Big data: a normal accident waiting to happen?

Article

Accepted Version

Nunan, D. and Di Domenico, M. (2017) Big data: a normal accident waiting to happen? *Journal of Business Ethics*, 145 (3). pp. 481-491. ISSN 1573-0697 doi: <https://doi.org/10.1007/s10551-015-2904-x> Available at <https://centaur.reading.ac.uk/44808/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1007/s10551-015-2904-x>

Publisher: Springer

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Big data: A normal accident waiting to happen?

Daniel Nunan, Henley Business School, University of Reading, UK.

Maria Di Domenico, Surrey Business School University of Surrey, UK.

Abstract

Widespread commercial use of the internet has significantly increased the volume and scope of data being collected by organisations. ‘Big data’ has emerged as a term to encapsulate both the technical and commercial aspects of this growing data collection activity. To date, much of the discussion of big data has centred upon its transformational potential for innovation and efficiency, yet there has been less reflection on its wider implications beyond commercial value creation. This paper builds upon normal accident theory (NAT) to analyse the broader ethical implications of big data. It argues that the strategies behind big data require organisational systems that leave them vulnerable to normal accidents, that is to say some form of accident or disaster that is both unanticipated and inevitable. Whilst NAT has previously focused on the consequences of physical accidents, this paper suggests a new form of system accident that we label *data accidents*. These have distinct, less tangible and more complex characteristics and raise significant questions over the role of individual privacy in a ‘data society’. The paper concludes by considering the ways in which the risks of such data accidents might be managed or mitigated.

Keywords: Big data, normal accident theory, privacy, system accidents

Introduction

“Don’t be evil. We believe strongly that in the long term, we will be better served — as shareholders and in all other ways — by a company that does good things for the world even if we forgo some short term gains.” Google IPO Prospectus, 2004

Over the last decade, as usage of the internet has grown so too has interest in exploitation of the data generated by an increasingly technologically connected society (George et al., 2014). In the 1998 film 'Enemy of the State' a rogue government agency is portrayed as having unlimited access to personal data from emails, telephone conversations and other data sources. At the time the scenario was so shocking to the FBI that a PR campaign was launched to reassure the public that the plot was purely fiction (Miller, 2013). Yet, as recent revelations about data and privacy have made clear, the forms of surveillance present in this fictional account are now very real. What has changed is the role of governments as collectors of data, and the increasing importance of commercial entities as the drivers of both data collection *and* analysis. Previous generations of information technology were dominated by firms with expertise in hardware or software, many leading internet firms now have commercial strategies built around the collection of data. For firms such as Google, Facebook and others, the collection of data has become an end in itself rather than a way of achieving other ancillary business goals. This comes at a time when there is increasing interest from researchers in the ethical implications of wider data use in areas such as data privacy (Hong & Thong, 2013), cyber hacking (Bambauer, 2014), government regulation (Fink et al., 2012) and intellectual property (Bateman et al., 2013). In addition there is concern over the extent

of government surveillance of commercial social networks as a result of leaks by Edward Snowden (Witte, 2013).

Within this context the term big data has gained popularity in both business and public policy circles as a label encapsulating the technical and commercial aspects of new types of personal data collection (Nunan & Di Domenico, 2013). To date, much of the discussion of big data has focused on the potential for positive impacts both for business and society (George et al., 2014). This goes beyond improvements to commercial efficiency and extends to claims about its transformational impact upon areas such healthcare and the delivery of public services (Manyika et al., 2011). Yet, there has been relatively little reflection on, or theoretical development of, the ethical issues raised by big data. This is perhaps unsurprising given the paucity of research on internet ethics (Schlegelmilch & Oberseder, 2010) and that understanding the ethical implications of new technologies lags behind their implementation (De George, 2003). This paper argues for reflection upon the wider ethical implications of big data, not least because the ‘don’t be evil’ philosophy adopted by Google in 2004 seems, a decade later, to be a naïve reflection of technology firms’ role in the collection and exploitation of personal data. At the same time this slogan captures the complex, and sometimes contradictory, ethical stances taken by organisations that have enabled this data explosion to occur.

This paper analyses the ethical issues around big data through the lens of normal accident theory (NAT). Normal accidents are normal in the sense that these negative events are inevitable and occur where organisational systems are both complex and tightly coupled (Perrow, 1984). NAT has been applied to the study of physical accidents including nuclear power stations (Perrow, 1981; 1984; Pidgeon, 2011), military ‘friendly fire’ (Snook, 2000),

plane crashes (Helmreich, 1997) and the recent failures in the financial system (Palmer & Maher, 2010). It is argued that the emerging forms of organisation that are enabling big data have the system characteristics identified by NAT. However, the consequences of normal accidents in these data-centric organisations are less immediately tangible than with a physical disaster, making identification and remedy more difficult. The paper proposes a form of *data accident* that extends the theory of normal accidents to take account of the scale, interconnectedness and uncertainty generated through the exploitation of data in contemporary society. In doing so an ethical dilemma is highlighted, whereby the consequences of such inevitable data accidents must be balanced with the utility gained by the public through the use of these technologies.

The paper is structured as follows. We begin by identifying the key aspects of big data and the types of organisation that enable, and are enabled by, this growth in data collection. Secondly, we outline normal accident theory and the characteristics of organisational systems that have been seen to enable such accidents. Thirdly, we extend NAT by suggesting both the characteristics and potential consequences of a data accident. Finally, we follow Perrow (1999) and consider both the question of what *should* be done and, given the increasing commercial significance of big data, what *can* be done to mitigate the unforeseen consequences of big data.

