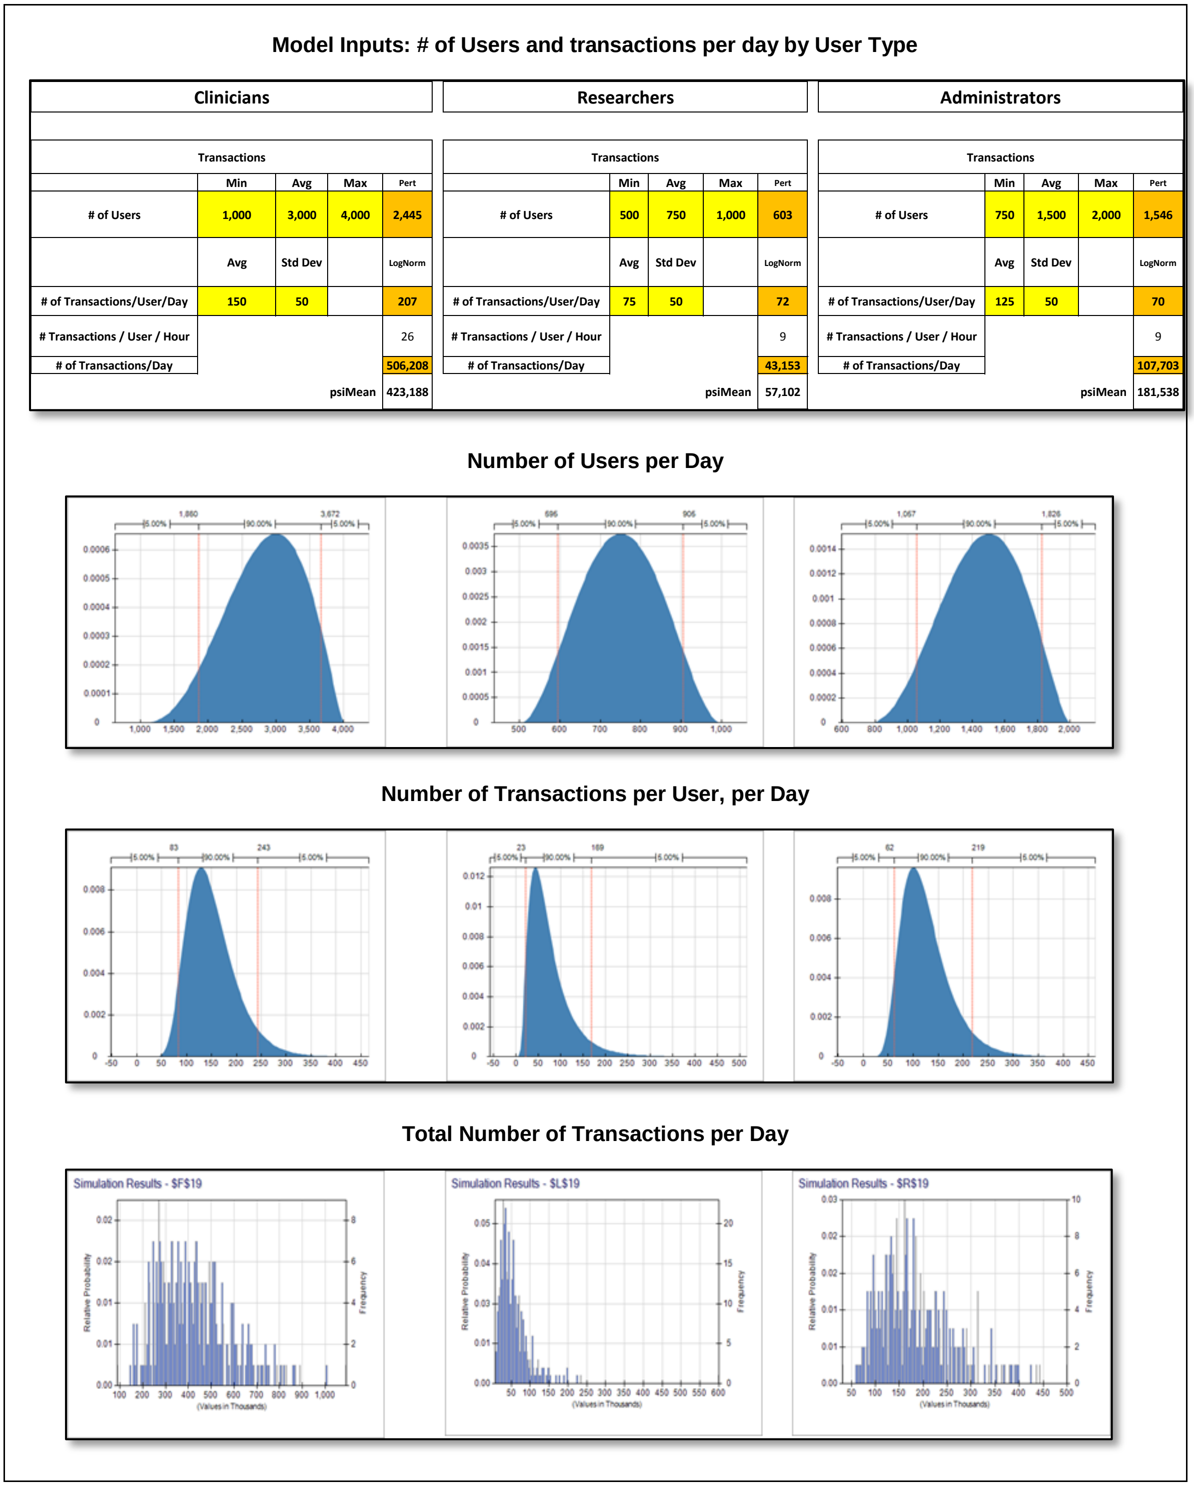

**Model Inputs: # of Users and transactions per day by User Type**

| Clinicians | | | | |
|---|---|---|---|---|
| | Transactions | | | |
| | Min | Avg | Max | Pert |
| # of Users | 1,000 | 3,000 | 4,000 | 2,445 |
| | Avg | Std Dev | | LogNorm |
| # of Transactions/User/Day | 150 | 50 | | 207 |
| # Transactions / User / Hour | | | | 26 |
| # of Transactions/Day | | | | 506,208 |
| | | | psiMean | 423,188 |

| Researchers | | | | |
|---|---|---|---|---|
| | Transactions | | | |
| | Min | Avg | Max | Pert |
| # of Users | 500 | 750 | 1,000 | 603 |
| | Avg | Std Dev | | LogNorm |
| # of Transactions/User/Day | 75 | 50 | | 72 |
| # Transactions / User / Hour | | | | 9 |
| # of Transactions/Day | | | | 43,153 |
| | | | psiMean | 57,102 |

| Administrators | | | | |
|---|---|---|---|---|
| | Transactions | | | |
| | Min | Avg | Max | Pert |
| # of Users | 750 | 1,500 | 2,000 | 1,546 |
| | Avg | Std Dev | | LogNorm |
| # of Transactions/User/Day | 125 | 50 | | 70 |
| # Transactions / User / Hour | | | | 9 |
| # of Transactions/Day | | | | 107,703 |
| | | | psiMean | 181,538 |

**Number of Users per Day**

**Number of Transactions per User, per Day**

**Total Number of Transactions per Day**

After calculating the number of transactions per user type, the model then attributes a business value per transaction for each user type. Here the average value of a clinician's transaction is deemed to be the highest, followed by research and administrative use. The nominal per transaction value, per day, by user type is then made stochastic using a log-normal specification based on a specified average and standard deviation:

**Figure 100 – System Per-Transaction Value Calculation (Research User Example)**

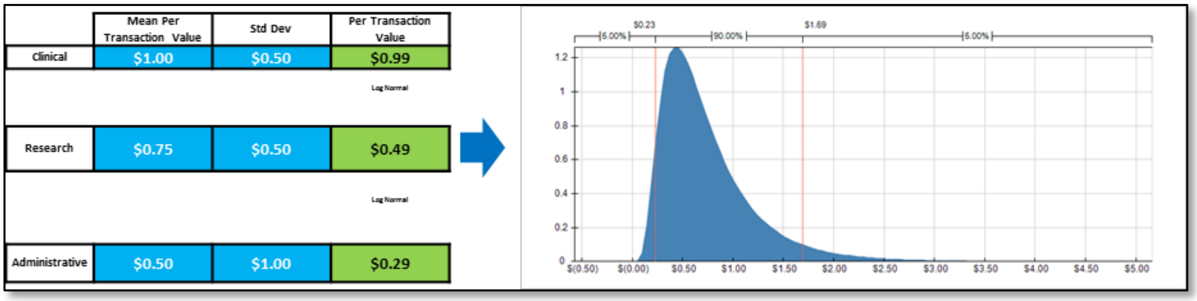

| | Mean Per Transaction Value | Std Dev | Per Transaction Value |
|---|---|---|---|
| Clinical | $1.00 | $0.50 | $0.99 |
| | | | Log Normal |
| Research | $0.75 | $0.50 | $0.49 |
| | | | Log Normal |
| Administrative | $0.50 | $1.00 | $0.29 |

The model then multiplies the number of transactions and the value per transaction with the previously calculated decrement in system Availability for each day to calculate the daily value of losses attributed to system Availability impacts:

**Figure 101 – Daily Loss Calculation – All Users**

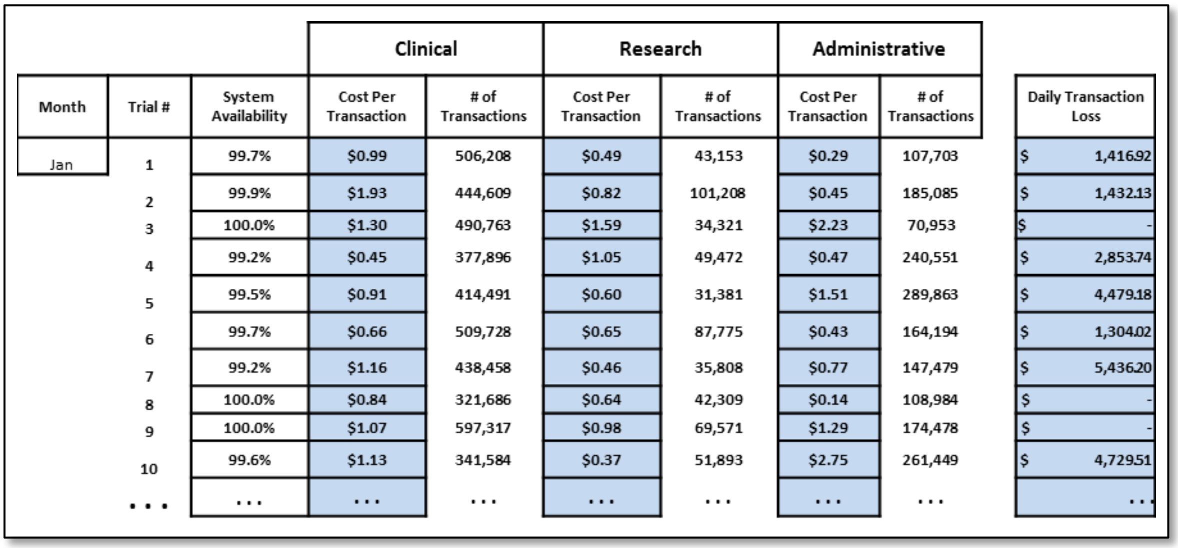

| | | | Clinical | | Research | | Administrative | | |
|---|---|---|---|---|---|---|---|---|---|
| Month | Trial # | System Availability | Cost Per Transaction | # of Transactions | Cost Per Transaction | # of Transactions | Cost Per Transaction | # of Transactions | Daily Transaction Loss |
| Jan | 1 | 99.7% | $0.99 | 506,208 | $0.49 | 43,153 | $0.29 | 107,703 | $ 1,416.92 |
| | 2 | 99.9% | $1.93 | 444,609 | $0.82 | 101,208 | $0.45 | 185,085 | $ 1,432.13 |
| | 3 | 100.0% | $1.30 | 490,763 | $1.59 | 34,321 | $2.23 | 70,953 | $ - |
| | 4 | 99.2% | $0.45 | 377,896 | $1.05 | 49,472 | $0.47 | 240,551 | $ 2,853.74 |
| | 5 | 99.5% | $0.91 | 414,491 | $0.60 | 31,381 | $1.51 | 289,863 | $ 4,479.18 |
| | 6 | 99.7% | $0.66 | 509,728 | $0.65 | 87,775 | $0.43 | 164,194 | $ 1,304.02 |
| | 7 | 99.2% | $1.16 | 438,458 | $0.46 | 35,808 | $0.77 | 147,479 | $ 5,436.20 |
| | 8 | 100.0% | $0.84 | 321,686 | $0.64 | 42,309 | $0.14 | 108,984 | $ - |
| | 9 | 100.0% | $1.07 | 597,317 | $0.98 | 69,571 | $1.29 | 174,478 | $ - |
| | 10 | 99.6% | $1.13 | 341,584 | $0.37 | 51,893 | $2.75 | 261,449 | $ 4,729.51 |
| | • • • | … | … | … | … | … | … | … | … |

## Discussion of Simulation Model Results

The simulation model presented in this section generates realistic daily business losses attributed to information security breach events based on the effectiveness of the security controls at preventing, detecting, countering and recovering from successful attacks at various levels of control. The resulting pattern of generated losses would be reasonably familiar to security practitioners regardless of the scale of the enterprise where the typical distribution of losses is bounded by zero on the left and heavily skewed to the right. The distributional characteristics of losses is a key feature of stochastic systems generally, where the higher moments of the distribution are of particular concern to decision makers seeking to avoid catastrophic loss (Acerbi and Tasche 2002; Krokhmal 2007). The losses are attributed to degradation in system ‘availability’, a measure of performance which practitioners directly seek to maximize in operation. Availability is measured as a percentage of maximal availability (100%), where incremental degradation negatively affects users’ ability to transact with the system at a normal rate either through interruption or incompletion of transactions or through a slowing of individual transaction completion time. ‘High’ or ‘Very High’ levels of overall control are expected to result in system Availability levels that meet 99.9% or greater availability levels (less than 9 hours of downtime per year) and which are typically a formal ‘service level requirement’ for large scale enterprises in which hundreds of thousands to millions of individual user system transactions are undertaken daily.

The logical relationship between the individual hardware and software components is based on an expert validated security model and is extensible in terms of the number and architecture of the components that comprise the overall ‘system’. The model produces a time series of events reflecting a ‘risk aware business process’ that is able to prevent, detect and recover from security breaches (Tjoa, Jakoubi et al. 2011). The security ‘threat and risk’ model for each logical component reflects industry accepted profiles for individual attacks and component vulnerabilities and utilizes a conditional, probabilistic approach to the determination of successful attacks on components, resulting in an inherently stochastic profile of component control effectiveness and therefore of attack success likelihood (Sommestad, Ekstedt et al. 2013). Since we are generally interested in enterprise-level effects, consideration for the variability of day-to-day business operations that are dependent on the system at risk have been taken into account, where the type and number of users, the number of transaction per user and the attributed value per transaction are allowed to vary stochastically. This permits the scaling of the model inputs and impacts to suit the business of concern, either in terms of the number of users or the attributed value of the transactions affected by system performance degradation. Attributed business losses have been restricted to those arising from system availability risks where the impact to transaction throughput is reasonably estimated based largely on the nominal availability level. In practice, significant business losses would also potentially be experienced due to both system Integrity and Confidentiality breaches, however the attribution of these types of losses requires further assumptions regarding, for example, business process rework for Integrity impacts or costs associated with breach recovery, reputational effects and direct or indirect costs (e.g.

litigation) over of loss of privacy for Confidentiality impacts (Thomas 2009). It should be noted that Availability risks may also feature all of these additional cost dimensions, however for my purposes Availability can be most directly modelled purely as a function of preventive, detective and recovery controls. In addition, there is a recognition that Confidentiality risks may be ultimately more appropriately compensated for by the use of insurance rather than self-protection where the anticipated impacts are not based primarily on the ability to transact but rather on breach recovery or compensation to affected individuals to whom the breached information pertains (Acquisti, Friedman et al. 2006).

For the purposes of this research, I also required a simulation model that reflected both these practitioner expectations, but which also allowed for the simulation of specific scenarios that controlled for specific quantitative performance differences between system control levels in order to support the decision making Games presented in the following sections. In Game 2, for example, where we seek to test individual subjective probability estimates through Bayesian updating, the user bets on which of two systems has produced a non-representative sample of losses. In this scenario, the relatively ‘Low Controls’ system was required to generate losses above $10,000 60% of the time, whereas the relatively ‘High Controls” system generates losses above $10,000 40% of the time. Similarly in Game 3, where the probability of losses across systems is unknown to the Participant, the systems were required to have specific loss percentages above a certain level. In Game 5 where we seek to test the willingness to buy cyber insurance and precautionary controls, the experiment also required the percentage of losses above a common insurable level to be specific across control levels. By tuning both the average and variance of control effectiveness in both systems, I was able to successfully simulate the required profiles in each case. The simulation is therefore both a practical tool for the demonstration of stochastic system security performance characteristics and a platform for generating controlled experiments. On this basis, I am able to propose the following experiments which (for Games 2 through 4) depend on system simulation to generate prospective loss scenarios in which Participants make context relevant choices under both risk (known probabilities) and uncertainty (unknown probabilities) of system performance.

# 7 - Methodological Motivation for Experiments – Overview of Experimental Approaches

The majority of the experiments undertaken for this research introduce the simulation of a ‘system at risk’ in which Participants make bets on prospective loss outcomes with varying types and degrees of control. A major feature of the simulation experiments is that the subject is not directly presented with the probabilities of risk occurrence and must subjectively infer these from the presented qualitative and quantitative data. This allows us to test both risk preference and risk *perception*: for example, whether the subject is using Bayes rules to accurately form posterior probabilities of risk and to what extent the strength and weight of the evidence presented affects the subjective probability estimates (Griffin and Tversky 1992; Antoniou, Harrison et al. 2015).

**Experimental Set-up**

Originally motivated by the experimental setup of ‘virtual reality’ simulation experiments which were applied to economic decision making for homeowner protection against fire damage (Fiore, Harrison et al. 2009; Sen 2010; George, Harrison et al. 2012). I have proceeded to correspondingly modify the simulation aspects of this and other decision experiments for a security risk management context while retaining Harrison’s overall approach for lab experiments which generate the requisite data on which to undertake the structural modeling outlined for each experiment below. To my knowledge this is the first study to undertake this type of simulation in support of decision experiments over security controls

Sen’s approach, for example, is an appropriate motivational analog for the security control context because the essential feature of the core preference elicitation experiment (aside from the simulation of a particular system at risk) is the ability for the subject to alter the risk profile of the system generating the lottery results in order to optimize the subjective economic outcomes. In Sen’s experiment, the subject decides whether and when to purchase a type of fire prevention control which prevents severity (i.e. impact, but not frequency) of a prospective wildfire and the resulting likelihood that their ‘in-game’ house burns down. The control selection model is based on *endogenous* risk models which more generally recognize the role of ‘self- protection’ and insurance (including ‘self-insurance’) in situations where the decision maker can influence the risk profile of the system (Ehrlich and Becker 1972; Shogren and Crocker 1991). Whereas Sen specifically designed decision making experiments in order to investigate certain aspects of subjective utility and multi-period decision making (McKee, Berrens et al. 2004; Talberth, Berrens et al. 2006; Fiore, Harrison et al. 2009), I further modify other Harrison-sponsored research as noted below to incorporate domain relevant control over both likelihood (self-protection) and impact (self-insurance) of a security incident across multiple control factors reflecting the multi-attribute nature of the security control problem.

There are several common features of the resulting experiments that draw motivation directly from the reviewed literature:

**First**, the subject is exposed generally to the factors that comprise the system simulation model and how the various combinations of threats, vulnerabilities and possible controls increase or decrease both the probability and impact on availability and the attributed business losses as discussed above.

**Second**, some type of graphical representation of business losses is presented and compared to representative control selection scenarios. This overall approach was fundamentally motivated by Harrison's numerous lab studies using a 'virtual reality' simulation of a forest fire in which the Participant variously chooses to protect by buying a fire reduction control (Fiore, Harrison et al. 2009). In Sen's experiments, for example, aside from the 3D computer representation of a simulated fire resulting in the house burning (or not burning), the subject is also shown a histogram of the percentage of acreage burned for the 48 possible simulation factors the system is able to model, and indicating the relative impact reduction effect of implementing a specific type of fire control. The example is clearly analogous to a security control selection decision under uncertainty where the general effect of the control is to improve the chances of no catastrophic loss, but not with certainty:

**Figure 102 – Example Simulation Outcomes represented as Histograms**

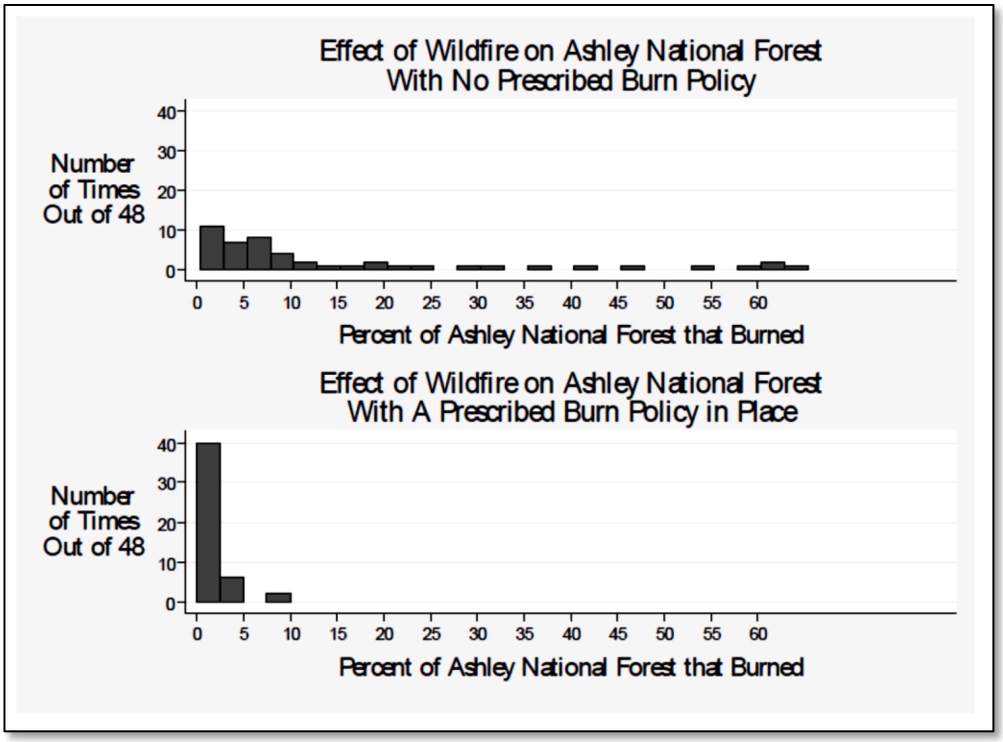

(Sen 2010)

**Third**, the subject is then exposed to simulations representing two control and two no-control scenarios. This is done to familiarize the subject to discrete outcomes for the simulated model over a single period. For example, the model is simulated assuming a random set of associated threats and vulnerabilities grouped into qualitative ranges (e.g. low, medium, high) and also simulated at these risk levels either with or without a control in place. This is done to provide context for the subsequent control choice without revealing the actual underlying probabilities for the threats, vulnerabilities and efficacy of the controls.

**Fourth** and finally, the subject is presented with betting scenarios in which the objective risk of the system may or may not be stated up front and the subject is asked to place bets as to whether there will be a loss above a certain threshold outcome. The betting device is generally a multiple price list (MPL) representing ranged payoffs from ‘bookies’ offering different odds on the occurrence of ‘catastrophic’ or some other arbitrary level of loss (e.g. the percentage of losses above a certain dollar level or the dollar level of losses representing a certain percentage of outcomes) where the odds offered are the inverse of the expected probability of the outcome (Harrison and Rutström 2008). Alternatively, Participants may be asked whether they elect to purchase a control at a range of offered prices. Once the betting selections are made, the system is simulated under low and high threat risk assumptions and the outcomes are determined by the bets made in each case or the losses incurred at the level of system control purchased.

The following table indicates Sen’s payouts based on the High risk scenario of fire burning the house:

**Figure 103 – Example MPL Betting Table**

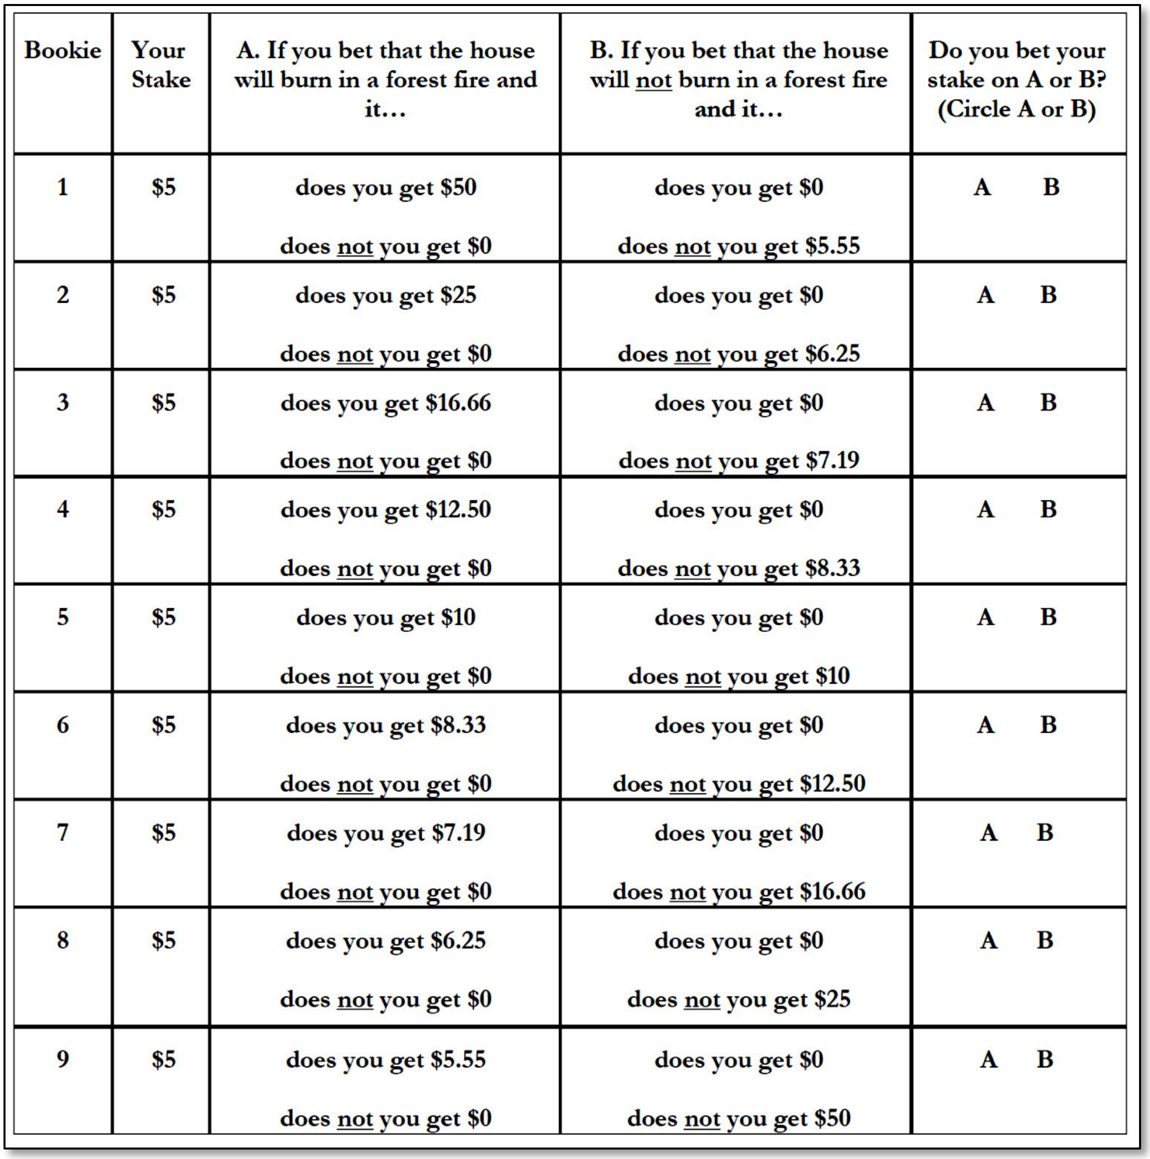

| Bookie | Your Stake | A. If you bet that the house will burn in a forest fire and it… | B. If you bet that the house will not burn in a forest fire and it… | Do you bet your stake on A or B? (Circle A or B) |
|---|---|---|---|---|
| 1 | $5 | does you get $50<br>does not you get $0 | does you get $0<br>does not you get $5.55 | A B |
| 2 | $5 | does you get $25<br>does not you get $0 | does you get $0<br>does not you get $6.25 | A B |
| 3 | $5 | does you get $16.66<br>does not you get $0 | does you get $0<br>does not you get $7.19 | A B |
| 4 | $5 | does you get $12.50<br>does not you get $0 | does you get $0<br>does not you get $8.33 | A B |
| 5 | $5 | does you get $10<br>does not you get $0 | does you get $0<br>does not you get $10 | A B |
| 6 | $5 | does you get $8.33<br>does not you get $0 | does you get $0<br>does not you get $12.50 | A B |
| 7 | $5 | does you get $7.19<br>does not you get $0 | does you get $0<br>does not you get $16.66 | A B |
| 8 | $5 | does you get $6.25<br>does not you get $0 | does you get $0<br>does not you get $25 | A B |
| 9 | $5 | does you get $5.55<br>does not you get $0 | does you get $0<br>does not you get $50 | A B |

(Sen 2010)

In Sen's example, the following table indicates the prospective outcomes based on the subject's *subjective* perception of the risk of the house burning down (outcome A) to be 75% or, conversely, the house not burning down (outcome B) at 25%:

**Figure 104 - Betting Task with Stake of $1 (subject believes A will occur with probability 0.75)**

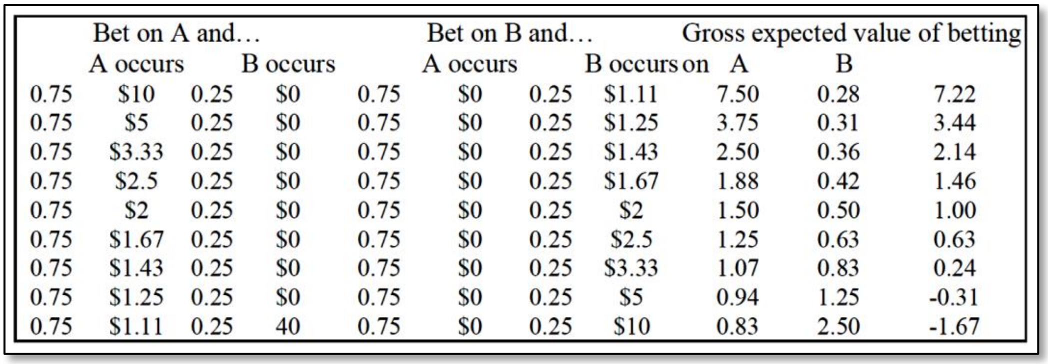

| | Bet on A and… | | | | Bet on B and… | | | Gross expected value of betting | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | A occurs | | B occurs | | A occurs | | B occurs | on A | B | |
| 0.75 | $10 | 0.25 | $0 | 0.75 | $0 | 0.25 | $1.11 | 7.50 | 0.28 | 7.22 |
| 0.75 | $5 | 0.25 | $0 | 0.75 | $0 | 0.25 | $1.25 | 3.75 | 0.31 | 3.44 |
| 0.75 | $3.33 | 0.25 | $0 | 0.75 | $0 | 0.25 | $1.43 | 2.50 | 0.36 | 2.14 |
| 0.75 | $2.5 | 0.25 | $0 | 0.75 | $0 | 0.25 | $1.67 | 1.88 | 0.42 | 1.46 |
| 0.75 | $2 | 0.25 | $0 | 0.75 | $0 | 0.25 | $2 | 1.50 | 0.50 | 1.00 |
| 0.75 | $1.67 | 0.25 | $0 | 0.75 | $0 | 0.25 | $2.5 | 1.25 | 0.63 | 0.63 |
| 0.75 | $1.43 | 0.25 | $0 | 0.75 | $0 | 0.25 | $3.33 | 1.07 | 0.83 | 0.24 |
| 0.75 | $1.25 | 0.25 | $0 | 0.75 | $0 | 0.25 | $5 | 0.94 | 1.25 | -0.31 |
| 0.75 | $1.11 | 0.25 | 40 | 0.75 | $0 | 0.25 | $10 | 0.83 | 2.50 | -1.67 |

(Sen 2010)

In the above example, the (risk neutral) subject would be expected to bet on A for every bookie offering odds that corresponded to a lower house probability than 0.75 of A winning, and then switch over to bet on B for every bookie offering odds that corresponded to a higher house probability than 0.75 of A winning. In the above example at the expected burn rate of 75%, a risk neutral subject would bet on event A for the first 7 bookies and then switch to event B.

In my version of this 2-part experiment (**Game 3**) presented below, the subject is fist provided with an initial stake amount and is asked whether they wish to prospectively purchase a control (Lottery A) across a range of incremental control values ranging from zero up to the value of the catastrophic loss under two separate relatively high and low stake/impact simulation scenarios. One of the control levels is then chosen at random. If the subject chose to pay for the control at that level, then the cost of the control is deducted from their stake and the control is applied to the simulation, otherwise no control is applied, they keep their entire stake, and the system is then simulated with or without the control in place. The following Table illustrates the payoff matrix that is implied by Sen’s fire simulation, assuming that the subject accurately infers the true probabilities of his own property burning when choosing the control (6%). Again, the subject does not know this percentage and must subjectively infer it from the exposure to the initial simulations of high and low fire risk scenarios:

**Figure 105 - Inferred WTP Instrument when Initial Credit is $40 and the House is Valued at $18**

| | Lottery A (control) | | | | Lottery B (no control) | | | | EVA | EVB | Difference |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Control Cost | p(burn) | | p(safe) | | p(burn) | | p(safe) | | | | |
| $0 | 0.06 | -18 | 0.94 | 0 | 0.29 | -18 | 0.71 | 0 | -1.08 | -5.22 | 4.14 |
| $2 | 0.06 | -20 | 0.94 | -2 | 0.29 | -18 | 0.71 | 0 | -3.08 | -5.22 | 2.14 |
| $4 | 0.06 | -22 | 0.94 | -4 | 0.29 | -18 | 0.71 | 0 | -5.08 | -5.22 | 0.14 |
| $6 | 0.06 | -24 | 0.94 | -6 | 0.29 | -18 | 0.71 | 0 | -7.08 | -5.22 | -1.86 |
| $8 | 0.06 | -26 | 0.94 | -8 | 0.29 | -18 | 0.71 | 0 | -9.08 | -5.22 | -3.86 |
| $10 | 0.06 | -28 | 0.94 | -10 | 0.29 | -18 | 0.71 | 0 | -11.08 | -5.22 | -5.86 |
| $12 | 0.06 | -30 | 0.94 | -12 | 0.29 | -18 | 0.71 | 0 | -13.08 | -5.22 | -7.86 |
| $14 | 0.06 | -32 | 0.94 | -14 | 0.29 | -18 | 0.71 | 0 | -15.08 | -5.22 | -9.86 |
| $16 | 0.06 | -34 | 0.94 | -16 | 0.29 | -18 | 0.71 | 0 | -17.08 | -5.22 | -11.86 |
| $18 | 0.06 | -36 | 0.94 | -18 | 0.29 | -18 | 0.71 | 0 | -19.08 | -5.22 | -13.86 |

(Sen 2010)

In the above example, the subject indicates whether they want to purchase control at each dollar cost between $0 and $18 in $2 increments. The researcher then randomly determines that the $4 control scenario will be played out, and it happens that the subject elected to spend $4 for control (in fact, it happens that they elected to spend at each available control level, since they subjectively desired to have control under any possible scenario). The resulting Lottery A bet in the 3rd line (green highlight), indicating selection of the 'control choice'  in Sen's experiment results in a 6% probability of the house burning down (at a cost of $18 for the house) and a 94% chance of nothing happening, but cost $4 to purchase in either case. Conversely, had the subject not selected control at the $ level, the resulting Lottery B in line 3 increases the probability of the house burning down to 29% (at a cost of $18 for the house), but with no loss whatsoever (for house or control cost) 71% of the time.

As the core motivational experiment involving simulation with control, I find this treatment somewhat unrealistic for the security context: presumably choosing control (at any dollar level other than zero) would have some *stochastic* effect on the security outcome, even if 'threshold' mitigation effects are applicable for certain types of controls (i.e. in order for the control to be minimally effective, it has to be deployed above a certain dollar cost – employee training for example). In my version of Sen's experiment, the subject will be presented with a one-shot decision of whether to implement a security control. For my research purposes, the control effectiveness will therefore be incorporated into the system model as a stochastic effect i.e. as long as the subject picks some level of control above zero, the control will have some mitigation effect on the simulation, although without certainty. This treatment is more realistically analogous to budgeting for a specific control level, implementing the control and then subsequently finding out whether the level of control chosen was effective for the purpose.

Game 3 introduces risk endogeneity in the form of a one-period control choice over a prospective loss, but in this case with known probabilities. In my version of Sen's experiment, the subject will be presented with a one-shot decision of whether to implement a security control which will, for example, absolutely (i.e. deterministically) lower productivity but will eliminate the entire impact of a prospective security incident. In Sen's experiment, the subject chooses between either Lottery A (a relatively safe bet) versus Lottery B (a relatively risky prospect) with known probabilities for each of 9 bets for each variation in task:

**Inferred WTP Instrument when Initial Credit is $40 and the House is Valued at $18** (Sen 2010)

| **Lottery A (control)** | | | | **Lottery B (no control)** | | | | **EVA** | **EVB** | **Difference** |
|---|---|---|---|---|---|---|---|---|---|---|
| **p(burn)** | | **p(safe)** | | **p(burn)** | | **p(safe)** | | | | |
| 0.06 | -18 | 0.94 | 0 | 0.29 | -18 | 0.71 | 0 | -1.08 | -5.22 | 4.14 |
| 0.06 | -20 | 0.94 | -2 | 0.29 | -18 | 0.71 | 0 | -3.08 | -5.22 | 2.14 |
| 0.06 | -22 | 0.94 | -4 | 0.29 | -18 | 0.71 | 0 | -5.08 | -5.22 | 0.14 |
| 0.06 | -24 | 0.94 | -6 | 0.29 | -18 | 0.71 | 0 | -7.08 | -5.22 | -1.86 |
| 0.06 | -26 | 0.94 | -8 | 0.29 | -18 | 0.71 | 0 | -9.08 | -5.22 | -3.86 |
| 0.06 | -28 | 0.94 | -10 | 0.29 | -18 | 0.71 | 0 | -11.08 | -5.22 | -5.86 |
| 0.06 | -30 | 0.94 | -12 | 0.29 | -18 | 0.71 | 0 | -13.08 | -5.22 | -7.86 |
| 0.06 | -32 | 0.94 | -14 | 0.29 | -18 | 0.71 | 0 | -15.08 | -5.22 | -9.86 |
| 0.06 | -34 | 0.94 | -16 | 0.29 | -18 | 0.71 | 0 | -17.08 | -5.22 | -11.86 |
| 0.06 | -36 | 0.94 | -18 | 0.29 | -18 | 0.71 | 0 | -19.08 | -5.22 | -13.86 |

In the above scenario, risk neutral subjects would be expected to forego the control for the first 3 bets and switch to the control choice in the fourth line and thereafter.

A testable hypothesis with this treatment is that the coefficient of risk aversion is equal in both the standard lottery loss task with *exogenous* risk and here where risk is made *endogenous* (i.e. controllable by the subject). Rejection of the hypothesis would indicate that some type of 'risk framing' effect is present when risk is presented endogenously (Shogren and Crocker 1991).

**Experiments, Treatments, Hypotheses and Models**

The following table indicates the motivating reference papers, my corresponding hypotheses and the resulting experiments undertaken to generate choice data that is appropriate for the estimation of various latent models of choice and the estimation of model parameters that reflect the indicated range of decisional biases in the participants' decision making processes:

**Table 8 - Experiment Treatments and Hypotheses (Games 1 – 5)**

| | |
|---|---|
| **Experiment # 1 – Choice Behaviour and Asset Integration:**<br>**Motivating Reference Paper:**<br>Andersen, S., G. W. Harrison, et al. (2006). Choice Behaviour, Asset Integration and Natural Reference Points. Working Paper 06-07, Department of Economics, College of Business Administration, University of Central Florida. | |
| **In Experiment 1**, participants are asked to choose between a risky lottery and a series of 'sure money' gain, loss and mixed gain/loss payoffs over 20 rounds, where the participant accumulates winnings over the course of the experiment. This experiment directly replicates the Andersen experiment and tests whether managers integrate a cumulative wealth reference point when making marginal decisions over both gain and loss prospects with objectively known risk. Hypotheses include: whether decision makers integrate accumulated wealth into marginal decisions; testing for probability weighting; and undertaking a finite mixture model which estimates weighting between EUT and PT latent data generating processes. | |
| **Treatment 1: Multi-period gain/loss tasks with asset integration**<br>Exogenous gain/loss/mixed gain and loss risk frame, without system simulation; tests for asset integration in the context of cumulative payout after 20 rounds of choice.<br>• **Hypothesis 6:** Prior outcomes affect a manager's current decisions.<br>• **Hypothesis 9:** Managers integrate choices over accumulated gains or losses.<br>• **Hypothesis 10:** A manager's reference point for risky choices is not the marginal value of the current choice. | **Modeling:**<br>• EUT, PT and Mixed EUT/PT models, with asset integration and SEU specification<br>• PT model incorporates gain, loss risk aversion, probability weights<br>**Selected Parameter tests**<br>• Model Chi-square test, log pseudolikelihood comparison and parameter estimate t-statistics<br>• Risk aversion parameters ≠ between gain and loss scenarios<br>• Probability weighting parameter estimate ≠ 1<br>• Loss aversion parameter < 1<br>• Risk attitude ≠ between SEU and non-SEU cases; r ≠ between baseline and joint estimation cases<br>• Loss aversion assuming risk neutrality < loss aversion when loss aversion and risk aversion are jointly estimated (i.e. joint estimation matters)<br>• EUT model weight ≠ PT model weight |
| **Experiment # 2 – Subjective Expected Utility and Bayesian Updating**<br>**Motivating Reference Paper:**<br>Antoniou, C., G. W. Harrison, et al. (2015). "Subjective Bayesian Beliefs." Journal of Risk and Uncertainty **50**(1): 35-54 (Antoniou, Harrison et al. 2015) | |
| **In Experiment 2** participants are asked to bet on whether samples of daily security losses, which vary in both 'strength' (the % of losses above a threshold) and 'weight' (the number of samples) are in fact drawn from either a 'Low Controls' or High Controls' system. This experiment replicates Antoniou's treatment of Griffin and Tversky's SEUT experiment using 'low risk' and 'high risk' versions of our simulation model. (Griffin and Tversky 1992; Antoniou, Harrison et al. 2010). The experiment specifically tests whether managers violate Bayes rule when making decisions that involve updating prior expectations under assumptions of subjective expected utility (SEU) with and without the assumption of risk neutrality and allows for testing of the participants' weighting of both the 'strength' and 'weight' of evidence. We also test for the prevalence of the SEUT model over the GT specification. | |
| **Treatment 1: Risk strength / weight testing** | **Modeling:** |

- Loss frame task using system simulation w. uncertain posterior probabilities and strength and weight variations

**Hypothesis 4A:** SEUT probability estimates are more accurate under an assumption of risk neutrality.

**Hypothesis 4B:** The strength and weight of quantitative estimates of security risk influence a manager's subjective expected utility over control decisions (GT model is superior to SEUT model)

- SEUT and SEUT probability estimates assuming risk neutrality, with and without demographic controls.
- GT model strength (alpha) and weight (beta) estimates
- GT model probability estimates

**Selected Parameter tests:**

- Model Chi-square test, log pseudolikelihood comparison and parameter estimate t-statistics
- Risk aversion parameters ≠ between gain and loss scenarios
- Probability weighting parameter estimate ≠ 1
- Loss aversion parameter < 1
- Strength and weight parameters ≠ 1

**Experiment # 3 – Decision Making under Endogenous Risk**

**Motivating Reference Papers:**

Fiore, S. M., G. W. Harrison, et al. (2009). "Virtual experiments and environmental policy." Journal of Environmental Economics and Management **57**(1): 65-86.(Fiore, Harrison et al. 2009)

Sen, S. (2010). Behavioural Response to Endogenous Risk in the Laboratory. Department of Economics. Orlando, Florida, University of Central Florida. **Doctor of Philosophy:** 245. (Sen 2010)

George, J. G., G. W. Harrison, et al. (2012). "Behavioural Responses towards Risk Mitigation: An Experiment with Wild Fire Risks.(George, Harrison et al. 2012)

**In Experiment 3** participants play two separate games. In the first game, participants place bets on whether a single day's loss drawn from a Low Controls versus a High Controls system, respectively, will exceed a certain pre-defined threshold. One bet is then selected at random and each participant is paid out based on whether they bet that the loss would exceed/not exceed the threshold. If they bet correctly, they earn whatever the bookie was offering for that particular bet. If they bet incorrectly, they earn nothing.

In the second game, participants are given a budget to spend on additional controls and then choose whether to upgrade a 'Low' controls system to a 'High' controls system across a range of uncertain costs. The actual cost of the controls is then chosen at random and either a 'Low' or 'High' cost system is then simulated for that participant based on whether the participant chose to upgrade controls at that cost. A single day's loss is then drawn randomly from the simulated system and the value of any controls purchased plus the loss on the day is subtracted from the starting budget. Participants are then paid out based on the net remaining budget. This experiment replicates a study by Fiore (Fiore, Harrison et al. 2009) and then Sen (Sen 2010) incorporating binary choice over both exogenous and endogenous risk using system simulation.

**Treatment 1: Simulation benchmarking**

- Exogenous risk loss frame task, with simulation
- Endogenous risk loss frame task, with simulation

**Hypothesis 1A:** Risk attitude is equal under conditions of Low vs. High Risk

**Hypothesis 1B:** Subjective estimates of security risk are equally accurate under conditions of Low vs. High Risk

**Hypothesis 1C:** Risk attitude varies according to demographic heterogeneity

**Hypothesis 1D:** Subjective estimates of security risk vary according to demographic heterogeneity.

**Modeling:**

- Exogenous case (Low and High systems): SEUT model with risk aversion and subjective probability, with and without demographics
- Endogenous case (Low Controls vs. High Controls systems): SEUT model with risk aversion and subjective probability, with and without demographics

**Selected Parameter tests**

- Model Chi-square test, log pseudolikelihood comparison and parameter estimate t-statistics
- Risk aversion parameters = between Low and High Risk scenarios
- Subjective probability estimates are not statistically different from actual probabilities.

**Hypothesis 1E:** Risk attitude is equal under conditions of exogenous and endogenous risk

**Hypothesis 1E:** Subjective probability estimates of security risk are equal under conditions of exogenous and endogenous risk (i.e. no source dependence)

## Experiment # 4 – Recovering Beliefs Over Continuous Loss Distributions

**Motivating Reference Papers:**

Andersen, S., J. Fountain, et al. (2009). "Estimating aversion to uncertainty." Unpublished discussion paper, College of Business, Univ. of Central Florida (May)(Andersen, Fountain et al. 2009)

Harrison, G. W., J. Martínez-Correa, et al. (2013). "Scoring rules for subjective probability distributions." Manuscript, Georgia State University.(Harrison, Martínez-Correa et al. 2013)

Harrison, G. W. and E. R. Ulm (2015). "Recovering Subjective Probability Distributions." Working Paper. (Harrison and Ulm 2015)

**In Experiment 4** participants play two separate games. In the first game participants undertake a 50-round binary lottery game which tests for the presence of Rank Dependent Utility (RDU), a violation of the common assumption of SEU. In the second game we use a Quadratic Scoring Rule device and a range of security loss simulations to allow participants to bet on subjective estimates of: 1) the probability of dollar losses above a stated threshold, and 2) the dollar value threshold over which a stated probability of losses would occur. Following Harrison et al, reports over continuous ranges are known to be close to latent subjective beliefs if we assume the individual obeys subjective expected utility (SEU) (Savage 1971; Matheson and Winkler 1976; Harrison, Martínez-Correa et al. 2013). The results of these two experiments permit the recovery of subjective beliefs over continuous ranges at the individual level under both an assumption of SEU and if the individual is known to distort probabilities into decision weights using Rank Dependent Utility (RDU) (Harrison and Ulm 2015).

**Treatment 1: RDU Binary Lottery Choices**

- Rank Dependent Binary Lottery Choice (n=50) over prospective gains

**Hypothesis 15A:** Individuals behave according to EUT as opposed to RDU for discrete risky choices.

**Treatment 2: Quadratic Scoring over Continuous Probability Distributions**

- Reports on percentage of security losses over a dollar amount using time series simulation representations.
- Reports on the dollar amount representing losses above a certain percentage threshold

**Hypothesis 15B:** Recovered beliefs over continuous probability distributions are conditional on the assumed model of individual risk preference (SEU vs. RDU).

**Hypothesis 15C:** Security managers are 'reasonably' able to accurately identify the moments of security loss time series and probability distribution representations.

**Modeling:**

- EUT model with Fechner error specification
- RDU model with decision weights:
- Power probability weighting (Quiggin 1982)

**Selected Parameter tests**

- Model Chi-square test, log pseudolikelihood comparison and parameter estimate t-statistics
- Individual RDU power weighting parameter ω(p) = $p^{\gamma}$, γ =1, (no rank dependency)

**Non-Parametric Tests:**

- Histograms of recovered EUT and RDU reports
- Profiles of report accuracy across all participants

## Experiment # 5 – Ambiguity Effects in Self-Protection and Self-Insurance for Security Risks

**Motivating Reference Papers:**
Bajtelsmit, V., J. C. Coats, et al. (2015). "The effect of ambiguity on risk management choices: An experimental study." Journal of Risk and Uncertainty 50(3): 249-280. (Bajtelsmit, Coats et al. 2015)

**Laury, S. K., M. M. McInnes, et al. (2009). "Insurance decisions for low-probability losses." Journal of Risk and Uncertainty 39(1): 17-44 (Laury, McInnes et al. 2009)**

**In Experiment 5** participants play three games. In the first game participants are given a budget and choose whether or not to purchase 'cyber insurance' against simulated daily business losses exceeding a stated threshold for a system with a set control level ('High' or 'Very High') across a range of insurance loadings. In the second game participants are given a budget and choose whether or not to purchase a higher level of system controls for either a 'Low' or 'High' system. In the third game participants are given a budget and choose whether or not to purchase insurance and/or a higher level of system controls. If a higher level of system controls is purchased, the associated insurance premium is adjusted based on the level of control. Participants play each Game 12 times. At the end of each round in each game, the computer selects a daily loss at random, and a payoff is calculated equal to the starting budget minus any insurance and/or controls purchased, minus the random loss. The results of this experiment replicate the experiment of Bajtelsmit (Bajtelsmit, Coats et al. 2015) and is based on and earlier experiment by Laury et al (Laury, McInnes et al. 2009)

**Treatment 1: Insurance only**

- Binary choice of insurance/no insurance over uncertain loss time series (n=12 rounds)

**Treatment 2: Precaution only**

- Binary choice of precaution/no precaution over uncertain loss time series (n=12 rounds)

**Treatment 3: Precaution and Insurance**

- Binary choice of precaution / no precaution and insurance/no insurance over uncertain loss time series (n=12 rounds)

**Hypothesis 16A:** Individuals who prefer lower risk will choose the more efficient risk management method (precaution vs. insurance) to accomplish this goal.

**Hypothesis 16B:** Risk mitigation decisions will be consistent in otherwise similar treatments with and without insurance.

**Hypothesis 16C** Risk mitigation decisions will be consistent in otherwise similar treatments with and without precaution.

**Hypothesis 16D:** Individuals will exercise more precaution when the probability of loss is more ambiguous.

**Hypothesis 16E:** The likelihood of insurance purchase will increase when the probability of loss is more ambiguous.

**Modeling:**

- Logit regression (dependent variable: decision to purchase insurance), as a function of:
  - Probability of loss and insurance premium, when precaution unavailable, w. and without demographics
  - Probability of loss and insurance premium, when precaution available, w. and without demographics
  - Risk Type (Reference = High Controls; 1 = Low Controls or Reference = Very High Controls; 1 = High Controls; 2 = Low Controls Risk = 1, Very High = 2)
  -
- Logit regression (dependent variable: decision to purchase insurance or controls when insurance available), as a function of:
  - Taking precaution when insurance available
  - Degree of ambiguity (variance of loss distribution)

**Selected Parameter tests**

- Non-parametric grouping of risk averse / risk taking decision makers
- Model Chi-square test, log pseudolikelihood comparison and parameter estimate t-statistics
- % buying insurance for lower probability vs. higher probability loss systems are equal holding constant insurance load[59] and expected loss (McNemar test)

**Non-Parametric Results:**

- Percentage of Participants Choosing Precaution at Each System Level
- Percentage of Participants Choosing Precaution or Insurance at Each System Level
- Percentage of Participants Choosing Precaution and Insurance at Each System Level

[59] 'Insurance load' refers to the markup on an otherwise actuarially fair insurance premium. We extend the loading factor from 1x to 6x the actuarially fair premium.

The following section details the general procedures of each experiment and then each of the specific treatments and tasks used to generate my research data. The subsequent section presents the results of the experiments and describes the analytical methods and estimation results of the structural models used in each case. The thesis concludes with a discussion of limitations and reflections on the work and further research opportunities.

# 8 – Experiment Procedures

### Sampling Frame and Sample Recruitment Considerations

For generalizability, careful consideration was given to the sample frame, attracting suitable professionals who have familiarity and affinity with the general context of the information security decision problems being posed since they are, in part, hypothesized to make these control decisions, in context, with biases based on both professional and personal attributes. The opportunity cost for these participants is assumed to be greater than that for students (who are typically recruited for these studies) and therefore needed to be offset with consideration for both increased participation fees and other appropriate incentives, including timing and location convenience. Since the labs were to be conducted in person at the participating site's location (to control for the performance of the computer experiments on researcher supplied laptops), candidates were selected from the Greater Toronto area in Ontario, Canada as a matter of travel convenience to the researcher.

In personal correspondence with Glenn Harrison, he also confirmed the need for panels of possibly upwards of 150 participants to support the essential mixed EUT/PT model specifications and the associated statistical tests which have been developed for these specifications (Vuong 1989; Clarke 2007), although individual experiments have featured as few as 60 participants. Harrison's clear guidance was to ensure an adequate number of participants across gain, loss and mixed treatments and a corresponding adequate number of choice tasks within each treatment which, after pooling the data, in turn provide sufficient model degrees of freedom across and within subjects.  He has cautioned that 'power' requirements for sampling purposes in these models are more closely linked to the design of the gain/loss treatment variations, followed by the number of subjects, and not the number of choices made by each subject. After these considerations and some additional pilot testing of the STATA models using as few as 12 participants, I decided to undertake 12 separate labs with 5 participants per lab, resulting in 60 participants in total.

A total of 60 participants were recruited from the information technology, privacy and security staff of 12 peer academic and community hospitals in the greater Toronto area during March to June 2016. A maximum of five participants were recruited from each individual facility. Participants were informed that the researcher was conducting research into 'information security' and that they would play a series of games on a computer that would involve decision making in an information security context. Each subject signed an informed consent form on arrival to a meeting room at the host facility and was told to sit at any of the 5 numbered laptop computers provided by the researcher. Each laptop had a mouse to facilitate ease of clicking on onscreen buttons within the games. The laptops were physically spaced apart by the researcher as was practical in each location provided by the host organization and participants were instructed to play the games as individuals without collaboration:

**Figure 106 - Typical Lab Setup**

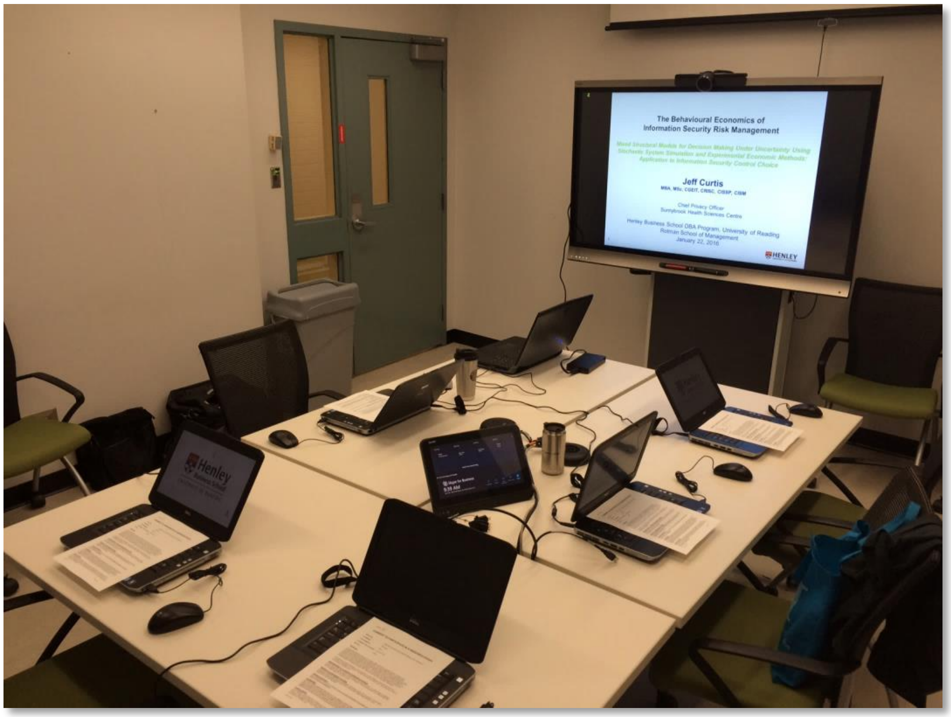
The Behavioural Economics of
Information Security Risk Management
Mixed Structural Models for Decision Making Under Uncertainty Using Stochastic System Simulation and Experimental Economic Methods: Application to Information Security Control Choice
Jeff Curtis
MBA, MSc, CGEIT, CRISC, CISSP, CISM
Chief Privacy Officer
Sunnybrook Health Sciences Centre
Henley Business School DBA Program, University of Reading
Rotman School of Management
January 22, 2016
HENLEY

All instructions were provided verbally by the researcher who also demonstrated each game demonstration of each game on a large screen at the front of the room prior to its being played.  Each game also contained an onscreen, written description of that game's instructions for individual reference during the play. Participants were instructed to indicate immediately after the instructions whether they did not understand how to play a particular game or if they were having difficulty interacting with the game interface during play.  Regardless of whether participants indicated that they understood the game instructions prior to play, some participants appeared to require several rounds of play before verbally indicating that they understood exactly how to play a game. In any case, no participant was allowed to restart or replay a game. In some cases the researcher was required to indicate the order of operation / button pressing in a game to a participant who appeared to be struggling or was not making significant progress on the first several rounds. Other than indicating the elements of the interface and the order of decision making tasks, no other game play guidance was generally provided to the participants.

Participants were instructed that they would be playing five (5) individual games over the course of approximately 2.5 hours and that, in the interest of completing all games within the time allotted in the session, each game was intended to be completed by all participants within approximately 20 minutes before moving on to the next game. Individual participants typically completed games at different times and those finishing before others were allowed to leave the room, check email etc. while waiting for others to finish. In some instances significantly slower participants were not able to complete a game in sufficient time and were instructed to stop playing a particular game in the interest of the group moving ahead together. In a few instances these participants were permitted to return to an incomplete game to finish as many rounds of that game as possible.

In Harrison's experiments in which multiple individual experiments or treatments are undertaken, the ordering of the experiments and treatments undertaken by individual participants is typically randomized to control for 'order effects' which might otherwise influence the way in which participants both learn to play and then play a series of games (Harrison, Johnson et al. 2005). Within this research, randomization of the order of the individual Games 1 through 5 across individuals was not considered practicable since participants were undertaking multiple complex decision games and in pilot sessions it became clear that the provision of common instructions at the start of each game immediately followed by playing the game just instructed and the provision of one-on-one guidance during the games was required to ensure all participants could complete each game within a reasonable period of time. This could only be practicably provided by me as the sole facilitator during the sessions. To have each participant sit through 5 different sets of instructions and then potentially undertake different games from their peers at the same time risked confusion and failure of the sessions and was decided against. Otherwise, the order of the treatments *within* Games 2, 3 and 5 were randomized by reversing the presentation order of the individual sub-games for approximately half of the participants. The order of the lotteries in Game 1 is inherently randomized, as are the Rank dependent Utility (RDU) pie charts lotteries and continuous distribution questions in Game 4. The specific randomization aspects of each game will be explained in the description for each game below. While the lack of randomization *between* games may be a possible confounding element of the data generation process within participants, we consider the nature of the games to be sufficiently diverse that any ordering effect may not be a significant factor overall. Comparatively, the randomization of rounds within each game is considered of primary importance in eliminating order effects for each treatment (i.e. we must be able to control whether the ordering of rounds *within a game* is confounding). Randomization within the sub-games is discussed later in the paper.

Games were designed to be 'incentive compatible' following consistent Harrison's guidance on this aspect of the experimental design (Harrison 2007). Participants were told that the games involved making bets on security control choices and that, if they chose, they might earn some money in the lab (although they might also earn nothing) depending on their choices within the games, although they could not lose any money that was not provided to them within the session. They were told that one game would be selected at random after the session to be paid out for real money in the form of a debit card for popular coffee shop to be emailed to them individually sometime after the session. Payment was indicated as optional, but that they should play the games 'as if' playing for the prospective earnings indicated in order to ensure that their choices were 'consistent' with a general desire to earn money where possible[60].

Participants were told that each game would be played in an Excel spreadsheet which had been modified to eliminate menus and other typical Excel screen features; that they should only use the mouse; that they should only click on 'buttons' visible onscreen as instructed by me before each game; and that they should otherwise not touch the keyboard or, for example, hit the 'Esc' button which would stop the game from running (in which case they would have to restart that game) and to raise their hand if they had any difficulty understanding or operating the game. I explained that I would provide instructions and demonstrate on a separate laptop how to play each game before participants commenced each game and that we would allow everyone to complete each game before proceeding together to the next game to ensure that everyone understood the instructions prior to starting a game. In practice, this meant that the actual time allotted to each game effectively depended on the slowest player and there was considerable diversity in the finishing times across players. To ensure that participants completed all five games, I occasionally had to encourage some laggard players to 'speed up' in order to allow everyone to proceed to the next game. Finally, they were advised that, where some of the games indicated choices based on numerical attributes (percentages or dollar values), in the interest of time, that they were not expected to make 'overly careful mathematical calculations' or use any calculator supports in order to make choices. This coincides with the typical approach in experiments that I have reviewed where the motivation for this restriction appears to align with the idea that in these experiments we are generally testing for the presence of decisional biases, including the use of heuristics, and not whether someone can accurately calculate an expected value or whether that actually affects their decision making (although that certainly is testable). Practically, 'permitting' calculators would involve either the provision of calculators to all participants (to avoid the potentially confounding effects of some participants using and some not using a calculator) of the use of an explicit control group with or without calculators – neither approach was practicable for this study from either a financial or sample size perspective (Birnbaum 2000; Greiner, Jacobsen et al. 2012; Kairies-Schwarz, Kokot et al. 2014; McNair and Feeney 2014; Grossman and Eckel 2015)

The participants were then instructed to open a folder on their laptop which presented a series of Excel files and were instructed to open the file labelled 'Demographic Survey' to begin the session:

[60] In practice, a significant percentage of Participants did not provide an email address, although no one was questioned as to whether they wished to be paid or why they did or did not provide an email address. A few participants indicated that they did not feel that payment was 'necessary' in order to induce them to play or to make decisions that were 'truthful' or consistent with their desire to achieve the indicated payments and accordingly declined payment. Interestingly, I can find no reference to this particular phenomenon in any of the hundreds of papers on experimental procedures I have reviewed to date, although Grossman mentions a particularly interesting potential confound: approximately 25% of participants practiced a religion that prohibited gambling (Grossman and Eckel 2015). This could clearly influence their choices even if they elected to play the games. It is also possible that some participants were sufficiently concerned about their personal privacy that they preferred to not be paid rather than provide an email address. On the other hand, while the games were carefully designed to be 'incentive compatible' (Harrison 2007), the random selection of each participant's game to be paid out and actual cash payment to each participant at the conclusion of the lab session was not practicable due to resource and time limitations at the time of the session and therefore post-session payment via email was undertaken instead..

**Figure 107 - Desktop Game Selection**

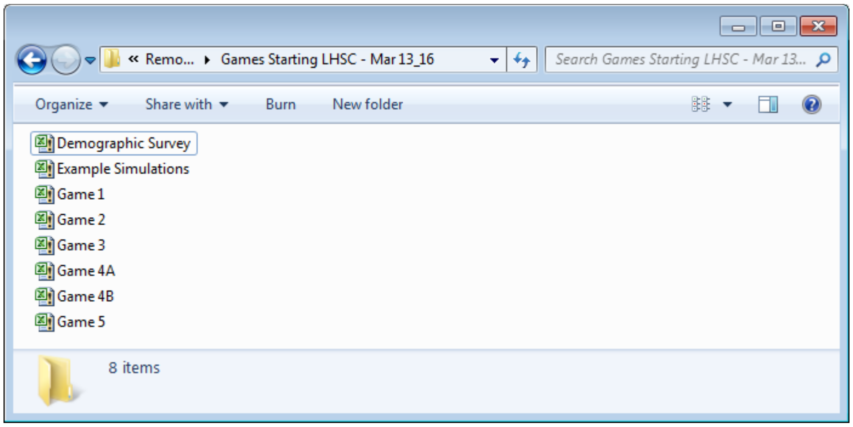
« Remo... ▸ Games Starting LHSC - Mar 13_16
Search Games Starting LHSC - Mar 13...
Organize
Share with
Burn
New folder

Demographic Survey
Example Simulations
Game 1
Game 2
Game 3
Game 4A
Game 4B
Game 5

8 items

### 1 - Demographic Survey

The Demographic Survey collects personal information and asks a series of questions related to general risk preference and several decision examples involving risk. The purpose of the survey was twofold: 1) to collect individual demographic information, some of which would be used as independent variables in the various econometric models including some qualitative information on their perceptions of their own risk preferences which could be used to benchmark individual attitudes towards risk outside of the games; and 2) to get participants used to using the mouse and the software interface by clicking on the screen ‘buttons’ (as opposed to traditional Excel or Windows interface elements) and to initiate their thinking towards risk-based decision making generally. They were instructed that, if they wished to be paid for the session as was indicated in the Instructions, they should include their email address on the Demographics Survey and that their email address would only be used to contact them regarding the results of the lab and for payment purposes. Participants were told that they could skip a question if they were unsure of the answer or preferred not to answer for privacy reasons. Two Participants asked to re-complete the survey after the session after reconsidering their willingness to reveal certain personal information.

The following diagram illustrates the form of the survey; a complete listing of the demographic survey questions is included in the Appendix:

**Figure 108 - Demographic Survey Interface**

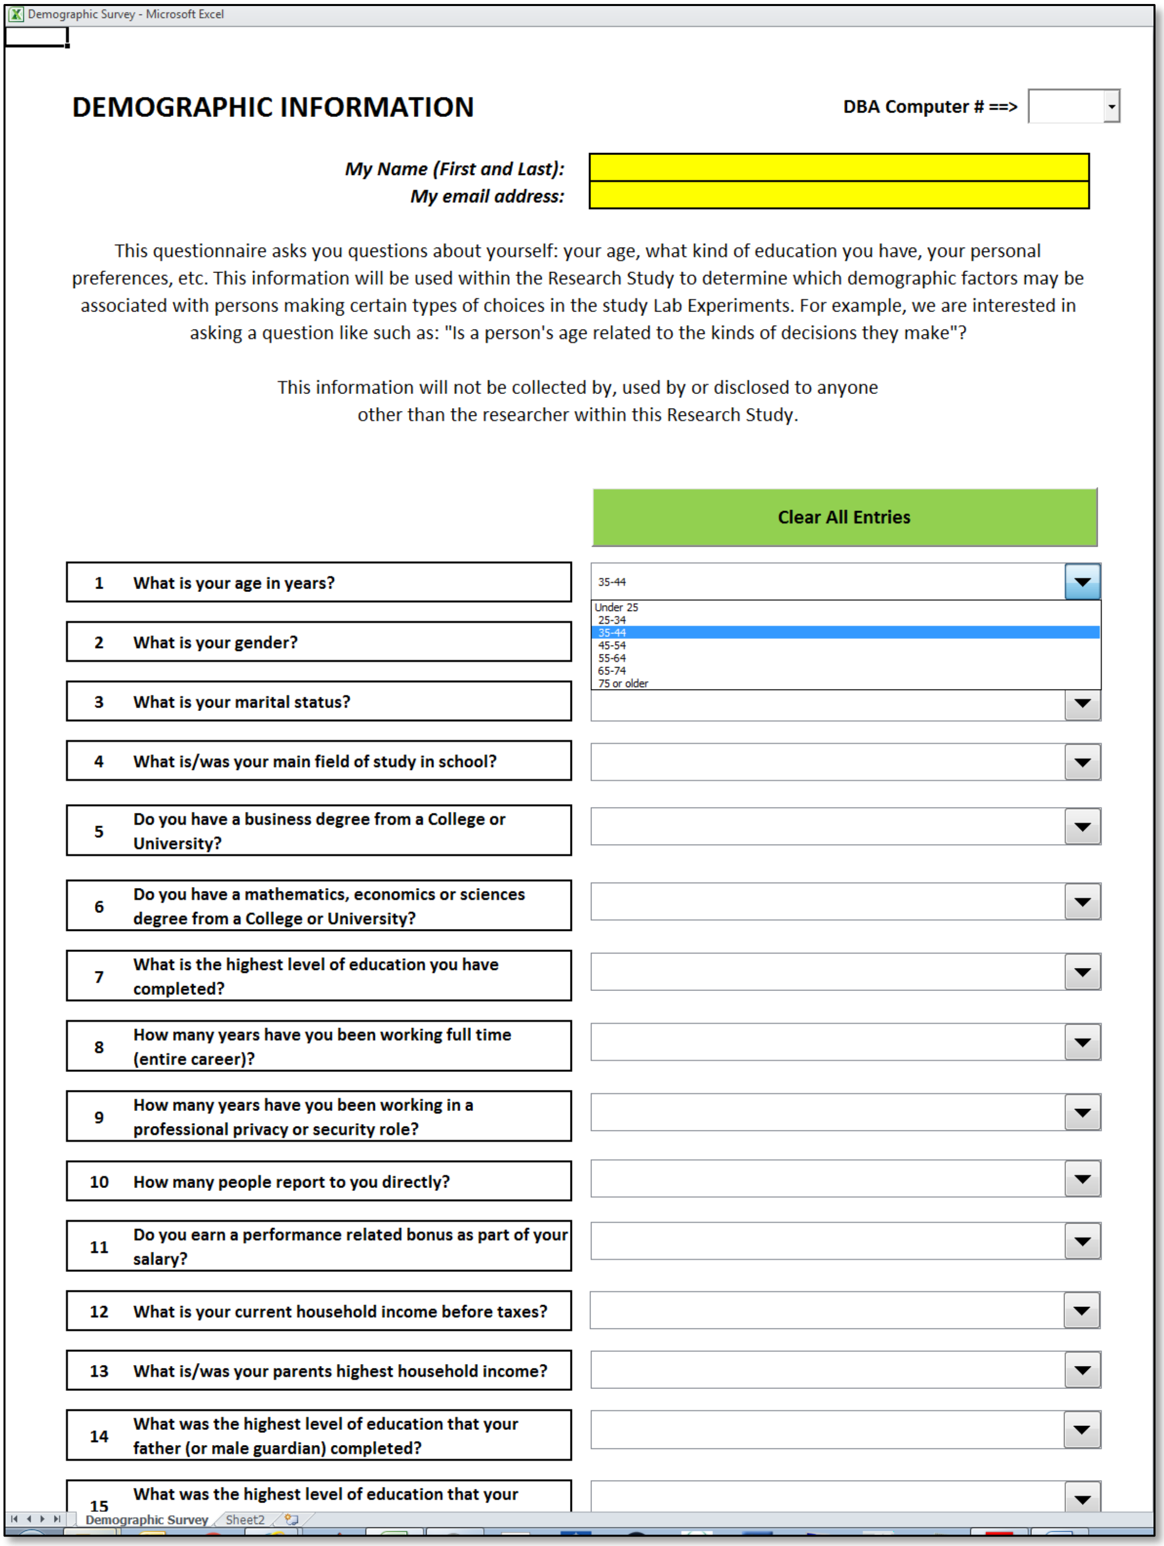

## DEMOGRAPHIC INFORMATION

DBA Computer # ==>

*My Name (First and Last):*

*My email address:*

This questionnaire asks you questions about yourself: your age, what kind of education you have, your personal preferences, etc. This information will be used within the Research Study to determine which demographic factors may be associated with persons making certain types of choices in the study Lab Experiments. For example, we are interested in asking a question like such as: "Is a person's age related to the kinds of decisions they make"?

This information will not be collected by, used by or disclosed to anyone other than the researcher within this Research Study.

| | | |
|---|---|---|
| 1 | What is your age in years? | 35-44<br>Under 25<br>25-34<br>35-44<br>45-54<br>55-64<br>65-74<br>75 or older |
| 2 | What is your gender? | |
| 3 | What is your marital status? | |
| 4 | What is/was your main field of study in school? | |
| 5 | Do you have a business degree from a College or University? | |
| 6 | Do you have a mathematics, economics or sciences degree from a College or University? | |
| 7 | What is the highest level of education you have completed? | |
| 8 | How many years have you been working full time (entire career)? | |
| 9 | How many years have you been working in a professional privacy or security role? | |
| 10 | How many people report to you directly? | |
| 11 | Do you earn a performance related bonus as part of your salary? | |
| 12 | What is your current household income before taxes? | |
| 13 | What is/was your parents highest household income? | |
| 14 | What was the highest level of education that your father (or male guardian) completed? | |
| 15 | What was the highest level of education that your | |

**2 – Game #1: Asset Integration**

Participants are then instructed to open the ‘Game 1’ file. All rounds of the game are played on a single screen (see figure next page). In this game, each participant is initially endowed with a stake, selected at random from a uniform distribution between $1 and $6 in discrete $1 intervals. The subject then faces 20 lottery choices in a sequence. These lotteries offer the possibility of some gains, but also offer the possibility of some losses. Some lotteries mix positive and negative outcomes. The specific lotteries presented to each subject are selected ‘more or less’ at random from a set of 36 probability combinations and 12 payout combinations. This set replicates virtually all of the lotteries used by Tversky and Kahneman (Tversky and Kahneman 1992) as well as several lotteries designed to have outcomes with different signs. The selection is non-random only to the extent that we ensure that the first three lotteries all offer non-negative outcomes, to allow subjects to accumulate some earnings before facing the prospect of losses. After these three, each subject receives a random sequence of pie chart draws, with repetition, from the possible probability and payout combinations:

**Figure 109 - Game 1 User Interface**

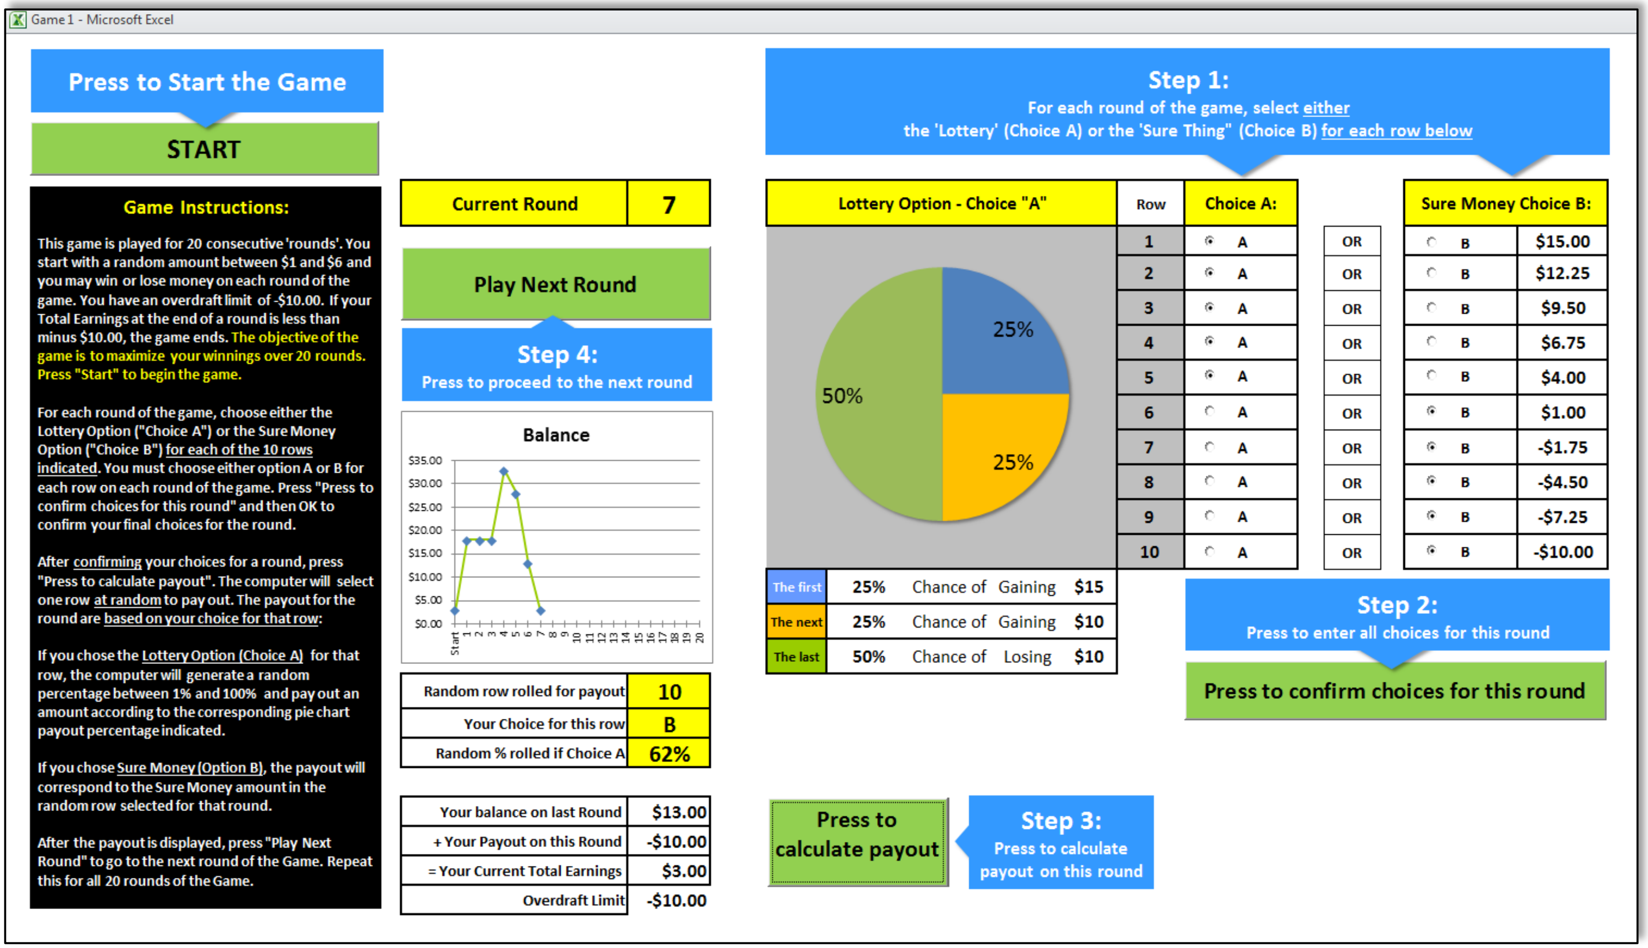

**Game Instructions:**

This game is played for 20 consecutive 'rounds'. You start with a random amount between $1 and $6 and you may win or lose money on each round of the game. You have an overdraft limit of -$10.00. If your Total Earnings at the end of a round is less than minus $10.00, the game ends. The objective of the game is to maximize your winnings over 20 rounds. Press "Start" to begin the game.

For each round of the game, choose either the Lottery Option ("Choice A") or the Sure Money Option ("Choice B") for each of the 10 rows indicated. You must choose either option A or B for each row on each round of the game. Press "Press to confirm choices for this round" and then OK to confirm your final choices for the round.

After confirming your choices for a round, press "Press to calculate payout". The computer will select one row at random to pay out. The payout for the round are based on your choice for that row:

If you chose the Lottery Option (Choice A) for that row, the computer will generate a random percentage between 1% and 100% and pay out an amount according to the corresponding pie chart payout percentage indicated.

If you chose Sure Money (Option B), the payout will correspond to the Sure Money amount in the random row selected for that round.

After the payout is displayed, press "Play Next Round" to go to the next round of the Game. Repeat this for all 20 rounds of the Game.

| Current Round | 7 |
|---|---|

**Step 4:** Press to proceed to the next round

| | |
|---|---|
| Random row rolled for payout | 10 |
| Your Choice for this row | B |
| Random % rolled if Choice A | 62% |

| | |
|---|---|
| Your balance on last Round | $13.00 |
| + Your Payout on this Round | -$10.00 |
| = Your Current Total Earnings | $3.00 |
| Overdraft Limit | -$10.00 |

**Step 1:** For each round of the game, select either the 'Lottery' (Choice A) or the 'Sure Thing" (Choice B) for each row below

Lottery Option - Choice "A"

| The first | 25% | Chance of | Gaining | $15 |
|---|---|---|---|---|
| The next | 25% | Chance of | Gaining | $10 |
| The last | 50% | Chance of | Losing | $10 |

| Row | Choice A: | | Sure Money Choice B: | |
|---|---|---|---|---|
| 1 | ● A | OR | ○ B | $15.00 |
| 2 | ● A | OR | ○ B | $12.25 |
| 3 | ● A | OR | ○ B | $9.50 |
| 4 | ● A | OR | ○ B | $6.75 |
| 5 | ● A | OR | ○ B | $4.00 |
| 6 | ○ A | OR | ● B | $1.00 |
| 7 | ○ A | OR | ● B | -$1.75 |
| 8 | ○ A | OR | ● B | -$4.50 |
| 9 | ○ A | OR | ● B | -$7.25 |
| 10 | ○ A | OR | ● B | -$10.00 |

**Step 2:** Press to enter all choices for this round

**Step 3:** Press to calculate payout on this round

Each lottery choice is actually a choice between a specific pie chart lottery and an ordered list of 10 non-stochastic amounts based on the minimum and maximum payout for that round. The list was ordered in a linear manner from the lowest lottery prize to the highest lottery prize. In the example above, consider a pie chart lottery with an outcome of gaining $15 with probability 25% (i.e. for the first 1%-25%), an outcome of gaining $10 with probability 25% (i.e. for the next 26%-50%) and an outcome of losing $10 with probability 50% (i.e. for the last 51%-100%). The first fixed choice would be between the pie chart lottery and gaining $15 for sure. The second would be between the pie chart lottery and gaining $12.25 for sure, the third would be between the pie chart lottery and gaining $9.50 for sure, and the tenth would be between the pie chart lottery and losing $10 for sure.

