

The face value of arguments with and without manipulation

Article

Published Version

Creative Commons: Attribution 4.0 (CC-BY)

Open Access

Felgenhauer, M. and Xu, F. ORCID: <https://orcid.org/0000-0002-3489-6824> (2020) The face value of arguments with and without manipulation. *International Economic Review*. ISSN 0020-6598 doi: <https://doi.org/10.1111/iere.12479> Available at <http://centaur.reading.ac.uk/92336/>

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/10.1111/iere.12479>

Publisher: Wiley


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



THE FACE VALUE OF ARGUMENTS WITH AND WITHOUT MANIPULATION*

BY MIKE FELGENHAUER AND FANGYA XU¹

Brunel University London, UK; University of Reading, UK

A sender wishes to persuade a receiver with a (surprising) result that challenges the prior belief. The result stems either from sequential private experimentation or manipulation. The incentive to experiment and to manipulate depends on the quality threshold for persuasion. Higher thresholds make it harder to find a surprising outcome via experimentation and may encourage manipulation. Suppose there are observable nonmanipulable and manipulable research methods. For the decision quality, the quality threshold for persuasion for non-manipulable methods should be higher than for manipulable methods. We discuss philosophy of science implications, such as field contingent quality standards and P -value adjustments.

1. INTRODUCTION

Many proposals for raising hurdles to the publication of experimental results, as in Benjamin et al. (2018), ignore the effect that such proposals will have on experimenter behavior. Higher quality requirements could discourage potentially valuable experimentation and encourage manipulation. The net effect could be to lower the value of submitted papers.

Suppose a researcher (sender) can sequentially run private experiments and knows that if he reveals a surprising outcome that goes against the prior belief, then an editor (receiver) publishes this outcome. The researcher cares about publication, though less so if the publication is a result of a false positive. The editor wants to make the correct decision. The sender can also privately manipulate to achieve the desired outcome. The difference between manipulation and experimentation is that manipulation produces an outcome that is unrelated to an (unknown) decision-relevant state of the world, whereas an experiment yields an informative outcome. Manipulation is costly to the sender. In practice, manipulation costs can result from expected punishment costs and they depend on the research method. For example, manipulating privately collected data for a regression can hardly be detected and expected punishment costs are low. If a regression is run on publicly available data instead, then manipulation is easier to detect and expected punishment costs should be higher.

We study how the face value of the evidence required for publication affects the sender's behavior and the quality of the publication process. The face value of a revealed outcome corresponds to the precision of an experiment that can generate this outcome.² For example, when reviewing a paper containing a regression, the editor can assess the quality of the regression in the manuscript (i) conditional on this regression being run on nonmanipulated data and (ii) conditional on this specification being the only regression run. The face value

*Manuscript received February 2019; revised June 2020.

¹ Please address correspondence to: Mike Felgenhauer, Brunel University London, Marie Jahoda, Uxbridge, UB8 3PH, UK. E-mail: mike.felgenhauer@brunel.ac.uk.

² Formally, an experiment's precision is the probability with which its outcome correctly predicts the state of the world. For a nonmanipulable research method the face value of the revealed outcome is equal to the precision of the experiment that generated this outcome. For a manipulable method the revealed outcome may either stem from an experiment with a certain precision or from manipulation, which pretends that it stems from an experiment with this precision. For a manipulable research method the face value of a revealed outcome is equal to this (actual or faked) precision.

of any revealed outcome is observable, by reading the manuscript, but whether the revealed outcome stems from an experiment or manipulation is not observable. We assume that the sender can only run experiments with the face value that is required for publication and that the receiver only publishes a surprising outcome with this face value.³

The article characterizes the face value requirement for publication that maximizes the quality of the receiver's decision. We find that the decision quality may not be maximized with a demanding high face value. High-quality requirements make it hard to find a surprising result that goes against the prior belief via experimentation. Not experimenting or manipulation may then be more attractive for the sender.

We find that an increase of the face value may encourage a switch from experimentation to manipulation or vice versa. In spite of this ambiguity, our major result shows that the face value requirements for manipulable research methods should be weakly below those for non-manipulable methods.⁴ In order to illustrate the intuition, suppose that manipulation costs are sufficiently low such that manipulation at the outset is optimal for the sender at the face value that maximizes the decision quality for nonmanipulable outcomes. The question is how the face value can be adjusted to make informative experimentation more attractive than uninformative manipulation. If the sender's benefit from experimentation decreases in the face value, for example, because he opportunistically cares more for a favorable decision than about the state of the world, then experimenting becomes more profitable if the face value is reduced: With a lower face value it is more likely to obtain a surprising experimental outcome by chance that can be used for persuasion. If the sender's benefit from experimentation increases in the face value instead, for example, because he cares more about the right decision, then the face value should be as high as possible for nonmanipulable outcomes. Increasing the face value for manipulable outcomes in order to encourage experimentation is then not possible.

We argue that there should be different quality standards in scientific fields that use different research methods: The decision quality maximizing publication standards should depend on experimentation and manipulation costs associated with these methods. We also argue that a reduction of the default P -value threshold for statistical significance for claims of new discoveries, as proposed by 72 authors in Benjamin et al. (2018), may deteriorate the quality of the publication process by increasing the number of published manipulated articles.

2. LITERATURE

This article is part of the persuasion literature in which a sender discloses information to a receiver who then makes a decision that affects the sender's well-being (e.g., Jovanovic, 1982; Milgrom and Roberts, 1986; Glazer and Rubinstein, 2001; Dzuida, 2011). The article belongs to a branch of this literature that combines persuasion with information acquisition via experimentation.⁵ Many of these papers study public experimentation, where the receiver observes the experimentation history (e.g., Kamenica and Gentzkow, 2011; Henry and Ottaviani, 2019). Henry and Ottaviani (2019) find that it may be in the receiver's interest to commit to a low approval standard to achieve desirable stopping behavior of the sender. This is related to our point that a too challenging face value requirement for persuasion may deter experimentation and encourage manipulation. Manipulation is not considered in Henry and Ottaviani (2019).

³ It can be shown that the results are not affected if the receiver sets the face value requirement and the sender makes history-dependent precision choices.

⁴ The research method is observable (by reading the manuscript). A method can be viewed as nonmanipulable if manipulation costs are sufficiently high such that manipulation does not occur at any face value. For a manipulable method, manipulation occurs for some face values. Many papers in the literature assume nonmanipulable methods and we consider them as an interesting benchmark.

⁵ The article relates to strategic experimentation as in Rothchild (1974), Aghion, et al. (1991), Bolton and Harris (1990), Keller, et al. (2005), and Rosenberg, et al. (2007). A survey on these "bandit problems" is Bergemann and Välimäki (2008). Experimentation is also studied in the literature on the classical problem of sequential analysis (as in Wald, 1947; Moscarini and Smith, 2001).

Other contributions, as this article, focus on private experimentation, where the receiver cannot observe the experimentation history (e.g., Brocas and Carillo, 2007; Henry, 2009; Felgenhauer and Schulte, 2014; Felgenhauer and Loerke, 2017). Private experimentation with selective information revelation is a natural assumption for logical arguments or a regression analysis. Henry (2009) and Brocas and Carillo (2007) study settings where the receiver knows or deduces the number of the sender's experiments. Skeptical beliefs as in Milgrom and Roberts (1986) induce unravelling and the receiver obtains access to the same information as under public experimentation. If private experimentation is sequential instead, that is, the decision to continue experimenting is history dependent (as in Celik, 2003; Felgenhauer and Schulte, 2014; Felgenhauer and Loerke, 2017; and here), then the receiver in general cannot deduce the actual number of experiments, even though she anticipates the experimentation plan.⁶ Skeptical beliefs are not always helpful and, in general, communication is not fully revealing. Most closely related are Felgenhauer and Schulte (2014) and Felgenhauer and Loerke (2017), in the following FS and FL, respectively.

FS investigate nonmanipulable arguments that are generated via sequential private experimentation. They study how many arguments with a given face value are used for persuasion. In an application they rationalize restrictions on scientific methods. This article studies research methods where manipulation is possible. It characterizes how the quality of the receiver's decision depends on the face value required for persuasion. In contrast to FS, the focus here is on the relation between experimentation and manipulation and a comparison of different types of research methods.