Introducing Big Data

Although data has become closely associated with information technology, managing and making sense of data is an age old problem. Historians, politicians and military leaders have relied on information as a source of power for centuries, whilst preventing access to information has long been a lever for removing power. However, what has changed is the

volume of data under consideration. When, in the 1850s, managers of US railroads sought new organisational designs to help overcome the challenges of the ‘data avalanche’ in a growing business, they were talking about a volume of data that did not even add up to a single megabyte in modern terms (Rosenthal, 2014). By contrast, a key characteristic of data in contemporary business is the constant, and almost exponential, growth in its volume. To use one contemporary statistic, 90% of all data in existence has been created in the last two years (IBM, 2015).

In this context big data initially emerged as a term to describe the technical innovations underpinning the massive increase in data being collected (Jacobs, 2009). Beyond enabling more data to be collected, the technology has also changed the *velocity* and *variety* of data that can be collected and analysed (IBM, 2015). Velocity refers to the speed with which data can be collected and analysed, with real-time analysis becoming a possibility with even very large data sets. Variety is significant as it signals a shift away from simply collecting data in text form towards data in video, audio and image formats (Kuechler, 2007). More recently, big data has moved beyond its technical roots to encompass a broad range of commercial opportunities enabled through the analysis of such data (Manyika et al., 2011). As it has done so, the term has been adopted by politicians as the means of achieving economic growth. In the UK big data has become one of eight key government priorities (HM Government, 2013) and in the US Barrack Obama’s use of campaign data has led to him being referred to as the ‘big data President’ (Hurwitz, 2012).

This technological innovation is paired with a number of economic factors. First, the cost of storing data has reduced to the point that it may now be economically viable to store all forms of data, even if there is no immediate use for it. A second economic factor is the widespread

availability of both the necessary software and hardware. Much of the software that supports big data has emerged not from traditional technology companies, but from organisations where necessity has driven them to find solutions to their own big data problems. Crucially, for the wider use of big data, much of this technology has been made available through open source licences, enabling other organisations to more easily use and adapt the technology for their own needs (Nunan & Di Domenico, 2013).

Generating value from big data

Clearly, big data creates the potential for a wide range of commercial possibilities and innovations. The most widely known examples originate from the major internet companies for whom data collection and analysis is a core competency. For example, as Marcus (2012) illustrates, Google was able to develop a more effective spell checker not through knowledge of natural language analysis or the psychology of spelling, but through collecting and analysing a massive database of actual spelling corrections:

“What did users most often type next after failing to find what they wanted with the word “people”? Aha, “people.”... The lesson, it seemed, was that with a big enough database and fast enough computers, human problems could be solved without much insight into the particulars of the human mind”

Other examples that work on similar principles include the recommendation engines used by Amazon and Netflix that leverage large databases of consumer preferences on books or films to make recommendations for future watching. However, despite the high profile of such cases the commercial benefits of big data are not limited to internet firms. Indeed, the McKinsey report that served to popularise big data as a commercial strategy (Manyika et al.,

2011) highlights the potential in healthcare, public service delivery in local government and financial services. For example, in a healthcare environment characterized by fragmented datasets big data has been proposed as both a mechanism to help control spiraling costs and a way of speeding up the R&D process for new pharmaceutical techniques (Groves et al., 2013). An aspect of the commercial application of big data that has, perhaps, resulted in many gains being less public is the incremental nature of the improvements to existing processes. Put another way, many of the cases of big data uses are about improving existing processes rather than inventing new businesses. A powerful example of this comes from the shift in airlines from relying on pilots to provide ETA (estimated time of arrival) information through to using a data driven system combining multiple sources including weather, radar and flight schedules (McAfee & Brynjolfsson, 2012). Whilst busy pilots would previously make estimates that were often generally accurate, at least 30% were more than five minutes out. By combining multiple data points through an automated algorithm, airlines have been able to virtually eliminate gaps between estimated and actual arrival times, saving individual airports millions of dollars per year (McAfee & Brynjolfsson, 2012).

What draws these examples together and makes them 'big data' is that value is created through large scale analysis of data, not just collection, and the combination of multiple data sets. Although these examples help to demonstrate the positive commercial benefits of big data, the question of the wider implications for big data, resulting through this commercial exploitation, remain ambiguous. Although the focus is often on the implications of the volume of information being collected, big data is less about size that is big than it is about a capacity to search, aggregate, and analyse large data sets (boyd & Crawford, 2012). That is not to say that there is not a high level of concern over the privacy implications of new technology. Concern about privacy implications of new technologies is nothing new whether

it is photography (Warren & Brandeis, 1890), personal computers (Zuboff, 1988) or the internet (Nissenbaum, 2004). However, the nature of this concern has typically lagged behind the reality of the way that new technologies use and analyse data (Mundie, 2014). More specifically, concerns over privacy have been driven by fears over active and visible data collection rather than the sorts of passive and autonomous data collection allied to big data. Essentially, privacy concerns that relate directly to big data are less well articulated because users of technology are less able to appropriately understand or contextualise the ways in which the data is being collected and analysed. In turn, as this paper explores, this lack of context is driven by the unknown nature of the privacy risks. Returning to a core premise of big data, data is being collected because it is economical to do so and the potential for valuable analysis in the future exists. Because the use of the data is unknown, so too are the broader social implications.

Normal Accident Theory

“What seems certain... is that the problem of technological determinism - that is, of the impact of machines on history-will remain germane until there is forged a degree of public control over technology far greater than anything that now exists.” –

Heilbroner, 1967:65

The advancement of technology has always created a tension between economic benefits and social impact. As Heilbroner argues, the question is not whether technology in some ways defines societies, but the extent to which it does so. Although data may now seem a key factor in the debate over the ethical implications of computer technology, this was not always the case. In arguing for a form of ethical overload Coates (1982) identifies the very wide

range of ethical issues on which computer technology might touch. This includes speech recognition, biotechnology and the impact on workplace structure of home working, although there is little mention of the role of data itself. This highlights the challenge of predicting the consequences of technology, even when the technology itself is well understood, a challenge that the author captures in the statement that in this position “the new immorality is to act in ignorance of future consequences” (Coates, 1982:239).

