Strategic Private Experimentation*

Mike FELGENHAUER and Elisabeth SCHULTE†

March 10, 2011

Abstract

In this paper we view models as hard evidence obtained from experimentation. The possibility of private experimentation and selective information revelation dilutes their informational value. The scientific value of a research result depends on the experimentation costs as well as on institutional incentives. We show that "cheap" science does not have a value. Restrictions on modeling assumptions may be valuable even if other assumptions are available which better describe reality. We provide a rationale for a barrier to new methods.

Keywords: Experimentation, information provision, information acquisition.

JEL classification: D82, D83

*We would like to thank Felix Bierbrauer, Pierre Boyer, Thomas Gall, Sanjeev Goyal, Hans Peter Grüner, Martin Hellwig, Navin Kartik, Paul Milgrom, Andras Niedermayer, Georg Noeldecke, Wojciech Olszewski, Marco Ottaviani, Nicola Persico, Andrew Postlewaite, Mike Riordan, Klaus Schmidt, Ernst-Ludwig von Thadden, Michael Ziegelmeyer and the audiences at the University of Pennsylvania, MPI (Bonn), Silvaplana, ESWC 2010 (Shanghai), EEA 2010 (Glasgow) 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 Mannheim. Correspondence address: Mike Felgenhauer, University of Mannheim, Department of Economics, D-68131 Mannheim, Germany; email: felgenha(at)staff.mail.uni-mannheim.de; Tel: 0049 (0)621 1811911.

I'll tell you the value of your regression, if you tell me how many specifications you tried ...

1 How often did you try?

In science the arguments that are brought forward to support an idea are only considered as admissible if they satisfy certain standards. In theoretical economics for example an idea needs to be supported by a reasonable model. But does it make sense to impose restrictions on theories?[1] Popper (1959) argues that a theory should be formulated as restrictively as possible in order to have a high scientific value. A theory which does not exclude anything cannot be proven wrong. The restriction distinguishes the metaphysical from the scientific world.[2] Kuhn (1970) introduces the notion of a "paradigm".[3] The restrictions inherent in a paradigm have a value because (i) the researcher becomes aware of anomalies and puzzles and (ii) after observing an anomaly, the researcher investigates whether there is a fundamental problem with the paradigm before trying a new approach.[4]

This paper complements these views by taking into account the effect of institutional incentives on researchers' behavior. We argue that the scientific community imposes restrictions on admissible arguments to deter opportunistic behavior that is a by-product of institutional incentives. Academic researchers are intelligent individuals who are aware of institutional incentives. We think that it is naive to expect that incentives do not matter and that taking them into account when judging the scientific value of an argument is fundamental. Yet the strategic repercussions of institutional incentives on researchers' output have not yet been sufficiently explored. We will show that imposing a set of restrictions on arguments in response to institutional incentives can increase their scientific value.

We think that a typical researcher proceeds as follows. First he has to have a good idea,

[1]An overview on epistemology can be found in Chalmers (1999). For a critique on methods see Feyerabend (1993).

[2]There are some problems with falsification though. Consider the famous "tower argument". A critique of adherents of Ptolomy's geocentrical view of the world to Copernikus' theory that the earth turns around its own axis and the sun was as follows. If a stone is dropped from a tower, then the stone should touch the ground at a certain distance from the tower, since the earth moves while the stone falls. As this was not the case, Copernikus apparently was falsified. It took a long time before inertia was understood, which explains why the stone falls down close to the tower. The lesson is that "falsification" itself can be based on a theory which may turn out to be wrong.

[3]In a nutshell, a paradigm consists of general theoretical assumptions and concepts as well as rules on how to apply them which are generally accepted by the scientific community.

[4]Later Lakatos (1978) considers "research programmes" which are a set of theories and experimental techniques collected over time and share a common idea. These are also restrictive, since a researcher involved in a research programme must not attempt to alter the common idea. He may shield the common idea from falsification by a protective belt of auxiliary hypotheses. A researcher attacking the common idea is not part of the programme.

which we take as given.[5] In the next step the researcher tries to find an argument to support his idea. This activity is an act of strategic private experimentation and yields hard, imperfect information about whether the researcher's idea is true or false. The experimental outcome is only observed by the researcher, but there is no scope for manipulation. For instance, if a hypothesis is supported by an abstract model, then the model can be viewed as hard (imperfect) evidence that the hypothesis is true. The assumptions underlying the model are observable and the results follow logically from the assumptions. Since all models abstract, the evidence is an imperfect description of reality. It is well known that abstraction has a value, because it helps to focus the mind on what we consider as the essential features driving an effect. Formal models also help to structure arguments and facilitate communication with fellow researchers. In this paper we claim that there is an additional feature: A modeling approach helps to make arguments expensive and we will see that this can increase their scientific value.[6]

Typically, a hypothesis is not summarily rejected if the first modeling attempt is not successful. Successively other sets of assumptions are explored. It is unusual to report previous failures. Instead, experimentation results are reported selectively. The scientific value of the evidence provided depends on the researcher's experimentation plan. This is where the institutional incentives come into play. If a researcher is rewarded only if his idea is published, then this suggests that he has an incentive to excessively try to confirm his ideas.