FL build on FS and endogenize the design of the experiments.⁷ FL determine the set of equilibria with persuasion under private experimentation that are not Pareto dominated. These equilibria differ regarding the precision of the experiments. For expositional convenience, this article instead assumes that the face value required for persuasion can be set. It addresses the normative question what the face value requirement should be for different research methods.

Manipulation is related to the cheap talk literature (e.g., Crawford and Sobel, 1982), as a manipulated outcome does not have an inherent meaning. Strulovici (2017) studies the impact of compensation schemes on experimentation and manipulation, that is, manipulable methods in our sense. We instead derive face value requirements for persuasion that maximize the decision quality. Felgenhauer and Xu (2019) extend and complement this article by studying the informative content of manipulable results contingent on the state of the debate.

Finally, this article is related to the economic literature on academic research (e.g., Stern, 2004; Aghion et al., 2008; Lewis and Ottaviani, 2008; Olszewski and Sandroni, 2011) and the philosophy of science (e.g., Popper, 1959; Kuhn, 1970).

3. ASSUMPTIONS

There is a state of the world $\omega \in \Omega$, with $\Omega = \{\omega_1, \omega_2\}$. The ex ante probability that the state is ω_1 is $\text{prob}\{\omega = \omega_1\} = \mu_0$, with $\mu_0 \in (0, 1/2]$. There is a sender and a receiver. The receiver chooses action $a \in A$, with $A = \{a_1, a_2\}$.

⁶ Non-history-dependent experimentation implies a commitment problem. If the sender commits to a number of experiments, but finds too many adverse outcomes, then he anticipates that persuasion is impossible by running the final experiments. As experimentation is private, it is unclear why he should run the remaining costly experiments. Baliga and Ely (2016) on the other hand investigate a repeated receiver-sender framework where the receiver has to make history-dependent decisions.

⁷ Felgenhauer (2019) introduces costly verification and shows that communication breaks down under public experimentation, but it is possible under private experimentation.

3.1. Sender's Gross Payoff. The sender obtains gross utility

	$\omega = \omega_1$	$\omega = \omega_2$
$a = a_1$	1	θ
$a = a_2$	0	0

with $\theta \in [0, 1]$ being his type.⁸ The sender prefers action a_1 regardless of the state. However, if $\theta < 1$ and the state is ω_2 , then he does not benefit from a_1 as much as in state ω_1 . In an academic context θ could be interpreted as a researcher's integrity. If $\theta = 1$, then he values a publication ($a = a_1$) the same regardless of whether the result is "true" or "false", that is, regardless of whether it reflects the state. If $\theta < 1$, then he feels less comfortable with a publication if the published outcome is "false" ($\omega = \omega_2$) than if the outcome is "true" ($\omega = \omega_1$).⁹ θ may depend on the career stage and affiliation of the researcher, as publication pressure depends on both.

3.2. Experimentation, Manipulation, and Messages. The sender has access to an experimentation technology that can generate signals about ω . He can run as many experiments as desired. The outcome of an experiment τ is $\sigma_\tau \in \{s_1, s_2\}$. s_1 is called a "positive outcome" and s_2 an "adverse outcome." The precision of each experiment that the sender may run is π , with $\pi = \text{prob}\{\sigma_\tau = s_i \mid \omega = \omega_i\}$ and $\pi \in (1/2, 1]$ for all $i \in \{1, 2\}$. The experimentation technology is, thus, symmetric in the sense that the probability that the outcome of an experiment correctly predicts the state is the same in both states. All experiments have the same precision.

Let $h_t = \{\sigma_j\}_{j=1, \dots, t}$ be an experimentation history that the sender observes and which contains the outcomes of t experiments and let $h_0 = \emptyset$ be the history if no experiment is run. The posterior probability that the state is ω_1 given some history h_t is $\text{prob}\{\omega = \omega_1 \mid h_t\}$. Let μ_t be the posterior $\text{prob}\{\omega = \omega_1 \mid h_t\}$ if h_t exclusively contains t adverse outcomes. Running an experiment costs $c_E > 0$. Experimentation costs have to be subtracted from the sender's gross payoff.

The sender sends a message $m \in \{s_1, s_2, \emptyset\}$. Given history h_t it is feasible to send $m = s_i$ if $\sigma_j = s_i$ for some $\sigma_j \in h_t$. The sender may also manipulate at costs $c_M > 0$. If the sender manipulates, then he creates an outcome s_i that is not informative regarding ω and it is feasible to send message $m = s_i$. Manipulation is not observable.

We exclude parameters for which the sender prefers doing nothing to experimentation for all π , which is the case if experimentation costs are too high.¹⁰

ASSUMPTION A. Assume $c_E \leq \max(\mu_0/2 + (1 - \mu_0)\theta/2, \mu_0)$.

The sender's behavior does not depend on π if $c_E > \max(\mu_0/2 + (1 - \mu_0)\theta/2, \mu_0)$.

3.3. Receiver Behavior and Decision Quality. We assume a nonstrategic receiver who chooses a_1 if she observes message $m = s_1$ with face value π and a_2 otherwise. The observable face value π is, thus, also a quality requirement for persuasion. The objective is to determine the face value requirements for persuasion (considered as exogenous by the players, for

⁸ The utility from choosing a particular action also depends on the state in Herresthal (2017), who compares public and private experimentation (without manipulation).

⁹ We do not consider $\theta < 0$, in which case the researcher would be better off without publishing his results $a = a_2$ if they are false $\omega = \omega_2$ compared to publishing them. This reduces the number of case distinctions. Such types do not manipulate if they know that $\omega = \omega_2$, however, they may manipulate if there is uncertainty regarding the state.

¹⁰ Assumption A ensures that this is the case by excluding parameters for which inequality (2) below is reversed at $\pi = 1/2$ and in addition at $\pi = 1$. This assumption is not critical for the results.

simplicity) that maximize the decision quality. Decision quality is defined as the probability that the receiver's action matches the state, $prob\{a_i = \omega_i\} \equiv prob\{\omega = \omega_1\}prob\{m = s_1 \mid \omega = \omega_1\} + prob\{\omega = \omega_2\}prob\{m \neq s_1 \mid \omega = \omega_2\}$. Denote the face value that maximizes the decision quality for nonmanipulable methods by π_E and the face value that maximizes the decision quality for manipulable methods by π_M .

Maximizing the decision quality corresponds to maximizing the receiver's expected utility if she has the following ex post utility:

	$\omega = \omega_1$	$\omega = \omega_2$
$a = a_1$	1	0
$a = a_2$	0	1.

The source of conflict between the players is that the receiver prefers the "correct" decision not to publish when the state is ω_2 and the sender prefers publication in this state. As $\mu_0 \in (0, 1/2]$, state ω_1 is considered as less likely than ω_2 . A positive experimental outcome is thus also less likely (or more "surprising") than an adverse outcome. The receiver's decision rule only to publish a positive outcome reflects that an editor may be more willing to publish a surprising result that goes against the prior.

3.4. Timing. The sender moves first. For each experimentation history h_t he makes the history-dependent choice to run a further experiment or to stop experimenting.¹¹ If he stops experimenting at h_t , then he chooses whether to manipulate. Then, the sender sends message m . Finally, the receiver chooses a .

4. ONE-SHOT EXPERIMENTATION

We now illustrate key effects if the sender can run at most one experiment and only manipulate at the outset. We discuss π_E and π_M and show that $\pi_M \leq \pi_E$.

4.1. Nonmanipulable Methods. The decision quality $prob\{a_i = \omega_i\}$ depends on whether the sender experiments or not. If the receiver's decision is based on an experimental outcome, then $prob\{a_i = \omega_i\} = \pi$ and the decision quality increases in π . If there is no experiment and the receiver chooses against the sender a_2 , then $prob\{a_i = \omega_i\} = 1 - \mu_0$. The experiment should only be run if $\pi \geq 1 - \mu_0$.

The sender's expected utility from running one experiment is $EU^1 = \mu_0\pi + (1 - \mu_0)(1 - \pi)\theta - c_E$ and he experiments if his participation constraint $EU^1 \geq 0$ is satisfied.

OBSERVATION 1. EU^1 decreases in π if $\theta \geq \frac{\mu_0}{1 - \mu_0}$ and it increases otherwise.

