

Strategic Private Experimentation*

Mike FELGENHAUER[†] and Elisabeth SCHULTE[‡]

November 6, 2014

Abstract

We consider a model of persuasion in which an agent who tries to persuade a decision maker can sequentially acquire imperfect signals. The agent's information acquisition is unobservable and he has the option to hide unfavorable signals. Nevertheless, if the signal precision is sufficiently high, he can persuade the decision maker by revealing a sufficiently large number of favorable signals. When the number of signals that can be transmitted to the decision maker is limited, persuasion is impossible if the agent's stakes are too high.

Keywords: Experimentation, information provision, information acquisition.

JEL classification: D82, D83

*We thank Thomas Gall, Sanjeev Goyal, Hans Peter Grüner, Martin Hellwig, Johannes Hörner, Navin Kartik, George Mailath, Paul Milgrom, Timofy Mylovanov, Andras Niedermayer, Georg Noeldecke, Wojciech Olszewski, Marco Ottaviani, Nicola Persico, Andrew Postlewaite, Mike Rior-dan, Klaus Schmidt, Heiner Schumacher, Joel van der Weele, Ernst-Ludwig von Thadden, Michael Ziegelmeyer and the audiences at the University of Pennsylvania, MPI (Bonn), Silaplana, ESWC 2010 (Shanghai), EEA 2010 (Glasgow), and the anonymous referees for their valuable comments and suggestions. Elisabeth Schulte gratefully acknowledges support from the Collaborative Research Center 884. Part of this research was done while Mike Felgenhauer was visiting the University of Pennsylvania. Their hospitality is gratefully acknowledged.

[†]University of Plymouth, School of Management, Cookworthy Building, Drake Circus, Plymouth, Devon, PL4 8AA, UK; email: mike.felgenhauer@plymouth.ac.uk.

[‡]Philipps-University of Marburg, Marburg Centre for Institutional Economics, School of Economics & Business Administration, D-35037 Marburg, Germany; email: elisabeth.schulte@wiwi.uni-marburg.de; Tel: 0049 (0)6421 2823080.

1 Introduction

This paper studies a situation in which an agent tries to persuade a decision maker to choose his preferred action by means of argumentation. We view arguments as verifiable signals that are imperfectly informative about a decision-relevant state of the world. We explore the effect of the agent’s strategic behavior on the value of his arguments for the decision maker, and we illustrate how the decision maker can extract valuable information from the agent.

To highlight the key properties of argumentation, consider the following examples.

- (A) A lobby of car manufacturers wants to prevent stricter emission standards. The policy maker wants to impose stricter regulation if it is technologically and economically feasible to meet the standard, but refrain from stricter regulation otherwise. The lobby argues: “There is an upward trend in average size and weight of car drivers. Moving a larger mass and simultaneously reducing emissions is not feasible.”¹
- (B) A girl wants to go out to a party. Her parents want to allow her to go only if the party is at a safe place. The girl argues: “Ann and Susan are allowed to go.”

These arguments are not cheap talk, as their content is verifiable. They suggest that the respective state is the one in which sender and receiver agree about the appropriate course of action. However, the evidence that is provided is imperfect. In Example A, the weight of the driver makes up only a part of the weight that the car has to move. In Example B, Ann’s and Susan’s parents’ decisions only reveal that they think the party is at a reasonably safe place, but not how sure they are about this assessment.

The acquisition of arguments, in particular the process of thinking, occurs in private. Hence, unfavorable arguments can be concealed. The policy maker in Example A may reasonably assume that the lobby has conducted additional research with potential relevance for predicting emissions, and the girl’s parents in Example B may reasonably assume that her daughter has asked several other friends if they are allowed to go to the party. If the decision maker believes that the agent may have intensively searched for arguments, she should be skeptical if the results of such efforts are not presented. As the decision maker cannot observe the agent’s evidence acquisition efforts, the evidence

¹An argument of this kind has been brought forward by the ACEA in its answer to the European Commission’s consultation on *CO*₂ emission standards in 2007, which is available at the commission’s web page: http://ec.europa.eu/clima/consultations/0001/organisations/acea_en.pdf.

that the agent brings forward may not be persuasive though taken at face value it could be.

The process of thinking typically does not immediately stop if the agent does not find a favorable argument at his first attempt. Rather, the incentive to continue or to stop the search for arguments depends on the evidence already collected. The value of the agent's arguments for the decision maker is affected by (i) repeated private attempts to find favorable arguments, (ii) selective revelation of the results, and (iii) a history-dependent decision to stop the search for arguments.

In our model, the agent obtains arguments by running experiments. Examples for experiments in the sense in which the word is used in this paper are the development of a theoretical model, a regression analysis, drawing a random sample, or the agent's exploration of his knowledge base as in Aragonés et al. (2005). An argument can be deductive, i.e., a logical inference from a set of assumptions, or inductive, i.e., a reason supporting the probable truth. It is in the nature of arguments that they are imperfectly informative. A deductive argument is valid only within a set of restrictive assumptions that are an imperfect description of reality. The imperfection of inductive arguments is evident. We therefore assume that an experiment is subject to both types of errors, false positives and false negatives. An experiment may yield a favorable argument in a state in which it is not in the decision maker's best interest to choose the agent's favorite action.² It may as well yield an unfavorable argument in a state in which the decision maker and the agent in fact agree.³ By running experiments, the agent learns about the state of the world.

The question that this paper explores is whether it is possible to persuade a Bayesian decision maker given that the agent acquires and reveals his arguments strategically. We show that the agent can persuade her by bringing forward a sufficiently large number of arguments. If a large number of arguments is needed to persuade the decision maker, this deters experimentation in the state in which the decision maker and the agent disagree, where it is less likely that an experiment yields a favorable argument. As the agent learns from previous failed experiments, he stops experimentation unsuccessfully when his interim belief becomes too pessimistic. Therewith, experimentation obtains an informational value: The fact that the agent acquired the number of arguments needed to be persuasive is in fact persuasive.

The value of an argument depends on the stakes of the speaker. An interested party

²Ann's and Susan's parents may be mistaken in their judgement of the safety of the place.

³The total weight of a car may decline and still it may be infeasible to reach the emission standard.

with larger stakes or lower cost of experimentation has to provide more arguments in order to be persuasive. We apply our model to the interaction between a researcher and an editor, in which the former tries to persuade the latter to publish his work, and we extend our model in two directions. First, we study a symmetric version of the problem. A special feature of this application is that the decision maker may be persuaded to choose the agent’s favorite action (to publish his paper) if she is sufficiently convinced that the paper’s proposition is true, independently of the state of the world that the proposition indicates. In such a problem, it is unambiguously beneficial for the decision maker if the agent has strong incentives for experimentation. Second, in the context of our application, it is reasonable to assume that the number of arguments that can be transmitted to the editor is limited. In this case, if the researcher’s stakes are high or the cost of experimenting is low, persuasion is impossible. The existence of a bound on the number of transmittable arguments can rationalize restrictive scientific standards, and it may trigger a barrier to new, objectively superior, methods.

2 Related literature

Our model relates to the literature on strategic experimentation, where the basic underlying trade-off is between benefits from exploration and benefits from exploitation, as in Rothschild (1974), Aghion et al. (1991), Bolton and Harris (1999), Keller et al. (2005) and Rosenberg et al. (2007).⁴ Similar to these bandit problems the agent here dynamically updates his belief about an exogenous state. At the same time our model belongs to the class of persuasion games (e.g., Jovanovich, 1982, Milgrom and Roberts 1986, Glazer and Rubinstein 2001, 2004, 2006, Kamenica and Gentzkow, 2011) in which the production of evidence is endogenous (e.g., Milgrom 2008). Papers that share features of both strands of the literature and that are closely related to our work are Celik (2003), Brocas and Carillo (2007) and Henry (2009).

Brocas and Carillo (2007) illustrate how an interested party can exert influence by controlling a flow of public experiments.⁵ They also consider a case in which experimentation is private, the agent sends a report about his posterior, and the report is verifiable. In this setting, skeptical beliefs à la Milgrom and Roberts (1986) induce unraveling, and the same information is available to the decision maker as under public

⁴For a survey see Bergemann and Välimäki (2008).

⁵Gul and Pesendorfer (2012) study competing interested parties who provide a flow of public experiments. Felgenhauer and Loerke (2014) build on the present framework and compare sequential public and private experimentation, but with an endogenous precision of the experiments.

experimentation. Henry (2009) studies private experimentation in a framework in which the interested party *ex ante* chooses how much to invest in experimentation. Again an unraveling argument applies and the agent’s report is fully revealing as the decision maker in equilibrium deduces the optimal amount of experimentation. The assumption that the agent is not able to adjust his decision to continue experimenting in the experimentation phase is central for the result. In our paper the decision to continue experimenting is history-dependent. As information acquisition occurs in private, this is a natural assumption. In such a framework the decision maker can only anticipate the optimal experimentation plan but not the actual number of experiments conducted. The decision maker’s beliefs are not always degenerate such that skeptical beliefs are not always helpful and there is no unraveling. In a similar setting, but with a continuous action space for the decision maker and a Poisson evidence production technology, Celik (2003) shows that no productive fully revealing equilibrium exists if the agent is *ex ante* uninformed about the state. A feature of his evidence production technology is that the agent’s interim belief evolves continuously between two experimental successes, which occur randomly in time. Hence, it is impossible for the decision maker to perfectly deduce the agent’s posterior from the presented evidence. In our model, the decision maker’s action space is binary and evidence production is modeled as a series of Bernoulli trials. Persuasion is possible, but in general communication is not fully revealing. Only if the presentation of a single successful experiment suffices to persuade the decision maker and a single failure deters further experimentation can the decision maker deduce the agent’s posterior perfectly.

Glazer and Rubinstein (2001, 2004, 2006) analyze debates, assuming that the debaters are endowed with hard evidence and that they can only report a subset of their evidence. If a debater cannot respond “appropriately” to his opponent’s argument, then this suggests that his overall endowment of evidence is unfavorable. As a consequence, evidence of the same quality (but different endogenous “appropriateness”) may have a different value. In our paper instead, information acquisition is endogenous and we find that exactly the same hard evidence (and not only evidence of the same quality) may have a different value depending on who presents the evidence.

The agent’s reporting space in our model is similar to that in Dzuida (2011). In Dzuida’s paper, a pool of arguments is exogenously given to the agent, but the decision maker does not know how many arguments the agent has. Like in our paper, in order to persuade the decision maker, the agent has to provide a sufficiently high number of arguments in his favor. Dzuida shows that an opportunistic agent also reveals coun-

terarguments, but these are ignored by the decision maker when updating about the state. Two-sided argumentation is triggered by the incentive to pool with an honest type.⁶ In the context of our model, an equilibrium with two-sided argumentation is either Pareto-inferior or payoff equivalent to an equilibrium in which counterarguments are not revealed.

Our work also relates to the literature on informational lobbying with evidence that is to some extent hard (e.g., Austen-Smith, 1994). Related to our specific application, there is an economic literature on academic research, e.g., Lewis and Ottaviani (2008), Olszewski and Sandroni (2011), Aghion et al. (2008) and Stern (2004). To the literature on the philosophy of science, e.g., Popper (1959) and Kuhn (1970), we add a rationalization of argumentation within a restrictive framework and of a possible barrier to new methods.

3 Model

We model information transmission from an interested agent to a decision maker with state-dependent preferences. There are two ex ante equally likely states of the world, $s \in \{0, 1\}$. The decision maker has the choice between two actions, $x \in \{0, 1\}$. The decision maker's utility is

$u(x, s)$	$s = 1$	$s = 0$
$x = 1$	0	$-p_d$
$x = 0$	$-(1 - p_d)$	0

with $p_d > 1/2$. We call p_d the decision maker's "threshold of doubt".⁷ She maximizes her expected utility and prefers $x = 1$ if the probability that $s = 1$ passes her threshold of doubt, and $x = 0$ otherwise. The agent prefers $x = 1$ regardless of the state of the world. In case $x = 1$, the agent's gross utility is $U > 0$ and otherwise it is 0. The agent maximizes expected gross utility minus the expected cost of experimentation. The agent and the decision maker hold the same prior belief about s .