After the participant chooses between either the pie chart or the Sure Money option for each of the 10 rows, they click ‘Press to Confirm Choices’ which ensures that they have made choices in all rows, and then click ‘Press to Calculate Payout’. The computer then generates a random number between 0 and 10 to select a row to play out. If the participant selected the pie chart option for that row, the computer generates a random number between 1% and 100% (inclusive), and the result determines their payout from the pie chart. If they chose the sure money option for the selected row, the sure money amount is the payout for that round. The payout determines the change to their accumulated earnings (‘Balance’) and they are informed of the result onscreen. Participants are told that they will be provided an “overdraft limit” of $10 and may continue to play as long as their balance does not exceed minus $10 on any round; if their accumulated earnings drop below minus $10, their game is over. Players continue make choices on each round for 20 rounds at which point the Game is over. When completed, participants are instructed to close the Excel window which saves the choice results for that player.

**Example Simulations**

Participants are then instructed to open the file “Example Simulations’:

Figure 110 - Game 1: Example Simulations Introduction

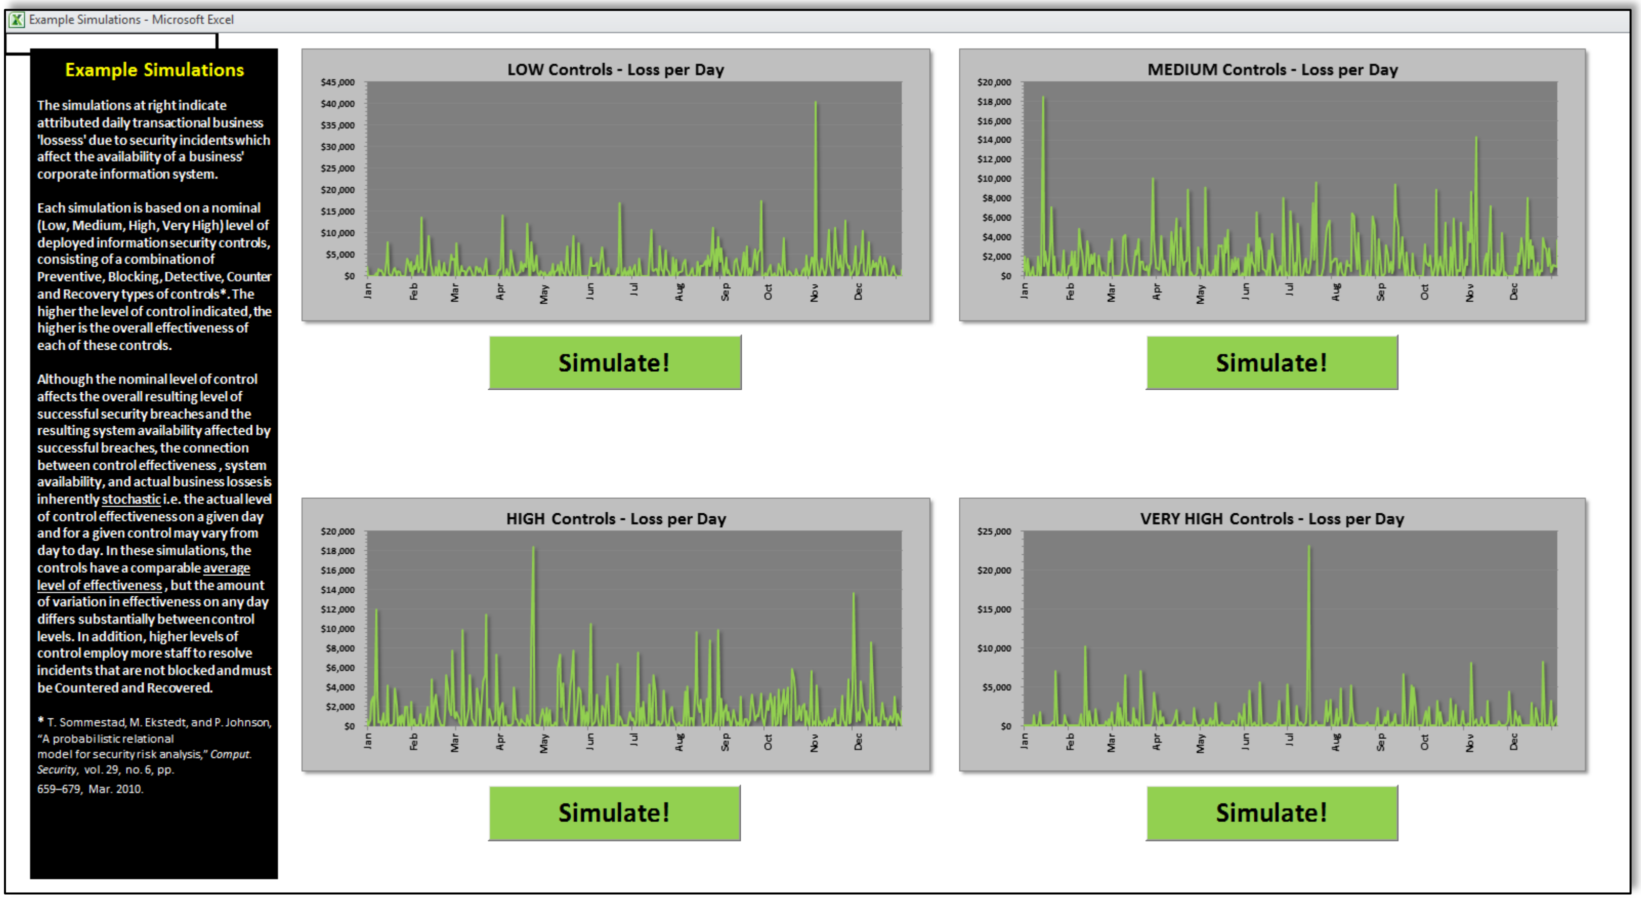

Example Simulations - Microsoft Excel

## Example Simulations

The simulations at right indicate attributed daily transactional business 'lossess' due to security incidents which affect the availability of a business' corporate information system.

Each simulation is based on a nominal (Low, Medium, High, Very High) level of deployed information security controls, consisting of a combination of Preventive, Blocking, Detective, Counter and Recovery types of controls*. The higher the level of control indicated, the higher is the overall effectiveness of each of these controls.

Although the nominal level of control affects the overall resulting level of successful security breaches and the resulting system availability affected by successful breaches, the connection between control effectiveness , system availability, and actual business losses is inherently stochastic i.e. the actual level of control effectiveness on a given day and for a given control may vary from day to day. In these simulations, the controls have a comparable average level of effectiveness , but the amount of variation in effectiveness on any day differs substantially between control levels. In addition, higher levels of control employ more staff to resolve incidents that are not blocked and must be Countered and Recovered.

* T. Sommestad, M. Ekstedt, and P. Johnson, "A probabilistic relational model for security risk analysis," *Comput. Security*, vol. 29, no. 6, pp. 659–679, Mar. 2010.

Participants are then told that the remainder of the games (including the current game) involve the ‘simulation’ of an information system modelled at different levels of security control which is at risk of security breaches and the daily ‘business losses’ attributed to the breaches due to the resulting lack of availability of the system. Participants were encouraged to click on the ‘Simulate’ button for one or more of the indicated systems to get a sense of how the system performed under the various control levels and were told that they would be able to do similarly within each of the games to come as required. They were told that as the indicated levels of control increased from ‘Low Controls’ through to ‘Very High Controls’, the average daily loss for that system would generally decrease, although the loss on any particular day could exceed a loss level for a lower level of control. In practice, some participants at this point ask questions regarding how the availability or the attributed losses are calculated, but generally the participants appeared to accept the premise of business or operational losses attributed (in some manner) to the nominal security control posture of the indicated information system. Participants were asked whether they understood the explanation of the simulation of losses before proceeding and we did not proceed until all participants indicated that they understood the concept and use of the charts.

As noted above, results for each of the control level simulations were precomputed daily for 30 years and were stored as data values which could be interactively selected in annual samples (n=365 days) and displayed as required onscreen in real time. This eliminated the need to interactively compute Monte Carlo simulations, taking up to a minute or more to compute, which would have unacceptably interfered with the immediacy of the games and ensured that the resulting simulation samples were being drawn from the same ‘population’ of precomputed simulation results across participants.

### 3 – Game #2: Subjective Bayesian Updating and Strength / Weight of Evidence

Participants are then instructed to open the ‘Game 2’ file. This game is played over a series of Excel ‘tabs’ with each tab representing one round of the game. The game starts with an overview of an information system simulated under ‘relatively Low security controls’ and, alternatively, under ‘relatively High security controls’:

## Figure 111 - Game 2: Example Simulations

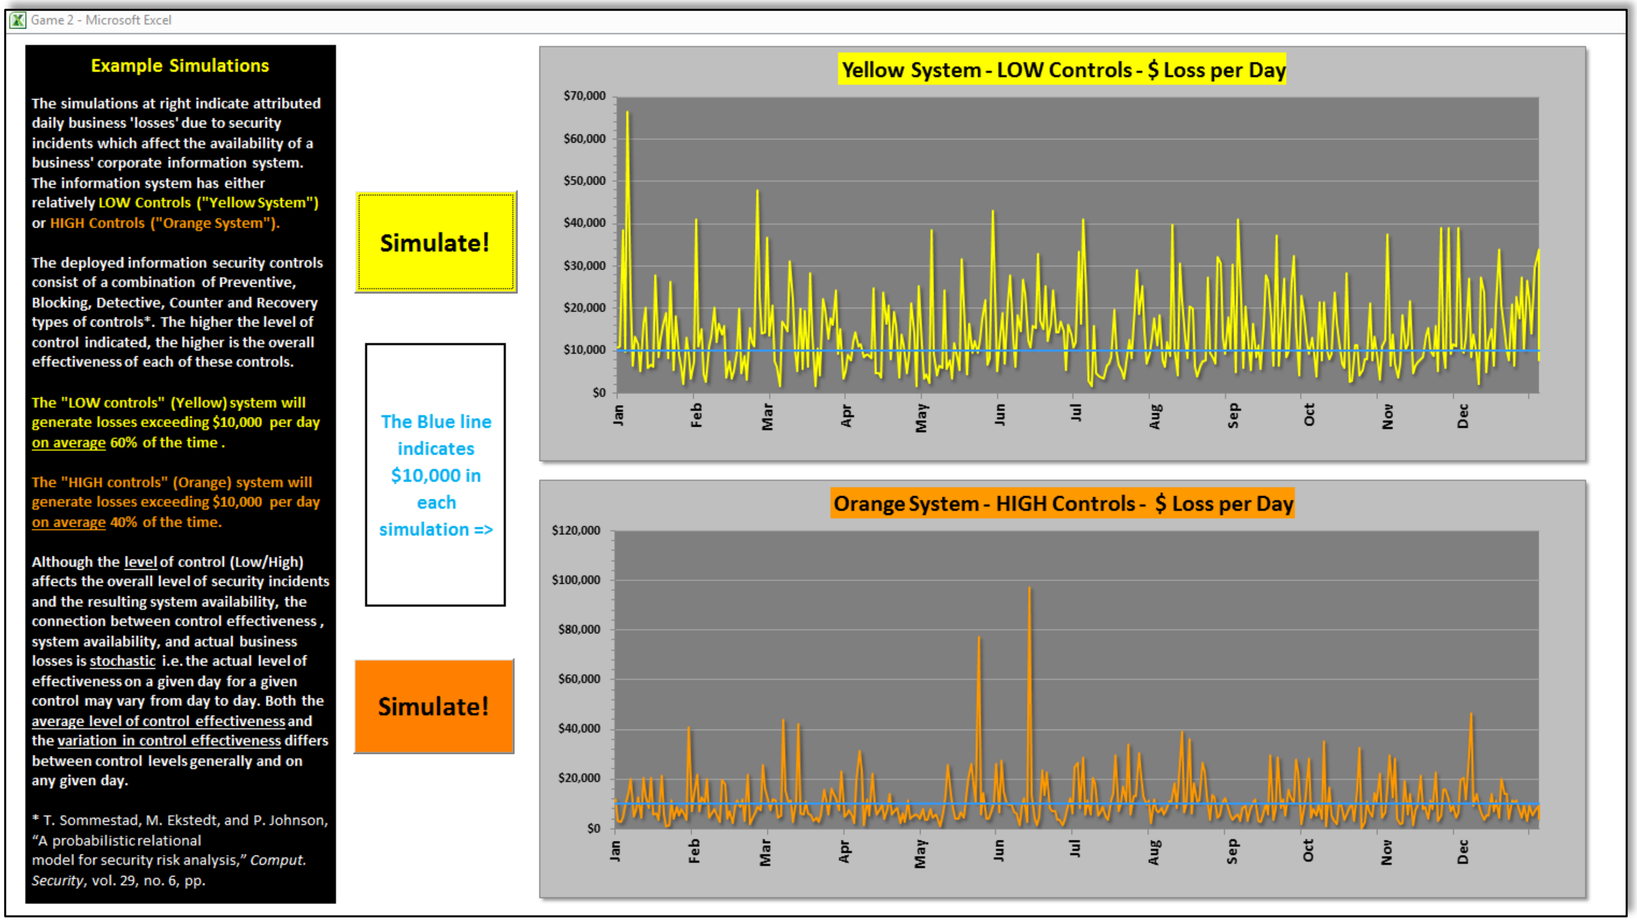

**Example Simulations**

The simulations at right indicate attributed daily business 'losses' due to security incidents which affect the availability of a business' corporate information system. The information system has either relatively LOW Controls ("Yellow System") or HIGH Controls ("Orange System").

The deployed information security controls consist of a combination of Preventive, Blocking, Detective, Counter and Recovery types of controls*. The higher the level of control indicated, the higher is the overall effectiveness of each of these controls.

The "LOW controls" (Yellow) system will generate losses exceeding $10,000 per day on average 60% of the time .

The "HIGH controls" (Orange) system will generate losses exceeding $10,000 per day on average 40% of the time.

Although the level of control (Low/High) affects the overall level of security incidents and the resulting system availability, the connection between control effectiveness , system availability, and actual business losses is stochastic i.e. the actual level of effectiveness on a given day for a given control may vary from day to day. Both the average level of control effectiveness and the variation in control effectiveness differs between control levels generally and on any given day.

* T. Sommestad, M. Ekstedt, and P. Johnson, "A probabilistic relational model for security risk analysis," *Comput. Security*, vol. 29, no. 6, pp.

The Blue line indicates $10,000 in each simulation =>

Participants are told that the ‘Low Controls’ system generates losses exceeding $10,000, 60% of the time and that the ‘High Controls’ system generates losses exceeding $10,000, 40% of the time and a common $10,000 baseline indicator is displayed within the simulated time series for each system. As previously indicated, Participants were reminded that they are free to click on the respective ‘Simulate’ buttons to observe the simulated business losses of the system under each set of controls, and are told that they can return to this tab at any time during the ensuing rounds of the game to re-examine the loss behaviour of the systems as they may require.

Participants then proceed to the first game tab which is representative of each of the ensuing game rounds:

Figure 112 - Game 2: User Interface

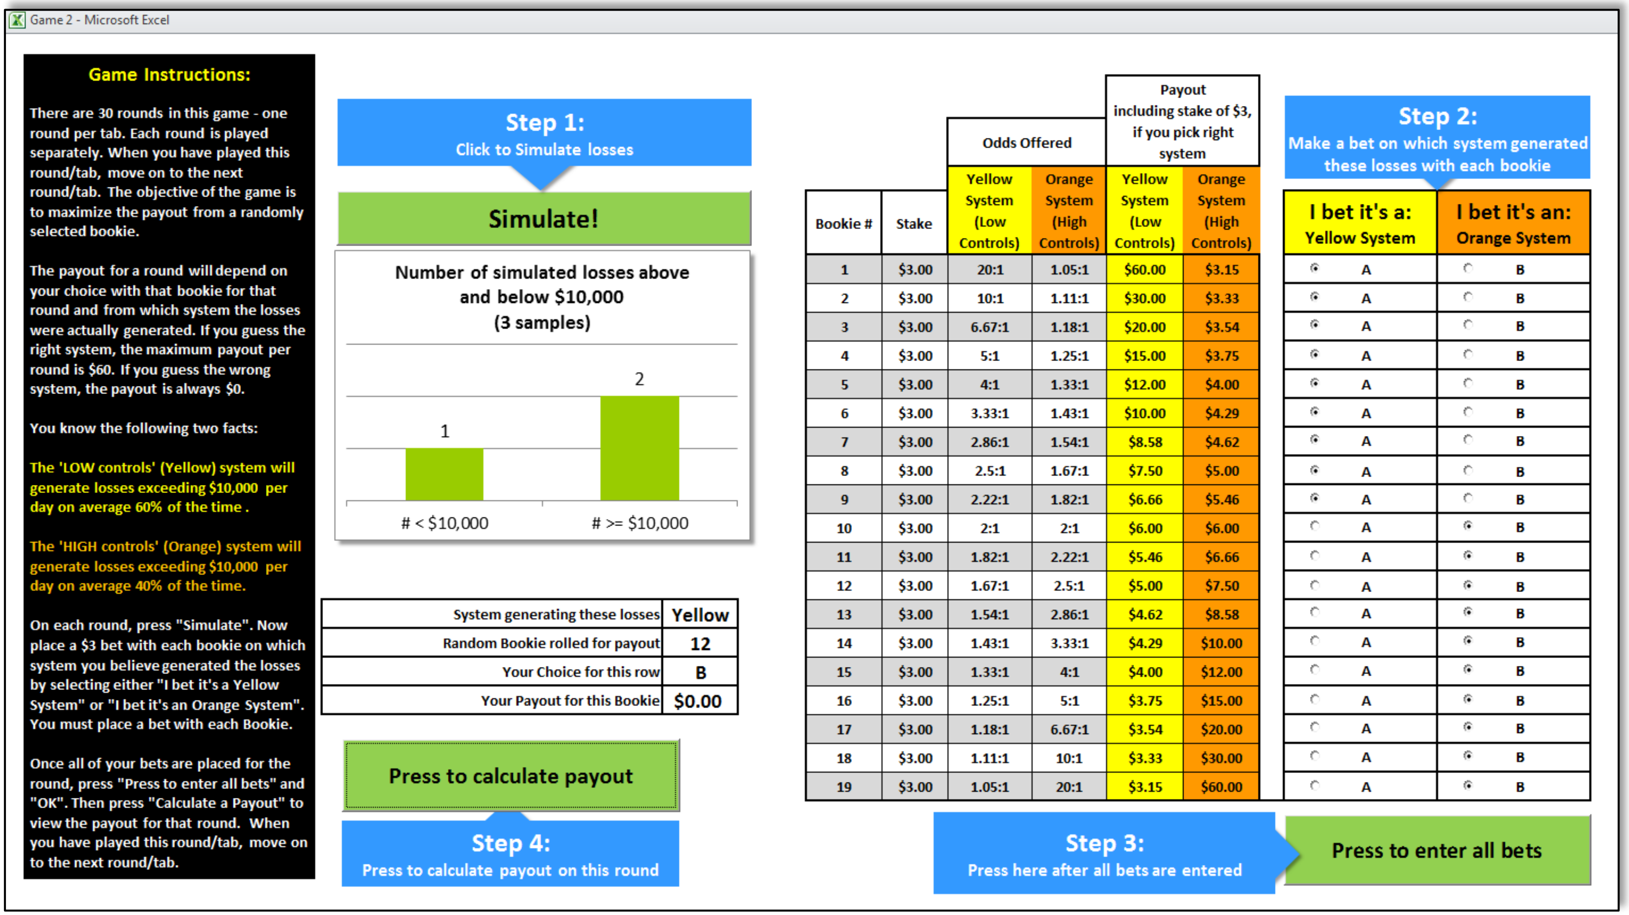

Game 2 - Microsoft Excel

**Game Instructions:**

There are 30 rounds in this game - one round per tab. Each round is played separately. When you have played this round/tab, move on to the next round/tab. The objective of the game is to maximize the payout from a randomly selected bookie.

The payout for a round will depend on your choice with that bookie for that round and from which system the losses were actually generated. If you guess the right system, the maximum payout per round is $60. If you guess the wrong system, the payout is always $0.

You know the following two facts:

The 'LOW controls' (Yellow) system will generate losses exceeding $10,000 per day on average 60% of the time .

The 'HIGH controls' (Orange) system will generate losses exceeding $10,000 per day on average 40% of the time.

On each round, press "Simulate". Now place a $3 bet with each bookie on which system you believe generated the losses by selecting either "I bet it's a Yellow System" or "I bet it's an Orange System". You must place a bet with each Bookie.

Once all of your bets are placed for the round, press "Press to enter all bets" and "OK". Then press "Calculate a Payout" to view the payout for that round. When you have played this round/tab, move on to the next round/tab.

**Step 1:**
Click to Simulate losses

Number of simulated losses above and below $10,000 (3 samples)

| | |
|---|---|
| System generating these losses | Yellow |
| Random Bookie rolled for payout | 12 |
| Your Choice for this row | B |
| Your Payout for this Bookie | $0.00 |

**Step 4:**
Press to calculate payout on this round

**Step 2:**
Make a bet on which system generated these losses with each bookie

| | | Odds Offered | | Payout including stake of $3, if you pick right system | | | |
|---|---|---|---|---|---|---|---|
| Bookie # | Stake | Yellow System (Low Controls) | Orange System (High Controls) | Yellow System (Low Controls) | Orange System (High Controls) | I bet it's a: Yellow System | I bet it's an: Orange System |
| 1 | $3.00 | 20:1 | 1.05:1 | $60.00 | $3.15 | ● A | ○ B |
| 2 | $3.00 | 10:1 | 1.11:1 | $30.00 | $3.33 | ● A | ○ B |
| 3 | $3.00 | 6.67:1 | 1.18:1 | $20.00 | $3.54 | ● A | ○ B |
| 4 | $3.00 | 5:1 | 1.25:1 | $15.00 | $3.75 | ● A | ○ B |
| 5 | $3.00 | 4:1 | 1.33:1 | $12.00 | $4.00 | ● A | ○ B |
| 6 | $3.00 | 3.33:1 | 1.43:1 | $10.00 | $4.29 | ● A | ○ B |
| 7 | $3.00 | 2.86:1 | 1.54:1 | $8.58 | $4.62 | ● A | ○ B |
| 8 | $3.00 | 2.5:1 | 1.67:1 | $7.50 | $5.00 | ● A | ○ B |
| 9 | $3.00 | 2.22:1 | 1.82:1 | $6.66 | $5.46 | ● A | ○ B |
| 10 | $3.00 | 2:1 | 2:1 | $6.00 | $6.00 | ○ A | ● B |
| 11 | $3.00 | 1.82:1 | 2.22:1 | $5.46 | $6.66 | ○ A | ● B |
| 12 | $3.00 | 1.67:1 | 2.5:1 | $5.00 | $7.50 | ○ A | ● B |
| 13 | $3.00 | 1.54:1 | 2.86:1 | $4.62 | $8.58 | ○ A | ● B |
| 14 | $3.00 | 1.43:1 | 3.33:1 | $4.29 | $10.00 | ○ A | ● B |
| 15 | $3.00 | 1.33:1 | 4:1 | $4.00 | $12.00 | ○ A | ● B |
| 16 | $3.00 | 1.25:1 | 5:1 | $3.75 | $15.00 | ○ A | ● B |
| 17 | $3.00 | 1.18:1 | 6.67:1 | $3.54 | $20.00 | ○ A | ● B |
| 18 | $3.00 | 1.11:1 | 10:1 | $3.33 | $30.00 | ○ A | ● B |
| 19 | $3.00 | 1.05:1 | 20:1 | $3.15 | $60.00 | ○ A | ● B |

**Step 3:**
Press here after all bets are entered

For each round of the game, the computer selects either the yellow (Low Controls) or the Orange (High Controls) system and then randomly draws a set of observations from a randomized 365 day period within 30 years of simulation data from that system and displays the number of observations equal to or greater than $10,000, and less than $10,000 respectively. In the above example, the computer has selected 1 observation below $10,000 and 2 observations equal to or greater than $10,000. The Participants then places a bet with each of 19 'bookies' offering different odds and corresponding prospective payouts over whether the actual source of the observations is a Low or High Controls system. Participants are 'staked' $3 for each bet and must place a bet with each bookie. The odds and payouts for each bookie are the same in each game. The Participant then confirms their choices and the computer selects one row/bookie to play out for money. The Participant then proceeds to the next round tab and completes the remaining rounds in the same manner. As in Antoniou, each subject participated for 30 rounds of this betting task: 4 rounds with N=3 system samples[61], 14 rounds with N=5 system samples, 6 rounds with N=9 system samples, and 6 rounds with N=17 system samples. This distribution of sample sizes was chosen to ensure that we observe roughly the same number of "extreme" samples. To control for possible order effects within the game, in half of the sessions I varied the presentation of sample sizes in ascending order (i.e., first 4 rounds of 3, then 14 rounds of 5, etc.) and in descending order (6 rounds of 17, 6 rounds of 9, etc.).

This is a general form of a betting procedure for eliciting choices that depend on subjective probabilities (Savage 1971; Antoniou, Harrison et al. 2015): "The recovery of subjective probabilities and beliefs requires formal theoretical and parametric assumptions…The essential logic is that this decision sheet is a "multiple price list," just like the decision sheet used to infer discount rates by (Harrison, Lau et al. 2002), the decision sheet used to infer risk attitudes by (Holt 2002), and the decision sheet used to infer valuations for goods by (Andersen, Harrison et al. 2006). The general "multiple price list" (MPL) betting interface employed here was first used by (Fiore, Harrison et al. 2009)." The method is also used to test the degree to which the 'strength' of evidence (i.e. the proportion of samples indicating one system or the other) and the 'weight' of evidence (the number of samples) affects the Participants' ability to update their prior expectation of the probability of choosing either the Low or High system:

> Griffin and Tversky [1992] suggest that, contrary to Bayes Rule, individuals are more sensitive to two aspects of the sample evidence. One is the "strength" of the evidence, defined here as how many white or blue sides came up as a proportion of the total number of dice that were rolled. The other aspect of the sample evidence is its "weight," defined here as how many dice were rolled in total. Our experimental design, directly extending the design of Griffin and Tversky [1992], tests this conjecture by comparing subjects' betting patterns for realized dice patterns that have the same posterior but differ in strength and weight.

[61] To generate his 'samples', Antoniou employed two sets of primarily blue- and primarily white-faced dice, where each set of dice had a 60% and 40% chance of rolling blue/white faces as representing the two 'systems at risk'. The dice were randomly selected from one of the two sets without informing the participant which set they were drawn from and then rolled in front of the participant with the resulting faces acting as the 'samples' for each round of the game. I have essentially recreated his experiment in a security context using a simulated system at risk in which the loss outcomes are similarly constrained between the 'Low' and 'High' controls systems. Instead of rolling multiple dice for each round of the game, I select multiple corresponding 'samples' from the selected simulated system population of results.

For example, if we roll N=5 dice and get 4 dice with a white face and 1 dice with a blue face, the posterior probability of the white box being used is exactly the same (0.77) as when we toss N=9 dice and get 6 white and 3 blue. Bayes Rule predicts that the elicited subjective probabilities in these two cases should be identical. Under SEU, and assuming the correct application of Bayes Rule, the realization from the N=5 case and the N=9 dice are equally informative. In each case there is a posterior subjective probability, with no uncertainty: the individual does not behave as if he thinks that the 0.77 is more likely in the N=9 case than the N=5 case[62].

Griffin and Tversky [1992] observe that the strength of the N=5 pattern is greater (since [4/5] > [6/9]) and the weight smaller (since 5 < 9) in this example. They hypothesize that subjective beliefs will be updated differently when we get 4 white and 1 blue as opposed to 6 white and 3 blue, and the two cases will generate different subjective probabilities. Specifically, they hypothesize that decision makers will hold subjective probabilities that overshoot Bayes Rule when there is "high strength" and "low weight," and undershoot Bayes Rule when there is "low strength" and "high weight."

[62] Antoniou: "We say "behave as if" not out of rhetorical ritual. It is apparent that [{Savage, 1971 #31}; p.56ff] was well aware that individuals might have uncertain beliefs about some event (e.g., inferring that it could be 0.75 with probability 1/3, 0.77 with probability 1/3, and 0.79 with probability 1/3), but he insisted that if they behaved consistently with the axioms he proposed that they would behave identically to someone that had the average of that distribution of uncertain beliefs (in this example, 0.77). In other words, the two possibilities were observationally equivalent, and one might then impose Occam's Razor or some other a priori argument to whittle down the hypotheses to be tested. (Antoniou, Harrison et al. 2010)

### 4 – Game #3: Decisions Under Exogenous and Endogenous Risk

Participants are then instructed to open the 'Game 3' file. This game is played in 2 separate sub-games: in the first sub-game players place bets on to whether a random daily loss will exceed a defined threshold for a Low Controls System, and in the second round of the sub-game make a similar bet on a 'High Controls' System. In the second sub-game, players decide whether to buy security controls that lower the probability of a random loss exceeding a defined threshold where the cost of the controls varies by means of a 'multiple price list'. Following Sen, using both of these tasks we can jointly estimate subjective perception of risk and risk attitude under condition involving endogenous risk (Sen 2010).

The game starts with an overview an information system again simulated under 'relatively Low security controls' and, alternatively, under 'relatively High security controls':

Figure 113 - Game 3: Example Simulations

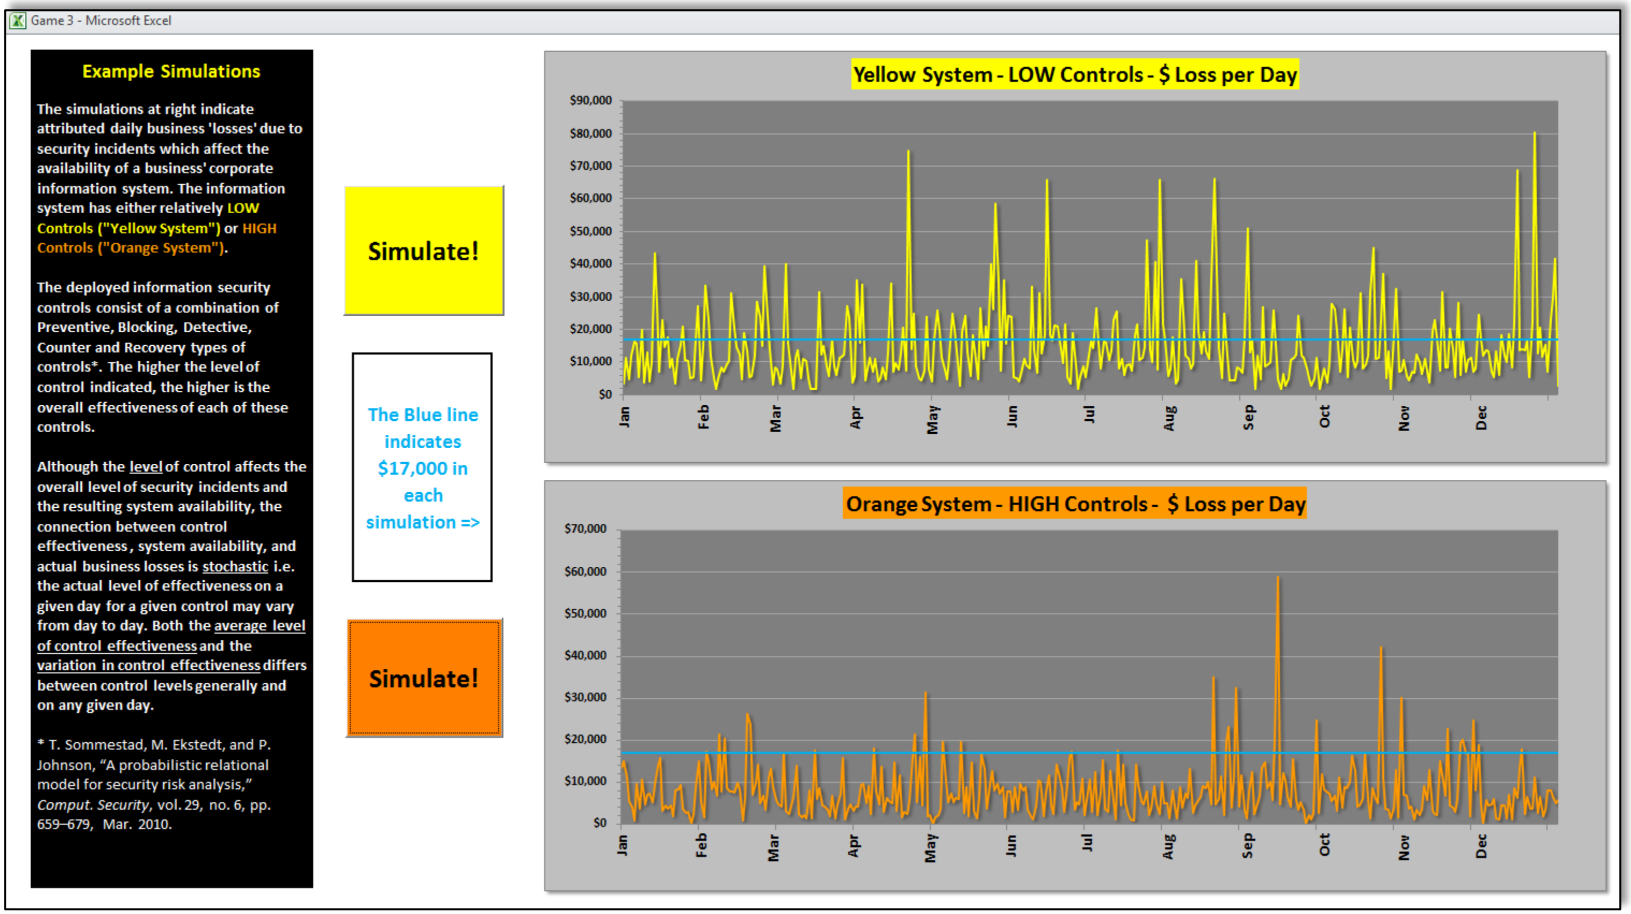

### Example Simulations

The simulations at right indicate attributed daily business 'losses' due to security incidents which affect the availability of a business' corporate information system. The information system has either relatively LOW Controls ("Yellow System") or HIGH Controls ("Orange System").

The deployed information security controls consist of a combination of Preventive, Blocking, Detective, Counter and Recovery types of controls*. The higher the level of control indicated, the higher is the overall effectiveness of each of these controls.

Although the level of control affects the overall level of security incidents and the resulting system availability, the connection between control effectiveness, system availability, and actual business losses is stochastic i.e. the actual level of effectiveness on a given day for a given control may vary from day to day. Both the average level of control effectiveness and the variation in control effectiveness differs between control levels generally and on any given day.

[*] T. Sommestad, M. Ekstedt, and P. Johnson, "A probabilistic relational model for security risk analysis," *Comput. Security*, vol. 29, no. 6, pp. 659–679, Mar. 2010.

The Blue line indicates $17,000 in each simulation =>

Participants are told that the ‘Low Controls’ system typically generates losses greater than the ‘High Controls’ system, but are told nothing otherwise about the relative loss probability distributions between the two simulations. Participants were reminded that they are free to click on the respective ‘Simulate’ buttons to observe the simulated business losses of the system under each set of controls, and are told that they can return to this tab at any time during the ensuing rounds of the game to re-examine the loss behaviour of the systems as they may require.

Participants then proceed to the first of two game tabs which is representative of the first two sub-game rounds:

## Figure 114 - Game 3A: User Interface

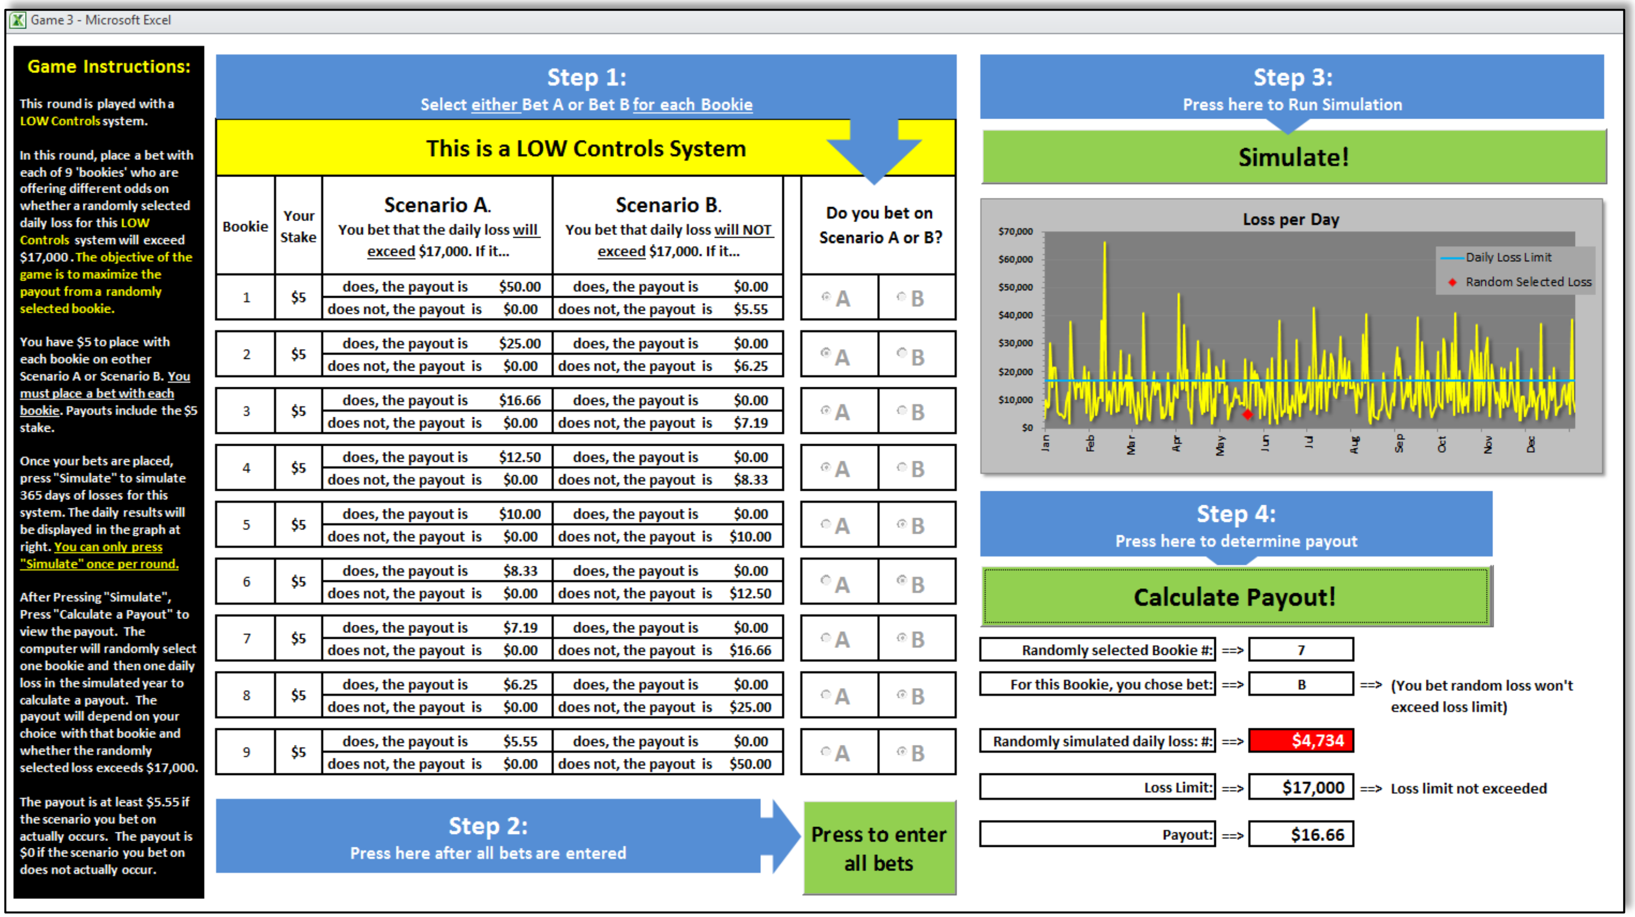

**Game Instructions:**

This round is played with a LOW Controls system.

In this round, place a bet with each of 9 'bookies' who are offering different odds on whether a randomly selected daily loss for this LOW Controls system will exceed $17,000. The objective of the game is to maximize the payout from a randomly selected bookie.

You have $5 to place with each bookie on either Scenario A or Scenario B. You must place a bet with each bookie. Payouts include the $5 stake.

Once your bets are placed, press "Simulate" to simulate 365 days of losses for this system. The daily results will be displayed in the graph at right. You can only press "Simulate" once per round.

After Pressing "Simulate", Press "Calculate a Payout" to view the payout. The computer will randomly select one bookie and then one daily loss in the simulated year to calculate a payout. The payout will depend on your choice with that bookie and whether the randomly selected loss exceeds $17,000.

The payout is at least $5.55 if the scenario you bet on actually occurs. The payout is $0 if the scenario you bet on does not actually occur.

**Step 1:**
Select either Bet A or Bet B for each Bookie

**This is a LOW Controls System**

| Bookie | Your Stake | Scenario A. You bet that the daily loss will exceed $17,000. If it... | | Scenario B. You bet that daily loss will NOT exceed $17,000. If it... | | Do you bet on Scenario A or B? | |
|---|---|---|---|---|---|---|---|
| 1 | $5 | does, the payout is | $50.00 | does, the payout is | $0.00 | ● A | ○ B |
| | | does not, the payout is | $0.00 | does not, the payout is | $5.55 | | |
| 2 | $5 | does, the payout is | $25.00 | does, the payout is | $0.00 | ● A | ○ B |
| | | does not, the payout is | $0.00 | does not, the payout is | $6.25 | | |
| 3 | $5 | does, the payout is | $16.66 | does, the payout is | $0.00 | ● A | ○ B |
| | | does not, the payout is | $0.00 | does not, the payout is | $7.19 | | |
| 4 | $5 | does, the payout is | $12.50 | does, the payout is | $0.00 | ● A | ○ B |
| | | does not, the payout is | $0.00 | does not, the payout is | $8.33 | | |
| 5 | $5 | does, the payout is | $10.00 | does, the payout is | $0.00 | ○ A | ● B |
| | | does not, the payout is | $0.00 | does not, the payout is | $10.00 | | |
| 6 | $5 | does, the payout is | $8.33 | does, the payout is | $0.00 | ○ A | ● B |
| | | does not, the payout is | $0.00 | does not, the payout is | $12.50 | | |
| 7 | $5 | does, the payout is | $7.19 | does, the payout is | $0.00 | ○ A | ● B |
| | | does not, the payout is | $0.00 | does not, the payout is | $16.66 | | |
| 8 | $5 | does, the payout is | $6.25 | does, the payout is | $0.00 | ○ A | ● B |
| | | does not, the payout is | $0.00 | does not, the payout is | $25.00 | | |
| 9 | $5 | does, the payout is | $5.55 | does, the payout is | $0.00 | ○ A | ● B |
| | | does not, the payout is | $0.00 | does not, the payout is | $50.00 | | |

**Step 2:**
Press here after all bets are entered

**Step 3:**
Press here to Run Simulation

**Step 4:**
Press here to determine payout

Randomly selected Bookie #: ==> 7

For this Bookie, you chose bet: ==> B ==> (You bet random loss won't exceed loss limit)

Randomly simulated daily loss: #: ==> $4,734

Loss Limit: ==> $17,000 ==> Loss limit not exceeded

Payout: ==> $16.66

In the first sub-game round the screen indicates that the system is a ‘Low Control’ system. The Participants then places a bet with each of 9 ‘bookies’ offering different prospective payouts over whether a randomly selected daily loss will or will not exceed a $17,000 threshold. For each bookie Participants bet whether the random loss will exceed or will not exceed the threshold. Participants are ‘staked’ $5 for each bet and must place a bet with each bookie. The Participant then confirms their choices and then press ‘Simulate’ to simulate one year’s losses for the Low controls system and they cannot simulate the system again. They then press ‘Calculate Payout’ and the computer selects one row/bookie to play out for money. In the above example, the computer randomly selected Row 7 and for that row the Participant had bet that the random loss would not exceed $17,000. The randomly selected daily loss was $4,734 which is displayed on the graph and numerically. In this case since the Participant bet that the loss limit would not be exceeded and it was not actually exceeded, the player wins the amount indicated for that scenario for that row/bookie, $16.66. This sub-game round is now over and the player proceeds to the next tab/round and plays the same game with a ‘High Controls’ system.

Participants then proceed to the second sub-game which is played over three rounds. The following screen is representative of the next three sub-game rounds:

Figure 115 - Game 3B: User Interface

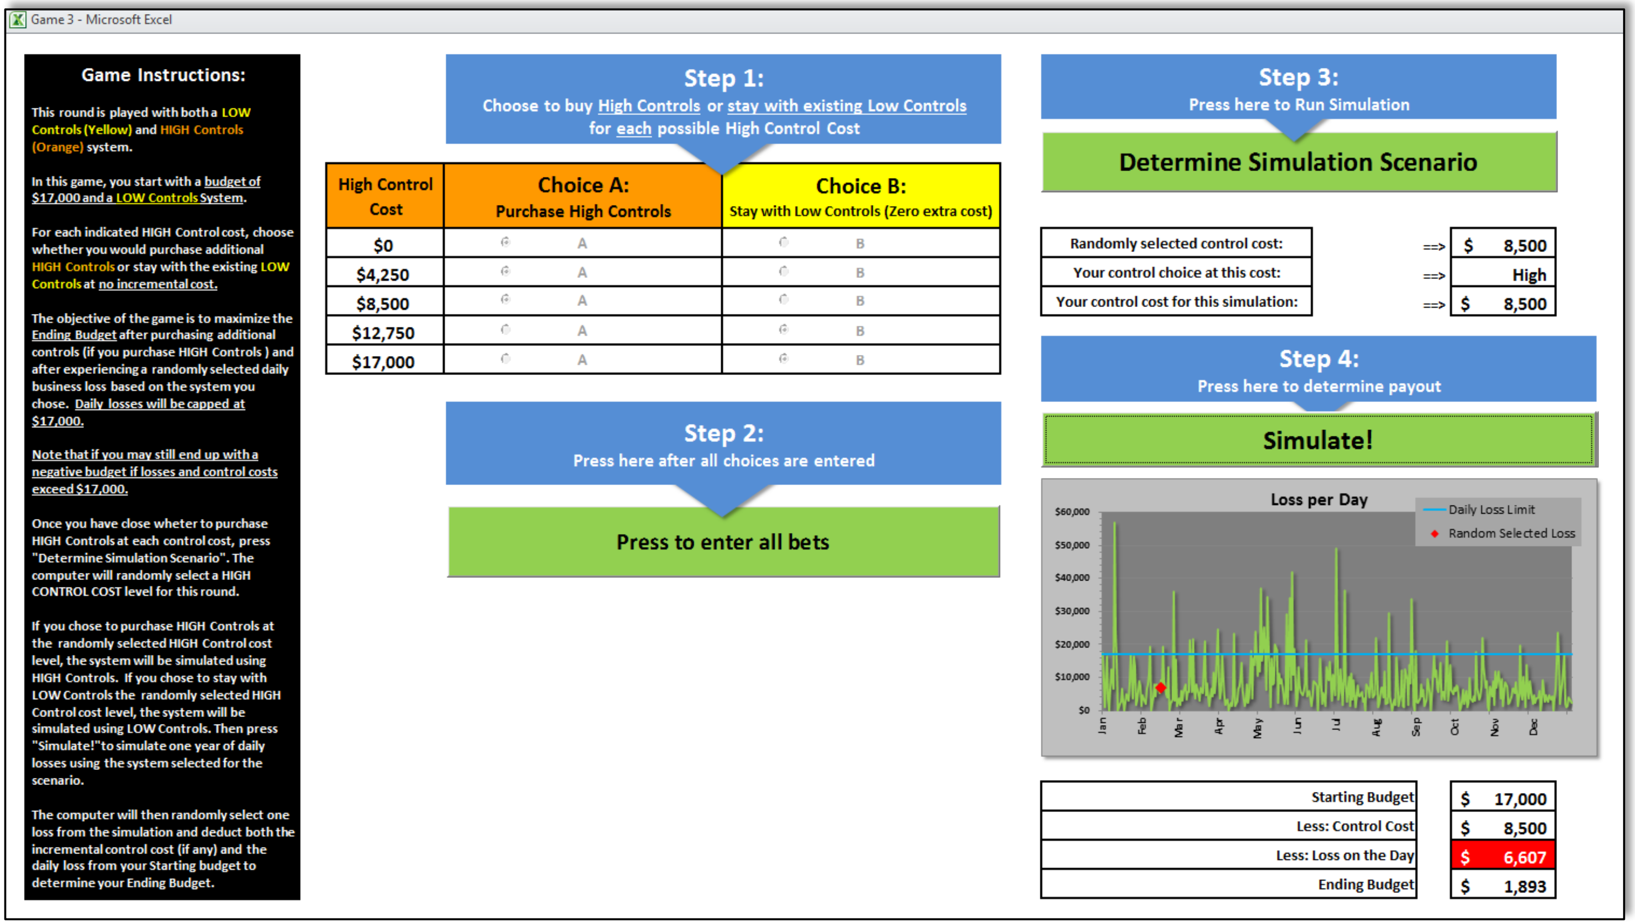

**Game Instructions:**

This round is played with both a LOW Controls (Yellow) and HIGH Controls (Orange) system.

In this game, you start with a budget of $17,000 and a LOW Controls System.

For each indicated HIGH Control cost, choose whether you would purchase additional HIGH Controls or stay with the existing LOW Controls at no incremental cost.

The objective of the game is to maximize the Ending Budget after purchasing additional controls (if you purchase HIGH Controls ) and after experiencing a randomly selected daily business loss based on the system you chose. Daily losses will be capped at $17,000.

Note that if you may still end up with a negative budget if losses and control costs exceed $17,000.

Once you have close wheter to purchase HIGH Controls at each control cost, press "Determine Simulation Scenario". The computer will randomly select a HIGH CONTROL COST level for this round.

If you chose to purchase HIGH Controls at the randomly selected HIGH Control cost level, the system will be simulated using HIGH Controls. If you chose to stay with LOW Controls the randomly selected HIGH Control cost level, the system will be simulated using LOW Controls. Then press "Simulate!"to simulate one year of daily losses using the system selected for the scenario.

The computer will then randomly select one loss from the simulation and deduct both the incremental control cost (if any) and the daily loss from your Starting budget to determine your Ending Budget.

**Step 1:**
Choose to buy High Controls or stay with existing Low Controls for each possible High Control Cost

| High Control Cost | Choice A: Purchase High Controls | Choice B: Stay with Low Controls (Zero extra cost) |
|---|---|---|
| $0 | ● A | ○ B |
| $4,250 | ● A | ○ B |
| $8,500 | ● A | ○ B |
| $12,750 | ○ A | ● B |
| $17,000 | ○ A | ● B |

**Step 2:**
Press here after all choices are entered

Press to enter all bets

**Step 3:**
Press here to Run Simulation

Determine Simulation Scenario

| | | |
|---|---|---|
| Randomly selected control cost: | ==> | $ 8,500 |
| Your control choice at this cost: | ==> | High |
| Your control cost for this simulation: | ==> | $ 8,500 |

**Step 4:**
Press here to determine payout

Simulate!

| | |
|---|---|
| Starting Budget | $ 17,000 |
| Less: Control Cost | $ 8,500 |
| Less: Loss on the Day | $ 6,607 |
| Ending Budget | $ 1,893 |

In the first sub-game round Participants are told that they start with a budget of $17,000 and a 'Low Controls' system. The objective of the game is to maximize the amount of budget the Participant has left over after buying (or not buying) controls and then experiencing a random daily loss where the payout for the round equals the starting budget less any purchased controls less the random daily loss.

The Participants first choose whether to purchase 'High Controls' that convert the 'Low Controls ' system to a 'High Controls' system at different High Control cost levels between zero and $17,000, or spend nothing on High Controls and stay with the Low Controls system. Participants must select either purchase High Controls or stay with Low Controls at each indicated High Control cost level. They are told that after they make their control selections, the computer will select one control cost level at random which will determine whether a 'Low Controls' or a 'High Controls' system will be used to undertake the simulation for the rest of this round based on whether the Participant chose High or Low controls at that cost level. Once the type of system for the round is determined, the Participants click 'Simulate' and the computer selects either the Yellow (Low Controls) or the Orange (High Controls) system (based on the players choice of controls) and then randomly draws a set of loss observations from a randomized 365 day period within 30 years of simulation data for that system and displays the time series on screen. The computer then selects a random daily loss from the 365 day simulation, calculates the budget balance and displays the results onscreen.

In the above example, the computer randomly selected a High Control cost of $8,500 and, since the Participant had indicated they would buy High Controls at that cost level, the simulation for this round uses a High Controls system. Simulated losses are then generated for a High Controls system and a random daily random of $6,607 is drawn by the computer. The computer then displays the resulting balance: $17,000, minus control costs of $8,500, minus the loss of $6,607, netting a balance of $1,893. The payout for the player would be 1/1000 of this amount, or $1.89. This sub-game round is now over and the player proceeds to the next tab/round and plays the same game two more rounds with different cost spreads:

## Table 9 - Game 3 Control Cost Spreads

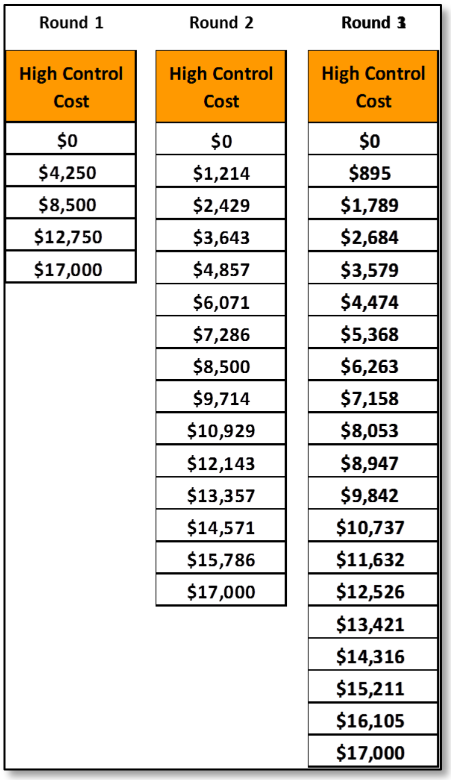

| Round 1 | Round 2 | Round 3 |
|---|---|---|
| High Control Cost | High Control Cost | High Control Cost |
| $0 | $0 | $0 |
| $4,250 | $1,214 | $895 |
| $8,500 | $2,429 | $1,789 |
| $12,750 | $3,643 | $2,684 |
| $17,000 | $4,857 | $3,579 |
|  | $6,071 | $4,474 |
|  | $7,286 | $5,368 |
|  | $8,500 | $6,263 |
|  | $9,714 | $7,158 |
|  | $10,929 | $8,053 |
|  | $12,143 | $8,947 |
|  | $13,357 | $9,842 |
|  | $14,571 | $10,737 |
|  | $15,786 | $11,632 |
|  | $17,000 | $12,526 |
|  |  | $13,421 |
|  |  | $14,316 |
|  |  | $15,211 |
|  |  | $16,105 |
|  |  | $17,000 |

### 5 – Game #4: Recovering Subjective Probability Estimates and Rank Dependent Utility Bias

The Participants then proceed to Game 4, which is composed of 2 sub-games. Following Harrison and Ulm, we attempt to recover latent subjective beliefs over continuous events (probabilities of loss thresholds and loss thresholds representing probability distribution moments) using a quadratic scoring rule (QSR) defined over monetary payoffs. The participant reports are evaluated under assumptions of both subjective expected utility and rank dependent utility to allow recovery of 'true' latent subjective beliefs.

In the first sub-game, the participant undertakes 50 binary lottery choices with objective probabilities using a standard pie chart interface where the lottery pairs have been specially designed for the purpose of testing whether the individual behaves consistently with EUT or RDU following Harrison (Harrison and Ulm 2015) as derived from Wakker et al (Wakker, Erev et al. 1994) and Wilcox (Wilcox 2010):

Figure 116 - Game 4A: User Interface

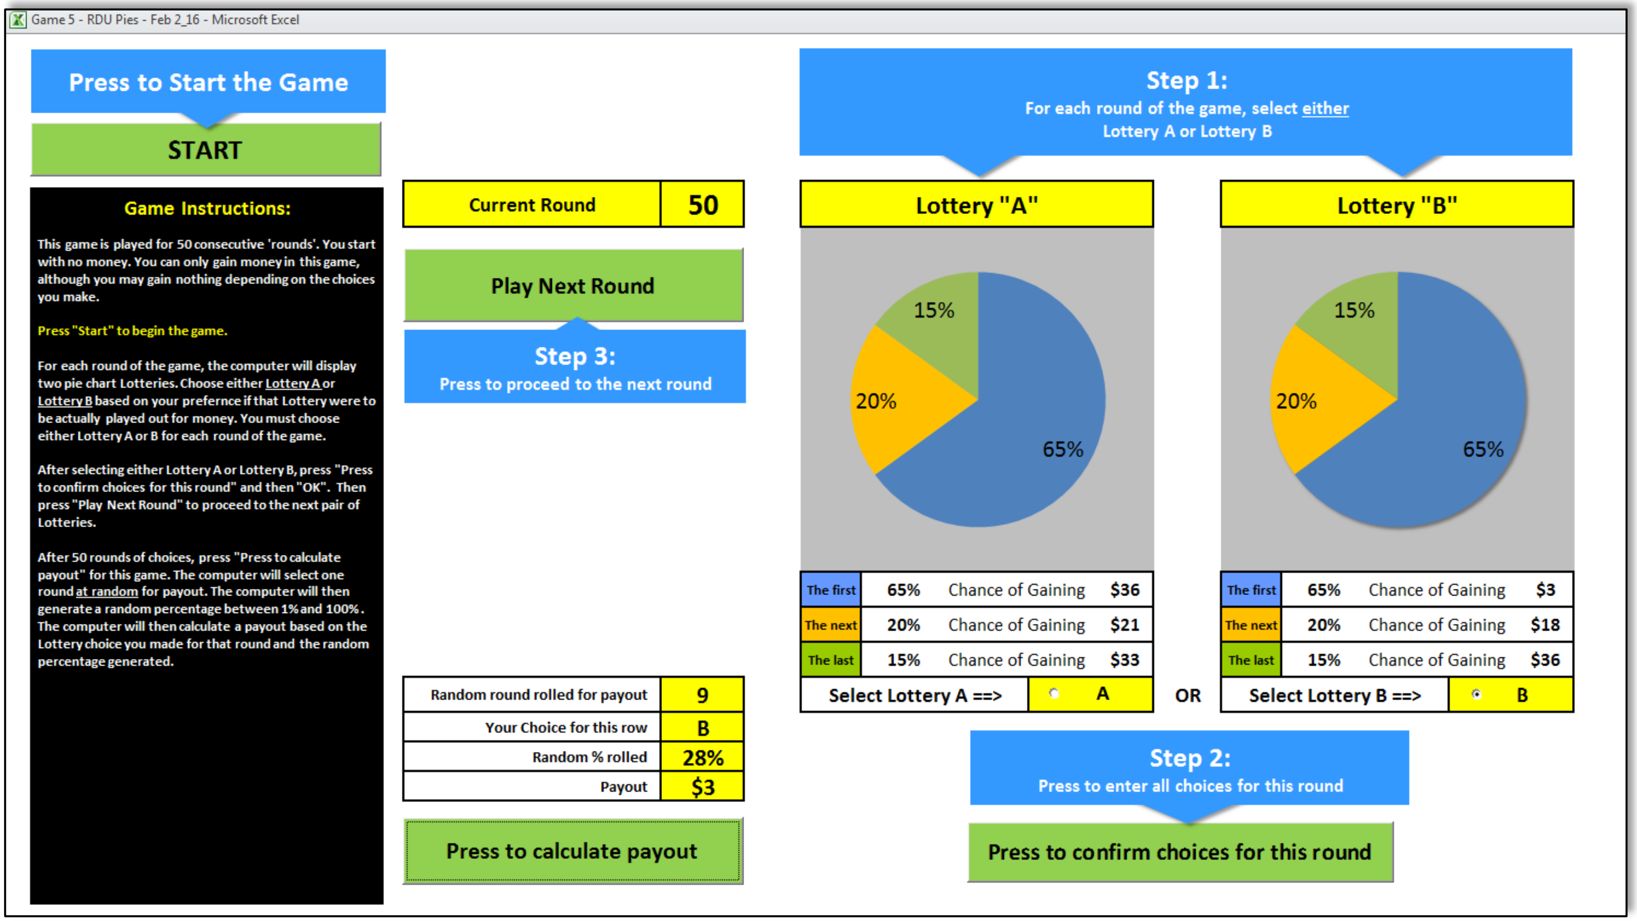

**Game Instructions:**

This game is played for 50 consecutive 'rounds'. You start with no money. You can only gain money in this game, although you may gain nothing depending on the choices you make.

Press "Start" to begin the game.

For each round of the game, the computer will display two pie chart Lotteries. Choose either Lottery A or Lottery B based on your prefernce if that Lottery were to be actually played out for money. You must choose either Lottery A or B for each round of the game.

After selecting either Lottery A or Lottery B, press "Press to confirm choices for this round" and then "OK". Then press "Play Next Round" to proceed to the next pair of Lotteries.

After 50 rounds of choices, press "Press to calculate payout" for this game. The computer will select one round at random for payout. The computer will then generate a random percentage between 1% and 100%. The computer will then calculate a payout based on the Lottery choice you made for that round and the random percentage generated.

| | |
|---|---|
| Current Round | 50 |

Play Next Round
Step 3:
Press to proceed to the next round
Step 1:
For each round of the game, select either Lottery A or Lottery B
Lottery "A"
15%
20%
65%
Lottery "B"
15%
20%
65%

| | | | |
|---|---|---|---|
| The first | 65% | Chance of Gaining | $36 |
| The next | 20% | Chance of Gaining | $21 |
| The last | 15% | Chance of Gaining | $33 |
| Select Lottery A ==> | | ○ | A |

OR

| | | | |
|---|---|---|---|
| The first | 65% | Chance of Gaining | $3 |
| The next | 20% | Chance of Gaining | $18 |
| The last | 15% | Chance of Gaining | $36 |
| Select Lottery B ==> | | ● | B |

| | |
|---|---|
| Random round rolled for payout | 9 |
| Your Choice for this row | B |
| Random % rolled | 28% |
| Payout | $3 |

Press to calculate payout
Step 2:
Press to enter all choices for this round
Press to confirm choices for this round

All lotteries are in the gain frame. In each round, Participants choose either the left or the right lottery and then proceed to the next lottery without calculating a payout until all lotteries have been played. Once all 50 have been played, the Participant clicks ‘Press to calculate payout’ and the computer selects one lottery at random to calculate a payout. The computer then generates a random percentage between 1% and 100% (inclusive) and based on the Participants choice for that round indicates the payout for that lottery. In the above example, the computer selected round #8 in which the Participant chose the right hand lottery which pays $3 with probability 65% (i.e. for 1%-65%), $18 with probability 20% (i.e. for 66% – 85%) and $36 with probability 15% (i.e. for the last 86%-100%). The computer generated 28% and the corresponding payout is $3. The results of this game are used to estimate the Participant’s risk preferences under assumption of both EUT and RDU which can be subsequently used to infer bias in their subjective reports in sub-game #2.

Participants then proceed to the second sub-game which is played over 16 rounds. In each round the Participant is shown two screens: one screen displays either a time series or a probability distribution of losses for either a single year or for multiple years, and a second screen in which they are asked to bet on the answer to a numerical question about the displayed time series/probability distribution. Players answer all 16 questions and then the computer selects one question at random to pay out.

The questions are always of one of two forms, based on whether the displayed losses are shown as a time series or as a probability distribution:

**Time Series:**
*For this [Low/Medium/High/Very High] system, [for this year/across all years], what percentage of daily loss amounts would exceed [$ dollar amount]?*

**Probability distribution:**
*For this [Low/Medium/High/Very High] system, [for this year/across all years], what daily loss amount would be exceeded [X%] of the time?*

The following illustrates a typical probability distribution display screen and the corresponding question screen:

Figure 117 - Game 4B: Low Controls Question Probability Distribution Representation

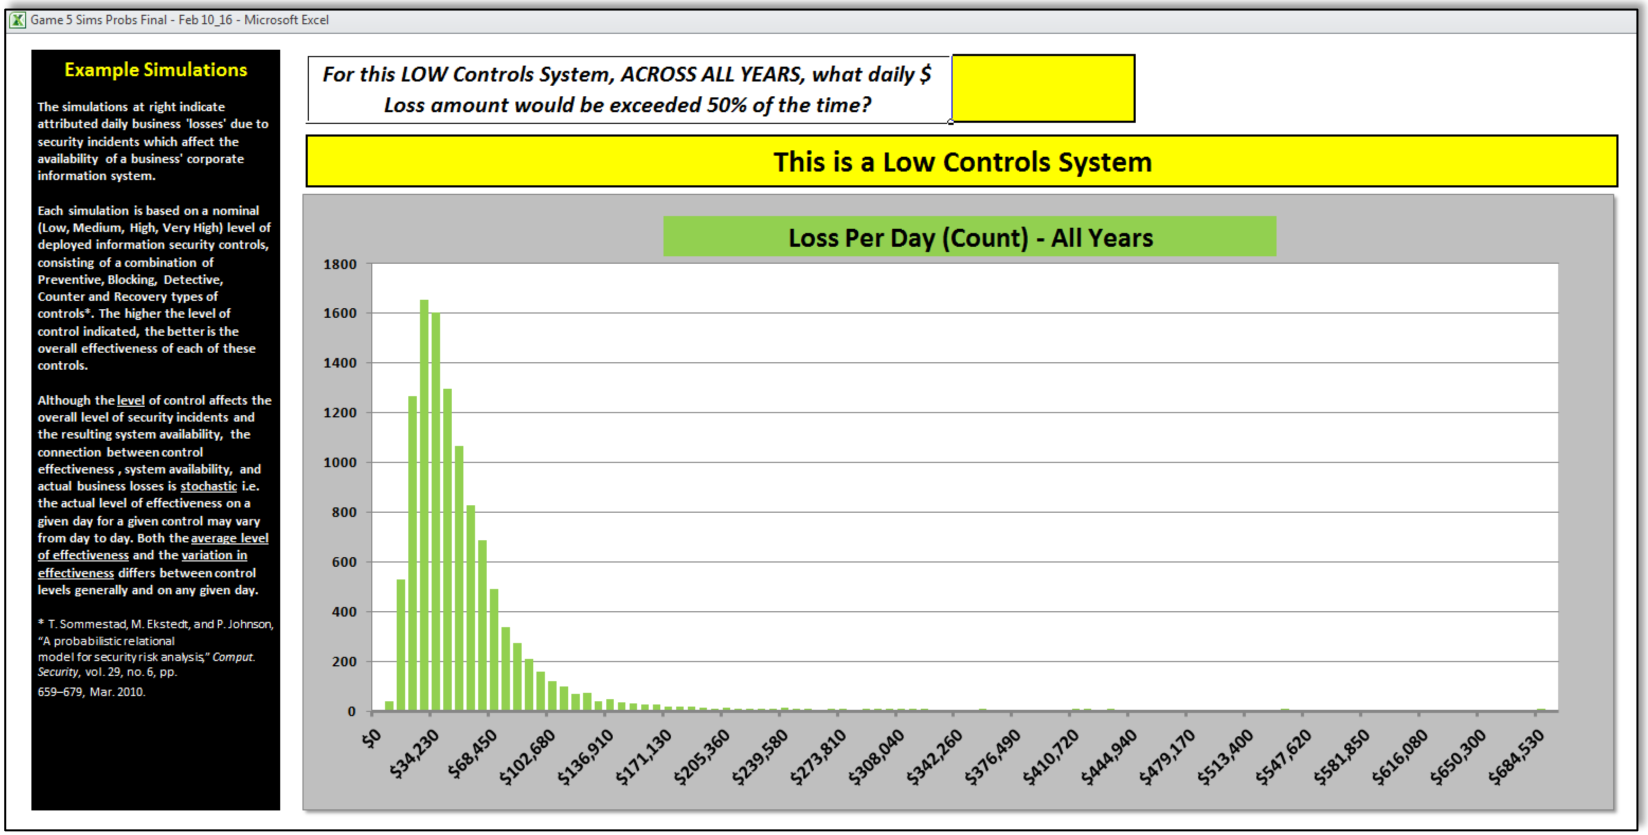

Game 5 Sims Probs Final - Feb 10_16 - Microsoft Excel

**Example Simulations**

The simulations at right indicate attributed daily business 'losses' due to security incidents which affect the availability of a business' corporate information system.

Each simulation is based on a nominal (Low, Medium, High, Very High) level of deployed information security controls, consisting of a combination of Preventive, Blocking, Detective, Counter and Recovery types of controls*. The higher the level of control indicated, the better is the overall effectiveness of each of these controls.

Although the level of control affects the overall level of security incidents and the resulting system availability, the connection between control effectiveness , system availability, and actual business losses is stochastic i.e. the actual level of effectiveness on a given day for a given control may vary from day to day. Both the average level of effectiveness and the variation in effectiveness differs between control levels generally and on any given day.

* T. Sommestad, M. Ekstedt, and P. Johnson, "A probabilistic relational model for security risk analysis," *Comput. Security*, vol. 29, no. 6, pp. 659–679, Mar. 2010.

*For this LOW Controls System, ACROSS ALL YEARS, what daily $ Loss amount would be exceeded 50% of the time?*

This is a Low Controls System

**Figure 118 - Game 4A: User Betting Interface**

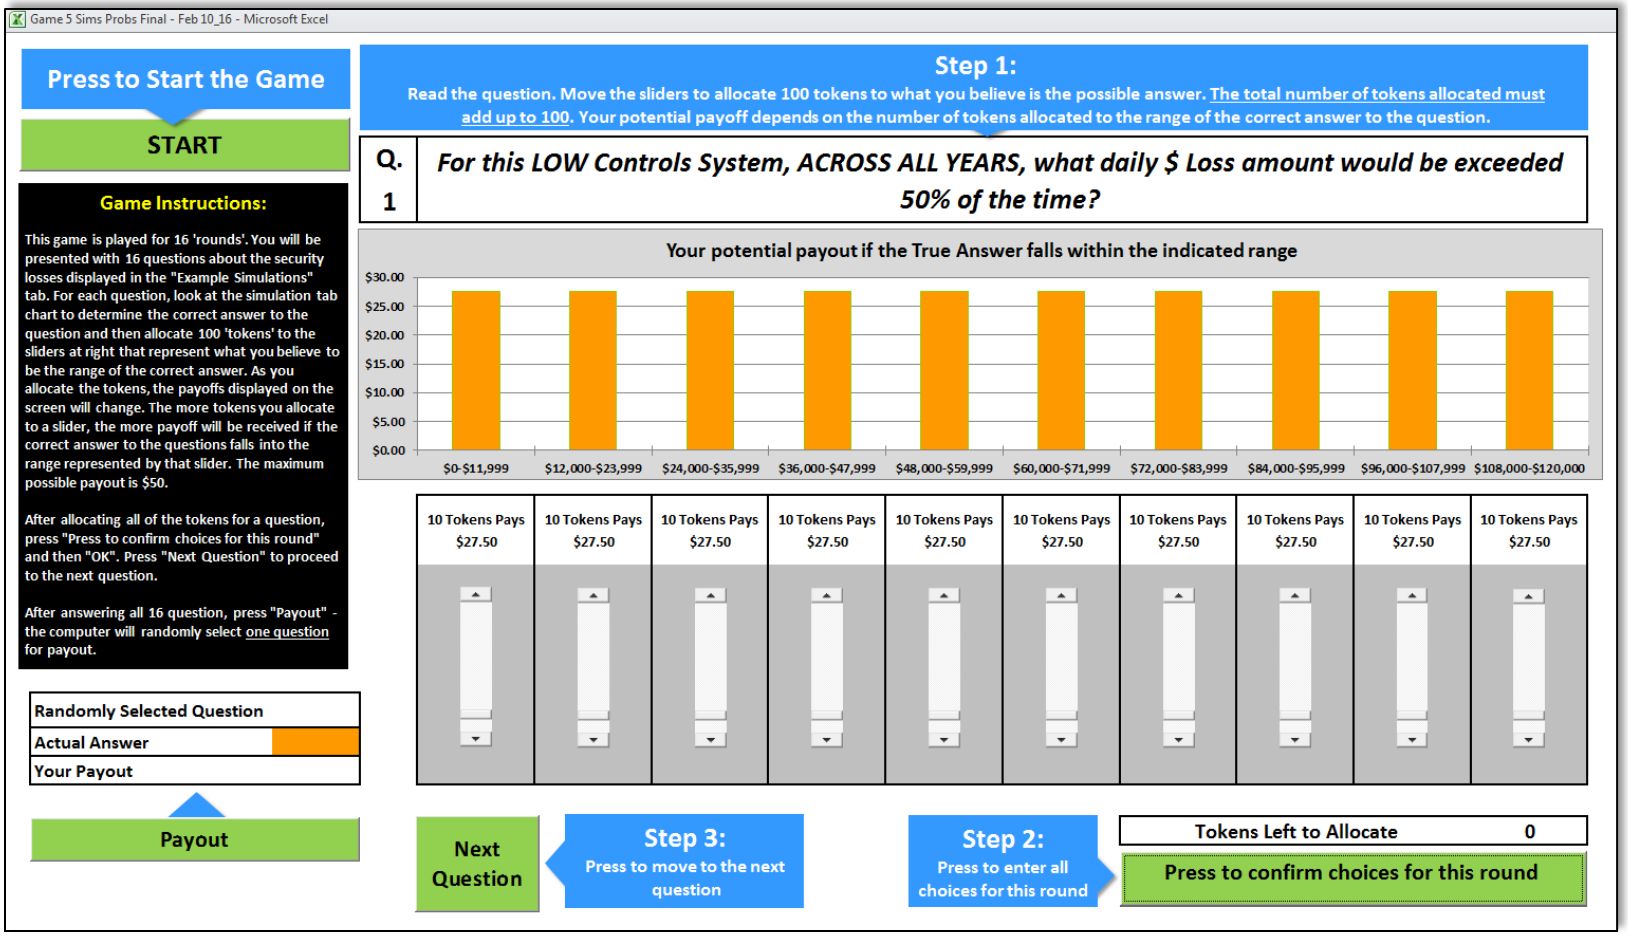

**Step 1:**

Read the question. Move the sliders to allocate 100 tokens to what you believe is the possible answer. The total number of tokens allocated must add up to 100. Your potential payoff depends on the number of tokens allocated to the range of the correct answer to the question.

**Q. 1** *For this LOW Controls System, ACROSS ALL YEARS, what daily $ Loss amount would be exceeded 50% of the time?*

**Game Instructions:**

This game is played for 16 'rounds'. You will be presented with 16 questions about the security losses displayed in the "Example Simulations" tab. For each question, look at the simulation tab chart to determine the correct answer to the question and then allocate 100 'tokens' to the sliders at right that represent what you believe to be the range of the correct answer. As you allocate the tokens, the payoffs displayed on the screen will change. The more tokens you allocate to a slider, the more payoff will be received if the correct answer to the questions falls into the range represented by that slider. The maximum possible payout is $50.

After allocating all of the tokens for a question, press "Press to confirm choices for this round" and then "OK". Press "Next Question" to proceed to the next question.

After answering all 16 question, press "Payout" - the computer will randomly select one question for payout.

Your potential payout if the True Answer falls within the indicated range

| $0-$11,999 | $12,000-$23,999 | $24,000-$35,999 | $36,000-$47,999 | $48,000-$59,999 | $60,000-$71,999 | $72,000-$83,999 | $84,000-$95,999 | $96,000-$107,999 | $108,000-$120,000 |
|---|---|---|---|---|---|---|---|---|---|
| 10 Tokens Pays<br>$27.50 | 10 Tokens Pays<br>$27.50 | 10 Tokens Pays<br>$27.50 | 10 Tokens Pays<br>$27.50 | 10 Tokens Pays<br>$27.50 | 10 Tokens Pays<br>$27.50 | 10 Tokens Pays<br>$27.50 | 10 Tokens Pays<br>$27.50 | 10 Tokens Pays<br>$27.50 | 10 Tokens Pays<br>$27.50 |

| | |
|---|---|
| Randomly Selected Question | |
| Actual Answer | |
| Your Payout | |

**Step 3:** Press to move to the next question

**Step 2:** Press to enter all choices for this round

Tokens Left to Allocate 0

In this example round, the loss display indicates a probability distribution of losses for a ‘Low’ Controls system over ‘All Years’ i.e. all possible simulation years. Based on this information, the Participant is then asked to bet on their belief regarding the loss amount above which losses would occur 50% of the time (i.e. the average loss level). In order to bet on the correct answer, the Participant is told that they have been given 100 ‘tokens’ that they can allocate to the range, or several ranges, that they believe corresponds to the correct answer to the question using the ‘slider’ interface. The payoff for each range, should the actual answer fall within that range, is based on the number of allocated tokens to the range and is interactively displayed onscreen as the Participant allocates their tokens between ranges. Participants are told that the more tokens they allocate to the correct range, the greater is the payoff. The following example allocation indicates that the Participant believes the correct answer most likely falls within one range, allocating 60 tokens to that range, but they are somewhat unsure and so allocate 20 tokens to each of the ranges immediately to the right and left of the middle range:

**Figure 119 - Game 4B: User Interface (3 Bin Betting Example)**

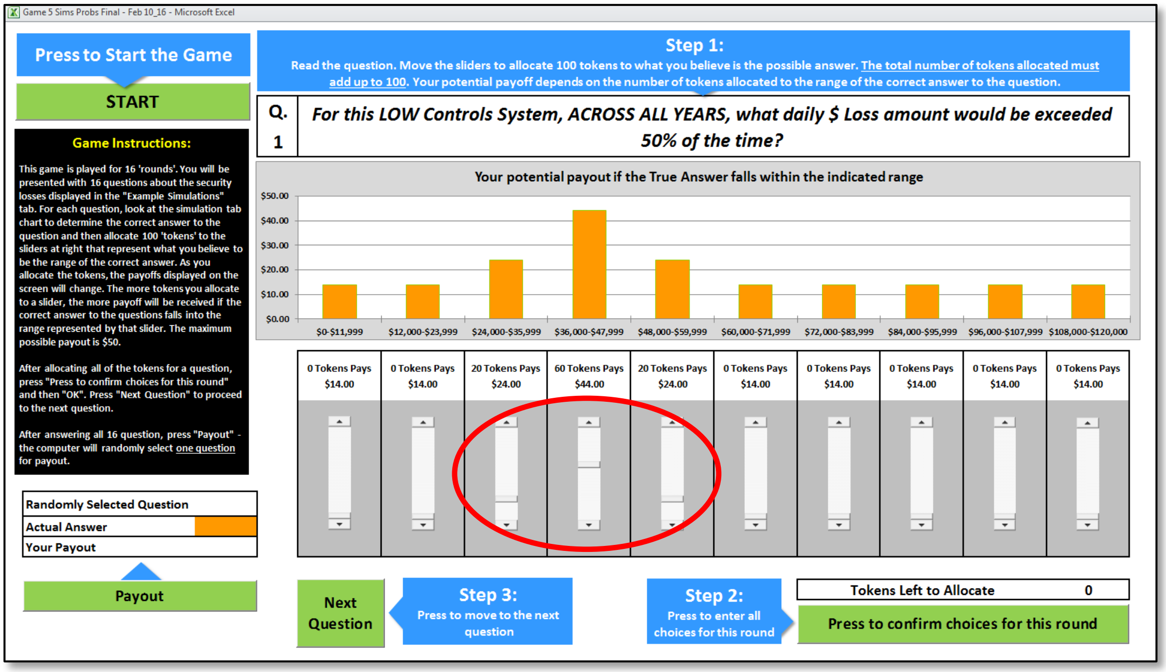

Game 5 Sims Probs Final - Feb 10_16 - Microsoft Excel

Step 1:
Read the question. Move the sliders to allocate 100 tokens to what you believe is the possible answer. The total number of tokens allocated must add up to 100. Your potential payoff depends on the number of tokens allocated to the range of the correct answer to the question.

Q. 1 For this LOW Controls System, ACROSS ALL YEARS, what daily $ Loss amount would be exceeded 50% of the time?

Game Instructions:

This game is played for 16 'rounds'. You will be presented with 16 questions about the security losses displayed in the "Example Simulations" tab. For each question, look at the simulation tab chart to determine the correct answer to the question and then allocate 100 'tokens' to the sliders at right that represent what you believe to be the range of the correct answer. As you allocate the tokens, the payoffs displayed on the screen will change. The more tokens you allocate to a slider, the more payoff will be received if the correct answer to the questions falls into the range represented by that slider. The maximum possible payout is $50.

After allocating all of the tokens for a question, press "Press to confirm choices for this round" and then "OK". Press "Next Question" to proceed to the next question.

After answering all 16 question, press "Payout" - the computer will randomly select one question for payout.