Increasing an experiment's precision π makes it more likely that publication occurs in the less likely state ω_1 , which the sender desires, but makes it less likely that it occurs in the more likely state ω_2 , which he also desires at θ . Thus, if θ is high or if ω_2 is very likely (μ_0 is low), increasing π is detrimental to the sender. Otherwise, the sender benefits from an increase of π . A sender with a high θ ($\geq \frac{\mu_0}{1 - \mu_0}$) can be interpreted as opportunistic. A sender with a low θ ($< \frac{\mu_0}{1 - \mu_0}$) can be viewed as sincere. The players' preferences are misaligned if the sender is opportunistic and they are aligned if he is sincere.

The decision quality is maximized at $\pi_E = 1$ if experimentation costs are sufficiently low such that the sender's participation constraint $EU^1 \geq 0$ is satisfied at $\pi_E = 1$ (regardless of

¹¹ If the sender never stops experimenting, then his gross payoff is zero.




FIGURE 1

MANIPULATION AT ALL π WITH A SINCERE SENDER

whether he is opportunistic or sincere). This is the case if $\mu_0 \geq c_E$. Suppose in the following $\mu_0 < c_E$.

Consider a sincere sender, whose EU^1 increases in π . If his participation constraint is violated at $\pi = 1$ (which is the case if $\mu_0 < c_E$), then it is also violated at all lower π . Parameters for which the sender does not experiment for all π are excluded by Assumption A. Assumption A and $\mu_0 < c_E$, therefore, imply that the sender has to be opportunistic.

Consider an opportunistic sender, whose EU^1 decreases in π . In the set of π for which $EU^1 \geq 0$ (the precisions for which the sender experiments), the decision quality $\text{prob}\{a_i = \omega_i\} = \pi$ is maximized at the π where the participation constraint is binding. We have $EU^1 = 0$ at $\pi = \frac{c_E - (1 - \mu_0)\theta}{\mu_0 - (1 - \mu_0)\theta}$. But should the receiver's decision be based on an experiment with this precision? The prior is against the sender and this π may be very low. The decision quality if there is an experiment is only higher than the decision quality without an experiment if $\pi \geq 1 - \mu_0$. Therefore, we have $\pi_E = \frac{c_E - (1 - \mu_0)\theta}{\mu_0 - (1 - \mu_0)\theta}$ if $\frac{c_E - (1 - \mu_0)\theta}{\mu_0 - (1 - \mu_0)\theta} \geq 1 - \mu_0$. Otherwise, experimentation should be deterred with a high $\pi_E \in (\frac{c_E - (1 - \mu_0)\theta}{\mu_0 - (1 - \mu_0)\theta}, 1]$ such that the participation constraint is violated.

The more opportunistic the sender (higher θ), the higher is $\pi_E = \frac{c_E - (1 - \mu_0)\theta}{\mu_0 - (1 - \mu_0)\theta}$, as he then has a greater benefit from persuasion and is indifferent between experimentation and not experimenting at a higher π . Furthermore, $\pi_E = \frac{c_E - (1 - \mu_0)\theta}{\mu_0 - (1 - \mu_0)\theta}$ increases in μ_0 (the less surprising a positive outcome): A higher μ_0 increases EU^1 , treating π_E as a parameter, which relaxes the participation constraint. Only an increase of π_E can make it binding again for an opportunistic sender.¹²

4.2. Manipulable Methods. The sender's expected utility from manipulation is $EU^M = \mu_0 + (1 - \mu_0)\theta - c_M$. If manipulation costs c_M are sufficiently high such that $EU^M \leq \mu_0\pi_E + (1 - \mu_0)(1 - \pi_E)\theta - c_E$, then we have $\pi_M = \pi_E$. Suppose manipulation is optimal for the sender at π_E . Let us check whether and how a change of π may encourage a switch to informative experimentation.

Consider a sincere sender, whose EU^1 increases in π . If manipulation is optimal at $\pi_E = 1$, then there is no $\pi \in (1/2, 1)$ that encourages experimentation, as illustrated in Figure 1.

Next, consider an opportunistic sender and suppose that manipulation is not optimal for all π . Since EU^1 decreases in π for an opportunistic sender, only a reduction of π may

¹² $\frac{d\pi_E}{d\theta} = \frac{(c_E - \mu_0)(1 - \mu_0)}{(\mu_0 - \theta + \theta\mu_0)^2} > 0$, as $\mu_0 < c_E$. $\frac{d\pi_E}{d\mu_0} = \frac{\theta(1 - c_E) - c_E}{(\mu_0 - \theta + \theta\mu_0)^2} \geq 0$: According to Assumption A we have $c_E \leq \max(\mu_0/2 + (1 - \mu_0)\theta/2, \mu_0)$. As $\mu_0 < c_E$, this simplifies to $c_E \leq \mu_0/2 + (1 - \mu_0)\theta/2 \Leftrightarrow \theta \geq \frac{2c_E - \mu_0}{1 - \mu_0}$. Substituting $\theta \geq \frac{2c_E - \mu_0}{1 - \mu_0}$ into the numerator of $\frac{d\pi_E}{d\mu_0}$ yields $\theta(1 - c_E) - c_E \geq \frac{2c_E - \mu_0}{1 - \mu_0}(1 - c_E) - c_E = \frac{(1 - 2c_E)(c_E - \mu_0)}{1 - \mu_0} \geq 0$, as $c_E \leq 1/2$ by Assumption A and $\mu_0 < c_E$.




FIGURE 2

 $\pi_M < \pi_E$ WITH AN OPPORTUNISTIC SENDER

induce experimentation. A lower precision increases the chance to obtain a positive outcome from an experiment given the adverse prior, but the sender's benefit from manipulation does not depend on π . Therefore, $\pi_M < \pi_E$. The decision quality is maximized at the π that solves $EU^M = EU^1$, that is, at $\pi_M = \frac{\mu_0 + c_E - c_M}{\mu_0 - \theta(1 - \mu_0)}$.

In Figure 2, higher manipulation costs shift EU^M downward and the intersection of EU^M and EU^1 moves right. Thus, the decision quality maximizing face value requirement for persuasion increases in manipulation costs.

Analogous to the previous section, the receiver would be best off by choosing against the sender if the sender only experiments at π that are too low given the adverse prior belief or if c_M are sufficiently low such that $EU^M > EU^1$ for all π .

5. SEQUENTIAL EXPERIMENTATION AND MANIPULATION

We now show $\pi_M \leq \pi_E$ if there are no restrictions on the number of experiments and with history-dependent experimentation and manipulation choices.

5.1. Nonmanipulable Methods.

5.1.1. *Sender behavior.* Sequential rationality implies that the sender stops experimenting if he faces some experimentation history h_t that contains a positive outcome. He has then found persuasive evidence and induces his preferred decision a_1 by sending message $m = s_1$. It may be sequentially rational to continue experimenting if the history contains exclusively adverse outcomes. Define EU_t^1 as the sender's continuation utility at history h_t from running one further experiment and then stopping after either outcome. Consider any history h_t that contains exclusively adverse outcomes. Stopping unsuccessfully at h_t , yielding decision a_2 , is sequentially rational if

$$(1) \quad EU_t^1 = \mu_t \pi + (1 - \mu_t)(1 - \pi)\theta - c_E < 0,$$

with $\mu_t = \frac{\mu_0(1-\pi)^t}{\mu_0(1-\pi)^t + (1-\mu_0)\pi^t}$. With each additional adverse outcome, the posterior that the state is ω_1 decreases. If this posterior decreases, then the probability to obtain a positive outcome with the next experiment also decreases, which makes it less attractive to experiment further. The sender stops experimenting at the lowest t for which the inequality holds. Denote the number of adverse outcomes until the sender stops experimenting unsuccessfully by T_E . The sender never stops experimenting unsuccessfully if inequality (1) is violated at the worst posterior from the sender's perspective, that is, if the posterior is zero. The sender does not experiment at all ($T_E = 0$) if running a single experiment at the prior belief is worse than not

experimenting, that is, for parameters for which $EU_0^1 < 0$. T_E is strictly positive and finite for all other parameters. In this case information about the state of the world is generated.¹³

We now study how the sender's behavior depends on π . The sender's continuation utility from running one more experiment and then stopping after either outcome is described in (1). Lemma 1 finds that T_E decreases in π (conditional on $T_E \geq 1$) as an increase of π lowers EU_t^1 . In (1), the positive effect of a higher π in the unlikely state ω_1 is overcompensated by the adverse reduction of the posterior μ_t and the negative effect in the likely state ω_2 .¹⁴

LEMMA 1. *Consider parameters such that $T_E \geq 1$. (i) T_E weakly decreases in π . (ii) $T_E = 1$ if π is sufficiently high.*

The reduction of excessive experimentation that is implied by a higher π suggests that the π_E that maximizes the decision quality in the one-shot experimentation model remains optimal if more experiments are possible. We confirm this intuition below.