Although accidents and disasters have always been a characteristic of society, the complexity and embeddedness of modern technologies create an ever-greater need for understanding (Leveson et al., 2009). Normal accident theory (NAT) has its origins in an attempt to explain the consequences of such complex technologies and the need for better understanding of the causes of the catastrophic accidents they cause. The root of the theory goes back to when sociologist Charles Perrow was asked to provide a background report for the *President's Commission on The Accident at Three Mile Island*, following the 1979 nuclear accident at the power station on the site. The report found that the accident was not caused, as might have been expected, by an isolated technical fault or human error but by a series of organisational factors present in systems, such as those in a nuclear power station, that are both complex and tightly coupled (Perrow, 1981; 1984). The focus of the report was therefore not on finding *who* was to blame, but on *how* organisational systems had developed to enable such accidents to take place.

The types of systems in which normal accidents occur have two key characteristics: interactive complexity and tight coupling (Pidgeon, 2011). Interactive complexity refers to chains of events that occur in sequences that are unfamiliar or unplanned. This is as opposed to linear interactivity that occurs in an expected, familiar, visible and planned way (Perrow,

1984). Linearity implies not that the system is simple (i.e. lacking in complexity) but rather that events happen in a linear fashion. Thus, production lines for pharmaceuticals or the flight of an airplane are linear in the sense that they are processes that can be explained, but they are by no means simple. Accidents can, and do, occur in such environments, but through their linear nature the impacts of such accidents can be more easily identified and remedied. In linear systems, such as a production line, when unplanned events happen, they can be easily located and remedied by employees (Perrow, 1984). On the other hand, complex interactivity occurs where interactions are not fully understood, at least by those who have to make the time critical decisions required to mitigate against accidents (Perrow, 1984).

The second component of a normal accident, independent from complex interactivity, is the requirement for tight coupling. Where tight coupling occurs, interactions occur quickly and in a way that is unobstructed as components within a system impact each other, allowing incidents to escalate into accidents. In this context 'accident' has a specific definition associated with a major systems failure, as opposed to a more minor and routine failure that Perrow (1984) labels an incident. As such, an accident is not simply a failure of part of a system but of a system as a whole. The significance of this is that for a system accident to occur multiple failures must happen.

“It is not the source of the accident that distinguishes the two types, since both start with component failures; it is the presence or not of multiple failures that interact in unanticipated ways.” (Perrow, 1984:71)

Critiques of normal accident theory

One challenge with NAT is that the use of the term *system* isn't formally defined by Perrow, despite the concept of a system being at the heart of the theory (Shrivastava et al., 2009).

However, in using this term NAT is able to shift the focus from individual failure to the failure of systems. Another way of putting this is that individuals no longer become the cause of normal accidents, whilst instead it is the organisational systems in which the accidents occur that are to blame (Cummings, 1984).

The term normal accident is both memorable and a little misleading as an accident is only 'normal' in the sense that in a certain set of organisational circumstances accidents become inevitable. This inevitability of accidents has been the source of discussion and controversy, as high reliability theorists (HRT) have argued that it is possible to design organisations which are complex and tightly coupled yet are able to survive normal accidents (Roberts, 1990). The debate between NAT and HRT has continued since (Sagan, 1993; Rijpma, 1997; Shrivastava et al., 2009; Perrow, 2008), with multiple, largely unsuccessful, attempts to 'break through the deadlock' (Rijpma, 1997:15). Perrow's response is that most types of accidents, even complex accidents involving multiple failures, can be prevented whilst normal accidents are inevitable and therefore a feature of the system (Perrow, 1999). The central message of NAT is therefore not about risk prevention but of dealing with the consequences of normal accidents. This is a challenging message, one that calls into question the overall relationship between technology and society. However, Perrow remains unmoved on this central point:

“I have a simple message: disasters from natural, industrial and technological sources, and from deliberate sources such as terrorism, are inevitable, and increasing. We may prevent some and mitigate some, but we cannot escape them.” (Perrow, 2008:733)

It is this message of inevitability and inescapability that distinguishes NAT from other theories of technology risk, and it is this point that creates particular salience when applying to the issues around big data.

Normal accidents and big data

The argument of this paper is not that big data *causes* normal accidents, but rather that it is central in creating forms of organisation, and organisational systems, in which normal accidents are likely to occur. To make this argument we analyse big data through the core components of NAT - tight coupling and complexity. In addition, we extend the theory through consideration of the specific organisational context in which big data occurs.

The first of these components is that big data has the characteristics of a tightly coupled system. This may appear to be a counterintuitive argument considering that the internet itself is an archetype of a loosely coupled system, designed in a cold war era to survive the destruction of any one of its parts. However, behind big data lies tightly connected infrastructure and firms that bind together the organisations reliant on this technology. This is characterised by a shift away from organisations exerting full control over their technology ‘stack’ and towards cloud computing, where storage and processing power is treated as a utility. Although organisations may develop the software that runs, and collects data, to maximise efficiency, they are increasingly reliant on the use of large data centres run by firms such as Amazon, Microsoft or Google. For all but the very largest firms, the efficiency

and perceived resilience offered by these multi-billion dollar data-centres make it a necessity to use third party storage for technology. Yet, when there is a failure in data-centres, the consequences are far more widespread, unexpected and unpredictable than when firms had more control over their own infrastructures. For example, an accidental deletion of a small amount of data by a developer in the software that balances traffic between different servers using Amazon Web Services¹ on Christmas Eve 2012 resulted in a chain of events that rendered online services, including Netflix, unavailable on Christmas morning (Cockcroft, 2012). In their apology for the incident (Amazon, 2012) Amazon highlighted that only a very small number of developers had access to this data, the developer did not initially realize the mistake and, for the first few hours, the technical teams were puzzled by the error messages being generated by their system. This issue was fixed not through technical means, but by implementing a process to ensure changes to systems were double-checked to avoid accidental deletion in the future.