The agent has access to an experimentation technology which can generate signals y_i , $y_i \in \{0, 1\}$. He can conduct as many experiments as he wants. Conditional on the

⁶In a setting with uncertainty about the agent's preferences and about his information endowment, and with an exogenous limit on the number of arguments that can be transmitted, Le Quement (2012) shows that the agent can use counter-arguments to signal a large (exogenous) endowment with signals.

⁷A similar formulation of preferences is used in the literature on committee decision making (e.g., Feddersen and Pesendorfer, 1998).

state, experimental outcomes are drawn independently. We assume $\text{prob}\{y_i = s' | s = s'\} = p$, $p \in (1/2, 1)$. If $y_i = 1$, experiment i yields an argument (or evidence) in favor of the agent, whereas $y_i = 0$ is an argument against him. He can neither manipulate nor invent experimental outcomes. In that sense, y_i is “hard” information.⁸ However, the agent can conceal experimental outcomes and he cannot prove that he did not conduct a particular experiment. Each experiment costs c . Unless stated otherwise, we assume $c > 0$. Let y_t denote the outcome of the t^{th} experiment. Denote with $h_t = (y_1, \dots, y_t)$ the experimentation history after the first t experiments. The number of conducted experiments cannot be observed by the decision maker. Denote with $n^0(h_t)$ the number of experiments with outcome 0 and with $n^1(h_t)$ the number of experiments with outcome 1 after the t^{th} experiment. Note that $n^0(h_t) + n^1(h_t) = t$. Hence, (n^0, n^1) , with obvious notation, summarizes all relevant information in a history h_t .

After each experiment, the agent updates his assessment of the probability distributions regarding the state of the world and future experimental outcomes and decides whether to continue or to stop experimenting.

The experimentation phase is modeled as a time interval. If the agent conducts an experiment at any given point in time, then he may still carry out as many experiments as desired before the experimentation phase ends. Therewith, we exclude the possibility of inferring information from the length of the experimentation phase.⁹ After the experimentation phase the agent publishes a report. Finally, the decision maker observes the announcement and chooses x .

A strategy for the agent consists of an experimentation plan and an announcement plan. For each possible experimentation history h_t , the strategy specifies whether to continue or to stop experimenting and, in the case of stopping, what (if anything) to reveal to the decision maker. Any announcement available to the agent can be summarized by $\hat{n} = (\hat{n}^0, \hat{n}^1)$, where \hat{n}^0 and \hat{n}^1 are the numbers of the announced unfavorable and favorable results, respectively. The case that the agent does not make any announcement is captured by $\hat{n} = (0, 0)$. He cannot manipulate experimental outcomes nor invent arguments. Hence, if the agent stops experimenting after t experiments, then the announcement has to satisfy $\hat{n}^0 \leq n^0(h_t)$ and $\hat{n}^1 \leq n^1(h_t)$. A strategy for the decision maker is to choose $x \in \{0, 1\}$ for each possible \hat{n} . In equilibrium, players’ strategies are

⁸An argument, e.g., as a logical inference from a set of assumptions, cannot be manipulated.

⁹If the decision maker can deduce something from the time elapsed until he receives information, a longer period suggests many failed experiments (see Hopenhayn and Squintani, 2011, in a different context). Our model abstracts from these issues. In many situations experiments may differ with respect to the time they require until completed and deducing failure from the time elapsed is difficult.

sequentially rational and players update their beliefs according to Bayes' Law whenever possible.

4 Equilibrium analysis

We start our analysis with an introduction of our notion of equilibrium standards of evidence. Then, we derive the agent's optimal experimentation plan and subsequently turn to the persuasiveness of the agent's announcements. The proofs of our statements can be found in the appendix.

4.1 Decision rule

The decision maker chooses $x = 1$ if and only if the probability that this is the correct decision passes her threshold of doubt:

$$x^* = \begin{cases} 1, & \text{if } \text{prob}\{s = 1 | \hat{n}\} \geq p_d \\ 0, & \text{if } \text{prob}\{s = 1 | \hat{n}\} < p_d. \end{cases}$$

The above decision rule maximizes her expected utility, taking into account all the information available to her. She forms beliefs regarding the experimental outcomes hidden from her according to the agent's presumed strategy, applying Bayes' Law whenever possible. Upon observing an event which occurs with probability zero in equilibrium, arbitrary beliefs are allowed. Ex ante, we have $\text{prob}\{s = 1\} = 1/2 < p_d$. Without access to additional information, the decision maker chooses $x = 0$. If she is confronted with an announcement \hat{n} , the decision maker takes into account the agent's experimentation and announcement strategy when assessing the informational value of \hat{n} . An announcement \hat{n} has an informational value if the decision maker's posterior $\text{prob}\{s = 1 | \hat{n}\}$ is different from her prior $\frac{1}{2}$. The informational value of an announcement \hat{n} is the higher, the larger $|\text{prob}\{s = 1 | \hat{n}\} - \frac{1}{2}|$.

There is a class of equilibria in which the decision maker always chooses $x = 0$. She may, e.g., believe that for any argument in favor of the agent that he reveals, he hides an argument against him. Thus, $\text{prob}\{s = 1 | \hat{n}\} \leq 1/2$ for all \hat{n} . As the decision maker cannot be persuaded to choose $x = 1$, there is no point for the agent to collect (costly) evidence. Hence, all announcements $\hat{n} \neq (0, 0)$ are out-of-equilibrium-events and the decision maker's beliefs are consistent. Note that there are many other out-of-equilibrium-beliefs which support this equilibrium behavior.¹⁰

¹⁰This class of equilibria is similar to the "babbling"-equilibria in cheap talk games. There, the decision

We are interested in equilibria in which the decision maker can sometimes be persuaded to choose in favor of the agent.

Definition 1 *An equilibrium is a persuasion equilibrium if $x^* = 1$ for at least one announcement \hat{n} that is announced with positive probability.*

Due to the power that out-of-equilibrium-beliefs have in this game, multiple persuasion equilibria may exist.

Definition 2 *A persuasion equilibrium in which the decision maker uses the decision rule*

$$x^* = \begin{cases} 1, & \text{if } \hat{n}^0 = 0, \hat{n}^1 \geq n^* \\ 0, & \text{else.} \end{cases} \quad (1)$$

is an equilibrium with a standard of evidence n^ .*

In the following we refer to n^* as the “standard of evidence”, and we restrict attention to the class of equilibria with a standard of evidence n^* . We consider this class of equilibria as particularly relevant because the following proposition holds.

Proposition 1 *(i) Any persuasion equilibrium coexists with a payoff-equivalent or Pareto-dominant equilibrium with a standard of evidence n^* .*

(ii) If an equilibrium with a standard of evidence n^ exists, it Pareto-dominates any equilibrium in which the decision maker always chooses $x = 0$.*

Focussing on the class of equilibria with a standard of evidence n^* allows us to identify an equilibrium with a single number n^* . That way, we avoid case distinctions in the proofs and qualifications in our statements. Note that the restriction to the announcement $\hat{n}^0 = 0$ is harmless because it is available to the agent for any information endowment.

4.2 Optimal experimentation

Assume that the decision maker applies decision rule (1), and note that it is optimal for the agent to release all favorable arguments and to conceal all unfavorable arguments maker believes that the agent sends the same message for any possible information endowment (giving the sender no incentive to do otherwise). Observing a message off the equilibrium path, the decision maker thinks that the message is not informative. Here, off the equilibrium path the decision maker believes that the agent has searched too often and thus any hard evidence announced does not contain sufficient information.

that he has acquired. The agent's experimentation strategy maps all possible histories of experimentation outcomes h_t into either of the two actions "stop experimenting" or "continue experimenting". His optimal experimentation plan can be characterized by two sets of experimentation histories that occur with positive probability on the equilibrium path, a set $H_s(n^*)$ of histories that yield persuasive evidence, and a set $H_f(n^*)$ of histories after which experimentation is stopped unsuccessfully.¹¹ For each element h_t of $H_s(n^*)$ and $H_f(n^*)$ it is optimal to continue experimenting after each sub-history of h_t .¹² Conditional on having acquired a certain number of favorable and adverse outcomes, the sequence of experimental outcomes does not matter for the agent's continuation decision. The relevant information at h_t can be summarized by the number of unsuccessful and successful experiments (n^0, n^1) .

Denote with $v_{n^*}(n^0, n^1)$ the continuation value when the agent is endowed with a stock of experimental evidence (n^0, n^1) . The continuation value is the maximum of $v_{n^*}^s(n^0, n^1)$ and $v_{n^*}^c(n^0, n^1)$, which denote the agent's continuation value when stopping and when continuing experimentation, respectively.

If $n^1 \geq n^*$, the decision maker can be successfully persuaded. Further experimentation is costly and does not yield a benefit. It is optimal for the agent to stop experimenting at h_t if $n^1 = n^*$. $h_t \in H_s(n^*)$ if it is optimal for the agent to continue experimenting for each sub-history of h_t .

Consider (n^0, n^1) with $n^1 < n^*$. As the acquired evidence does not persuade the decision maker, the agent's continuation utility when stopping experimentation is zero at such an experimentation history. Continuing experimentation yields continuation utility

$$v_{n^*}^c(n^0, n^1) = \beta(n^0, n^1)v_{n^*}(n^0, n^1 + 1) + (1 - \beta(n^0, n^1))v_{n^*}(n^0 + 1, n^1) - c,$$

where $\beta(n^0, n^1) = \gamma(n^0, n^1)p + (1 - \gamma(n^0, n^1))(1 - p)$ denotes the probability that the next experiment yields a favorable outcome given the experimental evidence (n^0, n^1) , with $\gamma(n^0, n^1) = \text{prob}\{s = 1 | (n^0, n^1)\} = \frac{1}{1 + (\frac{1-p}{p})^{n^1 - n^0}}$. If $v_{n^*}^c(n^0, n^1) < 0$ at h_t , it is optimal for the agent to stop experimenting. $h_t \in H_f(n^*)$ if it is optimal for the agent to continue experimenting for each sub-history of h_t .

The agent's incentive to continue experimenting depends only on the number of arguments still to be acquired, $n^* - n^1$, and his assessment of the probability distribution

¹¹If it is optimal for the agent not to start experimenting, i.e., $H_s(n^*) = \emptyset$, $H_f(n^*)$ contains the "history" at which the experimentation phase starts, i.e., $H_f(n^*) \neq \emptyset$.

¹²The set of sub-histories of some history h_t contains the empty history and the sequences of experimental outcomes in h_t up to the ξ^{th} experiment, $\xi = 1, \dots, t - 1$.

of experimental outcomes, which is determined by the “net evidence”, $n^1 - n^0$. If both are the same at any two histories in any two equilibria, the agent’s continuation utility is the same, i.e., $v_{n^*+1}^c(n^0 + 1, n^1 + 1) = v_{n^*}^c(n^0, n^1)$.

Ceteris paribus, the more evidence the agent has to collect to meet the decision maker’s standard of evidence, the lower is the continuation value as expected experimentation costs increase, i.e., $v_{n^*+1}^c(n^0, n^1) < v_{n^*}^c(n^0, n^1)$.

If an experiment succeeds, the number of arguments still to be acquired to persuade the decision maker decreases, and the probability that an experiment yields a favorable argument increases. Hence, the agent’s experimentation incentives improve after a successful experiment, i.e., $v_{n^*}^c(n^0, n^1 + 1) > v_{n^*}^c(n^0, n^1)$.