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| 0 Tokens Pays $14.00 | 0 Tokens Pays $14.00 | 20 Tokens Pays $24.00 | 60 Tokens Pays $44.00 | 20 Tokens Pays $24.00 | 0 Tokens Pays $14.00 | 0 Tokens Pays $14.00 | 0 Tokens Pays $14.00 | 0 Tokens Pays $14.00 | 0 Tokens Pays $14.00 |

| | |
|---|---|
| Randomly Selected Question | |
| Actual Answer | |
| Your Payout | |

Tokens Left to Allocate 0

Step 3: Press to move to the next question

Step 2: Press to enter all choices for this round

The scoring device which calculates the payoffs is based on a quadratic scoring rule (Matheson and Winkler 1976) as described for this experiment in Harrison (Harrison and Ulm 2015). In this scoring rule, the subject is rewarded for accuracy of their beliefs, but if that accuracy misses the true interval, “…the penalty is severe”. In contrast to the above example, if the Participant was confident that the true answer fell within one range, they might allocate all 100 tokens to that single range. The payoff in that case is rewards accuracy, but severely penalizes error:

**Figure 120 - Game 4B: User Interface (1 Bin Betting Example)**

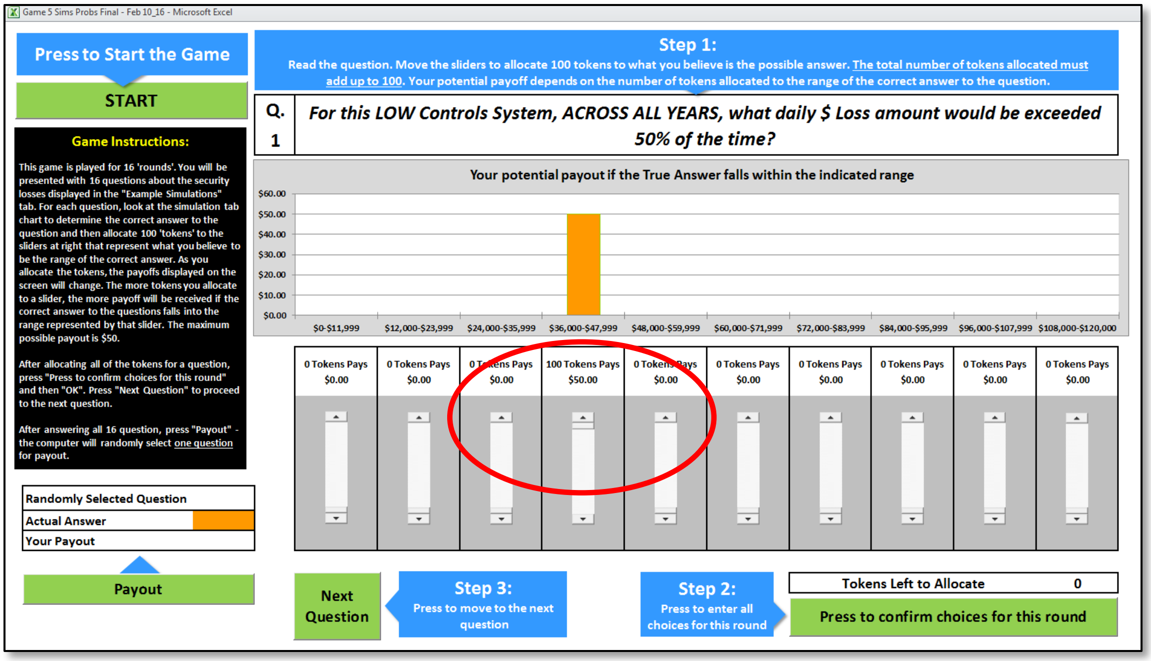

This game is played for 16 'rounds'. You will be presented with 16 questions about the security losses displayed in the "Example Simulations" tab. For each question, look at the simulation tab chart to determine the correct answer to the question and then allocate 100 'tokens' to the sliders at right that represent what you believe to be the range of the correct answer. As you allocate the tokens, the payoffs displayed on the screen will change. The more tokens you allocate to a slider, the more payoff will be received if the correct answer to the questions falls into the range represented by that slider. The maximum possible payout is $50.

After allocating all of the tokens for a question, press "Press to confirm choices for this round" and then "OK". Press "Next Question" to proceed to the next question.

After answering all 16 question, press "Payout" - the computer will randomly select one question for payout.

Step 1:
Read the question. Move the sliders to allocate 100 tokens to what you believe is the possible answer. The total number of tokens allocated must add up to 100. Your potential payoff depends on the number of tokens allocated to the range of the correct answer to the question.

Q. 1 *For this LOW Controls System, ACROSS ALL YEARS, what daily $ Loss amount would be exceeded 50% of the time?*

| 0 Tokens Pays | 0 Tokens Pays | 0 Tokens Pays | 100 Tokens Pays | 0 Tokens Pays | 0 Tokens Pays | 0 Tokens Pays | 0 Tokens Pays | 0 Tokens Pays | 0 Tokens Pays |
|---|---|---|---|---|---|---|---|---|---|
| $0.00 | $0.00 | $0.00 | $50.00 | $0.00 | $0.00 | $0.00 | $0.00 | $0.00 | $0.00 |

Gneiting et al describe the general formulation of 'proper scoring rules' (Gneiting and Raftery 2007)[63], and Andersen et al (Offerman, Sonnemans et al. 2009; Andersen, Fountain et al. 2014) describe the use of 'quadratic scoring rules' specifically for eliciting subjective beliefs over continuous probability distributions and indicate the required correction for risk attitude when recovering the corresponding subjective beliefs:

> Subjective probabilities about some event are operationally defined as those probabilities that lead an agent to make certain choices over outcomes that depend on that event. These choices could be as natural as placing a bet on a horse race, or as structured as responding to the payoffs provided by some scoring rule. In order to infer subjective probabilities from observed choices of this kind, however, one either has to make some strong assumptions about risk attitudes or jointly estimate risk attitudes and subjective probabilities…
>
> If inferred subjective probabilities are conditioned on knowing risk attitudes, then any statistical uncertainty in the estimation of risk attitudes would be expected to "propagate" into some additional uncertainty about subjective probabilities. Joint estimation allows these effects to occur, providing more reliable estimates of subjective probabilities, even if those estimates have large standard errors. In other words, it is possible that the choice task for eliciting subjective probabilities generates a point response that appears to be quite precise by itself, but which is actually not a very precise estimate of the latent subjective probability when one properly accounts for uncertainty over risk attitudes…
>
> An important example is the response to a proper scoring rule, such as the quadratic scoring rule (QSR). A respondent might select 67% when faced with a QSR, and that would be the exact

[63] "Scoring rules assess the quality of probabilistic forecasts, by assigning a numerical score based on the forecast and on the event or value that materializes. A scoring rule is proper if the forecaster maximizes the expected score for an observation drawn from the distribution F if she issues the probabilistic forecast F, rather than G 6= F. It is strictly proper if the maximum is unique. In prediction problems, proper scoring rules encourage the forecaster to make careful assessments and to be honest. In estimation problems, strictly proper scoring rules provide attractive loss and utility functions that can be tailored to the scientific problem at hand." (Gneiting and Raftery 2007)

> subjective probability if the subject were known to be exactly risk neutral. But if the subject was estimated to have risk attitudes in some interval, there would be a corresponding range of subjective probabilities to be inferred from that 67% response to the scoring rule task. If the estimate of the subject's risk attitudes spanned the possibility of risk neutrality, then the inferred subjective probability would include 67%, but would also include subjective probabilities on either side. If the estimate of the subject's risk attitudes revealed him to be clearly risk averse, with no statistically significant chance of being risk neutral, then the risk-adjusted subjective beliefs would actually be strictly higher than 67%...
>
> Formalizing this [intuition] requires that one obtain estimates of risk attitudes from a task with objective probabilities as well as from a task whose outcomes depend on some subjective probability, and then *untangle the effects of risk attitudes and subjective probability with a structural model of choice.* (Andersen, Fountain et al. 2014)

Alternatively, the Participant also faces questions regarding the time series representation of losses. Generally, we are also interested in testing whether Participants are better able to estimate the characteristics of a stochastic system using either time series or probability distributions, particularly where the accurate identification of the likelihood of extreme events is important and where evidence indicates that persons may not accurately perceive the probability of unlikely events, even when presented with distributional information about those events (Barron and Erev 2003; Hertwig, Barron et al. 2004; Weston 2014). The following illustrates the presentation of a typical time series based question:

**Figure 121 - Game 4B: High Controls Question Time Series Representation**

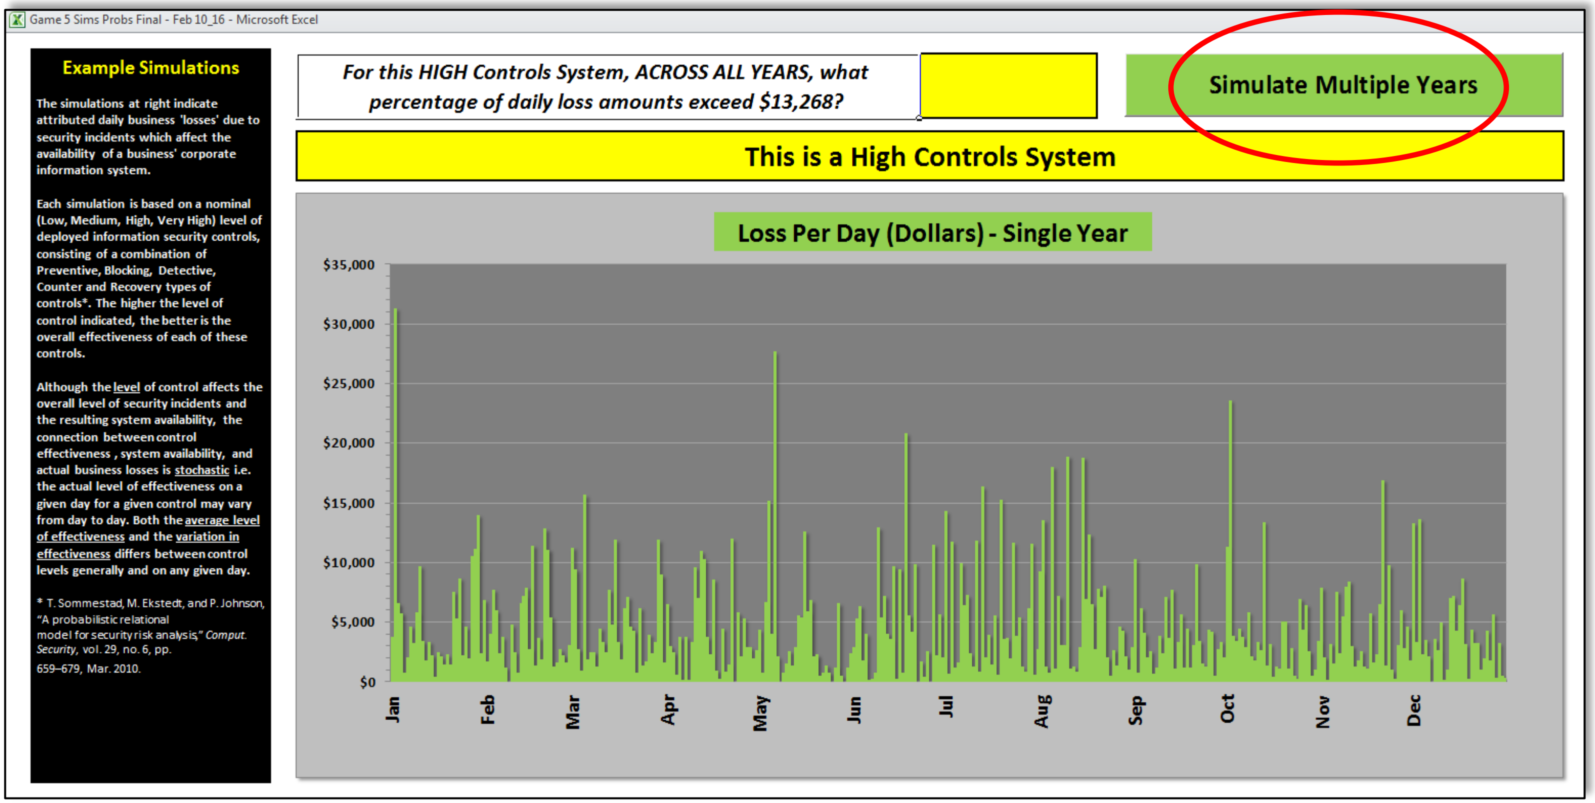

Game 5 Sims Probs Final - Feb 10_16 - Microsoft Excel

**Example Simulations**

The simulations at right indicate attributed daily business 'losses' due to security incidents which affect the availability of a business' corporate information system.

Each simulation is based on a nominal (Low, Medium, High, Very High) level of deployed information security controls, consisting of a combination of Preventive, Blocking, Detective, Counter and Recovery types of controls*. The higher the level of control indicated, the better is the overall effectiveness of each of these controls.

Although the level of control affects the overall level of security incidents and the resulting system availability, the connection between control effectiveness, system availability, and actual business losses is stochastic i.e. the actual level of effectiveness on a given day for a given control may vary from day to day. Both the average level of effectiveness and the variation in effectiveness differs between control levels generally and on any given day.

* T. Sommestad, M. Ekstedt, and P. Johnson, "A probabilistic relational model for security risk analysis," *Compute. Security*, vol. 29, no. 6, pp. 659–679, Mar. 2010.

*For this HIGH Controls System, ACROSS ALL YEARS, what percentage of daily loss amounts exceed $13,268?*

**This is a High Controls System**

In this case, the Participant must estimate the percentage of losses exceeding a certain dollar threshold. In addition, the information is presented as a series of annual simulations viewable one year at a time, where the Participant must click ‘Simulate Multiple Years’ in order to observe some (but not all) potential simulations of the ‘High Controls’ system. In this case, the information is potentially ‘doubly ambiguous’: not only does the observable profile of losses (maximum, average, etc.) change as more information is observed, but the Participant cannot be sure that they have observed the entire population of simulated outcomes (Camerer and Weber 1992). This unique multi-year simulation ‘reveal’ approach is employed for both multiple single year time series and multiple single year probability distribution presentations. I propose that this context likely corresponds to a practitioner’s real world experience where, even when some operating loss information may be available to the decision maker, the manager cannot be certain (and should be aware that it is uncertain) as to whether the observed data represents a valid *sample* of the range of possible outcomes. We are therefore interested in and able to test the extent to which the presentation of partial information (which is reflective of increased ambiguity) may affect the subjective probability report.

Participants proceed to answer 16 questions in total covering 4 x 4 x 2 = 16 treatments:

## Table 10 - Game 4B Question Set

| System Type | Single or Multi-Year Simulation | Time Series or Probability Distribution and Display Method | Question | Answer |
|---|---|---|---|---|
| **Low Controls** | **Single Year** | **Single Time Series** | **For this LOW Controls System, FOR THIS YEAR, what percentage of daily loss amounts exceed $56,444?** | **25%** |
| | **Multiple Years** | **Single Frequency Distribution** | **For this LOW Controls System, ACROSS ALL YEARS**, what daily $ Loss amount would be exceeded 50% of the time? | **$36,353** |
| | **Multiple Years** | **Multiple Time Series** | **For this LOW Controls System, ACROSS ALL YEARS, what percentage of daily loss amounts exceed $98,525?** | **5%** |
| | **Multiple Years** | **Multiple Frequency Distributions** | **For this LOW Controls System, ACROSS ALL YEARS**, what daily $ Loss amount would be exceeded 5% of the time? | **$98,525** |
| **Medium Controls** | **Single Year** | **Single Time Series** | **For this MEDIUM Controls System, FOR THIS YEAR, what percentage of daily loss amounts exceed $19,138?** | **25%** |
| | **Multiple Years** | **Single Frequency Distribution** | **For this MEDIUM Controls System, ACROSS ALL YEARS**, what daily $ Loss amount would be exceeded 50% of the time? | **$11,799** |
| | **Multiple Years** | **Multiple Time Series** | **For this MEDIUM Controls System, ACROSS ALL YEARS, what percentage of daily loss amounts exceed $36,222?** | **5%** |
| | **Multiple Years** | **Multiple Frequency Distributions** | **For this MEDIUM Controls System, ACROSS ALL YEARS**, what daily $ Loss amount would be exceeded 5% of the time? | **$36,222** |
| **High Controls** | **Single Year** | **Single Time Series** | **For this HIGH Controls System, FOR THIS YEAR, what percentage of daily loss amounts exceed $5,755?** | **25%** |
| | **Multiple Years** | **Single Frequency Distribution** | **For this HIGH Controls System, ACROSS ALL YEARS**, what daily $ Loss amount would be exceeded 50% of the time? | **$3,026** |
| | **Multiple Years** | **Multiple Time Series** | **For this HIGH Controls System, ACROSS ALL YEARS, what percentage of daily loss amounts exceed $13,268?** | **5%** |
| | **Multiple Years** | **Multiple Frequency Distributions** | **For this HIGH Controls System, ACROSS ALL YEARS**, what daily $ Loss amount would be exceeded 5% of the time? | **$13,268** |
| **Very High Controls** | **Single Year** | **Single Time Series** | **For this VERY HIGH Controls System, FOR THIS YEAR, what percentage of daily loss amounts exceed $1,048?** | **25%** |
| | **Multiple Years** | **Single Frequency Distribution** | **For this VERY HIGH Controls System, ACROSS ALL YEARS**, what daily $ Loss amount would be exceeded 25% of the time? | **$901** |
| | **Multiple Years** | **Multiple Time Series** | **For this VERY HIGH Controls System, ACROSS ALL YEARS, what percentage of daily loss amounts exceed $3,938?** | **5%** |
| | **Multiple Years** | **Multiple Frequency Distributions** | **For this VERY HIGH Controls System, ACROSS ALL YEARS**, what daily $ Loss amount would be exceeded 5% of the time? | **$3938** |

### 6 – Game #5: Effect of Risk and Ambiguity on Precaution vs. Insurance Choices

The Participants then proceed to Game 5, which is composed of 3 sub-games. Following Bajtelsmit, we present Participants with choices under uncertainty involving 1) the purchase of 'cyber insurance'; 2) the purchase of additional security controls; and 3) the purchase of both additional controls and insurance, in the context of simulated security losses. In this experiment we are interested in testing hypotheses regarding the Participant's risk preference for protection vs. insurance and to what extent these preferences change when the ambiguity of the prospects increase (Bajtelsmit, Coats et al. 2015). The experiment is based on an extension by Bajtelsmit of an original lab experiment done by Laury et al (Laury, McInnes et al. 2009) extended to include the choice of 'self-protection' (i.e. controls affecting the endogenous risk of the system) as well as insurance and is considered unique in presentation for this research:

> Previous laboratory studies have considered the effect of ambiguity on insurance purchases, and on underinsurance against low frequency, high severity losses…To our knowledge, we design the first experiment that allows for a choice between exercising a level of precaution to achieve a desired level of risk of a loss versus insuring against the loss…Laboratory experiments on insurance purchase decisions under different risk and ambiguity conditions have been conducted under a wide variety of designs and protocols and the results are highly inconclusive (Jaspersen 2015)…However, the differences in designs, procedures, and parameters employed across the studies limits the ability to generalize conclusions from their results. The Laury *et al*. experimental design [adopted for this experiment] implements a choice task to investigate the phenomenon of under insurance for low-probability, high-severity losses, and produces results that are counter to the notion that individuals ignore very low probabilities…Many studies attempt to explain insurance markets by designing the experiments as auctions rather than choice tasks. Although this design may work well as a mechanism for eliciting willingness to pay for insurance, and under a double auction, studying both sides of the insurance markets, the results are not necessarily generalizable to the insurance marketplace in which consumers face choice tasks rather than pricing tasks. (Bajtelsmit, Coats et al. 2015).

All sub-games are played over 12 rounds each. In the first sub-game the Participant is shown the following interface:

Figure 122 - Game 5: User Interface (Insurance Only Treatment)

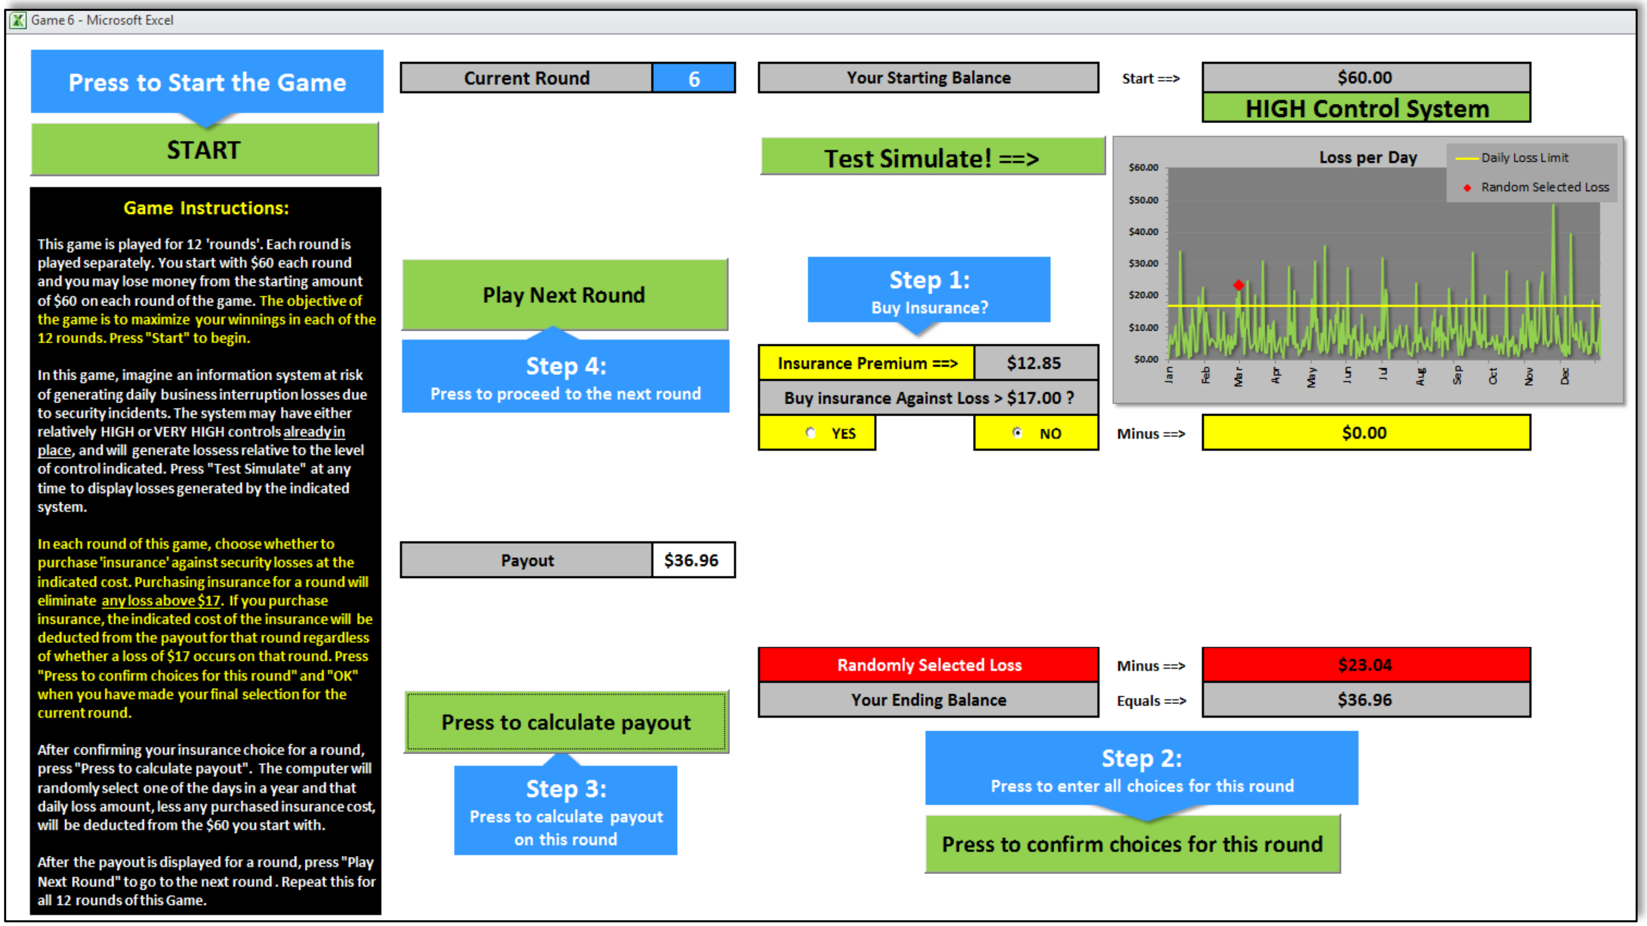

**Game Instructions:**

This game is played for 12 'rounds'. Each round is played separately. You start with $60 each round and you may lose money from the starting amount of $60 on each round of the game. The objective of the game is to maximize your winnings in each of the 12 rounds. Press "Start" to begin.

In this game, imagine an information system at risk of generating daily business interruption losses due to security incidents. The system may have either relatively HIGH or VERY HIGH controls already in place, and will generate lossess relative to the level of control indicated. Press "Test Simulate" at any time to display losses generated by the indicated system.

In each round of this game, choose whether to purchase 'insurance' against security losses at the indicated cost. Purchasing insurance for a round will eliminate any loss above $17. If you purchase insurance, the indicated cost of the insurance will be deducted from the payout for that round regardless of whether a loss of $17 occurs on that round. Press "Press to confirm choices for this round" and "OK" when you have made your final selection for the current round.

After confirming your insurance choice for a round, press "Press to calculate payout". The computer will randomly select one of the days in a year and that daily loss amount, less any purchased insurance cost, will be deducted from the $60 you start with.

After the payout is displayed for a round, press "Play Next Round" to go to the next round. Repeat this for all 12 rounds of this Game.

| | |
|---|---|
| Current Round | 6 |
| Payout | $36.96 |

| | | | |
|---|---|---|---|
| Your Starting Balance | | Start ==> | $60.00 |
| | | | HIGH Control System |
| Insurance Premium ==> | $12.85 | | |
| Buy insurance Against Loss > $17.00 ? | | | |
| ○ YES | ● NO | Minus ==> | $0.00 |
| Randomly Selected Loss | | Minus ==> | $23.04 |
| Your Ending Balance | | Equals ==> | $36.96 |

In this sub-game, the Participant is given a budget of $60 for each round and is presented with either a 'High' or a 'Very High' controls system and is asked whether they wish to purchase insurance at the indicated premium against a randomly selected single daily loss exceeding $17. The Participant can interactively simulate the system by clicking 'Test Simulate' which, similar to the above experimental treatments, randomly draws a year's worth of loss observations from a randomized 365 day period within 30 years of simulation data for that system and displays the time series on screen. Once the participant chooses to either buy or decline insurance, the computer then randomly selects a 365 day simulation for that system and a random daily loss from that sample and calculates the budget balance ($60, less any insurance costs, less the loss on the day) and displays the results onscreen. The Participant proceeds to the next round and makes 12 insurance-only choices in total. The type of system presented in each round is randomly determined, and the insurance premium offered varies with the type of system presented. Insurance is generally more expensive for the High Controls system than it is for the Very High controls system, although the 'insurance load' also varies over the rounds as well where insurance load refers to a multiple of the 'actuarially fair' insurance premium for the simulated system type[64].

In the second sub-game, the Participant is again given a budget of $60 for each round and is presented with either a 'Low' or 'High' controls starting system and is asked whether they wish to purchase additional controls at the indicated cost. In the 'Low' controls starting system case, they can purchase two higher levels of control, converting the system into either a 'High' or a 'Very High' controls system; in the 'High' controls starting system case, they can only converting the system into a 'Very High' controls system:

[64] "In the absence of the ability to take precaution against accident, theory suggests that risk-averse individuals will fully insure when actuarially fair insurance is available. In situations where insurance is not fairly priced or where precaution is an alternative, the optimal choice depends on risk aversion, load and the cost of precaution. In this paper, we show theoretically that, when the probability of loss is more ambiguous, the demand for insurance increases. However, ambiguous increases in the probability of loss may increase or decrease expenditure on precaution, depending on assumptions related to the cost and benefit of precautionary spending. We test these results empirically in a laboratory experiment in which participants make decisions about insurance and precaution under different ambiguity conditions. (Bajtelsmit, Coats et al. 2015)

Figure 123 - Game 5: User Interface (Precaution Only Treatment)

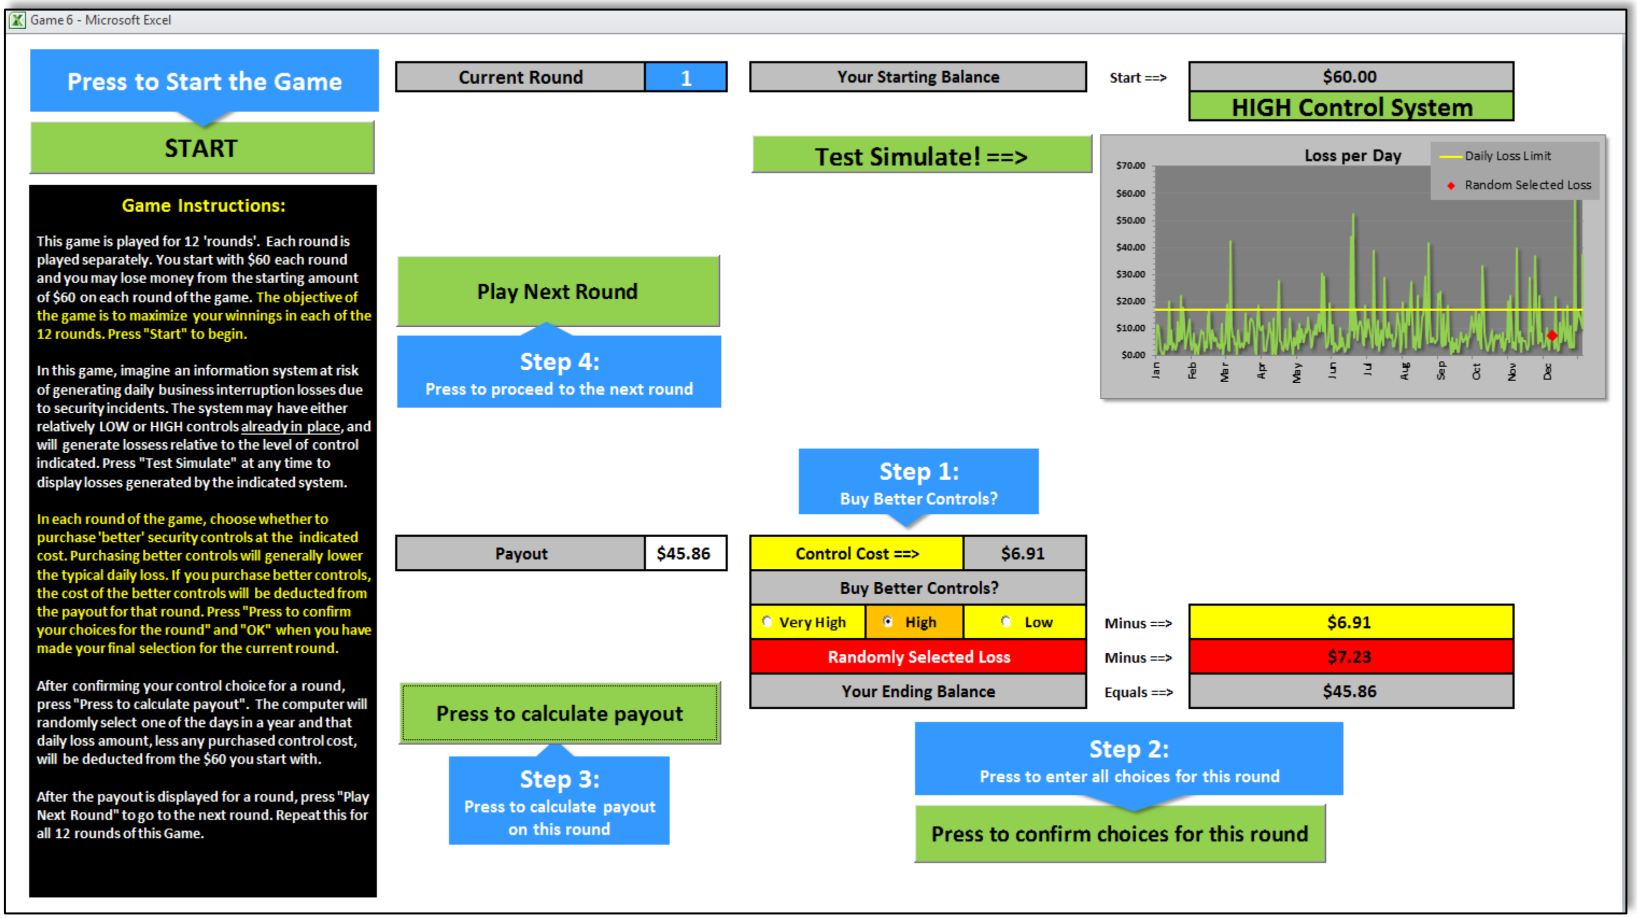

Game 6 - Microsoft Excel

Press to Start the Game

START

Game Instructions:

This game is played for 12 'rounds'. Each round is played separately. You start with $60 each round and you may lose money from the starting amount of $60 on each round of the game. The objective of the game is to maximize your winnings in each of the 12 rounds. Press "Start" to begin.

In this game, imagine an information system at risk of generating daily business interruption losses due to security incidents. The system may have either relatively LOW or HIGH controls already in place, and will generate lossess relative to the level of control indicated. Press "Test Simulate" at any time to display losses generated by the indicated system.

In each round of the game, choose whether to purchase 'better' security controls at the indicated cost. Purchasing better controls will generally lower the typical daily loss. If you purchase better controls, the cost of the better controls will be deducted from the payout for that round. Press "Press to confirm your choices for the round" and "OK" when you have made your final selection for the current round.

After confirming your control choice for a round, press "Press to calculate payout". The computer will randomly select one of the days in a year and that daily loss amount, less any purchased control cost, will be deducted from the $60 you start with.

After the payout is displayed for a round, press "Play Next Round" to go to the next round. Repeat this for all 12 rounds of this Game.

Current Round | 1

Play Next Round

Step 4:
Press to proceed to the next round

Payout | $45.86

Press to calculate payout

Step 3:
Press to calculate payout on this round

Your Starting Balance | Start ==> | $60.00

HIGH Control System

Test Simulate! ==>

Step 1:
Buy Better Controls?

Control Cost ==> | $6.91

Buy Better Controls?

○ Very High | ● High | ○ Low | Minus ==> | $6.91

Randomly Selected Loss | Minus ==> | $7.23

Your Ending Balance | Equals ==> | $45.86

Step 2:
Press to enter all choices for this round

Press to confirm choices for this round

As in the first sub-game, the Participant can interactively simulate the chosen system by clicking 'Test Simulate'. Once the participant chooses to either buy additional controls or stay with the controls of the starting system, the computer then randomly selects a 365 day simulation for that system and a random daily loss from that sample and calculates the budget balance ($60, less any controls purchased, less the loss on the day) and displays the results onscreen. The Participant proceeds to the next round and makes 12 control-only choices in total. The type of starting system presented in each round is randomly determined, and the control costs offered varies with the type of system presented. The control costs are otherwise consistent between rounds.

In the third sub-game, the Participant is again given a budget of $60 for each round and is presented with either a 'Low' or 'High' controls starting system and is asked whether they wish to purchase additional controls and/or insurance at the indicated cost. The control purchase options are the same as in sub-game 2:

Figure 124 - Game 5: User Interface (Insurance + Precaution Treatment)

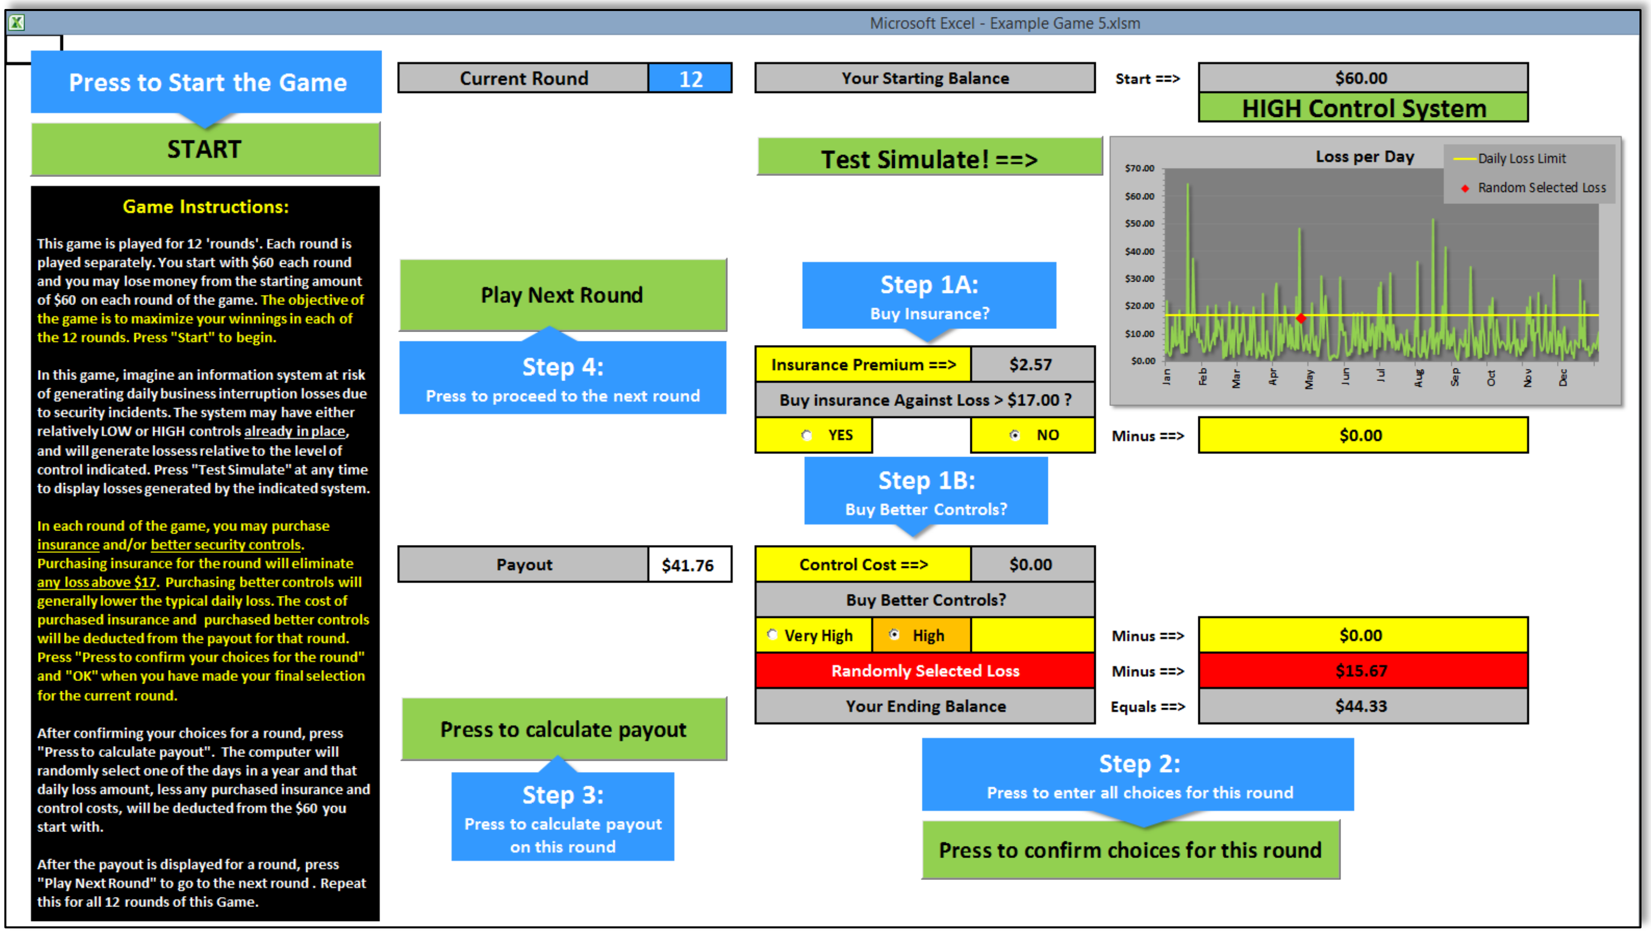

**Game Instructions:**

This game is played for 12 'rounds'. Each round is played separately. You start with $60 each round and you may lose money from the starting amount of $60 on each round of the game. The objective of the game is to maximize your winnings in each of the 12 rounds. Press "Start" to begin.

In this game, imagine an information system at risk of generating daily business interruption losses due to security incidents. The system may have either relatively LOW or HIGH controls already in place, and will generate losses relative to the level of control indicated. Press "Test Simulate" at any time to display losses generated by the indicated system.

In each round of the game, you may purchase insurance and/or better security controls. Purchasing insurance for the round will eliminate any loss above $17. Purchasing better controls will generally lower the typical daily loss. The cost of purchased insurance and purchased better controls will be deducted from the payout for that round. Press "Press to confirm your choices for the round" and "OK" when you have made your final selection for the current round.

After confirming your choices for a round, press "Press to calculate payout". The computer will randomly select one of the days in a year and that daily loss amount, less any purchased insurance and control costs, will be deducted from the $60 you start with.

After the payout is displayed for a round, press "Play Next Round" to go to the next round . Repeat this for all 12 rounds of this Game.

Current Round: 12

Payout: $41.76

Your Starting Balance — Start ==> $60.00

HIGH Control System

Insurance Premium ==> $2.57

Buy insurance Against Loss > $17.00 ?

○ YES ● NO — Minus ==> $0.00

Control Cost ==> $0.00

Buy Better Controls?

○ Very High ● High — Minus ==> $0.00

Randomly Selected Loss — Minus ==> $15.67

Your Ending Balance — Equals ==> $44.33

The insurance purchase options and the associated premium loadings are the same as in sub-game 1 with two exceptions: 1) insurance can now be purchased for the ‘Low’ controls system; and 2) the insurance premium offered is further dependent on the resulting control level of the of system as chosen by the Participant e.g. if they buy controls converting a Low system to a High/Very High system, the offered premiums are associated with a High/Very High system, with insurance loading as in sub-game 1.

As in the first sub-game, the Participant can interactively simulate the chosen system by clicking ‘Test Simulate’. Once the participant chooses to either buy additional controls and /or additional insurance, the computer then randomly selects a 365 day simulation for that system and a random daily loss from that sample and calculates the budget balance ($60, less any controls purchased, less the loss on the day) and displays the results onscreen. The Participant proceeds to the next round and makes 12 control-only choices in total.

# 9 - Results

### Analytical Methods

As was noted in Section 3 above, various structural utility models can be estimated using maximum likelihood methods which estimate the most likely value of the parameters given the data that is observed. (Quinn 2011). This differs from OLS methods, for example, which seek to estimate the parameters by minimizing the sum of squared errors between independent variable observations $y_i$ and $\beta_{\mathrm{HAT}}x_i$, but may not represent the most *likely* value of the parameters given the data observed. This is an important *analytical* distinction since the empirical objective here is to estimate a sufficiently descriptive theoretical model of choice under uncertainty for the subjects under study by allowing for observable professional and demographic characteristics and 'non-traditional' heterogeneity in the data generating process and using competing structural model(s) of choice which describe the subjects' specific risk attitudes and biases.

For my Pilot Study, Harrison provided his own study data and the corresponding STATA code with additional personal guidance which I used to learn the rudiments of maximum likelihood estimation of his particular EUT/PT mixed model specifications, along with documented support from his published papers. (Andersen, Harrison et al. 2007; Harrison and Rutström 2009). This benchmarking was successful to the point of confirming similar statistical results to his own and thereby allowing me to fully understand how to specify and estimate this form of model in STATA for my own purposes. I then replicated the model using my own specifications and data from my Pilot. In addition, another pilot was conducted in February 2016 using MBA students to generate choice data for all of the indicated experiments. The resulting data from the Asset integration experiment were used to run test analyses in STATA for EUT, PT and EUT/PT mixture models in order to confirm the ability to extract, clean and code my own experimental data and execute the STATA scripts for the baseline Harrison models. This was successful and the resulting STATA scripts were used for the analyses presented here.

### Game 1 Findings

### Analysis

Following Andersen (Andersen, Harrison et al. 2006) we assume that the observed choices were generated by latent data-generating processes, each process corresponding to a different theoretical model of choice behaviour. We examine the two most popular models of choice under uncertainty: expected utility theory (EUT) and cumulative prospect theory (CPT). We estimate a *finite mixture model* in which both processes are allowed to explain the observed choice behaviour, and allow the data to tell us the relative fraction of the sample that each process accounts for. Our analysis assumes that there is a representative EUT decision-maker and a representative CPT decision-maker across all participants, in the sense that we do not allow any heterogeneity within each type of decision-maker. This focuses attention on the heterogeneity of the processes themselves. Andersen considered three further extensions, the first which was undertaken and two others which were not undertaken in this paper: 1) Allowing for endogenous reference points for determining if some outcome is a gain or a loss; 2) Allowing for heterogeneity over time in the mix between EUT and CPT; 3) Allowing for heterogeneity within each of the decision-making processes, so that we allow some EUT (CPT) decision-makers to have difference preferences than other EUT (CPT) decision-makers. The first extension allows for variability in both the weighting between EUT and PT choice models and variability in the model parameter estimates between subjects based on their specific demographics which permits testing of a core hypothesis of this paper. For a full review of the other approaches we refer the reader to Andersen's original paper.

### A. Expected Utility Theory (EUT) Specification

We assume that utility of income is defined by the Constant Relative Risk Aversion (CRRA) specification

$$U(s,x) = (\omega s + x)^r for\ (\omega s + x) \geq 0 \tag{9.1}$$

and

$$U(s,x) = -[-(\omega s + x)]^r for\ (\omega s + x) < 0 \tag{9.2}$$

where *s* is the cumulative earnings of the subject at each choice, *x* is the lottery prize, and *ω* and *r* are parameters to be estimated. The CRRA for this specification is *r*, with *r*=1 for risk-neutral subjects, *r*<1 for risk-averse subjects, and *r*>1 for risk-loving subjects. When *ω*=0 this specification collapses to assuming that utility is defined solely over the prize for a particular choice; when *ω*=1 it collapses to assuming that utility is defined solely over the cumulative income for the sequence of tasks. Following the general arguments of Cox and Sadiraj (Cox and Sadiraj 2006) we also allow *ω* to take on any values between 0 and 1, so that we can estimate what arguments of the utility function best account for observed behaviour without imposing one or other assumption a priori (Andersen, Harrison et al. 2006).

Probabilities for each outcome k, p(k), are those that were induced by the experimental setup (i.e. for each choice round, either those probabilities pertaining to the pie chart or the corresponding 100% probability of the “sure money’ choices) so expected utility is simply the probability weighted utility of each outcome in each lottery. Since there were up to two implicit outcomes in each lottery i,

$$\mathrm{EU}_i = \sum_{k=1}^{n} [p(k) \ * \ U(k)] \text{ for n= 2} \tag{9.3}$$

A simple stochastic specification was used to specify likelihoods conditional on the model. The EU for each lottery pair was calculated for a candidate estimate of *r* and *ω*, and the index

$$\nabla\mathrm{EU} = (\mathrm{EU}_R - \mathrm{EU}_L) / \mu \tag{9.4}$$

calculated, where $\mathrm{EU}_L$ is the left lottery (pie chart) in the display, $\mathrm{EU}_R$ is the right lottery (‘Sure Money’), and μ is a Fechner noise parameter following Hey and Orme (Hey and Orme 1994). In our case, one of the lotteries can be degenerate representing a non-stochastic amount of money (the ‘sure money’ choice). Following Harrison, the index ∇EU is then used to define the cumulating probability of the observed choice using the cumulative standard normal distribution function:

$$\mathrm{G}(\nabla\mathrm{EU}) = \Phi(\nabla\mathrm{EU}) \tag{9.5}$$

Thus the likelihood, conditional on the EUT model being true, depends on the estimates of *r* and *ω* given the above specification and the observed choices. The conditional log-likelihood is

$$\ln L^{EUT}(r, \omega; y, X) = \sum_i l_i^{EUT} = \sum_i [(\ln \mathrm{G}(\nabla\mathrm{EU}) \mid y_i = 1) + \ (\ln(1 - \mathrm{G}(\nabla\mathrm{EU})) \mid y_i = 0) \tag{9.6}$$

Where $y_i$ = 1(0) denotes the choice of the right(left) lottery in task i, and X is a vector of individual demographic characteristics (e.g. male/female, math degree, smoker, etc.)

We allow each parameter to be a linear function of the observed individual characteristics of the participant. This is the X vector referred to above. Out of the demographics surveyed prior to the choice experiments, we consider six characteristics: binary (dummy) variables to identify females, the age of the participant, participants indicating having ‘business degree’, participants indicating a having ‘mathematics, economics or sciences degree’, the number of years the participant has been working in their career, the number of IT certifications held, and the participant’s current annual salary level. The estimates of each parameter in the likelihood function entails estimation of the coefficients of a linear function of these characteristics. So the estimate of r, $\hat{r}$, to be tested in the model would be

$$\hat{r} = \hat{r}_0 + (\hat{r}_{FEMALE} \times FEMALE) + (\hat{r}_{AGE} \times AGE) + (\hat{r}_{BUSINESS} \times BUSINESS) + (\hat{r}_{MATH} \times MATH) + (\hat{r}_{CERTS} \times CERTS) + (\hat{r}_{SALARY} \times SALARY)$$

where $\hat{r}_0$ is the estimate of the constant over all demographic dummy variables. The same linear function is used for ω (weighting parameter for asset integration). If we collapse this specification by dropping all individual characteristics, we would simply be estimating the constant terms for each of r (risk aversion) and ω asset integration.

The maximum likelihood estimates also allow for the possibility of correlation between responses by the same subject[65], so the standard errors on estimates are corrected for the possibility that the responses are clustered for the same subject. The use of clustering to allow for "panel effects" from unobserved individual effects is common in the statistical survey literature[66].

**Results:**

Figure 125 and Table 11 indicate the evolution of the cumulative earnings of the 57 participants over 20 rounds of choices. There is considerable variation in the earnings between participants and 7 of the participants went bankrupt. Since the first three lotteries were all in the gain frame, by construction, no subject could go bankrupt until the fourth choice.

Average earnings for the entire sample, which in this experiment represents all of the participants since no one went bankrupt, are $64.23, with a standard deviation of $31.16. There is considerable evolution of the distribution of income over the life of the experiment, as expected, and this evolution is ideal for our experimental objective of identifying the effect of changes in earnings on reference points i.e. if everyone had the same income path, there would be little variation to test.

[65] In my data sets, each individual is labelled with a unique identifier ('id') comprised of the concatenation of the hospital site (1-12) and a player number (1-5). This permits identification of unique individuals within the data set. In STATA, this permits control over errors within each individual using the 'cluster(id) option in the ml command.

[66] See (Andersen, Harrison et al. 2006) For the original discussion of this treatment in the context of the mixture model specification approach, see Harrison (Harrison and Rutström 2009)

**Figure 125 - Game 1 Cumulating Earnings per Participant**

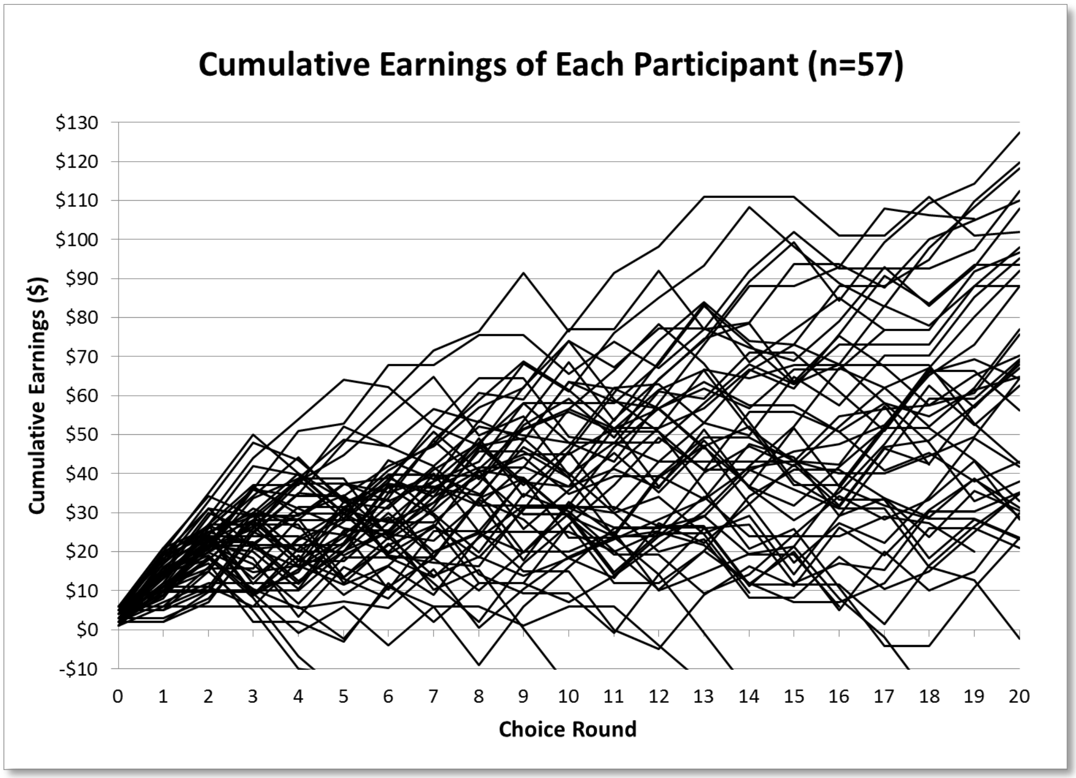

**Table 11 - Game 1 Cumulative Earnings Across All Participants**

| Round | Median | Mean | Std. Dev. | Minimum | Maximum |
|---|---|---|---|---|---|
| Start | $6.00 | $6.00 | $0.00 | $6.00 | $6.00 |
| 1 | $12.50 | $12.74 | $4.70 | $2.00 | $21.00 |
| 2 | $22.13 | $21.38 | $6.84 | $6.00 | $35.00 |
| 3 | $25.00 | $23.92 | $11.34 | $2.00 | $50.00 |
| 4 | $25.00 | $24.55 | $13.87 | -$10.00 | $54.00 |
| 5 | $28.25 | $26.31 | $15.03 | -$16.75 | $64.00 |
| 6 | $28.75 | $30.68 | $13.65 | $5.50 | $67.75 |
| 7 | $34.00 | $33.22 | $16.06 | $2.00 | $71.50 |
| 8 | $38.50 | $36.95 | $18.18 | -$9.00 | $76.50 |
| 9 | $41.50 | $40.61 | $18.87 | $0.25 | $91.50 |
| 10 | $39.00 | $41.19 | $19.39 | -$14.75 | $77.00 |
| 11 | $42.13 | $40.99 | $20.24 | $0.00 | $91.50 |
| 12 | $43.00 | $44.54 | $23.02 | -$5.00 | $98.25 |
| 13 | $46.63 | $48.16 | $23.79 | $9.00 | $111.00 |
| 14 | $43.00 | $47.44 | $26.47 | $8.25 | $111.00 |
| 15 | $43.25 | $48.35 | $27.00 | $7.00 | $111.00 |
| 16 | $43.00 | $48.66 | $26.58 | $5.00 | $101.00 |
| 17 | $52.13 | $52.06 | $28.08 | -$4.25 | $108.00 |
| 18 | $54.75 | $53.41 | $29.34 | -$17.00 | $111.00 |
| 19 | $57.50 | $59.86 | $28.32 | $10.75 | $114.25 |
| 20 | $64.25 | $64.23 | $31.16 | $21.00 | $127.50 |

Panels A - D report the maximum likelihood estimates of choice models defined by the EUT specification. We report results initially with no demographic controls, constraining the effect of asset integration to zero (ω=0) in Panel A, ω=1 in Panel B and then, in Panel C, allowing the data to sort out the degree of weighting between the accumulated winnings and the marginal prospect letting ω ϵ (0,1). We then report results in Panel D when ω ϵ (0,1) and demographic controls are added for both ω and r:

**Table 12 – Game 1 Panel A1: Maximum Likelihood Estimates Assuming EUT**

**Assuming that utility is only defined over prizes: ω=0, with no Fechner error term**

```
                                                 Number of obs   =      10400
                                                 Wald chi2(0)    =          .
Log pseudolikelihood = -6485.3642                Prob > chi2     =          .

                                    (Std. Err. adjusted for 57 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       _cons |   .2247362   .0430865     5.22   0.000     .1402882    .3091843
------------------------------------------------------------------------------
```

In Table 13, assuming no error process within the decision model, we find that the CRRA estimates relatively strong risk aversion (r=0.225<1)[67] when utility is defined over prizes only. The 95% confidence interval for CRRA under this specification is between 0.14 and 0.309.

Adjusting to allow for structural errors by incorporating the Fechner error specification noted above, in Panel A2 we re-estimate the model with the error term 'mu' and the resulting estimate of risk aversion flips to moderately risk *seeking* (r =1.116,  > 1) although the estimate is not statistically different than 1:

**Table 13 - Game 1 Panel A2: Maximum Likelihood Estimates Assuming EUT**

**Assuming that utility is only defined over prizes: ω=0, with Fechner error term**

```
                                                 Number of obs   =      10400
                                                 Wald chi2(0)    =          .
Log pseudolikelihood = -5644.2565                Prob > chi2     =          .

                                    (Std. Err. adjusted for 57 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
       _cons |   1.115974    .073428    15.20   0.000     .9720583    1.259891
-------------+----------------------------------------------------------------
mu           |
```

[67] Recall, in the Power utility model, U(x) = $x^r$, *r*=1 for risk-neutral subjects, *r*<1 for risk-averse subjects, and *r*>1 for risk-seeking subjects.

```
       _cons |   10.72508   2.052223     5.23   0.000     6.702792    14.74736
------------------------------------------------------------------------------
```

The incorporation of a parameter allowing for noise or errors in the structural model of choice (‘mu’ or $\mu$) is considered to improve the estimation of the risk aversion coefficient and is retained throughout the rest of the analysis unless otherwise indicated. In Panel A3 we introduce a single dummy variable indicating whether the participant is female (sex = 1). The resulting parameter is not statistically significant, but the parameter estimate indicates that females may be somewhat more risk averse than males:

**Table 14 - Game 1 Panel A3: Maximum Likelihood Estimates Assuming EUT**

**Assuming that utility is only defined over prizes: ω=0, and demographic dummy (sex (Female) = 1)**

```
                                                  Number of obs   =      10400
                                                  Wald chi2(1)    =       2.54
Log pseudolikelihood = -5621.3974                 Prob > chi2     =     0.1110

                                    (Std. Err. adjusted for 57 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
         sex |  -.1062383   .0666565    -1.59   0.111    -.2368827    .0244061
       _cons |   1.142971   .0721315    15.85   0.000     1.001596    1.284346
-------------+----------------------------------------------------------------
mu           |
       _cons |   10.28048   2.032033     5.06   0.000     6.297765    14.26319
------------------------------------------------------------------------------
```

In Panel B, we next examine the same data through the lens of an EUT model that assumes that the argument of utility is cumulative income, initially with ω=1. Here infer that subjects remain risk-seeking although no more so than if cumulative income is ignored as in Panels A above: the point estimate of *r* in this case is 1.15, and the 95% confidence interval does not quite include r=1. The decrease in the log-likelihood value in Panel B vs. Panel A2 suggests that the inclusion of cumulative income in the total prize amounts is a marginally better specification from an EUT perspective.

**Table 15 - Game 1 Panel B: Maximum Likelihood Estimates Assuming EUT**

**Assuming that utility is defined over cumulative income and prizes: ω=1**

```
                                                  Number of obs   =      10400
                                                  Wald chi2(0)    =          .
Log pseudolikelihood = -5630.6826                 Prob > chi2     =          .

 ( 1)  [omega]_cons = 1
                                    (Std. Err. adjusted for 57 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
```

```
r            |
       _cons |   1.151501   .0671945    17.14   0.000     1.019802      1.2832
-------------+----------------------------------------------------------------
omega        |
       _cons |          1          .        .       .            .           .
-------------+----------------------------------------------------------------
mu           |
       _cons |   12.75556   3.240903     3.94   0.000     6.403504    19.10761
------------------------------------------------------------------------------
```

In Panel C, ω is allowed to take on any value between 0 and 1. The estimated parameter (‘omega_’) is actually the log odds of the term we are interested in (‘omega’), in order to ensure that the resulting weight falls between zero and 1, where:

`omega' = 1 / (1+exp(`omega_'))     (9.7)

The log odds parameter estimate is then converted back to the weight in the range (0,1) from the using the STATA command ‘nlcom’. This convention is also used throughout wherever a term is required to fall within the range (0,1).[68]

**Table 16 - Game 1 Panel C: Maximum Likelihood Estimates Assuming EUT**
**Assuming that utility is defined over a weighted average of cumulative income and prizes: ω϶(0,1)**

```
                                                  Number of obs   =      10400
                                                  Wald chi2(0)    =          .
Log pseudolikelihood = -5628.9519                 Prob > chi2     =          .

                                     (Std. Err. adjusted for 57 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
       _cons |   1.155826   .0688995    16.78   0.000     1.020785    1.290866
-------------+----------------------------------------------------------------
omega_       |
       _cons |    .701709   .0324723    21.61   0.000     .6380645    .7653535
-------------+----------------------------------------------------------------
mu           |
       _cons |   13.25913   3.505291     3.78   0.000     6.388891    20.12938
```

[68] For example, in Games 2 and 3 where we are interested in estimating the subjective probability estimate of the participants, the estimated probability is similarly constrained using the log odds form of the parameter within the model, and converted back from the log value to evaluate the resulting estimated parameter which is expected to fall between 0 and 1. According to Harrison:

> “…we allow non-linear transforms of the core parameter to be estimated. This allows STATA to happily maximize over some variable that ranges from -∞ to +∞, but for the underlying parameter in the economic model to be constrained. The error term [for example] can be constrained to be non-negative, by estimating the (natural) log of the parameter and then converting to the underlying parameter we want using generate double `mu' = exp(`LNmu'). Similarly, we can constrain the estimates of [probability estimates] to lie in the open interval (0,1) by estimating the log odds transform of the parameters we are really interested in, and converting using generate double `u5' = 1/(1+exp(`u5_'))…This transform is one of several used by statisticians; for example, see Rabe-Hesketh and Everitt [2004; p. 244].” {Harrison, 2008 #48}

```
------------------------------------------------------------------------------

. nlcom (omega: 1/(1+exp([omega_]_b[_cons])))

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       omega |   .3314334   .0071954    46.06   0.000     .3173307    .3455361
------------------------------------------------------------------------------
```

In this result participants are similarly risk seeking (r = 1.16, >1) but the estimated omega term (.331) indicates that approximately 33% of the decision weight is attributed to the accumulated earnings as opposed to the value of the prizes presented on each choice round.

Panel D extends the general specification in Panel C to allow for demographic variation in both ω and r. After controlling for demographic effects, participants still appear to be slightly risk *seeking* (i.e. r_ = 1.11, > 1). While only age and sex are significant at the 95% level, age, sex (female), business degrees, and IT certifications all increase risk aversion, while a math degree, increased work years and income all lower risk aversion on average.  After controlling for demographics, the average weighting for cumulative earnings drops to .122, however there is substantial demographic heterogeneity with respect to both the risk aversion parameter r and the weighting parameter ω:

**Table 17 - Game 1 Panel D: Maximum Likelihood Estimates Assuming EUT**

**Assuming that utility is defined over a weighted average of cumulative income and prizes: ωϵ(0,1) with demographics**

```
                                                  Number of obs   =      10400
                                                  Wald chi2(7)    =      11.57
Log pseudolikelihood = -5541.0479                 Prob > chi2     =     0.1156

                                     (Std. Err. adjusted for 57 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
         age |  -.0974969   .0447677    -2.18   0.029      -.18524   -.0097539
         sex |  -.1796082   .0668553    -2.69   0.007    -.3106423   -.0485742
         bus |  -.0470725   .0464241    -1.01   0.311    -.1380621    .0439172
        math |   .0138011   .0425507     0.32   0.746    -.0695968     .097199
   workyears |   .0460034   .0250433     1.84   0.066    -.0030806    .0950873
       certs |  -.0215968   .0146979    -1.47   0.142    -.0504042    .0072106
      income |   .0040181   .0379658     0.11   0.916    -.0703934    .0784297
       _cons |    1.10624   .0905453    12.22   0.000      .928774    1.283705
-------------+----------------------------------------------------------------
omega_       |
         age |   56.00383   9.768628     5.73   0.000     36.85767    75.14999
         sex |   20.10691   2.927977     6.87   0.000     14.36818    25.84564
         bus |   15.08322   1.751977     8.61   0.000     11.64941    18.51703
        math |  -10.31344   3.496679    -2.95   0.003     -17.1668   -3.460074
   workyears |  -20.24997   2.930177    -6.91   0.000    -25.99301   -14.50693
       certs |  -30.53077   5.313834    -5.75   0.000     -40.9457   -20.11585
      income |  -30.85692   5.314535    -5.81   0.000    -41.27322   -20.44063
```

```
       _cons |   1.971869   .1051358    18.76   0.000     1.765807    2.177932
-------------+----------------------------------------------------------------
mu           |
       _cons |   6.736277    1.53528     4.39   0.000     3.727183     9.74537
------------------------------------------------------------------------------

. nlcom (omega: 1/(1+exp([omega_]_b[_cons])))

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       omega |   .1221883   .0112767    10.84   0.000     .1000864    .1442902
------------------------------------------------------------------------------
```

Figures 126 and 127 illustrate the kernel density functions of the predicted values of CRRA and Omega across all participants, indicating significant diversity across participants. CRRA estimates indicate participants range from slightly risk seeking to very risk averse. Omega is clearly multi-modal where a substantial proportion of participants effectively do not substantially consider cumulative earnings when making bets (omega < .2), whereas the next highest proportion of participants weight cumulative earnings in excess of 90% (omega >.9):

**Figure 126 - Kernel density of the Game 1 CRRA estimates, with demographics**

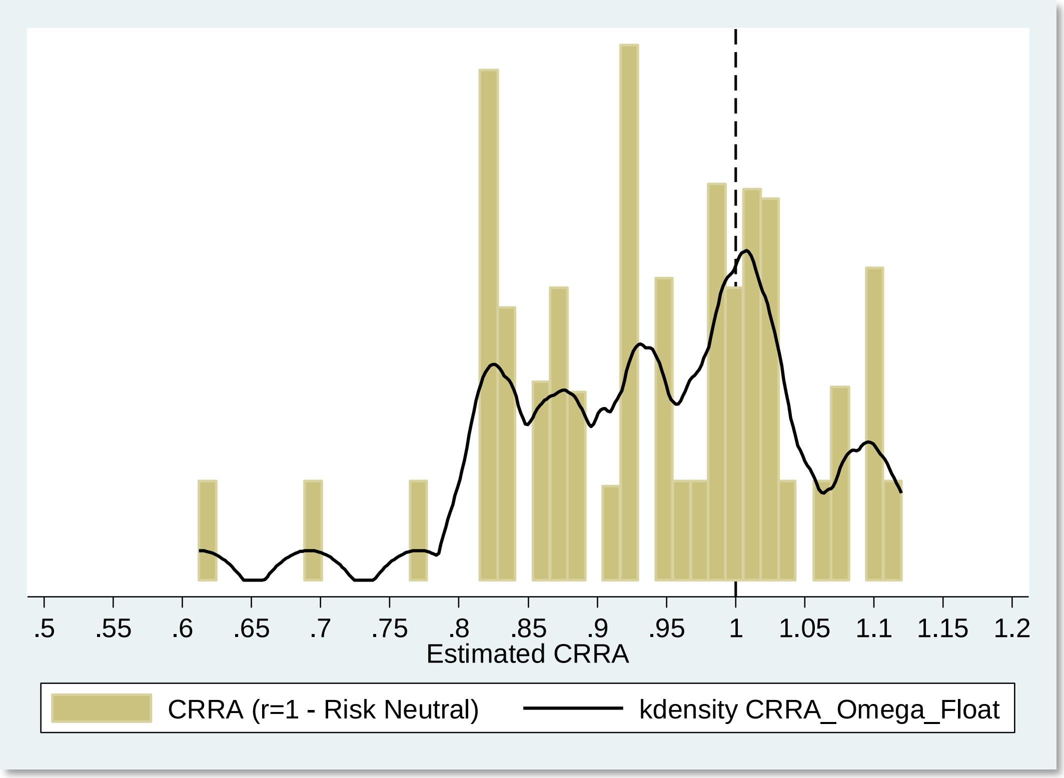

**Figure 127 - Kernel density of the Game 1 Omega estimates, with demographics**

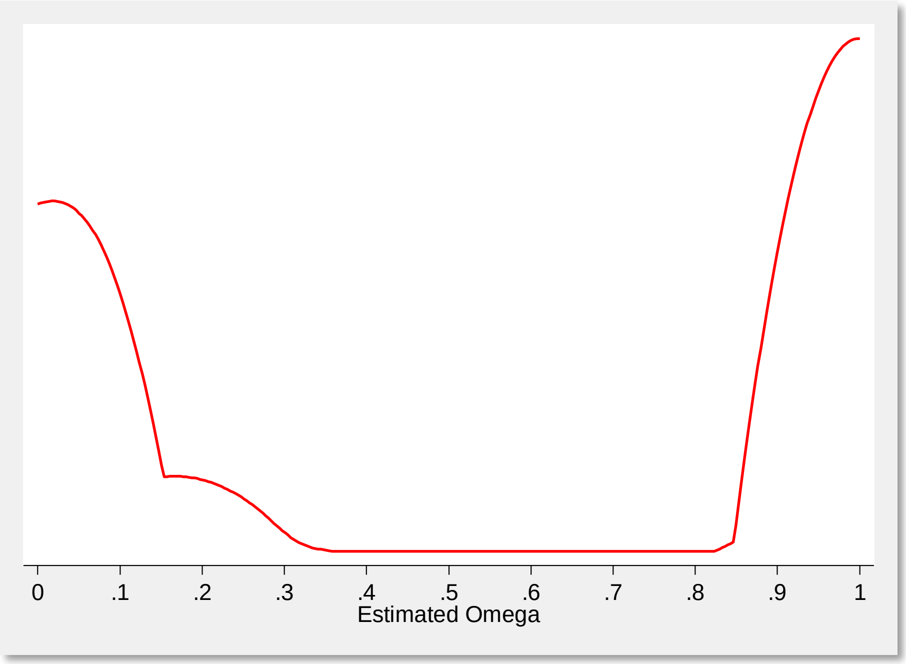

## Cumulative Prospect Theory (CPT) Specification

The Cumulative Prospect Theory (CPT) model is an extension of the EUT model to account for probability weighting and loss aversion relative to a reference point. We assume that utility is defined by the CRRA specification

$$U(x) = x^{\alpha} \text{ for } x \geq \chi \qquad (9.8)$$

and

$$U(x) = -\lambda(-x)^{\beta} \text{ for } x \geq \chi \qquad (9.9)$$

where **alpha** (α) is the CRRA parameter in the gain domain, **beta** (β) is the CRRA parameter in the loss domain, **lambda** (λ) is the loss aversion parameter, and χ is the reference point used to define if lottery prize x is a gain or a loss. When χ=0 this specification collapses to assuming that utility is defined solely over the prize for a particular choice, and that the subject determines gain and loss frames directly from the sign of the prize: that is, the reference point is $0. When χ>0 we allow the subject to characterize prizes x as gains or losses even if x>0. We discuss the specification of χ further below.

In the PT model the decision maker is assumed to employ weighted probabilities rather than the probabilities induced by the experimenter, although a special case is where the objective probabilities are not weighted. Tversky and Kahneman(Tversky and Kahneman 1992) posit a simple *weighting function* proposed by

$$\text{w(p)} = p^{\gamma} /[p^{\gamma} + (1-p)^{\gamma}]^{1/\gamma} \tag{9.10}$$

for induced probabilities p and parameter **gamma** (γ). When γ=1 this function collapses to the standard EUT specification that w(p) = p. When γ<1, the usual case, the decision maker exhibits overweighting of low probabilities and underweighting of higher probabilities, with fixed point w(p)=p at p=1/3.

Assuming that SPT is the true model, prospective utility PU is defined in much the same manner as when EUT is assumed to be the true model. The PT utility function is used instead of the EUT utility function, and w(p) is used instead of p, but the steps are otherwise identical. Under CPT, however, there is an additional step required to transform the w(p) values into decision weights. The weighting function is given by w(p) above, but the prospective utility of lottery i under CPT is now defined as

$$\text{PU}_i = [\ \text{w}(\text{p}_2)\text{U}(\text{k}_2) + (1\text{- w}(\text{p}_2))\text{U}(\text{k}_1)] \tag{9.11}$$

where $k_2 > k_1$. The difference in prospective utilities under CPT is then defined in the same way as under EUT. The index

$$\nabla\text{PU} = (\text{PU}_R - \text{PU}_L) / \mu \tag{9.12}$$

is calculated, where PU replaces EU in the EUT specification. Again, this index of differences in prospective utility can be transformed into a cumulative distribution using

$$\text{G}(\nabla\text{PU}) = \Phi(\nabla\text{PU}) \tag{9.12}$$

The likelihood, conditional on the CPT model being true, depends on the estimates of α, β, λ and γ given the above specification and the observed choices. The conditional log-likelihood is

$$\ln L^{PT}(\alpha, \beta, \lambda, \gamma\, ; y, X) = \textstyle\sum_i l_i^{PT} = \sum_i[(\ln \text{G}(\nabla\text{PU}) \mid y_i = 1) + \ (\ln(1 - \text{G}(\nabla\text{PU})) \mid y_i = 0)] \tag{9.13}$$

Again, each parameter can be estimated as a linear function of a vector of characteristics, X. This includes the individual demographic characteristics used for the EUT model.

It should be noted that PT differs from EUT in many respects, but one that is important for present purposes is the treatment of reference points. PT is quite clear that the argument of the utility function is to be the prospect over which the choice is being made and not some broader concept of income or lifetime wealth. Thus, allowing the reference point to be non-zero ($\chi \geq 0$) is distinct from assuming that the argument of the utility function is something other than the immediate prospect.

Panel E reports maximum likelihood estimates of the CPT model assuming homogenous agents (no demographics) and a zero wealth reference point:

**Table 18 – Game 1 Panel E: Maximum Likelihood Estimates Assuming CPT**

**Assuming zero is the wealth reference point: ω=0**

```
                                                  Number of obs   =       9610
                                                  Wald chi2(0)    =          .
Log pseudolikelihood = -5747.0906                 Prob > chi2     =          .

                                    (Std. Err. adjusted for 52 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
alpha        |
       _cons |   .3264526   .0353169     9.24   0.000     .2572329    .3956724
-------------+----------------------------------------------------------------
beta         |
       _cons |   1.396031   .1464688     9.53   0.000     1.108957    1.683104
-------------+----------------------------------------------------------------
lambda       |
       _cons |   .0246151   .0127038     1.94   0.053     -.000284    .0495142
-------------+----------------------------------------------------------------
gamma        |
       _cons |   .9271024   .0411127    22.55   0.000      .846523    1.007682
------------------------------------------------------------------------------
```

The first result is that subjects have coefficients for α and β that are generally consistent with them being risk averse over gains ($\alpha < 1$) and risk seeking over losses ($\beta > 1$). The second result is that the probability weighting parameter gamma (γ) is estimated to be 0.927, indicating that participants probability weight in the usual fashion: overweighting low probabilities and underweighting higher probabilities, although the estimate is only marginally different than 1 at the 95% level. In this model the loss aversion parameter lambda (λ) is estimated to be not statistically different than zero indicating no loss aversion.

Panel F reports maximum likelihood estimates of the CPT model assuming heterogeneous agents and allowing for demographic variation in the estimated parameters. The maximum likelihood model specified over all possible demographic variables did not converge, so the model was re-estimated for a subset of the demographic attributes: sex, bus, math, and income:

**Table 19 - Game 1 Panel F: Maximum Likelihood Estimates Assuming CPT Allowing for Demographic Heterogeneity**

```
                                                  Number of obs   =       9610
                                                  Wald chi2(4)    =       8.08
Log pseudolikelihood = -5601.8098                 Prob > chi2     =     0.0888

                                    (Std. Err. adjusted for 52 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
alpha        |
         sex |  -.1397724   .0724272    -1.93   0.054     -.281727    .0021823
         bus |  -.0823576   .0821773    -1.00   0.316    -.2434221     .078707
        math |   .0744085   .0737446     1.01   0.313    -.0701283    .2189454
      income |   -.098969   .0564576    -1.75   0.080    -.2096239    .0116858
       _cons |   .5175003   .0855237     6.05   0.000     .3498769    .6851236
-------------+----------------------------------------------------------------
beta         |
         sex |   .5593686   .3446937     1.62   0.105    -.1162188    1.234956
         bus |  -.3117955   .2474982    -1.26   0.208     -.796883     .173292
        math |   .1162828   .2070025     0.56   0.574    -.2894347    .5220002
      income |   .1251098   .1465526     0.85   0.393    -.1621281    .4123476
       _cons |     1.0502   .2762887     3.80   0.000     .5086845    1.591716
-------------+----------------------------------------------------------------
lambda       |
         sex |  -.0472633   .0249525    -1.89   0.058    -.0961692    .0016427
         bus |    .024558   .0221931     1.11   0.268    -.0189397    .0680557
        math |   .0040453   .0052592     0.77   0.442    -.0062625    .0143532
      income |  -.0125178   .0114439    -1.09   0.274    -.0349473    .0099118
       _cons |   .0665895    .030895     2.16   0.031     .0060365    .1271425
-------------+----------------------------------------------------------------
gamma        |
         sex |   .0140913   .1129323     0.12   0.901    -.2072519    .2354345
         bus |  -.0018316   .0900635    -0.02   0.984    -.1783527    .1746895
        math |  -.1613296   .1167301    -1.38   0.167    -.3901163    .0674572
      income |   .0916806   .0604511     1.52   0.129    -.0268014    .2101627
       _cons |   .8313479   .1213022     6.85   0.000        .5936    1.069096
 ------------------------------------------------------------------------------
```

Here allowing for demographic heterogeneity, coefficients for α and β draw closer together and now indicating risk aversion in gains (α < 1). Risk aversion in losses has decreased and is now not statistically different than 1 (β = 1). The probability weighting parameter gamma (γ) has also decreased indicating no overweighting/underweighting of low/high probabilities respectively. The loss aversion parameter lambda (λ) has increased and is now statistically greater than zero indicating a slight tendency towards loss aversion.

Figure 128 illustrate the kernel density functions for alpha and beta estimates across all participants, indicating significant diversity across participants:

**Figure 128 - Kernel density for Game 1 CPT alpha and beta estimates, with demographics**

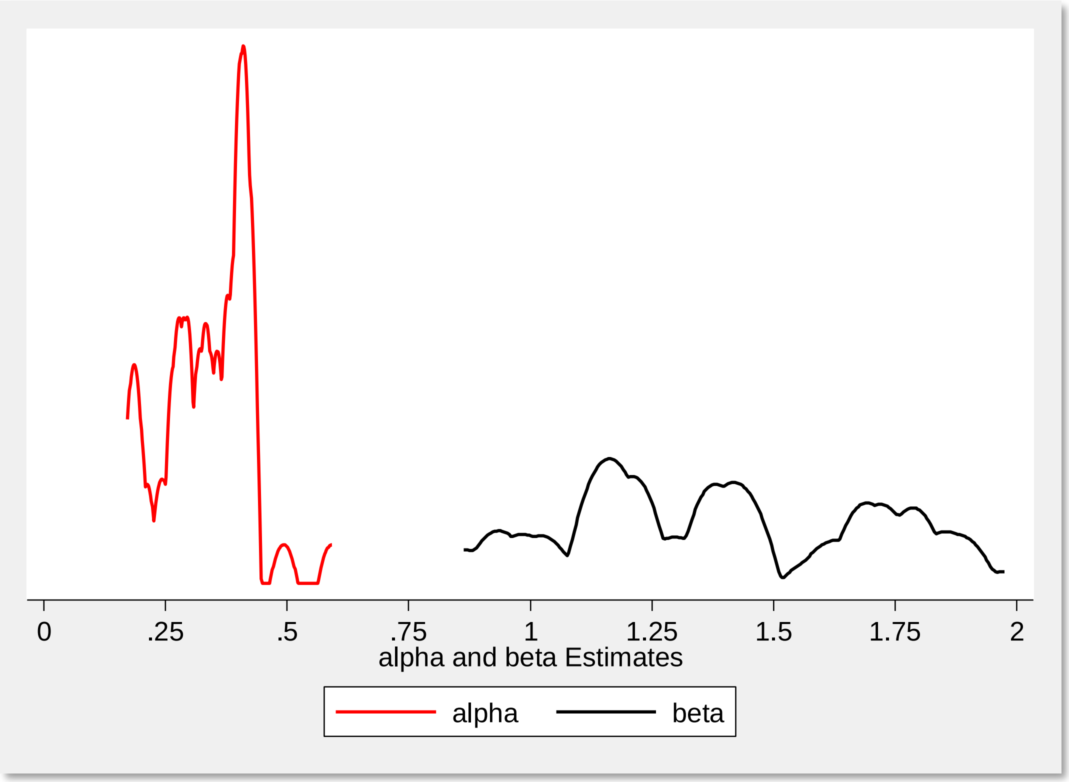

## B. Mixture (EUT + CPT) Model Specification

If we let $\pi^{EUT}$ denote the probability that the EUT model is correct, and $\pi^{PT} = (1\text{-}\pi^{EUT})$ denote the probability that the PT model is correct, a *grand likelihood* can be written as the probability weighted average of the conditional likelihoods. The likelihood for the overall model is then defined by

$$\ln L(r, w, \alpha, \beta, \lambda, \gamma, \chi, \pi^{EUT}; y, X) = \sum_i ln[(\pi^{EUT} * l_i^{EUT}) + (\pi^{PT} * l_i^{PT})] \qquad (9.14)$$

This log-likelihood can be maximized to find estimates of the parameters.

This approach assumes that any one observation can be generated by both models, although it also permits extremes in which one or other model wholly generates the observation. One could alternatively define a grand likelihood in which observations or subjects are classified as following one model or the other on the basis of the latent probabilities $\pi^{EUT}$ and $\pi^{PT}$. We will interpret estimates of $\pi^{EUT}$ and hence $\pi^{PT}$ as reflecting the fraction of the observations consistent with each model, and can interpret these probabilities as the fraction of subjects consistent with each model. These interpretations are perfectly consistent with the statistical specification above, but they are interpretations and there are others that are possible.