Secondly, related to this loose coupling, the distributed nature of the way in which big data is collected and stored creates inherent complexity with multiple organisations and multiple technologies. Going back to the example given previously of the Netflix failure, what is even more significant is that the Netflix was reliant on another organisation's infrastructure to deliver much of its content. In this case an organisation with which it competes directly, Amazon's Instant Video service. Whilst a few firms can achieve a degree of vertical integration most are reliant on an increasing web of third parties. Taking a theoretical example of a mid-size online retailer in addition to hosting the website, reliance upon other websites might include payment services (such as credit cards Paypal, Apple Pay),

¹ Amazon Web Services is a cloud storage service run by Amazon that allows other firms to make use of Amazon's data-centres and computer processing power for their websites and online services.

integration with social media, content delivery networks for video content, integration with a third party CRM, courier and delivery services, web analytics tools, and advertising servers, to name but a few. This is on top of the complexity inherent in the data itself, a complexity reflected in the increasing concern over the ability to find suitably skilled employees with the ability to effectively analyse such data (Brown et al., 2014).

Thirdly, a distinguishing characteristic of normal accidents and big data is the lack of a shared understanding over risk. Organisations that have previously been used as cases of normal accidents can be characterised by a clear understanding of what constitutes an accident, even if they may disagree with the extent to which such accidents can be considered as ‘normal’. Furthermore, previous organisations that collected very large amounts of data were often governments who had clear incentives to maintain privacy because, in the absence of a commercial incentives, ensuring limited access to the data also meant increasing levels of control. By contrast, the commercial demands of many of the firms that enable big data are built around both the indiscriminate collection of data and encouraging sharing of data. When Facebook CEO Mark Zuckerberg says ‘privacy is dead’, even looking beyond the rhetoric, they are signalling that the underlying culture relating to data risk and privacy is different. This creates a tension in the types of activity that generate commercial value that can also be responsible for the forms of normal accidents we describe in this paper.

To summarise this argument we refer back to the four types of organisational system in NAT, defined by whether there is tight or loose coupling and where interactions are complex or linear. Table 1 provides a quadrant that illustrates this categorisation of systems together a scenario of types of data-based failures that might occur in each of these categories. It is the top right-hand corner that links to normal accident theory where systems are both complex

and tightly coupled. Reflecting on this quadrant, one of the difficulties is that the more abstract nature of the concept of data makes characterisations based upon normal ‘physical’ accidents less effective. To accommodate the characteristics of big data we therefore suggest that there exists an alternative form of *data accident*.

Table 1 goes about here.

Data Accidents

It is proposed that normal accidents involving big data have certain characteristics that distinguish them from the types of physical accident accounted for by existing theories of normal accidents (Perrow, 1999; Snook, 2000). The term *data* rather than *information* is used deliberately as a key facet of big data is the collection of unstructured rather than structured data. The application of structure to data means, in technical terms, defining the characteristics of the data to be collected before the data collection process begins. For example, specifying that certain data is an image, or sound or text, or that it is a name, time or currency. The act of structuring data before collection implies that some thought has been given to the eventual use of the data. This may seem like a semantic point, but it serves to reinforce one of the key technological norms of big data: that data can be collected regardless of, and potentially without knowledge of, the purpose for which it is to be finally used.

Following from Perrow (1984) we consider data accidents through the lens of two recent examples. These are the two widely publicised cases of information leaks occurring through two individuals associated with the US government - Bradley Manning leaking US diplomatic cables to the Wikileaks site, and Edward Snowden leaking classified NSA data to

various media organisations. These examples are used not because they are direct examples for the commercial use of big data – they are not – but because they allow the exploration of questions around the characteristics of data accidents.

One point that requires clarification is the use of the term ‘accident’ to describe something that, in the Manning/Snowden case the original data leaks were not caused by activities that could be described as accidental. This use of the term ‘accident’ is one that has caused some confusion in the normal accident/high-reliability organisation debate (Leveson et al., 2009). However, the sorts of accident being referred to here are a form of system accident caused by unpredictable interactions within the system. In other words, the accident referred to is a form of failure within the system itself, not an individual act. Thus, in the case of the Snowden scenario the accident refers to the failure in a system designed, in all circumstances, to prevent leaks of data rather than the intent of Snowden himself.

As with the example given previously of Amazon (Cockcroft, 2012), these data accidents occurred through a combination of unforeseen events. The Wikileaks incident occurred through a low level employee, Bradley Manning, copying data, primarily diplomatic cables, to a CD-R and then contacting the media (Leigh, 2010). The exact mechanism through which Edward Snowden removed data from the NSA is unknown, although it involved downloading data from a PC to a simple USB thumb drive (Waterman, 2013). This apparently simple mechanism was not considered as computers containing sensitive information are supposed to have the USB ports disabled (Waterman, 2013).

Critically, the existence of both accidents was not apparent until the data actually came to be used. At time of writing, the full significance of the Snowden leaks is unknown due to

uncertainty over its extent, and whether data remains to be released. In both cases the authorities were not aware of the loss of data until it was released to the media for analysis. Additionally, it was not the work of a foreign government or a criminal conspiracy involving multiple actors. Rather, it was the result of a single - relatively junior - employee who was able to leverage possibly unavoidable characteristics of the organisation together with the ability to quickly access and transfer large amounts of data. The key question that is raised by these examples is: If individual members of staff in a high security organisational environment are able to remove and publicly distribute large amounts of unstructured data, much of it, at least in the Snowden case, classified 'Top Secret', then to what extent can it be assumed that commercial applications of big data will not suffer from the same issue?

Characteristics of data accidents

This question of the potential for commercial data accidents is difficult to estimate directly as, unlike the Snowden case, most data accidents are not public by nature. Rather, they have a number of characteristics which differentiate them from the more traditional physical accidents previously associated with NAT. These characteristics make such data accidents more difficult to observe, prevent and remedy.