If an experiment fails, the probability that future experiments succeed decreases, and the expected number of experiments yet to be conducted in order to be able to persuade the decision maker increases. As a consequence, continuing experimentation becomes less attractive the more experiments already failed, i.e., $v_{n^*}^c(n^0 + 1, n^1) < v_{n^*}^c(n^0, n^1)$. In particular, for each stock of favorable arguments n^1 , there exists $\tilde{n}(n^1)$ (possibly infinity) so that the agent continues the search upon observing a sequence of unfavorable outcomes if and only if $n^0 < \tilde{n}(n^1)$. $\tilde{n}(n^1)$ is weakly higher in the stock of successful experimental outcomes.

The number of experiments to be conducted to meet the standard of evidence follows a negative binomial distribution with success probability p in state $s = 1$ and success probability $1 - p$ in state $s = 0$. The expected number of experiments to meet the standard of evidence given interim belief $\gamma(n^0, n^1)$, is $\frac{(1-2p)\gamma(n^0, n^1)+p}{p(1-p)}(n^* - n^1)$. Hence, the continuation value of continuing experimentation given a stock of arguments (n^0, n^1) is at least $U - \frac{n^*-n^1}{1-p}c$, the expected utility associated with continuing the evidence acquisition until the standard of evidence is met if the agent is sure that $s = 0$, i.e. $\gamma(\cdot) = 0$.

We summarize our findings in the following Lemma.

Lemma 1 Consider $n^1 < n^*$. (i) $v_{n^*+1}^c(n^0+1, n^1+1) = v_{n^*}^c(n^0, n^1)$. (ii) $v_{n^*+1}^c(n^0, n^1) < v_{n^*}^c(n^0, n^1)$. (iii) $v_{n^*}^c(n^0, n^1 + 1) > v_{n^*}^c(n^0, n^1)$. (iv) $v_{n^*}^c(n^0 + 1, n^1) < v_{n^*}^c(n^0, n^1)$. (v) $v_{n^*}^c(n^0, n^1) \geq U - \frac{(1-2p)\gamma(n^0, n^1)+p}{p(1-p)}(n^* - n^1)c$.

Part (v) of the lemma states a lower bound for the value of continuing experimentation. Note that as soon as the agent has obtained $n^1 \geq n^* - \frac{(1-p)U}{c}$ favorable arguments, he stops experimentation only when he meets the standard of evidence. On the other hand, for any (n^0, n^1) with $n^1 < n^* - \frac{(1-p)U}{c}$, there exists a n' such that if the next n'

experimental outcomes are unfavorable, then the continuation value from experimenting drops below zero and the agent stops experimentation unsuccessfully.

Resolving indifference in favor of continuing experimentation, the optimal experimentation plan induced by a decision rule with a standard of evidence n^* is unique. As the agent stops experimenting only when he meets the standard of evidence if $n^1 \geq n^* - \frac{(1-p)U}{c}$, we have $H_f(n^*) = \emptyset$ for $n^* \leq \frac{(1-p)U}{c}$. If $n^* > \frac{(1-p)U}{c}$, we have $H_f(n^*) \neq \emptyset$.

The plan to unconditionally continue experimentation until the agent meets the standard of evidence defines a lower bound for the expected utility from engaging in experimentation: $v^c(0,0) \geq U - \frac{n^*}{2p(1-p)}c$. Hence, a sufficient condition for $H_s(n^*) \neq \emptyset$ is $n^* \leq \frac{2p(1-p)U}{c}$.

In the Appendix we describe an algorithm to determine the optimal experimentation plan on the equilibrium path, and the associated sets $H_s(n^*)$ and $H_f(n^*)$ for a standard of evidence n^* .

4.3 Persuasiveness

The statement $\hat{n} = (0, n^*)$ has an informational value if and only if neither $H_f(n^*)$ nor $H_s(n^*)$ are empty. In this case the decision maker can rule out some experimentation histories upon the observation of $\hat{n} = (0, n^*)$. Histories for which the agent optimally stops experimenting without meeting the standard of evidence are more likely to occur if $s = 0$ than if $s = 1$. If $H_f(n^*) = \emptyset$ instead, then no experimentation history yielding the standard of evidence can be ruled out. The probability that the agent experiments until he has acquired a set of evidence that allows the statement $\hat{n} = (0, n^*)$ is one in both states and the decision maker's posterior is equal to her prior. In particular, a standard of evidence $n^* < \frac{(1-p)U}{c}$ cannot be persuasive.

It is in the agent's interest to sometimes stop the search for arguments which renders his arguments informative. Given that ex ante, the agent has an incentive to search for n^* arguments, the failure to provide them means that the probability that the state of the world is in his favor is actually lower than the prior. Consequently, a successful collection of the evidence n^* boosts the posterior above the prior. Persuasive evidence has the property that the agent stops experimenting often enough such that the successful collection of the evidence indicates the favorable state with a probability that passes the

decision maker's threshold of doubt, i.e.:

$$prob\{s = 1 | \hat{n} = (0, n^*)\} = \frac{\sum_{h_t \in H_s(n^*)} prob\{s = 1 \cup h_t\}}{\sum_{h_t \in H_s(n^*)} prob\{h_t\}} \geq p_d. \quad (2)$$

Any n^* which induces an experimentation plan such that $H_s(n^*) \neq \emptyset$ and (2) holds specifies an equilibrium. The next question to address is under which conditions there is such an n^* .

If $U/c < 2$, $H_s(n^*) = \emptyset$ for any $n^* \geq 1$. Even if a single argument suffices to persuade the decision maker, it is too expensive to acquire. For these parameters there is no persuasion equilibrium (see Proposition 2 (i) below). A necessary condition for the agent to be willing to engage in experimentation is $U/c \geq 2$. A necessary condition for the evidence to be persuasive is that the agent sometimes stops experimentation unsuccessfully. Suppose that $U/c \geq 2$ and one favorable argument suffices to persuade the decision maker. If $U/c < 1/2p(1-p)$, the agent optimally experiments once and stops after observing a failure. Ex ante, the probability to find a persuasive argument is high enough to make the investment worthwhile. After an initial failure, the probability that the next experiment yields a success is too low. Suppose the agent reveals a favorable argument. The decision maker updates $prob\{s = 1 | \hat{n} = (0, 1)\} = p$. If the agent does not bring forward a favorable argument, the decision maker updates $prob\{s = 1 | \hat{n} = (0, 0)\} = 1-p$. If $p_d \leq p$, she chooses $x = 1$ if and only if the agent reveals a favorable argument. An equilibrium with a standard of evidence $n^* = 1$ exists for such a parameter constellation. (Only) in this persuasion equilibrium, the decision maker can perfectly deduce the agent's private information. Indeed, for any finite $U/c \geq 2$, there is a threshold for the signal's precision such that the agent optimally stops searching if the first experiment failed given that p exceeds the threshold, as it then becomes too unlikely to obtain a favorable outcome. If in addition $p \geq p_d$, an equilibrium with a standard of evidence $n^* = 1$ exists (see Proposition 2 (ii) below).

If one argument suffices to persuade the decision maker and $U/c \geq 1/2p(1-p)$, it is optimal for the agent to continue experimentation after an initial failure. As a consequence, the decision maker's posterior is smaller than p when the agent reveals a favorable argument. However, if $U/c < 1/(1-p)$, the agent has an incentive to eventually stop experimenting as he becomes more and more convinced that $s = 0$ and further experimentation becomes too costly. Hence, the decision maker's posterior exceeds her prior when the agent presents a favorable argument. If p_d is sufficiently close to $1/2$, an equilibrium in which the decision maker is persuaded by a single favorable

argument exists as long as $U/c < 1/(1-p)$. The higher U/c , the stronger the agent's incentive for experimentation, and the (weakly) lower the decision maker's posterior upon the report of a successful experiment. Hence, the higher U or the lower c , the (weakly) lower p_d must be for an equilibrium with a standard of evidence $n^* = 1$ to exist. If $U/c \geq 1/(1-p)$, $H_f(1) = \emptyset$, and an equilibrium with a standard of evidence $n^* = 1$ fails to exist.

We know from Lemma 1 (ii) that the decision maker can deter excessive experimentation by requiring a larger number of arguments. If the agent has to acquire a larger number of arguments in order to persuade the decision maker, this depresses his experimentation incentives at all experimentation histories. Hence, a too high standard of evidence may completely deter experimentation. From Lemma 1 (v) we can deduce that if $U/c \geq n^*/2p(1-p)$, a standard of evidence n^* induces experimentation on part of the agent. As $n^*/2p(1-p) < n^*/(1-p)$, we can always find U/c such that $H_s(n^*) \neq \emptyset$ and $H_f(n^*) \neq \emptyset$. If p_d is sufficiently close to $1/2$, the fact that $H_f(n^*) \neq \emptyset$ suffices to render the evidence n^* persuasive. For such parameters an equilibrium with a standard of evidence n^* exists (see Proposition 2 (iii)).

On the other hand, as Proposition 2 (iv) shows, if U/c is too high to render n^* persuasive, then an equilibrium with a higher standard of evidence exists if p is sufficiently high. Increasing n^* deters excessive private experimentation by rendering the search more expensive after adverse histories. However, p needs to be sufficiently high to eventually allow for separation, i.e., to deter experimentation at some, but not at all experimentation histories. If p is close to $1/2$, the ad interim success probability of the next experiment is close to $1/2$ at all experimentation histories. Hence, the agent's incentives to continue experimenting are almost exclusively driven by the number of arguments (still) to be collected. For certain parameter constellations where p is too low, any standard of evidence may only either trigger the agent's unconditional experimentation until the set of evidence is complete or induce no experimentation at all. For such a parameter constellation, no persuasion equilibrium exists.¹³

A special case is $c = 0$. If experimentation is costless, the agent could identify the state almost with certainty at zero costs. However, there is no equilibrium in which the decision maker can be persuaded with hard evidence (Proposition 2 (v)). Suppose in contrast that she could be persuaded with some set of evidence. The agent's best response is to search until he meets the standard of evidence, which happens almost with

¹³Remember that due to Proposition 1, non-existence of an equilibrium with a standard of evidence n^* implies that persuasion is impossible.

certainty as $p < 1$. The decision maker in such a hypothetical equilibrium anticipates that the standard of evidence is met regardless of the state. But then the evidence does not have an informational value and should not be persuasive, leading to a contradiction.

Proposition 2 summarizes the results of the above discussion.

Proposition 2 (i) For $U/c < 2$, there is no equilibrium with some standard of evidence n^* .

(ii) For $2 \leq U/c < \infty$, $p_d \leq p$ and $p \geq p'$, $p' < 1$, an equilibrium with a standard of evidence $n^* = 1$ exists.

(iii) For each n^* , there are \mathcal{U} and \mathcal{C} , such that an equilibrium with a standard of evidence n^* exists if $U \in \mathcal{U}$, $c \in \mathcal{C}$, and p_d is sufficiently close to $\frac{1}{2}$.

(iv) For $2 \leq U/c < \infty$ and p_d sufficiently close to $\frac{1}{2}$, if there is no equilibrium with a standard of evidence $n^* = 1$, there is an equilibrium with a standard of evidence $n^* > 1$ if $p \geq p''$, $p'' < 1$.

(v) For $c = 0$, there is no persuasion equilibrium.

If there are multiple integers between $(1-p)U/c$ and $2p(1-p)U/c$ and p_d is sufficiently close to $\frac{1}{2}$, multiple equilibria with different standards of evidence exist. A higher standard of evidence renders persuasion less likely, but is associated with a higher posterior that $s = 1$ if the agent succeeds to persuade.

Proposition 3 The probability that $s = 1$ if $x = 1$ increases in the equilibrium standard of evidence.