Panel G indicates the maximum likelihood estimates, assuming that the EUT and PT agents are each homogeneous where the only heterogeneity that is allowed in this model is with respect to the latent decision generating model EUT or PT:

**Table 20 - Game 1 Panel G: Maximum Likelihood Estimates for EUT/PT Mixture Model Assuming wealth point is a weighted average of cumulative income and prizes: ωϵ(0,1)**

```
                                                  Number of obs   =       6650
                                                  Wald chi2(0)    =          .
Log pseudolikelihood = -3319.9808                 Prob > chi2     =          .

                                    (Std. Err. adjusted for 37 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
alpha        |
       _cons |   1.64e-09          .        .       .            .           .
-------------+----------------------------------------------------------------
beta         |
       _cons |   .0613272   .2947925     0.21   0.835    -.5164554    .6391099
-------------+----------------------------------------------------------------
lambda       |
       _cons |  -.4498638   .2762334    -1.63   0.103    -.9912712    .0915436
-------------+----------------------------------------------------------------
gamma        |
       _cons |   1.013418   .0152242    66.57   0.000     .9835788    1.043257
-------------+----------------------------------------------------------------
r            |
       _cons |    .911611    .051465    17.71   0.000     .8107414    1.012481
-------------+----------------------------------------------------------------
omega_       |
       _cons |  -36.86966   10.72313    -3.44   0.001    -57.88662    -15.8527
-------------+----------------------------------------------------------------
kappa        |
       _cons |  -.6327859   .2171831    -2.91   0.004    -1.058457   -.2071149
------------------------------------------------------------------------------
       omega:  1/(1+exp([omega_]_b[_cons]))
        pEUT:  1/(1+exp([kappa]_b[_cons]))
         pPT:  exp([kappa]_b[_cons])/(1+exp([kappa]_b[_cons]))
------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       omega |          1          .        .       .            .           .
        pEUT |   .6531209   .0492037    13.27   0.000     .5566834    .7495584
         pPT |   .3468791   .0492037     7.05   0.000     .2504416    .4433166
------------------------------------------------------------------------------
* test probEUT = probPT
. testnl 1/(1+exp([kappa]_b[_cons])) -
exp([kappa]_b[_cons])/(1+exp([kappa]_b[_cons])) = 0

chi2(1) =        9.68
Prob > chi2 =    0.0019
```

The first result is that the EUT model is estimated to have 65% support (pEUT), indicating that the latent data generating process assumed by EUT accounts for approximately 2/3 of the choice observations ,with the CHI square test of the difference between the EUT and PT weightings significant at the 95% level.

The second result is that EUT subjects tend to be risk averse, with a CRRA of .912. This is lower than the estimates shown in Panel C ($r = 1.14$), which assumed that EUT characterized every observation and not just a percentage of the observations. The third result is that the EUT subjects tend to consider the prizes they are faced in conjunction with cumulative income when making decisions (omega ~= 1), again in marked qualitative contrast to the findings when we assumed that EUT characterized every observation. From Panel C we have estimates of $\omega$ that are comparatively low (52%), implying that subjects consider the utility from both the prizes and cumulative earnings about equally in each lottery. From Panel G we have an estimate of $\omega$ that implies that a weight of approximately 100% is placed on cumulative income. The logic of this difference is immediate from the finding that a large percentage of the observations are better characterized by EUT: the analysis that required that 100% of the observations "fit into the EUT model" effectively required that the parameters of the EUT model account for the 34% of the data that was 'actually' generated by the PT model. In that case the implicit weight on cumulative income in the utility function is 0, by definition, so the EUT model needed to account for those observations as well.

Turning to the parameters from the PT data-generating process estimated in Panel G, we find some robust results with respect to the earlier estimates of the PT model when it was assumed to characterize 100% of the observations. First, the risk aversion parameters consistently suggest risk aversion over gains and losses: $\alpha = 0$ and $\beta$ not statistically different than 0. Second, there is no evidence of significant probability weighting, with gamma ($\gamma$) now statistically closer to 1 than it was in Panel E when PT was assumed to characterize the complete sample. Some shift would have been expected since $\gamma$=1 for EUT subjects, and PT was required to explain their behaviour in Table 18 Panels E and F; in Table 20 it is "free" to just characterize the PT subjects. Third, there is no consistent evidence of loss aversion and the estimated parameter would actually indicate loss seeking (lambda = $-.44 < 1$). Even though the estimate of lambda is less than 1 (indicating loss-seeking), this parameter has a standard error of 0.276 and a 95% confidence interval between -0.99 and 0.09 and is not significantly different than zero. Thus we find no evidence of loss aversion in this case, although the imprecision of the estimate of $\lambda$ may be due to a combination of the participant heterogeneity that is assumed away in this specification and the relatively smaller number of observations attributed to PT decision making in a loss frame. The maximum likelihood model was, unfortunately, not able to converge using demographic explanatory variables and no further decomposition of these estimated variables to account for participant heterogeneity was possible.

### Game 2 Findings

### 1 – Eliciting Subjective Bayesian Beliefs

### Analysis

Following Antoniou (Andersen, Harrison et al. 2006) we assume responses to the belief elicitation task can be used to draw estimates about the belief that each subject holds if we are willing to assume something about how they make decisions under risk. To allow for the general case in which we have risk aversion, we jointly estimate the subjective probability and the parameters of the utility function as noted previously.

As described above,

$$U(x) = x^{1-r} / (1-r) \tag{9.15}$$

where "r" is the 'coefficient of risk aversion' and "x" is the monetary value of the outcome (>0), where r=0 corresponds to risk neutrality, r<0 to risk loving, and r>1 to risk aversion. The coefficient of risk aversion is constant for the power function family. 'Expected utility' is then specified as the *probability conditioned* utility,

$$EU_i = \sum_{k=1,K}(p_k \times U(x)_k) \tag{9.16}$$

where p is the probability of the outcome and U(x) is the CRRA specified utility function.

In each round, the participant bet as to whether the sample of daily losses were generated by a 'Low Controls System' or a 'High Controls System'. Table 21 sets out the lotteries over which the participant makes bets with 19 bookies:

Table 21 - Game 2 Betting Table

| | | Odds Offered | | Payout including stake of $3, if you pick right system | |
|---|---|---|---|---|---|
| Bookie # | Stake | Yellow System (Low Controls) | Orange System (High Controls) | Yellow System (Low Controls) | Orange System (High Controls) |
| 1 | $3.00 | 20:1 | 1.05:1 | $60.00 | $3.15 |
| 2 | $3.00 | 10:1 | 1.11:1 | $30.00 | $3.33 |
| 3 | $3.00 | 6.67:1 | 1.18:1 | $20.00 | $3.54 |
| 4 | $3.00 | 5:1 | 1.25:1 | $15.00 | $3.75 |
| 5 | $3.00 | 4:1 | 1.33:1 | $12.00 | $4.00 |
| 6 | $3.00 | 3.33:1 | 1.43:1 | $10.00 | $4.29 |
| 7 | $3.00 | 2.86:1 | 1.54:1 | $8.58 | $4.62 |
| 8 | $3.00 | 2.5:1 | 1.67:1 | $7.50 | $5.00 |
| 9 | $3.00 | 2.22:1 | 1.82:1 | $6.66 | $5.46 |
| 10 | $3.00 | 2:1 | 2:1 | $6.00 | $6.00 |
| 11 | $3.00 | 1.82:1 | 2.22:1 | $5.46 | $6.66 |
| 12 | $3.00 | 1.67:1 | 2.5:1 | $5.00 | $7.50 |
| 13 | $3.00 | 1.54:1 | 2.86:1 | $4.62 | $8.58 |
| 14 | $3.00 | 1.43:1 | 3.33:1 | $4.29 | $10.00 |
| 15 | $3.00 | 1.33:1 | 4:1 | $4.00 | $12.00 |
| 16 | $3.00 | 1.25:1 | 5:1 | $3.75 | $15.00 |
| 17 | $3.00 | 1.18:1 | 6.67:1 | $3.54 | $20.00 |
| 18 | $3.00 | 1.11:1 | 10:1 | $3.33 | $30.00 |
| 19 | $3.00 | 1.05:1 | 20:1 | $3.15 | $60.00 |

Using the betting table above, the subject that believes the loss data was generated by the Low Controls system (event “A” or the ‘left hand’ binary choice) and selecting “Low Controls System” from a given bookie ‘b’ receives EU

$$EU_{\text{LOW CONTROLS SYSTEM}} = \pi_{\text{LOW CONTROLS SYSTEM}} \times U(\text{payout if the system generating the observations was a Low Controls system} \mid \text{bet that the system was a Low Controls system}) + (1- \pi_{\text{LOW CONTROLS SYSTEM}}) \times U(\text{payout if the system generating the observations was a High Controls system} \mid \text{bet that the system was a Low Controls system}) \tag{9.17}$$

where $\pi_{\text{LOW CONTROLS SYSTEM}}$ is the subjective probability that the system generating the losses was a Low Controls System. The payouts that enter the utility function are defined by the odds that each bookie offers, and are shown in the betting table. For the bet offered by the first bookie, for example, these payouts are \$60 and \$0, so we have

$$EU_{\text{LOW CONTROLS SYSTEM}} = \pi_{\text{LOW CONTROLS SYSTEM}} \times U(\$60) + (1- \pi_{\text{LOW CONTROLS SYSTEM}}) \times U(\$0) \tag{9.18}$$

Where the participant potentially earns \$60 is correct but nothing if incorrect.

We similarly define the EU received from a bet on a High Controls System (event “B” or the ‘right hand’ binary choice) as the complement of event A:

$$EU_{\text{HIGH CONTROLS SYSTEM}} = \pi_{\text{LOW CONTROLS SYSTEM}} \times U(\text{payout if the system generating the observations was a Low Controls system} \mid \text{bet that the system was a High Controls system}) + (1- \pi_{\text{LOW CONTROLS SYSTEM}}) \times U(\text{payout if the system generating the observations was a High Controls system} \mid \text{bet that the system was a High Controls system}) \tag{9.19}$$

and this translates for the first bookie in Table 21 into payouts of \$0 and \$3.15 (i.e. the participant potentially earns \$0 if incorrect as usual, but \$3.15 if correct, so we have

$$EU_{\text{HIGH CONTROLS SYSTEM}} = \pi_{\text{LOW CONTROLS SYSTEM}} \times U(\$0) + (1- \pi_{\text{LOW CONTROLS SYSTEM}}) \times U(\$3.15) \tag{9.20}$$

for this particular bookie and bet. We observe the bet made by the subject for a range of odds, so we can calculate the likelihood of that choice given values of r (risk aversion), $\pi_{\text{LOW CONTROLS SYSTEM}}$ and $\mu$.

The rest of the structural specification is exactly the same as for the choices over lotteries with objective probabilities. Given the specifications for $EU_{LOW\ CONTROLS\ SYSTEM}$ and $EU_{HIGH\ CONTROLS\ SYSTEM}$ above, we can define the latent index as

$$\nabla EU = (EU_{LOW} - EU_{HIGH}) \tag{9.21}$$

for each of the 'Low Controls System' and 'High Controls System' bets from Table 21. We then define the probability of choosing a Low Controls System as

$$\text{prob(choose B)} = \Phi\ [\ (\nabla EU) / \nu\ ) / \mu\ ] \tag{9.22}$$

where $\nu$ is a normalizing term for the value of each bookie choice pair and $\mu > 0$ is a structural "noise parameter" for the belief choices that is used to allow some errors from the perspective of the deterministic SEU model[69].

Writing out the complete likelihood function, we have

$$\ln L(r, \pi_{LOW\ CONTROLS\ SYSTEM}, \mu; y, X) = \Sigma_i\ [\ (\ln \Phi(\nabla EU) \times I(y_i = 1)) + (\ln (1-\Phi(\nabla EU)) \times I(y_i = -1))\ ] \tag{9.23}$$

for the observed choices in the experiment defined over subjective probabilities.

It is useful to see how the estimation procedure maps back to the economics of the SEU model. Ignoring the behavioural error term $\mu$, we need r to evaluate the utility function, we need $\pi_{LOW\ CONTROLS\ SYSTEM}$ to calculate the EU once we know the utility values, and we need both of them to calculate the latent index that generates the probability of observing the choice of bet "Low Controls System" or bet "High Controls System". The joint maximum likelihood problem is to find the values of these parameters that best explain observed choices in the game.

This formal analysis assumes that we are estimating one subjective probability $\pi_{LOW\ CONTROLS\ SYSTEM}$. There are two simple extensions that allow us to consider a complete design, which involves the 6 posterior probabilities shown in Table 22:

[69] Antoniou: "The normalizing term $\nu$ is defined as the maximum utility over all prizes in this lottery pair minus the minimum utility over all prizes in this lottery pair, and ensures that the normalized EU difference $[(EUR - EUL)/\nu]$ remains in the unit interval. As $\mu \rightarrow \infty$ this specification collapses $\nabla EU$ to 0 for any values of EUR and EUL, so the probability of either choice converges to ½. So a larger $\mu$ means that the difference in the EU of the two lotteries, conditional on the estimate of r, has less predictive effect on choices. Thus $\mu$ can be viewed as a parameter that flattens out, or "sharpens," the link functions implicit in (4). This is just one of several different types of error story that could be used, and Wilcox (2008) provides a masterful review of the implications of the strengths and weaknesses of the major alternatives." (Antoniou, Harrison et al. 2010)

**Table 22 - Game 2 Observations (n=25,327)**

**Objective Posterior Probabilities of a Low System generating a given the number of Loss Observations vs. the number of observed Loss Observations**

| Posterior Probability of the Losses being generated by a "Low Controls System" | Total Number of Loss Observations Drawn from 365 days of simulated losses | Actual Number of Loss Observations > $17,000 | Actual Number of Loss Observations <= $17,000 | Number of Rounds played with this combination (n = 25,327 rounds total) |
|---|---|---|---|---|
| **.12** | 5 | 0 | 5 | 760 |
| | 9 | 2 | 7 | 399 |
| | 17 | 6 | 11 | 722 |
| **.23** | 3 | 0 | 3 | 589 |
| | 5 | 1 | 4 | 2185 |
| | 9 | 3 | 6 | 817 |
| | 17 | 7 | 10 | 475 |
| **0.4** | 3 | 1 | 2 | 1349 |
| | 5 | 2 | 3 | 3743 |
| | 9 | 4 | 5 | 1292 |
| | 17 | 8 | 9 | 722 |
| **0.6** | 3 | 2 | 1 | 1425 |
| | 5 | 3 | 2 | 3629 |
| | 9 | 5 | 4 | 969 |
| | 17 | 9 | 8 | 570 |
| **0.77** | 3 | 3 | 0 | 551 |
| | 5 | 4 | 1 | 2033 |
| | 9 | 6 | 3 | 798 |
| | 17 | 10 | 7 | 722 |
| **0.88** | 5 | 5 | 0 | 399 |
| | 9 | 7 | 2 | 532 |
| | 17 | 11 | 6 | 646 |

The first is to assume symmetry, in the sense that the estimate of $\pi_{\text{LOW CONTROLS SYSTEM}}$ is treated as (1-$\pi_{\text{LOW CONTROLS SYSTEM}}$) when evaluating the choices made for the corresponding task. In other words, if we only use the choices for the betting task that has posterior probability 0.60, we could directly apply the existing formal analysis. But we can also include the choices for the betting task that has posterior probability 0.40, and assume that the subjective probability for those choices is one minus the subjective probability for the choices with posterior probability 0.60. Therefore we only need to estimate one subjective probability. This seems to be an innocuous assumption, and is directly testable. The second extension is then to introduce two extra subjective probability parameters, so we have one for the betting task with posterior probability 0.6 (and 0.4 by symmetry), one for the betting task with posterior probability 0.77 (and 0.23), and one for the betting task with posterior probability 0.88 (and 0.12). Hence we have to estimate one risk attitude parameter and three subjective probabilities, along with the one behavioural error term.

**Results**

**1 – Eliciting Subjective Bayesian Beliefs**

Panel A1 reports the maximum likelihood estimates of the SEU probabilities with no demographic controls:

**Table 23 - Game 2 Panel A1: Maximum Likelihood Estimates Assuming SEUT Subjective Expected Utility Model with Fechner error term**

```
                                                Number of obs   =      25327
                                                Wald chi2(0)    =          .
Log pseudolikelihood = -14642.796               Prob > chi2     =          .

                                    (Std. Err. adjusted for 53 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
       _cons |   .0178005   .0319052     0.56   0.577    -.0447326    .0803336
-------------+----------------------------------------------------------------
sprob1_      |
       _cons |  -.4237312   .1079295    -3.93   0.000    -.6352691   -.2121934
-------------+----------------------------------------------------------------
sprob2_      |
       _cons |  -.9073699   .2043991    -4.44   0.000    -1.307985   -.5067551
-------------+----------------------------------------------------------------
sprob3_      |
       _cons |  -.8613967   .2559962    -3.36   0.001     -1.36314   -.3596533
-------------+----------------------------------------------------------------
LNmuB        |
       _cons |  -.6046584   .2033545    -2.97   0.003    -1.003226    -.206091
------------------------------------------------------------------------------

.     nlcom (sprob1: 1/(1+exp([sprob1_]_b[_cons]))) (sprob2:
1/(1+exp([sprob2_]_b[_cons]))) (sprob3: 1/(1+exp([sprob3_]_b[_cons])))

      sprob1:  1/(1+exp([sprob1_]_b[_cons]))
      sprob2:  1/(1+exp([sprob2_]_b[_cons]))
      sprob3:  1/(1+exp([sprob3_]_b[_cons]))

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      sprob1 |   .6043757   .0258065    23.42   0.000     .5537958    .6549557
      sprob2 |   .7124617   .0418732    17.01   0.000     .6303917    .7945316
      sprob3 |   .7029524   .0534547    13.15   0.000     .5981832    .8077216
------------------------------------------------------------------------------
```

In Panel A1, we find evidence of risk *seeking* with r=0.017, >0 but statistically insignificant from zero[70]. The estimated subjective posterior probabilities are sprob1 = 0.604 (vs. 0.6); sprob2 = .712 (vs. 0.77) and sprob 3 = 0.703 (vs. .88), with only sprob3 being statistically different than the true Bayesian posterior

[70] To check this result we also ran a model with r constrained to zero – the results are not reported here but were essentially identical to Panel A.

probability at a 95% confidence level. There is general *underestimation* of the Bayesian posterior probabilities that increases with the objective probability but with significant dispersion of the estimates across individuals and with precision declining as the posterior gets larger. Correspondingly, these results also imply a systematic *overestimation* of the true Bayesian posterior probability when the posterior is *less* than 0.5, becoming larger as the posterior decreases from 0.4 to 0.23 and 0.12.

If we rerun this model under the assumption of risk neutrality (r=.001), the results do not change substantially, indicating that there is no substantial difference in subjective probability estimates under an assumption of Risk neutrality:

**Figure 129 - Game 2 Panel A2: Maximum Likelihood Estimates Assuming SEUT and Risk Neutrality**

```
                                                   Number of obs   =      25327
                                                   Wald chi2(0)    =          .
Log pseudolikelihood = -14643.255                  Prob > chi2     =          .

 ( 1)  [r]_cons = .001
                                     (Std. Err. adjusted for 53 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
       _cons |       .001          .        .       .            .           .
-------------+----------------------------------------------------------------
sprob1_      |
       _cons |   -.423196   .1073301    -3.94   0.000    -.6335592   -.2128327
-------------+----------------------------------------------------------------
sprob2_      |
       _cons |  -.9061668   .2032307    -4.46   0.000    -1.304492   -.5078418
-------------+----------------------------------------------------------------
sprob3_      |
       _cons |  -.8608141   .2556783    -3.37   0.001    -1.361934   -.3596939
-------------+----------------------------------------------------------------
LNmuB        |
       _cons |  -.4961234   .1060705    -4.68   0.000    -.7040177    -.288229
------------------------------------------------------------------------------

.     nlcom (sprob1: 1/(1+exp([sprob1_]_b[_cons]))) (sprob2:
1/(1+exp([sprob2_]_b[_cons]))) (sprob3: 1/(1+exp([sprob3_]_b[_cons])))

      sprob1:  1/(1+exp([sprob1_]_b[_cons]))
      sprob2:  1/(1+exp([sprob2_]_b[_cons]))
      sprob3:  1/(1+exp([sprob3_]_b[_cons]))

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      sprob1 |   .6042478   .0256661    23.54   0.000     .5539431    .6545524
      sprob2 |   .7122151   .0416551    17.10   0.000     .6305726    .7938577
      sprob3 |   .7028307   .0534009    13.16   0.000     .5981669    .8074945
------------------------------------------------------------------------------
```

This result is contrary to Antonio's findings where the qualitative effect of risk aversion should follow from theory: the more risk averse the subject, the more likely he is to bet as if his subjective probability is 0.5, since this reduces the dispersion in final outcomes from all bets[71]. In our case, we may fail to see any effect since, as noted above, the participants appear to be risk seeking to being with. In any case, the log-likelihood of the risk neutral model (-14643) is nearly identical to the log-likelihood for the general model (--14642), and we do not appear to obtain coefficients of sprob shifted to the left allowing for risk aversion.

Panels B1 and B2 extend the specifications in Panels A1 and A2 to allow for demographic variation in r and each of the subjective probability estimates. In the model allowing for estimation of r, after controlling for demographic effects, participants appear to be slightly risk *seeking* (r = -.129 <0). While none of the demographic parameter estimates are significant at the 95% level, *age*, *sex* (female), a *business degree*, *IT certifications*, and increased *income* all increase risk aversion, while a *math degree* and increased number of career *work years* lower risk aversion on average, results which are similar to Game 1:

**Table 24 - Game 2 Panel B1: Maximum Likelihood Estimates Assuming SEUT Subjective Expected Utility Model with Fechner error term and demographics**

```
                                                  Number of obs   =      25327
                                                  Wald chi2(7)    =      39.31
Log pseudolikelihood = -14131.072                 Prob > chi2     =     0.0000

                                    (Std. Err. adjusted for 53 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
         age |   .0416305   .0351939     1.18   0.237    -.0273483    .1106092
         sex |   .0371306   .0549041     0.68   0.499    -.0704794    .1447406
         bus |   .0390964   .0918374     0.43   0.670    -.1409017    .2190945
        math |  -.0095796   .0281189    -0.34   0.733    -.0646916    .0455324
   workyears |  -.0100221   .0312938    -0.32   0.749    -.0713567    .0513126
       certs |   .0127775   .0196882     0.65   0.516    -.0258106    .0513656
      income |   .0627921   .0672722     0.93   0.351    -.0690589    .1946431
       _cons |  -.1287453   .0442458    -2.91   0.004    -.2154655   -.0420251
-------------+----------------------------------------------------------------
sprob1_      |
         age |   .1046971   .2381458     0.44   0.660    -.3620601    .5714542
         sex |   .2633251   .2120888     1.24   0.214    -.1523613    .6790115
         bus |  -.1191181   .3954357    -0.30   0.763    -.8941578    .6559217
        math |   .0944121   .1644827     0.57   0.566    -.2279682    .4167924
   workyears |  -.0843451   .1777815    -0.47   0.635    -.4327904    .2641001
       certs |   .0634325   .0702458     0.90   0.367    -.0742468    .2011118
      income |  -.2849189   .3547057    -0.80   0.422    -.9801293    .4102914
       _cons |  -.3018998   .1448808    -2.08   0.037     -.585861   -.0179387
-------------+----------------------------------------------------------------
sprob2_      |
         age |   .0261102   .6174798     0.04   0.966    -1.184128    1.236348
         sex |   .8075164   .3517658     2.30   0.022     .1180681    1.496965
         bus |  -.2125568   1.693283    -0.13   0.900     -3.53133    3.106216
        math |  -.0298328   .4302981    -0.07   0.945    -.8732015    .8135359
   workyears |   .0529048   .2698478     0.20   0.845    -.4759871    .5817967
```

[71] This qualitative result is exactly the same as one finds using a Quadratic scoring rule: see results from my Game# 4 below and generally from Andersen et al (Andersen, Fountain et al. 2014).

```
       certs |  -.0436988    .394462    -0.11   0.912    -.8168302    .7294325
      income |   -.412683   .5061199    -0.82   0.415     -1.40466    .5792937
       _cons |  -1.008586   .3829995    -2.63   0.008    -1.759252    -.257921
-------------+----------------------------------------------------------------
sprob3_      |
         age |  -1.044596   1.714596    -0.61   0.542    -4.405143     2.31595
         sex |   .6343824   1.344349     0.47   0.637    -2.000492    3.269257
         bus |   .8931119   1.753905     0.51   0.611    -2.544479    4.330703
        math |   .2364082   .7039048     0.34   0.737     -1.14322    1.616036
   workyears |   .4890464   .8595074     0.57   0.569    -1.195557     2.17365
       certs |  -.1264412   .3856509    -0.33   0.743     -.882303    .6294206
      income |   .0972284   1.241316     0.08   0.938    -2.335706    2.530162
       _cons |  -1.162614   .9313578    -1.25   0.212    -2.988042    .6628134
-------------+----------------------------------------------------------------
LNmuB        |
       _cons |  -.7000786   .3679674    -1.90   0.057    -1.421281    .0211244
------------------------------------------------------------------------------


      sprob1:  1/(1+exp([sprob1_]_b[_cons]))
      sprob2:  1/(1+exp([sprob2_]_b[_cons]))
      sprob3:  1/(1+exp([sprob3_]_b[_cons]))

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      sprob1 |   .5749069   .0354073    16.24   0.000     .5055099    .6443039
      sprob2 |   .7327434    .075003     9.77   0.000     .5857402    .8797465
      sprob3 |   .7618074   .1690013     4.51   0.000      .430571    1.093044
------------------------------------------------------------------------------
```

**Table 25 - Game 2 Panel B2: Maximum Likelihood Estimates Assuming SEUT and Risk Neutrality**

**Subjective Expected Utility Model with Fechner error term and demographics**

```
                                                  Number of obs   =      25327
                                                  Wald chi2(7)    =      29.57
Log pseudolikelihood = -14171.009                 Prob > chi2     =     0.0001

 ( 1)  [r]_cons = .001
                                 (Std. Err. adjusted for 53 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
         age |   .0495754   .0500063     0.99   0.321    -.0484352     .147586
         sex |   .0304329   .0830262     0.37   0.714    -.1322954    .1931612
         bus |   .0516867   .0953419     0.54   0.588    -.1351801    .2385535
        math |  -.0210517   .0449941    -0.47   0.640    -.1092386    .0671351
   workyears |   -.014644   .0469215    -0.31   0.755    -.1066084    .0773204
       certs |   .0117682   .0194961     0.60   0.546    -.0264434    .0499799
      income |   .0760817   .0872093     0.87   0.383    -.0948453    .2470088
       _cons |       .001          .        .       .            .           .
-------------+----------------------------------------------------------------
sprob1_      |
         age |   .1093948   .2332036     0.47   0.639    -.3476759    .5664655
         sex |   .2627839   .2291259     1.15   0.251    -.1862947    .7118624
         bus |  -.1115489   .2813028    -0.40   0.692    -.6628923    .4397945
        math |   .0813937   .1754596     0.46   0.643    -.2625009    .4252883
   workyears |  -.0877681   .1700601    -0.52   0.606    -.4210799    .2455436
       certs |   .0731711   .0591794     1.24   0.216    -.0428183    .1891605
      income |  -.2513605   .3226661    -0.78   0.436    -.8837744    .3810533
       _cons |  -.3207558   .1500612    -2.14   0.033    -.6148704   -.0266412
-------------+----------------------------------------------------------------
sprob2_      |
         age |   .0704023   .4862412     0.14   0.885    -.8826129    1.023418
         sex |   .7914908   .4161795     1.90   0.057     -.024206    1.607188
         bus |  -.2627463   .9519252    -0.28   0.783    -2.128485    1.602993
        math |  -.0725428   .4167179    -0.17   0.862    -.8892949    .7442093
   workyears |   .0292257   .2518767     0.12   0.908    -.4644437     .522895
       certs |  -.0185482     .18492    -0.10   0.920    -.3809848    .3438884
      income |  -.3419466   .4500247    -0.76   0.447    -1.223979    .5400856
       _cons |  -1.022145   .4343749    -2.35   0.019    -1.873504   -.1707864
-------------+----------------------------------------------------------------
sprob3_      |
         age |  -.7826717   1.678502    -0.47   0.641    -4.072476    2.507132
         sex |   .4688754   1.095068     0.43   0.669    -1.677418    2.615169
         bus |   .8024294   1.295483     0.62   0.536     -1.73667    3.341529
        math |   .1623243   .5395229     0.30   0.764    -.8951212     1.21977
   workyears |   .3452675   .8535333     0.40   0.686    -1.327627    2.018162
       certs |  -.1110328   .1971383    -0.56   0.573    -.4974168    .2753512
      income |   .0857451   .8684467     0.10   0.921    -1.616379    1.787869
       _cons |  -.9554302   .6609176    -1.45   0.148    -2.250805    .3399445
-------------+----------------------------------------------------------------
LNmuB        |
       _cons |  -1.472139   .3927333    -3.75   0.000    -2.241882   -.7023958
------------------------------------------------------------------------------

      sprob1:  1/(1+exp([sprob1_]_b[_cons]))
      sprob2:  1/(1+exp([sprob2_]_b[_cons]))
      sprob3:  1/(1+exp([sprob3_]_b[_cons]))

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      sprob1 |   .5795084   .0365667    15.85   0.000     .5078391    .6511778
      sprob2 |   .7353903   .0845256     8.70   0.000     .5697231    .9010575
      sprob3 |   .7222059   .1325963     5.45   0.000      .462322    .9820899
------------------------------------------------------------------------------
```

Interestingly, after controlling for demographics, the *predicted* values of the posterior probabilities after accounting for the demographics of the participants generally improves for sProb 2 (.77) and sProb 3 (.88) in the model allowing for risk aversion, but do not improve in the model assuming risk neutrality:

**Table 26 - Game 2 Panel B3: Predicted Probabilities**

**Assuming SEUT (Sprob) and SEUT w. Risk Neutrality (RNsprob)**

```
-----------------------------------------------------------------------------
Variable |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
-----------------------------------------------------------------------------
Sprob  1 |   25327    .6203263    .0004004    .0637177    .6195415     .621111
-----------------------------------------------------------------------------
Sprob  2 |   25327    .7440988    .0006926    .1102165    .7427413    .7454562
-----------------------------------------------------------------------------
Sprob  3 |   25327    .7526894    .0008068    .1283984     .751108    .7542708

-----------------------------------------------------------------------------
Variable |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
```

```
---------+-----------------------------------------------------------------
RNsprob1 |    25327     .6122906     .0003858     .0613988     .6115344     .6130468
---------+-----------------------------------------------------------------
RNsprob2 |    25327     .7285915     .0006574     .1046183      .727303       .72988
---------+-----------------------------------------------------------------
RNsprob3 |    25327     .7295159     .0007222     .1149335     .7281004     .7309315
```

Figures 130 and 131 illustrate the kernel density functions for the estimated subjective probabilities across all participants for the SEU and SEU with risk neutrality case, indicating significant diversity across participants:

**Figure 130 - Kernel density for Game 2 Subjective Probability Estimates under SEUT, with demographics**

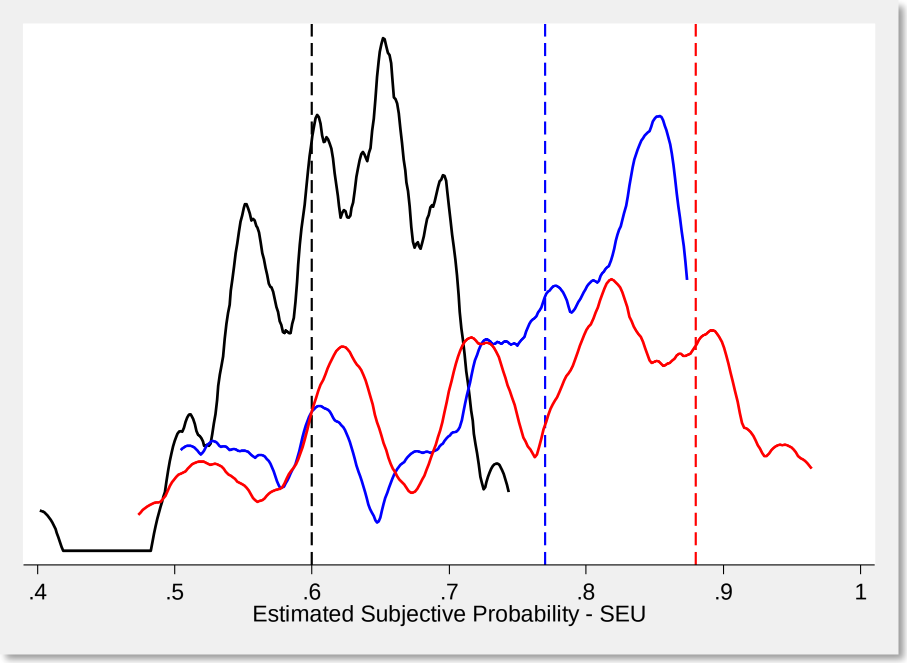

**Figure 131 - Kernel density for Game 2 Subjective Probability Estimates under SEUT assuming Risk Neutrality, with demographics,**

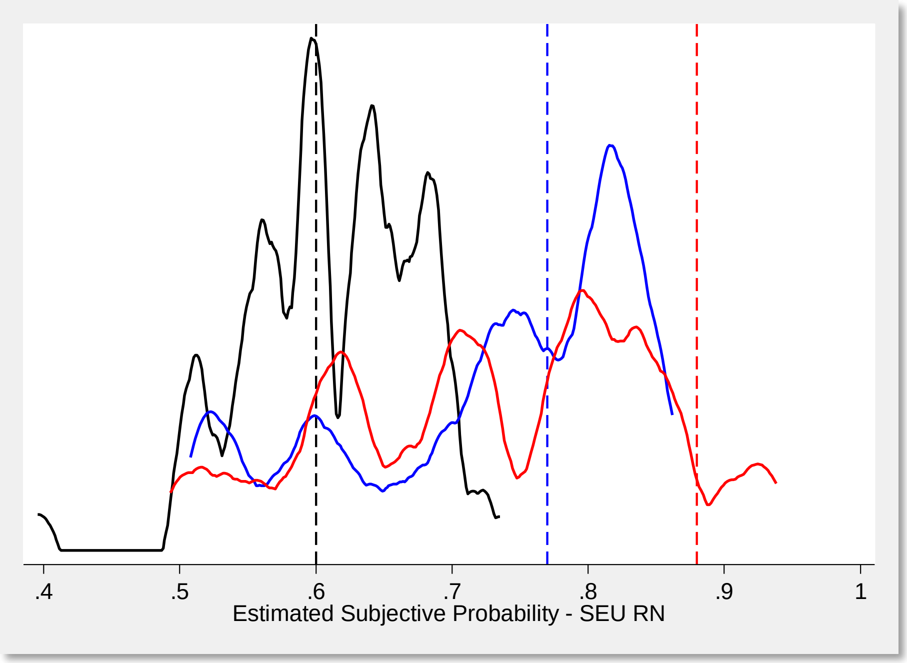

Following Antoniou, Figure 132 displays the estimated subjective probabilities, pooling the symmetric cases, with reference lines showing the corresponding posterior probability using Bayes Rule. The dispersion in these distributions reflects the standard errors in the parameter estimates of the subjective probabilities, as well as variations across subjects due to differences in demographic characteristics:

**Figure 132 – Estimates of Game 2 Subjective Probabilities under SEUT – pooled observations**

Black is 0.60 posterior, Blue is 0.77 posterior and Red is 0.88 posterior

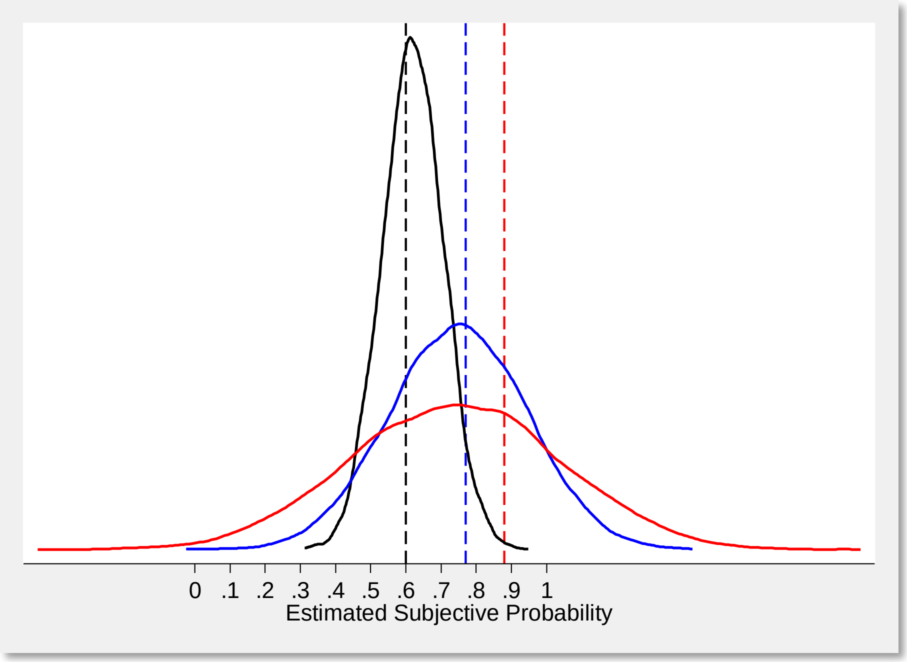

Results here are similar to Antoniou: underestimation of the true Bayesian posterior probability which becomes larger as the posterior increases from 0.6 to 0.77 and 0.88. In addition, the precision of the subjective probability estimate also declines as the posterior gets larger.

Figure 133 shows the corresponding estimates of subjective probability when we impose risk neutrality (r=.001). Finally, Figure 134 superimposes the two sets of results for comparison. In this case, contra Antoniou, Bayes Rule does *not* do better when we allow the data to determine the risk attitudes of subjects. In the risk neutral case the subjective estimate of the 0.6 posterior probability is similar, but the estimates for the 0.77 and 0.88 posterior probabilities, while not closer to the posterior on average, exhibit less dispersion around the Bayesian posterior when we assume risk neutrality:

**Figure 133 - Estimates of Game 2 Subjective Probabilities under SEUT assuming Risk Neutrality – pooled observations**

Black is 0.60 posterior, Blue is 0.77 posterior and Red is 0.88 posterior

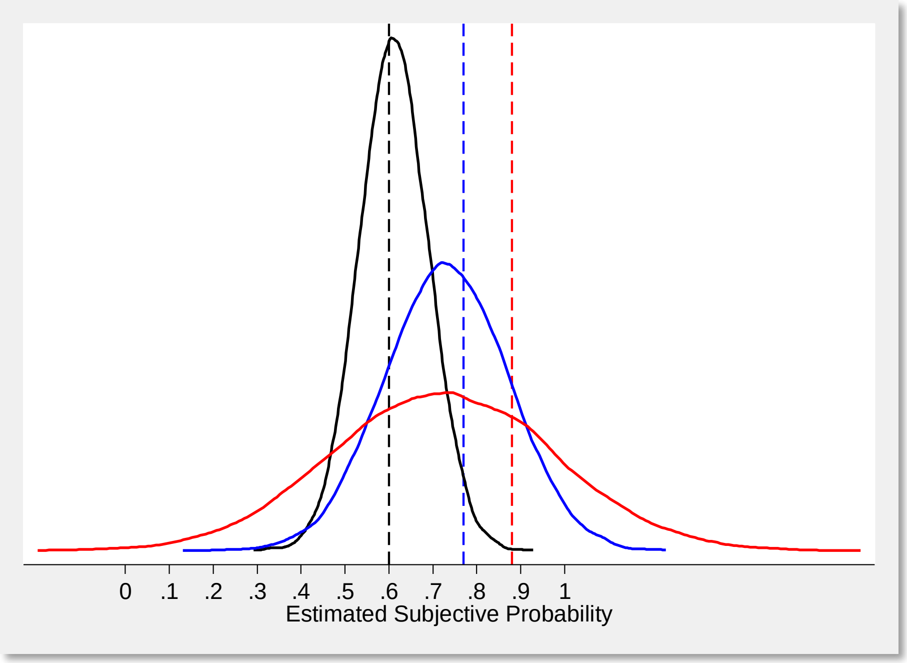

**Figure 134 - Estimates of Game 2 Subjective Probabilities Under SEU: The Effect of Risk Aversion – pooled observations**

Solid line is general SEU model, and dashed line is SEU assuming risk neutrality
Black is 0.60 posterior, blue is 0.77 posterior and red is 0.88 posterior

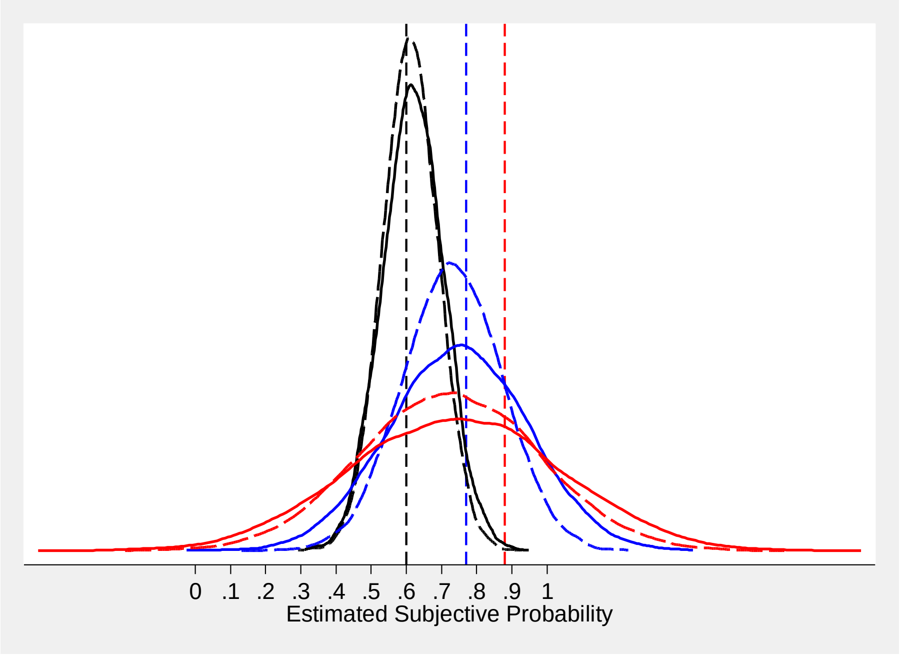

### 2 – Estimating the Effect of Strength and Weight of Evidence on Beliefs

**Analysis**

Following Antoniou (Andersen, Harrison et al. 2006) we now estimate the effect of the *strength* vs. the *weight* of evidence on subjective probability estimates.

In this model we use an example concerning two hypotheses: A (that the observed losses were generated by a Low Controls system); and B (that the observed losses were generated by a High Controls system) where P(A) + P(B) = 1. The task therefore is to determine the likelihood of A in comparison to the likelihood of B. Ceteris paribus our expectation of A is P(A) and of B, P(B). These probabilities in a Bayesian framework are called the priors.

Suppose that we now observe an information signal C, which is the number of losses above and below $10,000 (i.e., X observations above $10,000 and n-X observations below $10,000). Suppose further that we know the prior probability of the Low System generating losses above $10,000 is 60% and the corresponding probability for the High Controls system is 40%. Bayes rule is the method with which the priors *should be* combined with the new data, to get a posterior belief in terms of the hypotheses A and B. In other words the rule reveals how the occurrence of C should influence our expectation of A and B.

Formally the rule states that:

$$P(A|C) = \frac{P(A)*P(C|A)}{P(C)} \tag{9.24}$$

If we divide the above with the corresponding probability of the likelihood of P(B/C) we get:

$$\frac{P(A|C)}{P(B|C)} = \frac{P(C|A)}{P(C|B)} \tag{9.24}$$

This equation shows the *likelihood ratio*. Each hypothesis is distributed binomially with parameters *n* (the number of observations), and *p* (the probability of observing a loss above or below $10,000). In terms of our experiment, C is the sample of losses observed from the chosen system. The two competing hypotheses are that the observations came from the Low Controls System or the High Controls System.

$$P(C|A) = \binom{n}{w} 0.6^{x} * \ 0.4^{n-x} \tag{9.25}$$

This equation yields the probability of x losses (i.e. losses greater than $10,000) and n-x (y) losses (i.e. losses less than or equal to $10,000) from the Low Controls System. P(B/C) is the corresponding probability of the particular pattern coming from the High Controls System. Dividing P(C/A) by P(C/B) gives:

$$\frac{P(A|C)}{P(B|C)} = \frac{P(C|A)}{P(C|B)} = \frac{\binom{n}{w}0.6^{x} * 0.4^{n-x}}{\binom{n}{w}0.6^{n-x} * 0.4^{x}} = \frac{0.6^{-n+2x}}{0.4^{-n+2x}} \tag{9.26}$$

Taking logs of both sides:

$$\text{Log}\ \frac{\text{P(A|C)}}{\text{P(B|C)}} = (-\text{n} + 2\text{x})\log\left(\frac{0.6}{0.4}\right) \tag{9.27}$$

Because n = x + y, substituting for –n yields:

$$\text{Log}\ \frac{\text{P(A|C)}}{\text{P(B|C)}} = (\text{x} - \text{y})\log\left(\frac{0.6}{0.4}\right) \tag{9.28}$$

Now multiply and divide by n yields:

$$\text{Log}\ \frac{\text{P(A|C)}}{\text{P(B|C)}} = \text{n}\,\frac{\text{w}-\text{x}}{\text{n}}\log\left(\frac{0.6}{0.4}\right) \tag{9.29}$$

The right hand side is the probabilities elicited from subjects' responses modified to log-odds form. If they are Bayesian they should be explained by the right hand side, which was derived from Baye's Theorem. As noted above, the two dimensions that Griffin and Tversky suggest are sample size (weight), n, and strength (x – y)/n, i.e., the number of observed losses greater than $10,000 minus the number of observed losses less than or equal to $10,000, all divided by the total number observations. In order to decompose we take logs of both sides, which yields:

$$\text{Log}\left[\frac{\text{Log}\,\frac{\text{P(A|C)}}{\text{P(B|C)}}}{\text{Log}\left(\frac{0.6}{0.4}\right)}\right] = \text{a}\log(\text{n}) + \ \beta\log\left(\frac{\text{x}-\text{y}}{\text{n}}\right) \tag{9.30}$$

The coefficients a and β that are derived from the above regression are in effect the importance that the individuals have given to strength and weight of the evidence presented to them in the experiment in each round. Under the null hypothesis the coefficients should be equal to 1, because Bayes Rule predicts that they should affect judgment equally. Under the GT hypothesis α > β. We test this hypothesis directly using the choices made in Game 2.

**Results**

Panel C1 reports the maximum likelihood estimates of alpha and beta, the coefficient of risk aversion and the behavioural error term. We also test whether they are statistically different from each other:

**Table 27 - Game 2 Panel C1: Maximum Likelihood Estimates Assuming SEUT**
**Coefficients of strength (alpha) and weight (beta) of evidence and risk aversion**

```
                                                  Number of obs   =      25327
                                                  Wald chi2(0)    =          .
Log pseudolikelihood = -14659.344                 Prob > chi2     =          .

                                    (Std. Err. adjusted for 53 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
       _cons |   .0154356   .0318822     0.48   0.628    -.0470523    .0779235
-------------+----------------------------------------------------------------
alpha        |
       _cons |   .5834962   .1444606     4.04   0.000     .3003586    .8666338
-------------+----------------------------------------------------------------
beta         |
       _cons |   .5793732   .1578205     3.67   0.000     .2700507    .8886956
-------------+----------------------------------------------------------------
LNmuB        |
       _cons |  -.5988478   .2082122    -2.88   0.004    -1.006936   -.1907593
------------------------------------------------------------------------------

beta_alpha:  [beta]_b[_cons] - [alpha]_b[_cons]

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
  beta_alpha |  -.0041231   .1093722    -0.04   0.970    -.2184886    .2102425
------------------------------------------------------------------------------

testnl ([alpha]_b[_cons]-[beta]_b[_cons]=0)

  (1)  [alpha]_b[_cons]-[beta]_b[_cons] = 0

chi2(1) =        0.00
Prob > chi2 =    0.9699
```

The estimated coefficient of risk aversion r is not significantly different than zero indicating risk neutrality. The coefficients of both strength (alpha) and weight (beta) are, however, significantly different from 0 indicating that both strength and weight of evidence affect subjective probability estimates. The strength and weight parameters are not, however, significantly different from each other as predicted by the GT hypothesis, i.e., that α is generally > β. If we allow demographic diversity in the estimates, Panel C2 reports the difference in strength and weight effects between participants and Panel C3 the resulting overall predicted alpha and beta values:

### Table 28 - Game 2 Panel C2: Maximum Likelihood Estimates under Strength/Weight Heuristic (GT Model), w. demographics

### Coefficients of strength (alpha) and weight (beta) of evidence and risk aversion

```
                                                  Number of obs   =      25327
                                                  Wald chi2(5)    =      12.75
Log pseudolikelihood = -14316.316                 Prob > chi2     =     0.0258

                                    (Std. Err. adjusted for 53 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
         age |   .0478322    .018737     2.55   0.011     .0111084     .084556
         sex |   .0425012   .0513579     0.83   0.408    -.0581584    .1431608
         bus |   .0478112   .0537972     0.89   0.374    -.0576294    .1532517
        math |  -.0029667   .0468035    -0.06   0.949    -.0946999    .0887665
       certs |    .015561   .0126961     1.23   0.220    -.0093229     .040445
       _cons |  -.1077724   .0556866    -1.94   0.053    -.2169161    .0013714
-------------+----------------------------------------------------------------
alpha        |
         age |  -.1762316   .4895896    -0.36   0.719     -1.13581    .7833464
         sex |  -.8248275   1.975588    -0.42   0.676    -4.696909    3.047254
         bus |  -.3880153   .7100424    -0.55   0.585    -1.779673    1.003642
        math |  -.3748799   1.130587    -0.33   0.740    -2.590789    1.841029
       certs |   .0775537   .1297196     0.60   0.550    -.1766919    .3317994
       _cons |   1.209929   1.743057     0.69   0.488    -2.206401    4.626259
-------------+----------------------------------------------------------------
beta         |
         age |  -.2613067   .5018223    -0.52   0.603     -1.24486     .722247
         sex |  -.4092002    1.91641    -0.21   0.831    -4.165295    3.346894
         bus |  -.6702154   .6238307    -1.07   0.283    -1.892901    .5524703
        math |  -.2566558   .9903567    -0.26   0.796    -2.197719    1.684408
       certs |   .1366344   .1268109     1.08   0.281    -.1119105    .3851792
       _cons |   1.194852    1.74027     0.69   0.492    -2.216014    4.605717
-------------+----------------------------------------------------------------
LNmuB        |
       _cons |  -.7497431   .5320434    -1.41   0.159    -1.792529    .2930428
------------------------------------------------------------------------------
```

### Table 29 – Game 2 Panel C3 Predicted strength (alpha) and weight (beta) of evidence

```
------------------------------------------------------------------------------
Variable |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
---------+--------------------------------------------------------------------
 GTalpha |   25327    .4271979    .0035553    .5658103    .4202292    .4341665
---------+--------------------------------------------------------------------
  GTbeta |   25327      .45317    .0035591    .5664106     .446194     .460146
------------------------------------------------------------------------------
```

Figure 135 reports the kernel densities of the estimated alpha and beta values across participants, allowing for demographic heterogeneity. Once again we see a wide range of values based on the diversity of the participants over the demographic factors indicated above:

**Figure 135 – Kernel Density for Game 2 strength (alpha) and weight (beta) of evidence, with demographics**

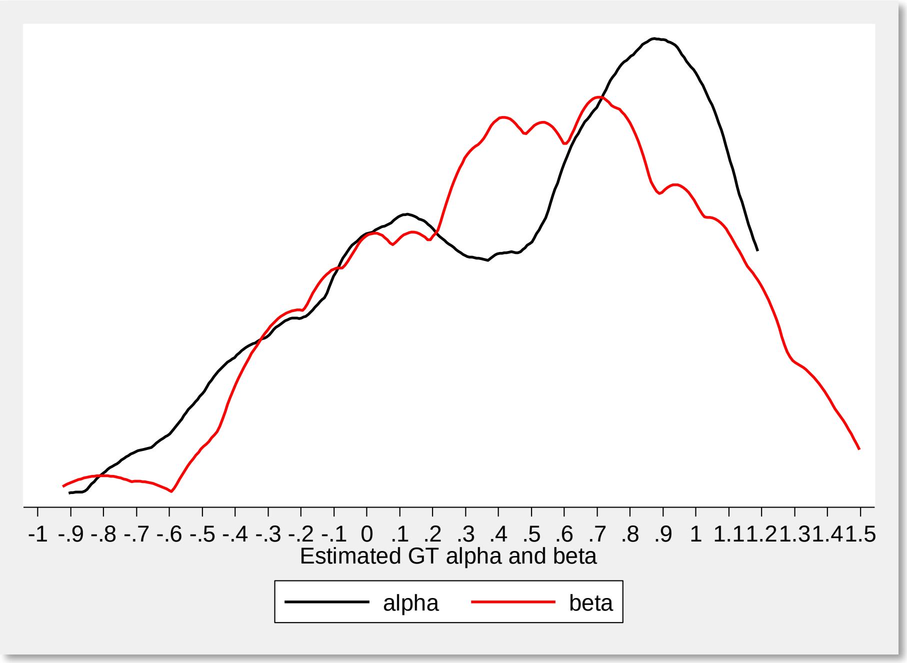

Figure 136 reports the associated estimated subjective probabilities under the strength/weight heuristic:

**Figure 136 – Estimated Subjective Probabilities using Strength/Weight Heuristic, w. demographics**

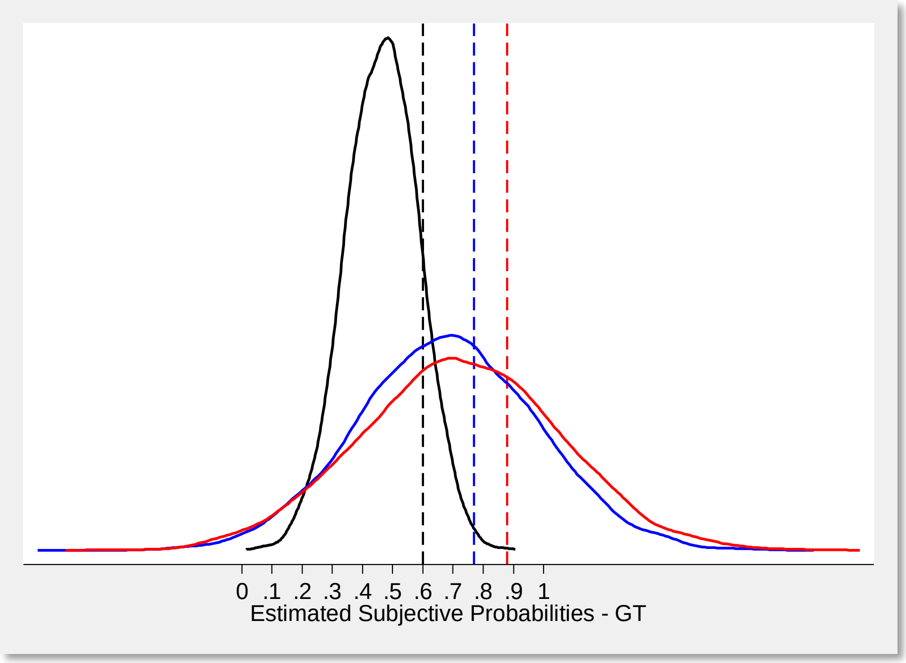

Finally, if we restrict the demographic factors to gender only, we can more easily compare the central tendencies of the three subjective probabilities under either SEU or the Strength/Weight heuristic. Figure 137 reports the kernel densities of the probability estimates for each of the models:

**Figure 137 - Estimated Subjective Probabilities using SEU and Strength/Weight Heuristic, (sex only)**

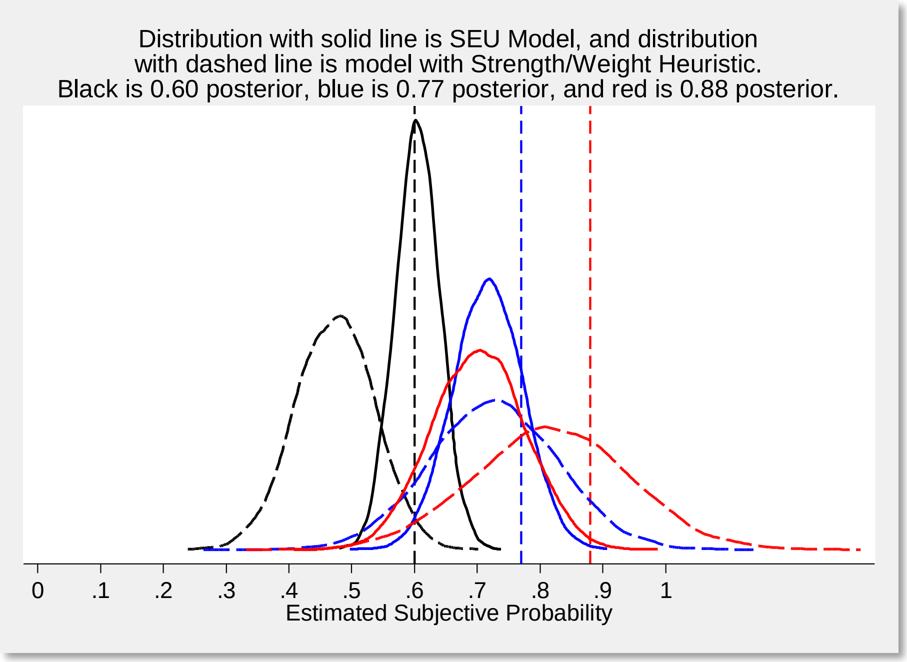

Here we see that the strength/weight model generally performs worse than the SEU model for p=.6 and .77, but strength/weight is superior to SEU at p=.88.

### Clarke and Vuong Statistical Tests of Model Superiority

In Game 1 we specified a 'mixed' structural model' of decision making that lets the data generating process determine the degree to which decision makers are characterized as either operating under expected utility or prospect theory following methods developed by Harrison et al {Harrison, 2009 #9}. The discrimination between Subjective Utility and the Strength/Weight 'heuristic' models of latent decision making within Game 2 can also be undertaken, here using the Vuong and Clarke statistical tests described in Harrison (Harrison 2008) and Andersen (Harrison 2008):

> Whenever one considers two non-nested models, readers expect to see some comparative measures of goodness of fit. …Common measures include R2, pseudo-R2, a "hit ratio," some other scalar appropriate for choice models (e.g., Hosmer and Lemeshow [2000; ch.5]), and formal likelihood-ratio tests of one model against another (e.g., Cox [1961][1962] or Vuong [1989]). From the perspective [of determining the possible mixture of data generating processes], the interpretation of these tests suffers from the problem of implicitly assuming just one data-generating process. In effect, the mixture model provides a built-in comparative measure of goodness of fit - the mixture probability itself. If this probability is close to 0 or 1 by standard tests, one of the models is effectively rejected, in favor of the hypothesis that there is just one data-generating process.
>
> In fact, if one traces back through the literature on non-nested hypothesis tests, these points are "well known." That literature is generally held to have been started formally by Cox [1961], who proposed a test statistic that generalized the usual *likelihood ratio test* (LRT). His test compares the difference between the actual LRT of the two models with the expected LRT, suitably normalized by the variance of that difference, under the hypothesis that one of the models is the true data-generating process. The statistic is applied symmetrically to both models, in the sense that each takes a turn at being the true model, and leads to one of four conclusions: one model is the true model, the other model is the true model, neither model is true, or both models are true.
>
> Perhaps the most popular modern variant of the generalized LRT approach of Cox [1961][1962] is due to (Vuong 1989). He proposes the null hypothesis that both models are the true models, and then allows two one-sided alternative hypotheses[72]. The statistic he derives takes observation-specific ratios of the likelihoods under each model, so that …the ratio for observation i is the likelihood of observation i under SEUT divided by the likelihood of observation I under [the strength/weight heuristic]. It then calculates the log of these ratios, and tests whether the expected value of these log-ratios over the sample is zero. Under reasonably general conditions a normalized version of this statistic is distributed according to the standard normal, allowing test criteria to be developed.
>
> Clarke (Clarke 2007) proposes a non-parametric sign test be applied to the sample of ratios. Clarke demonstrates that when the distribution of the log of the likelihood ratios is normally distributed that the Vuong test is better in terms of asymptotic efficiency. But if

[72] Andersen: "Some have criticized the Vuong test because the null hypothesis is often logically impossible, but it can also be interpreted as the hypothesis that one cannot say which model is correct." (Harrison 2008)

> this distribution exhibits sharp peaks, in the sense that it is mesokurtic, then the non-parametric version is better. … Many of the likelihood ratios we deal with have the latter shape…so it is useful to be armed with both tests.

In Game 2 we undertake both the Vuong and Clark Tests to determine whether the SEUT or GT model better describes the data generating process. The SEUT model was rerun using the same demographics used in Panel C2 for the GT model for direct comparison73 and the non-nested model tests were undertaken in STATA based on Harrison's supplied code. The revised SEUT model output (to be compared to the Panel C2 output) and the test results are displayed below [74]:

**Table 30 – Game 2 Panel D1: Maximum Likelihood Estimates under SEUT, w. comparative demographics**

```
                                                Number of obs   =      25327
                                                Wald chi2(5)    =       9.09
Log pseudolikelihood = -14242.108               Prob > chi2     =     0.1054

                                  (Std. Err. adjusted for 53 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
         age |   .0472371   .0176056     2.68   0.007     .0127309    .0817434
         sex |   .0477897   .0858943     0.56   0.578    -.1205601    .2161394
         bus |   .0774373   .0999445     0.77   0.438    -.1184502    .2733248
        math |    .008527   .0458962     0.19   0.853    -.0814279     .098482
       certs |   .0187484   .0307287     0.61   0.542    -.0414788    .0789756
       _cons |  -.1042465   .0891244    -1.17   0.242    -.2789272    .0704342
-------------+----------------------------------------------------------------
sprob1_      |
         age |  -.1170676   .0974391    -1.20   0.230    -.3080448    .0739095
         sex |   .2550122   .2900993     0.88   0.379    -.3135721    .8235964
         bus |  -.2738571   .5701508    -0.48   0.631    -1.391332    .8436179
        math |   .0562761   .2635795     0.21   0.831    -.4603303    .5728825
       certs |   .0418312   .1162265     0.36   0.719    -.1859685    .2696308
       _cons |  -.3499684   .2654682    -1.32   0.187    -.8702766    .1703398
-------------+----------------------------------------------------------------
sprob2_      |
         age |   .0763212    .250247     0.30   0.760    -.4141538    .5667963
         sex |   .8626605   .5970009     1.44   0.148    -.3074398    2.032761
         bus |  -.7367642   2.930545    -0.25   0.801    -6.480527    5.006999
        math |  -.0704714   1.092504    -0.06   0.949    -2.211739    2.070797
       certs |  -.1375253   .8359328    -0.16   0.869    -1.775923    1.500873
       _cons |  -1.206342   .5860682    -2.06   0.040    -2.355014    -.057669
-------------+----------------------------------------------------------------
sprob3_      |
         age |  -.0779608    .274798    -0.28   0.777     -.616555    .4606334
         sex |   .1804514   .5533543     0.33   0.744     -.904103    1.265006
         bus |   .4541235   1.040351     0.44   0.662    -1.584927    2.493174
        math |   .2326474    .591559     0.39   0.694     -.926787    1.392082
```

[73] The SEUT model was able to run using *age, sex, bus, math, workyears and income* demographic dummies, but the GT ml model would not converge without removing *workyears* and *income*. The SEUT model was therefore rerun removing workyears and income in order to directly compare SEUT with GT using the non-nested tests described in this section.

[74] All STATA results are included in the Appendix.

```
       certs |  -.1653186   .2637817    -0.63   0.531    -.6823214    .3516841
       _cons |  -.7711871   .5228387    -1.47   0.140    -1.795932     .253558
-------------+----------------------------------------------------------------
LNmuB        |
       _cons |  -.8331817   .3516986    -2.37   0.018    -1.522498   -.1438652
```

Figure 138 indicates the resulting ratio of log likelihoods from the GT and SEU models vs. a normal distribution. These results are similar to those found by Antoniou and Harrison for this experiment: we see that the empirical distribution (ratio of GT to SEU log likelihoods)  is sharp-peaked compared to a Normal distribution fit to the same data and confirms here that the GT model better explains the data:

**Figure 138- Distribution of Log Likelihood Ratios for Non-Nested Tests of GT vs. SEU Model**

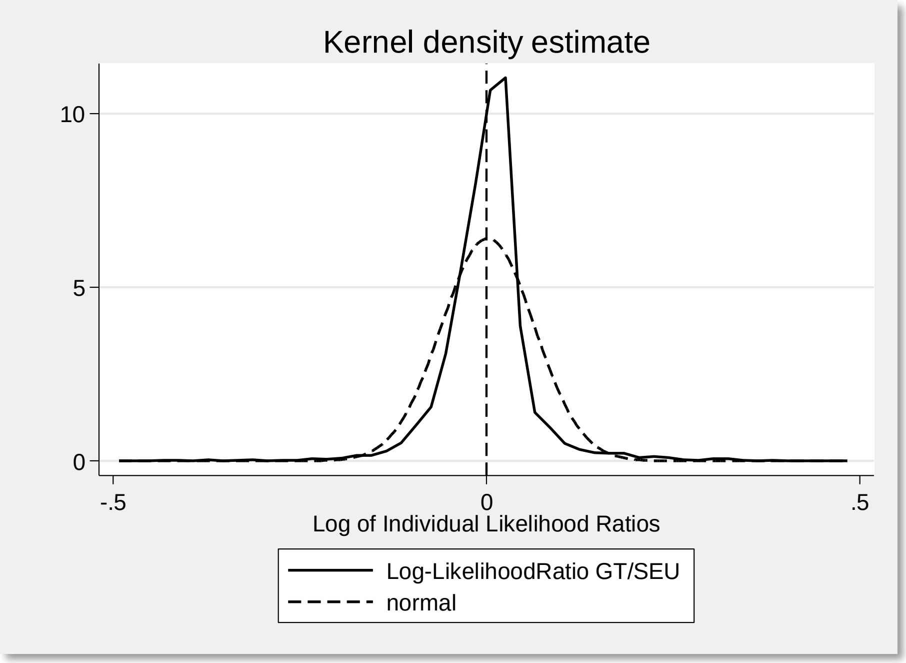

Table 31 indicates the result of the two tests and confirms that GT is the better model:

**Table 31 – Vuong and Clarke Test Results for GT vs. SEU Model Superiority**

```
Log-likelihood for SEU is -14242.10756685881 and for GT is -14316.31637639697

Vuong statistic for test of SEU and GT is 4.404718065925073 and favors SEU over
GT if positive (p-value = 5.29607719557e-06 of SEU not being the better model)

Clarke statistic for test of SEU and GT is 13779 and favors SEU over GT if
greater than 12663.5 (p-value = 5.76207345444e-45 of SEU not being the better
model)
```

**Game 3 Findings**

**Exogenous versus Endogenous Risk Beliefs**

**Analysis**

Following Sen (Andersen, Harrison et al. 2006) we conduct an experiment that involves both the estimation of subjective probabilities over exogenous risks and the consideration of these probability estimates when the risk is endogenous i.e. in situations where the decision maker can alter the objective probabilities of an outcome by taking some action with uncertain effect. Harrison and Rutström (2008) provide an extensive review of the procedures used to elicit and estimate risk attitudes in the exogenous risk setting, and we build on the toolkit reviewed there to consider what modeling issues have to be considered as one moves to model endogenous risk and risk management choice options.

The first extension of the modeling undertaken in Games 1 and 2 above therefore is to consider the possible role of endogeneity on risk attitudes, in the sense that attitudes towards risk might be somehow "source dependent" on the type of stochastic process generating the outcome (Loewenstein and Issacharoff 1994; Tversky and Wakker 1995). It is one thing to make risk averse bets over processes that cannot be controlled; it is another to do so when you can (possibly, but uncertainly) alter the stochastic process itself.

The second extension is to consider the implications of risk mitigation on subjective probabilities. I propose that, particularly in the context of information security control decisions, it is rare that the agent has a reliable expectation regarding the means by which different mitigation actions change the objective probabilities of losses, even if the agent could assume that the initial, no-mitigation probabilities were in fact objective. So we have to consider the elicitation of subjective probabilities. In principle, so far, this involves no radical extension of EUT, interpreted in the sense of Subjective EUT (SEUT).

For both of the following treatments, as in Game 2, we assume that the utility of income is defined by

$$U(x) = x^{1-r} / (1-r) \qquad (9.30)$$

where "r" is the 'coefficient of risk aversion' and "x" is the monetary value of the outcome (>0), where r=0 corresponds to risk neutrality, r<0 to risk loving, and r>1 to risk aversion. The coefficient of risk aversion is constant for the power function family. 'Expected utility' is then specified as the *probability conditioned* utility,

$$EU_i = \sum_{k=1,K} (p_k \times U(x)_k)$$

(9.31)

where p is the probability of the outcome and U(x) is the CRRA specified utility function.

**Exogenous Risk Treatment**

In the first task, the participant bets as to whether a randomly selected daily loss generated by a Low Control System will exceed $17,000. Unlike Game 2 the participant does not know the relative probabilities of each system generating losses and these must be inferred subjectively from test simulations of the system undertaken prior to the betting task. Table 32 sets out the lotteries over which the participant makes bets that a daily loss will exceed $17,000 with 9 bookies:

**Table 32 - Game 3A: Betting Table**

| Bookie | Your Stake | Scenario A.<br>You bet that the daily loss will exceed $17,000. If it… | | Scenario B.<br>You bet that daily loss will NOT exceed $17,000. If it… | |
|---|---|---|---|---|---|
| 1 | $5 | does, the payout is | $50.00 | does, the payout is | $0.00 |
| | | does not, the payout is | $0.00 | does not, the payout is | $5.55 |
| | | | | | |
| 2 | $5 | does, the payout is | $25.00 | does, the payout is | $0.00 |
| | | does not, the payout is | $0.00 | does not, the payout is | $6.25 |
| | | | | | |
| 3 | $5 | does, the payout is | $16.66 | does, the payout is | $0.00 |
| | | does not, the payout is | $0.00 | does not, the payout is | $7.19 |
| | | | | | |
| 4 | $5 | does, the payout is | $12.50 | does, the payout is | $0.00 |
| | | does not, the payout is | $0.00 | does not, the payout is | $8.33 |
| | | | | | |
| 5 | $5 | does, the payout is | $10.00 | does, the payout is | $0.00 |
| | | does not, the payout is | $0.00 | does not, the payout is | $10.00 |
| | | | | | |
| 6 | $5 | does, the payout is | $8.33 | does, the payout is | $0.00 |
| | | does not, the payout is | $0.00 | does not, the payout is | $12.50 |
| | | | | | |
| 7 | $5 | does, the payout is | $7.19 | does, the payout is | $0.00 |
| | | does not, the payout is | $0.00 | does not, the payout is | $16.66 |
| | | | | | |
| 8 | $5 | does, the payout is | $6.25 | does, the payout is | $0.00 |
| | | does not, the payout is | $0.00 | does not, the payout is | $25.00 |
| | | | | | |
| 9 | $5 | does, the payout is | $5.55 | does, the payout is | $0.00 |
| | | does not, the payout is | $0.00 | does not, the payout is | $50.00 |

Considering Table 32, the participant that believes the loss will exceed $17,000 from a given system (event "A" or the 'left hand' binary choice) and selecting "A" from a given bookie 'b' receives EU

$$\begin{gathered} EU_{\text{WILL EXCEED \$17K}} = \pi_{\text{WILL EXCEED \$17K}} \times U(\text{payout if the loss exceeds \$17,000} \mid \text{bet that the loss exceeds} \\ \text{\$17,000}) + \\ (1- \pi_{\text{WILL EXCEED \$17K}}) \times U(\text{payout if the loss does not exceed \$17,000} \mid \text{bet that the loss exceeds \$17,000}) \end{gathered} \tag{9.32}$$

where $\pi_{\text{WILL EXCEED \$17K}}$ is the subjective probability that the loss exceeds \$17,000. The payouts that enter the utility function are shown in Table 32 above. For the bet offered by the first bookie, for example, these payouts are \$50 and \$0, so we have

$$EU_{\text{EXCEEDS \$17K}} = \pi_{\text{WILL EXCEED \$17K}} \times U(\$50) + (1- \pi_{\text{WILL EXCEED \$17K}}) \times U(\$0) \qquad (9.33)$$

Where the participant potentially earns \$50 if correct but nothing if incorrect.

We similarly define the EU received from a bet that the loss will exceed not exceed \$17,000 (event “B” or the ‘right hand’ binary choice) as the complement of event A:

$$\begin{gathered} EU_{\text{WILL NOT EXCEED \$17K}} = \pi_{\text{WILL EXCEED \$17K}} \times U(\text{payout if the loss exceeds \$17,000 } \mid \text{bet that the loss will not} \\ \text{exceed \$17,000}) + \\ (1- {}_{\text{WILL EXCEED \$17K}}) \times U(\text{payout if the loss does not exceed \$17,000} \mid \text{bet that the loss will not exceed} \\ \text{\$17,000}) \end{gathered} \qquad (9.34)$$

and this translates for the first bookie in Table 32 into payouts of \$0 and \$5.55 (i.e. the participant potentially earns \$0 if incorrect as usual, but \$5.55 if correct, so we have

$$EU_{\text{WILL NOT EXCEED \$17K}} = \pi_{\text{WILL EXCEED \$17K}} \times U(\$0) + (1- {}_{\text{WILL EXCEED \$17K}}) \times U(\$3.15) \qquad (9.35)$$

for this particular bookie and bet. We observe the bet made by the subject for a range of payouts, so we can calculate the likelihood of that choice given values of r (risk aversion), $\pi_{\text{WILL EXCEED \$17K}}$ and $\mu$.

The rest of the structural specification is exactly the same as for the choices over lotteries with objective probabilities. Given the specifications for $EU_{\text{WILL EXCEED \$17K}}$ and $EU_{\text{WILL NOT EXCEED \$17K}}$ above, we can define the latent index as

$$\nabla EU = (EU_{\text{WILL EXCEED \$17K}} - EU_{\text{WILL NOT EXCEED \$17K}}) \qquad (9.36)$$

for each of the ‘Will Exceed’ and ‘Will Not Exceed’ bets from Table 32. We then define the probability of choosing “Will Exceed” as

$$\text{prob(choose B)} = \Phi\, [\, (\nabla EU) / \nu\, ) / \mu\, ] \qquad (9.37)$$

where ν is a normalizing term for the value of each bookie choice pair and μ >0 is a structural "noise parameter" for the belief choices that is used to allow some errors from the perspective of the deterministic SEU model. The normalizing term *ν* is defined as the maximum utility over all prizes in a lottery pair minus the minimum utility over all prizes in a lottery pair, and ensures that the normalized EU difference

$$(EU_{\text{Will Exceed}} - EU_{\text{Wont Exceed}}) / \nu \tag{9.38}$$

remains in the unit interval. This normalization allows one to define robust measures of "stochastic risk aversion," in parallel to the deterministic concepts from traditional theory. Notice that when μ = 1 we return to the original index specification without error. As μ increases, the index in (9.37) falls until, at μ = ∞, the index collapses to zero, so that the probability of either choice becomes ½. In other words, as the noise in the data increases, the model has less and less predictive power until at the extreme the prediction collapses to fifty-fifty or equal likelihood of both choices[75].

Writing out the complete likelihood function, we have

$$\ln L(r, \pi_{\text{WILL EXCEED \$17K}}, \mu; y, X) = \Sigma_i \left[ (\ln \Phi(\nabla EU)\times I(y_i = 1)) + (\ln (1\text{-}\Phi(\nabla EU))\times I(y_i = -1)) \right] \tag{9.39}$$

for the observed choices in the experiment defined over subjective probabilities.

**Endogenous Risk Treatment**

In the second task, the participant starts with a Low Controls system and is given a budget of $17,000 with which they can choose to invest in upgraded security controls to convert to a High Controls System at different price points between $0 and $17,000. After selecting whether to upgrade at each potential price point, one of the price points is chosen at random and a single daily loss is generated from either the Low or High Controls system based on whether the player's chose to upgrade at that price point. The player is instructed that they will get to keep the net amount of the budget left over after deducting the cost of any control purchased less the random loss on the day. Similar to the first treatment, the participant does not know the relative probabilities of each system generating losses and these must be inferred subjectively from test simulations of the system undertaken prior to the betting task, but in this case with the payoff endogenously affected by the player's choice of system at unknown cost. The game is played in 3 rounds at 5, 15 and 20 levels of control cost between $0 and $17,000. Table 33 illustrates the control costs in the

[75] Antoniou: "The normalizing term ν is defined as the maximum utility over all prizes in this lottery pair minus the minimum utility over all prizes in this lottery pair, and ensures that the normalized EU difference [(EUR - EUL)/ν] remains in the unit interval. As μ → ∞ this specification collapses ∇EU to 0 for any values of EUR and EUL, so the probability of either choice converges to ½. So a larger μ means that the difference in the EU of the two lotteries, conditional on the estimate of r, has less predictive effect on choices. Thus μ can be viewed as a parameter that flattens out, or "sharpens," the link functions implicit in (4). This is just one of several different types of error story that could be used, and Wilcox (2008) provides a masterful review of the implications of the strengths and weaknesses of the major alternatives." (Antoniou, Harrison et al. 2010)

15-level treatment over which the participant must choose to either upgrade to a High Controls System or stay with the Low Controls system at zero incremental cost:

**Table 33 - Game 3B: 15-Level Control Pricing Table**

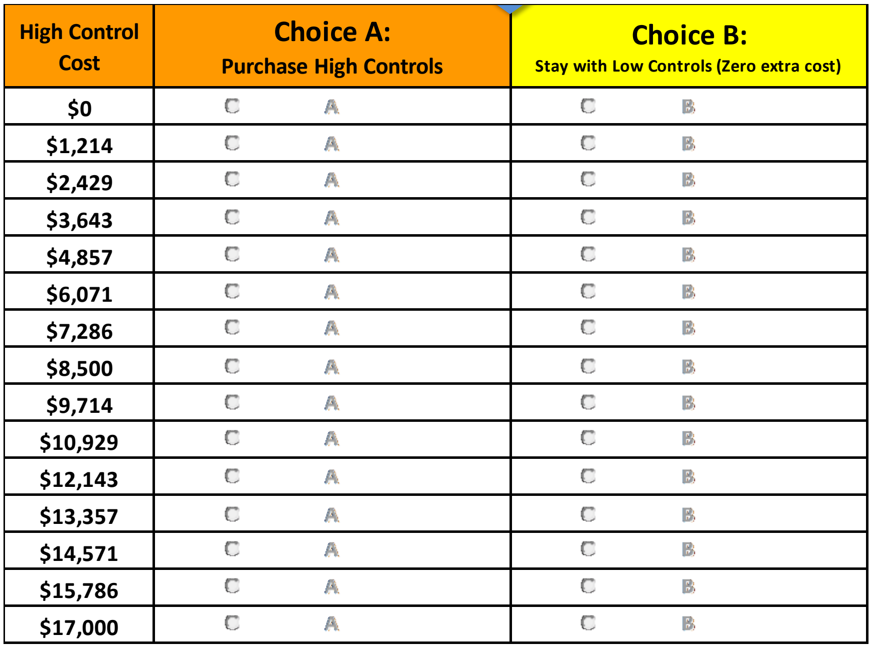

| High Control Cost | Choice A:<br>Purchase High Controls | Choice B:<br>Stay with Low Controls (Zero extra cost) |
|---|---|---|
| $0 | ○ A | ○ B |
| $1,214 | ○ A | ○ B |
| $2,429 | ○ A | ○ B |
| $3,643 | ○ A | ○ B |
| $4,857 | ○ A | ○ B |
| $6,071 | ○ A | ○ B |
| $7,286 | ○ A | ○ B |
| $8,500 | ○ A | ○ B |
| $9,714 | ○ A | ○ B |
| $10,929 | ○ A | ○ B |
| $12,143 | ○ A | ○ B |
| $13,357 | ○ A | ○ B |
| $14,571 | ○ A | ○ B |
| $15,786 | ○ A | ○ B |
| $17,000 | ○ A | ○ B |

In this game $\pi^{Safe}_{\text{LOSS EXCEEDS \$17K}}$ is the subjective probability of the loss exceeding $17,000 when High Controls are purchased, and $\pi^{Risky}_{\text{LOSS EXCEEDS \$17K}}$ is the subjective probability of the loss exceeding $17,000 when the player stays with the Low Controls system. More generally:

$$\begin{aligned}\text{EU}_{\text{Safe}} = {} & \pi^{Safe}_{\text{LOSS EXCEEDS \$17K}} \times \text{U(payout net of cost of control upgrade if the loss exceeds \$17,000)} + \\ & (1- \ \pi^{Safe}_{\text{LOSS EXCEEDS \$17K}}) \times \text{U(payout net of cost of control upgrade if the loss does not exceed \$17,000)}\end{aligned} \tag{9.40}$$

and

$$\begin{aligned}\text{EU}_{\text{Risky}} = {} & \pi^{Risky}_{\text{LOSS EXCEEDS \$17K}} \times \text{U(payout with no control costs if the loss exceeds \$17,000)} + \\ & (1 - \pi^{Risky}_{\text{LOSS EXCEEDS \$17K}}) \times \text{U(payout with no control costs if the loss does not exceed \$17,000)}\end{aligned} \tag{9.41}$$

The latent index in this problem is the difference in EU from paying for upgraded controls and not paying for upgraded controls, using the error specification form described for the exogenous risk treatment:

$$\nabla\text{EU} = [(\text{EU}_{\text{Risky}} - \text{EU}_{\text{Safe}}) / \nu\ ] / \mu \tag{9.42}$$

An increase in this index should increase the likelihood of selecting the risky option, i.e. of not paying for High Controls. Apart from r, *v* we now need to estimate the two beliefs $\pi^{Safe}$ and $\pi^{Risky}$.