The first difference relates to the lack of physical artifacts, making it difficult to identify when an accident has happened. In the examples above, as in many others, the first sign of a data accident occurs when the data has been put to use. Thus, although the incidents of a nuclear accident can be quantified by the physical location, a data accident only makes its presence felt by the subsequent (mis)use of the data.

Secondly, the impact of data accidents is typically neither geographically specific nor geographically located. Organisations collecting data may be nationally based, at least in normative legal terms, but the nature of the internet allows them to collect data on individuals from around the world with few limitations. Where data is lost, legally or otherwise, it can spread around the world quickly. This lack of geographical bounding has a number of implications in terms of jurisdiction when it comes to both limiting the impact of such accidents and preventing their spread (Allen & Overy, 2013).

Thirdly, the timeframes in which the impacts of data accidents are felt can be hard to predict. Once it has occurred, the impacts of a typical physical accident can, to an extent, be estimated. However, with data accidents the impact is only felt through the analysis of data and the dissemination of this analysis. Even here, the impact can be extended by weeks, months or even years into the future as new means of data analysis become available (Nunan & Di Domenico, 2013).

Limitations within normal accident theory

To effectively build upon the theories of normal accidents, it is necessary to acknowledge certain limitations within NAT. The first is the apparent absence of the real world forms of normal accidents suggested by the theory, suggesting that there many exist highly reliable organisations that can design-out such accidents (Leveson, et al., 2009). Shrivastava et al. (2009) suggest that the absence of the sort of system accidents posited by Perrow (1984) demonstrates that such use of examples is, in part, an attempt by Perrow to fit the data to support a theory. This is an extension of the argument made by Cummings (1984) that, in arguing for NAT, Perrow goes beyond theoretical scholarship to instead challenge us as to the extent to which we need systems that may be so complex as to be uncontrollable. One

example is the “Y2K” problem where, writing nine months before the turn of the new millennium, and thus without the benefit of hindsight, Perrow (1999:390) states:

"Y2K could be the quintessential Normal Accident of both the 20th and 21st centuries...Y2K has the potential for making a linear, loosely coupled system more complex and tightly coupled than anyone had any reason to anticipate."

Not only was the Y2K problem a normal accident that failed to happen, but also one that has at least some of the same underlying characteristics as big data. For many researchers and observers Y2K provided an archetype case for considering the moral and social impacts of technology. Yet the real world impact of Y2K was minimal (MacGregor, 2003). Thus in one sense, the argument for normal accidents is undermined by the absence of the accident (La Porte & Consolini, 1991). More recently Perrow reflects upon this in the context of the nuclear industry (Perrow, 2008). One answer is that public pressure in the face of those few accidents that happened, resulted in a regulatory regime that made it difficult for nuclear power to be established and, with the exception of France and Japan, nuclear power has remained a relatively limited power source. This regulatory structure also served to create a significantly different control structure with high costs, moratoria on the construction of new nuclear power stations, and a public ownership structure. An alternative approach, in the light of the Fukushima nuclear accident in Japan, would argue that Perrow had insufficiently long time frames. But the core contradiction remains, that limiting the potential of powerful technologies because of a hypothesised and unpredictable future is an unrealistic argument.

Responses to normal accidents

If we accept the theoretical proposition that there are at least some forms of accidents that cannot be prevented, we are returned to asking the question, as Perrow himself does: "What must be done?" (Perrow, 1984:304). For systems where the risk is catastrophic, where the consequences of failure far outweigh potential benefits, Perrow suggests abandoning such systems whilst also acknowledging that these recommendations are unlikely to be practical in the real world. However, this assumes that it is possible to evaluate the risk of new technologies. Given the essential intangibility of data accidents, evaluating risks in a way that is likely to be useful, or utilisable, is somewhere between difficult and impossible. At the same time, the increasing embeddedness of data collection within public life means that the likely direction of big data is toward even bigger data. Yet, as we have argued in this paper, it is not the bigness of big data that creates the risk of information accidents, nor is it the existence of technology that enables more efficient storage and analysis of data. Rather, the risk emerges from the ambiguity surrounding the collection of data, leading to a modern day 'gold-rush' where the company with the most data wins, even where the commercial value of that data is uncertain.

One outcome would be the acceptance of both the costs that arise from a reduction in personal privacy, as well as any benefits from the reduction in general societal information asymmetry. However, this cannot be done without ignoring the legal and regulatory frameworks governing the collection and use of customer data. The history of regulating commercial activities relating to consumers suggests that regulation happens with a rear view mirror, as for example with the regulation of pharmaceuticals and tobacco. Thus, any adaptations to the regulatory environment will happen *post hoc*. Whilst the pace of regulation is often driven by the lobbying power of the industries involved and the effectiveness of

enforcement, the relationship between data collection and regulation is a more difficult complex. Firstly, there is the recognition amongst policy makers that existing regulations are insufficient in an environment where consumers themselves are complicit in so much of the generation and use of data. Secondly, there is the question of enforceability of regulations where data has increasingly cross-border characteristics and the jurisdiction is less clear. Thirdly, as social networks and other online services begin to resemble utilities, and indeed seek to describe themselves as utilities, they begin to resemble the types of organisations that are normally regulated. Finally, there is the changing environment of regulation itself where regulation is not only becoming increasingly common, but also the nature of regulation is changing (Bygrave, 2014).