Consider a parameter constellation such that an equilibrium with some standard of evidence n^* exists and denote with \underline{n} the lowest standard of evidence and with \bar{n} the highest standard of evidence attainable in equilibrium for this parameter constellation. In an equilibrium with \underline{n} , the decision maker's posterior just passes her threshold of doubt when the agent announces $\hat{n} = (0, \underline{n})$. In an equilibrium with \bar{n} the agent just has an incentive to start experimentation.

Proposition 4 Any natural number between \underline{n} and \bar{n} is an equilibrium standard of evidence.

The agent unambiguously prefers equilibria with lower standards of evidence, because he persuades the decision maker with a higher probability at a lower expected cost. As the agent fails to persuade the decision maker more often in an equilibrium with a higher standard of evidence, one may conjecture that $x = 0$ is the wrong decision more often. Example 1 below illustrates that this need not always be the case.

Example 1 Let $p = \frac{7}{8}$, let $U/c \in (\frac{206}{57}, \frac{32}{7})$ and let $p_d < p$.

(a) There is an equilibrium with a standard of evidence $n^* = 1$.

$H_f(1) = \{(0)\}$, $H_s(1) = \{(1)\}$. $\text{prob}\{x = 1\} = 1/2$, $\text{prob}\{s = 1|x = 1\} = p$, $\text{prob}\{s = 1|x = 0\} = 1 - p$.

(b) There is an equilibrium with a standard of evidence $n^* = 2$.

$H_f(2) = \{(0), (1, 0, 0)\}$, $H_s(2) = \{(1, 1), (1, 0, 1)\}$. $\text{prob}\{x = 1\} = \frac{1}{2}(1 - p(1 - p))$, $\text{prob}\{s = 1|x = 1\} = \frac{p^2(2-p)}{p^2(2-p)+(1-p)^2(1+p)}$, $\text{prob}\{s = 1|x = 0\} = 1 - p$.

From an ex ante point of view, the agent's experimentation induces a lottery over the decision maker's posteriors with one outcome above p_d and one outcome below p_d . As the decision maker's expected utility is weakly convex in her posterior (it is piecewise linear with a kink at p_d), she strictly prefers lotteries with more extreme outcomes. Hence, she strictly prefers the equilibrium with a higher standard of evidence in Example 1. In Example 2, a higher standard of evidence gives rise to a trade-off: Decreasing the probability of wrongfully choosing $x = 1$ comes at the cost of increasing the probability to wrongfully choose $x = 0$. As the decision maker is more averse to the first type of error, she strictly prefers the higher standard of evidence in Example 2.

Example 2 Let $p = \frac{7}{8}$, let $U/c \in (\frac{32}{7}, \frac{50}{7})$ and let $p_d < \frac{21}{26}$.

(a) There is an equilibrium with a standard of evidence $n^* = 1$.

$H_f(1) = \{(0, 0)\}$, $H_s(1) = \{(1), (0, 1)\}$. $\text{prob}\{x = 1\} = 1/2 + p(1 - p)$, $\text{prob}\{s = 1|x = 1\} = \frac{p(2-p)}{1+2p(1-p)}$, $\text{prob}\{s = 1|x = 0\} = \frac{(1-p)^2}{(1-p)^2+p^2}$.

(b) There is an equilibrium with a standard of evidence $n^* = 2$.

$H_f(2) = \{(0), (1, 0, 0, 0)\}$, $H_s(2) = \{(1, 1), (1, 0, 1), (1, 0, 0, 1)\}$. $\text{prob}\{x = 1\} = 1/2(p^2(1 - p(1 - p)) + (1 - p)^2(1 + p))$, $\text{prob}\{s = 1|x = 1\} = \frac{p^2(2-p+(1-p)^2)}{p^2(2-p+(1-p)^2)+(1-p)^2(1+p+p^2)}$, $\text{prob}\{s = 1|x = 0\} = \frac{1-p+p(1-p)^3}{1+p(1-p)^3+p^3(1-p)}$.

In both examples, the decision maker and the agent have diametrically opposed preferences over the set of equilibria. In Example 1 (a) the agent's information provision is unbiased in the sense that from an ex ante viewpoint, the decision maker can be persuaded to choose the agent's preferred option with a probability equal to the ex ante probability that this decision is best for her. In Examples 1 (b) and 2 (b), information provision is biased against the agent, and it exhibits a bias in favor of the agent in Example 2 (a). Note that for given parameters the higher n^* , the more often the agent stops his evidence acquisition unsuccessfully. Hence, the higher n^* , the less often the agent persuades the decision maker. If there is bias in favor of the agent for $n^* = \underline{n}$ and a bias against him for $n^* = \bar{n}$, then there is a threshold standard of evidence, below which

optimal experimentation yields a bias in favor of the agent and above which information provision is biased against him.

U/c can be viewed as an indicator for the size of an interested party. We now study how the informational value of $\hat{n} = (0, n^*)$ depends on U/c for a given $n^* \in [\underline{n}, \bar{n}]$:

Proposition 5 *If n^* is an equilibrium standard of evidence for $U/c = U''/c''$ and $U/c = U'/c'$, with $U''/c'' > U'/c'$, then $\text{prob}\{s = 1 | \hat{n} = (0, n^*)\}$ is weakly lower if $U/c = U''/c''$ than if $U/c = U'/c'$.*

The higher the agent's stakes U or the lower his cost of experimentation c are, the stronger are his experimentation incentives. Hence, the maximum standard of evidence that he is willing to provide (weakly) increases in U/c . At the same time, additional experimentation by the agent, in particular after having observed adverse experimentation histories, reduces the informational value of $\hat{n} = (0, n^*)$. Consequently, the minimum standard of evidence that persuades the decision maker (weakly) increases in U/c . The following proposition summarizes these findings.

Proposition 6 *\underline{n} and \bar{n} are weakly increasing in U/c .*

In the attempt to convince a decision maker of his position, an agent with a high valuation for his preferred action typically has to provide more evidence in support for his position than an agent with a low valuation. For instance, a small group of environmentalists can be more persuasive with a given set of evidence than a big firm with the same set of evidence. Our analysis offers an explanation for this phenomenon. The value of the evidence that an interested party provides depends on the experimentation incentives. An interested party with a high valuation for the preferred decision has stronger incentives to acquire favorable arguments even if it already encountered a lot of counterarguments during the search. As we have seen above, this dilutes the value of the arguments that are finally presented to the decision maker. Consequently, the value of the same hard evidence may be different, depending on the type of interested party presenting it. If the valuations of interested parties differ sufficiently, then the corresponding intervals $[\underline{n}, \bar{n}]$ do not overlap and an interested party with a high valuation has to provide strictly more evidence for persuasion in order to deter excessive experimentation. An analogous argument holds with respect to experimentation costs.

5 Extensions

In this section, we study two extensions of our model in the context of a particular application. We consider a researcher who tries to persuade the editor of a scientific journal to publish his work in her journal. We assume that the researcher faces high-powered publication-based incentives such that he always prefers his work to be published ($x = 1$). The editor, on the other hand, needs to be sufficiently convinced that the researcher’s hypothesis is true in order to be willing to publish his work.

In the context of this application, we can distinguish “symmetric” problems from “asymmetric” ones. An example for an asymmetric problem is the question whether a certain surprising, unanticipated effect exists. Rejecting existence typically does not merit publication in a top journal. Such asymmetries are widespread. An asymmetric problem gives rise to the persuasion game that we already studied. We can directly apply the results of our analysis above. If the researcher cares sufficiently about institutional incentives (i.e., if U is high), then due to the opportunity of private experimentation and selective revelation of experimental outcomes, the arguments brought forward cannot be taken at face value. Our comparative statics results suggest that *ceteris paribus* more arguments have to be provided for successful persuasion if the researcher’s stakes are higher and/or the experimentation costs are lower. Technical innovations like the internet, faster computers and better mathematics programs have decreased experimentation costs. As a consequence, more arguments (e.g., robustness checks, or various specifications of a regression analysis) may have to be provided to persuasively support a hypothesis.¹⁴

In a symmetric problem, the acceptance and the rejection of the hypothesis under consideration are equally interesting, and both findings may be considered for publication in a top journal. In our first extension, we study the game under the alternative specification of the decision maker’s preferences.

Applying our model to the publication of research articles, it is plausible to assume an upper bound on the number of arguments that the agent can bring forward in favor of his findings. For instance, there may be a limit to the number of pages of a research article, or the editor may not be willing to handle an unlimited number of arguments. In our second extension, we impose an upper bound on the number of arguments that can be transmitted.

¹⁴The appendix of a paper published in *Econometrica* in 1985 constitutes on average roughly 6% of the paper. In 2009, it is roughly 21%. The average length of an article in *Econometrica* in 1985 is roughly 19 pages and in 2009, it is roughly 37.

5.1 Symmetric problems

We adjust the model as follows. After the experimentation phase, the agent writes a paper in which he either claims that $s = 1$ or he claims that $s = 0$. If the agent claims that $s = 1$, the decision maker's utility when choosing x in state s is:

$u_1(x, s)$	$s = 1$	$s = 0$
$x = 1$	0	$-p_d$
$x = 0$	$-(1 - p_d)$	0

If the agent claims that $s = 0$, the decision maker's utility when choosing x in state s is:

$u_0(x, s)$	$s = 1$	$s = 0$
$x = 1$	$-p_d$	0
$x = 0$	0	$-(1 - p_d)$

It is optimal for the decision maker to choose $x = 1$ if her assessment of the probability that $s = 1$ exceeds p_d and the agent argues in favor of $s = 1$, or if her assessment of the probability that $s = 0$ exceeds p_d and the agent argues in favor of $s = 0$. Otherwise, it is optimal for her to choose $x = 0$. The agent, as before, is only interested in x . If the decision rule is symmetric, the agent chooses to argue for the state that he considers more likely. We focus on equilibria in which counterarguments are not required in order to persuade the decision maker and in which she treats evidence in both directions symmetrically. The decision maker uses the rule

$$x^* = \begin{cases} 1, & \text{if } \hat{n}^0 = 0, \hat{n}^1 \geq n^* \text{ or } \hat{n}^0 \geq n^*, \hat{n}^1 = 0 \\ 0, & \text{else.} \end{cases}$$

Consider a given n^* and suppose that it is optimal for the agent to engage in experimentation. Then, he will stop experimentation if and only if $\max\{n^0, n^1\} = n^*$. For any experimentation history, the chances to meet the standard of evidence are better than ex ante as the stock of arguments that will be used for persuasion is higher. Hence, the number of arguments to be acquired is lower, and the probability to acquire a favorable outcome with the next trial is (weakly) higher as well. Thus, the optimal experimentation plan is to either keep on experimenting until the standard of evidence is met or not to start experimenting at all. The expected cost of the former experimentation plan is increasing in n^* . There exists a $c'(n^*)$, decreasing in n^* , such that it is optimal to engage in experimentation if and only if $c < c'(n^*)$.

The agent stops if and only if he has successfully acquired n^* arguments (pro or contra). Unlike in the asymmetric case, n^* does not obtain its informational value from

the agent's stopping behavior, i.e., by ruling out experimentation histories which do not occur under optimal experimentation. Instead, n^* has an informational value because the agent chooses to argue for the position for which he has acquired more arguments. The agent's posterior that he is arguing for the truth is strictly increasing in the "net" evidence $|n^1 - n^0|$ he is endowed with. The probability that the decision maker assigns to the agent arguing for the truth is equal to the ex ante expected value of all the agent's possible posteriors. The higher n^* , the more probability mass is on large realizations of $|n^1 - n^0|$, i.e., the higher is the informational value of the agent's announcement. If the informational value of an announcement n^* exceeds p_d and $c < c'(n^*)$, an equilibrium with a standard of evidence n^* exists. In particular, if $p_d < p$ and $c \leq U$, there is an equilibrium in which the agent conducts a single experiment and reports the result to the decision maker, who is persuaded by the evidence. The decision maker strictly prefers equilibria with higher standards of evidence, whereas the agent strictly prefers lower standards.