**Results**

**Exogenous Risk Treatment**

Panel A reports the maximum likelihood estimates of the coefficient of risk aversion, the subjective probability of a loss exceeding $17,000 for a Low Controls System and the behavioural error term:

**Table 34 - Game 3 Panel A: Maximum Likelihood Estimates Assuming SEUT – Low Controls System**

**Subjective probability of a daily loss exceeding $17,000 under exogenous risk**

```
                                                  Number of obs   =        504
                                                  Wald chi2(0)    =          .
Log pseudolikelihood =  -251.7711                 Prob > chi2     =          .

                                   (Std. Err. adjusted for 48 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
       _cons |   .1188801   .1213645     0.98   0.327    -.1189899      .35675
-------------+----------------------------------------------------------------
sprob_       |
       _cons |  -.7446101   .1873644    -3.97   0.000    -1.111838   -.3773826
-------------+----------------------------------------------------------------
LNmuRA       |
       _cons |  -1.027557   .2329254    -4.41   0.000    -1.484083   -.5710318
------------------------------------------------------------------------------

       sprob:  1/(1+exp([sprob_]_b[_cons]))

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       sprob |   .6780031   .0409044    16.58   0.000     .5978319    .7581744
------------------------------------------------------------------------------
```

For the Low Controls System, we see moderate risk aversion (r > 0) although the estimate is not statistically different than zero. The subjective probability of a loss exceeding $17,000 is .67, which is more than double the actual probability for the Low Controls system (.3).

Panel B reports the maximum likelihood estimates of the coefficient of risk aversion, the subjective probability of a loss exceeding $17,000 for a High Controls System and the behavioural error term:

**Table 35 - Game 3 Panel B: Maximum Likelihood Estimates Assuming SEUT**

**Subjective probability of a daily loss exceeding $17,000
for a High Controls System under exogenous risk**

```
                                                  Number of obs   =        504
                                                  Wald chi2(0)    =          .
Log pseudolikelihood =  -275.6472                 Prob > chi2     =          .

                                    (Std. Err. adjusted for 48 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
       _cons |  -.1308428   .0755151    -1.73   0.083    -.2788498    .0171641
-------------+----------------------------------------------------------------
sprob_       |
       _cons |   1.257398   .4949204     2.54   0.011     .2873714    2.227424
-------------+----------------------------------------------------------------
LNmuRA       |
       _cons |  -.0954333   .3203516    -0.30   0.766     -.723311    .5324443
------------------------------------------------------------------------------

       sprob:  1/(1+exp([sprob_]_b[_cons]))

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       sprob |   .2214222   .0853215     2.60   0.009     .0541951    .3886493
------------------------------------------------------------------------------
```

For the High Controls System, we see marginal risk seeking (r < 0) although the estimate is not statistically different than zero. The subjective probability of a loss exceeding $17,000 is .22 which, similar to the Low case above, is more than double the actual probability for the High Controls system (.10).

We assume that heterogeneity between the participants might account for the variance in both risk aversion and the degree of overestimation in the subjective probabilities and rerun the estimates using participant demographics as dummy variables:

**Table 36 - Game 3 Panel C: Maximum Likelihood Estimates Assuming SEUT**

**Subjective probability of a daily loss exceeding $17,000
for a Low Controls System under exogenous risk, with demographics**

```
                                                  Number of obs   =        504
                                                  Wald chi2(6)    =      18.68
Log pseudolikelihood =  -231.6402                 Prob > chi2     =     0.0047

                                    (Std. Err. adjusted for 48 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
         sex |   -.177615   .3184661    -0.56   0.577    -.8017972    .4465672
         bus |  -.1029661   .2835908    -0.36   0.717    -.6587937    .4528616
        math |  -.0108591   .2707406    -0.04   0.968    -.5415009    .5197828
   workyears |    .092381   .0568834     1.62   0.104    -.0191085    .2038705
       certs |   .0026191    .121615     0.02   0.983    -.2357419      .24098
      income |  -.4822955   .1513289    -3.19   0.001    -.7788948   -.1856962
       _cons |   .6607168   .2789322     2.37   0.018     .1140197    1.207414
-------------+----------------------------------------------------------------
sprob_       |
         sex |   .3464918   .6296705     0.55   0.582    -.8876396    1.580623
         bus |   .2416479   .3089241     0.78   0.434    -.3638323     .847128
        math |   .2782034   .4126236     0.67   0.500     -.530524    1.086931
   workyears |   .0160379   .0968538     0.17   0.868     -.173792    .2058679
       certs |  -.0442161   .1860682    -0.24   0.812     -.408903    .3204709
      income |   .0849485   .4910882     0.17   0.863    -.8775666    1.047464
       _cons |  -.9959447   1.137945    -0.88   0.381    -3.226276    1.234387
-------------+----------------------------------------------------------------
LNmuRA       |
       _cons |  -1.347072   .5162032    -2.61   0.009    -2.358812   -.3353322
------------------------------------------------------------------------------

       sprob:  1/(1+exp([sprob_]_b[_cons]))

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       sprob |   .7302605   .2241526     3.26   0.001     .2909296    1.169591
------------------------------------------------------------------------------
```

**Table 37 - Game 3 Panel D: Maximum Likelihood Estimates Assuming SEUT**

**Subjective probability of a daily loss exceeding $17,000
for a High Controls System under exogenous risk, with demographics**

```
                                                  Number of obs   =        504
                                                  Wald chi2(2)    =          .
Log pseudolikelihood = -248.68661                 Prob > chi2     =          .

                                     (Std. Err. adjusted for 48 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
         sex |  -.0047875          .        .       .            .           .
         bus |  -.4316819   .1536781    -2.81   0.005    -.7328854   -.1304784
        math |  -.1147367   .1436011    -0.80   0.424    -.3961897    .1667164
   workyears |   .2521705          .        .       .            .           .
       certs |   .0003455          .        .       .            .           .
      income |  -.0668787          .        .       .            .           .
       _cons |  -.4472493          .        .       .            .           .
-------------+----------------------------------------------------------------
sprob_       |
         sex |   .1332954   .8934244     0.15   0.881    -1.617784    1.884375
         bus |  -1.239598   1.104524    -1.12   0.262    -3.404426     .925229
        math |   1.168815   .7903422     1.48   0.139    -.3802275    2.717857
   workyears |      .4911   .4855314     1.01   0.312    -.4605241    1.442724
       certs |  -.3928655    .250548    -1.57   0.117    -.8839305    .0981995
      income |    .047526   .9640091     0.05   0.961    -1.841897    1.936949
```

```
       _cons |    1.12994   1.552742     0.73   0.467    -1.913378    4.173257
-------------+----------------------------------------------------------------
LNmuRA       |
       _cons |   -.163069   .2623071    -0.62   0.534    -.6771816    .3510435
------------------------------------------------------------------------------

       sprob:  1/(1+exp([sprob_]_b[_cons]))

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       sprob |   .2441723   .2865619     0.85   0.394    -.3174787    .8058232
------------------------------------------------------------------------------
```

Figure 139 illustrates the kernel density function of the CRRA estimates across all participants for the Low Controls and High Controls treatments respectively. The results indicate significant diversity in the CRRA across participants, ranging from very risk seeking to risk averse:

**Figure 139 - Kernel density of the Game 3 CRRA estimates**

**High Controls vs. Low Controls System under exogenous risk, with demographics**

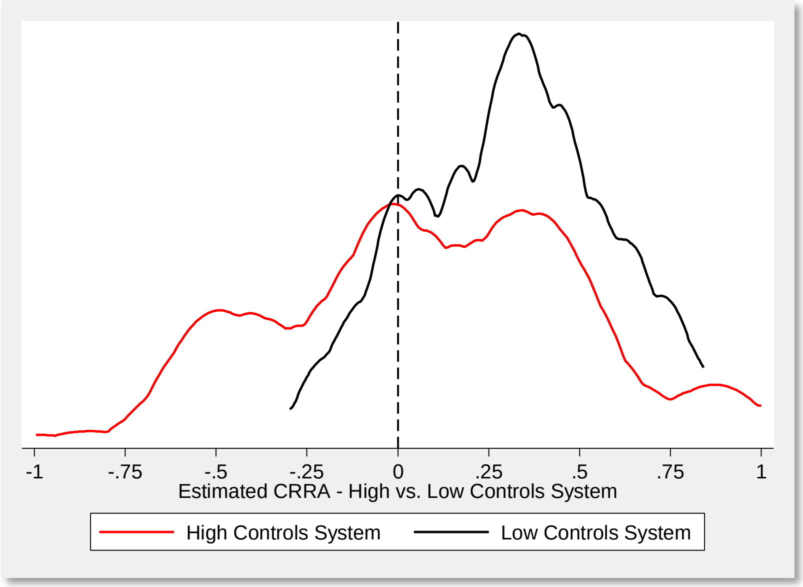

Figure 140 illustrates the kernel density of the subjective probability estimates across all participants for the Low Controls and High Controls treatments respectively functions. The results indicate significant diversity in the subjective probabilities across participants:

**Figure 140 - Kernel density of the Game 3 Subjective Probability Estimates**

**High Controls vs. Low Controls System under exogenous risk, with demographics**

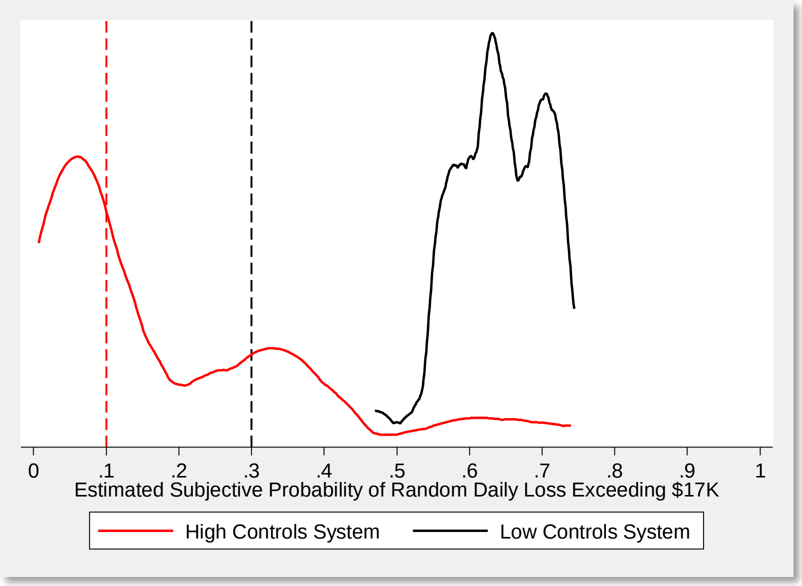

These results indicate that while some participants correctly estimate the probability of the High Controls system generating losses above $17,000, most over-estimate the probability by a factor of 2 to 4 times which is much higher than the results found by Sen. All participants over-estimate the Low Controls probability of generating losses above $17,000. There may be several explanations for this. It clear from the raw data, for example, that some participants may not sufficiently understand the betting mechanism or, more importantly, the implications of the betting mechanism for their payout prospects, or may have simply have ignored the instructions when betting, or may be betting randomly in some cases. For example, some participants do not indicate a single 'switch point' in the multiple price list and many switch more than once between bookies. This is a commonly experienced issue with multiple price list elicitations which strictly violates the monotonicity of the bets and results in larger model errors. Harrison and others discuss the issue at length and, although they report that multiple switch points may not cause significant errors, I suspect that it may do so in my data set given the smaller number of participants and resulting observations undertaken (Holt 2002; Andersen, Harrison et al. 2006; Harrison and Rutström 2008; Dave, Eckel et al. 2010; Bruner 2011; Charness, Gneezy et al. 2013; Freeman, Halevy et al. 2015). Clear violations of expected utility maximization theory are apparent. In the Game 3 Endogenous Risk treatment below, for example, we find many participants electing to expend all of their control budget on a High Controls upgrade (in 30% of these choices) which absolutely ensures no net gain should that prospective outcome occur or, conversely, *not* electing to upgrade at no cost (in 18% of these choices) i.e. apparently preferring to pay for a control that is being offered, potentially, for free. In this way, we observe that participants are not consistent in what are otherwise 'obvious' choices, and we conclude that the data may be noisier than it otherwise should be. Controlling for this is not straightforward: while Harrison and other stress the importance of ensuring that Participants understand how to play the games, they generally stop short of providing any ‘strategic’ guidance that would indicate ‘logical’ or consistent choices to Participants, however objective that choice might be or seem to be. In these experiments I did not attempt to coach players other than indicating that some choices might be more obvious than others in terms of maximizing payouts.

**Endogenous Risk Treatment**

Panel E reports the maximum likelihood estimates of the coefficient of risk aversion, the subjective probability of a loss exceeding $17,000 for a Low Controls System (“sprobRisky”) and High Controls System (“sprobSafe”) respectively and the behavioural error term, where the participant can select to purchase High Controls and therefore endogenously affect the payouts:

**Table 38 - Game 3 Panel E: Maximum Likelihood Estimates Assuming SEUT**

**Subjective probability of a daily loss exceeding $17,000**
**for a High Controls System under endogenous risk**

```
                                                  Number of obs   =       2242
                                                  Wald chi2(0)    =          .
Log pseudolikelihood = -1230.4956                 Prob > chi2     =          .

                                   (Std. Err. adjusted for 58 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
       _cons |   .9996453   .0000436 22919.54   0.000     .9995598    .9997308
-------------+----------------------------------------------------------------
sprobRisky_  |
       _cons |  -1.957071   .0002403 -8143.10   0.000    -1.957542     -1.9566
-------------+----------------------------------------------------------------
sprobSafe_   |
       _cons |   1.958812          .        .       .            .           .
-------------+----------------------------------------------------------------
LNmuRA       |
       _cons |  -8.813299          .        .       .            .           .
------------------------------------------------------------------------------

  sprobRisky:  1/(1+exp([sprobRisky_]_b[_cons]))

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
  sprobRisky |   .8762156   .0000261        .   0.000     .8761645    .8762667
------------------------------------------------------------------------------

   sprobSafe:  1/(1+exp([sprobSafe_]_b[_cons]))

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
   sprobSafe |   .1235956          .        .       .            .           .
------------------------------------------------------------------------------
```

In this treatment we see relatively strong risk aversion (r > 0) and the estimate is statistically significant at the 95% level. The subjective probability of a loss exceeding $17,000 for the Low Controls (risky) System is now .88 which greatly exceeds the actual probability (.3). The subjective probability of a loss exceeding $17,000 for the High Controls (risky) System on the other hand is .12 which is very close to the actual probability (.1) although no standard error statistics are calculable. The interpretation of the subjective probabilities requires careful consideration in the endogenous case particularly since participants only opted for Higher Controls 32% of the time. The maximum likelihood estimator was unable to converge when controlling for demographic variable, and we are therefore only able to report average subjective probabilities for this treatment.

## Game 4 Findings

## Recovering Subjective Probability Distributions

### Analysis

Following Harrison and Ulm  (Harrison and Ulm 2015) we conduct an experiment that involves the estimation of subjective probabilities over continuous distributions. Participants report subjective beliefs over continuous events using a Quadratic Scoring Rule (QSR). Under some mild additional assumption, it has been known since Matheson and Winkler (Matheson and Winkler 1976) that these reports reflect latent subjective beliefs if the individual is risk neutral and obeys Subjective Expected Utility Theory (SEUT). It is also now known that these reports are "close" to latent subjective beliefs if the individual obeys SEU and has a concave utility function in the range observed over typical payments in experiments (Harrison, Martínez-Correa et al. 2013). Following Harrison, we extend these theoretical results in the context of recovering latent subjective beliefs about security outcome represented as times series and probability distributions if the individual is known to distort probabilities into decision weights using Rank Dependent Utility (RDU) theory (Quiggin 1982).

### Recovering Subjective probability Estimates Using Quadratic Scoring Rules

Let the decision maker report his subjective beliefs in a discrete version of a quadratic scoring rule (QSR) for continuous distributions (Matheson and Winkler 1976). Partition the domain of response into K intervals (here we use 10), and denote as $r_k$ the report of the participant of the likelihood that the event falls in interval k = 1, … , K. Assume for the moment that the decision maker is risk neutral, and that the full report consists of a series of reports for each interval, { $r_1, r_2, \ldots, r_k, \ldots, r_K$ } such that $r_k \geq 0\ \forall\ k$ and $\sum_{i = 1 \ldots K} (r_i) = 1$.

If *k* is the interval in which the actual value lies, then the payoff score is defined by Matheson and Winkler [1976; p.1088, equation (6)]:

$$S = (2 \times r_k) - \sum_{i = 1 \ldots K} (r_i)^2 \qquad (9.43)$$

So the reward in the score ($2 \times r_k$) is a doubling of the report allocated to the true interval, and the penalty depends on how these reports are distributed across the K intervals ( $\sum_{i = 1 \ldots K} (r_i)^2$ ).  The subject is rewarded for accuracy, but if that accuracy misses the true interval the punishment is severe. The punishment includes all possible reports, including the correct one.[76]

[76] Harrison "Take some examples, assuming K = 4. What if the subject has very tight subjective beliefs and allocates all of the weight to the correct interval? Then the score is $S = (2 \times 1) - (1^2 + 0^2 + 0^2 + 0^2) = 2 - 1 = 1$, and this is positive. But if the subject has tight subjective beliefs that are wrong, the score is $S = (2 \times 0) - (1^2 + 0^2 + 0^2 + 0^2) = 0 - 1 = -1$, and the score is negative. So we see that this score would have to include some additional "endowment" to ensure that the earnings are positive [in order to calculate the participant's EUT CRRA]. Assuming that the subject has very diffuse subjective beliefs and allocates 25% of the weight to each interval, the score is less than 1: $S = (2 \times ¼) - ((¼)^2 + (¼)^2 + (¼)^2 + (¼)^2 ) = ½ - ¼ = ¼ < 1$. So the tradeoff from the last case is that

To ensure complete generality, and avoid any decision maker facing losses, allow some endowment, α, and scaling of the score, β. We then get the following scoring rule for each report in interval *k*

$$\alpha + \beta\ [\ (2 \times r_k) - \sum_{\text{i = 1 …K}} (r_i)^2] \tag{9.44}$$

where we initially assumed α=0 and β=1. We can assume α>0 and β>0 to get the payoffs to any positive level and units we want. Let $p_k$ represent the underlying, true, latent subjective probability of an individual for an outcome that falls into interval *k*. Figures 141 and 142 illustrate the QSR which we use in the Game 4 experiment, for α = β = 25 and K=10:

one can always ensure a score of ¼, but there is an incentive to provide less diffuse reports, and that incentive is the possibility of a score of 1.

**Figure 141 - Game 4B: Quadratic Scoring Rule Response Device w. Multi-Bin report**

**Example Report ($r_4$ = 20%, $r_5$ = 60%, $r_6$ = 20%, all $r_k$ = 0%)**

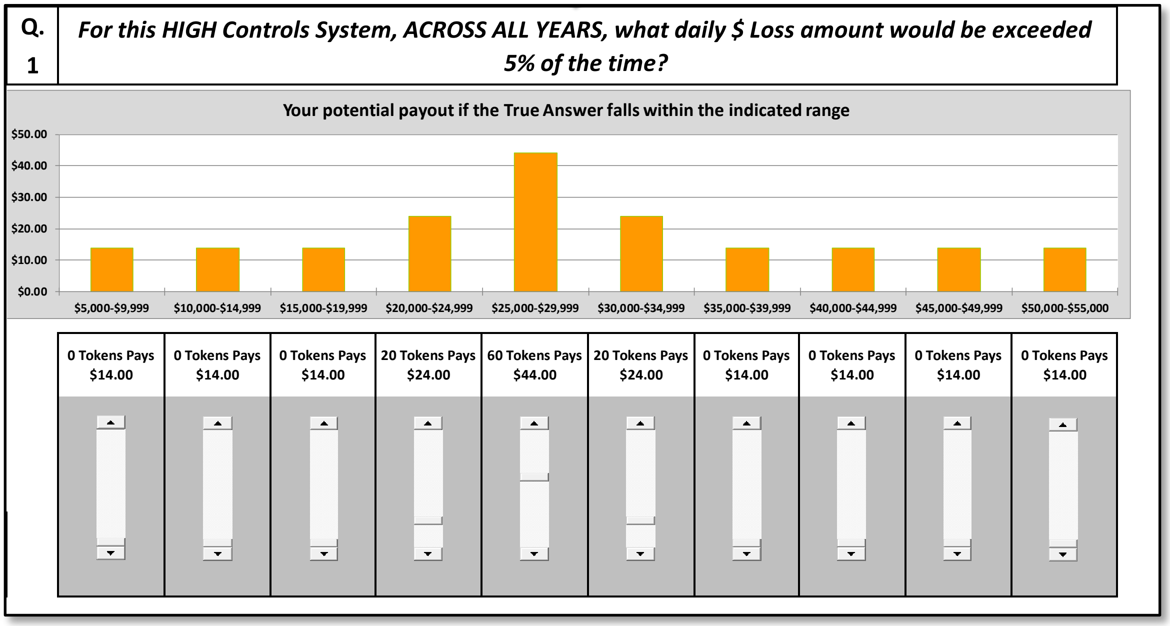

Q. 1

*For this HIGH Controls System, ACROSS ALL YEARS, what daily $ Loss amount would be exceeded 5% of the time?*

Your potential payout if the True Answer falls within the indicated range

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| 0 Tokens Pays $14.00 | 0 Tokens Pays $14.00 | 0 Tokens Pays $14.00 | 20 Tokens Pays $24.00 | 60 Tokens Pays $44.00 | 20 Tokens Pays $24.00 | 0 Tokens Pays $14.00 | 0 Tokens Pays $14.00 | 0 Tokens Pays $14.00 | 0 Tokens Pays $14.00 |

**Figure 142 - Game 4B: Quadratic Scoring Rule Response Device w. Single Bin Report**

**Example Report ($r_5$ = 100%, all $r_{k \neq 5}$ = 0%)**

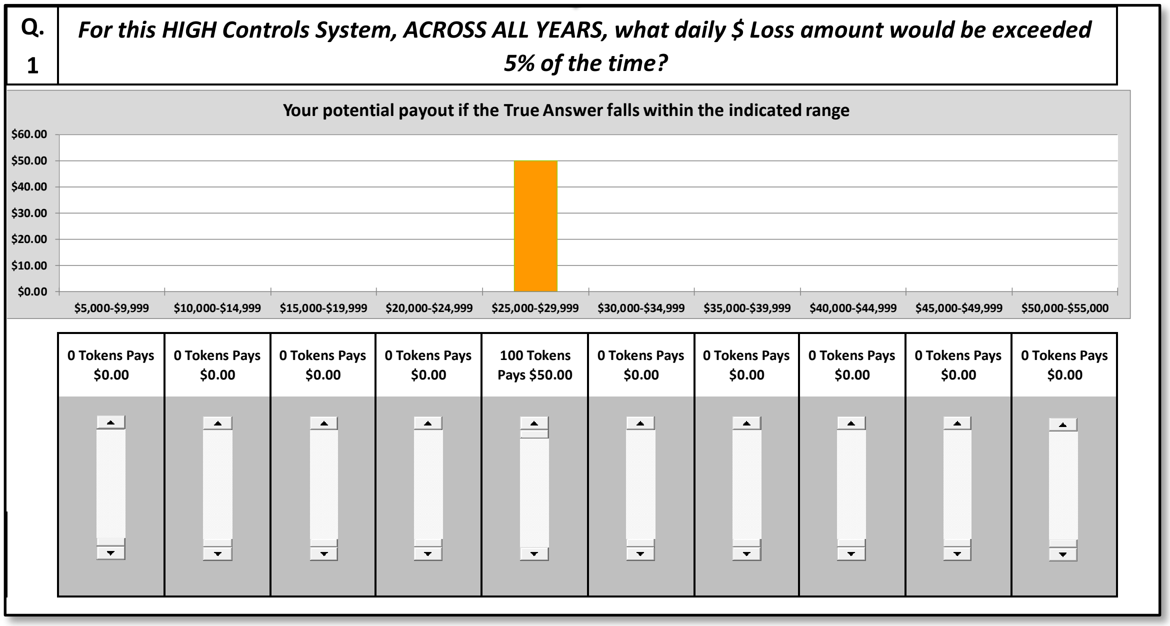

Q. 1

*For this HIGH Controls System, ACROSS ALL YEARS, what daily $ Loss amount would be exceeded 5% of the time?*

Your potential payout if the True Answer falls within the indicated range

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| 0 Tokens Pays $0.00 | 0 Tokens Pays $0.00 | 0 Tokens Pays $0.00 | 0 Tokens Pays $0.00 | 100 Tokens Pays $50.00 | 0 Tokens Pays $0.00 | 0 Tokens Pays $0.00 | 0 Tokens Pays $0.00 | 0 Tokens Pays $0.00 | 0 Tokens Pays $0.00 |

For the following treatments, as in Game 2, we assume that the utility of income is defined by

$$U(x) = x^{1-r} / (1-r) \tag{9.45}$$

where "r" is the 'coefficient of risk aversion' and "x" is the monetary value of the outcome (>0), where r=0 corresponds to risk neutrality, r<0 to risk loving, and r>1 to risk aversion.  'Expected utility' is then specified as the *probability conditioned* utility,

$$EU_i = \sum_{k=1,K} (p_k \times U(x)_k) \tag{9.46}$$

where p is the probability of the outcome and U(x) is the CRRA specified utility function.

**Rank Dependent Utility**

The Rank Dependent Utility (RDU) model of Quiggin (Quiggin 1982) extends the EUT model by allowing for *decision weights* on lottery outcomes i.e. where the decision maker possibly distorts subjective probability estimates based on the rank of the probabilities across outcomes, typically overweighting low probability events. The specification of the utility function is the same parametric specification considered above for EUT. To calculate decision weights under RDU one replaces expected utility with RDU:

$$RDU_i = \sum_{j=1,J} [\, w\left(p(M_j)\right) \times U(M_j)] = \sum_{j=1,J} [\, w_j \times U(M_j)] \tag{9.47}$$

where

$$w_j = \omega(p_j + \ldots + p_J) - \omega(p_{j+1} + \ldots + p_J) \tag{9.48}$$

for j=1, ... , J-1, and

$$w_j = \omega(p_j) \tag{9.49}$$

for j=J, with the subscript j ranking outcomes from worst to best (i.e. in a 'decumulative' ranking), and ω(.) is some probability weighting function.

Several probability weighting functions can be considered[77] – in this study I consider the “Power” probability weighting function popularized by Quiggin with curvature parameter $\gamma$:

$$\omega(p) = p^{\gamma} \tag{9.50}$$

So $\gamma \neq 1$ is consistent with a deviation from the conventional EUT representation. Convexity of the probability weighting function is said to reflect “pessimism” and generates, if one assumes for simplicity a linear utility function, a risk premium since $\omega(p) < p \; \forall \, p$ and hence the “RDU EV” weighted by $\omega(p)$ instead of p has to be less than the EV weighted by p.

**Recovering subjective probabilities under EUT and RDU:**

Let $p_k$ represent the underlying subjective probability of an individual for outcome $k$ and let $r_k$ represent the reported probability for outcome k in a given scoring rule. Let $\theta(k) = \alpha + \beta 2 r_k - \beta \sum_{i = 1 \ldots K} (r_i)^2$ be the (Quadratic) scoring rule that determines earnings $\theta$ if state $k$ occurs. Assume that the individual uses some probability weighting function $\omega(.)$, leading to decision weights $w(.)$ defined in the ‘decumulative’ fashion noted above. Assume that the individual behaves consistently with RDU, applied to subjective probabilities. If the individual has a utility function $u(.)$ that is continuous, twice differentiable, increasing and concave and maximizes rank dependent utility over weighted subjective probabilities, the actual and reported probabilities must obey the following system of equations:

$$w(p_k) \times \partial u/\partial\theta \,|_{\theta = \theta(k)} - \sum_{j = 1,K} \{ w(p_j) \times r_j \times \partial u/\partial\theta \,|_{\theta = \theta(j)} \} = 0, \forall \; k = 1, \ldots , K \tag{9.51}$$

The application of this system of equations is straightforward. If the reports $r_k$ are given from observation of experimental data, the partial derivatives are fixed and independent of the decision weights $w(p_k)$, so this is a linear system of equations in the unknown decision weights. These equations can be solved using standard linear algebra techniques. Although it turns out the equations are linearly dependent, we can replace any one of them with $\sum_{k = 1,K} \{ w(p_k) \} = 1$ to remove the redundancy and obtain a unique solution[78].

[77] Harrison and Ulm employ three utility functions: Power, “inverse-S” and Prelec (Harrison and Ulm 2015)

[78] Ibid. Harrison and Ulm present a numerical example for greater exposition – the reader is referred to their example for that detail.

**Results:**

**CRRA and RDU Weight Estimation Using Objective Lotteries**

Panel A reports the maximum likelihood estimates of the coefficient of risk aversion, and the behavioural error term elicited across all participants:

**Table 39 - Game 4A Panel A: Maximum Likelihood Estimates Assuming EUT**
**Estimate of the CRRA, all participants**

```
                                                  Number of obs   =       1750
                                                  Wald chi2(0)    =          .
Log pseudolikelihood =   -1205.26                 Prob > chi2     =          .

                                   (Std. Err. adjusted for 35 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
       _cons |   .3097792   .6381193     0.49   0.627    -.9409116     1.56047
-------------+----------------------------------------------------------------
LNmuRA       |
       _cons |    .259151   .3252643     0.80   0.426    -.3783552    .8966573
------------------------------------------------------------------------------

        muRA:  exp([LNmuRA]_cons)

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
        muRA |    1.29583   .4214871     3.07   0.002     .4697301    2.121929
------------------------------------------------------------------------------
```

Across undifferentiated participants we see moderate risk aversion (r >0) although the estimate is not statistically different than zero. In Panel B we assume that heterogeneity between the participants might account for the variance in risk aversion and rerun the estimates using participant demographics as dummy variables:

**Table 40 - Game 4A Panel B: Maximum Likelihood Estimates Assuming SEUT**
**Predicted values of the CRRA for all participants, with demographics**

```
                                                  Number of obs   =       1750
                                                  Wald chi2(2)    =          .
Log pseudolikelihood = -1202.2823                 Prob > chi2     =          .

                                   (Std. Err. adjusted for 35 clusters in id)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
         sex |   .3729342          .        .       .            .           .
         bus |  -.3419725   1.653118    -0.21   0.836    -3.582025     2.89808
        math |    .435237          .        .       .            .           .
```

```
   workyears |   .1050916           .           .        .             .             .
       certs |  -.2141195    .4998308     -0.43    0.668      -1.19377      .7655309
      income |  -.3187145           .           .        .             .             .
       _cons |    .302607           .           .        .             .             .
-------------+----------------------------------------------------------------
LNmuRA       |
       _cons |    .197197    .4905171      0.40    0.688     -.7641989      1.158593
------------------------------------------------------------------------------

        muRA:  exp([LNmuRA]_cons)

------------------------------------------------------------------------------
             |      Coef.   Std. Err.       z     P>|z|      [95% Conf. Interval]
-------------+----------------------------------------------------------------
        muRA |   1.217984    .5974419      2.04    0.041      .0470192      2.388949
------------------------------------------------------------------------------
```

Figure 143 illustrates the kernel density function of the CRRA estimates across all participants. The results indicate significant diversity in the CRRA across participants, ranging from risk seeking to risk averse:

**Figure 143 - Kernel density for Game 4A CRRA estimates for all participants, with demographics**

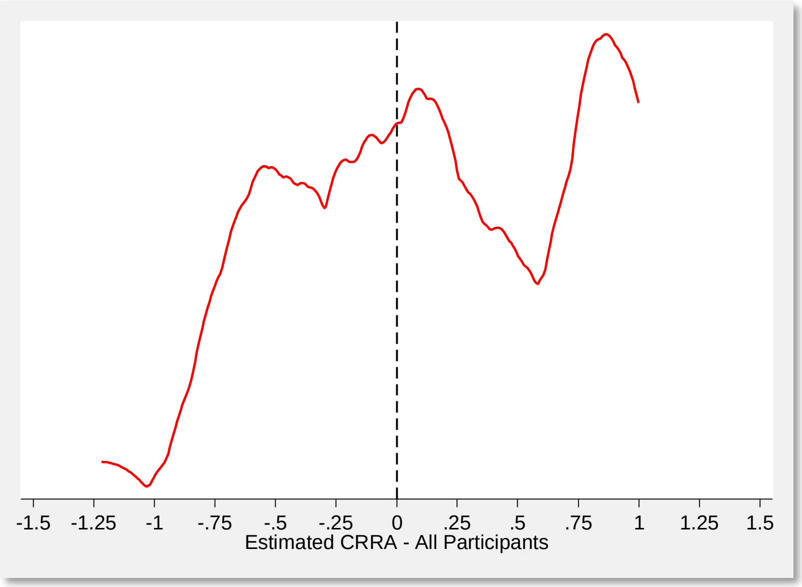

To evaluate RDU preferences, we estimate an RDU model across all participants and for each individual, following procedures explained by Harrison and Rutström (Harrison and Rutström 2008). We consider the CRRA utility function and the Power probability weighting function defined above. For our purposes, the only issue is at what statistical confidence level we can (or cannot) reject the EUT hypothesis that ω(p) =p. After estimating γ (gamma) for each individual, we run t-tests on the hypothesis ω(p) = p, or γ = 1. Unfortunately, only three of the 60 individuals generate plausible gamma values where the hypothesis could be rejected at the 10% level. Figure 144 indicates the kernel density function of the p-values for the RDU estimates – vertical red lines indicate p-values of .01, .05 and .1, left to right respectively:

**Figure 144 - Kernel density of Game 4A p-values of test of EUT**

**$H_0$: ($\omega(p) = p$, or $\gamma = 1$)**

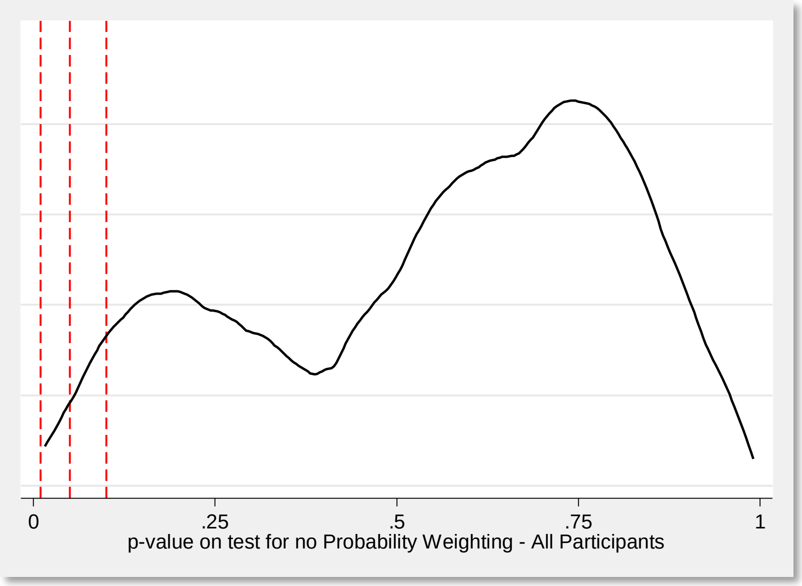

To demonstrate the potential effect of RDU on recovered beliefs, we select results from one of the participants where the gamma parameter estimate is significant at the 5% level. Panel C reports the maximum likelihood estimates of this participant's coefficient of risk aversion, and the behavioural error term:

**Table 41 - Game 4A Panel C: Maximum Likelihood Estimates Assuming SEUT Estimates of the CRRA, participant ID 52**

```
                                                Number of obs   =         50
                                                Wald chi2(0)    =          .
Log likelihood = -34.245068                     Prob > chi2     =          .

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
       _cons |   .9995994   3.643209     0.27   0.784    -6.140959    8.140158
-------------+----------------------------------------------------------------
muRA         |
       _cons |   .8203305   1.660168     0.49   0.621    -2.433539      4.0742
------------------------------------------------------------------------------
```

The CRRA estimate itself is not significant, however when we run an RDU model for this participant, we get a significant estimate for gamma at the 5% level. Panel D reports the maximum likelihood estimates of this participant's gamma coefficient (under Power Utility), the CRRA term and the behavioural error term:

**Table 42 - Game 4A Panel D: Maximum Likelihood Estimates Assuming SEUT Estimates of CRRA and Gamma (RDU Model), participant ID 52**

```
                                                   Number of obs   =         50
                                                   Wald chi2(0)    =          .
Log likelihood = -31.664484                        Prob > chi2     =          .

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
r            |
       _cons |  -3.978473   3.367166    -1.18   0.237      -10.578    2.621051
-------------+----------------------------------------------------------------
gamma        |
       _cons |   3.892763   1.611483     2.42   0.016     .7343152    7.051211
-------------+----------------------------------------------------------------
muRA         |
       _cons |   .2605455   .1565854     1.66   0.096    -.0463563    .5674473
------------------------------------------------------------------------------

. test [gamma]_b[_cons]=1

           chi2(  1) =    3.22
         Prob > chi2 =    0.0726
```

In order to demonstrate the effect of assuming different models of risk preference on recovering subjective probabilities, we use the gamma estimate, together with the CRRA estimate from Panel D to recover the subjective beliefs of Participant #52 based on their original reports for 3 questions regarding the characteristics of a particular distribution of business losses attributed to security breach. Our results confirm those of Harrison and Ulm:

> The figures illustrate a central theme that goes back to Savage (Savage 1971): one cannot recover subjective beliefs without making some assumptions about the underlying model of risk preferences. Those assumptions might take the form of designing an elicitation procedure that is assumed to "risk neutralize" the individual (e.g. (Koszegi and Rabin 2008) and (Karni 2009)), applying a payoff procedure that is assumed to "risk neutralize" the individual (e.g., (Smith 1961)] and (Harrison, Martínez-Correa et al. 2014)), or just assuming, contrary to the evidence, that individuals are risk neutral. (Harrison and Ulm 2015)

In this case there are positive reports for 3 or 4 'bins' in each question, and neither the EUT or the RDU model of risk preferences assigns any subjective belief to the bins that have zero reports. The first bar of each bin in the following figures shows the observed report. The second bar of each bin shows the recovered belief assuming that this individual behaved as if they were an EUT decision-maker, and further had a CRRA coefficient of .99. This coefficient was estimated for this individual from the separate task of

50 lottery choices noted above. Again, as expected from the theoretical results of Harrison et al (Harrison, Martínez-Correa et al. 2013), we do not see a significant difference between the beliefs recovered under EUT from the reported beliefs.

The third bar of each bin show the dramatic effect of assuming different RDU models, where the difference derives solely from different assumptions about the probability weighting function. Here we see a large shift in subjective probabilities if we assume the individual is an RDU decision maker with a power probability weighting function.

Hence the bottom line for this subject and his recovered subjective beliefs about these three questions regarding the probability distributions is to compare the reported belief to the final bar within each bin.

Figure 145 - Game 4B Sample Question #1

**Question #1:**

*"For this HIGH Controls System, ACROSS ALL YEARS, what* ***daily $ Loss amount*** *would be exceeded 50% of the time?*

**Actual Answer:** $3,026 (red line)

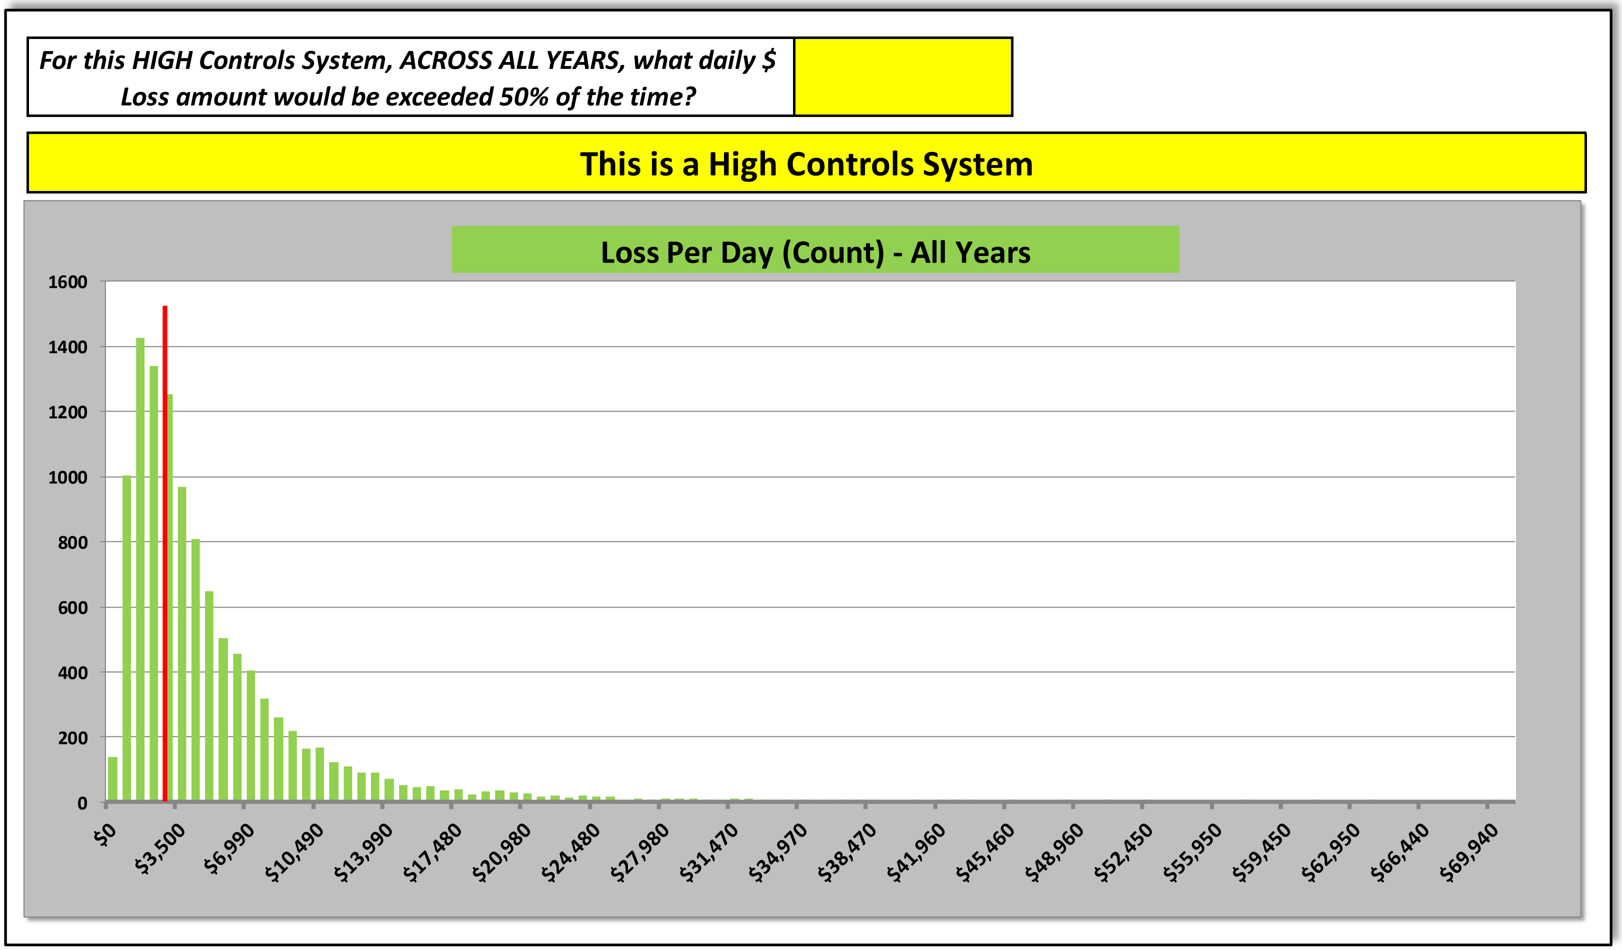

***For this HIGH Controls System, ACROSS ALL YEARS, what daily $ Loss amount would be exceeded 50% of the time?***

**This is a High Controls System**

**Figure 146 – Game 4B: Recovered Belief for Participant #52, Question #1**

***"For this HIGH Controls System, ACROSS ALL YEARS, what daily $ Loss amount would be exceeded 50% of the time?***

**Actual Answer: $3,026 (yellow bar)**

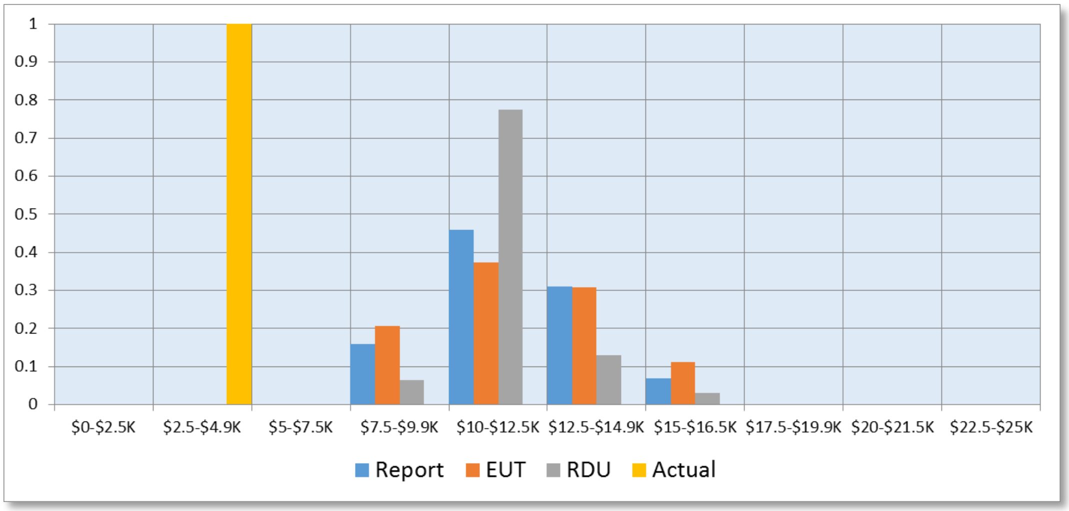

**EUT: r: .99 RDU: Power Utility r= -3.98, gamma = 3.9**

Here we see that the Participant missed the actual answer by a substantial margin and in fact did not report any weighting whatsoever to the true answer. The recovered beliefs based on EUT estimates are similar to those of the raw reports. However, assuming RDU probability weighting increases the recovered probabilities substantially, although retaining the overall shape of the reported values. Looking across all reporting Participants, we also note that most participants got the answer wrong, and many by a substantial margin of error. Many of the responses are in fact completely out of the range of plausible responses and we might question whether participants understood either or both the scoring device or the nature of the questions being asked. We report similar results for the other two question results profiled in detail for this experiment and discuss these results in the Conclusions section below:

**Figure 147 - Game 4B: Recovered Belief for All Participants, Question #1**

*"For this HIGH Controls System, ACROSS ALL YEARS,*
*what* ***daily $ Loss amount*** *would be exceeded 50% of the time?*

**Actual Answer: $3,026 (yellow bar)**

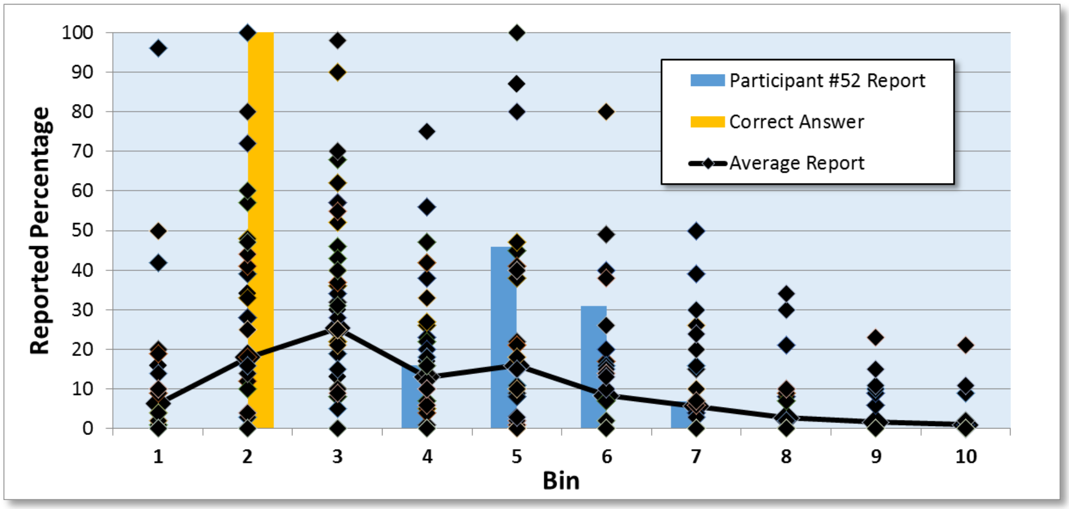

Figure 148 - Game 4B Sample Question #2

**Question #2:**

*For this MEDIUM Controls System, ACROSS ALL YEARS, what **daily $ Loss amount** would be exceeded 5% of the time?*

**Actual Answer:** $36,222 (red line)

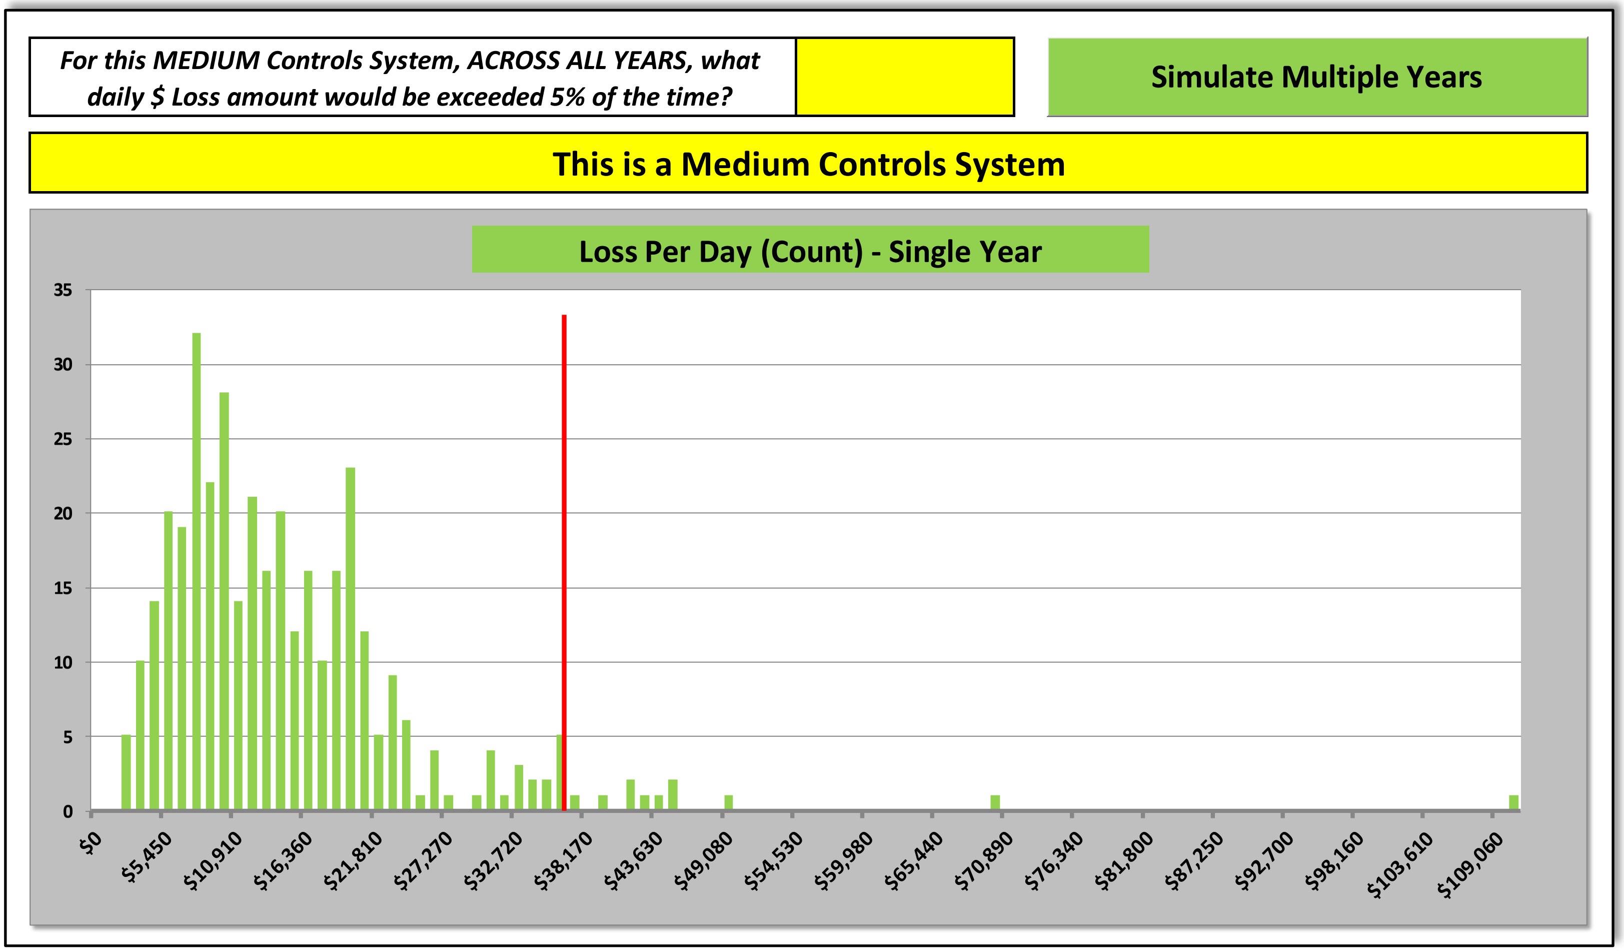

*For this MEDIUM Controls System, ACROSS ALL YEARS, what daily $ Loss amount would be exceeded 5% of the time?*

**Figure 149 – Game 4B: Recovered Belief for Participant #52, Question #2**

***For this MEDIUM Controls System, ACROSS ALL YEARS,***
***what daily $ Loss amount would be exceeded 5% of the time?***
**Actual Answer: $36,222 (yellow bar)**

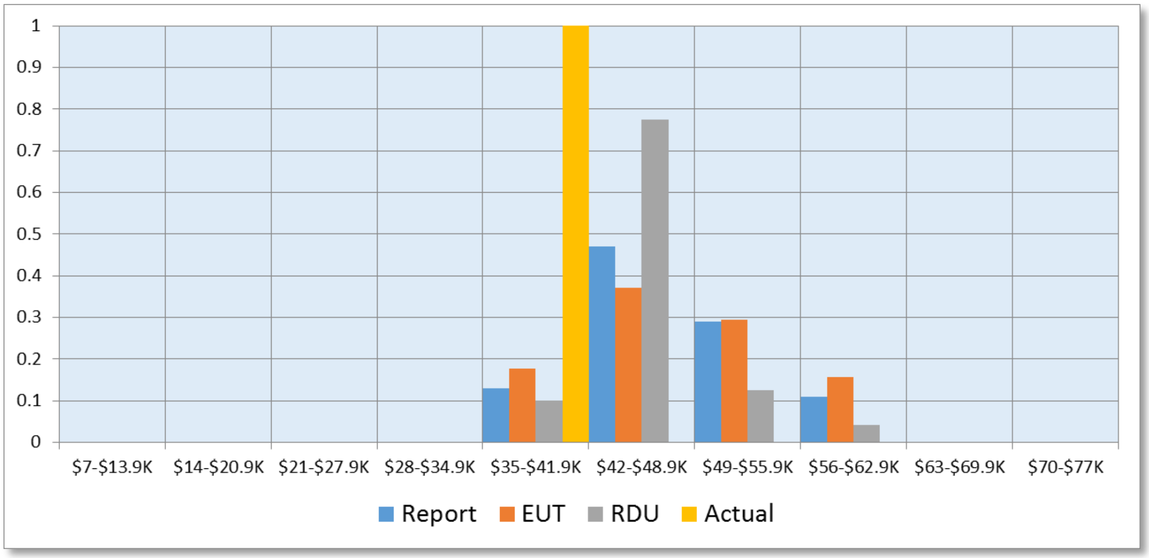

**EUT: r: .99   RDU: Power Utility r= -3.98, gamma = 3.9**

Here we see that the Participant reported much closer to the actual answer. The recovered beliefs based on EUT estimates are again similar to those of the raw reports. One again, the RDU probability weighting increases the recovered probabilities substantially, significantly modifying the overall shape of the reported values. Looking across all reporting Participants, we again note that most participants got the answer wrong, and even the pooled average maximal response was incorrect, again with a substantial margin of error within and across responses:

**Figure 150 - Game 4B: Recovered Belief for All Participants, Question #2**

***For this MEDIUM Controls System, ACROSS ALL YEARS,***
***what daily $ Loss amount would be exceeded 5% of the time?***

**Actual Answer: $36,222 (yellow bar)**

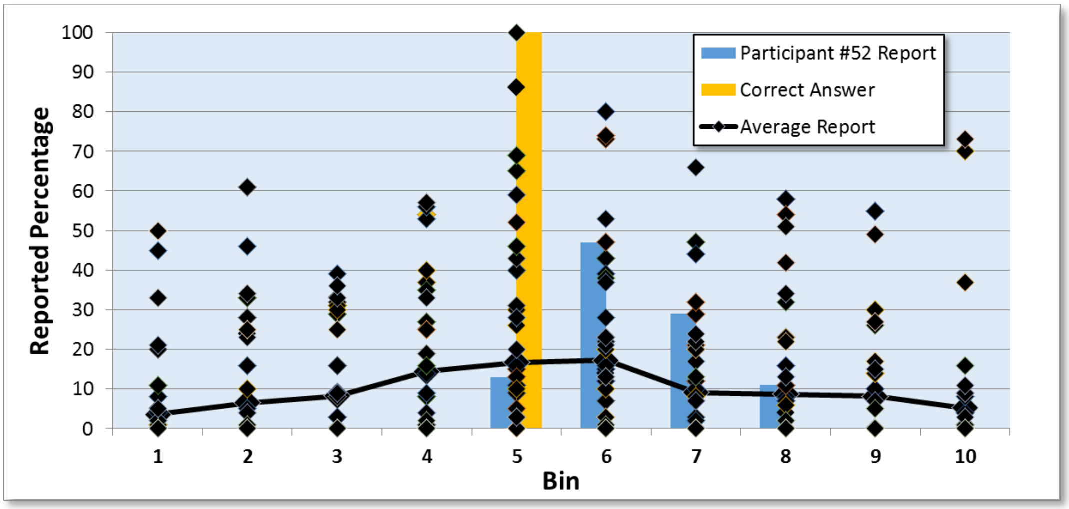

Figure 151 - Game 4B Sample Question #3

**Question #3:**

*For this HIGH Controls System, FOR THIS YEAR, what* ***percentage of daily loss amounts*** *exceed $5,755?*

**Actual Answer:** 25% (red line = $5,755)

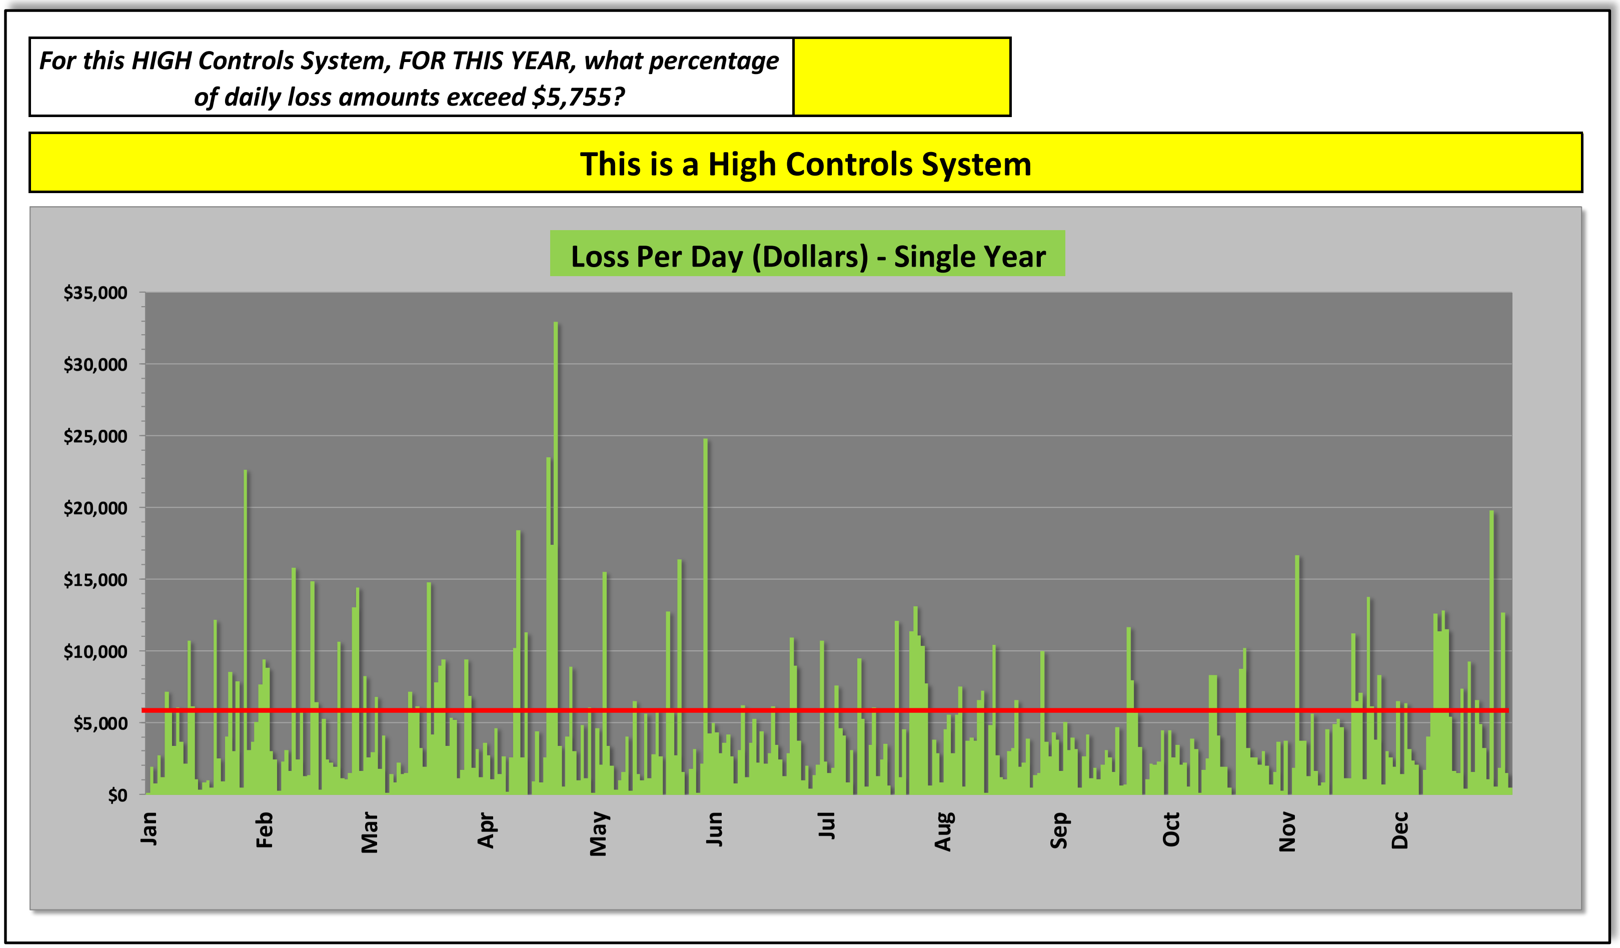

*For this HIGH Controls System, FOR THIS YEAR, what percentage of daily loss amounts exceed $5,755?*

This is a High Controls System

**Figure 152 – Game 4B: Recovered Belief for Participant #52, Question #3**

***For this HIGH Controls System, FOR THIS YEAR,***
***what <u>percentage of daily loss amounts</u> exceed $5,755?***

**Actual Answer: 25% (yellow bar)**

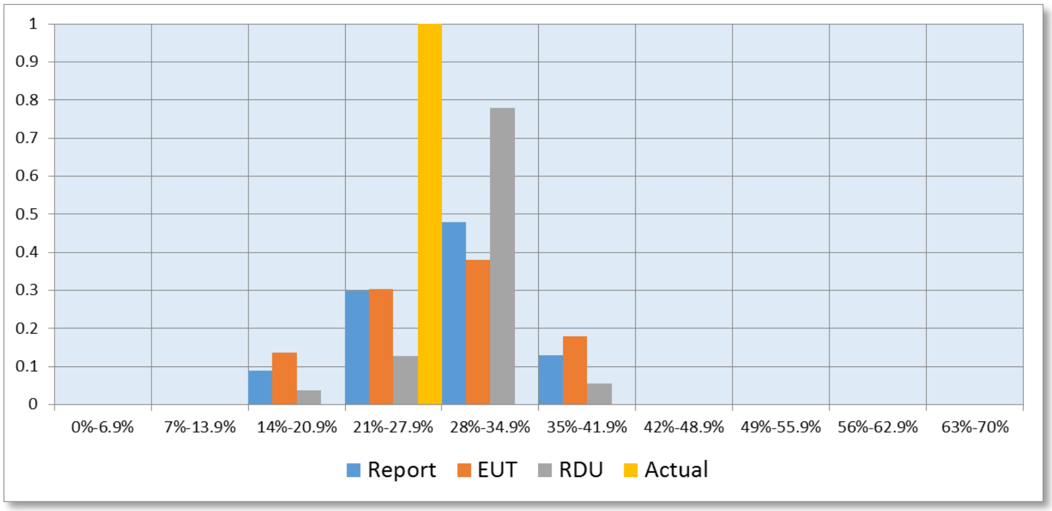

**EUT: r: .99   RDU: Power Utility r= -3.98, gamma = 3.9**

Here we again see that the Participant reported close to the actual answer, although as in each of the examples reported here, she was unable to select the correct answer. Once again, the RDU probability weighting increases the recovered probabilities substantially, significantly modifying the overall shape of the reported values. Looking across all reporting Participants, we note that most participants got the answer wrong, although pooled responses were right on average, but again with a substantial margin of error:

**Figure 153 - Game 4B: Recovered Belief for All Participants, Question #2**

***For this HIGH Controls System, FOR THIS YEAR,***
***what <u>percentage of daily loss amounts</u> exceed $5,755?***

**Actual Answer: 25% (yellow bar)**

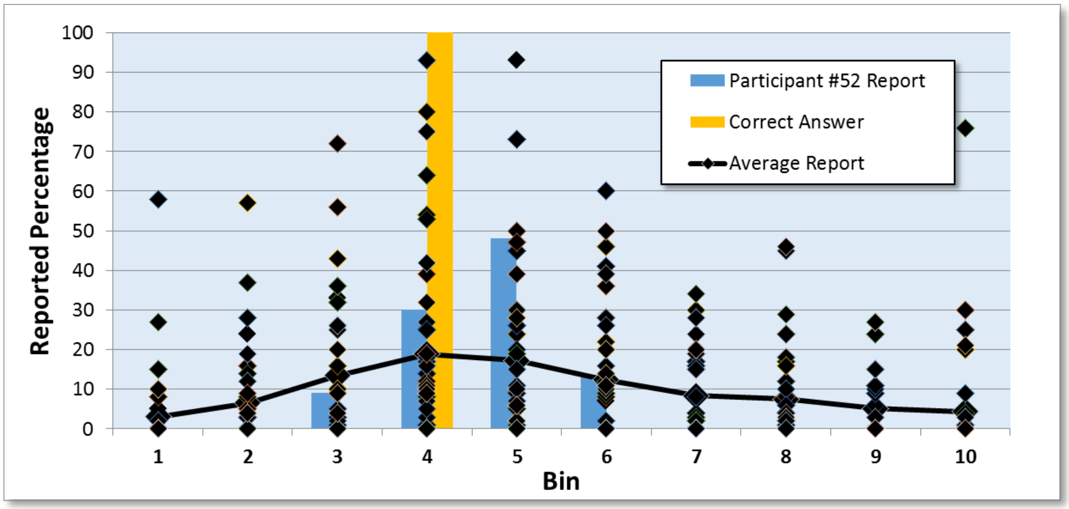

## Game 5 Findings

## The Effect of Ambiguity on Self Precaution and Cyber Insurance Choice

### Analysis

Following Bajtelsmit, Coates et al (Bajtelsmit, Coats et al. 2015) we conduct an experiment involving choices between self-precaution and insurance over security risks. In this experiment, we extend the literature on cyber-insurance by introducing a model which includes system performance and the level of security control as sources of ambiguity underlying the decision between precaution and insurance and by replicating an established experimental design to examine lab participant choices between protection and insurance. To my knowledge, this is the first study to explicitly incorporate the context of simulated security losses on precaution and insurance in this way. The experimental design allows us to test certain theoretical findings of the Bajtelsmit and Thistle (2008) model of the choice between precaution and insurance with high versus low probabilities of loss. We employ a design, parameters, and framing which allow us to contribute additional evidence to existing mixed results of experimental studies of the decision to insure against low-probability, high-severity losses.

In the absence of the ability to take precaution against accident, theory suggests that risk-averse individuals will fully insure when actuarially fair insurance is available. In situations where insurance is not fairly priced (and presumably where individuals can recognize this disparity) or where precaution is an alternative, the optimal choice depends on risk aversion, load and the cost of precaution. In this experiment, we confirm that, when the probability of loss is more ambiguous, the demand for insurance increases. However, ambiguous increases in the probability of loss may increase or decrease expenditure on precaution, depending on assumptions related to the cost and benefit of precautionary spending. We test these results empirically in a laboratory experiment in which participants make decisions about insurance and precaution under different ambiguity conditions represented by an IT system at risk of business losses attributed to security breach where the system is characterized as having relatively Low, High or Very High effectiveness security controls and where the outcomes are increasingly ambiguous the lower the level of control i.e. the apparent dispersion of results are greater and the probability of outcome is therefore harder to subjectively estimate.

The underlying theory is based on the standard model of accidents in the law and economics literature. In the absence of the ability to take precaution against accident, theory suggests that risk-averse expected utility maximizers will fully insure when actuarially fair insurance is available. In general, the assumption of risk aversion implies that individuals will be willing to pay some level of load or risk premium to avoid risk. Thus, when insurance is not fairly priced, the optimal choice depends on the level of risk aversion and the insurance loading factor. We assume that individuals are expected utility maximizers with increasing concave von Neumann Morgenstern utility *u*. Individuals have exogenous initial wealth *w* and face a potential loss d < w with probability $\pi$. Expenditure on precaution or care is denoted c ($c \geq 0$) and the risk of a loss is a deceasing, convex function of c. Individuals have either a high or low probability of loss, where $\pi_H(c) > \pi_L(c)$ for any expenditure on precaution. We assume $0 < \pi(c) < 1$, that is, it is possible to reduce the risk of loss but not predictably to zero[79]. We also assume precaution has a lower marginal impact on the probability of loss for low-probability risks (L) than for high probability risks (H), $0 > \pi'_L(c) > \pi'_H(c)$. We assume each person knows whether they face high or low risk and understands how the level of precaution 'generally' affects the probability of loss[80]. An insurance policy consists of a premium, $p_i$, paid whether or not loss occurs, and an indemnity, $q_i$, paid in the event that the loss occurs. In the complete coverage case, the first best levels of precaution are $c_i^* = \text{argmin } c_i + \pi_i(c_i)d$, i = H, L[81].

If insurance is not available, then expected utility is

$$U_i(c_i) = (1 - \pi_i(c_i))u(w - c_i) + \pi_i(c_i)u(w - c_i - d) \qquad (9.52)$$

The individual chooses the level of precaution, $c_i^0$, that maximizes expected utility. Because the individual is risk averse, she is willing to pay some amount $P_{iU}$ to avoid the risk of loss. The results in Bajtelsmit and Thistle (2008) imply that the willingness to pay to avoid the risk is given by $u(w - P_{iU}) = U_i(c_i^0)$. 'Willingness to pay' can be written as $P_{iU} = c_i^0 + \pi_i(c_i^0)d + \rho_{iU}$, where $\rho_{iU}$ is a risk premium.

Now assume that insurance is available, that insurers can determine risk type ex ante, and that the expenditure on precaution is observable. In general, the insurance premium can be written as $p_i = \lambda\pi_i(c_i)q_i$, where $\lambda$ is the premium 'loading factor' (i.e. X times the 'actuarially fair' premium and the insurance premium is actuarially fair if $\lambda = 1$). The individual who buys the insurance policy ($p_i$, $q_i$) and spends c*i* on care has expected utility given by

$$U_i(p_i, q_i, c_i) = (1 - \pi_i(c_i))u(w - p_i - c_i) + \pi_i(c_i)u(w - p_i - c_i - d + q_i) \qquad (9.53)$$

for i = H, L.

Generally, individuals will not choose to insure if the cost of doing so is greater than the cost of precaution which meets the insurance negligence standard. The size of the insurance loading factor relative to expected loss and cost of precaution therefore makes a difference in the predicted decision between

[79] This assumption differs from Bajtelsmit, Coates et al who assume $0 \leq \pi(c) < 1$, i.e. risk of loss may be reduced to zero through either 'perfect' control or insurance against complete loss. Although convenient for their study, I consider this unrealistic in the context of this study since we are incorporating system simulation in which the precautionary controls are inherently stochastic and the available insurance product at writing is unlikely to cover total loss under any circumstances. In general, I hypothesize that this this adds appropriate 'ambiguity' to the decision making context which is expected to be reflected in the empirical results.

[80] In our experiment, participants can interactively simulate the system at risk to experience the effect of the precaution and insurance choices on net losses at the chosen level of system control and insurance (i.e. their stake amount, less any purchased controls andéor insurance premium, less a randomly selected daily loss).

[81] *argmin[f,x]* gives a position $x_{min}$ at which function *f* is minimized. https://reference.wolfram.com/language/ref/ArgMin.html

insurance and precaution. For example, for low frequency, low severity risks, expected loss may be so small that even a modest insurer profit and risk charge will tilt the scale toward taking care instead of buying insurance. In most analyses of liability, as in the analysis described above, the probability of an accident is a function of care or precaution and is deterministic. Now suppose that it is possible for the system to experience security breaches that, despite expenditure on care, can result in, sometime catastrophic, business losses. This is analogous, in our context, to organization personnel committing operational errors or omissions affecting security control efficacy. Further, as noted above, the inherent effectiveness of the controls is assumed to be stochastic such that the vector of control performance attributes cannot be absolutely determined *ex ante*. Following Bajtelsmit et al, we therefore model the case in which individuals know that there is a random chance of system 'error' - in this case, controls that do not perform as expected relative to an alternate level of control - but they do not know exactly how it will impact the probability of loss or the loss amount. The fact that the probability of an error is unknown introduces ambiguity into the precaution vs. insurance decision making process. Bajtelsmit et al show that if an individual decision maker is ambiguity neutral, then a utility function incorporating a measure of the probability of 'mistake' reflecting ambiguity of the system outcome is linear and if the individual is ambiguity averse then the function is concave. An ambiguity-averse individual (like a risk averse individual) is therefore willing to pay to eliminate the risk and they show that ambiguity aversion increases the willingness to pay to avoid the risk (i.e. to increase control levels) and that ambiguity aversion is shown to increase the demand for insurance (i.e. we expect that participants will be more willing to pay for insurance the lower the control level).

On the other hand, the effect of ambiguity aversion on the optimal level of *precaution* is theoretically indeterminate and depends on the fine detail of the chosen theoretical model. Snow (Snow 2011) shows that if individuals have unbiased beliefs ($E\{\pi(c,\tilde{m})\}$, where $\pi$ equals the objective loss probability, c is the loss and $\tilde{m}$ is the occurance of a 'mistake'), then the loss probability must be either *multiplicatively separable* ($\pi(c, \tilde{m}) = \alpha(c)\pi(\tilde{m})$) or *additively separable* ($\pi(c, \tilde{m}) = \pi(\tilde{m}) + \beta(c)$). Snow further shows that multiplicative separability, implying ambiguity aversion, increases the expenditure on care. Snow and Alary, Gollier and Treich (Alary, Gollier et al. 2010) on the other hand, show that additive separability decreases the expenditure on care. The effect of ambiguity aversion on the expenditure on care is therefore an empirical question. However, decreased willingness to pay for small reductions in risk seems at odds with an increased willingness to pay to avoid the risk and implies a discontinuity in behaviour between small risk reductions and risk elimination. This suggests that ambiguity will lead to lower expenditures on care.

## Results

### 1 - Non-Parametric Analyses

Panels A1, A2, and A3 report the three treatments for the experiment with their respective system probabilities of insurable loss exceeding a stated threshold ($17), and the respective insurance loading factors and premia:

**Table 43 – Game 5 Panel A1: Treatment #1 (Insurance Only)***

| Insurance<br>Load Factor | High Controls System<br>Prob of Daily Loss >$17 = 10%<br>Average Daily loss = $8.20 | Very High Controls System<br>Prob of Daily Loss >$17 = 2.3%<br>Average Daily loss = $4.14 |
|---|---|---|
| 1x<br>(= ‘actuarially fair’ premium for the simulated system’s 30 year loss profile) | $2.57 | $0.57 |
| 2x | $5.14 | $1.13 |
| 3x | $7.71 | $1.70 |
| 4x | $10.28 | $2.27 |
| 5x | $12.85 | $2.83 |
| 6x | $15.42 | $3.40 |

* No insurance-only scheme was proposed for a Low Controls system in this experiment since this would effectively cover losses occurring 30% of the time and was considered unrealistic in this context.

**Table 44 - Game 5 Panel A2: Treatment #2 (Precaution Only)**

| Control Cost to Upgrade to Higher system profile* | LOW<br>CONTROLS SYSTEM<br>Prob of Daily Loss >$17 = 30%<br>Average Daily loss = $15.11 | HIGH<br>CONTROLS SYSTEM<br>Prob of Daily Loss >$17 = 10%<br>Average Daily loss = $8.20 | VERY HIGH<br>CONTROLS SYSTEM<br>Prob of Daily Loss >$17 = 2.3%<br>Average Daily loss = $4.14 |
|---|---|---|---|
| **Upgrade from Low to:** | N/A | $6.91 | $10.97 |
| **Upgrade from High to:** | N/A | N/A | $4.06 |
| **Upgrade from Very High to:** | N/A | N/A | N/A |

* Upgrade costs represent the differential in the average daily loss between the systems e.g. upgrading from a Low Controls system to a High Controls system = Average Loss/day LOW – Average Loss/Day High = $15.11 - $8.20 = $6.91. These are therefore ‘actuarially fair’ costs to change the expected daily loss of a system. No ‘downgrades’ were permitted.

**Table 45 - Game 5 Panel A3: Treatment #3 (Insurance + Precaution)**

| Control Cost to Upgrade to Higher system profile* | LOW<br>CONTROLS SYSTEM<br>Prob of Daily Loss >$17 = 30%<br>Average Daily loss = $15.11 | HIGH<br>CONTROLS SYSTEM<br>Prob of Daily Loss >$17 = 10%<br>Average Daily loss = $8.20 | VERY HIGH<br>CONTROLS SYSTEM<br>Prob of Daily Loss >$17 = 2.3%<br>Average Daily loss = $4.14 |
|---|---|---|---|
| **Upgrade from Low to:** | N/A | $6.91 | $10.97 |
| **Upgrade from High to:** | N/A | N/A | $4.06 |
| **Upgrade from Very High to:** | N/A | N/A | N/A |
| | | | |
| **Insurance Premium** | $8.62 | $2.57 | $0.57 |

* Upgrade costs represent the differential in the average daily loss between the systems e.g. upgrading from a Low Controls system to a High Controls system = Average Loss/day LOW – Average Loss/Day High = $15.11 - $8.20 = $6.91. These are therefore ‘actuarially fair’ costs to change the expected daily loss of a system. No ‘downgrades’ were permitted.

Figure 154 indicates the percentage of insurance choices for the High vs. Very High systems overall and the percentage of insurance choices at each insurance premium for the ‘Insurance Only’ treatment:

**Figure 154 – Game 5: Insurance Only Treatment - Percentage of Participants Choosing Insurance at Each Premium Load Factor**

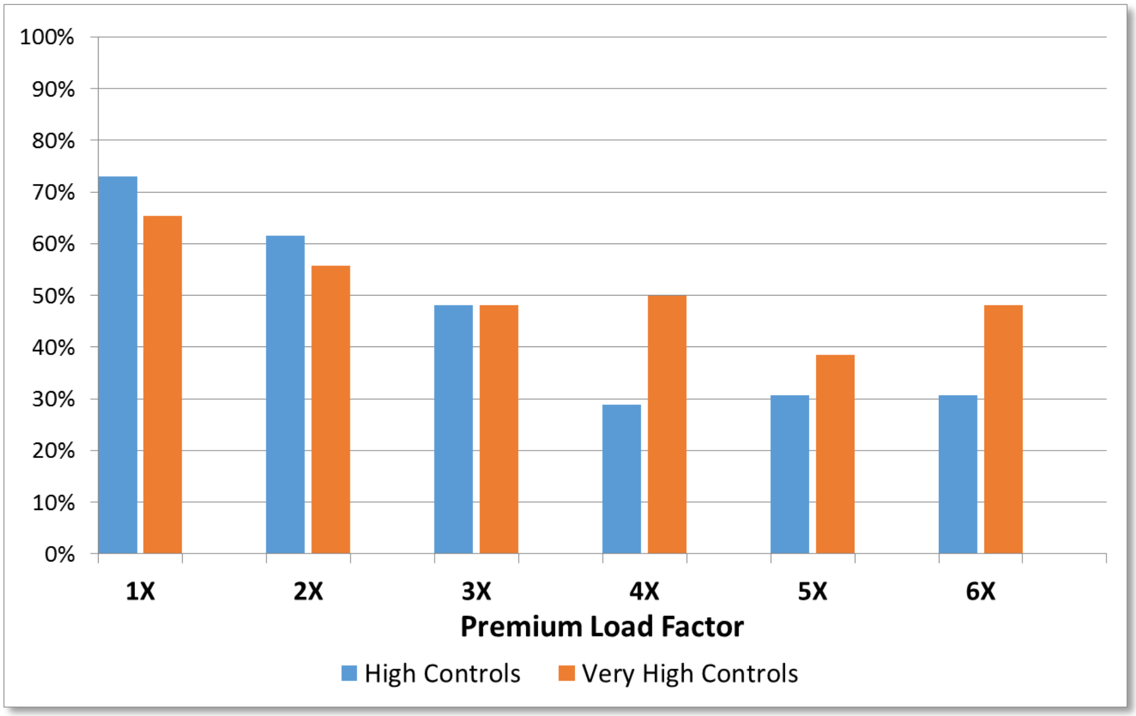

Panel B presents the data from the Insurance-Only treatment and McNemar tests for the differences in insurance purchases by individual participants under the Low Controls vs. High Controls cases:

**Table 46 – Game 5 Panel B: McNemar Test for Difference in Insurance Purchase based on System Control Level**

| Insurance Load | Controls | % Buying Insurance | McNemar | p value |
|---|---|---|---|---|
| 1 | High | 73% | 3.250 | 0.071 |
| 1 | Very High | 65% | | |
| 2 | High | 62% | 0.417 | 0.519 |
| 2 | Very High | 56% | | |
| 3 | High | 48% | 0.011 | 0.915 |
| 3 | Very High | 48% | | |
| 4 | High | 29% | 5.513 | 0.019 |
| 4 | Very High | 50% | | |
| 5 | High | 31% | 1.125 | 0.289 |
| 5 | Very High | 38% | | |
| 6 | High | 31% | 3.924 | 0.048 |
| 6 | Very High | 48% | | |

The results suggest that participants are more inclined to pay for insurance for the Lower control system as expected, but only until the loading for the premium on the lower controls system becomes exorbitant relative to the higher controls system, with the percentage of participants opting for insurance declining for both systems up to Load =  3. Thereafter, the percentage of participants opting for insurance is greater for the higher controls system than for the low controls system. As expected, participants do respond to the price of insurance in the predictable direction, generally purchasing less insurance as price (load) increases. While the percentage of choices opting for insurance declines as insurance load increases, participants continue to opt for insurance more than 30% of the time even at 6x load, where the expected value of the loss being avoided is much less than the cost of insurance. We conclude this represents substantial risk aversion on the part of participants opting for coverage at higher levels of load regardless of control levels.