Discussion & Mitigation

Having outlined potential consequences of big data, it is necessary to acknowledge the risk that, in effect, the paper appears to be attempting to predict the future. By using examples of accidents it is possible to make the cognitive leap towards imagining what such events would look like in a contemporary context, and make a further leap towards postulating about their impact. However, it is not the purpose of this paper to speculate about specific future scenarios, and to do so would misinterpret the uncertain and unpredictable nature of normal accidents. Accidents that can be predicted can be managed, and in some way mitigated. To repeat a previous point, the term 'normal' is used to indicate inevitability. Thus, the theory is suggesting that accidents will happen - not when or how they might happen. With this in mind, the question of mitigation refers not to specific types of incidents (e.g. “preventing” another Snowden leak) but rather addressing the underlying organisational context under which normal accidents occur. This would need a requirement to decouple and simplify these organisational contexts.

Despite these avenues for some form of risk reduction, we return to the basic paradox of big data and normal accident theory. The mitigating factors illustrated require both a shift in the level of acceptance, and movement in the direction of travel of corporate strategy, that acceptance would only be created through experiencing the consequences of these data accidents. Yet, the characteristics of these accidents are highly individual and ambiguous, and often unknown.

As such, big data creates the forms of externalities that would normally be dealt with through the regulatory process. However, these very characteristics of individuality and ambiguity, together with the technical and commercial boundaries that underlie big data, have challenged the creation of effective regulations. A clear example of this can be seen in the process of updating data protection legislation within the European Union to replace the existing outdated regulations from the 1990s, and standardise legislative approaches to privacy across Europe (Ashford, 2014). Firstly, the challenges of regulating across borders has led to attempts to apply legislation to major US social networks, in turn resulting in issues ranging from practical questions over the limitations of national jurisdiction through to threats of a trade war (Farivar, 2014). Secondly, attempts to regulate the ability of firms to actually collect data in a format that enables big data analysis will raise questions over the extent to which citizens are actively opposed to such data collection. For example, if new EU legislation resulted in Facebook limiting features or even withdrawing from the European market, would consumers be likely to be grateful for their enhanced privacy, or chastise the European Union for infringing upon their online ‘lives’?

Understanding the consequences of big data is made more difficult by the question of ownership over data or, more specifically, the usage rights granted when firms collect data. Ownership itself typically depends on the mechanism through which the data is collected, and much of the discussion in terms of regulation seeks to address the forms of asymmetry that are perceived to be occurring here. For example, when individuals sign-up to online services such as social media they provide a form of blanket consent through which they agree for their data to be used for an often unspecified range of purposes for an unspecified length of time. Such sign-ups provide a form of legal consent to enable the use of data, and thus *de facto* ownership for the period that an individual is a member of that service, but are unlikely to meet the requirements of traditional research based upon informed consent. Additionally, whilst the public might expect some level of ownership rights over the data they have collected, one of the characteristics of big data that we have discussed refers to the autonomous collection of data generated *about* individuals rather than *by* individuals. The value of this autonomously collected data, for example by sensors in cars, homes or buildings, might not be immediately clear even to those organisations collecting it. Furthermore, can there be an ownership debate over commercial data that the public may not even be aware is being collected? Finally, we have the overarching issue of whether the data is personal or not, as data protection legislation typically does not give individuals rights over their data when it has been aggregated and anonymized. One of the paradoxes of big data is that the same techniques that are used to analyse data through the combination of data sets can also be used to deanonymise existing data sets, creating personal data and ownership issues where none previously existed. Despite the core role that anonymisation plays in legislation, unlike in traditional forms of data analysis – such as market research – where the goal is to produce aggregate insights, the commercial benefits of big data are driven by the need for individual level data and analysis.

Two other avenues for mitigation must be considered. One intriguing possibility is change in behaviour by the commercial organisations that are the source of many of the big data challenges identified in this paper. Here a distinction must be drawn between large technology firms in general and the specific characteristics of those firms which have data collection as an end, rather than a means to an end. For those firms with strategies that do not require them to become what we might refer to as 'data conglomerates', there is an opportunity to develop a strategic positioning away from big data. A final, potentially limiting, factor comes down to the role of consumer behaviour itself. Firms such as Facebook have presumably shifted from, to quote Mark Zuckerberg again, a 'privacy is dead' philosophy to one that at least partially recognises the importance of privacy in consumer decision making. This reflects research highlighting that consumers do take privacy into account when sharing data online, at least when it is with distrusted strangers (Johnson et al., 2012). Given the commercial drivers behind big data, this highlights that it is a reduction in acceptance from consumers that would drive the most effective longer term shift away from the big data economy.

In presenting this discussion, it is necessary to acknowledge some potential limitations. The first is that normal accident theorists have previously been proved wrong, most notably with the Y2K incident. Yet this is a weakness not of the theory, but of attempts to use the theory to predict rather than explain. In his defence, Perrow argues that many people got their analysis of big data wrong, and that more broadly this demonstrates the challenges of analyzing risks associated with business technology without sufficient understanding of the technical aspects that underpin it. A second limitation relates to whether 'big data' as a construct is useful for developing theories, or whether it is an ill-defined phrase, sufficiently vague in use to allow

for associations of technological progress without requiring precise knowledge of the nature of the progress. We accept that the use of the term 'big data' has perhaps outpaced its understanding; the term has arisen out of a very real need to understand the changes taking place in the way that data is collected and analysed in society. Thus our use of the term reflects our attempts to explain activities and behaviours that are already happening, where the term is already used, rather than attempting to invent new terms.

Conclusion

This paper discusses the increasingly important role of big data as part of organisational strategies and theorises the potential consequences of big data. The paper suggests that big data enables the forms of organisational systems that have traditionally been associated with those susceptible to normal accidents. Additionally, the nature of accidents involving data results in a number of unique and complex characteristics that can be differentiated from the forms of physical accidents that have featured in NAT research to date.