In the symmetric case, a higher experimentation cost c deters experimentation if the number of arguments needed to persuade is too high. However, the experimentation cost has no effect on the informational value of an announcement n^* . Our conclusion is that for symmetric problems, low experimentation costs are desirable in order to make engagement in research attractive.

5.2 A bound on n

In this subsection, we introduce an exogenous maximum amount of evidence N that the agent can transmit to the decision maker. In the context of the asymmetric problem, we have pointed out that excessive experimentation can be counteracted by demanding more evidence. However, if the amount of evidence that can be brought forward is limited, demanding more evidence becomes infeasible at that point.

If $c \leq \frac{U(1-p)}{N}$, no persuasion equilibrium exists. Suppose that there is a $n' \leq N$ such that the decision maker chooses $x = 1$ if the agent announces $\hat{n} = (0, n')$. If $c < \frac{U(1-p)}{n'}$, the agent has an incentive to experiment until he has acquired n' favorable arguments even if he is sure that the state of the world is 0. Thus, a necessary condition for n' to have an informational value is that $c > \frac{U(1-p)}{n'}$. If $c \leq \frac{U(1-p)}{N}$, this condition cannot be satisfied. Hence, the announcement n' is equally likely in both states such that $\text{prob}\{s = 1 | \hat{n} = (0, n')\} = 1/2$ and the decision maker is better off choosing $x = 0$.

If the amount of evidence that is necessary to be persuasive exceeds the limit due to low c or high U , then no persuasion equilibrium exists. Excessive private experimenta-

tion can be deterred by making the acquisition of arguments sufficiently expensive. In the context of our application, the scientific community can increase the costs, e.g., by imposing restrictions on arguments that are admissible. It tends to be harder to find arguments given such restrictions. This reduces private experimentation and increases the value of evidence such that the limit is not binding.

Our analysis offers an explanation for the observation that the use of new methods tends not to be well received by the scientific community.¹⁵ If a researcher deviates from standard methods, then the additional degree of freedom offers more scope for excessive private experimentation. Even if new methods have a higher explanatory power at face value, the lower costs of finding arguments can decrease their overall value due to the incentive to experiment excessively.¹⁶

In his influential book, Kuhn (1970) states that in “normal times” researchers solve puzzles with the methods inherent in the current paradigm. When addressing a research question, researchers expect certain solutions which match with the paradigm. In the context of our model, we may want to view normal times as times in which researchers predominantly work on “asymmetric problems”. In normal times, new methods are not welcome. Kuhn further says that as over time more and more anomalies appear, doubt is cast on the current paradigm and eventually a scientific revolution is triggered. New schools emerge and battle each other until a new paradigm evolves. The appearance of too many anomalies makes the scientific community reflect more on the methods that are used. At this stage the “threshold of doubt” may be lowered and/or problems may be viewed as more symmetric. Both would encourage the use of new methods.

6 Conclusion

When trying to find arguments for the preferred course of action, private experimentation and a selective revelation of the results are common practice. The process of thinking about arguments typically occurs in private. Due to an incentive to hide counterarguments, arguments cannot always be taken at face value.

Excessive private experimentation can be deterred by requiring a sufficiently large number of arguments. With each counterargument that an interested agent encounters, he becomes more and more pessimistic that he can acquire a set of persuasive arguments within the next few trials. As each trial is costly, the agent stops experimentation if the

¹⁵Kuhn (1970) documents ample evidence that new methods are not welcome in what he calls “normal times”.

¹⁶A formal argument supporting this claim can be found in the appendix.

number of arguments that he still needs to acquire in order to persuade is too large and at the same time the probability to encounter a favorable argument is too small. The fewer experiments are conducted in private, the more valuable are the agent's arguments for the decision maker. The informational value of a certain number of favorable arguments is the higher, the higher the experimentation costs and the lower the agent's stakes.

Our theory contributes to explaining why the interpretation of hard evidence may depend on whether it is presented by an agent with high stakes or by an agent with low stakes. If the former gives up experimentation unsuccessfully less often, the value of the (identical) hard evidence provided is lower.

In applications, in which time or capacity constraints impose an upper bound on the number of arguments that can be submitted to the decision maker, excessive private experimentation can depress the informational value of any feasible number of arguments so much that persuasion is impossible. Our paper offers an explanation why restrictive standards on arguments can have a value in the presence of such capacity constraints. Restrictions tend to increase experimentation costs and thereby reduce private experimentation. This can increase the arguments' informational value such that they eventually become persuasive.

In this paper, we model arguments as imperfect evidence for a certain state of the world that are acquired privately and sequentially. In future work it would be interesting to further explore the key properties of arguments and their implications in settings of strategic information transmission. In particular, it could be fruitful to take account of the possibility to support one's position with logical arguments in a mechanism design framework.

APPENDIX

Proof of Proposition 1. Part (i) follows from Lemmata A1–A3 below. Part (ii): In an equilibrium in which the decision maker always chooses $x = 0$ the agent's expected payoff is zero. In an equilibrium with a standard of evidence n^* , zero payoff is attainable by not experimenting and announcing nothing. As the agent does not choose this action, he must be weakly better off. The agent's possible announcements induce a (possibly degenerate) lottery over the decision maker's posteriors. Her expected utility is piecewise linear and weakly convex in the posterior, and exhibits a kink at p_d . She is indifferent between all equilibria in which only posteriors smaller than (or equal to) p_d realize with a positive probability, and she strictly prefers an equilibrium in which a posterior greater than p_d realizes with a positive probability. Q.E.D.

Lemma A 1 *Suppose there exists a persuasion equilibrium in which $x^* = 1$ if $\hat{n}^0 = a$, $\hat{n}^1 = b$, with $a > 0$. Then there exists a persuasion equilibrium in which $x^* = 1$ if $\hat{n}^0 = 0$, $\hat{n}^1 = b$. The latter is either payoff-equivalent or Pareto-dominant.*

Proof. Whenever the agent's information endowment allows him to make the announcement (a, b) , he can also make the announcement $(0, b)$. If his information endowment allows him to make the announcement $(0, b)$, but not (a, b) , then it is more likely that the state is 1 than if the latter announcement is available. Thus, if the decision maker is persuaded if the agent announces (a, b) , it should also be possible to persuade her with the announcement $(0, b)$. If the decision maker is persuaded by $(0, b)$ in the supposed equilibrium, then the implication in the lemma as well as payoff-equivalence immediately follows. Suppose that she is not persuaded upon the announcement $(0, b)$. Then the announcement $(0, b)$ must be an out-of-equilibrium-event attached with adverse beliefs. Then, there exists another equilibrium in which the decision maker chooses $x = 1$ for all the announcements for which she does so in the original equilibrium, and, in addition, for the announcement $(0, b)$. In the latter equilibrium, the agent makes persuasive announcements more often, and the probability that $s = 1$ conditional on the agent making a persuasive statement is also higher. Hence, the decision maker obtains a higher expected payoff. The agent is better off because he persuades the decision maker more often to choose his preferred alternative and incurs a lower experimentation cost. Q.E.D.

Lemma A1 allows us to focus on equilibria in which counterarguments are not needed in order to convince the decision maker. The next lemma further reduces the set of equilibria under consideration to those where the announcement of counterarguments would be harmful.

Lemma A 2 *Suppose there exists a persuasion equilibrium in which $x^* = 1$ iff $\hat{n}^0 \in N^0$, $\hat{n}^1 \geq b$, where N^0 is a set of natural numbers including 0. Then there exists a payoff-equivalent persuasion equilibrium in which $x^* = 1$ iff $\hat{n}^0 = 0$, $\hat{n}^1 \geq b$.*

Proof. The agent does not experiment more than necessary to persuade the decision maker. He stops (latest) if he has found b arguments in his favor. Hence, regarding experimentation, he best-responds in the same way to both decision rules. If he finds arguments against him during that search, he does not prefer any (feasible) announcement to $\hat{n}^0 = 0$. Hence, his best responses to both decision rules yield the same utility. As the agent's search behavior is identical and his announcement behavior

equivalent, the decision maker makes the same inferences (now attaching adverse beliefs to out-of-equilibrium-announcements $\hat{n}^0 > 0$). Hence, if the first decision rule is a best response, then the second one is a best response as well. The decision maker attains the same payoff in both cases. Q.E.D.

The last step is to identify a persuasion equilibrium with the minimum number of arguments needed to convince the decision maker.

Lemma A 3 *Suppose there exists a persuasion equilibrium in which $x^* = 1$ iff $\hat{n}^0 = 0$, $\hat{n}^1 \in N^1$, where N^1 is a set of natural numbers and n^* is the smallest of them. Then there exists a payoff-equivalent persuasion equilibrium in which $x^* = 1$ iff $\hat{n}^0 = 0$, $\hat{n}^1 \geq n^*$.*

Proof. Given that n^* arguments are enough to persuade the decision maker, the agent never collects more than n^* arguments in equilibrium. The decision rule for announcements $\hat{n}^1 > n^*$ is not relevant neither for the agent's experimentation and announcement strategy nor for the players' payoffs. Q.E.D.

Proof of Lemma 1. (i) The number of arguments still to be acquired is the same, the agent's assessment of the probability distribution of experimental outcomes is the same. Hence, the continuation utility is the same.

(ii) Suppose n^* arguments suffice to persuade the decision maker. Suppose the agent executes the experimentation plan that is optimal for acquiring $n^* + 1$ arguments at histories $h_t : n^1 < n^*$, and stops at histories $h_t : n^1 = n^*$. This (possibly suboptimal) experimentation plan yields a higher expected utility for the agent than that associated with optimal experimentation if he needs $n^* + 1$ arguments to persuade the decision maker as he persuades her with a (weakly) higher probability and faces strictly lower expected costs of experimentation.

(iii) With a larger stock of favorable arguments, there are less arguments still to be acquired. Moreover, ceteris paribus, favorable experimental outcomes are more likely.

(iv) The probability that an experiment yields a favorable outcome is ceteris paribus lower if the number of unfavorable arguments is higher. Hence, expected experimentation costs are higher and/or the probability to persuade is lower.

(v) The number of experiments to be conducted to meet the standard of evidence follows a negative binomial distribution with success probability p in state $s = 1$ and success probability $1 - p$ in state $s = 0$. With success probability π , the expected number of experiments to be conducted until $n^* - n^1$ successes are obtained is $\frac{n^* - n^1}{\pi}$.




Figure 1: Steps of the algorithm in case $n^* = 4$

Ad interim, the probability that the success probability is p is $\gamma(n^0, n^1)$ at history (n^0, n^1) . With probability $1 - \gamma(n^0, n^1)$, the success probability is $1 - p$. Hence, the interim expected number of experiments to be conducted until the standard of evidence is met is $\gamma(n^0, n^1) \frac{n^* - n^1}{p} + (1 - \gamma(n^0, n^1)) \frac{n^* - n^1}{1 - p} = \frac{(1 - 2p)\gamma(n^0, n^1) + p}{p(1 - p)} (n^* - n^1)$. Optimal experimentation yields at least the continuation utility $U - \frac{(1 - 2p)\gamma(n^0, n^1) + p}{p(1 - p)} (n^* - n^1)c$. Q.E.D.