Figure 155 indicates the percentage of choices opting for a system precaution upgrade at both the Low and High Controls base system levels for the ‘Precaution Only’ treatment

**Figure 155 – Game 5: Precaution Only Treatment - Percentage of Participants Choosing Precaution at Each System Level**

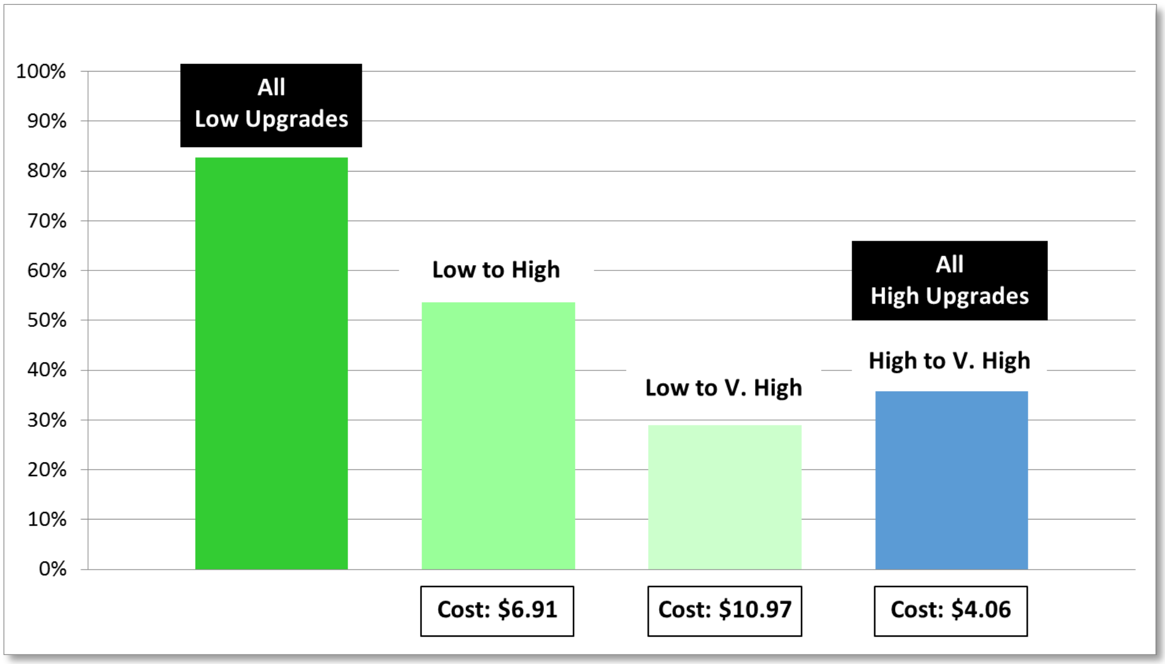

Here the percentage of choices opting for a precaution upgrade is greater for the Low Controls System than for the High Controls System, indicating that participants are more likely to take precaution as loss ambiguity increases. This is also shown in the marginal cases for Low upgrades where participants choose precaution nearly twice as often for the Low-to-High vs. the Low-to-Very High or the High-to-Very High upgrade (the cheapest upgrade). Since the cost of the upgrade is equal to the expected value of the upgrade, we conclude that participants appear to be essentially risk *averse* in losses i.e. they are generally willing to pay more than the expected value of the upgrade to avoid losses.

Figure 156 indicates the percentage of choices opting for either a system precaution upgrade or insurance at both the Low and High Controls base system levels for the 'Insurance + Precaution' treatment:

**Figure 156 - Game 5: Insurance + Precaution Treatment - Percentage of Participants Choosing Precaution or Insurance at Each System Level**

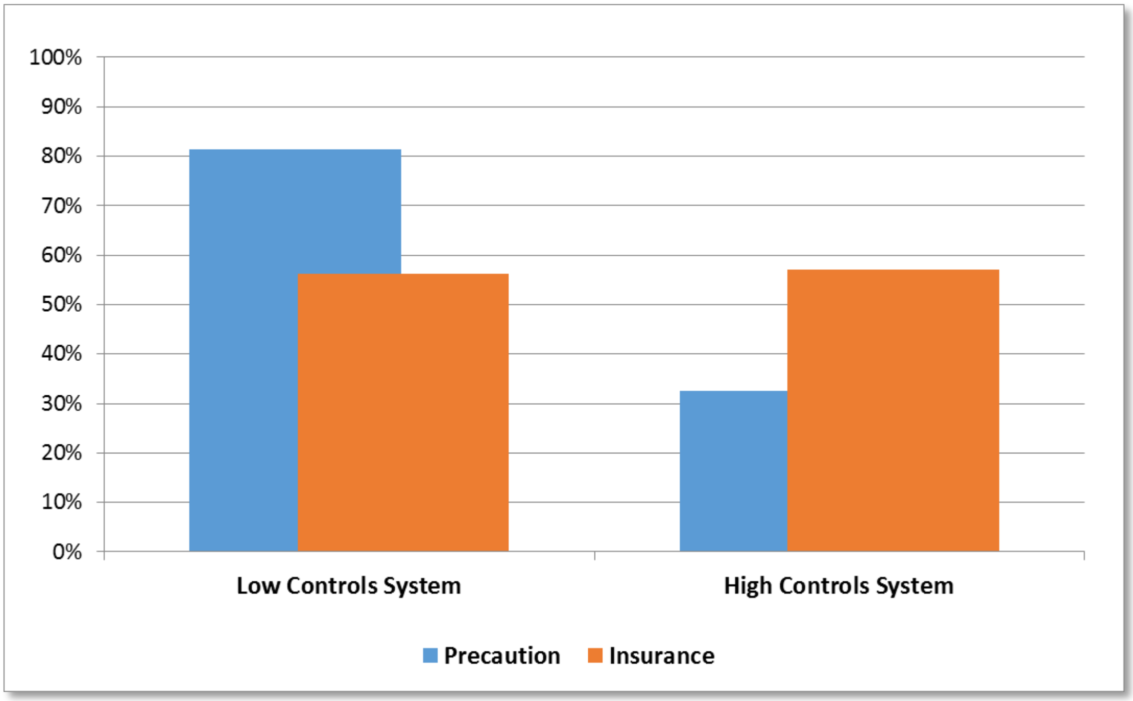

Here the percentage of choices opting for a precaution upgrade continues to be greater for the Low Controls System (81%) than for the High Controls System (32%) and is consistent with the Precaution only case, indicating that participants are more likely to take precaution as loss ambiguity increases even if insurance is available. The percentage of choices opting for insurance is nearly identical between the Low System (56%) vs. the High System (57%) although slightly higher than the insurance-only case, indicating that participants are not more likely to take insurance as loss ambiguity increases where precaution is available. This appears to contradict Bajtelsmit’s Hypothesis 2 but confirms our hypothesis that Participants’ level of risk mitigation (i.e. precaution) will be different across treatments with and without insurance. On the other hand, the presence of precaution does not appear to strongly affect the likelihood of insurance purchase. Figure 157 indicates the percentage of choices opting for both system precaution and insurance at both the Low and High Controls base system levels for the ‘Insurance + Precaution’ treatment:

**Figure 157 - Game 5: Insurance + Precaution Treatment - Percentage of Participants Choosing Precaution and Insurance at Each System Level**

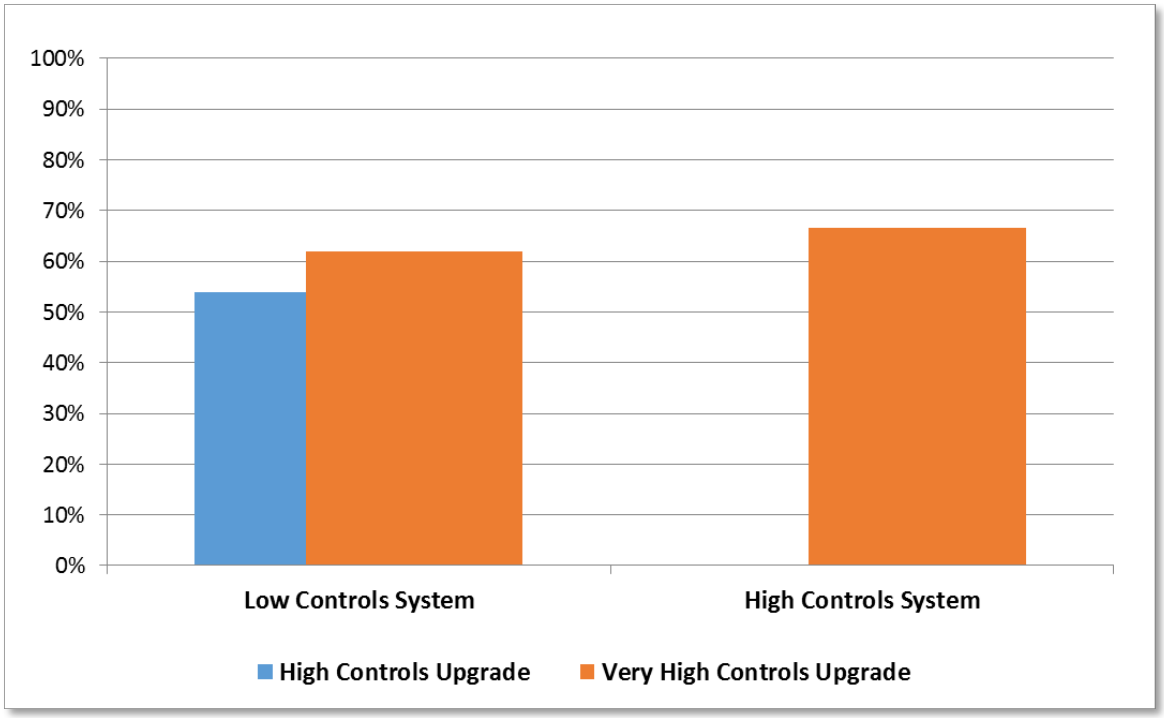

We conclude that precaution and insurance are likely seen as complements by Participants rather than substitutes for lowering the expected losses.

### 2 - Parametric Analyses

### 1 - Insurance-only Treatment

We perform a series logit regressions in which the dependent variable is the decision to purchase insurance (variable: ‘buyins’, 1 =Yes, 0 = No) in the treatments where it is available. All estimated coefficients represent the change in the odds of buying insurance with a one unit change in the respective dependent variable. Significant coefficients > 1 indicate that a one unit increase in the independent variable increases the odds of buying insurance, while significant coefficients < 1 indicate that a one unit increase in the independent variable decreases the odds of buying insurance. All regressions report a Chi-square statistic and an associated p-value which tests whether the model is significantly explanatory compared to an intercept only model. Models are considered statistically significant at the 95% level where the reported p-value is less than .05. The reported log-likelihood value is used to compare the goodness of fit between models.

For the Insurance-Only treatment, we include an independent categorical variable for the risk level of the (probloss: 0 = High Controls, 1 = Very High Controls) and the insurance premium (insprem), which varies by load and the risk level of the system. Panel C reports the results of this model:

**Table 47 – Game 5 Panel C1: Logit Model - Insurance Only, All participants**

**buyins = *f* (probloss, insprem)**

```
Logistic regression                               Number of obs   =        642
                                                  LR chi2(2)      =      35.99
                                                  Prob > chi2     =     0.0000
Log likelihood = -426.47933                       Pseudo R2       =     0.0405

------------------------------------------------------------------------------
      buyins | Odds Ratio   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
    probloss |   2.268318   .5563492     3.34   0.001     1.402588    3.668408
     insprem |   .8590246   .0233938    -5.58   0.000     .8143757    .9061214
------------------------------------------------------------------------------
```

This logistic regression model reports a Chi-square value of 35.99 indicating that the model is significant at the 95% confidence level. Across undifferentiated participants we see that Low System case (probloss = 1) increases the odds of choosing insurance by a factor of 2.3 and a one dollar increase in the insurance premium lowers the odds of choosing insurance by a factor of .86, or about -15%. Both coefficients are significant at the 95% level. This is in line with the expectation that, in the absence of other controls, higher risk drives a need for insurance while higher insurance costs lowers the demand for insurance. We assume that heterogeneity between the participants might account for the variance in the decision to purchase insurance and rerun the model using participant demographics as dummy variables:

**Table 48 - Game 5 Panel C2: Logit Model - Insurance Only, w. Demographics**

**buyins = *f* (probloss, insprem, demographics)**

```
Logistic regression                               Number of obs   =        642
                                                  LR chi2(11)     =      89.33
                                                  Prob > chi2     =     0.0000
Log likelihood = -399.80694                       Pseudo R2       =     0.1005

------------------------------------------------------------------------------
      buyins | Odds Ratio   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
    probloss |   2.442494   .6257121     3.49   0.000     1.478342    4.035453
     insprem |   .8474133   .0241621    -5.81   0.000     .8013554    .8961183
         age |   1.378752   .2389122     1.85   0.064     .9817227     1.93635
         sex |   1.844341   .3729744     3.03   0.002     1.240812    2.741427
     married |   1.158529   .2621331     0.65   0.515     .7435513    1.805107
         bus |   .6250157   .1176931    -2.50   0.013     .4321217    .9040154
        math |   1.217469   .2245995     1.07   0.286     .8480625    1.747786
   workyears |   .9177905   .0953831    -0.83   0.409     .7486534    1.125139
   privyears |   1.003724   .0690874     0.05   0.957     .8770514    1.148691
       certs |   .8750753   .0443052    -2.64   0.008     .7924082    .9663666
      income |   .9973649   .1487469    -0.02   0.986     .7445708    1.335987
------------------------------------------------------------------------------
```

Including demographic variables appears to add explanatory power to the model with increased Chi-square, significant at the 95% level, and a lower log-likelihood value. Across demographically differentiated participants we find similar direction and levels of effect on the odds of selecting insurance (buyins) for both the risk level of the system (probloss) and the price of insurance (insprem). In addition, increased age, and gender (female) also increase the odds of choosing insurance while a business degree and increasing levels of IT professional certification lower the odds of choosing insurance. All other demographic variables are not statistically significant at the 95% level and will be excluded form subsequent models.

Figure 158 indicates the resulting predicted probability of buying insurance across both low and high risk scenarios after taking into account participant demographics:

**Figure 158 – Game 5: Predicted probability of Buying Insurance, All Cases, Insurance only**

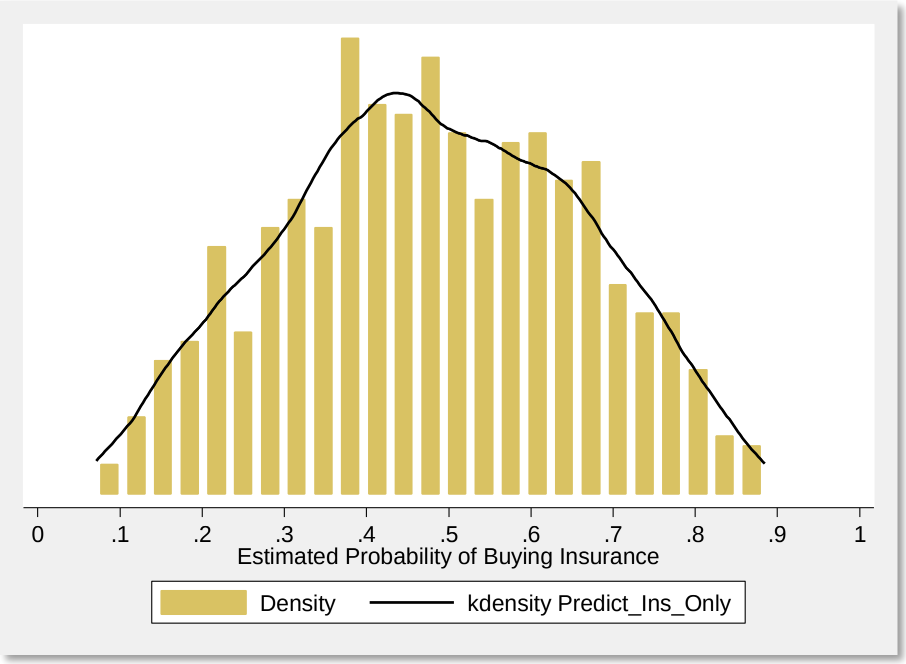

Figure 159 indicates the predicted probability of buying insurance in the low vs. the high control case. Here we see clearly that there is a slight increased central tendency to purchase insurance across participants in the high controls case after controlling for premium load and demographics, although the distribution of purchase probabilities is fairly similar between the Low and High Controls case:

**Figure 159 - Game 5: Kernel Densities of Predicted Probability of Buying Insurance - Insurance only, Low vs. High Controls**

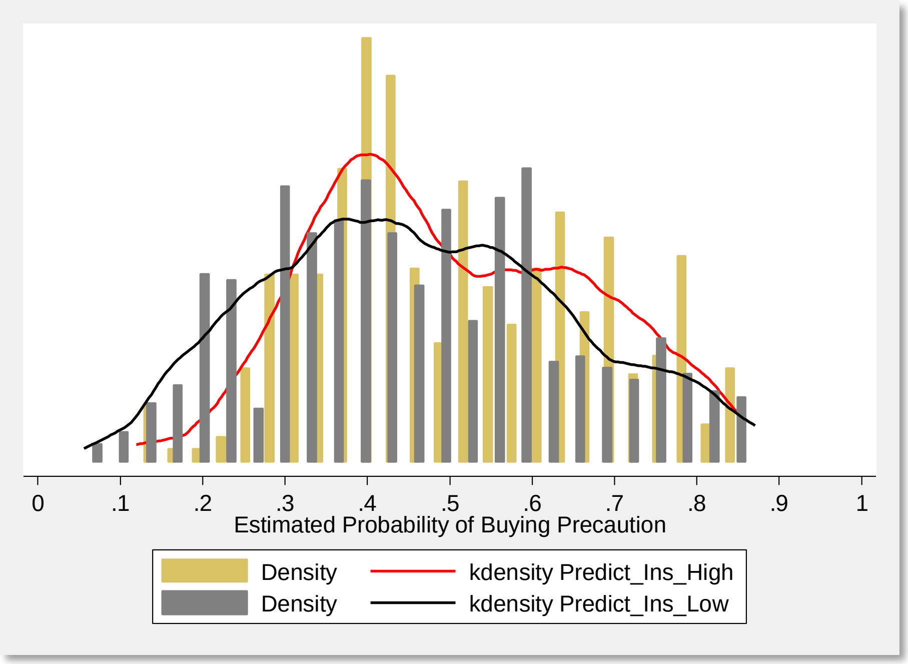

## 2 - Precaution Only Treatment:

For the Precaution-only treatment, we include an independent categorical variable for the presence of any system precaution upgrade (upgrade: 0 = no upgrade; 1 = some upgrade), and include all demographic variables. Panels C3 and C4 and C5 report the results of this model run for all cases and for the Low and High controls cases respectively. Figures 161, 162 and 163 then indicate the corresponding histograms and kernel densities for the probability of buying precautions under each case. Here we see clearly that the probability of buying precautions is bimodal based on the risk level of the system and confirm that there is a much greater tendency to purchase precautions in the Low Controls (more ambiguous ) case after controlling for demographics:

**Table 49 - Game 5 Panel C3: Logit Model – Precaution Only, All Cases, w. Demographics**

```
Logistic regression                               Number of obs   =        671
                                                  LR chi2(10)     =     175.77
                                                  Prob > chi2     =     0.0000
Log likelihood = -367.31451                       Pseudo R2       =     0.1931

------------------------------------------------------------------------------
     upgrade | Odds Ratio   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
vhighcontr~t |   1.377059   .0384525    11.46   0.000     1.303719    1.454525
         age |   .9936376   .1822209    -0.03   0.972     .6936304    1.423403
         sex |   1.258921   .2800605     1.04   0.301     .8140285    1.946962
     married |   1.789422   .4456738     2.34   0.019     1.098278    2.915501
         bus |   .8909425   .1799588    -0.57   0.568     .5996778    1.323675
        math |     1.0038   .1980015     0.02   0.985     .6819394    1.477571
   workyears |   1.035613   .1130438     0.32   0.749      .836148     1.28266
   privyears |   .7951459   .0591998    -3.08   0.002     .6871849    .9200682
       certs |   .9280657   .0491361    -1.41   0.159      .836589    1.029545
      income |   1.036535    .167644     0.22   0.824     .7549439    1.423157
------------------------------------------------------------------------------
```

**Table 50 - Game 5 Panel C4: Logit Model – Precaution Only Low Case, w. Demographics**

**upgradelow = f (age sex married bus math workyears privyears certs income)**

```
Logistic regression                               Number of obs   =        317
                                                  LR chi2(9)      =      18.70
                                                  Prob > chi2     =     0.0278
Log likelihood = -136.91118                       Pseudo R2       =     0.0639

------------------------------------------------------------------------------
  upgradelow | Odds Ratio   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         age |   1.087061    .342384     0.27   0.791     .5863501    2.015352
         sex |   .9919333   .3832403    -0.02   0.983     .4651734    2.115193
     married |   2.042215   .8746359     1.67   0.095     .8821664    4.727727
         bus |   .5889586   .1910038    -1.63   0.103      .311913     1.11208
        math |   .6164988   .2060777    -1.45   0.148     .3201844    1.187037
   workyears |   .8967541   .1723109    -0.57   0.571     .6153403    1.306867
   privyears |   .8314714    .107633    -1.43   0.154     .6451495    1.071604
       certs |   .9478136   .0846532    -0.60   0.548     .7956069    1.129139
      income |   1.935894   .5530141     2.31   0.021     1.105922    3.388744
------------------------------------------------------------------------------
```

**Table 51 - Game 5 Panel C5: Logit Model – Precaution Only High Case, w. Demographics**

**upgradehigh = f (age sex married bus math workyears privyears certs income)**

```
Logistic regression                               Number of obs   =        354
                                                  LR chi2(9)      =      21.89
                                                  Prob > chi2     =     0.0092
Log likelihood = -222.33727                       Pseudo R2       =     0.0469

------------------------------------------------------------------------------
 upgradehigh | Odds Ratio   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         age |   .9476658   .2239956    -0.23   0.820      .596294    1.506087
         sex |   1.315247   .3642411     0.99   0.322     .7643249    2.263273
     married |   1.728906   .5474998     1.73   0.084     .9294328    3.216066
         bus |   1.186526   .3110221     0.65   0.514     .7098279    1.983359
```

```
        math |   1.353779    .341551     1.20   0.230      .8256469    2.219736
   workyears |   1.120518   .1558329     0.82   0.413      .8531799    1.471626
   privyears |   .7731091   .0723806    -2.75   0.006      .6435006    .9288223
       certs |    .906085   .0630851    -1.42   0.157      .7905058    1.038563
      income |   .7146272   .1508554    -1.59   0.111      .4724912     1.08085
------------------------------------------------------------------------------
```

**Figure 160 – Game 5: Predicted probability of Buying Precaution, All Cases, Precaution only**

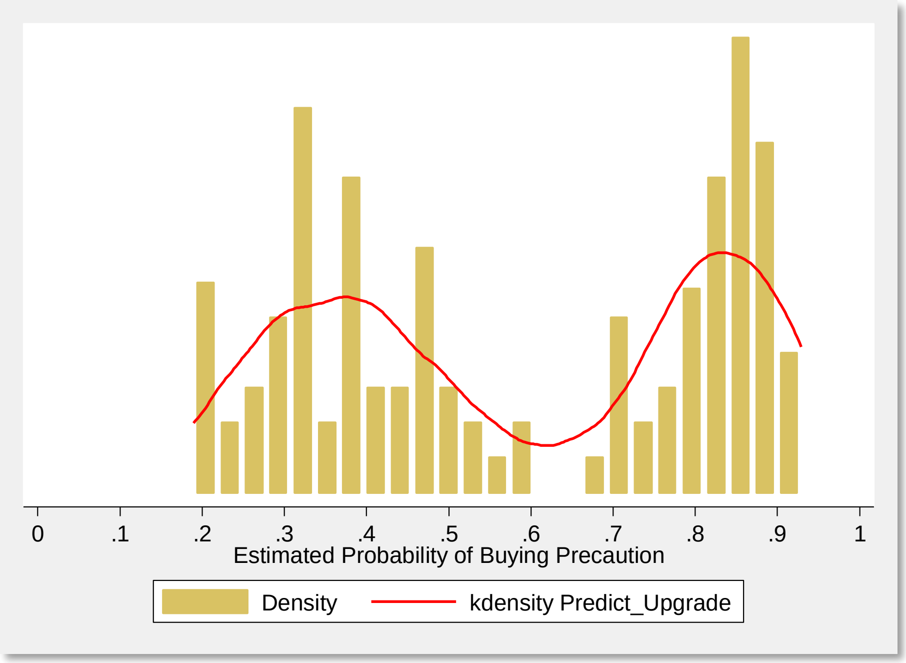

**Figure 161 - Game 5: Kernel Densities of Predicted Probability of Buying Precaution – Precaution only, Low vs. High Controls**

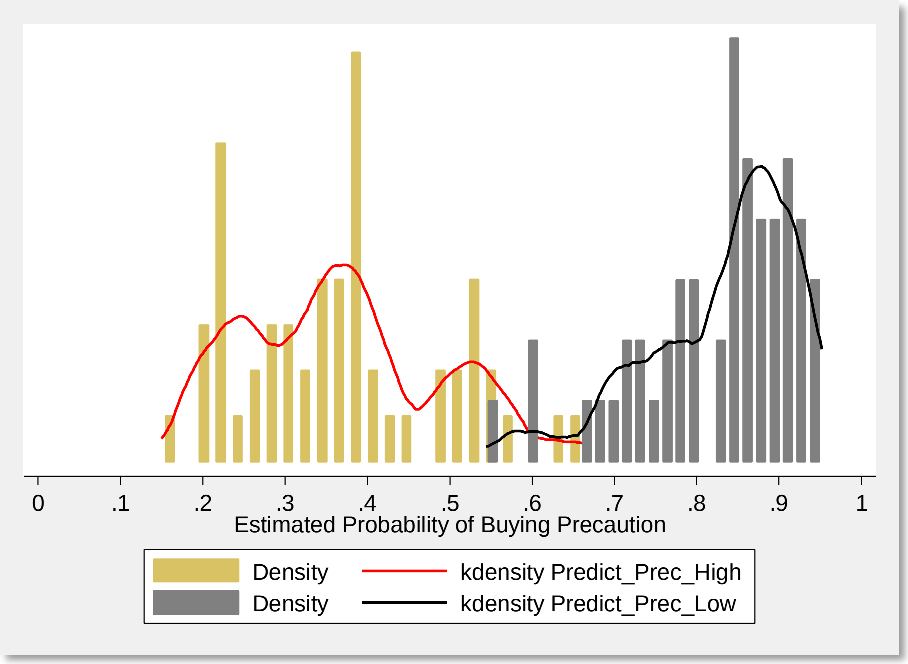

### 3 - Insurance + Precaution Treatment

For the Insurance+Precaution treatment, we include an independent categorical variable for the presence of any system precaution upgrade (upgrade: 0 = no upgrade, 1 = some upgrade), the insurance premium (insprem), and the cost of the Very High upgrade (vhigh_cost) which varies depending on whether you are upgrading from a Low system or a High system and include all demographic variables.

Panel D1 reports the results of this model where the dependent variable is the decision to buy insurance:

**Table 52 - Game 5 Panel D1: Logit Model – Insurance + Precaution, w. Demographics**

**buyins = *f* (upgrade, vhigh_cost, insprem, demographics)**

```
Logistic regression                                 Number of obs   =        624
                                                    LR chi2(12)     =      65.82
                                                    Prob > chi2     =     0.0000
Log likelihood = -393.94161                         Pseudo R2       =     0.0771

------------------------------------------------------------------------------
      buyins | Odds Ratio   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
     upgrade |   1.770722   .3647647     2.77   0.006      1.18251    2.651526
  vhigh_cost |   .9512445   .0270293    -1.76   0.079     .8997162    1.005724
     insprem |   1.022668   .0283346     0.81   0.419     .9686144    1.079738
         age |   1.711166    .297443     3.09   0.002     1.217116    2.405758
         sex |   2.821718   .6117119     4.79   0.000     1.844948    4.315619
     married |   1.130525   .2609771     0.53   0.595     .7190909    1.777365
```

```
         bus |   1.340957    .260539     1.51   0.131     .9162883    1.962446
        math |   1.005219   .1887956     0.03   0.978      .695652    1.452545
   workyears |   .8773843   .0970691    -1.18   0.237     .7063459    1.089839
   privyears |   .9864618   .0701631    -0.19   0.848     .8580997    1.134026
       certs |   1.072008   .0551602     1.35   0.177      .969169     1.18576
      income |   1.146994   .1768121     0.89   0.374      .847903    1.551587
------------------------------------------------------------------------------
```

Figures 163 and 164 indicate the resulting predicted probability of buying insurance when the option of buying precaution is also present as compared to the previously estimated probabilities of buying insurance when no precaution was available. We compare the overall probability regardless of Low or High Controls cases and then between Low and High Controls cases after taking into account participant demographics:

**Figure 162 – Comparison of Predicted Probability to Purchase Insurance: with and w/o Precaution, All Cases**

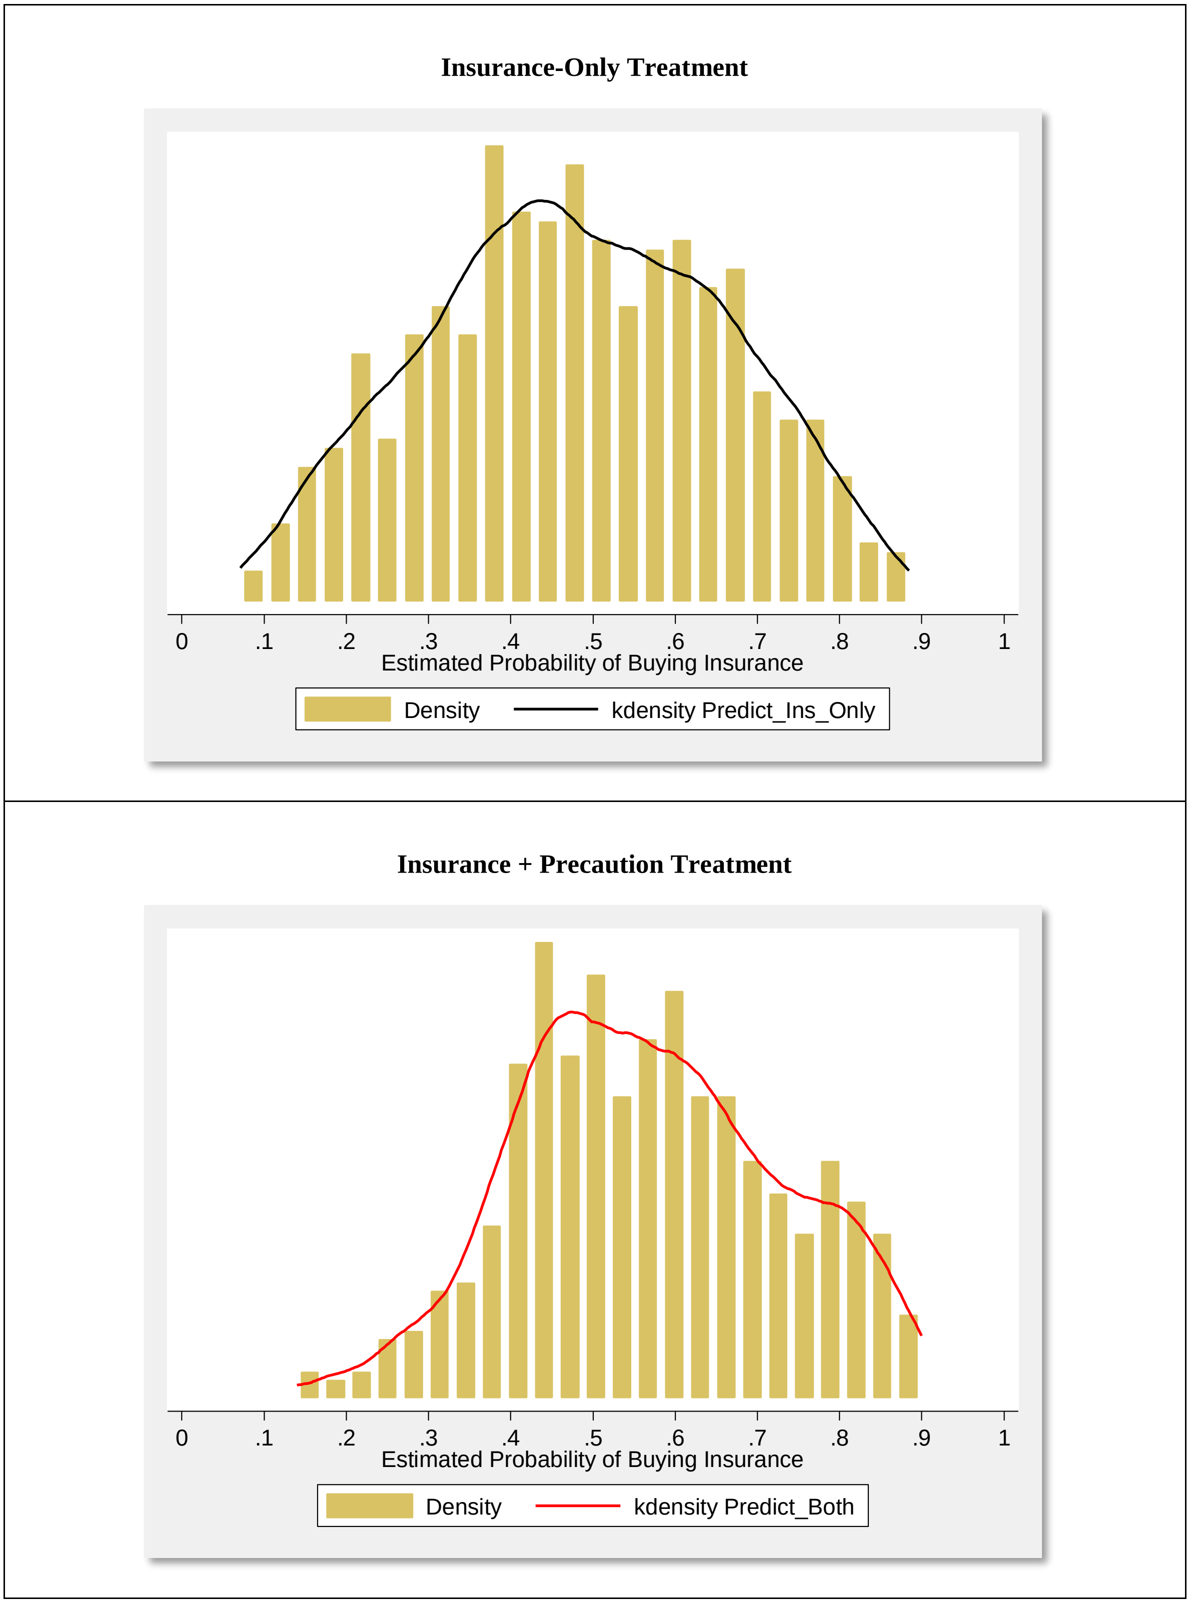

Insurance-Only Treatment

Insurance + Precaution Treatment

**Figure 163 - Comparison of Predicted Probability to Purchase Insurance: with and w/o Precaution, Low vs. High Controls Case**

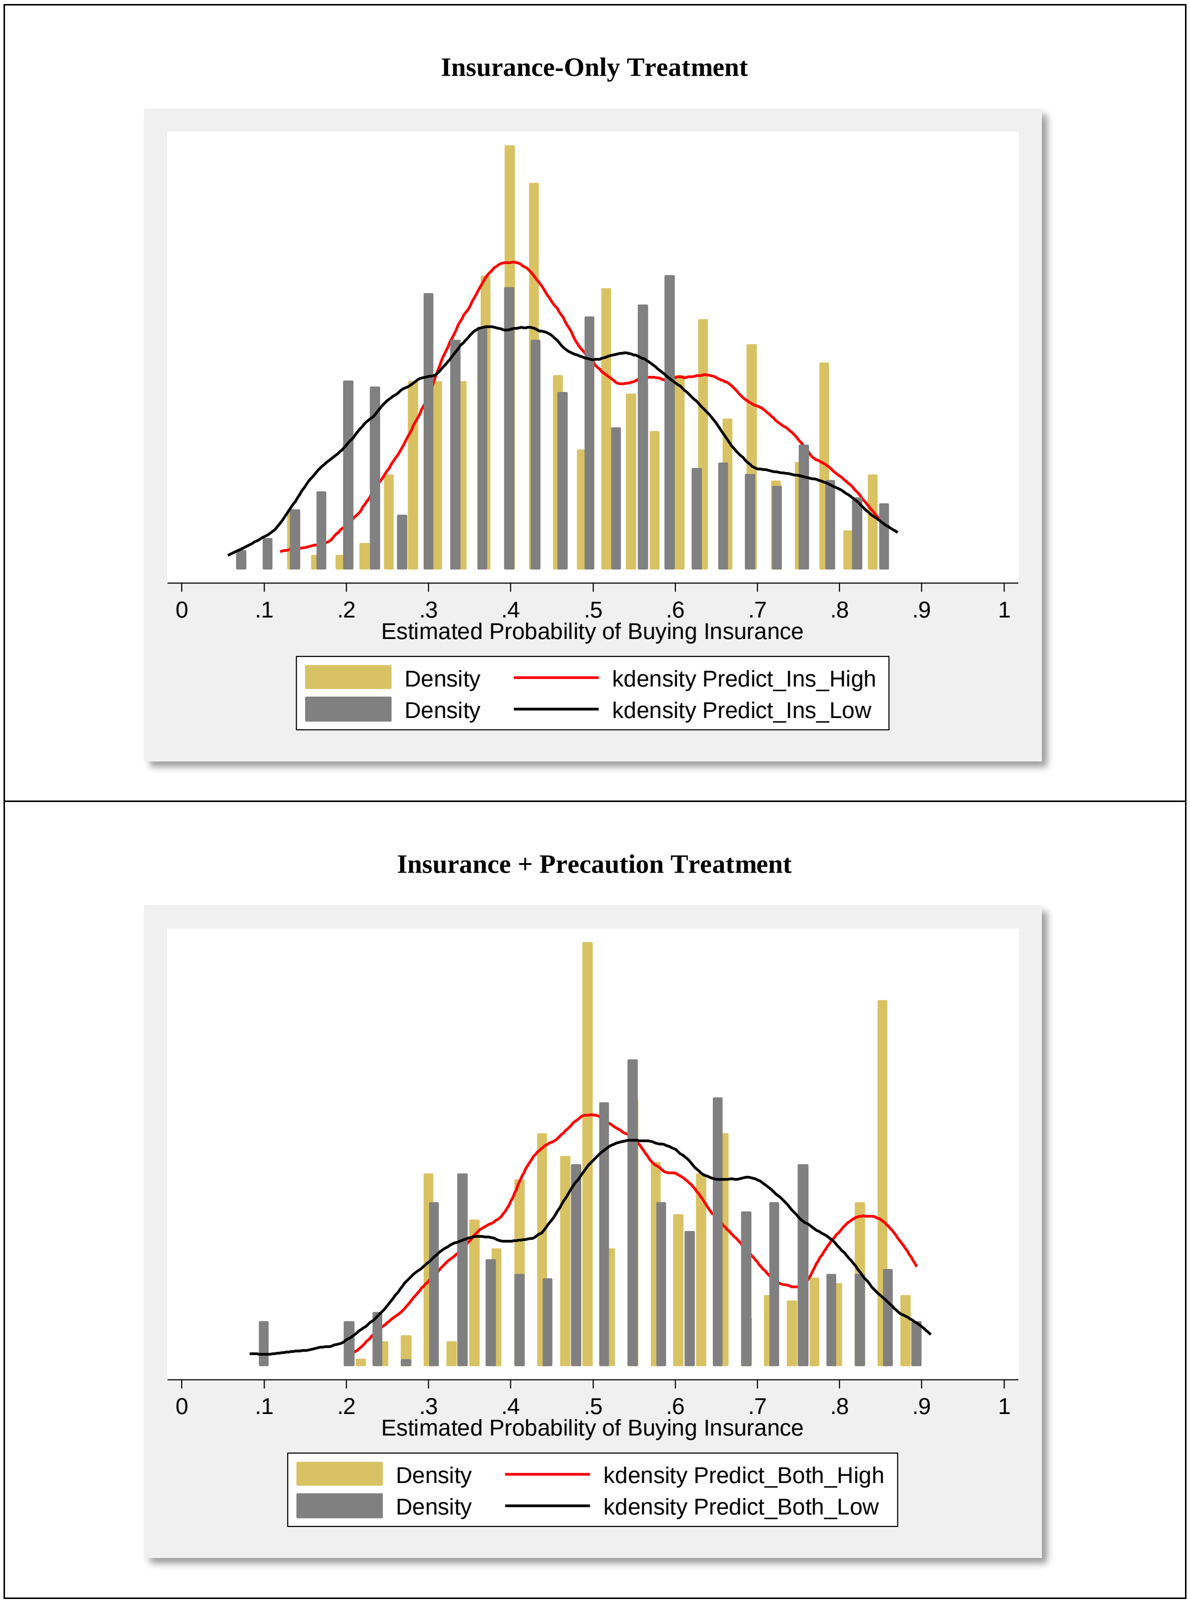

Insurance-Only Treatment

Insurance + Precaution Treatment

Here we see that, in a situation where both precaution and insurance options are available, the presence of a precaution upgrade increases the odds of choosing insurance, although there continues to be a wide dispersion of probabilities across participants. This result is reflective of the percentage of participants choosing both precaution and insurance in both the Low and High system cases in Figure 157. Variation in the cost of the very high upgrade (which applies to the decision to upgrade from either the Low or the High Controls system) does not appear to significantly affect the odds of choosing insurance, perhaps reflecting the fact that the upgrades costs are equal to the expected value of the loss prevented at each upgrade level, and indicating that, in the combined Precaution+Insurance case, participants may be risk neutral about upgrades overall. The cost of the insurance premium is also not significant. This may be because only the Load=1 level of premium was offered in this experiment - we might expect that increased premium load would have the same overall negative effect as was shown in the insurance only treatment.

Turning to the decision to buy precaution in the context of available insurance, Panel D2 reports the results of the model where the dependent variable is the decision to buy precaution:

```
Logistic regression                                 Number of obs   =        624
                                                    LR chi2(12)     =     192.92
                                                    Prob > chi2     =     0.0000
Log likelihood =  -323.2546                         Pseudo R2       =     0.2298

------------------------------------------------------------------------------
     upgrade | Odds Ratio   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      buyins |   1.726156   .3535383     2.67   0.008     1.155429    2.578794
  vhigh_cost |   1.387045   .0420044    10.80   0.000     1.307113    1.471864
     insprem |   1.013632   .0318835     0.43   0.667      .953029    1.078089
         age |   .6867515   .1358047    -1.90   0.057     .4660956    1.011869
         sex |    1.08657   .2645092     0.34   0.733     .6742879    1.750935
     married |   1.718685   .4576586     2.03   0.042     1.019847    2.896395
         bus |   .5968645   .1307994    -2.35   0.019     .3884542    .9170893
        math |    1.03445    .219483     0.16   0.873     .6825069    1.567877
   workyears |   1.450012   .1821674     2.96   0.003     1.133532    1.854853
   privyears |   .7185959   .0574013    -4.14   0.000     .6144563    .8403855
       certs |   .9154705   .0510632    -1.58   0.113     .8206651    1.021228
      income |   .8450392   .1467473    -0.97   0.332      .601256    1.187666
------------------------------------------------------------------------------
```

Figures 165 and 166 indicate the resulting predicted probability of buying precaution when the option of buying insurance is also present as compared to the previously estimated probabilities of buying precaution when no insurance was available. We compare the overall probability regardless of Low or High Controls cases and then between Low and High Controls cases after taking into account participant demographics:

**Figure 164 - Comparison of Predicted Probability to Purchase Precaution: with and w/o Insurance, All Cases**

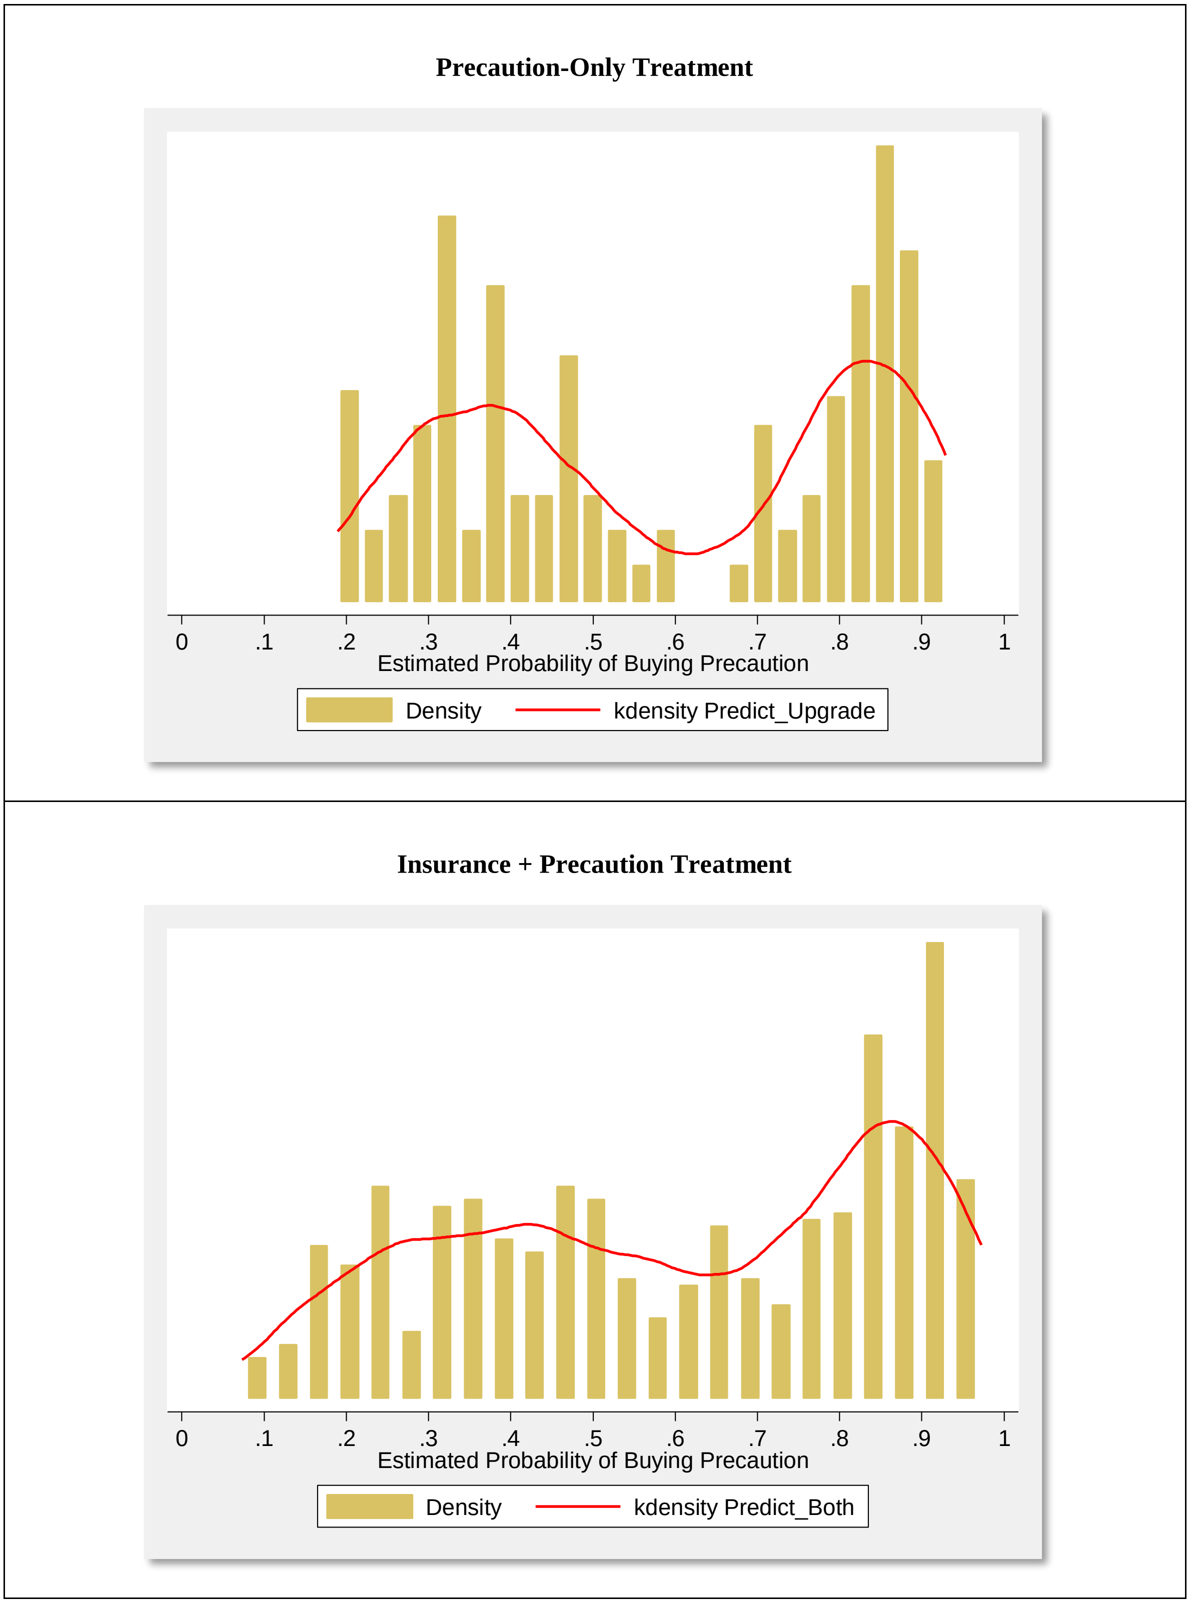

**Precaution-Only Treatment**

**Insurance + Precaution Treatment**

**Figure 165 - Comparison of Predicted Probability to Purchase Precaution with Insurance, Low vs. High Controls Case**

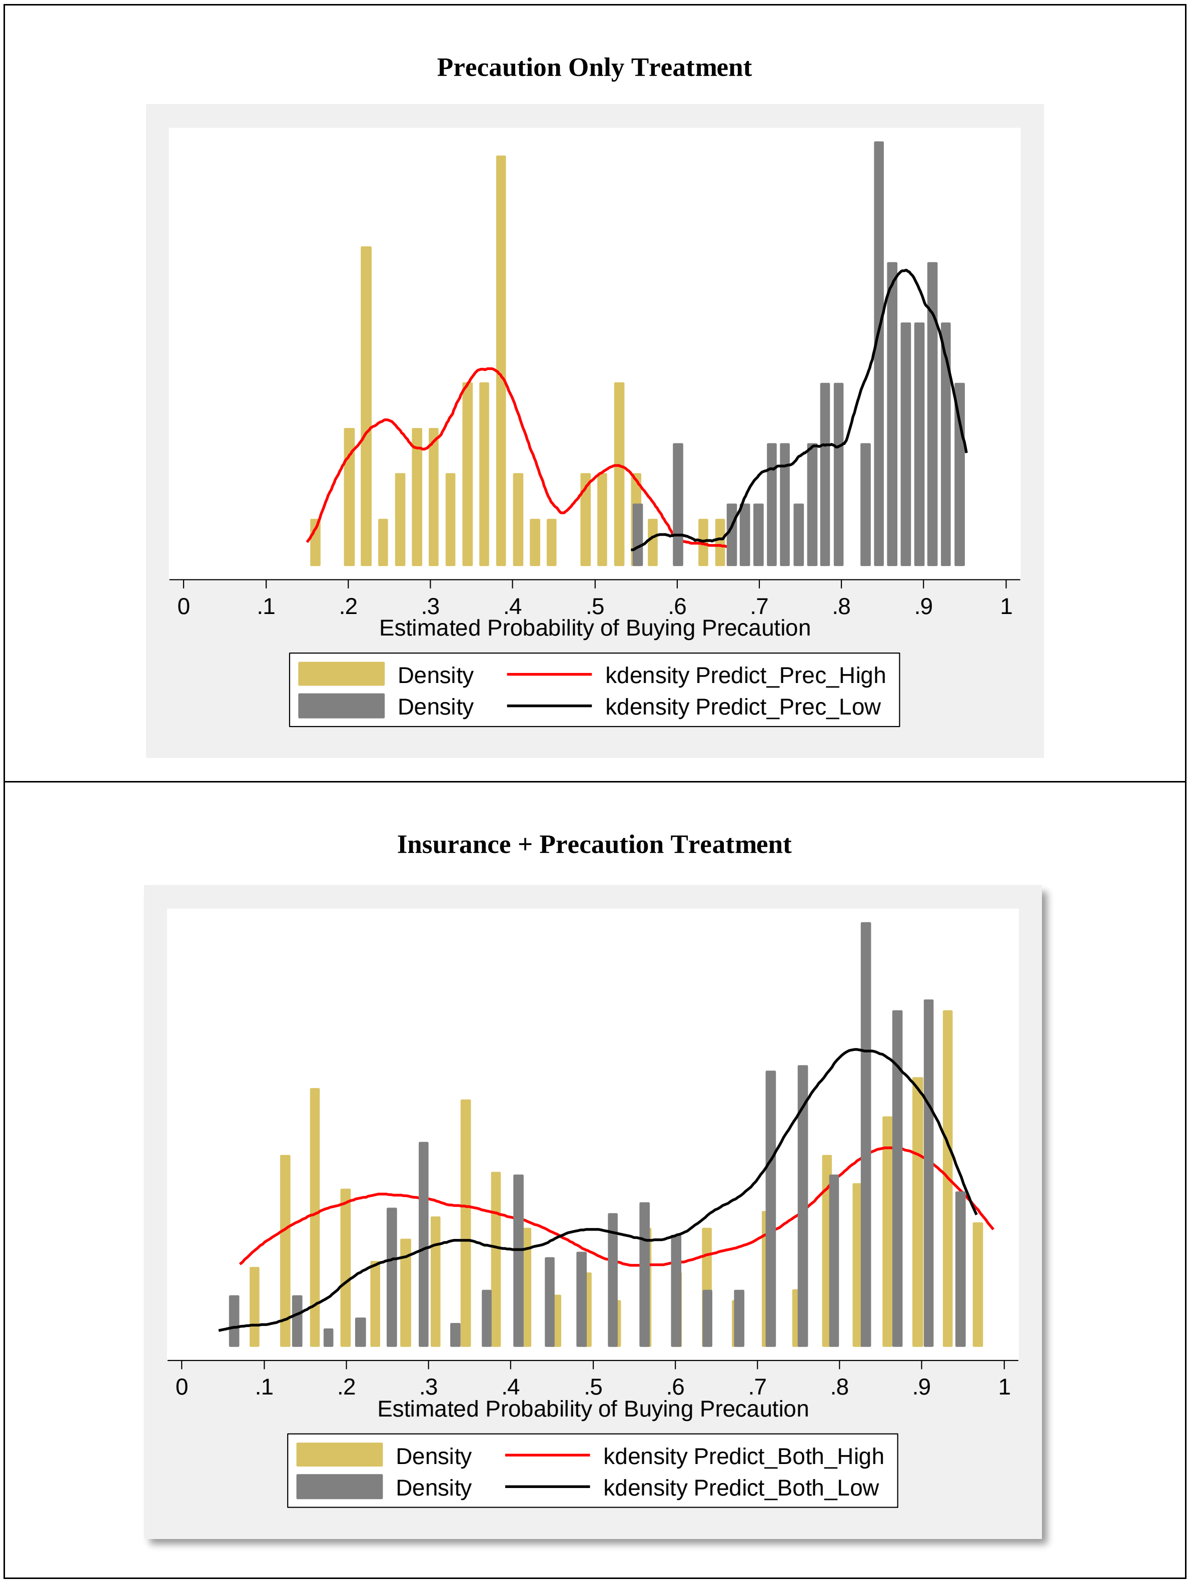

Precaution Only Treatment

Insurance + Precaution Treatment

# 10 – Conclusions and Reflection on Results

### Game 1 Conclusions

In Game 1 we test whether and to what degree security managers may integrate accumulated asset/wealth positions into marginal choice decisions over outcomes under assumptions of both EUT and CPT decision making, allowing for both gain and loss prospects.

In a model assuming EUT decision making over marginal outcomes only, we confirm relatively strong risk aversion ($r=0.225<1$) although the model is sensitive to the specification of 'errors' where risk aversion flips to moderately risk *seeking* ($r = 1.116, > 1$). If we assume that the argument of utility is only cumulative income, we infer that subjects remain risk-seeking although no more so than if cumulative income is ignored. These results suggest that the inclusion of cumulative income in the total prize amounts is a marginally better specification from an EUT perspective.

If we allow the data to decide the weighting between asset integration and marginal outcomes, participants are similarly risk seeking ($r = 1.16, >1$) but the estimated weighting term (.331) indicates that approximately 33% of the decision weight is attributed to the accumulated earnings as opposed to the value of the prizes presented on each choice round. After controlling for demographic effects, participants still appear to be slightly risk *seeking* (i.e. r_ = 1.11, > 1). While only age and sex are significant at the 95% level, age, sex (female), business degrees, and IT certifications all increase risk aversion, while a math degree, increased work years and income all lower risk aversion on average. After controlling for demographics, the average weighting for cumulative earnings drops to .122, however there is substantial demographic heterogeneity with respect to both the risk aversion parameter r and the weighting parameter $\omega$.

These results have potential implications for the composition of security teams based on the way individual practitioners frame business losses (regardless of whether losses are defined as pure gains, pure losses or mixed results) where some may focus on accumulated performance over time versus others who may pay more attention to the marginal decision. Potentially substantial (i.e. catastrophic) marginal loss prospects might be treated differently by the former decision maker without compensating controls for short term performance outcomes and management should therefore carefully consider how performance outcomes are framed in this context.

In a model assuming CPT decision making, we allow for separation of decision biases over gains versus losses including loss aversion and probability weighting which are both violations of EUT model assumptions. Using this model, we confirm most experimental results in which participants are risk averse over gains ($\alpha < 1$) and risk seeking over losses ($\beta > 1$). Furthermore, participants probability weight in the usual fashion: overweighting low probabilities and underweighting higher probabilities, although the estimated probability weights are only marginally different than 1 at the 95% level. In this model the loss aversion parameter lambda (λ) is estimated to be not statistically different than zero indicating no loss aversion. Allowing for demographic heterogeneity, average coefficients for α and β draw closer together, indicating risk aversion in gains ($\alpha < 1$) but no significant risk aversion in losses and no overweighting/ underweighting of low/high probabilities respectively. Loss aversion increases and is statistically greater than zero, however there is substantial diversity in individual estimates for all three parameters. These results therefore also have potential implications for the framing or decomposition of security outcomes where some individuals may treat decisions over outcomes framed as net gains (e.g. management focused on productivity improvements supported by security controls) differently than personnel who may frame daily performance in terms of pure losses (e.g. security operations personnel who tend to focus on losses as declines from 100% system availability etc.). Divergence in control decisions could occur if these perspectives are not aligned or decisions are permitted solely based on one frame or the other.

Following Harrison, we are also interested in whether there is diversity in the assumed model of decision making itself, so we allow for weighting between EUT and CPT models in the same specification. In the case of homogeneous agents we confirm that the EUT model accounts for approximately 2/3 of the choice observations. We also confirm that participants characterized by EUT tend to be risk averse but less so than when we assumed EUT characterized every observation and not just a percentage of the observations. We also confirm that EUT subjects tend to consider the marginal outcomes they are faced in conjunction with cumulative income when making decisions (omega ~= 1), in contrast to results that assume EUT characterizes every observation where subjects consider the utility from both the prizes and cumulative earnings about equally. This improves the models estimates for EUT by not forcing the EUT specification to account for observations that are better characterized by CPT. This result will be revisited in the context of participant reports over continuous time series and probability distributions of security losses where perspectives on catastrophic losses may be biased depending on the assumption of EUT versus a specification permitting probability weighting. The potential for distortion of low probability events such as catastrophic security losses is therefore an important consideration when practitioners are making decisions over low probability risky prospects.

For the PT part of the mixture model, the risk aversion parameters consistently suggest risk aversion over gains and losses: $\alpha = 0$ and β not statistically different than 0. There is also no evidence of significant probability weighting, with gamma (γ) now statistically closer to 1 than it was when PT was assumed to characterize the complete sample. Some shift would have been expected since $\gamma=1$ for EUT subjects, and PT was required to explain their behaviour in Table 18 Panels E and F; in Table 20 it is “free” to just characterize the PT subjects. There is also no consistent evidence of loss aversion and the estimated parameter would actually indicate loss *seeking* (lambda = -.44 < 1). Even though the estimate of lambda is less than 1, this parameter is not significantly different than zero and we can conclude find no evidence of loss aversion in this case, The imprecision of the estimate of λ may be due to a combination of the participant heterogeneity that is assumed away in this specification and the relatively smaller number of observations attributed to PT decision making in a loss frame.

Overall, we confirm Hypotheses 6, 9 and 10 that prior outcomes affect a manager's current decisions. We draw three major conclusions: First, expected utility theory accounts for a large fraction of the observations despite this experimental setting providing a reasonable for testing prospect theory and accounts for between 1/2 and 2/3 of the observations, depending on controls for demographic characteristics. Second, subjects behaving in accordance with expected utility theory appear to have utility functions defined over their cumulative income over the sequence of tasks, rather than defined over the prizes in each individual lottery choice. Finally, we identify demographic characteristics which differentiate risk aversion over gains vs. losses. Women are much more risk averse in gains and in losses than men, as are those with math degrees and higher incomes. Our results provide insight into the domain of applicability of each of the major choice models, rather than assuming one to be the sole, true model and the results are relevant for a security contexts which can be characterized by management from both a gains and a loss perspective.

### Game 2 Conclusions

Following Antoniou, we establish a "base camp" for exploring the fundamental characteristic of subjective Bayesian beliefs in the context of uncertainties regarding simulated IT system outcomes. We consider that the stochastic process that generates posterior probabilities (in this experiment a simulated IT system at risk of security breach) should be viewed as more uncertain than a process that generates risk when the probability of final outcomes is directly given. In other words, the posterior probability should be viewed as a subjective probability which may be seen by the decision-maker as subject to "uncertainty aversion" that exacerbates the effect of traditional "risk aversion." If this hypothesis is correct then the decision-maker will make choices that differ from those that would be made if she was neutral towards uncertainty. Consequently, the subjective posterior probability inferred from observed choices will differ depending on whether one allows for the possibility of uncertainty aversion. Previous analyses of subjective Bayesian decision-making, including our own here, have assumed that the subject is neutral towards the uncertainty that is involved in the use of an inferred posterior probability.

Quoting Antoniou:

> "Our approach is motivated by the same puzzle that has spurred the development of models towards uncertainty. It seems intuitive that behaviour towards an event that has a 50% chance of occurring with probability 0 and a 50% probability of occurring with probability 1 could reasonably differ from behaviour towards an event that has a 100% chance of occurring with probability 0.5. Yet, by some readings and axioms, these are not even two different states of the world, even though one can easily envisage distinct physical processes for each. Our version of this puzzle is that one could reasonably expect behaviour to differ when a decision-maker is credibly told that the probability of

> some event is 0.87 compared to when he is credibly given priors and sample realizations that imply, under Bayes Rule, a posterior probability of 0.87. We hypothesize that the reasons for differences in behaviour in these two puzzles are fundamentally the same."

We report SEU probability estimates with no demographic controls and find evidence of risk *seeking* behaviour, although not statistically different than risk neutrality. While participants do a reasonable job of estimating three posterior probabilities greater than 50%, there is consistent *underestimation* of the Bayesian posterior probabilities that increases with the objective probability and with significant dispersion of the estimates across individuals and with precision declining as the posterior gets larger. These results also imply a systematic *overestimation* of the true Bayesian posterior probability when the posterior is *less* than 0.5, becoming larger as the posterior decreases from 0.4 to 0.23 and 0.12 and is reflective of small probability overweighting. This result is contrary to Antonio's findings where the qualitative effect of risk aversion should follow from theory: the more risk averse the subject, the more likely he is to bet as if his subjective probability is 0.5, since this reduces the dispersion in final outcomes from all bets. In our case, we may fail to see any effect since, as noted above, the participants appear to be risk seeking to being with. There is no substantial difference in these results if we assume risk neutrality. After controlling for demographic effects, predicted values for the posterior probabilities improve although participants still appear to be slightly risk *seeking if we allow for risk aversion*, where *age*, *sex* (female), a *business degree*, *IT certifications,* and increased *income* all increase risk aversion, while a *math degree* and increased number of career *work years* lower risk aversion on average, results which are similar to Game 1 findings. In contrast to Antoniou's findings, Bayes Rule does *not* do better when we allow the data to determine the risk attitudes of subjects. In the risk neutral case the subjective estimate of the 0.6 posterior probability is similar, but the estimates for the 0.77 and 0.88 posterior probabilities, while not closer to the posterior on average, exhibit less dispersion around the Bayesian posterior than when we do not assume neutrality.

We also estimate the effect of the *strength* (alpha) vs. the *weight* (beta) of evidence on subjective probability estimates and find support for risk neutrality. The coefficients of both strength (alpha) and weight (beta) are, however, significantly different from 0 indicating that both strength and weight of evidence, while affecting subjective probability estimates are not significantly different from each other as predicted by the GT hypothesis, i.e., that α is generally > β. Allowing for demographic diversity in the estimates results in a wide range of predicted values for alpha and beta. Non-nested test results of model superiority are similar to those found by Antoniou and Harrison for this experiment: we see that the distribution of the ratio of GT to SEU log likelihoods is sharp-peaked compared to a Normal distribution fit to the same data and confirms here that the GT model better explains the data: strength and weight matter when estimating subjective probabilities.

**Game 3 Conclusions**

Virtually all information security settings in which there is risk involve what is referred to as endogenous risk whereby the decision maker can affect the risk outcome by making a control, in contrast to exogenous risk that is not controllable by the decision maker but is more commonly studied. Our design contributes a security practitioner context involving uncertain but controllable risks in which the structural decision factors can be untangled. Following Sen, our results provide evidence that participant risk attitudes and subjective probability estimates are different in these two environments. In the exogenous risk case for a Low Controls System, we see moderate risk aversion although the estimate is not statistically different than zero and the estimated subjective probability of loss is more than double the actual probability of loss for the system. For the High Controls System, we see marginal risk seeking although again the estimate is not statistically different than zero and the subjective probability of a loss is more than double the actual probability. After controlling for demographic variation, we see a wide range of risk attitudes from risk aversion to risk seeking. Subjective probability estimates also exhibit wide distribution around the mean, although the errors are more pronounced for the High Controls case. These are much greater dispersions than reported by Sen although similar (if reversed for the two cases) to those reported by Fiore.

In the Endogenous risk treatment we see relatively strong risk aversion ($r > 0$) and the estimate is statistically significant at the 95% level, similar to Fiore's findings but contrary to Sen's results which found no difference in risk aversion between the two treatments. On the other hand we confirm that the subjective probability for the High Controls (risky) System on the other hand is .12 which is very close to the actual probability (.1), reflecting an improved ability to estimate risk when risk is endogenous. Some of this may be due to the 'source dependence' of the risk: which is resulting in a 'focusing' effect by the participant in the context of this specific experiment which allows for participant control, in contrast to Treatment 1 and to all of Game 2 which similarly involves system simulation but where risk is exclusively exogenous. In practice, security risk is generally 'source dependent' and the implications are mixed for practitioners concerned about making better security choices. On the one hand, stronger risk aversion implies that participants may be over spending on controls (i.e. willing to spend more than required to obtain the desired lower subjective loss percentage). On the other hand, we see evidence that subjective probability estimates for 'low probability risks', which should be of particular concern to security professionals, may be more accurate than 'average' loss estimates. As noted in this thesis, the average happens all of the time in operations and likely should not be the particular concern of managers seeking to avoid risk, rather they should be paying attention to controlling tail risks by choosing controls that specifically lower higher end losses. The ability of managers to estimate the degree of prospective loss in the low probability cases versus the cost of control without bias is therefore a key insight supported by the findings of this thesis.

### Game 4 Conclusions

Following Harrison and Ulm, we demonstrate how to recover latent subjective beliefs if an individual is known to distort probabilities into decision weights using Rank Dependent Utility theory. Using a standard binary lottery choice experiment, across undifferentiated participants we see moderate risk aversion (r >0) although the estimate is not statistically different than zero. Allowing for demographic diversity, the results indicate significant diversity in the CRRA across participants, ranging from risk seeking to risk averse, where females, those with math degrees and increased work years all increase risk aversion. The diversity in individual risk attitudes once again illustrates the benefit of specifying decision models that account for this diversity before inferring general risk attitudes. We also estimate a Rank Dependent Utility specification to detect probability weighting since we are interested in how decision maker may distort subjective beliefs over the moments of security loss distributions. Indication of report distortion based on rank dependent probability weighting should signal to management that decision makers may not either be correctly perceiving loss probabilities or reporting subjective probabilities accurately. This has implications for security control decisions taken in the context of quantified operational risks (which in and of itself is a positive approach to decision making) where the perception of losses across a distribution of outcomes is an important factor in evaluating and choosing controls.

We confirm Harrison's results for the individual example undertaken: The effect on recovered beliefs assuming SEUT is minimal, but the probability distortions assuming RDU is shown to be significant, with large changes in the location and shape of subjective belief distributions. Our results allow the recovery of subjective belief distributions over security operations metrics that exhibit long tails which are increasingly important for security practitioners making control decisions that depend on accurate perception of tail risk. This experiment demonstrates that we should consider alternative models of risk preference and measure the effects of probability weighting when recovering subjective beliefs over the distributional characteristics of stochastic system attributes. Errors in reports versus subjective beliefs may be significant between decision makers and correcting for these perceptions would be important in situations involving accurate subjective perceptions of risk.

We also note that the data resulting from this experiment is surprisingly rather noisy: 16 of 60 participants failed to complete all 50 RDU lottery choice rounds, making some RDU model estimations not numerically feasible or valid. Furthermore, only three of the remaining individuals generate plausible rank dependency parameter estimates where the hypothesis of EUT could be rejected at the 10% level. This is not unexpected based on similar reports from Harrison's benchmark experiment for this Game where he notes, in some cases, '…none of the RDU models can be estimated for numerical reasons' and where certain extreme RDU Power function gamma results are rejected as implausible (gamma >5), in his case resulting in 65 of 71 available reports being usable. At only three of 44, our rejection rate is clearly much higher. I suspect this may be partially due to the nature of the binary lotteries used for the experiment which, while only being in the gain domain, are specifically designed to detect rank dependence by carefully varying either only the payouts or the probabilities, but rarely both between individual lottery pairs in order to induce decisions that possibly reflect probability weighting. While a rough calculation of expected value may still be possible for a typical participant, my observation of participants indicated that many found the lotteries challenging to complete compared to their experience in Game 1. I conclude that some participants may have been making arbitrary choices, resulting in difficult-to-fit RDU models for many individuals. For the purposes of confirming the effects of RDU weighting, we nonetheless recover subjective beliefs for one of these individuals under both EUT and RDU assumptions and successfully confirm our hypothesis that RDU distorts reports over continuous distributions of both time series and probability distributions for security losses.

Another issue identified with this Game is the potential reliability of the QSR method and the reported results. As noted, the dispersion of report results within any single question across participants is surprisingly large, where only one or two participants was able to identify (if not accurately report) the true answer across most questions. We also detect violations of transitivity (preference reversals between ranked answers) and therefore monotonicity violations in the reports of most (but not all) participants. This may be an indication of one or more confounds in this experiment:

1. **Participants simply did not care enough about providing ‘correct’ or ‘consistent’ responses i.e. responses were possibly arbitrary.** This is a possible confound for any Game but may have been a particular factor for this Game where, at this stage of the overall lab session (approximately 1.75 hours in) the Participants were possibly running out of energy for this level of Game concentration and complexity. As noted above, we did not vary the order of the Games to ensure that ‘simpler games’ were played first as a teaching method, so this remains a particular possible confound.
2. **Participants were seeking to report indifference over individual bin reports**: this is possible, but the QSR device has a built in way to permit this: Participants should simply allocate equal weightings to all 10 bins, or to a relevant sub-range of bins containing what they believe to be the actual. Examination of the data does not reveal any consistent equalization across bins for any particular question.
3. **Insufficient participant understanding of the nature of the question and therefore where an error is made in formulating a dominant, plausible answer.** This does not include those who understood the question but who may be truly uncertain and therefore sought to report a bi-modal answer for example, or for those who are simply mistaken about the answer, but we cannot be sure of this based solely on the reports themselves. The implication here is that the participants may be particularly ‘innumerate’ in this task, at least in estimating the moments of probability distributions and time series. If this is true, it represents an interesting finding for this research and is a possible area

for further research, as it directly relates to the ability to perceive probability distributions generally and tail risk circumstances specifically (Weston 2014).

4. **Insufficient understanding of the QSR scoring device mechanics or how to translate their answers into reports.** This is a possible confound of the QSR device noted by Harrison and others (Loomes, Starmer et al. 1991; Harrison and Rutström 2008). According to Artinger:

> The accurate elicitation of subjective expectations has become an important methodological concern in experimental economics. The Quadratic Scoring Rule (henceforth QSR) is currently the most widely used elicitation method. Part of the success of the rule is its theoretical incentive compatibility2 satisfying the principles set by Induced Value Theory, namely, monotonicity, salience and dominance Smith (1976). Of particular focus in this paper is the issue of salience: that subjects understand the payoff consequences of their actions given the rules of the game. Due to the mathematical formulation underlying the QSR, it is not a trivial matter for the experimental subjects to infer what a certain action means for their payoffs. However, should the subjects not understand the payoff consequences resulting from the formula, the comparative advantage of the mechanism no longer holds. When using the QSR, researchers face a trade-off between keeping the experiment as simple as possible vs. maintaining salience of the payoff mechanism…The possible loss of experimental control due to methodological complexity is not a new concern. Albeit theoretically incentive compatible, experimental subjects often find it difficult to correctly understand the incentives these mechanisms offer.3 With respect to the QSR, the old concern is illustrated anew in (Read 2005), pg. 273]: "I suspect that participants either more-or-less ignore the rule [QSR] or else get so caught up in understanding it that it becomes the focus of their activity". (Artinger, Exadaktylos et al. 2010)

The QSR device (and the associated presentation of evidence in the form of probability distributions) is arguably the most complex game interface in the entire experiment. My observation of the Participants indicated that they likely had the most difficulty with the QSR interface across all of the games, despite its being late in the lab session overall where we might expect participants to be more adept at manipulating the onscreen elements. Simplification of this Game and its components

### Game 5 Conclusions

Following Bajtelsmit, Coats et al (Bajtelsmit, Coats et al. 2015) we undertake a model of the decision between precaution and insurance under an ambiguous probability of loss and replicate an experimental design to test its predictions.

We generally confirm the results of Bajtelsmit across Insurance Only, Precaution Only and Insurance with Precaution treatments with noted exceptions. For the Insurance Only treatment, our results confirm Bajtelsmit's Hypothesis 5 and our Hypothesis that participants are more inclined to pay for insurance for the Lower control (more ambiguous) system as expected, but only until the loading for the insurance premium becomes exorbitant and absolutely much greater relative to premia for the higher controls system, with the percentage of participants opting for insurance declining for both systems up to a loading factor of 3x the actuarially fair premium. Thereafter, the percentage of participants opting for insurance is greater for the Higher controls (less ambiguous) system than for the Low controls system. As expected, participants do respond to the price of insurance in the predictable direction, generally purchasing less insurance as premium load increases. While the percentage of choices opting for insurance declines as insurance load increases, participants continue to opt for insurance more than 30% of the time even at 6x load, where the expected value of the loss being avoided is much less than the cost of insurance. We conclude this represents substantial risk aversion on the part of participants, opting for coverage at higher levels of load regardless of loading or control levels. Controlling for expected loss under insurance-only treatments, we find that participants neither underweight nor overweight low probability-high severity losses. Ambiguous increases in loss probability increase insurance uptake by more than objective increases in loss probability, suggesting evidence in favor of ambiguity aversion. Our results suggest that the tendency to over-insure against liability rather than meet a 'standard of care' through precaution may be partially explained, as suggested by our model, by sources of ambiguity surrounding liability losses. The results also lend further support to the Laury et al. (2009) findings that probability misperceptions is not an adequate explanation for observed under-insurance against catastrophe which is increasingly important in the context of information security breach.

For the Precaution Only treatment, we see clearly that the probability of buying precautions is bimodal based on the risk level of the system and confirm that there is a much greater tendency to purchase precautions in the Low Controls (more ambiguous ) case after controlling for demographics:  This is also shown in the marginal cases for Low upgrades where participants choose precaution nearly twice as often for the Low-to-High vs. the Low-to-Very High or the High-to-Very High upgrade (the cheapest upgrade). Since the cost of the upgrade is equal to the expected value of the upgrade, we conclude that participants appear to be essentially risk *averse* in losses i.e. they are generally willing to pay more than the expected value of an upgrade to avoid losses.

For the Insurance + Precaution treatment, we confirm the presence of a precaution upgrade option increases the odds of choosing insurance, although there continues to be a wide dispersion of probabilities to purchase insurance across participants. This result is reflective of the percentage of participants choosing both precaution and insurance in both the Low and High system cases. On the other hand, we do not confirm Bajtelsmit's Hypothesis 2: in our experiment the probability of opting for a precaution upgrade is greater in the absence of insurance than when insurance is available, and is also markedly different for the Low Controls System than for the High Controls System in both cases, indicating that participants are more likely to take precaution as loss ambiguity increases. We conclude that participants perceive that precautions and insurance are complementary at higher levels of control, possibly since the likelihood that the insurance will be needed is lower, even assuming zero load i.e. although the insurance is priced efficiently, its value is somehow greater than at lower system risk levels. This also reflects an emerging market perception for cyber insurance as complementary to precaution controls generally (Böhme 2005;

Baer and Parkinson 2007; Herath and Herath 2011; Johnson, Böhme et al. 2011) [82]. Variation in the cost of the very high upgrade (which applies to the decision to upgrade from either the Low or the High Controls system) does not appear to significantly affect the odds of choosing insurance, perhaps reflecting the fact that the upgrades costs are equal to the expected value of the loss prevented at each upgrade level, and indicating that, in the combined Precaution+Insurance case, participants may be risk neutral about upgrades overall. The cost of the insurance premium is also not significant. This may be because only the Load=1 level of premium was offered in this experiment - we might expect that increased premium load would have the same overall negative effect as was shown in the insurance only treatment.

As also noted by Bajtelsmit, there are important policy implications for cases in which security practitioners (and the firms they represent) may substitute cyber risk liability insurance in place of meeting a precautionary standard of care for security control. High transparency and consistency regarding compliance with a standard of care, when possible, may increase precaution and decrease the risk of loss due to accidents, whereas unclear standards and relatively unpredictable enforcement may deter expenditure on loss prevention. This is important, especially under arising cyber security loss liability, where investment in precaution may be more expensive and damages more extensive, and yet liability standards are currently relatively unclear and vary greatly by jurisdiction.

[82] See” Cyber insurance is a good complement to a high level of information security controls, says Aon Risk Solutions”, Computer Weekly, July 2014 http://www.computerweekly.com/news/2240224437/Cyber-insurance-complements-security-controls-says-Aon and “Cyber Insurance as one element of the Cyber risk management strategy, Deloitte LLC, downloaded July 2016 http://www2.deloitte.com/content/dam/Deloitte/lu/Documents/risk/lu-cyber-insurance-cyber-risk-management-strategy-03032015.pdf

# 11 – Limitations and Further Research

This research has approached the specific problem of evaluating decision making under risk and uncertainty for security control selection using lab experiments to generate individual choice data. The experiments permit the analysis of important but very specific hypotheses drawn from a potentially much richer range of questions and issues relevant to business managers operating enterprise information systems at risk of security attack and breach. Indeed, paraphrasing Harrison and others, we question whether and to what extent lab results regarding security control choice 'in the small' have relevance 'in the large'. We propose that the additional hypotheses noted in Section 1 can and should be explored using the experiments and methods developed here given more time and resources to refine and extend these analyses. I suggest that the combination of system simulation and experimental choice tools developed here combined with additional elements of managerial decision making found in practice are a clear opportunity to extend this research.

This research also focuses intentionally on the quantification of risk since quantification of business operations is considered to improve the objective evaluation and management of risk including security risks. On the other hand, qualitative security risk management approaches continue to dominate in practice (particularly in healthcare in my experience) because they often make more intuitive sense to busy managers. Some further consideration of the qualitative setting in which security risk management actually occurs, and the translation of quantified results by management into action would therefore extend and complement this study, notwithstanding the direct benefits of the illustrated quantified results and insights on security risk quantification presented here. Overall, we consider that the tools and methods developed here may be better applied *in the field* rather than the lab to reflect specific institutional circumstances and system architectures at risk, deepening both the simulation and decision making analyses to address more complex real-world environments.

This research also does not directly address more complex security 'optimization' problems that necessarily involve the modelling and trade-offs between corporate level security budgets and business outcomes, although some of these approaches have been well reviewed within this work in the preparation of the simulation modeling. When confronted with 'system simulation' and the associated quantification of security outcomes beyond typical system 'uptime' statistics, we expect that some or most practitioners would want to know what the model suggests as an 'optimal' security posture along the lines of Wang or Sawik {Wang, 2008 #13}{Sawik, 2013 #91}, and regardless of individual choice i.e. minimally, from a risk neutral posture. We suspect that, in practice, many managers might not welcome information that they are in fact biased and might prefer to know what the 'risk neutral system says' and leave it at that, and separating the simulation aspects from the decisional aspects might be a strong consideration (indeed this is what appears to be happening with many of the simulation platforms explored for this research).