In doing so, the paper highlights a key change in the mechanism through which privacy becomes enforceable. Previously privacy could be maintained, at an individual level, through limiting access to the private sphere. For example, the system of surveillance undertaken by the Stasi in pre-1990s East Germany is often held up as a 'model' of how state surveillance was able to permeate this private sphere. Yet, what is notable is both the exceptional cost - both financial and social - and the extent to which it was unusual, even by the standards of totalitarian governments (Sebestyen, 2009). With big data, managing privacy is about limiting access to data that already exists. Because the data being collected and their use are often unknown, maintaining privacy is dependent on the legitimacy and effectiveness of

organisational processes that prevent damaging forms of sharing. The argument of this paper is that the ethical issues with big data lie not so much with its collection but with the weaknesses in organisational processes and systems that enable it.

However appealing, the use of “Orwellian” language in the context of discussions is misleading as it belies the increasing strategic importance of data to organisations. In reality big data is an inevitable by-product of modern technology-enabled consumer society. Accepting the core tenets of NAT (Perrow, 2009), data accidents have an inevitability. Avoiding them becomes not a matter of managing risk but of altering the fabric upon which modern consumer society is based. In presenting the argument in this paper we recognise the limitations of NAT but also seek to strengthen the theory within the contemporary context of accidents where technology, and the organisational systems that enable them, are increasingly digital, virtual and dispersed. By doing so the paper furthers NAT in providing a lens through which to unpack the role of new technologies upon society. As the social perception of new technologies such as nuclear power turned from wonder to concern, we suggest the same will be the case for internet technology. Above all, the concept of NAT serves to remind organisations that they retain the choice of what data to collect, when to retain it, and how to build up stores of trust with their stakeholders to better navigate the consequences of any data accidents.

References

- Altheim, M. (2011) The data privacy debate in social media market research: a legal perspective, *Market Research in the Mobile World*, 18 October. Available at: <http://www.mrmw.net/MRMW-News-Blogs/the-data-privacy-debate-in-social-media-market-research-a-legal-perspective.html>
- Amazon (2012). Summary of the December 24, 2012 Amazon ELB Service Event in the US-East Region. Available at: <http://aws.amazon.com/message/680587/> [Accessed 2 April 2015]
- Ashford, W. (2014). *Infosec 2014: Act now, but no new EU data protection law before 2017, says ICO*. [online] Available at: <http://www.computerweekly.com/news/2240219908/Infosec-2014-Act-now-but-no-new-EU-data-protection-law-before-2017-says-ICO> [Accessed 10 Jun. 2014].
- Allen & Overy, (2013). *Big data - Annual review 2013 - Allen & Overy*. [online] Available at: <http://www.allenoverly.com/publications/en-gb/annualreview2013/global-local/Pages/Big-data.aspx> [Accessed 10 Jun. 2014].
- Bambauer, D. (2014). Ghost in the Network. *University of Pennsylvania Law Review*, 162(5).
- Bateman, C., Valentine, S. and Rittenburg, T. (2013). Ethical Decision Making in a Peer-to-Peer File Sharing Situation: The Role of Moral Absolutes and Social Consensus. *Journal of Business Ethics*, 115(2), pp.229-240.

boyd, d. and Crawford, K. (2012). Critical Questions for Big Data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, Communication & Society*. 15(5), pp. 662-679.

Brown, B., Court, D. and McGuire, T. (2014). *Views from the front lines of the data-analytics revolution*. [online] Mckinsey.com. Available at: http://www.mckinsey.com/insights/business_technology/views_from_the_front_lines_of_the_data_analytics_revolution [Accessed 10 Jun. 2014].

Bygrave, L. (2014). *Data Privacy Law: An International Perspective*. Oxford: Oxford University Press.

Coates, J. (1982). Computers and business ? A case of ethical overload. *Journal of Business Ethics* 1(3), pp.239-248.

Cockcroft, A. (2012). *The Netflix Tech Blog: A Closer Look At The Christmas Eve Outage*. [online] Available at: <http://techblog.netflix.com/2012/12/a-closer-look-at-christmas-eve-outage.html> [Accessed 10 Jun. 2014].

Cummings, L. (1984). Normal Accidents: Living with High-Risk Technologies. Book review. *Administrative Science Quarterly*. 29(4). pp.630-632.

De George, R. (2003) *The Ethics of Information Technology and Business*. Blackwell Publishing. Oxford, UK.

Farivar, C. (2014). *EU data protection reform could start 'trade war', US official says (Wired UK)*. [online] Available at: <http://www.wired.co.uk/news/archive/2013-02/01/eu-data-protection-us-trade-war> [Accessed 10 Jun. 2014].

Fink, M., Harms, R. and Hatak, I. (2012). Nanotechnology and Ethics: The Role of Regulation Versus Self-Commitment in Shaping Researchers' Behavior. *Journal of Business Ethics*, 109(4), pp.569--581.

George, G., Haas, M. and Pentland, A. (2014). Big Data and Management. *Academy of Management Journal*, 57(2), pp.321-326.

Google IPO Prospectus (2004). 2004 Founders' IPO Letter. Available online at: <https://investor.google.com/corporate/2004/ipo-founders-letter.html> (accessed 2nd June 2015).

Groves, P., Kayyali, B., Knott, D. and Van Kuiken, S. (2013). The 'big data' revolution in healthcare. Accelerating value and innovation. *Center for US Health System Reform, McKinsey & Company*.

Heilbroner, R. (1967). Do machines make history? *Technology and Culture*, 8(3), pp.335-345.

Helmreich, R. (1997). Managing human error in aviation. *Scientific American*, 276(5), p.62.

HM Government. (2013). *UK data capability strategy: seizing the data opportunity - Publications - GOV.UK*. [online] Available at: <https://www.gov.uk/government/publications/uk-data-capability-strategy> [Accessed 10 Jun. 2014].

Hong, W. and Thong, J. (2013). Internet Privacy Concerns: An Integrated Conceptualization and Four Empirical Studies. *MIS Quarterly*, 37(1), pp.275-298.