5.1.2. The decision quality. We now study the face values that maximize the decision quality $\text{prob}\{a_i = \omega_i\}$. Conditional on experimentation being sequentially rational, Lemma 2 finds that an increase of π has a positive impact on the decision quality. This is due to the positive direct effect of a higher π and the positive indirect effect via a reduction of excessive private experimentation according to Lemma 1.

LEMMA 2. *Consider an increase of the face value π and a sender that runs at least one experiment before and after the increase of π . The decision quality $\text{prob}\{a_i = \omega_i\}$ weakly increases in π .*

Lemmas 1 and 2 imply that $\text{prob}\{a_i = \omega_i\}$ is greater for each high π for which $T_E = 1$ is sequentially rational than for any low π for which $T_E > 1$ is sequentially rational. For the face values for which $T_E = 1$ is sequentially rational, we have that $\text{prob}\{a_i = \omega_i\}$ also increases in π . Running a single experiment is better for the sender than not experimenting if

$$(2) \quad \mu_0\pi + (1 - \mu_0)(1 - \pi)\theta - c_E \geq 0,$$

which corresponds to the participation constraint under one-shot experimentation.

We now argue that the sender does not want to continue experimenting upon observing an adverse outcome of the first experiment with precision π_E under one-shot experimentation. It then follows that the face values that maximize the decision quality under sequential experimentation are the same as under one-shot experimentation.

If $\pi_E = 1$ under one-shot experimentation, then the sender here does not run more than one experiment, as an adverse outcome confirms the adverse state with certainty and it is impossible to obtain a positive outcome by chance. If $\pi_E < 1$ under one-shot experimentation, then the participation constraint, which is analogous to inequality (2) here, is binding. The sender prefers to stop experimenting after observing an adverse outcome of an experiment with $\pi_E < 1$, as he is indifferent to run an experiment at the more favorable prior belief. Proposition 1 directly follows.

¹³ Wald (1947) discusses a statistician's ("sender's") decision function, which implies when to stop experimenting. Here, the sender's decision function is in addition endogenously influenced by the incentives provided by the receiver's decision rule (only to publish a positive outcome) and the information structure (adverse outcomes can be hidden).

¹⁴ Note that the prior μ_0 does not depend on π , whereas posterior μ_t deteriorates if π increases. As there is no detrimental reduction of μ_0 in response to an increase of π , we may have that EU_0^1 increases in π . EU_t^1 instead decreases in π if $t > 0$, which yields Lemma 1.

PROPOSITION 1. *The face value π_E that maximizes the decision quality $\text{prob}\{a_i = \omega_i\}$ is (i) $\pi_E = 1$ if $\mu_0 \geq c_E$, (ii) $\pi_E = \frac{c_E - (1 - \mu_0)\theta}{\mu_0 - (1 - \mu_0)\theta} < 1$ if $\mu_0 < c_E$ and $\frac{c_E - (1 - \mu_0)\theta}{\mu_0 - (1 - \mu_0)\theta} \geq 1 - \mu_0$, and (iii) $\pi_E \in (\frac{c_E - (1 - \mu_0)\theta}{\mu_0 - (1 - \mu_0)\theta}, 1]$ if $\mu_0 < c_E$ and $\frac{c_E - (1 - \mu_0)\theta}{\mu_0 - (1 - \mu_0)\theta} < 1 - \mu_0$. In cases (i) and (ii) a single experiment is run and the decision quality is $\text{prob}\{a_i = \omega_i\} = \pi_E$. In case (iii) there is no experimentation and the decision quality is $\text{prob}\{a_i = \omega_i\} = 1 - \mu_0$.*

Under the conditions of Proposition 1 (ii) and (iii), an opportunistic sender type may be deterred from running experiments if the face value is too high. It then becomes too hard to find a positive outcome with a high quality experiment given the unfavorable prior belief to justify experimentation costs.

5.2. Manipulable Methods.

5.2.1. *Sender behavior.* If manipulation is possible, then the sender may run experiments and eventually manipulate. Suppose the sender observes a history h_t containing exclusively adverse outcomes. The continuation utility from manipulation at history h_t is

$$(3) \quad EU_t^M \equiv \mu_t + (1 - \mu_t)\theta - c_M,$$

which decreases in the number of adverse outcomes that history h_t contains.

The sender does not manipulate if for all $t \leq T_E$ the continuation utility from manipulation EU_t^M is below the continuation utility EU_t^C from experimenting without eventual manipulation. Otherwise, eventual manipulation occurs if the sender does not find a positive outcome by experimentation. Denote by T_M the number of adverse outcomes after which such a sender type manipulates.¹⁵ If $T_M = 0$, then this type manipulates at the outset. Manipulation occurs if c_M is sufficiently low.

Let us again study how the sender's behavior depends on π .

PROPOSITION 2. *In response to an increase of π , there exist parameters such that a type switches from*

- (i) *eventual manipulation to experimenting with eventual unsuccessful stopping;*
- (ii) *experimenting with eventual unsuccessful stopping to eventual manipulation.*

The sender's incentive to manipulate at a given μ_t does not depend on π . The incentive to run an additional experiment and then to stop after either outcome without manipulation at a given μ_t may increase or decrease in π . In addition, an increase of π in general changes the posterior where the sender stops experimenting via an adjustment of the experimentation behavior and the change of the informative value of each outcome. Suppose, for example, that the number of adverse outcomes is reduced from some t' to some t'' . We may have $\mu_{t'} > \mu_{t''}$ (as each adverse outcome is more informative) or $\mu_{t'} < \mu_{t''}$ (as the number of adverse outcomes is lower with t''). Both, the benefit from experimenting without manipulation and the benefit from manipulation depend on μ_t , and parameters determine which option the sender prefers.

5.2.2. *The decision quality.* The following observation confirms the intuition that manipulation is detrimental for the decision quality:

OBSERVATION 2. *Experimentation with a strictly positive and finite T_E yields a higher $\text{prob}\{a_i = \omega_i\}$ than manipulation.*

¹⁵ In the Appendix we describe an algorithm for the derivation of EU_t^C and T_M .

We now compare the face value π_E that maximizes the decision quality for nonmanipulable outcomes with the face value π_M that maximizes the decision quality if manipulation is possible. Let manipulation costs c_M be sufficiently low such that manipulation at the outset is sequentially rational for the sender at π_E .¹⁶ The argument for why $\pi_M \leq \pi_E$ now is analogous to one-shot experimentation.

For parameters described in Proposition 1 (i), the sender may be opportunistic or sincere and we have $\pi_E = 1$. An increase of the face value is not possible and, thus, $\pi_M \leq \pi_E$. For parameters described in Proposition 1 (ii) and (iii), the sender is opportunistic and, consequently, his ex ante benefit from running a single experiment without manipulation decreases in the face value. Hence, informative experimentation may only become more attractive than uninformative manipulation if the face value is reduced and it thereby becomes more likely to find a surprising positive experimental outcome that goes against the adverse prior. It follows that $\pi_M \leq \pi_E$ for the parameters in Proposition 1 (ii) and (iii).¹⁷ The next proposition, our major result, follows.

PROPOSITION 3. *Consider any face value π_E that maximizes $\text{prob}\{a_i = \omega_i\}$ for methods where manipulation is not possible. Suppose that manipulation is possible and that manipulation is optimal for the sender at π_E . The sender may only be encouraged to experiment with eventual unsuccessful stopping, which increases $\text{prob}\{a_i = \omega_i\}$, by lowering the face value below π_E .*

In the Appendix we show that this result does not depend on our assumption that $\theta \geq 0$ and also not on Assumption A.