Algorithm for the identification of the equilibrium experimentation plan and the associated sets of successful and unsuccessful histories

We illustrate the algorithm with the help of Figure 1, which depicts the steps of the algorithm for $n^* = 4$. Each “node” K_{n^0, n^1} in Figure 1 corresponds to a stock of failed and successful experiments (n^0, n^1) . Note that multiple experimentation histories can lead to a particular node. The posterior and the number of favorable outcomes still needed to persuade the decision maker at a particular K_{n^0, n^1} is the same for all histories leading to this K_{n^0, n^1} . Therefore, the continuation value at a particular K_{n^0, n^1} is the same for all histories leading to this K_{n^0, n^1} .

Figure 1 depicts all the histories in $H_s(n^* = 4)$ and $H_f(n^* = 4)$ associated with the agent’s optimal experimentation plan. The agent starts experimenting with a stock of arguments $(0, 0)$, i.e., at node $K_{0;0}$. The line down from $K_{0;0}$ indicates that the first

experiment fails. The line up (of the same length) indicates that the first experiment is successful. According to the experimentation plan depicted in Figure 1, the agent stops unsuccessfully in case the first experiment fails, i.e., at $K_{1;0}$. He continues searching if the first experiment is successful, i.e., at $K_{0;1}$. The following lines are interpreted analogously. As $n^* = 4$ in Figure 1, the agent stops searching successfully if he has acquired four favorable pieces of evidence. (These are the histories ending at upward-pointing lines.)

The first step is to identify the critical stock of failed experiments $n_{n^*-1}^0$, with the property that it is optimal for the agent to continue searching only once more for the last piece of favorable evidence, i.e., to continue searching at $K_{n_{n^*-1}^0;n^*-1}$ and to stop searching unsuccessfully at $K_{n_{n^*-1}^0+1;n^*-1}$.¹⁷ In order to enable us to determine optimal stopping for all $n^1 < n^* - 1$, we next calculate the continuation values at all $K_{n^0;n^*-1}$, with $n^0 < n_{n^*-1}^0$. This corresponds to step 2. We then need to identify the number of failed experiments $n_{n^*-2}^0$, where it is optimal for the agent to continue searching at $K_{n_{n^*-2}^0;n^*-2}$, and to stop searching if $K_{n_{n^*-2}^0+1;n^*-2}$. This corresponds to step 3. In order to determine optimal stopping for all $n^1 < n^* - 2$, we need to calculate the continuation values at all $K_{n^0;n^*-2}$, with $n^0 < n_{n^*-2}^0$. This corresponds to step 4. The procedure is then continued analogously for all $n^1 < n^* - 2$.

Step 1: Calculate $n_{n^*-1}^0$, the largest n^0 , such that the agent is just willing to continue searching once more if he lacks one last favorable piece of evidence for persuasion,¹⁸ i.e., the largest n^0 for which the following inequality holds.

$$\beta(n^0, n^* - 1)U + (1 - \beta(n^0, n^* - 1))0 - c \geq 0,$$

where $\beta(n^0, n^* - 1) = \gamma(n^0, n^* - 1)p + (1 - \gamma(n^0, n^* - 1))(1 - p)$ with $\gamma(n^0, n^* - 1) = \text{prob}\{s = 1 | (n^0, n^* - 1)\} = \frac{1}{1 + (\frac{1-p}{p})^{n^*-1-n^0}}$ and n^0 is the only unknown. The solution is $n_{n^*-1}^0$. The continuation value at $K_{n_{n^*-1}^0;n^*-1}$ is equal to $v_{n^*}^c(n_{n^*-1}^0, n^* - 1) =$

¹⁷For expositional convenience, we use the notation $n_{n^1}^0$ in this algorithm as the maximum number of failed experiments for which the agent is willing to conduct a further experiment given his stock of successful experiments n^1 . Note that if $\tilde{n}(n^1)$ is finite, then $n_{n^1}^0 = \tilde{n}(n^1) - 1$. If there is no finite number $n_{n^*-1}^0$, start with the largest stock of favorable outcomes $\tilde{n}^1 < n^* - 1$ for which the agent sometimes optimally stops the search unsuccessfully.

¹⁸Suppose that U , c and p are such that the agent wants to stop searching for the last remaining favorable outcome if the posterior is too low. Otherwise, the algorithm starts by identifying \tilde{n}^1 , the highest stock of favorable outcomes such that the agent stops searching unsuccessfully for some finite stock of unfavorable outcomes $\tilde{n}^0 + 1$ (see the previous footnote). With a stock of \tilde{n}^1 favorable outcomes and \tilde{n}^0 unfavorable ones, the continuation value if the next experiment is successful is then not U but equal to $U - \frac{(1-2p)\gamma(\tilde{n}^0, \tilde{n}^1+1)+p}{p(1-p)}(n^* - \tilde{n}^1 - 1)c$. The remaining part of the algorithm is analogous.

$$\beta(n_{n^*-1}^0, n^* - 1)U - c.$$

Step 2: Start with calculating the continuation value at $K_{n_{n^*-1}^0-1; n^*-1}$. The continuation value after a successful experiment is U and it is obtained with probability $\beta(n_{n^*-1}^0 - 1, n^* - 1)$. The continuation value after a failed experiment is equal to $v_{n^*}^c(n_{n^*-1}^0, n^* - 1)$, i.e., the continuation value derived in step 1. This value is reached with probability $(1 - \beta(n_{n^*-1}^0 - 1, n^* - 1))$. The continuation value at $K_{n_{n^*-1}^0-1; n^*-1}$ is $v_{n^*}^c(n_{n^*-1}^0 - 1, n^* - 1) = \beta(n_{n^*-1}^0 - 1, n^* - 1)U + (1 - \beta(n_{n^*-1}^0 - 1, n^* - 1))v_{n^*}^c(n_{n^*-1}^0, n^* - 1) - c$. The continuation values for all $K_{n^0; n^*-1}$ with $n^0 < n_{n^*-1}^0 - 1$ are successively determined analogously.

Step 3: Note that it cannot be optimal to continue searching at some $K_{n^0; n^*-2}$ with $n^0 > n_{n^*-1}^0$, as the agent has to find more evidence in his favor than above and the posterior that $s = 1$ is lower than at $K_{n_{n^*-1}^0; n^*-1}$. Start with the hypothesis that the agent is willing to continue searching at $K_{n_{n^*-1}^0; n^*-2}$. This holds true if

$$\beta(n_{n^*-1}^0, n^* - 2) \cdot v_{n^*}^c(n_{n^*-1}^0, n^* - 1) + (1 - \beta(n_{n^*-1}^0, n^* - 2)) \cdot 0 - c \geq 0,$$

where $v_{n^*}^c(n_{n^*-1}^0, n^* - 1)$ was calculated in step 2. Otherwise, successively move to smaller n^0 for given $n^1 = n^* - 2$, until the continuation value gets positive, using the continuation values $v_{n^*}^c(n^0, n^* - 1)$ determined in step 2.

Step 4: Continue analogously to step 2 for $n^1 = n^* - 2$ and n^0 smaller than the critical level derived in step 3. Continue analogously for $n^1 < n^* - 2$.

The sets of histories $H_s(n^*)$ and $H_f(n^*)$ directly follow from the above procedure: All histories that end with a successful experiment are elements of $H_s(n^*)$. All histories that end with an unsuccessful experiment are elements of $H_f(n^*)$.

Proof of Proposition 2. (i) The incentive to experiment is strongest if a single argument suffices to persuade the decision maker. The incentive is strongest at the first attempt, where the probability to succeed is $1/2$. In subsequent trials, the success probability is lower than $1/2$. For $U/c < 2$, the first attempt to acquire a favorable argument yields negative expected utility.

(ii) As $U/c > 2$, it pays to conduct an experiment which succeeds with probability $1/2$. Suppose the experiment fails. The probability that the next trial yields a success is $2p(1-p)$. The agent has no incentive for further experimentation if $2p(1-p)U < c$, i.e., if $p > \frac{1}{2} + \sqrt{\frac{1}{4} - \frac{c}{2U}} := p'$. Note that $p' < 1$ for finite U/c . If $p > p'$, $\text{prob}\{s = 1 | \hat{n} = (0, 1)\} = p \geq p_d$, such that the decision maker can be persuaded with one argument.

(iii) A sufficient condition for engaging in experimentation is $U/c \geq \frac{n^*}{2p(1-p)}$. A sufficient condition for $H_f(n^*) \neq \emptyset$ is $U/c < \frac{n^*}{(1-p)}$. $\frac{n^*}{2p(1-p)} < \frac{n^*}{(1-p)}$ ensures that some

U/c satisfies both conditions. $H_f(n^*) \neq \emptyset$ implies that $\text{prob}\{s = 1 | \hat{n} = (0, n^*)\} > \frac{1}{2}$ such that the decision maker is optimally persuaded by n^* if p_d is sufficiently close to $\frac{1}{2}$.

(iv) Consider first a standard of evidence $n' > 1$ and suppose that $H_f(n') = \emptyset$. We show that there is $n'' > n'$ such that the agent stops experimenting for some history. We need to make sure that $H_s(n'') \neq \emptyset$. If the standard of evidence is increased by one, $n^* = n' + 1$, the agent still has an incentive to engage in experimentation if $U/c \geq \frac{n'+1}{2p(1-p)}$. As $H_f(n') = \emptyset$, we know that $U/c \geq \frac{n'}{(1-p)}$. The latter inequality implies the former if $p \geq \frac{n'+1}{2n'}$. For $n' = 2$, the agent's incentive for excessive experimentation when he has to acquire two arguments implies that he is willing to engage in the search for three arguments if $p \geq \frac{3}{4}$. The higher n' , the lower the threshold above which p must lie to apply the argument. If $H_f(n' + 1) = \emptyset$, we can successively increase the standard of evidence without deterring experimentation completely until the agent sometimes stops experimentation unsuccessfully and the evidence becomes persuasive. Now, consider $n' = 1$ and suppose $H_f(1) = \emptyset$. Consider $n^* = 2$, and consider the agent's expected utility when he engages in experimentation, and (sub-optimally) stops as soon as he observes a negative outcome or he meets the standard of evidence $n^* = 2$. Then, the probability to persuade the decision maker is $\frac{1}{2}(1 - 2p(1 - p))$ and the expected cost of experimenting is $\frac{1}{2}(c + 2c)$. This experimentation plan is better than not engaging in experimentation if $U/c \geq \frac{3}{1-2p(1-p)}$. The inequality is implied by $U/c < \frac{1}{1-p}$ if $p > \sqrt{\frac{17}{16}} - \frac{1}{4} := p''$.

(v) If experimentation is costless, $H_f(n^*) = \emptyset$ for all n^* . $\text{prob}\{s = 1 | \hat{n} = (0, n^*)\} = \frac{1}{2}$ such that the decision maker cannot be persuaded. Q.E.D.

Preliminaries for the proof of Proposition 3. For the proof of Proposition 3 we need to define the probability that $s = 1$ given that the decision maker observes persuasive evidence and where she knows that the agent at some point had a stock of evidence (n^0, n^1) , i.e., that the history leading to the persuasive evidence passed through node K_{n^0, n^1} .¹⁹ Denote this probability by $\text{prob}\{s = 1 | h_t \in H_s(K_{n^0, n^1})\}$, where $H_s(K_{n^0, n^1})$ contains all the histories of $H_s(n^*)$ that pass through node K_{n^0, n^1} . $\tilde{n}(n^1)$ is the cutoff number of failed experiments such that the agent continues the search at node K_{n^0, n^1} if and only if $n^0 < \tilde{n}(n^1)$.

For the proof we need the following Lemma, which provides an order of $\text{prob}\{s = 1 | h_t \in H_s(K_{n^0, n^1})\}$ for subsequent nodes.

¹⁹The notation K_{n^0, n^1} was introduced in the algorithm for the identification of the equilibrium experimentation plan.