Incorporation of optimization components into the simulation platform developed here would be an obvious extension of this work to balance the perception that 'there is no optimum, it's all relative to risk attitude'. From my perspective, this emphasizes the value of security risk modeling itself since it is only when information systems are presented as being at quantifiable risk that we can start to ask meaningful questions about very quantified costs and operational outcomes. Particularly, the introduction of stochasticity and associated Monte Carlo methods in a practitioner context elevates the security risk discussion to a new level of analysis and reflection, however the concepts are not particularly native to the sector in my opinion, although that is now changing, in no small part because of regulation and the maturing of cyber insurance products and markets. Additional work should be undertaken in either the lab or the field to ensure that practitioners encountering these methods have a solid grounding in and appreciation for the quantitative nuances and elements of these approaches before the tools are introduced in a practitioner setting.

A number of the practical aspects of this experimental approach have required careful planning and limitations of scope to ensure feasible, adequate and reliable data collection within the resources available to this researcher. While we were successful in accomplishing the intended study, additional work regarding assuring the incentive compatibility of the experiments would be necessary to ensure valid results in alternative experimental settings and contexts.  In the first place, a particular feature of these behavioural economics lab experiments which distinguishes them from psychological decision making approaches is the explicit recognition of the need to ensure that decision makers are in fact making 'honest' (if not always economically or psychologically consistent) choices over risky prospects (Grether 1980; Grether 1992). This is formally intended to avoid what is termed 'hypothetical bias': the documented tendency for experimental subjects to make hypothetical choices which are inconsistent with how they would tend to behave under field or 'real world' circumstances (Neill, Cummings et al. 1994; Cubitt and Sugden 2001; Holt 2002; Ajzen, Brown et al. 2004; Harrison 2005). The experimental economics solution for this is to provide clear economic incentives within the experiment that are, at minimum, internally consistent with the decision choice tasks and treatment variations undertaken by the respondents.

There are documented theoretical and methodological considerations for this incentive compatible approach which would need to be addressed to extend this research. First, using cash payments is intended to generally incent participants to make non-arbitrary choices, although this is no absolute guarantee against arbitrary decision making within the lab setting and, as we report, may be a continuing source of error within the approach. One fix for that might be to make decisions consensual, reflecting the fact that, in practice, significant security control decisions are rarely made individually.  This would 'correct' for any arbitrary decisions made by individuals although clearly at the cost of individual risk attitude profiling. The benefit would be possibly a richer and more practical domain for 'risk decisioning' while still based on system simulation and with clear controls and measurement for the effects of uncertainty. Second, as noted above, the incentives employed are also necessarily nominal versus those which may be expected under typical field conditions involving control decisions perhaps worth tens of thousands of dollars (in either control costs or various business outcomes to be optimized) (Pratt 1964; Cox and Sadiraj 2008)). Scaling of experiment payoffs or simply extending the research scope to more sites would require significant funding and is an ongoing, particular consideration for this type of research. Third, one could question whether these experiments are validly testing the subject's actual preferences over analogous real world control decisions (since the amounts in play are nominal), and therefore whether the subject perhaps considers that they are just 'playing a game' and not making analogous 'consequential' control decisions, notwithstanding the monetary incentives, or indeed whether they are, in any objective sense, actually playing the 'game' that you intend them to be playing (Smith 1982; Harrison 2010; Smith 2010). Again, this could be possibly better controlled by placing the experiments in professional settings beyond the ones explored here to induce participants to make more considered choices. Fourth, bets placed using non-personal funds and 'windfall' gains (however nominal) may also be treated differently by subjects across experimental and field settings (Carlsson, He et al. 2009; Chakravarty, Harrison et al. 2011). In the context of enterprise information security, placing the simulations in a field setting reflective of an actual enterprise and asking participants to then make choices that may affect real business decisions and performance would likely provide an opportunity to test for appropriate practitioner incentives vs. incentives that are typically developed for use with either the 'general public or MBA students (no matter how much I appreciate MBA students). Fifth, the predictability of 'out-of-context' decisions (i.e. the ability to make predictions outside of the experimental setting, either within or across subjects) may also be lower than within experiment results (Wilcox 2010; Friedman, Isaac et al. 2014). This research was unable to explore this aspect of analysis and additional study could be undertaken with guidance from Harrison and others regarding the econometric robustness of the experimental results beyond the participating sample.

Lastly, we recognize that the specific system simulation environments developed here were necessarily a hybrid of the best practice platforms and approaches reviewed. Some work could be undertaken to 'replatform' the simulation to allow it to run natively in P2CySeMol so that maximal advantage is taken of the Bayesian network approach to component attack and incident determination while retaining the ability to translate incidents into stochastic business losses. I remain encouraged by comments received from Mathias Ekstedt during the development of this work who, when I described my research agenda concerning security decision making under uncertainty and its potential use of CySeMol, commented that this research would constitute "…a full bingo on his scorecard'. I can think of no better continued motivation for further research from a valued colleague.

# References

Abdellaoui, M., A. Baillon, et al. (2008). The rich domain of uncertainty, Working Paper.

Abdellaoui, M. and B. Munier (1998). "The risk-structure dependence effect:Experimenting with an eye to decision-aiding." Annals of Operations Research **80**(1): 237-252.

Acerbi, C. (2002). "Spectral measures of risk: a coherent representation of subjective risk aversion." Journal of Banking & Finance **26**(7): 1505-1518.

Acerbi, C., C. Nordio, et al. (2001). "Expected shortfall as a tool for financial risk management." arXiv preprint cond-mat/0102304.

Acerbi, C. and D. Tasche (2002). "On the coherence of expected shortfall." Journal of Banking & Finance **26**(7): 1487-1503.

Acquisti, A. (2004). Privacy and Security of Personal Information. Economics of Information Security. L. Camp and S. Lewis, Springer US. **12:** 179-186.

Acquisti, A. (2004). Privacy in electronic commerce and the economics of immediate gratification. Proceedings of the 5th ACM conference on Electronic commerce. New York, NY, USA, ACM**:** 21-29.

Acquisti, A. (2005). Uncertainty, Ambiguity and Privacy. 4th Annual Workshop on Economics and Information Security (WEIS 2005).

Acquisti, A. (2009). "Nudging Privacy: The Behavioral Economics of Personal Information." IEEE Security and Privacy **7**(6): 82-85.

Acquisti, A. (2010). The Economics of Personal Data and the Economics of Privacy. Working Party for Information Security and Privacy (WPISP) / Working Party on the Information Economy (WPIE) - Joint WPISP-WPIE Roundtable "The Economics of Personal Data and Privacy: 30 Years after the OECD Privacy Guidelines".

Acquisti, A. (2010). From the Economics to the Behavioral Economics of Privacy: A Note. Ethics and Policy of Biometrics. A. Kumar and D. Zhang, Springer Berlin / Heidelberg. **6005:** 23-26.

Acquisti, A., A. Friedman, et al. (2006). "Is there a cost to privacy breaches? An event study." ICIS 2006 Proceedings: 94.

Acquisti, A. and J. Grossklags (2003). Losses, Gains, and Hyperbolic Discounting: An Experimental Approach to Information Security Attitudes and Behavior. 2nd Annual Workshop on “Economics and Information Security”. UC Berkeley.

Acquisti, A. and J. Grossklags (2005). "Privacy and rationality in individual decision making " Security & Privacy, IEEE **3**(1).

Acquisti, A. and J. Grossklags (2006). What Can Behavioral Economics Teach Us About Privacy? Keynote Paper, ETRICS 2006.

Acquisti, A., L. John, et al. (2009). What is Privacy Worth? Twenty First Workshop on Information Systems and Economics (WISE)
Phoenix, AZ.

Ajzen, I., T. C. Brown, et al. (2004). "Explaining the Discrepancy between Intentions and Actions: The Case of Hypothetical Bias in Contingent Valuation." Personality and Social Psychology Bulletin **30**(9): 1108-1121.

Alary, D., C. Gollier, et al. (2010). "The effect of ambiguity aversion on insurance demand." Toulouse School of Economics.

Alcazar and Fenz (2012). "Mapping ISO 27002 into Security Ontology."

Allais, M. (1953). "Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école Américaine." Econometrica **21**: 503-546.

Andersen, S., J. Fountain, et al. (2009). "Estimating aversion to uncertainty." Unpublished discussion paper, College of Business, Univ. of Central Florida (May).

Andersen, S., J. Fountain, et al. (2009). Estimating Aversion to Uncertainty, University of Central Florida.

Andersen, S., J. Fountain, et al. (2014). "Estimating subjective probabilities." Journal of Risk and Uncertainty **48**(3): 207-229.

Andersen, S., G. W. Harrison, et al. (2007). "Behavioral Econometrics for Psychologists." SSRN eLibrary.

Andersen, S., G. W. Harrison, et al. (2006). Dual Criteria Decisions.

Andersen, S., G. W. Harrison, et al. (2006). "Elicitation using multiple price list formats." Experimental Economics **9**(4): 383-405.

Andersen, S., G. W. Harrison, et al. (2006). Choice Behavior, Asset Integration and Natural Reference Points. Working Paper 06-07, , Department of Economics,College of Business Administration, University of Central Florida.

Anderson, E. E. (2010). "Firm objectives, IT alignment, and information security." IBM Journal of Research and Development **54**(3): 5:1-5:7.

Anderson, R. (1996). Security in Clinical Information Systems. Proceedings of the 15 th IEEE Symposium on Security and Privacy.

Anderson, R. (2001). Why Information Security is Hard-An Economic Perspective. Proceedings of the 17th Annual Computer Security Applications Conference, IEEE Computer Society**:** 358.

Anderson, R. (2001). Why information security is hard - an economic perspective. Computer Security Applications Conference, 2001. ACSAC 2001. Proceedings 17th Annual.

Anderson, R., R. Böhme, et al. (2009). Security Economics and European Policy. Managing Information Risk and the Economics of Security, Springer US**:** 55-80.

Anderson, R. and T. Moore (2006). "The Economics of Information Security." Science **314**(5799): 610-613.

Anderson, R. and T. Moore (2007). Information security economics - and beyond. Proceedings of the 27th annual international cryptology conference on Advances in cryptology. Santa Barbara, CA, USA, Springer-Verlag**:** 68-91.

Antoniou, C. (2010). Three essays in behavioural finance: An examination into non- Bayesian Investment behaviour, Durham University.

Antoniou, C., G. W. Harrison, et al. (2010). Subjective Bayesian Beliefs, Durham Business School, Durham University, UK.

Antoniou, C., G. W. Harrison, et al. (2015). "Subjective Bayesian Beliefs." Journal of Risk and Uncertainty **50**(1): 35-54.

Appari, A. a. J., M. Eric (2006). Which Hospitals Are Complying with HIPAA: An Empirical Investigation of US Hospitals. Hanover NH, Center for Digital Strategies, Tuck School of Business at Dartmouth.

Appari, A. a. J., M. Eric (2009). HIPAA Compliance: An Institutional Theory Perspective. AMCIS 2009.

Appari, A. a. J., M. Eric (2010). "Information security and privacy in healthcare: current state of research." Int. J. Internet and Enterprise Management **6**(4).

Artinger, F., F. Exadaktylos, et al. (2010). "Applying quadratic scoring rule transparently in multiple choice setting: a note."

Artzner, P., F. Delbaen, et al. (1999). "Coherent Measures of Risk." Mathematical Finance **9**(3): 203-228.

Atkinson, A. C. (1970). "A Method For Discriminating Between Models." Journal of the Royal Statistical Society. Series B (Methodological) **32**(3): 323-353.

Avižienis, A., J.-C. Laprie, et al. (2004). "Basic concepts and taxonomy of dependable and secure computing." Dependable and Secure Computing, IEEE Transactions on **1**(1): 11-33.

Baddeley, M. (2010). Security: Foundations from Behavioural Economics, Security and Human Behaviour 2010.

Baer, W. S. and A. Parkinson (2007). "Cyberinsurance in it security management." IEEE Security and Privacy **5**(3): 50-56.

Bajtelsmit, V., J. C. Coats, et al. (2015). "The effect of ambiguity on risk management choices: An experimental study." Journal of Risk and Uncertainty **50**(3): 249-280.

Baldwin, A., Y. Beres, et al. (2011). Economic methods and decision making by security professionals. Tenth Workshop on Economics of Information Security (WEIS 2011). George Mason University, Virginia.

Baldwin, A., M. C. Mont, et al. (2009). Using Modelling and Simulation for Policy Decision Support in Identity Management. Policies for Distributed Systems and Networks, 2009. POLICY 2009. IEEE International Symposium on.

Banks, J. (1984). Discrete-event system simulation, Pearson Education India.

Bansal, G., F. M. Zahedi, et al. (2010). "The impact of personal dispositions on information sensitivity, privacy concern and trust in disclosing health information online." Decis. Support Syst. **49**(2): 138-150.

Barberis, N., A. Shleifer, et al. (1998). "A model of investor sentiment." Journal of Financial Economics **49**(3): 307-343.

Barnett, A. H. (2009). "The End of the Externality Revolution." Social Philosophy and Policy **26**: 130-150.

Barron, G. and I. Erev (2003). "Small feedback-based decisions and their limited correspondence to description-based decisions." Journal of Behavioral Decision Making **16**(3): 215-233.

Baumol, W. J. (1972). "On Taxation and the Control of Externalities." The American Economic Review **62**(3).

Beautement, A., R. Coles, et al. (2009). Modelling the Human and Technological Costs and Benefits of USB Memory Stick Security. Managing Information Risk and the Economics of Security, Springer US**:** 141-163.

Beautement, A. and D. Pym (2010). "Structured Systems Economics for Security Management." WEIS 2010.

Beautement, A., M. A. Sasse, et al. (2008). The compliance budget: managing security behaviour in organisations. Proceedings of the 2008 workshop on New security paradigms. Lake Tahoe, California, USA, ACM**:** 47-58.

Becker, G. M., M. H. Degroot, et al. (1964). "Measuring utility by a single-response sequential method." Behavioral Science **9**(3): 226-232.

Bellamy, C. a. R., Charles (2010). "Information-sharing dilemmas in public services: using frameworks from risk management." Policy & Politics **38**(3).

Ben-Tal, A. and A. Ben-Israel (1991). "A recourse certainty equivalent for decisions under uncertainty." Annals of Operations Research **30**(1): 1-44.

Benartzi, S. and R. H. Thaler (1995). "Myopic Loss Aversion and the Equity Premium Puzzle." The Quarterly Journal of Economics **110**(1): 73-92.

Bennett, C. J. (1992). Regulating Privacy: Data Protection and Public Policy in Europe and the United States, Cornell University Press.

Bennett, C. J. (2000). "The Political Economy of Privacy: A Review of the Literature." Center for Social and Legal Research, DOE Human Genome Project.

Beres, Y., M. Casassa Mont, et al. (2009). Using security metrics coupled with predictive modeling and simulation to assess security processes. Proceedings of the 2009 3rd International Symposium on Empirical Software Engineering and Measurement, IEEE Computer Society**:** 564-573.

Beres, Y., J. Griffin, et al. (2008). Analysing the Performance of Security Solutions to Reduce Vulnerability Exposure Window. Computer Security Applications Conference, 2008. ACSAC 2008. Annual.

Beresnevichiene, Y., D. Pym, et al. (2010). Decision support for systems security investment. Network Operations and Management Symposium Workshops (NOMS Wksps), 2010 IEEE/IFIP.

Berg (2005). "Risk preference instability across institutions: A dilemma." Proceedings of the National Academy of Sciences of the United States of America (PNAS) **102**(11): 4209-4214.

Berliner, B. (1985). "Large risks and limits of insurability." Geneva Papers on Risk and Insurance: 313-329.

Bernoulli, D. (1738 (1954)). "Specimen Theoriae Novae de Mensura Sortis, Commentarii Academiae Scientiarum Imperialis Petropolitanae,5, 175-192 English translation (1954):
Exposition of a New Theory on the Measurement of Risk." Econometrica **22**: 23-36.

Biener, C., M. Eling, et al. (2015). "Insurability of Cyber Risk: An Empirical Analysis†." The Geneva Papers on Risk and Insurance-Issues and Practice **40**(1): 131-158.

Birnbaum, M. H. (2000). "Decision making in the lab and on the Web." Psychological experiments on the Internet: 3-34.

Black, F. (1986). "Noise." The Journal of Finance **41**(3): 529-543.

Blakley, B., E. McDermott, et al. (2001). Information security is information risk management. Proceedings of the 2001 Workshop on New Security Paradigms. Cloudcroft, New Mexico, ACM**:** 97-104.

Boehmer, W. (2011). Dynamic systems approach to analyzing event risks and behavioral risks with game theory. Privacy, Security, Risk and Trust (PASSAT) and 2011 IEEE Third Inernational Conference on Social Computing (SocialCom), 2011 IEEE Third International Conference on, IEEE.

Boehmer, W. (2012). "Information Security Management Systems Cybernetics." Strategic and Practical Approaches for Information Security Governance: Technologies and Applied Solutions: Technologies and Applied Solutions: 223.

Böhme, R. (2005). Cyber-Insurance Revisited. WEIS.

Bonner, B., M. Chiasson, et al. (2009). "Restoring balance: How history tilts the scales against privacy. An Actor-Network Theory investigation." Information and Organization **19**(2): 84-102.

Bonner, W. and M. Chiasson (2005). "If fair information principles are the answer, what was the question? An actor-network theory investigation of the modern constitution of privacy." Information and Organization **15**(4): 267-293.

Bregman-Eschet, Y. (2006). "Genetic Databases and Biobanks: Who Controls our Genetic Privacy?" Santa Clara Computer & High Tech. L.J. **23**(1).

Brooke, G. T. F. (2010). "Uncertainty, Profit and Entrepreneurial Action: Franks Knight's contribution reconsidered." Journal of the History of Economic Thought **32**.

Bruner, D. M. (2011). "Multiple switching behaviour in multiple price lists." Applied Economics Letters **18**(5): 417-420.

Buchanan, J. M. a. S., Wm. Craig (1962). "Externality." Economica **29**(116).

Buschle, M. (2014). Tool Support for Enterprise Architecture Analysis - with application in cyber security.

Buschle, M., J. Ullberg, et al. (2011). "A Tool for Enterprise Architecture Analysis using the PRM formalism."

Camerer, C. and M. Weber (1992). "Recent developments in modeling preferences: Uncertainty and ambiguity." Journal of Risk and Uncertainty **5**(4): 325-370.

Camerer, C. F. (2000). Prospect theory in the wild: Evidence from the field. Advances in Behavioral Economics**:** 148-161.

Camp, L. J. (2006). Economics of Information Security. SSRN eLibrary, Indiana University Bloomington - School of Informatics.

Camp, L. J. (2009). "Mental Models of Privacy and Security." IEEE Technology & Society Magazine **28**(3): 37-46.

Carlsson, F., H. He, et al. (2009). "Easy come, easy go - The role of windfall money in lab and field experiments." Working Papers in Economics **374**.

Cavusoglu, H., B. Mishra, et al. (2004). "A model for evaluating IT security investments." Commun. ACM **47**(7): 87-92.

Cavusoglu, H., S. Raghunathan, et al. (2008). "Decision-Theoretic and Game-Theoretic Approaches to IT Security Investment." J. Manage. Inf. Syst. **25**(2): 281-304.

Cavusoglu, H. a. C., Hasan and Zhang, Jun (2006). Economics of Security Patch Management. The Fifth Workshop on the Economics of Information Security (WEIS 2006).

Chakravarty, S., G. W. Harrison, et al. (2011). "Are You Risk Averse over Other People's Money?" Southern Economic Journal **77**(4): 901-913.

Charness, G., U. Gneezy, et al. (2013). "Experimental methods: Eliciting risk preferences." Journal of Economic Behavior & Organization **87**: 43-51.

Charniak (1991). "Bayesian Networks without Tears."

Chellappa, R. K. and S. Shivendu (2006). "An Economic Model of Privacy: a Property Rights Approach to Regulatory Choices for Online Personalization." SSRN eLibrary.

Clarke, K. A. (2007). "A Simple Distribution-Free Test for Nonnested Model Selection." Political Analysis **15**(3): 347-363.

Coase, R. A. (1960). "The Problem of Social Cost." Journal of Law and Economics **3**(Oct., 1960): 1-44.

Coles, S., J. Bawa, et al. (2001). An introduction to statistical modeling of extreme values, Springer.

Collinson, M., B. Monahan, et al. (2009). "A Logical and Computational Theory of Located Resource." J. Log. and Comput. **19**(6): 1207-1244.

Collinson, M., B. Monahan, et al. (2010). Semantics for structured systems modelling and simulation. Proceedings of the 3rd International ICST Conference on Simulation Tools and Techniques. Torremolinos, Malaga, Spain, ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering)**:** 1-8.

Collinson, M. and D. Pym (2009). "Algebra and logic for resource-based systems modelling." Mathematical Structures in Computer Science **19**(05): 959-1027.

Conrad, J. R. (2005). "Analyzing the Risks of Information Security Investments with Monte-Carlo Simulations " IEEE Computer Society.

Conrad, J. R., P. Oman, et al. (2006). Managing Uncertainty in Security Risk Model Forecasts with RAPSA/MC. Security Management, Integrity, and Internal Control in Information Systems: IFIP TC-11 WG 11.1 & WG 11.5 Joint Working Conference. P. Dowland, S. Furnell, B. Thuraisingham and X. S. Wang. Boston, MA, Springer US**:** 141-156.

Cooper, G. F. (1990). "The computational complexity of probabilistic inference using Bayesian belief networks." Artificial intelligence **42**(2-3): 393-405.

Cox, D. R. (1961). Tests of Separate Families of Hypotheses. Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Berkeley, University of California Press.

Cox, D. R. (1962). "Further Results on Tests of Separate Families of Hypotheses." Journal of the Royal Statistical Society. Series B (Methodological) **24**(2): 406-424.

Cox, J. C. and V. Sadiraj (2006). "Small-and large-stakes risk aversion: Implications of concavity calibration for decision theory." Games and Economic Behavior **56**(1): 45-60.

Cox, J. C. and V. Sadiraj (2008). Risky Decisions in the Large and in the Small: Theory and Experiment. Experimental Economics Center Working Paper Series, Georgia State University. **2008**.

Crookall, D. (2010). "Serious Games, Debriefing, and Simulation/Gaming as a Discipline." Simulation & Gaming **41**(6): 898-920.

Cubitt, R. P., C. Starmer, et al. (1998). "Dynamic Choice and the Common Ratio Effect: An Experimental Investigation." The Economic Journal **108**(450): 1362-1380.

Cubitt, R. P. and R. Sugden (2001). "Dynamic Decision-Making Under Uncertainty: An Experimental Investigation of Choices Between Accumulator Gambles." Journal of Risk and Uncertainty **22**(2): 103-128.

D'Arcy, J. a. H., Anat (2009). An Integrative Framework for the Study of Information Security Management Research Handbook of Research on Information Security and Assurance.

Dagum, P. and M. Luby (1993). "Approximating probabilistic inference in Bayesian belief networks is NP-hard." Artificial intelligence **60**(1): 141-153.

Dahlman, C. J. (1979). "The Problem of Externality." Journal of Law and Economics
**22**(1).

Daniel, K., D. Hirshleifer, et al. (1998). "Investor Psychology and Security Market under- and Overreactions." The Journal of Finance **53**(6): 1839-1885.

Danielsson, J. (2002). "The emperor has no clothes: Limits to risk modelling." Journal of Banking &amp; Finance **26**(7): 1273-1296.

Dantzig, G. and G. Infanger (1993). "Multi-stage stochastic linear programs for portfolio optimization." Annals of Operations Research **45**(1): 59-76.

Dave, C., C. C. Eckel, et al. (2010). "Eliciting risk preferences: When is simple better?" Journal of Risk and Uncertainty **41**(3): 219-243.

Davison, A. C. and R. L. Smith (1990). "Models for exceedances over high thresholds." Journal of the Royal Statistical Society. Series B (Methodological): 393-442.

DeBondt, W. F. M. and R. Thaler (1985). "Does the Stock Market Overreact?" The Journal of Finance **40**(3): 793-805.

Dembo, R. S. and A. J. King (1992). "Tracking models and the optimal regret distribution in asset allocation." Applied Stochastic Models and Data Analysis **8**(3): 151-157.

DiMaggio, P. J. a. P., Walter W. (1983). "The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields." American Sociological Review **48**(2).

Duffie, D. and J. Pan (1997). "An overview of value at risk." The Journal of derivatives **4**(3): 7-49.

Dupuis, D. J. (1999). "Exceedances over High Thresholds: A Guide to Threshold Selection." Extremes **1**(3): 251-261.

Eagle, S. J. (2004). "Environmental Amenities, Private Property and Public Policy." Natural Resources Journal, Vol. 44, No. 2, pp. 425-444, Spring 2004.

Edwards, W. (1954). "The theory of decision making. ." Psychological Bulletin **51**(4): 380-417.

Edwards, W. (1961). "Behavioral Decision Theory." Annual Review of Psychology **12**: 473-498.

Ehrlich, I. and G. S. Becker (1972). "Market Insurance, Self-Insurance, and Self-Protection." Journal of Political Economy, Vol. 80, No. 4, pp. 623-648, July-August 1972.

Eisenhardt, K. M. (1989). "Agency Theory: An Assessment and Review." The Academy of Management Review **14**(1).

Ekelhart, Fenz, et al. (2009). AURUM: A Framework for Information Security Risk Management. HICSS'09. 42nd Hawaii International Conference on System Sciences, 2009. , IEEE.

Ekelhart, A., S. Fenz, et al. (2006). Security Ontology: Simulating Threats to Corporate Assets. International Conference on Information Systems Security, Springer.

Ekelhart, A., S. Fenz, et al. (2009). Ontology-based Decision Support for Information Security Risk Management. ICONS'09. Fourth International Conference on Systems, 2009., IEEE.

Ekstedt, M. (2004). Enterprise Architecture for IT Management: A CIO Decision Making Perspective on the Electric Power Industry.

El-Gamal, M. A. and D. M. Grether (1995). "Are People Bayesian? Uncovering Behavioral Strategies." Journal of the American Statistical Association **90**(432): 1137-1145.

Eling, M. and J. H. Wirfs (2015). "Modelling and Management of Cyber Risk."

Ellis, B. and W. H. Wong (2008). "Learning Causal Bayesian Network Structures from Experimental Data." Journal of the American Statistical Association **103**: 778-789.

Ellison, R. J., D. A. Fisher, et al. (1997). Survivable network systems: An emerging discipline, DTIC Document.

Ellsberg, D. (1961). "Risk, Ambiguity, and the Savage Axioms." The Quarterly Journal of Economics **75**(4).

Embrechts, P., S. I. Resnick, et al. (1999). "Extreme value theory as a risk management tool." North American Actuarial Journal **3**(2): 30-41.

Fama, E. F. and K. R. French (1993). "Common risk factors in the returns on stocks and bonds." Journal of financial economics **33**(1): 3-56.

Farquhar, P. H. (1984). "Utility Assessment Methods." Management Science **30**(11): 1283-1300.

Feiler, D. C., J. D. Tong, et al. (2013). "Biased judgment in censored environments." Management Science **59**(3): 573-591.

Felde, M. (2010). Analyzing Security Decisions with Discrete Event Simulation.

Fenton, N. E., M. Neil, et al. (2007). "Using Ranked Nodes to Model Qualitative Judgments in Bayesian Networks." IEEE Trans. Knowl. Data Eng. **19**(10): 1420-1432.

Fenz, S. (2012). "An ontology-based approach for constructing Bayesian networks." Data & Knowledge Engineering **73**: 73-88.

Fenz, S. (2012). "An ontology-based approach for constructing Bayesian networks." Data Knowl. Eng. **73**: 73-88.

Fenz, S. and A. Ekelhart (2009). Formalizing Information Security Knowledge. Proceedings of the 4th international Symposium on Information, Computer, and Communications Security, ACM.

Fenz, S., A. Ekelhart, et al. (2011). "Information Security Risk Management: In Which Security Solutions Is It Worth Investing?" Communications of the Association for Information Systems **28**(1): 329-356.

Fenz, S., A. Ekelhart, et al. (2012). "Information Security Risk Management: In Which Security Solutions Is It Worth Investing?" Communications of the Association for Information Systems **28**(1).

Fenz, S. and T. Neubauer (2009). How to determine threat probabilities using ontologies and Bayesian networks. Proceedings of the 5th Annual Workshop on Cyber Security and Information Intelligence Research: Cyber Security and Information Intelligence Challenges and Strategies, ACM.

Fenz, S. and A. M. Tjoa (2008). Ontology- and Bayesian-based Threat Probability Determination.

Fenz, S., A. M. Tjoa, et al. (2009). Ontology-based generation of Bayesian networks. CISIS'09. International Conference on Complex, Intelligent and Software Intensive Systems, 2009. , IEEE.

Fenz, S. and E. R. Weippl (2006). Ontology based IT-security planning. PRDC.

Fiore, S. M., G. W. Harrison, et al. (2009). "Virtual experiments and environmental policy." Journal of Environmental Economics and Management **57**(1): 65-86.

Fishburn, P. C. (1967). "Methods of Estimating Additive Utilities." Management Science **13**(7): 435-453.

Franke, U., W. R. Flores, et al. (2009). Enterprise Architecture Dependency Analysis using Fault Trees and Bayesian Networks. Proceedings of the 2009 Spring Simulation Multiconference, Society for Computer Simulation International.

Franke, U., D. Hook, et al. (2009). EAF2-a framework for categorizing enterprise architecture frameworks. Software Engineering, Artificial Intelligences, Networking and Parallel/Distributed Computing, 2009. SNPD'09. 10th ACIS International Conference on, IEEE.

Franke, U., P. Johnson, et al. (2014). "An architecture framework for enterprise IT service availability analysis." Software & Systems Modeling **13**(4).

Fréchette, G. R. and A. Schotter (2015). Handbook of experimental economic methodology, Oxford University Press, USA.

Freeman, D., Y. Halevy, et al. (2015). Eliciting risk preferences using choice lists, Vancouver School of Economics.

Friedman, D., R. M. Isaac, et al. (2014). Risky curves: On the empirical failure of expected utility, Routledge.

Friedman, N., L. Getoor, et al. (1999). Learning Probabilistic Relational Models. IJCAI.

George, J. G., G. W. Harrison, et al. (2012). Behavioral Responses towards Risk Mitigation: An Experiment with Wild Fire Risks.

Gheorghe, M. (2012). "Techniques and Simulation Models in Risk Management." Economia. Seria Management **15**(2): 354-362.

Gilboa, I. and D. Schmeidler (1989). "Maxmin expected utility with non-unique prior." Journal of Mathematical Economics **18**(2): 141-153.

Gilli, M. and E. Këllezi (2006). "An Application of Extreme Value Theory for Measuring Financial Risk." Computational Economics **27**(2): 207-228.

Gneiting, T. and A. E. Raftery (2007). "Strictly proper scoring rules, prediction, and estimation." Journal of the American Statistical Association **102**(477): 359-378.

Gollier, C. (2016). Explaining rank-dependent utility with regret and rejoicing, Institut d'Économie Industrielle (IDEI), Toulouse.

Gordon, L. and M. Loeb (2006). "Economic aspects of information security: An emerging field of research." Information Systems Frontiers **8**(5): 335-337.

Gordon, L. A. and M. P. Loeb (2002). "The economics of information security investment." ACM Trans. Inf. Syst. Secur. **5**(4): 438-457.

Gordon, L. A., M. P. Loeb, et al. (2003). "Information Security Expenditures and Real Options: A Wait-and-See Approach." Computer Security Journal **XIX**(2).

Gordon, L. A., M. P. Loeb, et al. (2008). "Cybersecurity, Capital Allocations and Management Control Systems." European Accounting Review **17**(2): 215-241.

Gordon, L. A. and R. Richardson (2004). "The New Economics of Information Security." InformationWeek(982): 52-57.

Gray, J. and D. P. Siewiorek (1991). "High-availability computer systems." Computer **24**(9): 39-48.

Grechuk, B., A. Molyboha, et al. (2009). "Maximum entropy principle with general deviation measures." Mathematics of Operations Research **34**(2): 445-467.

Greiner, B., H. A. Jacobsen, et al. (2012). The Virtual Laboratory Infrastructure for Controlled Online Experiments in Economics, Max Planck Institute for Research into Economic Systems, Strategic Interaction Group, Jena. http://www. billingpreis. mpg. de/hbp02/schmidt. pdf.

Grether, D. M. (1980). "Bayes Rule as a Descriptive Model: The Representativeness Heuristic." The Quarterly Journal of Economics **95**(3): 537-557

Grether, D. M. (1992). "Testing bayes rule and the representativeness heuristic: Some experimental evidence." Journal of Economic Behavior & Organization **17**(1): 31-57.

Grether, D. M. and C. R. Plott (1979). "Economic Theory of Choice and the Preference Reversal Phenomenon." The American Economic Review **69**(4): 623-638

Griffin, D. and A. Tversky (1992). "The weighing of evidence and the determinants of confidence." Cognitive Psychology **24**(3): 411-435.

Grossklags, J., N. Christin, et al. (2008). Secure or insure?: a game-theoretic analysis of information security games. Proceeding of the 17th international conference on World Wide Web. Beijing, China, ACM**:** 209-218.

Grossklags, J., N. Christin, et al. (2008). Security investment (failures) in five economic environments: A comparison of homogeneous and heterogeneous user agents, School of Information, University of California, Berkeley.

Grossklags, J., B. Johnson, et al. (2010). The Price of Uncertainty in Security Games. Economics of Information Security and Privacy. T. Moore, D. Pym and C. Ioannidis, Springer US**:** 9-32.

Grossklags, J., B. Johnson, et al. (2010). When Information Improves Information Security. Financial Cryptography and Data Security. R. Sion, Springer Berlin / Heidelberg. **6052:** 416-423.

Grossman, P. J. and C. C. Eckel (2015). "Loving the long shot: Risk taking with skewed lotteries." Journal of Risk and Uncertainty **51**(3): 195-217.

Grunske, L. and D. Joyce (2008). "Quantitative risk-based security prediction for component-based systems with explicitly modeled attack profiles." Journal of Systems and Software **81**(8): 1327-1345.

Habermas, J. (1967). On the Logic of the Social Sciences, MIT Press.

Haddock, D. D. (2003). Irrelevant Internalities, Irrelevant Externalities, and Irrelevant Anxieties. Northwestern Law & Economics Research Paper, Northwestern University - School of Law and Department of Economics.

Haddock, D. D. (2007). "Irrelevant Externality Angst." The Journal of Interdisciplinary Economics **19**: 3–18.

Hall (2010). "Property, Privacy, and the Pursuit of Interconnected Electronic Medical Records." Iowa L. Rev. **95**(631).

Hardy, C. and S. Maguire (2016). "Organizing Risk: Discourse, Power, and “Riskification”." Academy of Management Review **41**(1): 80-108.

Harrison, G. and E. Rutström (2009). "Expected utility theory and prospect theory: one wedding and a decent funeral." Experimental Economics **12**(2): 133-158.

Harrison, G. W. (1986). "An experimental test for risk aversion." Economics Letters **21**(1): 7-11.

Harrison, G. W. (1989). "Theory and Misbehavior of First-Price Auctions." The American Economic Review **79**(4): 749-762.

Harrison, G. W. (1990). "Risk Attitudes in First-Price Auction Experiments: A Bayesian Analysis." The Review of Economics and Statistics **72**(3): 541-546.

Harrison, G. W. (1994). "Expected utility theory and the experiments." Empirical Economics **19**(2): 223-253.

Harrison, G. W. (2005). "Experimental Evidence on Alternative Environmental Valuation Methods." SSRN eLibrary.

Harrison, G. W. (2007). Making Choice Studies Incentive Compatible. Valuing Environmental Amenities Using Stated Choice Studies: A Common Sense Approach to Theory and Practice. B. J. Kanninen. Dordrecht, Springer Netherlands**:** 67-110.

Harrison, G. W. (2008). Maximum Likelihood Estimation of Utility Functions Using Stata, University of Central Florida.

Harrison, G. W. (2010). "The behavioral counter-revolution." Journal of Economic Behavior & Organization **73**(1): 49-57.

Harrison, G. W., E. Johnson, et al. (2003). Individual Choice and Risk Aversion in the Laboratory: A Reconsideration. University of Central Florida, Department of Economics Working Paper**:** 3-18.

Harrison, G. W., E. Johnson, et al. (2005). "Risk Aversion and Incentive Effects: Comment." The American Economic Review **95**(3): 897-901.

Harrison, G. W., E. Johnson, et al. (2005). "Risk aversion and incentive effects: Comment." American Economic Review: 897-901.

Harrison, G. W., M. Lau, et al. (2010). Theory, Experimental Design and Econometrics Are Complementary (And So Are Lab and Field Experiments). The Methods of Modern Experimental Economics. G. F. a. A. Schotter. New York, Oxford University Press.

Harrison, G. W., M. I. Lau, et al. (2002). "Estimating individual discount rates in Denmark: A field experiment." The American Economic Review **92**(5): 1606-1617.

Harrison, G. W., J. Martínez-Correa, et al. (2014). "Eliciting subjective probabilities with binary lotteries." Journal of Economic Behavior & Organization **101**: 128-140.

Harrison, G. W., J. Martínez-Correa, et al. (2015). "Reduction of compound lotteries with objective probabilities: Theory and evidence." Journal of Economic Behavior & Organization **119**: 32-55.

Harrison, G. W., J. Martínez-Correa, et al. (2013). "Scoring rules for subjective probability distributions." Manuscript, Georgia State University.

Harrison, G. W. and M. McKee (1985). "Experimental Evaluation of the Coase Theorem." Journal of Law and Economics **28**(3): 653-670.

Harrison, G. W. and E. E. Rutström, Eds. (2008). Risk Aversion in the Laboratory. Risk Aversion in Experiments (Research in Experimental Economics), Emerald Group Publishing Limited.

Harrison, G. W. and E. R. Ulm (2015). "Recovering Subjective Probability Distributions." Working Paper.

Herath, H. S. B. and T. C. Herath (2011). "Copula-based actuarial model for pricing cyber-insurance policies." Insurance Markets and Companies: Analyses and Actuarial Computations **2**(1).

Herland, K. (2015). Information security risk assessment of smartphones using Bayesian networks.

Hermalin, B. and M. Katz (2004). "Privacy, property rights and efficiency: The economics of privacy as secrecy." Quantitative Marketing and Economics **4**(3): 209-239.

Hershey, J. C. and P. J. Schoemaker (1985). "Probability versus certainty equivalence methods in utility measurement: Are they equivalent?" Management Science **31**(10): 1213-1231.

Hertwig, R., G. Barron, et al. (2004). "Decisions from experience and the effect of rare events in risky choice." Psychological science **15**(8): 534-539.

Hey, J. D. and C. Orme (1994). "Investigating Generalizations of Expected Utility Theory Using Experimental Data." Econometrica **62**(6): 1291-1326

Hirshleifer, D. (2001). "Investor Psychology and Asset Pricing." The Journal of Finance **56**(4): 1533-1597.

Hirshleifer, J. (1971). "The Private and Social Value of Information and the Reward to Inventive Activity." The American Economic Review **64**(71).

Hirshleifer, J. (1980). "Privacy: Its Origin, Function, and Future." The Journal of Legal Studies **9**(4).

Holm, H. (2014). A Framework and Calculation Engine for Modeling and Predicting the Cyber Security of Enterprise Architectures.

Holm, H. and M. Ekstedt (2012). A metamodel for web application injection attacks and countermeasures. Trends in Enterprise Architecture Research and Practice-Driven Research on Enterprise Transformation, Springer**:** 198-217.

Holm, H., M. Ekstedt, et al. (2012). "Empirical Analysis of System-Level Vulnerability Metrics through Actual Attacks." IEEE Transactions on dependable and secure computing **9**(6): 825-837.

Holm, H., M. Ekstedt, et al. (2013). Effort estimates on web application vulnerability discovery. 46th Hawaii International Conference on System Sciences (HICSS), 2013 IEEE.

Holm, H., M. Ekstedt, et al. (2013). A Manual for the Cyber Security Modeling Language. T. R. Royal Institute of Technology (KTH).

Holm, H., K. Shahzad, et al. (2015). "P2CySeMoL: Predictive, Probabilistic Cyber Security Modeling Language." IEEE Transactions on Dependable and Secure Computing **12**(6): 626-639.

Holt, C. A., and Laury, S.K. (2002). "Risk Aversion and Incentive Effects." American Economic Review **95**(5): 1644-1655.

Holt, C. A. and A. M. Smith (2007). "An update on Bayesian updating." Journal of Economic Behavior & Organization **69**(2): 125-134.

Hsu, W. H. and R. Joehanes (2004). Relational decision networks. Proceedings of the ICML Workshop on Statistical Relational Learning.

Huang, C. D., Q. Hu, et al. (2006). Economics of Information Security Investment in the Case of Simultaneous Attacks. The Fifth Workshop on the Economics of Information Security (WEIS 2006).

Hui, K.-L. and I. P. Png (2005). Economics of Privacy. Handbook of Information Systems and Economics. e. Terry Hendershott, Elsevier.

Infanger, G. (2006). "Dynamic asset allocation strategies using a stochastic dynamic programming approach." Handbook of asset and liability management **1**: 199-251.

Ioannidis, C., D. Pym, et al. (2009). Investments and Trade-offs in the Economics of Information Security Financial Cryptography and Data Security. R. Dingledine and P. Golle, Springer Berlin / Heidelberg. **5628:** 148-166.

Ioannidis, C., D. Pym, et al. (2011). Fixed Costs, Investment Rigidities, and Risk Aversion in Information Security: A Utility-theoretic Approach. WEIS 2011.

Jaisingh, J. and J. Rees (2001). Value at risk: A methodology for information security risk assessment. Proceedings of the INFORMS Conference on Information Systems and Technology.

Jakoubi, Neubauer, et al. (2009). A Roadmap to Risk-Aware Business Process Management. IEEE Asia-Pacific Services Computing Conference, 2009. APSCC 2009. , IEEE.

Jakoubi, Tjoa, et al. (2009). A Survey of Scientific Approaches Considering the Integration of Security and Risk Aspects into Business Process Management. 2009 20th International Workshop on Database and Expert Systems Application, IEEE.

Jakoubi, S., G. Goluch, et al. (2008). Deriving Resource Requirements Applying Risk-Aware Business Process Modeling and Simulation. ECIS.

Jakoubi, S., S. Tjoa, et al. (2010). A Formal Approach Towards Risk-Aware Service Level Analysis and Planning. ARES'10 International Conference on Availability, Reliability, and Security, 2010. , IEEE.

Jakoubi, S., S. Tjoa, et al. (2007). Rope: A Methodology for Enabling the Risk-Aware Modelling and Simulation of Business Processes. ECIS.

Jaspersen, J. G. (2015). "Hypothetical Surveys and Experimental Studies of Insurance Demand: A Review." Journal of Risk and Insurance **83**(1): 217-255.

Johansson, E. (2005). Assessment of Enterprise Information Security, KTH, Royal Institute of Technology Stockholm, Sweden.

Johansson, E. (2005). Assessment of enterprise information security – how to make it credible and efficient.

Johnson, B., R. Böhme, et al. (2011). Security games with market insurance. International Conference on Decision and Game Theory for Security, Springer.

Johnson, B., J. Grossklags, et al. (2010). Are Security Experts Useful? Bayesian Nash Equilibria for Network Security Games with Limited Information. Computer Security – ESORICS 2010. D. Gritzalis, B. Preneel and M. Theoharidou, Springer Berlin / Heidelberg. **6345:** 588-606.

Johnson, E. J. and A. Tversky (1983). "Affect, generalization, and the perception of risk." Journal of Personality and Social Psychology **45**(1): 20-31.

Johnson, M. (2009). Data Hemorrhages in the Health-Care Sector. Financial Cryptography and Data Security. R. Dingledine and P. Golle, Springer Berlin / Heidelberg. **5628:** 71-89.

Johnson, P. (2002). Enterprise Software Sytem Integration: An Architectural Perspective.

Johnson, P., M. Ekstedt, et al. (2007). Using Enterprise Architecture for CIO Decision-Making: On the Importance of Theory. Second Annual Conference on Systems Engineering Research.

Johnson, P., M. E. Iacob, et al. (2013). Business Model Risk Analysis: Predicting the Probability of Business Network Profitability. Enterprise Interoperability: 5th International IFIP Working Conference, IWEI 2013, Enschede, The Netherlands, March 27-28, 2013. Proceedings. M. Sinderen, P. Oude Luttighuis, E. Folmer and S. Bosems. Berlin, Heidelberg, Springer Berlin Heidelberg**:** 118-130.

Johnson, P., E. Johansson, et al. (2007). A Tool for Enterprise Architecture Analysis. 11th IEEE International Enterprise Distributed Object Computing Conference, 2007. EDOC 2007. , IEEE.

Johnson, P., R. Lagerstrom, et al. (2013). IT Management with Enterprise Architecture.

Johnson, P., J. Ullberg, et al. (2013). P2amf: Predictive, probabilistic architecture modeling framework. International IFIP Working Conference on Enterprise Interoperability, Springer.

Jolls, C., C. R. Sunstein, et al. (1998). "A Behavioral Approach to Law and Economics." Stanford Law Review **July, 1998**.

Kagel, J. H. and A. E. Roth, Eds. (1995). The Handbook of Experimental Economics, Princeton University Press.

Kahn, C., M. , J. McAndrews, et al. (2000). A theory of transactions privacy. Working Paper 2000-22, Federal Reserve Bank of Atlanta.

Kahneman, D., J. L. Knetsch, et al. (1990). "Experimental Tests of the Endowment Effect and the Coase Theorem." The Journal of Political Economy **98**(6): 1325-1348

Kahneman, D. and A. Tversky (1972). "Subjective probability: A judgement of representativeness." Cognitive Psychology **3**: 430-454.

Kahneman, D. and A. Tversky (1984). "Choices, values, and frames." American psychologist **39**(4): 341.

Kahneman, D. a. T., Amos (1979). "Prospect Theory: An Analysis of Decision under Risk." Econometrica **4**(2).

Kairies-Schwarz, N., J. Kokot, et al. (2014). How Do Consumers Choose Health Insurance?-An Experiment on Heterogeneity in Attribute Tastes and Risk Preferences, Rheinisch-Westfälisches Institut für Wirtschaftsforschung.

Karni, E. (2009). "A mechanism for eliciting probabilities." Econometrica **77**(2): 603-606.

Keane, M. P. (2010). "Structural vs. atheoretic approaches to econometrics." Journal of Econometrics **156**(1): 3-20.

Keeney, R. L. (1982). "Decision Analysis: An Overview." Operations Research **30**(5): 803-838.

Keeney, R. L. a. R., H. (1993). Decisions with Multiple Objectives: Preferences and Value Tradeoffs. New York, Cambridge University Press.

Keren, G. and L. E. M. Gerritsen (1999). "On the robustness and possible accounts of ambiguity aversion." Acta Psychologica **103**(1-2): 149-172.

Kiesling, E., A. Ekelhart, et al. (2014). "Evolving Secure Information Systems through Attack Simulation."

Kiesling, E., C. Strauß, et al. (2012). "A multi-objective decision support framework for simulation-based security control selection."

Klinke, A. and O. Renn (2002). "A New Approach to Risk Evaluation and Management: Risk-Based, Precaution-Based, and Discourse-Based Strategies1." Risk Analysis **22**(6): 1071-1094.

Koenker, R. and G. Bassett Jr (1978). "Regression quantiles." Econometrica: journal of the Econometric Society: 33-50.

Koszegi, B. and M. Rabin (2008). "Revealed mistakes and revealed preferences." The foundations of positive and normative economics: a handbook: 125-154.

Krokhmal, P., M. Zabarankin, et al. (2011). "Modeling and optimization of risk." Surveys in Operations Research and Management Science **16**(2): 49-66.

Krokhmal, P. A. (2007). "Higher moment coherent risk measures." Quantitative Finance **7**(4): 373-387.

Kugler, T., E. E. Kausel, et al. (2010). "Are Groups more Rational than Individuals? A Review of Interactive Decision Making in Groups."

Kuhn, T. S. (1962). The Structure of Scientific Revolutions. Chicago, IL, University of Chicago Press.

Kunreuther, H. and G. Heal (2003). "Interdependent Security." Journal of Risk and Uncertainty **26**(2): 231-249.

Kunreuther, H. and E. Michel-Kerjan (2012). "Demand for multi-year insurance: experimental evidence." Working Manuscript. Center for Risk Management and Decision Processes. The Wharton School, University of Pennsylvania.

Kunreuther, H., N. Novemsky, et al. (2000). "Making Low Probabilities Useful." Journal of Risk and Uncertainty **23**(2): 103-120.

Kunreuther, H. and M. Pauly (2004). "Neglecting Disaster: Why Don't People Insure Against Large Losses?" Journal of Risk and Uncertainty **28**(1): 5-21.

Kunreuther, H. and M. Pauly (2015). Insurance Decision-Making For Rare Events: The Role Of Emotions. N. W. Paper.

Lagerström, R. and P. Johnson (2008). Using Architectural Models to Predict the Maintainability of Enterprise Systems. 12th European Conference on Software Maintenance and Reengineering, 2008. CSMR 2008. , IEEE.

Lampel, J., J. Shamsie, et al. (2009). "Experiencing the improbable: Rare events and organizational learning." Organization Science **20**(5): 835-845.

Larrañaga, P., H. Karshenas, et al. (2013). "A review on evolutionary algorithms in Bayesian network learning and inference tasks." Information Sciences **233**: 109-125.

Latour, B. (2005). Reassembling the Social: An Introduction to Actor-network-theory, Oxford University Press.

Laudon, K. (1993). Markets and Privacy. Information Systems Working Papers Series, NYU.

Laury, S. K., M. M. McInnes, et al. (2009). "Insurance decisions for low-probability losses." Journal of Risk and Uncertainty **39**(1): 17-44.

Law, J. (2009). Actor Network Theory and Material Semiotics, Wiley-Blackwell.

Lawson, L. L. and C. L. Lawson (2010). "Video Game-Based Methodology for Business Research." Simulation & Gaming **41**(3): 360-373.

Lessig, L. (2000). Code and Other Laws of Cyberspace http://codev2.cc/.

Lessig, L. (2002). "Privacy as Property." Social Research **69**(1): 247-269.

Levy, H. and M. Levy (2002). "Experimental test of the prospect theory value function: A stochastic dominance approach." Organizational Behavior and Human Decision Processes **89**(2): 1058-1081.

Levy, H. and M. Levy (2009). "The safety first expected utility model: Experimental evidence and economic implications." Journal of Banking & Finance **33**(8): 1494-1506.

Levy, H. and M. Sarnat (1972). "Safety first—an expected utility principle." Journal of Financial and Quantitative Analysis **7**(03): 1829-1834.

Li, X., P. Parker, et al. (2011). "A Stochastic Model for Quantitative Security Analyses of Networked Systems." Dependable and Secure Computing, IEEE Transactions on **8**(1): 28-43.

Lindlof, T. R. and B. C. Taylor (2002). Qualitative Communication Research Methods. Thousand Oaks, CA, Sage.

Litman, J. (2000). "Information Privacy/Information Property." Stanford Law Review
**52**(5).

Littlewood, B., S. Brocklehurst, et al. (1993). "Towards Operational Measures of Computer Security." Journal of Computer Security **2**(2): 211-229.

Liu, Y. and H. Man (2005). Network Vulnerability Assessment using Bayesian Networks. Defense and Security, International Society for Optics and Photonics.

Loewenstein, G. and S. Issacharoff (1994). "Source dependence in the valuation of objects." Journal of Behavioral Decision Making **7**(3): 157-168.

Löf, F., J. Stomberg, et al. (2010). "An Approach to Network Security Assessment based on Probalistic Relational Models."

Longstaff, T. A., C. Chittister, et al. (2000). "Are we forgetting the risks of information technology?" Computer **33**(12): 43-51.

Loomes, G., C. Starmer, et al. (1991). "Observing violations of transitivity by experimental methods." Econometrica: Journal of the Econometric Society: 425-439.

Loomes, G. and R. Sugden (1982). "Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty." The Economic Journal **92**(368): 805-824.

Lopes, L. L. (1982). "Doing the impossible: A note on induction and the experience of randomness." Journal of Experimental Psychology: Learning, Memory, and Cognition **8**(6): 626.

Lopes, L. L. and G. C. Oden (1999). "The role of aspiration level in risky choice: a comparison of cumulative prospect theory and SP/A theory." J. Math. Psychol. **43**(2): 286-313.

Machina, M. J. (1983). Generalized Expected Utility Analysis and the Nature of Observed Violations of the Independence Axiom. Foundations of Utility and Risk Theory with Applications. B. P. Stigum and F. Wenstøp. Dordrecht, Springer Netherlands**:** 263-293.

Machina, M. J. (1992). Choice under uncertainty: Problems solved and unsolved. Foundations of Insurance Economics, Springer**:** 49-82.

Machina, M. J. and D. Schmeidler (1992). "A More Robust Definition of Subjective Probability." Econometrica **60**(4): 745-780

Malcolm, D. G., J. H. Roseboom, et al. (1959). "Application of a technique for research and development program evaluation." Operations research **7**(5): 646-669.

Markowitz, H. (1952). "Portfolio selection." The journal of finance **7**(1): 77-91.

Matheson, J. E. and R. L. Winkler (1976). "Scoring rules for continuous probability distributions." Management science **22**(10): 1087-1096.

McKee, M., R. P. Berrens, et al. (2004). "Using Experimental Economics to Examine Wildfire Insurance and Averting Decisions in the Wildland–Urban Interface." Society & Natural Resources **17**(6): 491-507.

McNair, S. and A. Feeney (2014). "When does information about causal structure improve statistical reasoning?" The Quarterly Journal of Experimental Psychology **67**(4): 625-645.

Mehra, R. and E. C. Prescott (1985). "The equity premium: A puzzle." Journal of Monetary Economics **15**(2): 145-161.

Mendez, F. and M. Mendez (2010). "Comparing Privacy Regimes: Federal Theory and the Politics of Privacy Regulation in the European Union and the United States." Publius: The Journal of Federalism **40**(4): 617-645.

Meyer, J. W. and B. Rowan (1977). "Institutionalized Organizations: Formal Structure as Myth and Ceremony." The American Journal of Sociology **83**(2).

Miller, A. R. and C. Tucker (2009). "Privacy Protection and Technology Diffusion: The Case of Electronic Medical Records." Management Science **55**(7): 1077-1093.

Mullainathan, S. T., Richard H. (2000). "Behavioral Economics." NBER Working Papers.

Myerson, R. B. (1982). "Optimal coordination mechanisms in generalized principal–agent problems." Journal of Mathematical Economics **10**(1): 67-81.

Närman, P., U. Franke, et al. (2012). "Enterprise architecture availability analysis using fault trees and stakeholder interviews."

Närman, P., P. Johnson, et al. (2007). "Enterprise Architecture: A Framework Supporting System Quality Analysis."

Nawrocki, D. N. (1999). "A brief history of downside risk measures." The Journal of Investing **8**(3): 9-25.

Neapolitan, R. E. (2004). Learning bayesian networks, Prentice Hall Upper Saddle River.

Neches, R., R. Fikes, et al. (1991). "Enabling Technology for Knowledge Sharing." AI Magazine **Winter** 36-56.

Neill, H. R., R. G. Cummings, et al. (1994). "Hypothetical Surveys and Real Economic Commitments." Land Economics **70**(2): 145-154.

Neubauer, T., A. Ekelhart, et al. (2008). Interactive Selection of ISO 27001 Controls under Multiple Objectives. IFIP International Information Security Conference, Springer.

Neubauer, T. and C. Hartl (2009). On the singularity of valuating IT security investments. Eighth IEEE/ACIS International Conference on Computer and Information Science, 2009. ICIS 2009. , IEEE.

Neubauer, T., M. Klemen, et al. (2005). Business Process-based Valuation of IT-Security, ACM.

Neubauer, T. and C. Stummer (2007). Extending Business Process Management to Determine Efficient IT Investments. Proceedings of the 2007 ACM Symposium on Applied Computing, ACM.

Neubauer, T., C. Stummer, et al. (2006). Workshop-based Multiobjective Security Safeguard Selection. First International Conference on Availability, Reliability and Security (ARES'06), IEEE.

Nicol, D. M., W. H. Sanders, et al. (2004). "Model-based evaluation: from dependability to security." Dependable and Secure Computing, IEEE Transactions on **1**(1): 48-65.

Nicol, D. M., W. H. Sanders, et al. (2004). "Model-Based Evaluation: From Dependability to Security." IEEE Trans. Dependable Secur. Comput. **1**(1): 48-65.

NIST (October 1995). An Introduction to Computer Security - The NIST Handbook. Technical report, National Institute of Technology. **Special Publication 800-12.**

Noam, E. M. (2002) "Privacy in Telecommunications: Markets, Rights and Regulations, Part 3: Markets in Privacy." Technology Futures, Inc. **95**.

Offerman, T., J. Sonnemans, et al. (2009). "A truth serum for non-bayesians: Correcting proper scoring rules for risk attitudes." The Review of Economic Studies **76**(4): 1461-1489.

Ozcelik, Y. and J. Rees (2005). "A New Approach for Information Security Risk Assessment: Value at Risk." SSRN eLibrary.

Perkusich, M., G. Soares, et al. (2015). "A procedure to detect problems of processes in software development projects using Bayesian networks." Expert Systems with Applications **42**(1).

Pickands III, J. (1975). "Statistical inference using extreme order statistics." the Annals of Statistics: 119-131.

Pigou, A. C. (1932). The Economics of Welfare.

Poindexter, J., J. Earp, et al. (2006). "An experimental economics approach toward quantifying online privacy choices." Information Systems Frontiers **8**(5): 363-374.

Posner, R. A. (1978). "An Economic Theory of Privacy." Georgia Law Review **May/June, 1978**.

Posner, R. A. (1981). "The Economics of Privacy." The American Economic Review **71**(2).

Prado, J. L. P. (2009). Strategic Business and IT Alignment Assessment: A modelling approach associated with Enterprise Architecture.

Pratt, J. W. (1964). "Risk Aversion in the Small and in the Large." Econometrica **32**(1/2): 122-136.

Pratt, J. W. (1976). "F. Y. Edgeworth and R. A. Fisher on the Efficiency of Maximum Likelihood Estimation." The Annals of Statistics **4**(3): 501-514.

Prins, J. E. J. (2006). "Property and Privacy: European Perspectives and the Commodification of Our Identity." Information Law Series, Vol. 16, pp. 223-257, 2006.

Quiggin, J. (1982). "A theory of anticipated utility." Journal of Economic Behavior & Organization **3**(4): 323-343.

Quiggin, J. (2002). "Risk and Self-Protection: A State-Contingent View." Journal of Risk and Uncertainty **25**(2): 133-145.

Quinn, S. (2011). Maximum Likelihood and Structural Models in Microeconometrics, University of Oxford.

Rabin, M. and R. H. Thaler (2001). "Anomalies: Risk Aversion." The Journal of Economic Perspectives **15**(1): 219-232.

Read, D. (2005). "Monetary incentives, what are they good for?" Journal of Economic Methodology **12**(2): 265-276.

Richards, N. M. (2006). "The Information Privacy Law Project." Georgetown Law Journal, Vol. 94, p. 1087, 2006.

Richards, N. M. and D. J. Solove (2007). "Privacy's Other Path: Recovering the Law of Confidentiality." Georgetown Law Journal **96**.

Rivenbark, D. (2010). Uncertainty, Identification and Privacy: Experiments in Individual Decision-Making. Economics. Orlando, University of Central Florida. **Doctor of Philosophy:** 214.

Rivenbark, D. R. (2011). Experimentally Elicited Beliefs Explain Privacy Behavior, University of Central Florida.

Robey, D. and M.-C. Boudreau (1999). "Accounting for the Contradictory Organizational Consequences of Information Technology: Theoretical Directions and Methodological Implications." Information Systems Research **10**(2): 167-185.

Rockafellar, R. T. (2007). "Coherent approaches to risk in optimization under uncertainty." Tutorials in operations research **3**: 38-61.

Rockafellar, R. T. and J. O. Royset (2010). "On buffered failure probability in design and optimization of structures." Reliability Engineering & System Safety **95**(5): 499-510.

Rockafellar, R. T. and S. Uryasev (2000). "Optimization of conditional value-at-risk." Journal of risk **2**: 21-42.

Rockafellar, R. T. and S. Uryasev (2013). "The fundamental risk quadrangle in risk management, optimization and statistical estimation." Surveys in Operations Research and Management Science **18**(1): 33-53.

Rockafellar, R. T., S. Uryasev, et al. (2006). "Generalized deviations in risk analysis." Finance and Stochastics **10**(1): 51-74.

Rockafellar, R. T., S. P. Uryasev, et al. (2002). "Deviation measures in risk analysis and optimization." University of Florida, Department of Industrial & Systems Engineering Working Paper(2002-7).

Rockafellar, T. and S. Uryasev (2002). "Conditional Value-at-Risk for General Loss Distributions (2001)." Journal of Banking and Finance **26**(7).

Rodwin, M. A. (2010). "Patient Data: Property, Privacy & the Public Interest." American Journal of Law and Medicine **December, 2010**.

Romanosky, S. (2010). Data Breaches and Identity Theft: When is Mandatory Disclosure Optimal? The Ninth Workshop on the Economics of Information Security (WEIS 2010).

Romanosky, S. and A. Acquisti (2009). "Privacy Costs and Personal Data Protection: Economic and Legal Perspectives." Berkeley Technology Law Journal, Vol. 24, No. 3, 2009.

Romanosky, S., R. Telang, et al. (2008). "Do Data Breach Disclosure Laws Reduce Identity Theft?" Forthcoming in the Journal of Policy Analysis and Management, 2011.

Rosenthal, R. (1989). "A bounded-rationality approach to the study of noncooperative games." International Journal of Game Theory **18**(3): 273-292.

Rotenberg, M. (2001). "Fair Information Practices and the Architecture of Privacy (What Larry Doesn’t Get)." Stan. Tech. L. Rev. **1**.

Roy, A. D. (1952). "Safety first and the holding of assets." Econometrica: Journal of the Econometric Society: 431-449.

Sage, W. M. (1999). "Regulating Through Information: Disclosure Laws and American Health Care." Columbia Law Review **99**(7).

Sage, W. M. (2008). "Relational Duties, Regulatory Duties, and the Widening Gap between Individual Health Law and Collective Health Policy." Georgetown Law Journal, Vol. 96, No. 2, 2008.

Saha, A. (1993). "Expo-power utility: A 'flexible'form for absolute and relative risk aversion." American Journal of Agricultural Economics **75**(4): 905-913.

Samuelson, P. (2000). "Privacy As Intellectual Property?" Stanford Law Review **52**(5).

Savage, L. (1954). The foundations of statistics. New York, Wiley.

Savage, L. J. (1971). "Elicitation of Personal Probabilities and Expectations." Journal of the American Statistical Association **66**(336): 783-801.

Sawik, T. (2013). "Selection of optimal countermeasure portfolio in IT security planning." Decision Support Systems **55**(1): 156-164.

Schechter, S. E. (2004). Computer Security Strength & Risk: A Quantitative Approach. The Division of Engineering and Applied Sciences. Cambridge, Massachusetts, Harvard University. **PhD. Dissertation**.

Schechter, S. E. (2005). "Toward Econometric Models of the Security Risk from Remote Attack." IEEE Security and Privacy **3**(1): 40-44.

Schmittling, R. and A. Munns (2010). "Performing a security risk assessment." ISACA Journal **1**: 18.

Schriber, T. J., D. T. Brunner, et al. (2013). Inside discrete-event simulation software: how it works and why it matters. Simulation Conference (WSC), 2013 Winter, IEEE.

Schroeder, N. J. (2005). Using Prospect Theory to investigate Decision-Making Bias Within an Information Security Context Department of Systems and Engineering Management, Department of the Air Force Air University, Air Force Institute of Technology. **MSc Thesis**.

Schwartz, P. M. (1994). "European Data Protection Law and Restrictions on International Data Flows; ." Iowa L. Rev. **80**(471).

Schwartz, P. M. (1997). "Privacy and the Economics of Health Care Information." Texas Law Review **76**(1).

Schwartz, P. M. (2000). "Beyond Lessig's Code for Internet Privacy: Cyberspace Filters, Privacy Control and Fair Information Practices." Wisconsin Law Review, 2000.

Schwartz, P. M. (2000). Privacy and Democracy in Cyberspace, University of California, Berkeley - School of Law.

Schwartz, P. M. (2004). "Property, Privacy, and Personal Data." Harvard Law Review **117**(7).

Schwartz, P. M. (2010) "Managing Global Data Privacy." IAPP, The Privacy Advisor, January 2010.

Sen, S. (2010). Behavioural Response to Endogenous Risk in the Laboratory. Department of Economics. Orlando, Florida, University of Central Florida. **Doctor of Philosophy:** 245.

Shapiro, C., & Varian, H. R. (1997). U.S. Government information policy. Highlands Forum. Department of Defense Washington, DC., Office of the Assistant Secretary of Defense.

Shogren, J. F. and T. D. Crocker (1991). "Risk, self-protection, and ex ante economic value." Journal of Environmental Economics and Management **20**(1): 1-15.

Shogren, J. F. and T. D. Crocker (1994). "Rational risk valuation given sequential reduction opportunities." Economics Letters **44**(3): 241-248.

Simon, H. A. (1955). "A Behavioral Model of Rational Choice." The Quarterly Journal of Economics **69**(1): 99-118.

Sklavos, N. a. S., Panagiotis (2006). "Economic Models and Approaches in Information Security for Computer Networks." International Journal of Network Security **2**(1).

Smith, C. A. (1961). "Consistency in statistical inference and decision." Journal of the Royal Statistical Society. Series B (Methodological): 1-37.

Smith, V. L. (1976). "Experimental Economics: Induced Value Theory." The American Economic Review **66**(2): 274-279.

Smith, V. L. (1982). "Microeconomic Systems as an Experimental Science." The American Economic Review **72**(5): 923-955

Smith, V. L. (2010). "Theory and experiment: What are the questions?" Journal of Economic Behavior & Organization **73**(1): 3-15.

Smith, V. L. (2010). "Theory and experiment: What are the questions?" Journal of Economic Behavior &amp; Organization **73**(1): 3-15.

Snow, A. (2011). "Ambiguity aversion and the propensities for self-insurance and self-protection." Journal of Risk and Uncertainty **42**(1): 27-43.

Sommestad, T. (2012). A framework and theory for cyber security assessments, KTH, Royal Institute of Technology Stockholm, Sweden.

Sommestad, T., M. Ekstedt, et al. (2013). "The Cyber Security Modeling Language: A Tool for Assessing the Vulnerability of Enterprise System Architectures." Systems Journal, IEEE **7**(3): 363-373.

Sommestad, T., M. Ekstedt, et al. (2008). Combining Defense Graphs and Enterprise Architecture Models for Security Analysis.

Sommestad, T., M. Ekstedt, et al. (2009). Cyber Security Risks Assessment with Bayesian Defense Graphs and Architectural Models. 42nd Hawaii International Conference on System Sciences, 2009. HICSS'09. , IEEE.

Sommestad, T., M. Ekstedt, et al. (2010). "A probabilistic relational model for security risk analysis." Computers & Security **29**(6): 659-679.

Sommestad, T., M. Ekstedt, et al. (2010). "A case study applying the Cyber Security Modeling Language."

Sommestad, T., H. Holm, et al. (2011). "Threats and vulnerabilities, final report."

Sommestad, T., E. Mathias, et al. (2013). "The Cyber Security Modeling Language: A Tool for Assessing the Vulnerability of Enterprise System Architectures."

Soo Hoo, K. (2000). How Much Is Enough? A Risk-Management Approach to Computer Security. Department of Management Science and Engineering, Stanford University. **PhD. Dissertation**.

Sorescu, S. and A. Subrahmanyam (2006). "The Cross Section of Analyst Recommendations." The Journal of Financial and Quantitative Analysis **41**(1): 139-168.

Stigler, G. J. (1980). "An Introduction to Privacy in Economics and Politics." The Journal of Legal Studies **9**(4): 623-644

Studer, Benjamins, et al. (1998). "Knowledge Engineering: Principles and Methods.
." Data and Knowledge Engineering **25** 161-197.

Swartout, B., R. Patil, et al. (1997). "Toward Distributed Use of Large-Scale Ontologies." Ontological Engineerin **Spring Symposium Series**: 138-148.

Talberth, J., R. P. Berrens, et al. (2006). "Averting and Insurance Decisions in te hWildland-Urban Interface: Implications of Survey and Experimenatl Data for Wildlife Risk Reduction Policy." Contemporary Economic Policy **24**(2): 203-223.

Taylor, C., A. Krings, et al. (2002). Risk analysis and probabilistic survivability assessment (RAPSA): An assessment approach for power substation hardening. Proc. ACM Workshop on Scientific Aspects of Cyber Terrorism,(SACT), Washington DC.

Thaler, R. (1981). "Some empirical evidence on dynamic inconsistency." Economics Letters **8**(3): 201-207.

Thaler, R. H. and E. J. Johnson (1990). "Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice." Management Science, Vol. 36, No. 6, pp. 643-660, 1990.

Thomas, R. C. (2009). Total cost of security: a method for managing risks and incentives across the extended enterprise. Proceedings of the 5th Annual Workshop on Cyber Security and Information Intelligence Research: Cyber Security and Information Intelligence Challenges and Strategies, Oak Ridge, Tennessee.

Thomas, R. C., M. Antkiewicz, et al. (2013). "How bad is it?–a branching activity model to estimate the impact of information security breaches." A Branching Activity Model to Estimate the Impact of Information Security Breaches (March 11, 2013).

Thomas, R. C., M. Antkiewicz, et al. (2013). How Bad is it? - A Branching Activity Model to Estimate the Impact of Information Security Breaches.

Tjoa, S., S. Jakoubi, et al. (2011). "A Formal Approach Enabling Risk-Aware Business Process Modeling and Simulation." Services Computing, IEEE Transactions on **4**(2): 153-166.

Tjoa, S., S. Jakoubi, et al. (2010). Planning Dynamic Activity and Resource Allocations Using a Risk-Aware Business Process Management Approach. ARES'10 International Conference on Availability, Reliability, and Security, 2010. , IEEE.

Tjoa, S., S. Jakoubi, et al. (2008). Enhancing Business Impact Analysis and Risk Assessment applying a Risk-Aware Business Process Modeling and Simulation Methodolog ARES 08. Third International Conference on Availability, Reliability and Security, 2008. , IEEE.

Tversky, A. and D. Kahneman (1974). "Judgment under Uncertainty: Heuristics and Biases." Science **185(**4157): 1124-1131.

Tversky, A. and D. Kahneman (1992). "Advances in prospect theory: Cumulative representation of uncertainty." Journal of Risk and Uncertainty **5**(4): 297-323.

Tversky, A. and P. Wakker (1995). "Risk attitudes and decision weights." Econometrica: Journal of the Econometric Society: 1255-1280.

Uryasev, S. (2000). Introduction to the Theory of Probabilistic Functions and Percentiles (Value-at-Risk). Probabilistic Constrained Optimization: Methodology and Applications. S. P. Uryasev. Boston, MA, Springer US**:** 1-25.

Varian, H. R. (1996). "Economic Aspects of Personal Privacy." Privacy and Self-Regulation in the Information Age: 101-109.

Verendel, V. (2008). A Prospect Theory Approach to Security Goteborg, Sweden, Department of Computer Science and Engineering, Chalmers University of Technology, Goteborg University.

Verendel, V. (2009). Quantified security is a weak hypothesis: a critical survey of results and assumptions. Proceedings of the 2009 Workshop on New Security Paradigms. Oxford, United Kingdom, ACM**:** 37-50.

Vickrey, W. (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders." The Journal of Finance **16**(1): 8-37.

von Neumann, J. a. M., O. (1947). Theory of Games and Economic Behavior. Princeton NJ, Princeton University Press.

Vose, D. (2008). Risk analysis: a quantitative guide, John Wiley & Sons.

Vuong, Q. H. (1989). "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses." Econometrica **57**(2): 307-333.

Wakker, P., I. Erev, et al. (1994). "Comonotonic independence: The critical test between classical and rank-dependent utility theories." Journal of Risk and Uncertainty **9**(3): 195-230.

Wang, J., A. Chaudhury, et al. (2008). "A Value-at-Risk Approach to Information Security Investment." Info. Sys. Research **19**(1): 106-120.

Warren and Brandeis (1890). "The Right to Privacy." Harvard Law Review. **4**(5).

WEF (2015). "Partnering for Cyber Resilience: Towards the Quantification of Cyber Threats."

Welch, R. L. (1996). Real Time Estimation of Bayesian Networks. Proceedings of the Twelfth international conference on Uncertainty in artificial intelligence, Morgan Kaufmann Publishers Inc.

Westin, A. F. (1967). Privacy and Freedom, Atheneum Publishers.

Weston, S. L. (2014). Envisioning the Improbable: Judgment and Strategy in Heavy-Tailed Contexts. Academy of Management Proceedings, Academy of Management.

Wilcox, N. T. (2008). Stochastic models for binary discrete choice under risk: a critical primer and econometric comparison. Risk Aversion in Experiments (Research in Experimental Economics). G. W. H. James C. Cox, Emerald Group Publishing Limited. **12:** 197-292.

Wilcox, N. T. (2010). A comparison of three probabilistic models of binary discrete choice under risk, working paper.

Wilcox, N. T. (2010). Predicting Risky Choices Out-of-Context: A Monte Carlo Study. University of Houston Working Paper.

Winkelvos, T., C. Rudolph, et al. (2011). A property based security risk analysis through weighted simulation. Information Security South Africa (ISSA), 2011.

Wolter, K. and P. Reinecke (2010). Performance and Security Tradeoff
Formal Methods for Quantitative Aspects of Programming Languages. A. Aldini, M. Bernardo, A. Di Pierro and H. Wiklicky, Springer Berlin / Heidelberg. **6154:** 135-167.

Yaari, M. E. (1987). "The Dual Theory of Choice under Risk." Econometrica **55**(1): 95-115

Yin, R. K. (2013). Case study research: Design and methods, Sage publications.

Zeisberger, S. (2013). "The importance of the probability of losing in repeated decisions."

Zeisberger, S. (2014). "Do Investors Care Explicitly about Loss Probabilities?" Available at SSRN 2169394.

Zellner, A. (1986). "Bayesian Estimation and Prediction Using Asymmetric Loss Functions." Journal of the American Statistical Association **81**(394): 446-451.

Zhan, J. and V. Rajamani (2008). "The Economics of Privacy: People, Policy and Technology." International Journal of Security and its Applications **2**(3).

Zucker, L. G. (1987). "Institutional Theories of Organization." Annual Review of Sociology **13**.

# Appendix 1 - Game Instructions

## Game 1: Asset Integration

### Game Interface:

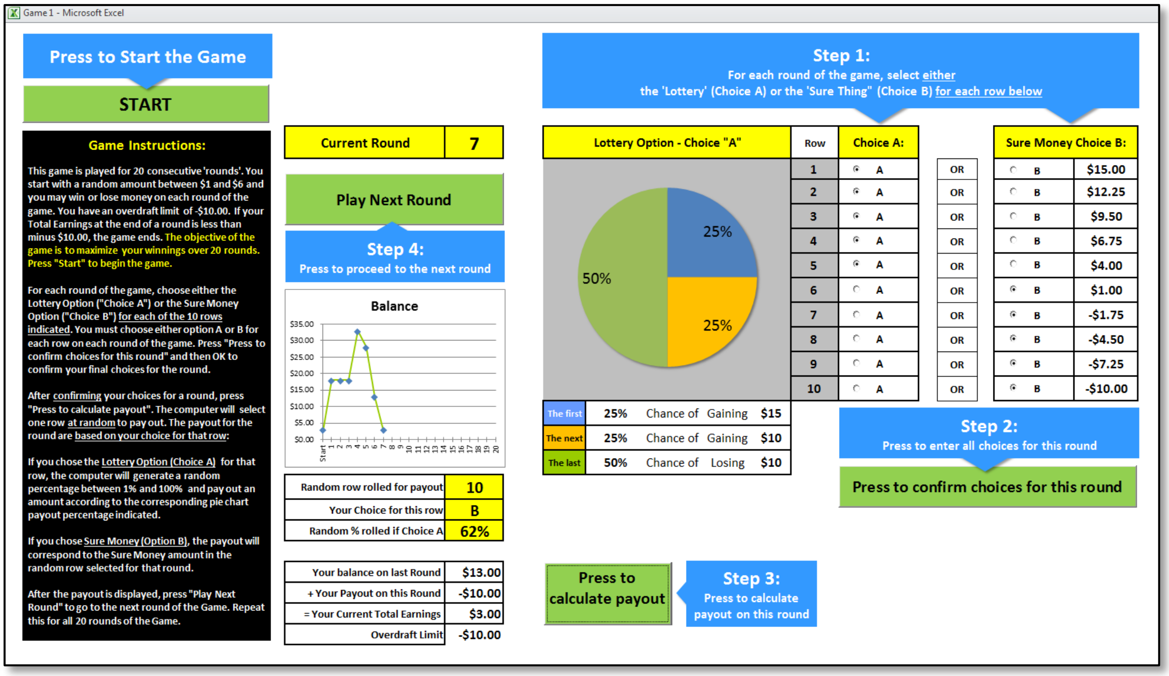
Game 1 - Microsoft Excel
Press to Start the Game
START
Play Next Round
Step 4:
Press to proceed to the next round
Balance
$35.00
$30.00
$25.00
$20.00
$15.00
$10.00
$5.00
$0.00
Start
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
Step 1:
For each round of the game, select either the 'Lottery' (Choice A) or the 'Sure Thing" (Choice B) for each row below
50%
25%
25%
Step 2:
Press to enter all choices for this round
Press to confirm choices for this round
Press to calculate payout
Step 3:
Press to calculate payout on this round

Game Instructions:

This game is played for 20 consecutive 'rounds'. You start with a random amount between $1 and $6 and you may win or lose money on each round of the game. You have an overdraft limit of -$10.00. If your Total Earnings at the end of a round is less than minus $10.00, the game ends. The objective of the game is to maximize your winnings over 20 rounds. Press "Start" to begin the game.

For each round of the game, choose either the Lottery Option ("Choice A") or the Sure Money Option ("Choice B") for each of the 10 rows indicated. You must choose either option A or B for each row on each round of the game. Press "Press to confirm choices for this round" and then OK to confirm your final choices for the round.

After confirming your choices for a round, press "Press to calculate payout". The computer will select one row at random to pay out. The payout for the round are based on your choice for that row:

If you chose the Lottery Option (Choice A) for that row, the computer will generate a random percentage between 1% and 100% and pay out an amount according to the corresponding pie chart payout percentage indicated.

If you chose Sure Money (Option B), the payout will correspond to the Sure Money amount in the random row selected for that round.

After the payout is displayed, press "Play Next Round" to go to the next round of the Game. Repeat this for all 20 rounds of the Game.

| | |
|---|---|
| Current Round | 7 |

| | |
|---|---|
| Random row rolled for payout | 10 |
| Your Choice for this row | B |
| Random % rolled if Choice A | 62% |

| | |
|---|---|
| Your balance on last Round | $13.00 |
| + Your Payout on this Round | -$10.00 |
| = Your Current Total Earnings | $3.00 |
| Overdraft Limit | -$10.00 |

| Lottery Option - Choice "A" | Row | Choice A: | | Sure Money Choice B: | |
|---|---|---|---|---|---|
| | 1 | ● A | OR | ○ B | $15.00 |
| | 2 | ● A | OR | ○ B | $12.25 |
| | 3 | ● A | OR | ○ B | $9.50 |
| | 4 | ● A | OR | ○ B | $6.75 |
| | 5 | ● A | OR | ○ B | $4.00 |
| | 6 | ○ A | OR | ● B | $1.00 |
| | 7 | ○ A | OR | ● B | -$1.75 |
| | 8 | ○ A | OR | ● B | -$4.50 |
| | 9 | ○ A | OR | ● B | -$7.25 |
| | 10 | ○ A | OR | ● B | -$10.00 |

| | | | | |
|---|---|---|---|---|
| The first | 25% | Chance of | Gaining | $15 |
| The next | 25% | Chance of | Gaining | $10 |
| The last | 50% | Chance of | Losing | $10 |

### Instructions:

“This game is played for 20 consecutive 'rounds'. You start with a random amount between $1 and $6 and you may win or lose money on each round of the game. You have an overdraft limit of -$10.00. If your Total Earnings at the end of a round is less than minus $10.00, the game ends. The objective of the game is to maximize your winnings over 20 rounds. Press "Start" to begin the game.

For each round of the game, choose either the Lottery Option ("Choice A") or the Sure Money Option ("Choice B") for each of the 10 rows indicated. You must choose either option A or B for each row on each round of the game. Press "Press to confirm choices for this round" and then OK to confirm your final choices for the round.

After confirming your choices for a round, press "Press to calculate payout". The computer will select one row at random to pay out. The payout for the round are based on your choice for that row:

If you chose the Lottery Option (Choice A) for that row, the computer will generate a random percentage between 1% and 100% and pay out an amount according to the corresponding pie chart payout percentage indicated.

If you chose Sure Money (Option B), the payout will correspond to the Sure Money amount in the random row selected for that round.

After the payout is displayed, press "Play Next Round" to go to the next round of the Game. Repeat this for all 20 rounds of the Game.”