A necessary condition for a strict increase of $\text{prob}\{a_i = \omega_i\}$ via a reduction of π below π_E (if manipulation occurs at π_E) is that the sender is opportunistic ($\theta \geq \frac{\mu_0}{1-\mu_0}$). Such a reduction can strictly increase $\text{prob}\{a_i = \omega_i\}$, for example, if $\pi_E < 1$ and if c_M is such that the sender just prefers manipulation at the outset: The sender's expected utility from running a single experiment with π_E without manipulation is zero. Manipulating at the outset by assumption (almost) yields the same expected utility. His continuation utility from manipulating after observing an adverse outcome of an experiment with π_E is strictly negative due to the worse posterior. A decrease of π ex ante renders running a single experiment without manipulation better than manipulation at the outset for an opportunistic sender. If the decrease of π is sufficiently small, then the benefit from manipulating after observing an adverse outcome of the experiment with the lower face value is still negative. Hence, such a decrease of π below π_E encourages experimentation without manipulation and strictly increases $\text{prob}\{a_i = \omega_i\}$. It is straightforward that $\pi_M < \pi_E$ does not only hold for knife edge cases.

Consider the parameters in Proposition 3 and in addition suppose that experimentation with unsuccessful stopping is sequentially rational for the sender for some face value. The face value that maximizes $\text{prob}\{a_i = \omega_i\}$ is the highest π such that $\mu_t \pi + (1 - \mu_t)(1 - \pi)\theta - c_E \geq \mu_t + (1 - \mu_t)\theta - c_M$ for all $t \leq T_E$, where μ_t and T_E are functions of π .

5.3. Welfare. Suppose the receiver obtains utility 1 if her decision matches the state and 0 otherwise. Suppose that the sender is small in the sense that his utility has a negligible impact on welfare. Welfare is then equal to the receiver's expected utility, which is equal to the

¹⁶ The case where manipulation occurs after observing an adverse outcome is discussed in the proof of Proposition 3. Consider a π_E , where the sender runs a single experiment under a nonmanipulable method. Eventual manipulation occurs if $c_M < \max\{\mu_0(1 - \pi_E) + (1 - \mu_0)\theta\pi_E + c_E, \mu_1 + (1 - \mu_1)\theta\}$, where μ_1 is based on π_E . Consider a π_E , where the sender does not experiment under a nonmanipulable method. Manipulation occurs if $c_M < \mu_0 + (1 - \mu_0)\theta$. If instead c_M is sufficiently high such that eventual manipulation does not occur at π_E , then $\pi_M = \pi_E$.

¹⁷ Note that for parameters in Proposition 1 (iii) no experimentation yields a higher $\text{prob}\{a_i = \omega_i\}$ than experimentation with $\pi \leq \frac{c_E - (1 - \mu_0)\theta}{\mu_0 - (1 - \mu_0)\theta}$. However, as noted above, no experimentation cannot be induced with any face value for these parameters. It would yield a higher decision quality if the decision maker chose against the sender for any face value.

decision quality $\text{prob}\{a_i = \omega_i\}$. The following proposition identifies circumstances where welfare decreases in π .

PROPOSITION 4. (i) *Suppose the sender switches from experimenting with eventual unsuccessful stopping to doing nothing in response to the increase of π . There is a threshold precision $\bar{\pi}$ such that welfare decreases if the initial precision is $\pi > \bar{\pi}$.*

(ii) *Suppose manipulation is possible and the sender switches from experimenting with eventual unsuccessful stopping to eventual manipulation in response to the increase of π . Welfare decreases in π .*

In part (i), the increase of π decreases $\text{prob}\{a_i = \omega_i\}$ if the initial face value is sufficiently high.¹⁸ If manipulation is possible (part (ii)), then there is the additional effect that an increase of π may induce a switch from experimentation without manipulation to manipulation, which lowers $\text{prob}\{a_i = \omega_i\}$.

If a social planner can make the decisions, then she can obtain the welfare maximum (at lowest costs) by running a single experiment with $\pi = 1$. Due to manipulation and strategic experimentation, the highest face value $\pi = 1$ is not necessarily socially optimal if the sender experiments.

6. APPLICATION—PHILOSOPHY OF SCIENCE ASPECTS

For many arguments in science, sequential private experimentation with selective information revelation is possible. For example, a researcher may privately run thought experiments for a logical argument or he may privately run different specifications on nonmanipulated data for an empirical argument. As argued in the introduction, manipulation is often possible and manipulation costs depend on the research method.

In the tradition of our discipline and based on empirical findings by replication studies, we think that incentives, such as publication-based rewards, matter for researcher behavior.¹⁹ For example, the Reproducibility Project (Open Science Collaboration 2015) tried to replicate the experimental results of 100 psychological studies published in three psychology journals. The replication effects were half the magnitude of the effects in the original studies. Whereas 97% of the original studies had significant results, only 36% of the replications had significant results. The Reproducibility Project suggests a substantial amount of excessive private experimentation with selective information revelation and/or manipulation in science.

6.1. Field Contingent Quality Requirements. Our article argues that the optimal face value depends on the scientific methods under consideration. In particular, it should depend on how easily manipulation is possible. For example, if we are interested in the decision quality for articles that go against the prior belief, then our analysis suggests that the optimal face value for empirical arguments that are based on privately collected data is below the optimal face value if the argument is based on publicly available data. Manipulating a logical deduction in science also tends to be rather costly. Logic per se cannot be manipulated, but mistakes may occur. An outsider cannot directly distinguish between an honest mistake and manipulation, that is a “mistake on purpose,” that yields the desired conclusion. However, when a logical argument is released, then the flawed analysis becomes public and there is a decent chance that

¹⁸ The intuition regarding the threshold $\bar{\pi}$ is as follows. If the initial precision is below $\bar{\pi}$ and the sender experiments, then it is likely that the wrong decision is made by chance. With the higher precision he does not experiment and induces $a = a_2$, where a prior $\mu_0 < 1/2$ suggests that this is the correct decision. If the initial precision is above $\bar{\pi}$ instead and the sender experiments, then it is likely that the correct decision is made. This effect can be stronger than the effect on the decision quality that inducing $a = a_2$ has in response to an increased precision.

¹⁹ Brodeur, et al. (2016) examine the distribution of 50,000 statistical tests published in the AER, JPE, and QJE and find evidence that researchers “inflate” the value of just-rejected tests by choosing significant specifications in order to get their results published.

the flaw is detected either by referees, an editor or the scientific community. This typically has negative reputation effects regarding ability, skills etc., which suggest high manipulation costs.

6.2. P-Value Adjustment. Benjamin et al. (2018) propose to reduce the default P -value threshold for statistical significance for claims of new discoveries from 0.05 to 0.005. In our terminology this resembles an increase of the face value.²⁰ They note that the rate of replication in recent replication projects in psychology and experimental economics is about twice as high for initial studies with $P < 0.005$ relative to initial studies with $0.005 < P < 0.05$.²¹ They also study the gains of reducing the P -value following a set of statistical assumptions.

Investigating reproducibility for different thresholds using the same set of articles in order to justify a higher face value is reasonable if the results of the Reproducibility Project are mainly due to a statistical anomaly, but it is problematic if the results are mainly due to scientific misconduct. Each face value triggers a behavioral response. Using the same set of articles to study the reproducibility for different thresholds may underestimate the behavioral response. According to our model, the impact of an increase of the face value on the quality of published articles is ambiguous. Depending on parameters, the sender may switch from eventual unsuccessful stopping to eventual manipulation or vice versa if the face value increases. The quality of published articles may therefore go up or down. Our model suggests that there are two risks from lowering the P -value threshold. First, the more sincere scientists may be deterred from experimenting. Second, some scientists may switch from experimenting to manipulation. With a threshold of 0.005, it becomes substantially harder to obtain surprising publishable results with honest scientific work and the incentive either not to experiment or, perhaps more importantly, to eventually manipulate may become stronger.

6.3. Manipulation and Experimentation Costs. FS study the effect of experimentation costs on the value of nonmanipulable methods. They use their findings to rationalize restrictions on admissible scientific methods. They interpret higher experimentation costs as more restrictions. For example, it tends to be harder (more costly) to rationalize an economic effect using the homo economicus assumption than allowing arbitrary utility functions. Lower costs encourage excessive private experimentation. More restrictions may, therefore, improve the value of scientific arguments.