Lemma A 4 Consider the nodes $K_{n^0;n^1}$ that are reached with a positive probability according to a given plan. We have

- (i) $\text{prob}\{s = 1|h_t \in H_s(K_{n^0;n^1})\}$ is strictly decreasing in n^0 .
- (ii) If $\tilde{n}(n^1-1)$ is finite and $n^0 = \tilde{n}(n^1-1)-1$, then $\text{prob}\{s = 1|h_t \in H_s(K_{n^0;n^1-1})\} = \text{prob}\{s = 1|h_t \in H_s(K_{n^0;n^1})\}$.
- (iii) If $\tilde{n}(n^1-1)$ is finite and $n^0 < \tilde{n}(n^1-1)-1$, then $\text{prob}\{s = 1|h_t \in H_s(K_{n^0;n^1-1})\} < \text{prob}\{s = 1|h_t \in H_s(K_{n^0;n^1})\}$.
- (iv) If $\tilde{n}(n^1-1)$ is infinite, then $\text{prob}\{s = 1|h_t \in H_s(K_{n^0;n^1-1})\} < \text{prob}\{s = 1|h_t \in H_s(K_{n^0;n^1})\}$.

Proof of Lemma A 4. (i) It is convenient to determine the relation between the conditional probabilities backwards in the experimentation plan as follows.

(1) Start with the nodes at which the agent has collected all the evidence $n^1 = n^*$ required for persuasion. Clearly, $\text{prob}\{s = 1|h_t \in \mathcal{H}_s(K_{n^0;n^*})\} = p^{n^*}(1-p)^{n^0}/(p^{n^*}(1-p)^{n^0} + (1-p)^{n^*}p^{n^0})$ is strictly decreasing in n^0 . This proves the statement for $n^* = 1$.

(2) For $n^* > 1$, consider next all nodes $K_{n^0;n^*-1}$. If the agent is at a node $K_{n^0;n^*-1}$, then two things may happen with a strictly positive probability. Either the next experiment is successful, in which case the conditional probability is $\text{prob}\{s = 1|h_t \in \mathcal{H}_s(K_{n^0;n^*})\}$. Or the next experiment is a failure, in which case the conditional probability is $\text{prob}\{s = 1|h_t \in \mathcal{H}_s(K_{n^0+1;n^*-1})\}$, which is weakly lower than $\text{prob}\{s = 1|h_t \in \mathcal{H}_s(K_{n^0+1;n^*})\}$. As $\text{prob}\{s = 1|h_t \in \mathcal{H}_s(K_{n^0;n^*})\} > \text{prob}\{s = 1|h_t \in \mathcal{H}_s(K_{n^0+1;n^*})\}$ according to Step (1), we have $\text{prob}\{s = 1|h_t \in \mathcal{H}_s(K_{n^0;n^*-1})\} > \text{prob}\{s = 1|h_t \in \mathcal{H}_s(K_{n^0+1;n^*-1})\}$.

(3) Analogously, for $n^* > 2$, we can derive the order of $\text{prob}\{s = 1|h_t \in \mathcal{H}_s(K_{n^0;n^1})\}$ at subsequent nodes for all other nodes working backwards through the experimentation plan.

(ii) If $\tilde{n}(n^1-1)$ is finite and $n^0 = \tilde{n}(n^1-1)-1$ we have $\text{prob}\{s = 1|h_t \in \mathcal{H}_s(K_{n^0;n^1-1})\} = \text{prob}\{s = 1|h_t \in \mathcal{H}_s(K_{n^0;n^1})\}$ as all the histories through node $K_{n^0;n^1-1}$ potentially leading to successful persuasion must pass through $K_{n^0;n^1}$.²⁰

(iii) Let $n^0 < \tilde{n}(n^1-1)-1$. At such a node $K_{n^0;n^1-1}$ two things may happen with a strictly positive probability. Either the next experiment is successful, in which case the conditional probability is $\text{prob}\{s = 1|h_t \in \mathcal{H}_s(K_{n^0;n^1})\}$ or the next experiment is not successful, in which case the conditional probability is lower than at $K_{n^0;n^1}$ according to (i). Hence, $\text{prob}\{s = 1|h_t \in \mathcal{H}_s(K_{n^0;n^1})\} > \text{prob}\{s = 1|h_t \in \mathcal{H}_s(K_{n^0;n^1-1})\}$.

²⁰This is illustrated in an example experimentation plan in Figure 2 below.




Figure 2: Relations between the nodes in experimentation plan P

(iv) Analogous to (iii).

Q.E.D.

Figure 2 illustrates the relative size of $\text{prob}\{s = 1 | h_t \in H_s(K_{n^0;n^1})\}$ at subsequent nodes in an example experimentation plan, where a “ $>$ ” between two nodes $K_{n^0;n^1}$ and $K_{n^{0'};n^{1'}}$ indicates that $\text{prob}\{s = 1 | h_t \in H_s(K_{n^0;n^1})\} > \text{prob}\{s = 1 | h_t \in H_s(K_{n^{0'};n^{1'}})\}$. A “ $=$ ” is interpreted analogously.

Proof of Proposition 3. Denote the optimal experimentation plan for a given standard of evidence n^* by $P(n^*)$. We construct a plan $P'(n^* + 1)$ that is identical to $P(n^*)$ up to the collection of n^* favorable outcomes, but then the agent searches for the $(n^* + 1)^{\text{th}}$ favorable outcome until he finds it. It follows that $\text{prob}\{s = 1 | h_t \in H_s(P(n^*))\} = \text{prob}\{s = 1 | h_t \in H_s(P'(n^* + 1))\}$.²¹ Note that for each node where the agent optimally continues experimenting in $P(n^* + 1)$, the agent continues experimenting in $P'(n^* + 1)$.²² The reverse however is not true. Indeed, there must be some histories for which $P'(n^* + 1)$ prescribes to continue, but it is optimal for the agent to stop experimenting according to $P(n^* + 1)$. We can approximate $P'(n^* + 1)$ with a sequence of modified experimentation plans based on $P(n^* + 1)$ and show that $\text{prob}\{s = 1 | h_t \in H_s(\cdot)\}$ decreases with each modification. Set $P^0 = P(n^* + 1)$. Each element P^τ of

²¹ $H_s(P(\cdot))$ is the set of histories containing persuasive evidence if the agent searches according to experimentation plan $P(\cdot)$.

²²See Lemma 1 (ii).

the sequence differs from its follower $P^{\tau+1}$ only at some node $K_{n^0;n^1}$ up to which the agent continues experimenting for both plans. At $K_{n^0;n^1}$, he stops searching according to P^τ but continues searching according to $P'(n^* + 1)$. $P^{\tau+1}$ is such that the agent now continues searching at $K_{n^0;n^1}$ as follows: (i) If the next experiment is a failure, then he stops searching. (ii) If the next experiment is successful, then he continues searching according to P^τ . Hence, all new histories leading to persuasive evidence must pass through node $K_{n^0;n^1+1}$ of P^τ and, from $K_{n^0;n^1+1}$ onwards, the agent's continuation decisions are as in P^τ .²³

Note that $\text{prob}\{s = 1|h_t \in H_s(K_{i;j})\}$ and, hence, the relations between the conditional probabilities depend on the experimentation plan under consideration.²⁴ As we aim to compare P^τ with the modified plan $P^{\tau+1}$, we need to make sure that we do not confuse $\text{prob}\{s = 1|h_t \in H_s(K_{i;j})\}$ (and hence the relations between them) from different plans. It is therefore important to note that the part of the experimentation plan $P^{\tau+1}$ that follows node $K_{n^0;n^1+1}$ equals the part of the experimentation plan P^τ that follows node $K_{n^0;n^1+1}$. As all the histories through $K_{n^0;n^1}$ that yield persuasive evidence according to $P^{\tau+1}$ must pass through $K_{n^0;n^1+1}$ we have $\text{prob}\{s = 1|h_t \in H_s(K_{n^0;n^1})\} = \text{prob}\{s = 1|h_t \in H_s(K_{n^0;n^1+1})\}$ where, importantly, the right hand side of the equality refers to plan P^τ .

We need to show that $\text{prob}\{s = 1|h_t \in H_s(P^{\tau+1})\} < \text{prob}\{s = 1|h_t \in H_s(P^\tau)\}$, where $\text{prob}\{s = 1|h_t \in H_s(P^\tau)\} = \text{prob}\{s = 1|h_t \in H_s(K_{0;0})\}$. As the set of successful histories of plan $P^{\tau+1}$ consist of the “old” successful histories originating from node $K_{0;0}$ of plan P^τ and the “new” (additional) successful histories through $K_{n^0;n^1}$, which pass through $K_{n^0;n^1+1}$ of plan P^τ , it is straightforward that the inequality holds if the posterior based on persuasive evidence of the “new” successful histories is lower than the posterior based on persuasive evidence of the “old” successful histories. I.e., the inequality holds if $\text{prob}\{s = 1|h_t \in H_s(K_{0;0})\} > \text{prob}\{s = 1|h_t \in H_s(K_{n^0;n^1+1})\}$, where both conditional probabilities in the latter inequality refer to plan P^τ .²⁵ By exploiting the relative size of the conditional probabilities according to Lemma A 4 (see also Figure 2) we indeed have $\text{prob}\{s = 1|h_t \in H_s(K_{0;0})\} > \text{prob}\{s = 1|h_t \in H_s(K_{n^0;n^1+1})\}$.

In order to ensure that the cutoff number of failed experiments at which the agent

²³Note that, as $P'(n^*+1)$ prescribes to continue searching for the final favorable outcome ad infinitum, we may need an infinite sequence of such modifications to construct $P'(n^* + 1)$.

²⁴For example, at node $K_{0;0}$, i.e., where experimentation starts, $\text{prob}\{s = 1|h_t \in H_s(K_{0;0})\}$ differs for $P(n^*)$ and $P(n^* + 1)$.

²⁵By construction $\text{prob}\{s = 1|h_t \in H_s(K_{n^0;n^1})\}$ of plan $P^{\tau+1}$ is equal to $\text{prob}\{s = 1|h_t \in H_s(K_{n^0;n^1+1})\}$ of plan P^τ .

stops experimenting unsuccessfully is monotonic in his stock of successful experiments, the sequence of modified plans is constructed as follows: We subsequently increase $\tilde{n}(n^1)$ by one at nodes where $P(n^* + 1)$ and $P'(n^* + 1)$ differ, starting at the node with a stock of evidence $(\tilde{n}(n^*), n^*)$ (assuming that $\tilde{n}(n^*)$ is higher under $P'(n^* + 1)$), then $\tilde{n}(n^* - 1)$ is increased by one (again assuming that the cutoff is higher under $P'(n^* + 1)$) and so on. Once $\tilde{n}(0)$ is reached, the second round is started by increasing $\tilde{n}(n^*)$ by one more unit, working backwards through the experimentation tree again and so on. Q.E.D.

Proof of Proposition 4. The proof proceeds in three steps. In part (i) it is shown that $\text{prob}\{s = 1 | \hat{n} = (0, n^*)\} \geq p_d$ in each hypothetical equilibrium with $n^* > \underline{n}$. In part (ii) we argue that the failure to provide the equilibrium standard of evidence n^* , $n^* > \underline{n}$ yields a decision against the agent. In part (iii) we show that the agent has an incentive to start searching in each hypothetical equilibrium with $n^* < \bar{n}$. These three parts directly imply the statement.

(i) As \underline{n} by definition constitutes an equilibrium standard of evidence, the provision of this evidence passes the decision maker's threshold of doubt in an equilibrium with $n^* = \underline{n}$. According to Proposition 3 the probability that $s = 1$ if $x = 1$ increases in the equilibrium standard of evidence.

(ii) In an equilibrium with $n^* > \underline{n}$, the failure to meet the standard of evidence implies that the agent stopped searching unsuccessfully, which only happens if his posterior belief suggests that state $s = 1$ is less likely than ex ante. The decision maker can thus deduce that $s = 0$ is more likely and it is optimal to decide against the agent.