## Game 2: Subjective Bayesian Updating and Strength / Weight of Evidence

### 1 - Example Simulations Screen:

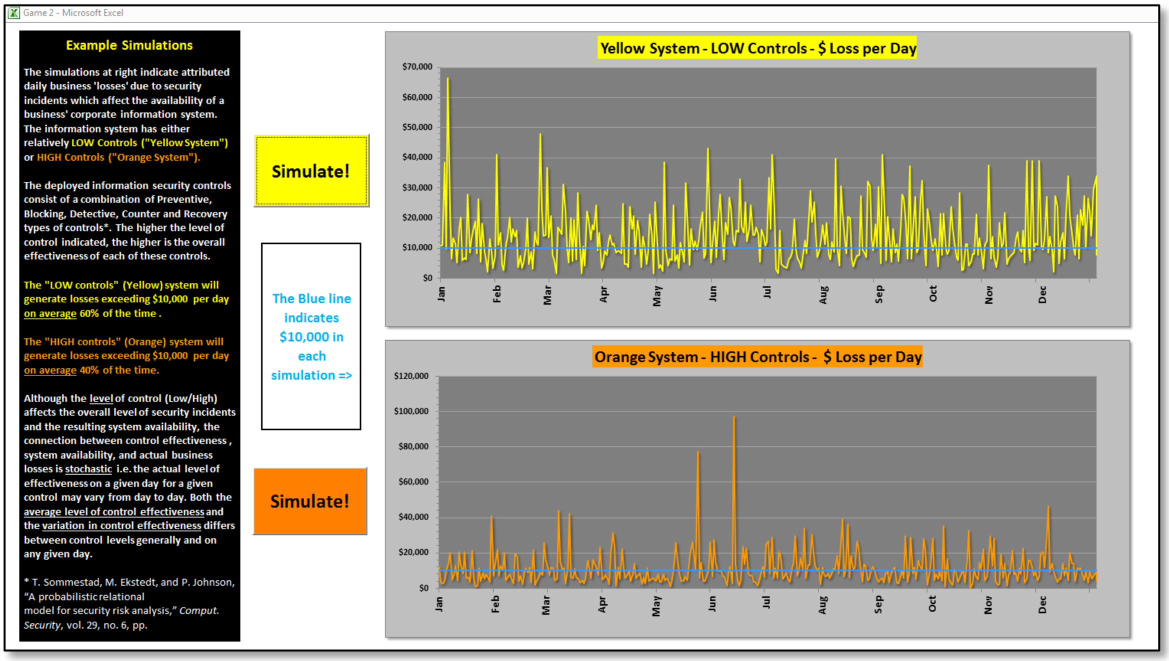

Example Simulations

The simulations at right indicate attributed daily business 'losses' due to security incidents which affect the availability of a business' corporate information system. The information system has either relatively LOW Controls ("Yellow System") or HIGH Controls ("Orange System").

The deployed information security controls consist of a combination of Preventive, Blocking, Detective, Counter and Recovery types of controls*. The higher the level of control indicated, the higher is the overall effectiveness of each of these controls.

The "LOW controls" (Yellow) system will generate losses exceeding $10,000 per day on average 60% of the time .

The "HIGH controls" (Orange) system will generate losses exceeding $10,000 per day on average 40% of the time.

Although the level of control (Low/High) affects the overall level of security incidents and the resulting system availability, the connection between control effectiveness , system availability, and actual business losses is stochastic i.e. the actual level of effectiveness on a given day for a given control may vary from day to day. Both the average level of control effectiveness and the variation in control effectiveness differs between control levels generally and on any given day.

* T. Sommestad, M. Ekstedt, and P. Johnson, "A probabilistic relational model for security risk analysis," *Comput. Security*, vol. 29, no. 6, pp.

The Blue line indicates $10,000 in each simulation =>

### Instructions:

"The simulations at right indicate attributed daily business 'losses' due to security incidents which affect the availability of a business' corporate information system. The information system has either relatively LOW Controls ("Yellow System") or HIGH Controls ("Orange System").

The deployed information security controls consist of a combination of Preventive, Blocking, Detective, Counter and Recovery types of controls*. The higher the level of control indicated, the higher is the overall effectiveness of each of these controls.

The "LOW controls" (Yellow) system will generate losses exceeding $10,000 per day on average 60% of the time .

The "HIGH controls" (Orange) system will generate losses exceeding $10,000 per day on average 40% of the time.

Although the level of control (Low/High) affects the overall level of security incidents and the resulting system availability, the connection between control effectiveness , system availability, and actual business losses is stochastic i.e. the actual level of effectiveness on a given day for a given control may vary from day to day. Both the average level of control effectiveness and the variation in control effectiveness differs between control levels generally and on any given day.

* T. Sommestad, M. Ekstedt, and P. Johnson, "A probabilistic relational model for security risk analysis," Comput. Security, vol. 29, no. 6, pp. 659–679, Mar. 2010."

## 2 – Game Interface:

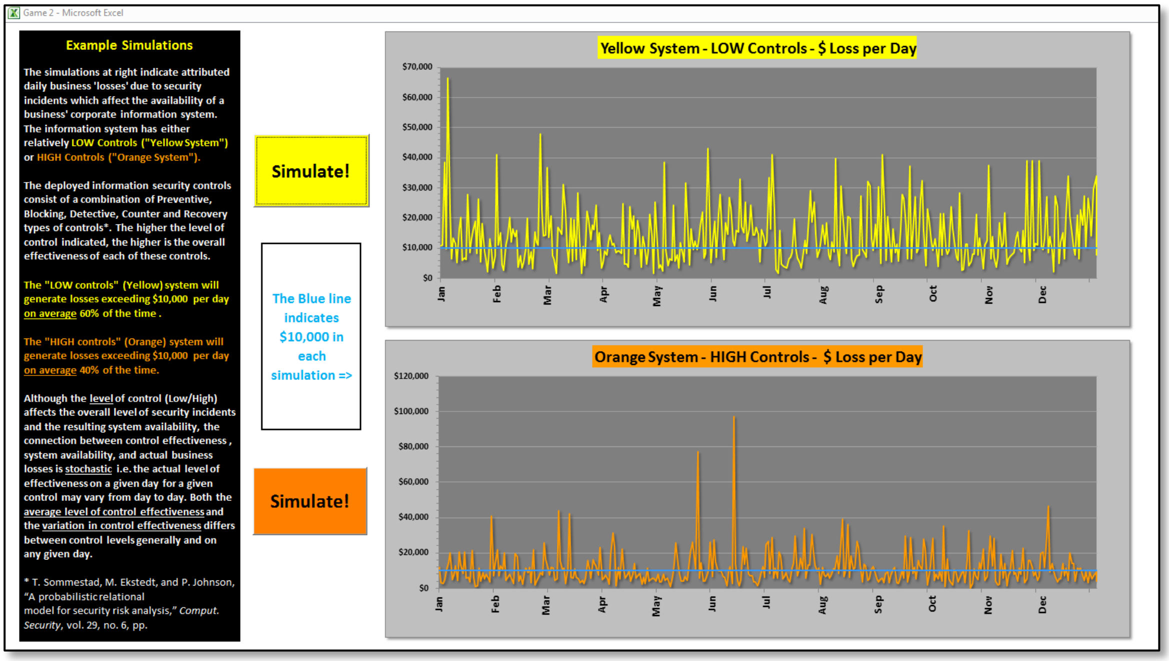

Game 2 - Microsoft Excel

**Example Simulations**

The simulations at right indicate attributed daily business 'losses' due to security incidents which affect the availability of a business' corporate information system. The information system has either relatively LOW Controls ("Yellow System") or HIGH Controls ("Orange System").

The deployed information security controls consist of a combination of Preventive, Blocking, Detective, Counter and Recovery types of controls*. The higher the level of control indicated, the higher is the overall effectiveness of each of these controls.

The "LOW controls" (Yellow) system will generate losses exceeding $10,000 per day <u>on average</u> 60% of the time .

The "HIGH controls" (Orange) system will generate losses exceeding $10,000 per day <u>on average</u> 40% of the time.

Although the <u>level</u> of control (Low/High) affects the overall level of security incidents and the resulting system availability, the connection between control effectiveness , system availability, and actual business losses is <u>stochastic</u> i.e. the actual level of effectiveness on a given day for a given control may vary from day to day. Both the <u>average level of control effectiveness</u> and the <u>variation in control effectiveness</u> differs between control levels generally and on any given day.

* T. Sommestad, M. Ekstedt, and P. Johnson, “A probabilistic relational model for security risk analysis,” *Comput. Security*, vol. 29, no. 6, pp.

The Blue line indicates $10,000 in each simulation =>

## Instructions:

“There are 30 rounds in this game - one round per tab. Each round is played separately. When you have played this round/tab, move on to the next round/tab. The objective of the game is to maximize the payout from a randomly selected bookie.

The payout for a round will depend on your choice with that bookie for that round and from which system the losses were actually generated. If you guess the right system, the maximum payout per round is $60. If you guess the wrong system, the payout is always $0.

You know the following two facts:

The 'LOW controls' (Yellow) system will generate losses exceeding $10,000 per day on average 60% of the time .

The 'HIGH controls' (Orange) system will generate losses exceeding $10,000 per day on average 40% of the time.

On each round, press "Simulate". Now place a $3 bet with each bookie on which system you believe generated the losses by selecting either "I bet it's a Yellow System" or "I bet it's an Orange System".  You must place a bet with each Bookie.

Once all of your bets are placed for the round, press "Press to enter all bets" and "OK". Then press "Calculate a Payout" to view the payout for that round. When you have played this round/tab, move on to the next round/tab.”

## Game 3: Subjective Bayesian Updating and Strength / Weight of Evidence

### 1 – Introductory Screen:

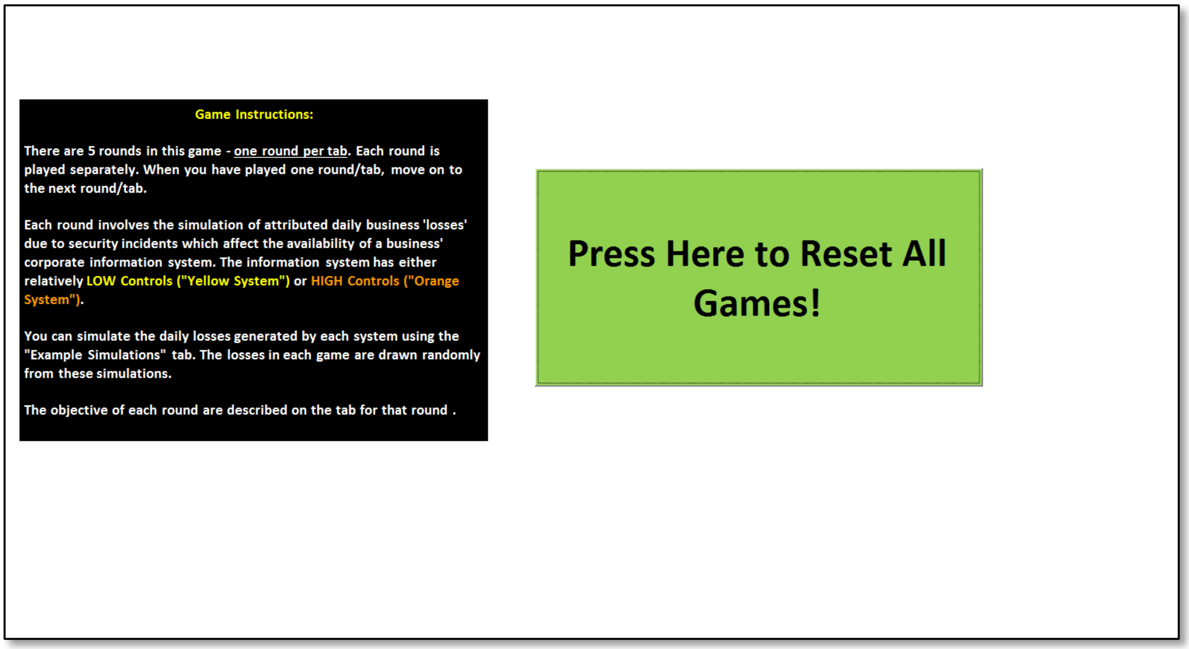

Game Instructions:

There are 5 rounds in this game - one round per tab. Each round is played separately. When you have played one round/tab, move on to the next round/tab.

Each round involves the simulation of attributed daily business 'losses' due to security incidents which affect the availability of a business' corporate information system. The information system has either relatively LOW Controls ("Yellow System") or HIGH Controls ("Orange System").

You can simulate the daily losses generated by each system using the "Example Simulations" tab. The losses in each game are drawn randomly from these simulations.

The objective of each round are described on the tab for that round .

### Instructions:

“There are 5 rounds in this game - one round per tab. Each round is played separately. When you have played one round/tab, move on to the next round/tab.

Each round involves the simulation of attributed daily business 'losses' due to security incidents which affect the availability of a business' corporate information system. The information system has either relatively LOW Controls ("Yellow System") or HIGH Controls ("Orange System").

You can simulate the daily losses generated by each system using the "Example Simulations" tab. The losses in each game are drawn randomly from these simulations.

The objective of each round are described on the tab for that round.”

## 2 - Example Simulations Screen:

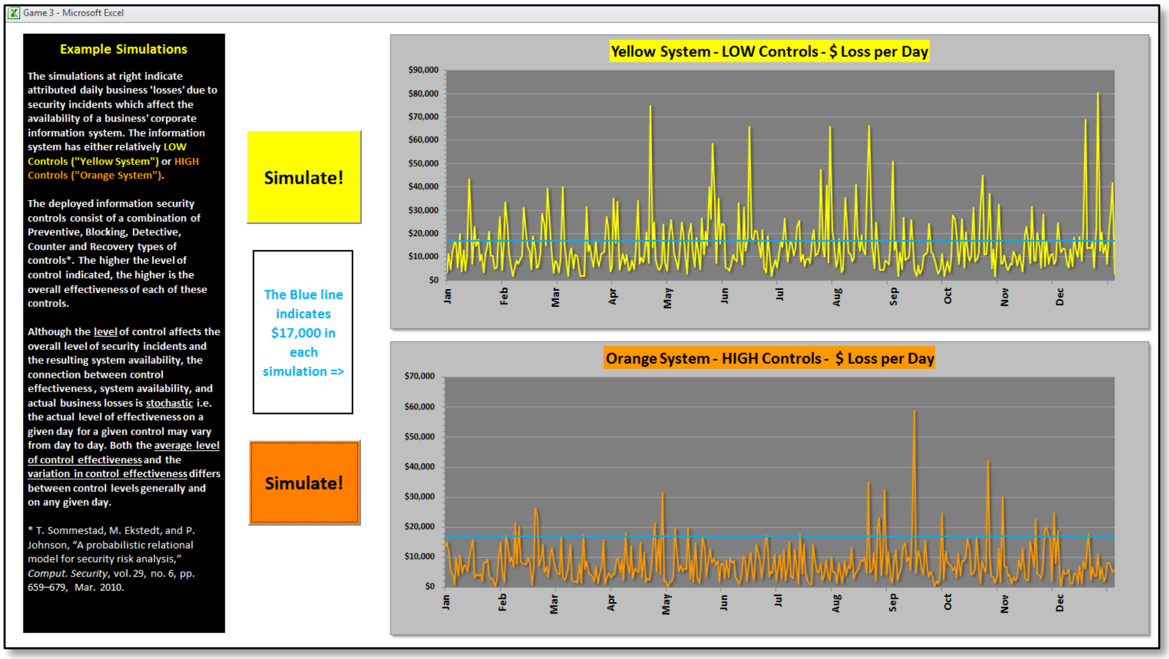

Game 3 - Microsoft Excel

**Example Simulations**

**The simulations at right indicate attributed daily business 'losses' due to security incidents which affect the availability of a business' corporate information system. The information system has either relatively LOW Controls ("Yellow System") or HIGH Controls ("Orange System").**

**The deployed information security controls consist of a combination of Preventive, Blocking, Detective, Counter and Recovery types of controls*. The higher the level of control indicated, the higher is the overall effectiveness of each of these controls.**

**Although the level of control affects the overall level of security incidents and the resulting system availability, the connection between control effectiveness, system availability, and actual business losses is stochastic i.e. the actual level of effectiveness on a given day for a given control may vary from day to day. Both the average level of control effectiveness and the variation in control effectiveness differs between control levels generally and on any given day.**

* T. Sommestad, M. Ekstedt, and P. Johnson, “A probabilistic relational model for security risk analysis,” *Comput. Security*, vol. 29, no. 6, pp. 659–679, Mar. 2010.

The Blue line indicates $17,000 in each simulation =>

### Instructions:

“The simulations at right indicate attributed daily business 'losses' due to security incidents which affect the availability of a business' corporate information system. The information system has either relatively LOW Controls ("Yellow System") or HIGH Controls ("Orange System").

The deployed information security controls consist of a combination of Preventive, Blocking, Detective, Counter and Recovery types of controls*. The higher the level of control indicated, the higher is the overall effectiveness of each of these controls.

Although the level of control affects the overall level of security incidents and the resulting system availability, the connection between control effectiveness , system availability, and actual business losses is stochastic i.e. the actual level of effectiveness on a given day for a given control may vary from day to day. Both the average level of control effectiveness and the variation in control effectiveness differs between control levels generally and on any given day.

* T. Sommestad, M. Ekstedt, and P. Johnson, “A probabilistic relational model for security risk analysis,” Comput. Security, vol. 29, no. 6, pp. 659–679, Mar. 2010.”

## 3 – Game Interface (Low Controls System)

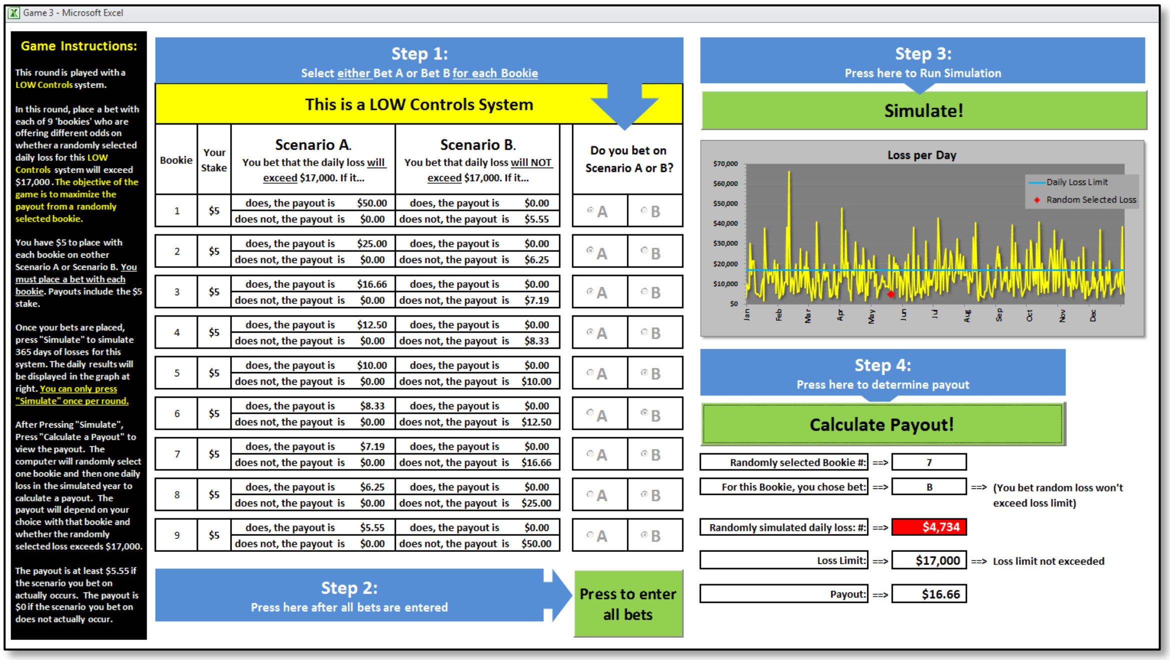

Game 3 - Microsoft Excel

**Game Instructions:**

This round is played with a LOW Controls system.

In this round, place a bet with each of 9 'bookies' who are offering different odds on whether a randomly selected daily loss for this LOW Controls system will exceed $17,000 . The objective of the game is to maximize the payout from a randomly selected bookie.

You have $5 to place with each bookie on eother Scenario A or Scenario B. You must place a bet with each bookie. Payouts include the $5 stake.

Once your bets are placed, press "Simulate" to simulate 365 days of losses for this system. The daily results will be displayed in the graph at right. You can only press "Simulate" once per round.

After Pressing "Simulate", Press "Calculate a Payout" to view the payout. The computer will randomly select one bookie and then one daily loss in the simulated year to calculate a payout. The payout will depend on your choice with that bookie and whether the randomly selected loss exceeds $17,000.

The payout is at least $5.55 if the scenario you bet on actually occurs. The payout is $0 if the scenario you bet on does not actually occur.

**Step 1:**
Select either Bet A or Bet B for each Bookie

**This is a LOW Controls System**

| Bookie | Your Stake | Scenario A. You bet that the daily loss will exceed $17,000. If it... | Scenario B. You bet that daily loss will NOT exceed $17,000. If it... | Do you bet on Scenario A or B? | |
|---|---|---|---|---|---|
| 1 | $5 | does, the payout is $50.00<br>does not, the payout is $0.00 | does, the payout is $0.00<br>does not, the payout is $5.55 | ● A | ○ B |
| 2 | $5 | does, the payout is $25.00<br>does not, the payout is $0.00 | does, the payout is $0.00<br>does not, the payout is $6.25 | ● A | ○ B |
| 3 | $5 | does, the payout is $16.66<br>does not, the payout is $0.00 | does, the payout is $0.00<br>does not, the payout is $7.19 | ● A | ○ B |
| 4 | $5 | does, the payout is $12.50<br>does not, the payout is $0.00 | does, the payout is $0.00<br>does not, the payout is $8.33 | ● A | ○ B |
| 5 | $5 | does, the payout is $10.00<br>does not, the payout is $0.00 | does, the payout is $0.00<br>does not, the payout is $10.00 | ○ A | ● B |
| 6 | $5 | does, the payout is $8.33<br>does not, the payout is $0.00 | does, the payout is $0.00<br>does not, the payout is $12.50 | ○ A | ● B |
| 7 | $5 | does, the payout is $7.19<br>does not, the payout is $0.00 | does, the payout is $0.00<br>does not, the payout is $16.66 | ○ A | ● B |
| 8 | $5 | does, the payout is $6.25<br>does not, the payout is $0.00 | does, the payout is $0.00<br>does not, the payout is $25.00 | ○ A | ● B |
| 9 | $5 | does, the payout is $5.55<br>does not, the payout is $0.00 | does, the payout is $0.00<br>does not, the payout is $50.00 | ○ A | ● B |

**Step 2:**
Press here after all bets are entered

Press to enter all bets

**Step 3:**
Press here to Run Simulation

Simulate!

**Step 4:**
Press here to determine payout

Calculate Payout!

Randomly selected Bookie #: ==> 7

For this Bookie, you chose bet: ==> B ==> (You bet random loss won't exceed loss limit)

Randomly simulated daily loss: #: ==> $4,734

Loss Limit: ==> $17,000 ==> Loss limit not exceeded

Payout: ==> $16.66

**Instructions:**

“This round is played with a LOW Controls system.

In this round, place a bet with each of 9 'bookies' who are offering different odds on whether a randomly selected daily loss for this LOW Controls system will exceed $17,000 . The objective of the game is to maximize the payout from a randomly selected bookie.

You have $5 to place with each bookie on either Scenario A or Scenario B. You must place a bet with each bookie. Payouts include the $5 stake.

Once your bets are placed, press "Simulate" to simulate 365 days of losses for this system. The daily results will be displayed in the graph at right. You can only press "Simulate" once per round.

After Pressing "Simulate", Press "Calculate a Payout" to view the payout. The computer will randomly select one bookie and then one daily loss in the simulated year to calculate a payout. The payout will depend on your choice with that bookie and whether the randomly selected loss exceeds $17,000.

The payout is at least $5.55 if the scenario you bet on actually occurs. The payout is $0 if the scenario you bet on does not actually occur.”

## 4 – Game Interface (High Controls System)

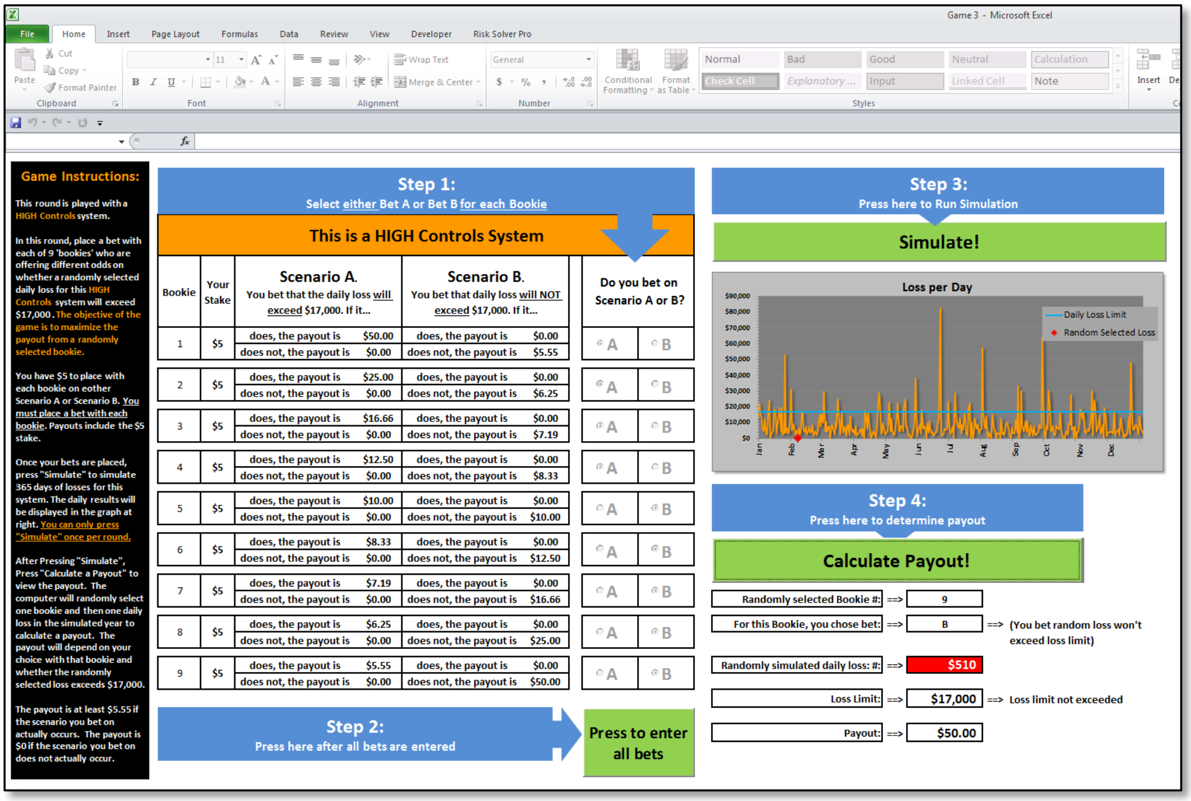

**Game Instructions:**

This round is played with a HIGH Controls system.

In this round, place a bet with each of 9 'bookies' who are offering different odds on whether a randomly selected daily loss for this HIGH Controls system will exceed $17,000 . The objective of the game is to maximize the payout from a randomly selected bookie.

You have $5 to place with each bookie on eother Scenario A or Scenario B. You must place a bet with each bookie. Payouts include the $5 stake.

Once your bets are placed, press "Simulate" to simulate 365 days of losses for this system. The daily results will be displayed in the graph at right. You can only press "Simulate" once per round.

After Pressing "Simulate", Press "Calculate a Payout" to view the payout. The computer will randomly select one bookie and then one daily loss in the simulated year to calculate a payout. The payout will depend on your choice with that bookie and whether the randomly selected loss exceeds $17,000.

The payout is at least $5.55 if the scenario you bet on actually occurs. The payout is $0 if the scenario you bet on does not actually occur.

**Step 1:**
Select either Bet A or Bet B for each Bookie

**This is a HIGH Controls System**

| Bookie | Your Stake | Scenario A.<br>You bet that the daily loss will exceed $17,000. If it... | | Scenario B.<br>You bet that daily loss will NOT exceed $17,000. If it... | | Do you bet on Scenario A or B? | |
|---|---|---|---|---|---|---|---|
| 1 | $5 | does, the payout is<br>does not, the payout is | $50.00<br>$0.00 | does, the payout is<br>does not, the payout is | $0.00<br>$5.55 | ● A | ○ B |
| 2 | $5 | does, the payout is<br>does not, the payout is | $25.00<br>$0.00 | does, the payout is<br>does not, the payout is | $0.00<br>$6.25 | ● A | ○ B |
| 3 | $5 | does, the payout is<br>does not, the payout is | $16.66<br>$0.00 | does, the payout is<br>does not, the payout is | $0.00<br>$7.19 | ● A | ○ B |
| 4 | $5 | does, the payout is<br>does not, the payout is | $12.50<br>$0.00 | does, the payout is<br>does not, the payout is | $0.00<br>$8.33 | ● A | ○ B |
| 5 | $5 | does, the payout is<br>does not, the payout is | $10.00<br>$0.00 | does, the payout is<br>does not, the payout is | $0.00<br>$10.00 | ○ A | ● B |
| 6 | $5 | does, the payout is<br>does not, the payout is | $8.33<br>$0.00 | does, the payout is<br>does not, the payout is | $0.00<br>$12.50 | ○ A | ● B |
| 7 | $5 | does, the payout is<br>does not, the payout is | $7.19<br>$0.00 | does, the payout is<br>does not, the payout is | $0.00<br>$16.66 | ○ A | ● B |
| 8 | $5 | does, the payout is<br>does not, the payout is | $6.25<br>$0.00 | does, the payout is<br>does not, the payout is | $0.00<br>$25.00 | ○ A | ● B |
| 9 | $5 | does, the payout is<br>does not, the payout is | $5.55<br>$0.00 | does, the payout is<br>does not, the payout is | $0.00<br>$50.00 | ○ A | ● B |

**Step 2:**
Press here after all bets are entered

Press to enter all bets

**Step 3:**
Press here to Run Simulation

Simulate!

**Step 4:**
Press here to determine payout

Calculate Payout!

| | | | | |
|---|---|---|---|---|
| Randomly selected Bookie #: | ==> | 9 | | |
| For this Bookie, you chose bet: | ==> | B | ==> | (You bet random loss won't exceed loss limit) |
| Randomly simulated daily loss: #: | ==> | $510 | | |
| Loss Limit: | ==> | $17,000 | ==> | Loss limit not exceeded |
| Payout: | ==> | $50.00 | | |

## Instructions:

“This round is played with a HIGH Controls system.

In this round, place a bet with each of 9 'bookies' who are offering different odds on whether a randomly selected daily loss for this HIGH Controls  system will exceed $17,000 . The objective of the game is to maximize the payout from a randomly selected bookie.

You have $5 to place with each bookie on eother Scenario A or Scenario B. You must place a bet with each bookie. Payouts include the $5 stake.

Once your bets are placed, press "Simulate" to simulate 365 days of losses for this system. The daily results will be displayed in the graph at right. You can only press "Simulate" once per round.

After Pressing "Simulate",  Press "Calculate a Payout" to view the payout.  The computer will randomly select one bookie and then one daily loss in the simulated year to calculate a payout.  The payout will depend on your choice with that bookie and whether the randomly selected loss exceeds $17,000.

The payout is at least $5.55 if the scenario you bet on actually occurs.  The payout is $0 if the scenario you bet on does not actually occur.”

## 5 – Game Interface (Multiple Price List Control Selection – 5 Price Levels Example)

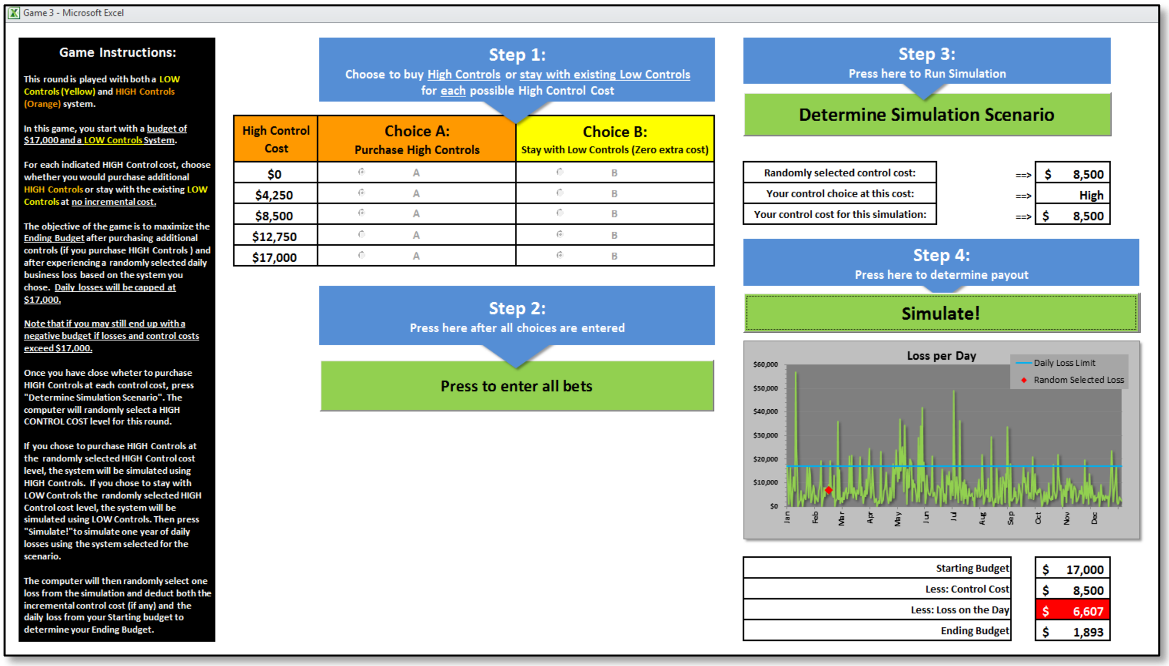

Game 3 - Microsoft Excel

**Game Instructions:**

This round is played with both a LOW Controls (Yellow) and HIGH Controls (Orange) system.

In this game, you start with a budget of $17,000 and a LOW Controls System.

For each indicated HIGH Control cost, choose whether you would purchase additional HIGH Controls or stay with the existing LOW Controls at no incremental cost.

The objective of the game is to maximize the Ending Budget after purchasing additional controls (if you purchase HIGH Controls ) and after experiencing a randomly selected daily business loss based on the system you chose. Daily losses will be capped at $17,000.

Note that if you may still end up with a negative budget if losses and control costs exceed $17,000.

Once you have close wheter to purchase HIGH Controls at each control cost, press "Determine Simulation Scenario". The computer will randomly select a HIGH CONTROL COST level for this round.

If you chose to purchase HIGH Controls at the randomly selected HIGH Control cost level, the system will be simulated using HIGH Controls. If you chose to stay with LOW Controls the randomly selected HIGH Control cost level, the system will be simulated using LOW Controls. Then press "Simulate!"to simulate one year of daily losses using the system selected for the scenario.

The computer will then randomly select one loss from the simulation and deduct both the incremental control cost (if any) and the daily loss from your Starting budget to determine your Ending Budget.

**Step 1:**
Choose to buy High Controls or stay with existing Low Controls for each possible High Control Cost

| High Control Cost | Choice A: Purchase High Controls | Choice B: Stay with Low Controls (Zero extra cost) |
|---|---|---|
| $0 | ● A | ○ B |
| $4,250 | ● A | ○ B |
| $8,500 | ● A | ○ B |
| $12,750 | ○ A | ● B |
| $17,000 | ○ A | ● B |

**Step 2:**
Press here after all choices are entered

Press to enter all bets

**Step 3:**
Press here to Run Simulation

Determine Simulation Scenario

| | | |
|---|---|---|
| Randomly selected control cost: | ==> | $ 8,500 |
| Your control choice at this cost: | ==> | High |
| Your control cost for this simulation: | ==> | $ 8,500 |

**Step 4:**
Press here to determine payout

Simulate!

| | |
|---|---|
| Starting Budget | $ 17,000 |
| Less: Control Cost | $ 8,500 |
| Less: Loss on the Day | $ 6,607 |
| Ending Budget | $ 1,893 |

## Instructions:

"This round is played with both a LOW Controls (Yellow) and HIGH Controls (Orange) system.

In this game, you start with a budget of $17,000 and a LOW Controls System.

For each indicated HIGH Control cost, choose whether you would purchase additional HIGH Controls or stay with the existing LOW Controls at no incremental cost.

The objective of the game is to maximize the Ending Budget after purchasing additional controls (if you purchase HIGH Controls) and after experiencing a randomly selected daily business loss based on the system you chose. Daily losses will be capped at $17,000.

Note that if you may still end up with a negative budget if losses and control costs exceed $17,000.

Once you have chosen whether to purchase HIGH Controls at each control cost, press "Determine Simulation Scenario". The computer will randomly select a HIGH CONTROL COST level for this round.

If you chose to purchase HIGH Controls at the randomly selected HIGH Control cost level, the system will be simulated using HIGH Controls. If you chose to stay with LOW Controls the randomly selected HIGH Control cost level, the system will be simulated using LOW Controls. Then press "Simulate!" to simulate one year of daily losses using the system selected for the scenario.

The computer will then randomly select one loss from the simulation and deduct both the incremental control cost (if any) and the daily loss from your Starting budget to determine your Ending Budget."

## Game 4: Recovering Subjective Probability Estimates and Rank Dependent Utility Bias

### 1 – Binary Lottery Game Interface

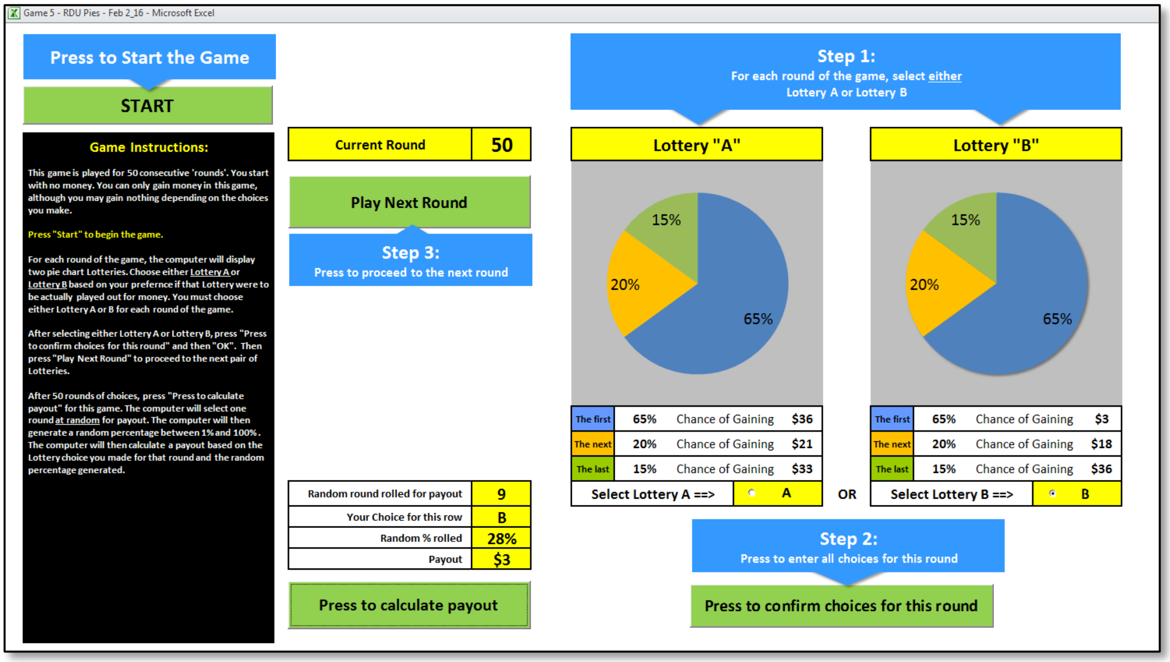
Game 5 - RDU Pies - Feb 2_16 - Microsoft Excel
Press to Start the Game
START
Play Next Round
Step 3:
Press to proceed to the next round
Press to calculate payout
Step 1:
For each round of the game, select either Lottery A or Lottery B
Lottery "A"
Lottery "B"
15%
20%
65%
15%
20%
65%
OR
Step 2:
Press to enter all choices for this round
Press to confirm choices for this round

Game Instructions:

This game is played for 50 consecutive 'rounds'. You start with no money. You can only gain money in this game, although you may gain nothing depending on the choices you make.

Press "Start" to begin the game.

For each round of the game, the computer will display two pie chart Lotteries. Choose either Lottery A or Lottery B based on your prefernce if that Lottery were to be actually played out for money. You must choose either Lottery A or B for each round of the game.

After selecting either Lottery A or Lottery B, press "Press to confirm choices for this round" and then "OK". Then press "Play Next Round" to proceed to the next pair of Lotteries.

After 50 rounds of choices, press "Press to calculate payout" for this game. The computer will select one round at random for payout. The computer will then generate a random percentage between 1% and 100%. The computer will then calculate a payout based on the Lottery choice you made for that round and the random percentage generated.

| | |
|---|---|
| Current Round | 50 |

| | |
|---|---|
| Random round rolled for payout | 9 |
| Your Choice for this row | B |
| Random % rolled | 28% |
| Payout | $3 |

| | | | |
|---|---|---|---|
| The first | 65% | Chance of Gaining | $36 |
| The next | 20% | Chance of Gaining | $21 |
| The last | 15% | Chance of Gaining | $33 |
| Select Lottery A ==> | | ○ | A |

| | | | |
|---|---|---|---|
| The first | 65% | Chance of Gaining | $3 |
| The next | 20% | Chance of Gaining | $18 |
| The last | 15% | Chance of Gaining | $36 |
| Select Lottery B ==> | | ● | B |

### Instructions:

“This game is played for 50 consecutive 'rounds'. You start with no money. You can only gain money in this game, although you may gain nothing depending on the choices you make.

Press "Start" to begin the game.

For each round of the game, the computer will display two pie chart Lotteries. Choose either Lottery A or Lottery B based on your preference if that Lottery were to be actually played out for money. You must choose either Lottery A or B for each round of the game.

After selecting either Lottery A or Lottery B, press "Press to confirm choices for this round" and then "OK". Then press "Play Next Round" to proceed to the next pair of Lotteries.

After 50 rounds of choices, press "Press to calculate payout" for this game. The computer will select one round at random for payout. The computer will then generate a random percentage between 1% and 100% . The computer will then calculate a payout based on the Lottery choice you made for that round and the random percentage generated.”

## 2A – QSR Scoring Device

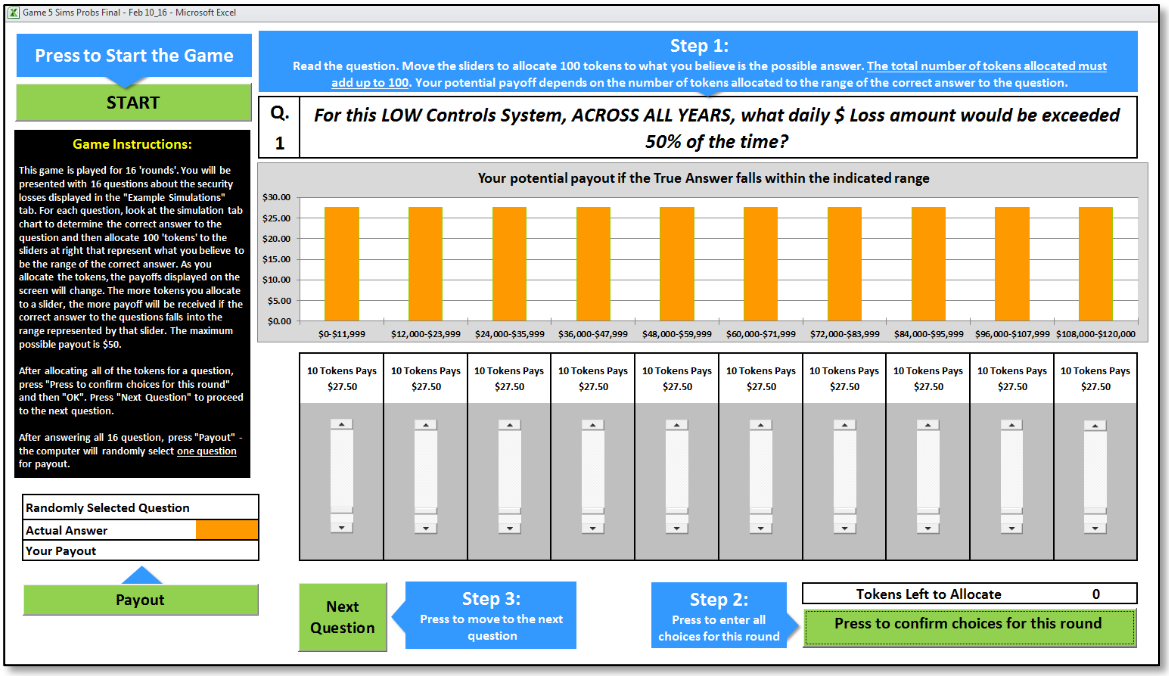

**Step 1:**
Read the question. Move the sliders to allocate 100 tokens to what you believe is the possible answer. The total number of tokens allocated must add up to 100. Your potential payoff depends on the number of tokens allocated to the range of the correct answer to the question.

**Q. 1** *For this LOW Controls System, ACROSS ALL YEARS, what daily $ Loss amount would be exceeded 50% of the time?*

Your potential payout if the True Answer falls within the indicated range

| $0-$11,999 | $12,000-$23,999 | $24,000-$35,999 | $36,000-$47,999 | $48,000-$59,999 | $60,000-$71,999 | $72,000-$83,999 | $84,000-$95,999 | $96,000-$107,999 | $108,000-$120,000 |
|---|---|---|---|---|---|---|---|---|---|
| 10 Tokens Pays $27.50 | 10 Tokens Pays $27.50 | 10 Tokens Pays $27.50 | 10 Tokens Pays $27.50 | 10 Tokens Pays $27.50 | 10 Tokens Pays $27.50 | 10 Tokens Pays $27.50 | 10 Tokens Pays $27.50 | 10 Tokens Pays $27.50 | 10 Tokens Pays $27.50 |

**Game Instructions:**

This game is played for 16 'rounds'. You will be presented with 16 questions about the security losses displayed in the "Example Simulations" tab. For each question, look at the simulation tab chart to determine the correct answer to the question and then allocate 100 'tokens' to the sliders at right that represent what you believe to be the range of the correct answer. As you allocate the tokens, the payoffs displayed on the screen will change. The more tokens you allocate to a slider, the more payoff will be received if the correct answer to the questions falls into the range represented by that slider. The maximum possible payout is $50.

After allocating all of the tokens for a question, press "Press to confirm choices for this round" and then "OK". Press "Next Question" to proceed to the next question.

After answering all 16 question, press "Payout" - the computer will randomly select one question for payout.

**Step 3:** Press to move to the next question

**Step 2:** Press to enter all choices for this round

**Instructions:**

“This game is played for 16 'rounds'. You will be presented with 16 questions about the security losses displayed in the "Example Simulations" tab. For each question, look at the simulation tab chart to determine the correct answer to the question and then allocate 100 'tokens' to the sliders at right that represent what you believe to be the range of the correct answer. As you allocate the tokens, the payoffs displayed on the screen will change. The more tokens you allocate to a slider, the more payoff will be received if the correct answer to the questions falls into the range represented by that slider. The maximum possible payout is $50.

After allocating all of the tokens for a question, press "Press to confirm choices for this round" and then "OK". Press "Next Question" to proceed to the next question.

After answering all 16 question,  press "Payout" - the computer will randomly select one question for payout.”

## 2B – Question Detail Screen (Probability Distribution Example)

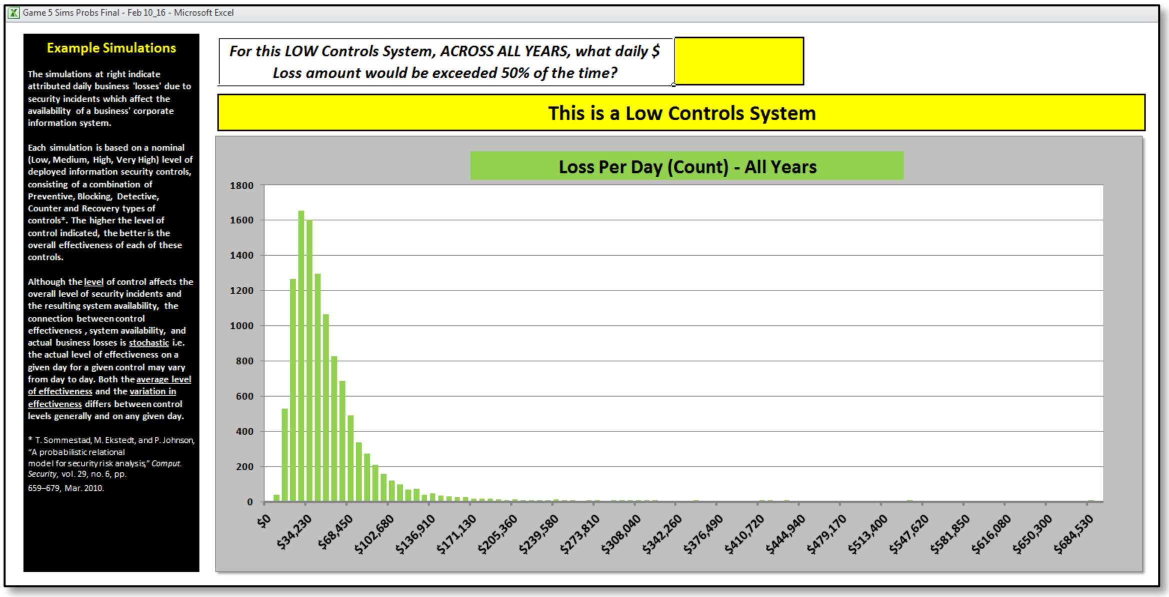

Game 5 Sims Probs Final - Feb 10_16 - Microsoft Excel

**Example Simulations**

The simulations at right indicate attributed daily business 'losses' due to security incidents which affect the availability of a business' corporate information system.

Each simulation is based on a nominal (Low, Medium, High, Very High) level of deployed information security controls, consisting of a combination of Preventive, Blocking, Detective, Counter and Recovery types of controls*. The higher the level of control indicated, the better is the overall effectiveness of each of these controls.

Although the level of control affects the overall level of security incidents and the resulting system availability, the connection between control effectiveness , system availability, and actual business losses is stochastic i.e. the actual level of effectiveness on a given day for a given control may vary from day to day. Both the average level of effectiveness and the variation in effectiveness differs between control levels generally and on any given day.

* T. Sommestad, M. Ekstedt, and P. Johnson, "A probabilistic relational model for security risk analysis," *Comput. Security*, vol. 29, no. 6, pp. 659–679, Mar. 2010.

*For this LOW Controls System, ACROSS ALL YEARS, what daily $ Loss amount would be exceeded 50% of the time?*

**This is a Low Controls System**

### Instructions:

"The simulations at right indicate attributed daily business 'losses' due to security incidents which affect the availability of a business' corporate information system.

Each simulation is based on a nominal (Low, Medium, High, Very High) level of deployed information security controls, consisting of a combination of Preventive, Blocking, Detective, Counter and Recovery types of controls*. The higher the level of control indicated, the better is the overall effectiveness of each of these controls.

Although the level of control affects the overall level of security incidents and the resulting system availability, the connection between control effectiveness , system availability, and actual business losses is stochastic i.e. the actual level of effectiveness on a given day for a given control may vary from day to day. Both the average level of effectiveness and the variation in effectiveness differs between control levels generally and on any given day.

* T. Sommestad, M. Ekstedt, and P. Johnson, "A probabilistic relational model for security risk analysis," Comput. Security, vol. 29, no. 6, pp. 659–679, Mar. 2010."

## Game 5: Effect of Risk and Ambiguity on Precaution vs. Insurance Choices

### 1 – Insurance-Only Treatment Interface

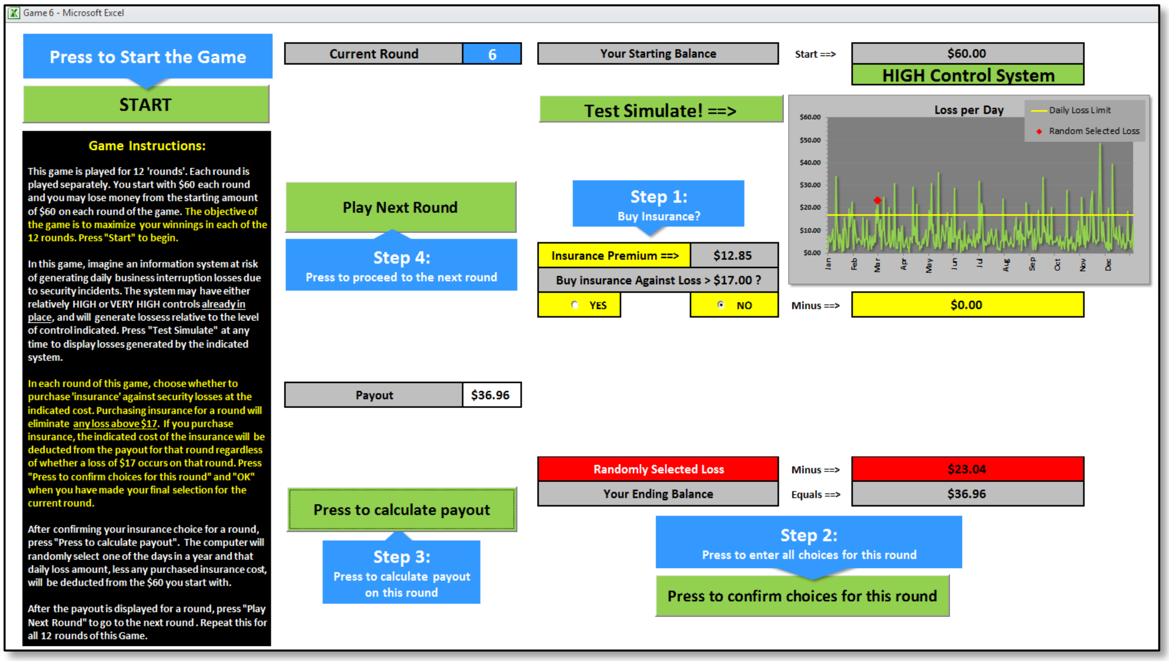

Game Instructions:

This game is played for 12 'rounds'. Each round is played separately. You start with $60 each round and you may lose money from the starting amount of $60 on each round of the game. The objective of the game is to maximize your winnings in each of the 12 rounds. Press "Start" to begin.

In this game, imagine an information system at risk of generating daily business interruption losses due to security incidents. The system may have either relatively HIGH or VERY HIGH controls already in place, and will generate lossess relative to the level of control indicated. Press "Test Simulate" at any time to display losses generated by the indicated system.

In each round of this game, choose whether to purchase 'insurance' against security losses at the indicated cost. Purchasing insurance for a round will eliminate any loss above $17. If you purchase insurance, the indicated cost of the insurance will be deducted from the payout for that round regardless of whether a loss of $17 occurs on that round. Press "Press to confirm choices for this round" and "OK" when you have made your final selection for the current round.

After confirming your insurance choice for a round, press "Press to calculate payout". The computer will randomly select one of the days in a year and that daily loss amount, less any purchased insurance cost, will be deducted from the $60 you start with.

After the payout is displayed for a round, press "Play Next Round" to go to the next round. Repeat this for all 12 rounds of this Game.

Current Round: 6

Payout: $36.96

Your Starting Balance — Start ==> $60.00

HIGH Control System

Insurance Premium ==> $12.85

Buy insurance Against Loss > $17.00 ?

○ YES ● NO — Minus ==> $0.00

Randomly Selected Loss — Minus ==> $23.04

Your Ending Balance — Equals ==> $36.96

**Instructions:**

“This game is played for 12 'rounds'. Each round is played separately. You start with $60 each round and you may lose money from the starting amount of $60 on each round of the game. The objective of the game is to maximize your winnings in each of the 12 rounds. Press "Start" to begin.

In this game, imagine an information system at risk of generating daily business interruption losses due to security incidents. The system may have either relatively HIGH or VERY HIGH controls already in place, and will generate losses relative to the level of control indicated. Press "Test Simulate" at any time to display losses generated by the indicated system.

In each round of this game, choose whether to purchase 'insurance' against security losses at the indicated cost. Purchasing insurance for a round will eliminate any loss above $17. If you purchase insurance, the indicated cost of the insurance will be deducted from the payout for that round regardless of whether a loss of $17 occurs on that round. Press "Press to confirm choices for this round" and "OK" when you have made your final selection for the current round.

After confirming your insurance choice for a round, press "Press to calculate payout". The computer will randomly select one of the days in a year and that daily loss amount, less any purchased insurance cost, will be deducted from the $60 you start with.

After the payout is displayed for a round, press "Play Next Round" to go to the next roun . Repeat this for all 12 rounds of this Game.”

## 2 – Precaution-Only Treatment Interface

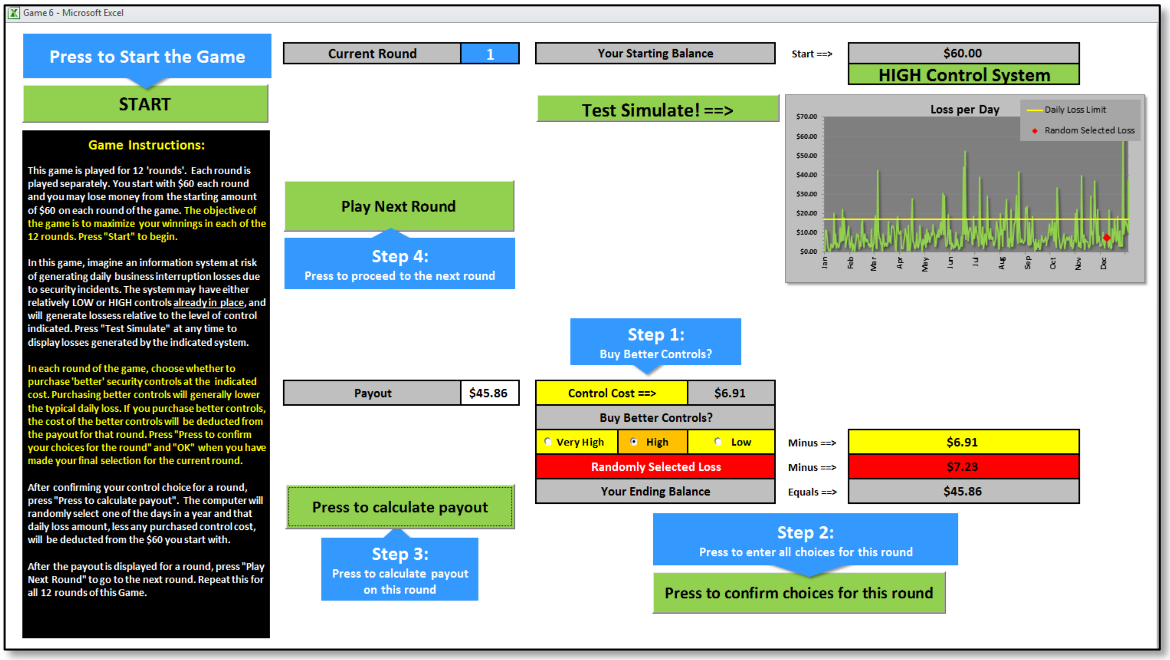

Game Instructions:

This game is played for 12 'rounds'. Each round is played separately. You start with $60 each round and you may lose money from the starting amount of $60 on each round of the game. The objective of the game is to maximize your winnings in each of the 12 rounds. Press "Start" to begin.

In this game, imagine an information system at risk of generating daily business interruption losses due to security incidents. The system may have either relatively LOW or HIGH controls already in place, and will generate lossess relative to the level of control indicated. Press "Test Simulate" at any time to display losses generated by the indicated system.

In each round of the game, choose whether to purchase 'better' security controls at the indicated cost. Purchasing better controls will generally lower the typical daily loss. If you purchase better controls, the cost of the better controls will be deducted from the payout for that round. Press "Press to confirm your choices for the round" and "OK" when you have made your final selection for the current round.

After confirming your control choice for a round, press "Press to calculate payout". The computer will randomly select one of the days in a year and that daily loss amount, less any purchased control cost, will be deducted from the $60 you start with.

After the payout is displayed for a round, press "Play Next Round" to go to the next round. Repeat this for all 12 rounds of this Game.

| | |
|---|---|
| Current Round | 1 |
| Payout | $45.86 |

| | | |
|---|---|---|
| Your Starting Balance | Start ==> | $60.00<br>HIGH Control System |
| Control Cost ==> | | $6.91 |
| Buy Better Controls? ○ Very High ● High ○ Low | Minus ==> | $6.91 |
| Randomly Selected Loss | Minus ==> | $7.23 |
| Your Ending Balance | Equals ==> | $45.86 |

**Instructions:**

“This game is played for 12 'rounds'. Each round is played separately. You start with $60 each round and you may lose money from the starting amount of $60 on each round of the game. The objective of the game is to maximize your winnings in each of the 12 rounds. Press "Start" to begin.

In this game, imagine an information system at risk of generating daily business interruption losses due to security incidents. The system may have either relatively LOW or HIGH controls already in place, and will generate losses relative to the level of control indicated. Press "Test Simulate" at any time to display losses generated by the indicated system.

In each round of the game, choose whether to purchase 'better' security controls at the indicated cost. Purchasing better controls will generally lower the typical daily loss. If you purchase better controls, the cost of the better controls will be deducted from the payout for that round. Press "Press to confirm your choices for the round" and "OK" when you have made your final selection for the current round.

After confirming your control choice for a round, press "Press to calculate payout". The computer will randomly select one of the days in a year and that daily loss amount, less any purchased control cost, will be deducted from the $60 you start with.

After the payout is displayed for a round, press "Play Next Round" to go to the next round. Repeat this for all 12 rounds of this Game.”

### 3 – Insurance + Precaution Treatment Interface

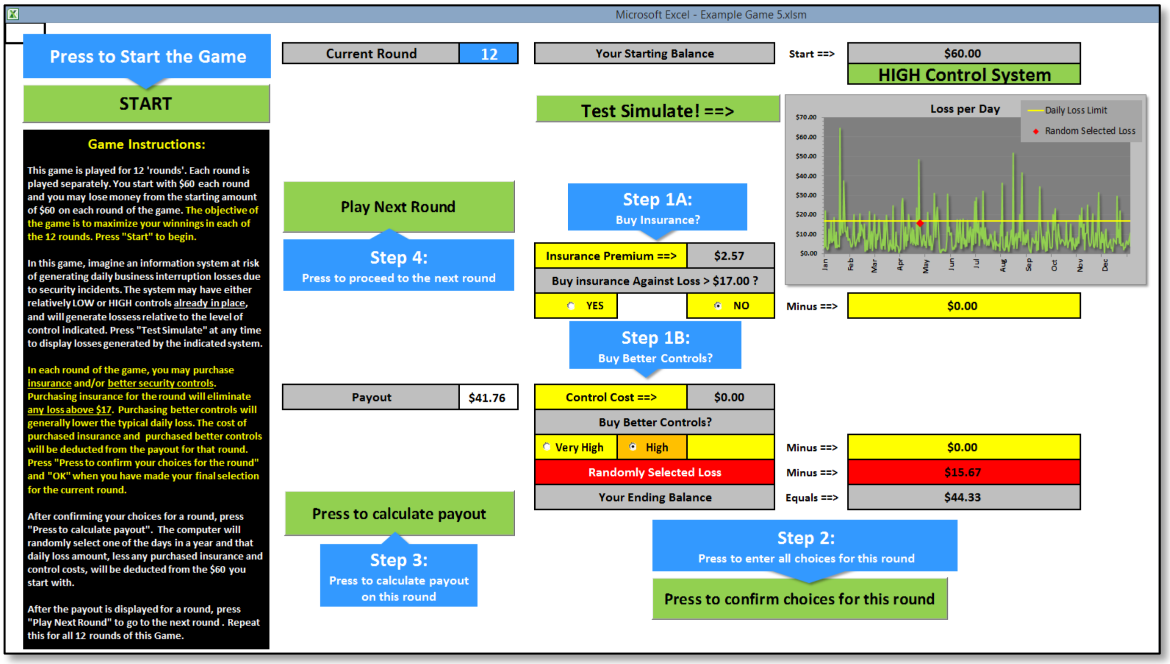

Current Round: 12

Payout: $41.76

| | | | |
|---|---|---|---|
| Your Starting Balance | | Start ==> | $60.00 |
| | | | HIGH Control System |
| Insurance Premium ==> | $2.57 | | |
| Buy insurance Against Loss > $17.00 ? | | | |
| ○ YES | ● NO | Minus ==> | $0.00 |
| Control Cost ==> | $0.00 | | |
| Buy Better Controls? | | | |
| ○ Very High | ● High | Minus ==> | $0.00 |
| Randomly Selected Loss | | Minus ==> | $15.67 |
| Your Ending Balance | | Equals ==> | $44.33 |

Game Instructions:

This game is played for 12 'rounds'. Each round is played separately. You start with $60 each round and you may lose money from the starting amount of $60 on each round of the game. The objective of the game is to maximize your winnings in each of the 12 rounds. Press "Start" to begin.

In this game, imagine an information system at risk of generating daily business interruption losses due to security incidents. The system may have either relatively LOW or HIGH controls already in place, and will generate lossess relative to the level of control indicated. Press "Test Simulate" at any time to display losses generated by the indicated system.

In each round of the game, you may purchase insurance and/or better security controls. Purchasing insurance for the round will eliminate any loss above $17. Purchasing better controls will generally lower the typical daily loss. The cost of purchased insurance and purchased better controls will be deducted from the payout for that round. Press "Press to confirm your choices for the round" and "OK" when you have made your final selection for the current round.

After confirming your choices for a round, press "Press to calculate payout". The computer will randomly select one of the days in a year and that daily loss amount, less any purchased insurance and control costs, will be deducted from the $60 you start with.

After the payout is displayed for a round, press "Play Next Round" to go to the next round . Repeat this for all 12 rounds of this Game.

## Instructions:

“This game is played for 12 'rounds'. Each round is played separately. You start with $60 each round and you may lose money from the starting amount of $60 on each round of the game. The objective of the game is to maximize your winnings in each of the 12 rounds. Press "Start" to begin.

In this game, imagine an information system at risk of generating daily business interruption losses due to security incidents. The system may have either relatively LOW or HIGH controls already in place, and will generate losses relative to the level of control indicated. Press "Test Simulate" at any time to display losses generated by the indicated system.

In each round of the game, you may purchase insurance and/or better security controls. Purchasing insurance for the round will eliminate any loss above $17. Purchasing better controls will generally lower the typical daily loss. The cost of purchased insurance and purchased better controls will be deducted from the payout for that round. Press "Press to confirm your choices for the round" and "OK" when you have made your final selection for the current round.

After confirming your choices for a round, press "Press to calculate payout".  The computer will randomly select one of the days in a year and that daily loss amount, less any purchased insurance and control costs, will be deducted from the $60 you start with.

After the payout is displayed for a round, press "Play Next Round" to go to the next round. Repeat this for all 12 rounds of this Game.”

## Appendix 2 - Demographic Survey Questions

| | | |
|---|---|---|
| 1 | What is your age in years? | Under 25<br>25-34<br>35-44<br>45-54<br>55-64<br>65-74<br>75 or older |
| | | |
| 2 | What is your gender? | Male<br>Female |
| | | |
| 3 | What is your marital status? | Never Married<br>Not married but living with a significant partner<br>Married<br>Separated or divorced<br>Widowed<br>Independent shared living arrangement |
| | | |
| 4 | What is/was your main field of study in school? | Accounting<br>Economics<br>Finance<br>Business Administration, other than Accounting, Economics, or Finance<br>Education<br>Engineering<br>Health and Medicine<br>Biological and Biomedical Sciences<br>Math, Computer Sciences, or Physical Sciences<br>Social Sciences or History<br>Law<br>Psychology<br>Modern Languages and Cultures<br>Other Fields |
| | | |
| 5 | Do you have a business degree from a College or University? | Yes<br>No |
| | | |
| 6 | Do you have a mathematics, economics or sciences degree from a College or University? | Yes<br>No |
| | | |
| 7 | What is the highest level of education you have completed? | High School<br>Community College Diploma<br>University Bachelor's degree<br>Master's degree<br>Doctoral degree |
| | | |
| 8 | How many years have you been working full time (entire career)? | Less than 1 year<br>1 - 5 years<br>6 - 10 years<br>11 - 15 years<br>16 - 20 years |

| | | |
|---|---|---|
| | | 21 - 30 years<br>more than 30 years |
| | | |
| 9 | How many years have you been working in a professional privacy or security role? | I have never worked in a privacy or security role<br>Less than 1 year<br>1 - 5 years<br>6 - 10 years<br>11 - 15 years<br>16 - 20 years<br>21 - 30 years<br>More than 30 years |
| | | |
| 10 | How many people report to you directly? | None<br>1 to 2<br>3 to 5<br>6 to 10<br>11 to 20<br>21 to 50<br>more than 50 |
| | | |
| 11 | Do you earn a performance related bonus as part of your salary? | Yes<br>No |
| | | |
| 12 | What is your current household income before taxes? | $15,000 or under<br>$15,001 - $35,000<br>$35,001 - $50,000<br>$50,001 - $75,000<br>$75,001 - $100,000<br>$100,001 - $150,000<br>$150,001 - $200,000<br>over $200,000 |
| | | |
| 13 | What is/was your parents highest household income? | $15,000 or under<br>$15,001 - $35,000<br>$35,001 - $50,000<br>$50,001 - $75,000<br>$75,001 - $100,000<br>$100,001 - $150,000<br>$150,001 - $200,000<br>over $200,000 |
| | | |
| 14 | What was the highest level of education that your father (or male guardian) completed? | Less than high school<br>High School Equivalency<br>High school<br>Vocational or trade school<br>College or university |
| | | |
| 15 | What was the highest level of education that your Mother (or female guardian) completed? | Less than high school<br>High School Equivalency<br>High school<br>Vocational or trade school<br>College or university |
| | | |
| 16 | Do you work part-time, full-time, or contract? | Part-time<br>Full-time |

| | | |
|---|---|---|
| | | Contract (full time or part time) |
| | | |
| 17 | Would you call yourself an avid player of video games? | Yes<br>No |
| | | |
| 18 | What is the ownership / legal structure of your current employer? | Government<br>Private Company<br>Publicly Traded Company<br>Other |
| | | |
| 19 | Which industry does your current employer belong to? Pick one only, and choose the industry representing the main revenue source if more than one category applies. | Publishing, Broadcasting, Communications<br>Information Technology and Related Services<br>Financial Services and Insurance<br>Professional, Scientific and Technical Services<br>Government<br>Educational Services<br>Health Care and Social Assistance<br>Retail and Wholesale Trade<br>Manufacturing<br>Utilities<br>Transportation and Warehousing<br>Mining, Agriculture, Forestry, Fishing and Hunting<br>Construction, Real Estate, Rental and Leasing<br>None of the Above |
| | | |
| 20 | How many employees does your organization have? | 1-49<br>50 - 249<br>250 - 499<br>500 - 749<br>1,000 - 2,499<br>2,500 - 4,999<br>5,000 - 9,999<br>10,000 - 19,999<br>20,000 - 49,000<br>50,000 or more<br>Don't Know |
| | | |
| 21 | What was your organizations approximate annual revenue or total budget last year? | < $1 million<br>$1 million-$24 million<br>$25 million-$99 million<br>$100 million-$499 million<br>$500 million-$999 million<br>$1 billion-$1.99 billion<br>$2 billion-$10 billion<br>> $10 billion<br>Don’t Know |
| | | |
| 22 | What percentage of your organization’s employees uses a laptop to access business applications/data? | <10%<br>11-20%<br>21 - 30%<br>31 - 40%<br>41 - 50%<br>51 - 60%<br>61 - 70%<br>71 - 80% |

| | | |
|---|---|---|
| | | 81 - 90%<br>91 - 100% |
| | | |
| 23 | What percentage of your organization’s employees uses a smartphone (e.g. Blackberry, iPhone, Android device) to access business applications/data? | <10%<br>11-20%<br>21 - 30%<br>31 - 40%<br>41 - 50%<br>51 - 60%<br>61 - 70%<br>71 - 80%<br>81 - 90%<br>91 - 100% |
| | | |
| 24 | Does your company allow the use of personal mobile devices (i.e. mobile devices not provisioned directly by the company to the employee)? | Yes<br>No |
| | | |
| 25 | Please choose the job title that most closely matches your own: | Chief Executive Officer<br>Chief Technology Officer<br>Chief Information Officer<br>Chief Privacy Officer<br>Chief Security Officer<br>Chief Information Security Officer<br>VP of IT or Security or Risk Management<br>Director<br>Manager<br>Security Analyst, Consultant, Auditor<br>System Administrator<br>Other |
| | | |
| 26 | How many IT-related professional certifications do you have? | None<br>1<br>2<br>3<br>4<br>5<br>more than 5 |
| | | |
| 27 | How long have you been with your current employer? | < 6 months<br>6 months to 1 year<br>1-3 years<br>4-6 years<br>7-9 years<br>10 years or more |
| | | |
| 28 | Which range contains your current annual salary (including any bonuses)? | $15,000 or under<br>$15,001 - $35,000<br>$35,001 - $50,000<br>$50,001 - $75,000<br>$75,001 - $100,000<br>$100,001 - $150,000<br>$150,001 - $200,000<br>over $200,000 |

| | | |
|---|---|---|
| | | I prefer not to answer |
| | | |
| 29 | Approximately how many full time equivalent staff (FTEs) does your organization devote to IT privacy and security (including operations, audit and policy functions)? | 0 FTEs<br>1 FTE<br>2-10 FTEs<br>11-50 FTEs<br>51-100 FTEs<br>101 - 200 FTEs<br>201 - 500 FTEs<br>>500 FTEs |
| | | |
| 30 | What is the % share of your organization's annual total budget dedicated to information technology? | 1% - 1.5%<br>1.6% - 2%<br>2.1% - 5%<br>5.1% - 10%<br>More than 10%<br>I don't know |
| | | |
| 31 | What is the % share of your organization's annual IT Budget dedicated to information security | 1% - 2%<br>2.1% - 5%<br>5.1% - 10%<br>10.1% - 20%<br>More than 20%<br>I don't know |
| | | |
| 32 | What is the most senior role in your organization dedicated to information security? | Chief Executive Officer<br>Chief Technology Officer<br>Chief Information Officer<br>Chief Privacy Officer<br>Chief Security Officer<br>Chief Information Security Officer<br>VP of IT or Security or Risk Management<br>Director<br>Manager<br>Security Analyst, Consultant, Auditor<br>System Administrator<br>Other |
| | | |
| 33 | Does your organization outsource any aspect of its Security Program or Security Operations? | Yes<br>No |
| | | |
| 34 | How satisfied are you with your organization's overall IT security posture? | Very dissatisfied<br>Dissatisfied<br>Neutral<br>Satisfied<br>Very satisfied<br>Not Sure/Don’t Know |
| | | |
| 35 | How would you rate your organization's 'people' security control performance? | Significant weakness<br>Slight disadvantage<br>Neutral<br>Slight advantage<br>Significant strength |
| | | |

| | | |
|---|---|---|
| 36 | How would you rate your organization's 'process' security control performance? | Significant weakness<br>Slight disadvantage<br>Neutral<br>Slight advantage<br>Significant strength |
| | | |
| 37 | How would you rate your organization's 'technology' security control performance? | Significant weakness<br>Slight disadvantage<br>Neutral<br>Slight advantage<br>Significant strength |
| | | |
| 38 | How would you rate your organization's 'prevention/blocking' security control performance? | Significant weakness<br>Slight disadvantage<br>Neutral<br>Slight advantage<br>Significant strength |
| | | |
| 39 | How would you rate your organization's 'detection/monitoring' security control performance? | Significant weakness<br>Slight disadvantage<br>Neutral<br>Slight advantage<br>Significant strength |
| | | |
| 40 | How would you rate your organization's 'counter/recovery' security control performance? | Significant weakness<br>Slight disadvantage<br>Neutral<br>Slight advantage<br>Significant strength |
| | | |
| 41 | Considering your IT security environment, staffing, budget, and mandate, how would you rate the overall complexity of your IT environment? | Very Low Complexity<br>Low Complexity<br>Medium Complexity<br>High Complexity<br>Very High Complexity |
| | | |
| 42 | How many data and application servers does your organization have in its environment? | 1 - 10<br>11 - 50<br>51 - 100<br>101 - 250<br>251 - 500<br>501 - 1000<br>1000 or more<br>Don't Know |
| | | |
| 43 | How many internally reportable privacy and security incidents do you estimate your organization has experienced in the past 12 months? | 1 - 5<br>6 - 20<br>21 - 50<br>51 - 100<br>100 - 250<br>250 - 500<br>more than 500 |
| | | |
| 44 | How confident are you that your organization could detect an | Very unconfident<br>Somewhat unconfident |

| | | |
|---|---|---|
| | information security breach as it was happening? | Confident<br>Very confident<br>Extremely confident |
| | | |
| 45 | How confident are you that your organization could detect an information security breach after it had happened? | Very unconfident<br>Somewhat unconfident<br>Confident<br>Very confident<br>Extremely confident |
| | | |
| 46 | How willing are you to take risks in your professional  life? | Very unlikely<br>Not likely<br>Somewhat Likely<br>Very Likely<br>Highly likely |
| | | |
| 47 | Does your job position allow you to make independent information security decisions? | Never<br>Very Rarely<br>Sometimes<br>Frequently<br>Almost all of the time |
| | | |
| 48 | How concerned are you that a severe or significant  security incident might materialize in your organization despite the existing protective measures? | Unconcerned<br>Slightly concerned<br>Somewhat concerned<br>Very concerned<br>Highly concerned |
| | | |
| 49 | Have you personally experienced any severe or significant security incident at work in the past? | Never<br>Once<br>A few times<br>More than a few times<br>Many Times |
| | | |
| 50 | How closely related do you think investment in information security is to the achievement of overall business objectives? | Not related<br>Slightly related<br>Somewhat related<br>Very related<br>Highly related |
| | | |
| 51 | How much does your organization focus on business operations versus security? | Always favours business operations over security considerations<br>Somewhat favours business operations over security considerations<br>Treats business operations and security about equally considerations<br>Somewhat favours security considerations over business operations<br>Usually favours security considerations over business operations |
| | | |
| 52 | How likely are you to take risks in your personal life? | Very unlikely<br>Not likely<br>Somewhat Likely<br>Very Likely |

|  |  |  |
|---|---|---|
|  |  | Highly likely |
|  |  |  |
| 53 | Do you currently smoke cigarettes? | Yes<br>No |
|  |  |  |
| 54 | Do you buy lottery tickets? | Yes<br>No |
|  |  |  |
| 55 | Do you engage in high risk sports or other recreational activities? | Never<br>I have once or twice<br>I do once in a while<br>Frequently<br>Very frequently |
|  |  |  |
| 56 | Do you think skydiving is a risky activity? | Not at all risky<br>Slightly risky<br>Somewhat risky<br>Very risky<br>Extremely risky |
|  |  |  |
| 57 | Do you think flying on a commercial aircraft is a risky activity? | Not at all risky<br>Slightly risky<br>Somewhat risky<br>Very risky<br>Extremely risky |
|  |  |  |
| 58 | Do you own a life insurance policy on yourself? | Yes<br>No |
|  |  |  |
| 59 | Do you own a life insurance policy on someone else? | Yes<br>No |
|  |  |  |
| 60 | When purchasing a computer or a smartphone for personal use, do you buy an 'extended warranty'? | Never<br>Occasionally<br>Usually<br>Most of the time<br>Always |
|  |  |  |
| 61 | In general, how would your best friend describe you as a risk taker? | A real gambler<br>Willing to take risks after completing adequate research<br>Cautious<br>A real risk avoider |
|  |  |  |
| 62 | You are on a TV game show and can choose one of the following. Which would you take? | $1,000 in cash<br>A 50% chance at winning $5,000<br>A 25% chance at winning $10,000<br>A 5% chance at winning $100,000 |
|  |  |  |
| 63 | You have just finished saving for a "once-in-a-lifetime" vacation. Three weeks before you plan to leave, you lose your job. You would: | • Cancel the vacation<br>• Take a much more modest vacation<br>• Go as scheduled, reasoning that you need the time to prepare for a job search<br>• Extend your vacation, because this might be your last chance to go first-class |

| | | |
|---|---|---|
| 64 | If you unexpectedly received $20,000 to invest, what would you do? | • Deposit it in a bank account, money market account, or an insured CD<br>• Invest it in safe high quality bonds or bond mutual funds<br>• Invest it in stocks or stock mutual funds |
| | | |
| 65 | In terms of experience, how comfortable are you investing in stocks or stock mutual funds? | Not at all comfortable<br>Somewhat comfortable<br>Comfortable<br>Very comfortable<br>Extremely comfortable |
| | | |
| 66 | When you think of the word "risk" which of the following words comes to mind first? | Loss<br>Uncertainty<br>Opportunity<br>Thrill |
| | | |
| 67 | Some experts are predicting prices of assets such as gold, jewels, collectibles, and real estate (hard assets) to increase in value; bond prices may fall, however, experts tend to agree that government bonds are relatively safe. Most of your investment assets are now in high-interest government bonds. What would you do? | • Hold the bonds<br>• Sell the bonds, put half the proceeds into money market accounts, and the other half into hard assets<br>• Sell the bonds and put the total proceeds into hard assets<br>• Sell the bonds, put all the money into hard assets, and borrow additional money to buy more |
| | | |
| 68 | Given the best- and worst-case returns in each of these investment choices, which would you prefer? | $200 gain best case; $0 gain/loss worst case<br>$800 gain best case; $200 loss worst case<br>$2,600 gain best case; $800 loss worst case<br>$4,800 gain best case; $2,400 loss worst case |
| | | |
| 69 | Choose between the following: | • A sure gain of $500<br>• A 50% chance to gain $1,000 and a 50% chance to gain nothing |
| | | |
| 70 | In addition to whatever you own, you have been given $2,000. You are now asked to choose between: | A sure loss of $1000<br>A 50% chance to lose the $2,000 and a 50% chance to lose nothing |
| | | |
| 71 | Suppose a relative left you an inheritance of $100,000, stipulating in the will that you invest ALL of the money in ONE of the following investment choices. Which one would you select? | A savings account or money market mutual fund<br>A mutual fund that owns stocks and bonds<br>A portfolio of 15 common stocks<br>Commodities like gold, silver, and oil |
| | | |
| 72 | If you had to invest $20,000, which of the following portfolios would you find most appealing? | • 60% low-risk investment, 30% medium-risk investment, 10% high-risk investment<br>• 30% low-risk investment, 40% medium-risk investment, 30% high-risk investment<br>• 10% low-risk investment, 40% medium-risk investment, 50% high-risk investment |

| | | |
|---|---|---|
| 73 | Your trusted friend and neighbor, an experienced geologist, is putting together a group of investors to fund an exploratory gold mining venture. The venture could pay back 50 to 100 times the investment if successful. If the mine is a bust, the entire investment is worthless. Your friend estimates the chance of success is only 20%. If you had the money, how much would you invest? | Nothing<br>One month's salary<br>Three month's salary<br>Six month's salary |