In our model the relative size of experimentation and manipulation costs matters for the value of manipulable methods. Low experimentation costs still encourage excessive experimentation, which has a negative impact on the quality of publications. Suppose, however, that a researcher type using a regression eventually manipulates his data for given experimentation and manipulation costs. This type's behavior is uninformative. A decrease of experimentation costs, for example, due to faster computer hardware or less restrictions, tends to make experimentation more attractive compared to manipulation. It is easy to construct cases where such a type starts to generate information if experimentation costs are lowered. As argued previously, there are scientific fields, such as psychology, where methods are used where manipulation is easy. Adjusting experimentation costs, for example, via the level of restrictions of admissible methods, and manipulation costs could be used to improve the quality of publications.

²⁰ Our Bayesian model and the classical perspective are not identical. The face value in this article cannot be directly translated into P -value, which is the calculated probability that the null hypothesis is true. However, we abstract from these differences. In both cases manipulation and excessive private experimentation with selective information revelation lower the real value of the analysis. If the required threshold of the P -value is lowered, the precision requirement is increased, which has a similar effect as an increase of π in our model.

²¹ The authors note though that the behavior of the researchers may change in response to a new threshold and that their behavior has to be monitored. On the other hand, if it were easy to monitor the behavior, then this could have been done with the 0.05 threshold as well.

7. CONCLUSION

This article studies persuasion with endogenous information acquisition and selective information revelation. We introduce the option to manipulate into a framework where an outcome may also be obtained by sequential private experimentation. We investigate how changing the face value required for persuasion affects the sender's behavior and its effect on the decision quality.

We describe the sender's behavior and derive the face value requirements for persuasion that maximize the decision quality for nonmanipulable methods. Opportunistic sender types, who do not worry much about persuasion with a false positive, may be deterred from informative experimentation if the face value requirements are too challenging. If the method is manipulable instead, then an increase of the face value requirement may encourage a switch from experimentation to manipulation or vice versa. Our major result shows that the face value that maximizes the decision quality for manipulable methods is below the face value that maximizes this quality for nonmanipulable methods.

We apply our analysis to scientific arguments. We find that increasing the face value required for publication may have a detrimental effect on the quality of the publication process, in particular for scientific methods that rely on privately collected data. A higher face value may, for example, make manipulation more attractive compared to experimentation. We are sceptical regarding a P -value adjustment, as proposed in Benjamin et al. (2018), since it makes it harder to find surprising publishable outcome with honest scientific work.

We conclude with a word of caution. Our result that the face value for manipulable methods should be *unambiguously* weakly lower than for nonmanipulable methods is model specific. In our view, we develop a reasonable intuition for why the face value should be lower for manipulable methods. But even in our model we can identify parameters where an increase of the face value from a suboptimal level induces a switch from manipulation to experimentation, as in Proposition 2 (i). Future research could explore potential limitations of this result.

APPENDIX A

PROOF OF LEMMA 1. The expected utility from running one further experiment at T_E can be written as $EU_{T_E}^1(\mu_{T_E-1}, \pi) = \frac{\mu_{T_E-1}(1-\pi)}{\mu_{T_E-1}(1-\pi)+(1-\mu_{T_E-1})\pi}\pi + \frac{(1-\mu_{T_E-1})\pi}{\mu_{T_E-1}(1-\pi)+(1-\mu_{T_E-1})\pi}(1-\pi)\theta - c_E$, with $\mu_{T_E} = \frac{\mu_{T_E-1}(1-\pi)}{\mu_{T_E-1}(1-\pi)+(1-\mu_{T_E-1})\pi}$. The proof proceeds in 3 steps. Step (1) shows that $EU_{T_E}^1$ decreases in π if μ_{T_E-1} is considered as a parameter that is not affected by the change of π . Step (2) argues that $EU_{T_E}^1$ increases in μ_{T_E-1} for a given π . Step (3) uses steps (1) and (2) and shows that $EU_{T_E}^1$ decreases if π increases and there is a weak decrease of μ_{T_E-1} .

(1) Suppose μ_{T_E-1} is a constant in $EU_{T_E}^1(\mu_{T_E-1}, \pi)$. Then,

$$\frac{dEU_{T_E}^1}{d\pi} = \frac{\theta + \mu_{T_E-1} - \theta\mu_{T_E-1}}{(\pi + \mu_{T_E-1} - 2\pi\mu_{T_E-1})^2}(\mu_{T_E-1}(1 - 2\pi) + \pi^2(2\mu_{T_E-1} - 1)) < 0, \quad \text{since } \pi > 1/2 \quad \text{and} \\ \mu_{T_E-1} < 1/2.$$

(2) Suppose π is a constant in $EU_{T_E}^1(\mu_{T_E-1}, \pi)$. Then, $\frac{dEU_{T_E}^1}{d\mu_{T_E-1}} > 0$.

(3) Consider a marginal increase of π from π' to π'' that implies that μ_{T_E-1} weakly decreases from μ'_{T_E-1} to μ''_{T_E-1} . According to step (2) we have $EU_{T_E}^1(\mu'_{T_E-1}, \pi') \geq EU_{T_E}^1(\mu''_{T_E-1}, \pi')$. According to step (1) we have $EU_{T_E}^1(\mu''_{T_E-1}, \pi') > EU_{T_E}^1(\mu''_{T_E-1}, \pi'')$. These two inequalities imply $EU_{T_E}^1(\mu'_{T_E-1}, \pi') > EU_{T_E}^1(\mu''_{T_E-1}, \pi'')$. \square

PROOF OF LEMMA 2. We have $prob\{a_i = \omega_i\} = \mu_0(1 - (1 - \pi)^{T_E}) + (1 - \mu_0)\pi^{T_E}$. The proof proceeds in 3 steps. Step (1) shows that $prob\{a_i = \omega_i\}$ increases in π if T_E is not affected by the

change of π . Step (2) shows that $\text{prob}\{a_i = \omega_i\}$ decreases in T_E for a given π if $T_E \geq 2$.²² Step (3) uses steps (1) and (2) and shows that the $\text{prob}\{a_i = \omega_i\}$ increases if π increases and there is a discrete decrease of T_E .

(1) Suppose that T_E does not change if π is marginally increased. In this case $\frac{d\text{prob}\{a_i = \omega_i\}}{d\pi} = T_E((1 - \mu_0)\pi^{T_E-1} + \mu_0(1 - \pi)^{T_E-1}) > 0$.

(2) Consider a discrete change of T_E without changing π . In $\text{prob}\{a_i = \omega_i\} = \mu_0(1 - (1 - \pi)^{T_E}) + (1 - \mu_0)\pi^{T_E}$, we have that $(1 - (1 - \pi)^{T_E})$ increases in T_E and π^{T_E} decreases in T_E . The decision quality $\text{prob}\{a_i = \omega_i\}$ is a convex combination of $(1 - (1 - \pi)^{T_E})$ and π^{T_E} , with μ_0 being the corresponding parameter. The critical case is the highest μ_0 , that is, $\mu_0 = 1/2$.²³ In this case the decision quality simplifies to $\text{prob}\{a_i = \omega_i\} = \frac{1}{2}(1 - (1 - \pi)^{T_E} + \pi^{T_E})$ for each T_E . We have

$$\begin{aligned} & \frac{1}{2}(1 - (1 - \pi)^{T_E} + \pi^{T_E}) \\ &= \frac{1}{2}(1 - (1 - \pi)^{T_E-1} + \pi(1 - \pi)(1 - \pi)^{T_E-2} + \pi\pi^{T_E-1}) \\ &\leq \frac{1}{2}(1 - (1 - \pi)^{T_E-1} + \pi(1 - \pi)\pi^{T_E-2} + \pi\pi^{T_E-1}) \\ &= \frac{1}{2}(1 - (1 - \pi)^{T_E-1} + \pi^{T_E-1}), \end{aligned}$$

where the inequality holds if $T_E \geq 2$, as $\pi > 1 - \pi$. Therefore, for a given π , we have that $\text{prob}\{a_i = \omega_i\}$ increases if T_E is decreased by 1, for all $\mu_0 \in [0, 1/2]$.

(3) Suppose that π increases from π' to π'' and that T_E decreases by 1 due to the change of π . In the critical case with $\mu_0 = 1/2$ we have $\frac{1}{2}(1 - (1 - \pi')^{T_E} + \pi'^{T_E}) < \frac{1}{2}(1 - (1 - \pi'')^{T_E} + \pi''^{T_E})$ as shown in step (1). Furthermore, we have $\frac{1}{2}(1 - (1 - \pi'')^{T_E} + \pi''^{T_E}) \leq \frac{1}{2}(1 - (1 - \pi'')^{T_E-1} + \pi''^{T_E-1})$ as shown in step (2). This implies $\frac{1}{2}(1 - (1 - \pi')^{T_E} + \pi'^{T_E}) < \frac{1}{2}(1 - (1 - \pi'')^{T_E-1} + \pi''^{T_E-1})$.