(iii) Consider two equilibria with standards of evidence $n^* = n'$ and $n^* = n''$ with $n' < n''$. In an equilibrium with n' the agent has the option to (sub-optimally) follow the same experimentation plan as in an equilibrium with n'' . Given this (suboptimal) plan he can make the announcement $\hat{n} = (0, n')$, if he has found n' favorable arguments. The optimal experimentation must yield a weakly higher ex ante payoff. Therefore, the incentive to start experimenting decreases in n^* . Q.E.D.

Verification of equilibria presented in Examples 1 and 2. (1) To verify that (a) and (b) describe equilibrium behavior, we check the agent's incentives to stop experimenting for each history in $H_f(n^*)$ and his incentive to continue experimenting for each sub-history up to the last experiment in $H_s(n^*)$ (note that this set of histories contains the set of sub-histories in $H_f(n^*)$). For the histories for which the agent stops unsuccessfully, we have $\gamma = 1 - p$. As $U/c \leq \frac{1}{2p(1-p)} = \frac{32}{7}$, it does not pay for the agent to try to find neither two nor one remaining piece of evidence. For the histories for

which he continues experimenting, γ is at least $1/2$. Hence, his incentives to continue experimenting stay intact for these histories given that he started. The incentive to start experimenting is weaker in (b) than in (a). It pays to execute the experimentation plan in (b) if $U/c \geq \frac{3+2p(1-p)}{1-p(1-p)} = \frac{206}{57}$. Note that $\text{prob}\{s = 1 | \hat{n} = (0, n^*)\} \geq p$ in (a) for $n^* = 1$ and (b) for $n^* = 2$. It follows from $p_d < p$ that the decision rule is optimal for the decision maker.

(2) Consider first (a). It is optimal for the agent to stop experimenting when the first two experiments fail if $\beta(2, 0)U - c < 0$, i.e., if $U/c < \frac{1-2p(1-p)}{p(1-p)} = \frac{50}{7}$. It is optimal to continue experimenting after one initial failure if $\beta(1, 0)U - c \geq 0$, which yields $\frac{U}{c} \geq \frac{1}{2p(1-p)} = \frac{32}{7}$. Given that it is optimal to continue after one failure, it is optimal to start ex ante. Hence, the experimentation plan is optimal for the parameter range under consideration. Given the experimentation plan, $\text{prob}(s = 1 | \hat{n} = (0, 1)) = \frac{\frac{1}{2}(p+(1-p)p)}{\frac{1}{2}(p+(1-p)p) + \frac{1}{2}((1-p)+p(1-p))} = \frac{p(2-p)}{1+2p(1-p)} = \frac{21}{26} > p_d$. Hence, it is a best response for the decision maker to choose in favor of the agent if he provides a single favorable outcome. Consider next (b). It pays to execute the suggested experimentation plan if $\frac{1}{2}(p^2 + (1-p)^2)(U - 2c) + \frac{1}{2}(p^2(1-p) + (1-p)^2p)(U - 3c) + \frac{1}{2}2p^2(1-p)^2(U - 4c) + \frac{1}{2}(-c) + \frac{1}{2}(p(1-p)^3 + (1-p)p^3)(-4c) \geq 0$, i.e., if $U/c \geq \frac{3(1+p(1-p))}{1-p(1-p)(1-2p(1-p))} = \frac{6816}{1873}$, which is satisfied for the given parameters. After an initial success, the agent's incentive to experiment stays intact. Note that after an initial success followed by a failure, the agent is in the same strategic situation as in Example 2(a). Hence, it is optimal for him to continue experimenting at history $(1, 0, 0)$ and to stop at history $(1, 0, 0, 0)$. It remains to verify that stopping is optimal after an initial failure. The agent anticipates that he optimally stops after history $(0, 0)$ (this is implied by the fact that it is optimal to stop the search for only one argument, as shown above), and that he optimally continues after history $(0, 1)$, subsequently following the plan as described in Example 2(a). Hence, his continuation utility after an initial failure is $(1-p)[p^2(U - 2c) + p^2(1-p)(U - 3c) - (1-p)c - p(1-p)^2 3c] + p[(1-p)^2(U - 2c) + p(1-p)^2(U - 3c) - pc - p^2(1-p)3c]$, which is negative if $U/c < \frac{1+3p(1-p)}{p(1-p)(1+2p(1-p))} = \frac{2720}{273}$. This is satisfied for the given parameters. Hence, it is optimal to stop experimenting after an initial failure. Given this experimentation plan, we have $\text{prob}\{s = 1 | \hat{n} = (0, 2)\} = \frac{p^2 + p^2(1-p) + p^2(1-p)^2}{p^2 + p^2(1-p) + p^2(1-p)^2 + (1-p)^2 + (1-p)^2 p + (1-p)^2 p^2} = \frac{4088}{4257}$, which is greater than p_d . Q.E.D.

Proof of Proposition 5. The larger the agent's utility and/or the lower the costs, the stronger is his incentive to engage (further) in experimentation at any history h_t . If $U/c = U''/c''$, he may continue experimentation at a history h_t with $n^0 > n^1$ for which

he stops experimentation if $U/c = U'/c'$. Thus, an analogous reasoning as in the proof of Proposition 3, i.e., to construct a sequence of experimentation plans, can be applied to prove Proposition 5. Q.E.D.

Proof of Proposition 6. Consider first \bar{n} . The larger the agent's stakes and/or the lower c , the stronger is his incentive to engage (further) in experimentation at any history h_t . Hence, if he continues experimentation if $U/c = U'/c'$ at a particular history h_t , then he does so if $U/c = U''/c''$. Consequently, the maximum standard of evidence that the agent is willing to provide is weakly higher.

Consider next \underline{n} . Consider U'/c' and U''/c'' , with $U''/c'' > U'/c'$. Due to Proposition 5, the probability that $s = 1$ conditional on reaching one of the histories in $H_s(P^*(n^*, U''/c''))$ is lower than the probability that $s = 1$ conditional on reaching one of the histories in $H_s(P^*(n^*, U'/c'))$, where $P^*(n^*, U/c)$ is the optimal experimentation plan given n^* and U/c . Hence, the lowest standard of evidence that persuades the decision maker is weakly higher, the higher U/c . Q.E.D.

Derivation of a possible barrier to new methods. Consider two experimentation technologies 1 and 2, with $p_1 < p_2$. Suppose upon observing an experimental outcome, the decision maker knows with which experimentation technology it was generated. We show that there are parameter constellations such that technology 1 can be used to persuade the decision maker, but technology 2 cannot be used to persuade the decision maker.

We prove the claim by construction. Assume that only one argument can be transmitted to the decision maker, i.e., $N = 1$. Let $p_d < p_1$, such that the decision maker can be persuaded if the agent stops experimenting unsuccessfully after the first failed trial using technology 1. Suppose that $c_1 > 2p_1(1 - p_1)U$, such that the agent indeed stops experimenting after the first failed trial. Suppose further that $c_1 < U/2$ such that the agent has an incentive to engage in experimentation. Hence, an equilibrium in which the agent persuades the decision maker with an argument acquired with technology 1 exists. If $c_2 < U(1 - p_2)$, the agent does not stop the search for a favorable argument with technology 2 until he has obtained one if the decision maker can be persuaded by such an argument. Hence, he obtains such an argument with probability one in either state of the world, such that it does not have an informational value. No equilibrium with a standard of evidence $n^* = 1$ exists.

REFERENCES

- AGHION, P., BOLTON, P., HARRIS, C. and JULLIEN, B. (1991): “Optimal learning by experimentation”, *Review of Economic Studies*, 58(4):621-654.
- AGHION, P., DEWATRIPONT, M. and STEIN, J. C. (2008): “Academic freedom, private-sector focus, and the process of innovation”, *RAND Journal of Economics*, 39(3):617-635.
- ARAGONES, E., GILBOA, I., POSTLEWAITE, A. and SCHMEIDLER, D. (2005): “Fact-free learning”, *The American Economic Review*, 95(5):1355-1368.
- AUSTEN-SMITH, D. (1994): “Strategic transmission of costly information”, *Econometrica*, 62(4):955-963.
- BERGEMANN, D. and VÄLIMÄKI, J. (2008): “Bandit problems”, *The New Palgrave Dictionary of Economics*, Steven Durlauf and Larry Blume, eds., 2nd ed. Macmillan Press.
- BOLTON, P. and HARRIS, C. (1999): “Strategic experimentation”, *Econometrica*, 67(2):349-374.
- BROCAS, I. and CARILLO, J. D. (2007): “Influence through ignorance”, *RAND Journal of Economics*, 38(4):931-947.
- CELIK, G. (2003): “Interested experts: Do they know more?”, *working paper*.
- DZUIDA, W. (2011): “Strategic argumentation”, *Journal of Economic Theory*, 146(4):1362 - 1397.
- FEDDERSEN, T. and PESENDORFER, W. (1998): “Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting”, *American Political Science Review*, 92(1):23-35.
- FELGENHAUER, M. and LOERKE, P. (2014): “Bayesian persuasion with private experimentation”, *working paper*.
- GLAZER, J. and RUBINSTEIN, A. (2001): “Debates and decisions: On a rationale of argumentation rules”, *Games and Economic Behavior*, 36:158-173.
- GLAZER, J. and RUBINSTEIN, A. (2004): “On optimal rules of persuasion”, *Econometrica*, 72(6):1715-1736.
- GLAZER, J. and RUBINSTEIN, A. (2006): “A study in the pragmatics of persuasion: A game theoretical approach”, *Theoretical Economics*, 1:395-410.
- GUL, F. and PESENDORFER, W. (2012): “The war of information”, *Review of Economic Studies*, 79(2):707-734
- HENRY, E. (2009): “Strategic disclosure of research results: the cost of proving your honesty”, *The Economic Journal*, 119:1036-1064.
- HOPENHAYN, H. and SQUINTANI, F. (2011): “Preemption games with private

information”, *Review of Economic Studies*, 78:667–692.

JOVANOVIC, B. (1982): “Truthful disclosure of information”, *Bell Journal of Economics*, 13(1):36-44.

KELLER, G., RADY, S. and CRIPPS, M. (2005): “Strategic experimentation with exponential bandits”, *Econometrica*, 73(1):39-68.

KUHN, T. S. (1970): “The structure of scientific revolutions”, Chicago: University of Chicago Press.

KAMENICA, E. and GENTZKOW, M. (2011): “Bayesian persuasion”, *American Economic Review*, 101(6):2590-2615.

LE QUEMENT, M. (2012): “Contrarian Corroboration”, *working paper*.

LEWIS, T. R. and OTTAVIANI, M. (2008): “Search agency”, *working paper*.

MILGROM, P. and ROBERTS, J. (1986): “Relying on the information of interested parties”, *The RAND Journal of Economics*, 17(1):18-32.

MILGROM, P. (2008): “What the seller won’t tell you: Persuasion and disclosure in markets”, *Journal of Economic Perspectives*, 22(2):115-131.

OLSZEWSKI, W. and SANDRONI, A. (2011): “Falsifiability”, *American Economic Review*, 101(2):788–818.

POPPER, K. R. (1959): “The logic of scientific discovery”, New York: Basic Books.

ROSENBERG, D., SOLAN, E. and VIEILLE, N. (2007): “Social learning in one-arm bandit problems”, *Econometrica*, 75(6):1591-1611.

ROTHSCHILD, M. (1974): “A two-armed bandit theory of market pricing”, *Journal of Economic Theory*, 9:185-202.

STERN, S. (2004): “Do scientists pay to be scientists”, *Management Science*, 50(6):835-853.

Copyright American Economic Association; reproduced with permission of the
American Economic Journal: Microeconomics