Hurwitz, J. (2012). *The Making of a (Big Data) President*. [online] Businessweek.com. Available at: <http://www.businessweek.com/articles/2012-11-14/the-making-of-a-big-data-president> [Accessed 10 Jun. 2014]

Johnson, M., Egelman, S. and Bellovin, S. (2012). Facebook and privacy: it's complicated. *Symposium on Usable Privacy and Security (SOUPS) 2012*, July 11–13.

La Porte, T., and Consolini, P. 1991 'Working in practice but not in theory: Theoretical challenges of high-reliability organizations'. *Journal of Public Administration Research and Theory* 1, pp.19–47.

Leigh, D. (2010) How 250,000 US embassy cables were leaked. *The Guardian*, 28 November 2010.

Leveson, N., Dulac, N., Marais, K. and Carroll, J. (2009). Moving beyond normal accidents and high reliability organizations: a systems approach to safety in complex systems. *Organization Studies*, 30(2-3), pp.227-249.

Linton, D. (1992). The Luddites: How did they get that bad reputation? *Labor History*, 33(4), pp.529-537.

IBM (2015). Bringing big data to the enterprise: what is big data? [online] Available at: <http://www.ibm.com/software/data/bigdata/> [accessed 9thth April 2015].

Jacobs, A. (2009) Pathologies of Big Data. *ACM Queue*, 7(6), July.

Kuechler, W. (2007). Business applications of unstructured text. *Commun. ACM*, 50(10), pp.86-93.

MacGregor, D. (2003). 10 Public response to Y2K: social amplification and risk adaptation: or, “how I learned to stop worrying and love Y2K”. in *The social amplification of risk*, ed Pidgeon, N., Kasperson, R and Slovic P. (New York, Cambridge University Press 2003) p.243.

Manyika, J., Chui, M., Brown, B., Bughin, J., Dobbs, R., Roxburgh, C. and Byers, A. (2011). Big data: The next frontier for innovation, competition, and productivity [online]. Available at: http://www.mckinsey.com/Insights/MGI/Research/Technology_and_Innovation/Big_data_The_next_frontier_for_innovation [Accessed 10 Jun. 2014].

Marcus, G. (2012). The web gets smarter. *The New Yorker*. [online]. Available at: <http://www.newyorker.com/culture/culture-desk/the-web-gets-smarter> [Accessed 2 April 2014].

McAfee, A. and Brynjolfsson, E. (2012). Big Data: The Management Revolution. *Harvard Business Review*. October. 90(10), pp. 60-68

Miller, Z. (2013). *Former NSA Chief Was Worried About "Enemy Of The State" Reputation* | *TIME.com*. [online] TIME.com. Available at: <http://swampland.time.com/2013/06/07/former-nsa-chief-was-worried-about-enemy-of-the-state-reputation/> [Accessed 10 Jun. 2014].

Mundie, C. (2014) Privacy pragmatism. *Foreign Affairs*. March/April. Vol. 93 Issue 2, pp. 28-38.

Nissenbaum, H. (2004) Privacy as Contextual Integrity. *Washington Law Review*, 79(1).

Nunan, D. and Di Domenico, M. (2013). Market Research and the Ethics of Big Data. *International Journal of Market Research*. 55(4). pp.2-13.

Palmer, D. and Maher, M. (2010). A normal accident analysis of the mortgage meltdown. *Research in the Sociology of Organizations*, pp.219-256.

Perrow, C. (1981). Normal accident at Three Mile Island. *Society*, 18(5), pp.17-26.

Perrow, C. (1984). *Normal Accidents: Living with High Risk Technologies*. New York: Basic Books.

Perrow, C. (1999) *Normal Accidents: Living with High Risk Technologies*. 2nd Edition. New Jersey: Princeton University Press.

Perrow, C. (2008). Disasters evermore? Reducing our vulnerabilities to natural, industrial, and terrorist disasters. *Social Research: An International Quarterly*, 75(3), pp.733-752.

Pidgeon, N. (2011). In retrospect: normal accidents. *Nature*, 477(7365), pp.404-405.

Rijpma, J. (1997). Complexity, tight--coupling and reliability: Connecting normal accidents theory and high reliability theory. *Journal of Contingencies and Crisis Management*, 5(1), pp.15-23.

Roberts, K. (1990). Some characteristics of high reliability organizations. *Organization Science*, 1, pp.160–177.

Rosenthal, C. (2014). *Big data in the age of the telegraph*. [online]. Available at: http://www.mckinsey.com/insights/organization/big_data_in_the_age_of_the_telegraph [Accessed 10 Jun. 2014].

Sagan, S. (1993). *The limits of safety: Organizations, accidents, and nuclear weapons*. Princeton, NJ: Princeton University Press.

Schlegelmilch, B. and Oberseder, M. (2010). Half a century of marketing ethics: Shifting perspectives and emerging trends. *Journal of Business Ethics*, 93(1), pp.1-19.

Sebestyen, V. (2009). *Revolution 1989*. 1st ed. New York: Pantheon Books.

Shrivastava, S., Sonpar, K. and Pazzaglia, F. (2009). Normal accident theory versus high reliability theory: a resolution and call for an open systems view of accidents. *Human Relations*, 62(9), pp.1357-1390.

Snook, S. (2000). *Friendly fire*. 1st ed. Princeton, N.J.: Princeton University Press

Warren, S. and Brandeis, L. (1890) The right to privacy. *Harvard Law Review*, IV(5).

Waterman, S. (2013). *NSA leaker Ed Snowden used banned thumb-drive, exceeded access*. [online] The Washington Times. Available at: <http://www.washingtontimes.com/news/2013/jun/14/nsa-leaker-ed-snowden-used-banned-thumb-drive-exce> [Accessed 10 Jun. 2014].

Witte, D. (2013). Privacy Deleted: Is It Too Late To Protect Our Privacy Online? *Journal Of Internet Law* 18(1), pp.1-28

Zuboff, S. (1988) *In the Age of the Smart Machine: Machine: The Future of Work and Power*. Basic Books, NY.