The same logic holds if T_E is decreased by more than 1. The proof if the sender never stops experimenting unsuccessfully before the increase of π is straightforward. \square

Algorithm for the derivation of EU_i^C and T_M

Define the continuation utility if the sender follows his optimal experimentation plan and if the state is ω_1 as $EU_i^{\omega_1} = \pi + (1 - \pi)EU_{i+1}^{\omega_1} - c_E$. Analogously define $EU_i^{\omega_2} = (1 - \pi)\theta + \pi EU_{i+1}^{\omega_2} - c_E$ in state ω_2 . We have $EU_i^{\omega_i} = 0$ for all $t \geq T_E$, as at all such t the sender does not run further experiments. $EU_i^{\omega_i}$ can be easily determined backwards starting at T_E (i.e., first determine it at T_E , then at $T_E - 1$, then at $T_E - 2$...). The continuation utility in t is then $EU_i^C = \mu_t EU_i^{\omega_1} + (1 - \mu_t) EU_i^{\omega_2}$.

For a type that manipulates, T_M can be determined as follows. We have $T_M \leq T_E$. First, form the hypothesis that $T_M = T$, with $T = T_E$. Then, calculate the continuation utilities for all $t \leq T$ from eventual manipulation given T . If “one stage deviations” at μ_t , that is, either manipulate at t or stop unsuccessfully at μ_t , are not profitable for all $t \leq T$, then the hypothesis is true. Otherwise, continue with the hypothesis that $T_M = T - 1$, then $T_M = T - 2$ and so on until a hypothesis is true.

PROOF OF PROPOSITION 2. (i) Table A1 shows that an increase of π may induce a switch from eventual manipulation to experimentation with eventual unsuccessful stopping.

The sender does not want to run a second experiment after observing an adverse outcome of the first experiment before and after the change of π . However, manipulating is just better than stopping unsuccessfully after observing an adverse outcome with the initial π . If π increases, then manipulating after observing an adverse outcome becomes worse than stopping unsuccessfully at this history, as the posterior is now worse due to the higher informational

²² Note that the case is considered, where T_E is finite and strictly greater than 0 before and after the change of π . Hence, if we have $T_E = 1$ before the change of π , then we have $T_E = 1$ after the change of π . If $T_E = 1$, then $\text{prob}\{a_i = \omega_i\} = \mu_0(1 - (1 - \pi)) + (1 - \mu_0)\pi = \pi$, which increases in π .

²³ As $(1 - (1 - \pi)^{T_E})$ increases in T_E , and as π^{T_E} decreases in T_E , the decision quality $\text{prob}\{a_i = \omega_i\}$ decreases in T_E for all μ_0 , if it decreases when the maximum weight $\mu_0 = 1/2$ is attached to $(1 - (1 - \pi)^{T_E})$.

TABLE A1
SWITCH FROM EXPERIMENTATION WITH MANIPULATION (PART A) TO EXPERIMENTATION WITHOUT MANIPULATION (PART B) FOR
PARAMETERS $\mu_0 = 0.3, c_E = 0.24, \theta = 0.5,$ AND $c_M = 0.5775$ IN RESPONSE TO AN INCREASE OF π

(A) $\pi_0 = 0.7$:	t	0	1	(B) $\pi_1 = 0.71$:	t	0	1
	μ_t	0.3	0.155		μ_t	0.3	0.149
	EU_t^1	0.075	-0.005		EU_t^1	0.0745	-0.011
	EU_t^C	0.075	0		EU_t^C	0.0745	0
	EU_t^M	0.0725	0.00009		EU_t^M	0.0725	-0.003

TABLE A2
OPTIMAL EXPERIMENTATION WITH $T_E = 7$ FOR PARAMETERS $\pi = 0.55, \mu_0 = 0.4, c_E = 0.35, \theta = 0.7,$ AND $c_M = 0.75$

t	0	1	2	3	4	5	6	7
μ_t	0.40	0.35	0.31	0.27	0.23	0.20	0.17	0.14
EU_t^1	0.059	0.048	0.038	0.028	0.019	0.011	0.004	-0.002
EU_t^C	0.099	0.078	0.059	0.042	0.026	0.013	0.004	0
EU_t^M	0.07	0.056	0.043	0.030	0.019	0.009	-0.000002	-0.008

TABLE A3
OPTIMAL MANIPULATION FOR PARAMETERS $\pi = 0.6, \mu_0 = 0.4, c_E = 0.35, \theta = 0.7,$ AND $c_M = 0.75$

t	0	1	2	3
μ_t	0.40	0.31	0.23	0.16
EU_t^1	0.058	0.028	0.003	-0.017
EU_t^C	0.074	0.030	0.003	0
EU_t^M	0.07	0.042	0.019	-0.001

value of the adverse outcome. At the prior, experimentation is sufficiently cheap such that (before and after the change of π) the sender prefers to gamble by running an experiment with a chance to obtain a positive outcome over obtaining a positive outcome with certainty by expensive manipulation.

(ii) Table A2 shows that manipulation is not optimal for some parameters.

Table A3 shows that eventual manipulation is optimal where all parameters except π , which is increased, are as in Table A2. \square

PROOF OF OBSERVATION 2. The decision quality from eventual manipulation is $prob\{a_i = \omega_i\} = \mu_0$, as $prob\{m = s_1 \mid \omega = \omega_1\} = 1$ and $prob\{m \neq s_1 \mid \omega = \omega_2\} = 0$.

The decision quality from experimentation with eventual unsuccessful stopping is $prob\{a_i = \omega_i\} = \mu_0(1 - (1 - \pi)^{T_E}) + (1 - \mu_0)\pi^{T_E}$.

We have $\mu_0(1 - (1 - \pi)^{T_E}) + (1 - \mu_0)\pi^{T_E} > \mu_0(1 - \pi^{T_E}) + (1 - \mu_0)\pi^{T_E} = \mu_0 + \pi^{T_E}(1 - 2\mu_0)$, where the inequality holds as $\pi > 1 - \pi$.

Finally, we have $\mu_0 \leq \mu_0 + \pi^{T_E}(1 - 2\mu_0)$, where $(1 - 2\mu_0) \geq 0$ due to $\mu_0 \leq 1/2$, which confirms that experimentation with unsuccessful stopping yields a higher decision quality than manipulation. \square

PROOF OF PROPOSITION 3. Suppose c_M is sufficiently low such that after observing an adverse outcome of an experiment with π_E the sender prefers manipulation to stopping unsuccessfully. The ex ante benefit from manipulation at the outset is then strictly greater than zero, as the prior belief is more favorable. This ex ante benefit does not depend on π . Consider an increase of π above π_E . An increase of π is only possible if $\pi_E < 1$. $\pi_E < 1$ implies that $\mu_0\pi + (1 - \mu_0)(1 - \pi)\theta - c_E$ decreases in π and we have $\mu_0\pi + (1 - \mu_0)(1 - \pi)\theta - c_E \leq 0$ if $\pi > \pi_E$. If $\pi > \pi_E$, it follows that ex ante manipulation at the outset is strictly better for the

sender than doing nothing, which is weakly better than running an experiment without manipulation. Regardless of whether manipulation at the outset or manipulation after observing an adverse outcome is ex ante best for the sender at a π exceeding π_E , the decision quality in both cases is the same. Hence, we have $\pi_M \leq \pi_E$.

This result does not depend on our assumption that $\theta \geq 0$. If $\theta < 0$, then $\mu_0\pi + (1 - \mu_0)(1 - \pi)\theta - c_E$ increases in π . Analogous to above, we then have $\pi_E = 1$ if $\mu_0 \geq c_E$ (and no experimentation occurs at any π if $\mu_0 < c_E$). With manipulation, an increase of the face value above $\pi_E = 1$ is not possible and we have $\pi_M \leq \pi_E$. Also note that for parameters excluded by Assumption A, for which the sender prefers no experimentation to experimentation for all $\pi \in (1/2, 1]$, we cannot have that a change of the face value induces a switch away from manipulation. In these cases neither the benefit from no experimentation nor from manipulation depend on π . \square

PROOF OF PROPOSITION 4. (i) Before the increase of π the sender experiments and after the increase he does not experiment. We now show that $\text{prob}\{a_i = \omega_i\}$ decreases in π for these parameters.

Consider an increase of π and a sender that switches from a finite and strictly positive T_E to $T_E = 0$. Consider the initial π . In this case $\text{prob}\{a_i = \omega_i\} = \mu_0(1 - (1 - \pi)^{T_E}) + (1 - \mu_0)\pi^{T_E}$.

Suppose π increases. Now the sender does not experiment and induces $a = a_2$. Hence, we have $\text{prob}\{a_i = \omega_i\} = \mu_0 * 0 + (1 - \mu_0) * 1 = (1 - \mu_0)$.

We have $\pi \in (1/2, 1)$. Setting $\pi = 1/2$ yields $(1 - \mu_0) > \mu_0(1 - (1 - \pi)^{T_E}) + (1 - \mu_0)\pi^{T_E}$ for any finite T_E , with $T_E > 0$. By continuity the inequality also holds for π greater but sufficiently close to $1/2$. Setting $\pi = 1$ yields $(1 - \mu_0) < \mu_0(1 - (1 - \pi)^{T_E}) + (1 - \mu_0)\pi^{T_E}$. By continuity the inequality also holds for π smaller but sufficiently close to 1. We have that $(1 - \mu_0)$ is independent of π and, as shown in the proof of Lemma 2, we have that $\mu_0(1 - (1 - \pi)^{T_E}) + (1 - \mu_0)\pi^{T_E}$ increases in π , which completes the proof.

(ii) After the increase of π the sender manipulates. As established in Observation 2, $\text{prob}\{a_i = \omega_i\}$ decreases in π for these parameters. \square

REFERENCES

- AGHION, P., P. BOLTON, C. HARRIS, AND B. JULLIEN, "Optimal Learning by Experimentation," *The Review of Economic Studies* 58 (1991), 621–54.
- , M. DEWATRIPONT, AND J. C. STEIN, "Academic Freedom, Private-Sector Focus, and the Process of Innovation," *RAND Journal of Economics* 39 (2008), 617–35.
- BALIGA, S., AND J. C. ELY, "Torture and the Commitment Problem," *Review of Economic Studies* 83 (2016), 1406–39.
- BENJAMIN, D. J., J. O. BERGER, M. JOHANNESSEN, B. A. NOSEK, E. J. WAGENMAKERS, R. BERK, K. A. BOLLEN, B. BREMBS, L. BROWN, C. CAMERER, D. CESARINI, C. D. CHAMBERS, M. CLYDE, T. D. COOK, P. DE BOECK, Z. DIENES, A. DREBER, K. EASWARAN, C. EFFERSON, E. FEHR, F. FIDLER, A. P. FIELD, M. FORSTER, E. I. GEORGE, R. GONZALEZ, S. GOODMAN, E. GREEN, D. P. GREEN, A. GREENWALD, J. D. HADFIELD, L. V. HEDGES, L. HELD, T. H. HO, H. HOIJTINK, D. J. HRUSCHKA, K. IMAI, G. IMBENS, J. P. A. IOANNIDIS, M. JEON, J. H. JONES, M. KIRCHLER, D. LAIBSON, J. LIST, R. LITTLE, A. LUPIA, E. MACHERY, S. E. MAXWELL, M. MCCARTHY, D. MOORE, S. L. MORGAN, M. MUNAFÓ, S. NAKAGAWA, B. NYHAN, T. H. PARKER, L. PERICCHI, M. PERUGINI, J. ROUDER, J. ROUSSEAU, V. SAVALEI, F. D. SCHÖNBRODT, T. SELLEKE, B. SINCLAIR, D. TINGLEY, T. VAN ZANDT, S. VAZIRE, D. J. WAITS, C. WINSHIP, R. L. WOLPERT, Y. XIE, C. YOUNG, J. ZINMAN, AND V. E. JOHNSON, "Redefine Statistical Significance," *Nature Human Behaviour* 2 (2018), 6–10.
- BERGEMANN, D., AND J. VÄLIMÄKI, "Bandit Problems," in S. N. Durlauf and L. E. Blume, eds., *The New Palgrave Dictionary of Economics*, 2nd ed. (Basingstoke: Palgrave Macmillan, 2008).
- BOLTON, P., AND C. HARRIS, "Strategic Experimentation," *Econometrica* 67 (1999), 349–74.
- BROCAS, I., AND J. D. CARILLO, "Influence through Ignorance," *The RAND Journal of Economics* 38 (2008), 931–47.
- BRODEUR, A., M. LÉ, M. SANGNIER, AND Y. ZYLBERBERG, "Star Wars: The Empirics Strike Back," *American Economic Journal: Applied Economics* 8 (2016), 1–32.
- CELIK, G., "Interested Experts: Do They Know More?," Mimeo, University of British Columbia, 2003.

- CRAWFORD, V. P., AND J. SOBEL, "Strategic Information Transmission," *Econometrica* 50 (1982), 1431–51.
- DZUIDA, W., "Strategic Argumentation," *Journal of Economic Theory* 146 (2011), 1362–97.
- FELGENHAUER, M., "Endogenous Persuasion with Costly Verification," *Scandinavian Journal of Economics* 121 (2019), 1054–87.
- , AND P. LOERKE, "Bayesian Persuasion with Private Experimentation," *International Economic Review* 58 (2017), 829–56.
- , AND E. SCHULTE, "Strategic Private Experimentation," *American Economic Journal: Microeconomics* 6 (2014), 74–105.
- , AND F. XU, "State of the Debate Contingent Arguments," *Economics Letters* 179 (2019), 46–8.
- GLAZER, J., AND A. RUBINSTEIN, "Debates and Decisions: On a Rationale of Argumentation Rules," *Games and Economic Behavior* 36 (2001), 158–73.
- HENRY, E., "Strategic Disclosure of Research Results: The Cost of Proving Your Honesty," *The Economic Journal* 119 (2009), 1036–64.
- , AND M. OTTAVIANI, "Research and the Approval Process: The Organization of Persuasion," *American Economic Review* 109 (2019), 911–55.
- HERRESTHAL, C., "Hidden Testing and Selective Disclosure of Evidence," Mimeo, University of Cambridge, 2017.
- JOVANOVIC, B., "Truthful Disclosure of Information," *Bell Journal of Economics* 13 (1982), 36–44.
- KAMENICA, E., AND M. GENTZKOW, "Bayesian Persuasion," *American Economic Review* 101 (2011), 2590–2615.
- KELLER, G., S. RADY, AND M. CRIPPS, "Strategic Experimentation with Exponential Bandits," *Econometrica* 73 (2005), 39–68.
- KUHN, T. S., *The Structure of Scientific Revolutions*, (Chicago, IL: University of Chicago Press, 1970).
- LEWIS, T. R., AND M. OTTAVIANI, "Search Agency," Mimeo, Università Bocconi, 2008.
- MILGROM, P., AND J. ROBERTS, "Relying on the Information of Interested Parties," *The RAND Journal of Economics* 17 (1986), 18–32.
- MOSCARINI, G., AND L. SMITH, "The Optimal Level of Experimentation," *Econometrica* 69 (2001), 1629–44.
- OLSZEWSKI, W., AND A. SANDRONI, "Falsifiability," *American Economic Review* 101 (2011), 788–818.
- OPEN, SCIENCE COLLABORATION, "Estimating the Reproducibility of Psychological Science," *Science*, 349 (2015), aac4716.
- POPPER, K. R., *The Logic of Scientific Discovery*, (New York: Basic Books, 1959).
- ROSENBERG, D., E. SOLAN, AND N. VIELLE, "Social Learning in One-Arm Bandit Problems," *Econometrica* 75 (2007), 1591–1611.
- ROTHSCHILD, M., "A Two-Armed Bandit Theory of Market Pricing," *Journal of Economic Theory* 9 (1974), 185–202.
- STERN, S., "Do Scientists Pay to Be Scientists," *Management Science* 50 (2004), 835–853.
- STRULOVICI, B., "Mediated Truth," Mimeo, Northwestern University, 2017.
- WALD, A., "Foundations of a General Theory of Sequential Decision Functions," *Econometrica* 15 (1947), 279–313.