How often a researcher searches crucially depends on the experimentation costs and on the probability with which he is eventually going to find favorable evidence. If he observes failure upon failure, then it becomes increasingly unlikely that his idea is true. This decreases the chance to find favorable evidence and hence the incentive to continue searching tends to decrease as well. At some point the researcher may stop experimenting after observing too many adverse outcomes. On the other hand the researcher stops as soon as he finds persuasive favorable evidence.[7]

When judging the provided evidence one has to form beliefs about the entire set of available

[5]Most researchers would agree that having a good idea is the most demanding part of their academic work. A good idea is the result of a brilliant moment which we do not attempt to model here. Such an idea may be the supposition of an interesting new trade-off, which no-one has noticed before.

[6]We use the language of theoretical economics although our model applies to empirical research as well: For instance, finding that a variable x shows up significantly in explaining a variable y in a regression can be seen as hard evidence that there is an impact of x on y. The evidence is imperfect as the data base is limited. The possibility of choosing different specifications and different subsets of data allows for strategic experimentation.

[7]Popper argues that there is an infinite number of theories which can "explain" a finite number of experimental observations and hence the probability that a particular theory is "true" is zero. In our paper a hypothesis can either be true or false, where the ex ante probability of both states is greater than zero. Conditional on the state there is a probability that a researcher finds an argument (a model or empirical evidence), yielding the correct prediction. We suppose that he finds support for the hypothesis with a higher probability if it is true than if it is false. An argument is therefore considered as an (imperfect) signal about the state.

evidence. The researcher's willingness to search for information has to be taken into account, which depends on his stakes and experimentation costs. If experimentation costs are low, then for any provided evidence the expected number of experiments is too high to be persuasive. Therefore this paper suggests that academic research must be costly in order to have a value. Indeed, if experimentation is costless, then, given that an editor can be persuaded, the researcher continues experimenting until he finds favorable evidence, which occurs almost with certainty if the evidence is imperfect. But then the evidence is worthless and the decision maker should not be persuaded. Our first major result is that "cheap" science does not yield persuasive evidence and hence does not have a value.

This leads to the important question how the scientific community can influence the costs of experimentation. Research can be made costly by restricting the set of assumptions from which researchers can draw their model. This can for example be accomplished by requiring that the assumptions belong to a set of established assumptions. The model presented by the researcher has to be *reasonable*. Our argument offers an explanation why many orthodox economists were skeptical and felt uneasy concerning the non-standard methods and assumptions deployed when behavioral economics entered the scene.[8] It is easier to support a hypothesis within a behavioral framework. Perhaps behavioral assumptions are better suited to describe real world behavior than the concept of "homo economicus". The advantage of imposing traditional microeconomic assumptions is that establishing consistency is harder, i.e., more expensive.[9] This makes it harder to confirm a hypothesis, thereby reduces the amount of private experimentation and hence increases the value of the information that is revealed.

An established set of assumptions has a certain drawing power. Because the scientific community is aware of strategic experimentation, deviating to a non-established set tends not to be well received. We agree with Kuhn that over time more and more anomalies appear, which makes a deviation to a new set more attractive. Still, due to the drawing power of the established assumptions, our paper suggests that researchers tend to stick too long to old methods. Even if an alternative set of assumptions is "better" in the sense that it is more "realistic" and cheaper, the old set of assumptions may be preferred. Hence researchers may be deterred from deploying new methods, which is our second major result.

The probability with which a set of assumptions confirming a hypothesis is found increases in the amount of private experimentation. This implies that the value of evidence in favor of the hypothesis decreases the more the researcher experiments on average. For a young researcher the publication decision is associated with high stakes. He should be willing to search longer for

[8]In the beginning there was also the concern that the effects found are not general and overall do not have a significant impact.

[9]As behavioral economics becomes more established, drawing assumptions from the set of established assumptions becomes more costly as well.

a set of assumptions which confirms his hypothesis than an established researcher. Hence, if an established researcher and a young researcher present the same model, then an editor should be more skeptical in the latter case even if both are considered as equally skilled.[10] Our third result is that the value of the same non-manipulable evidence depends on the utility associated with a favorable decision.[11]

The above arguments presuppose that the confirmation of an idea has a higher chance to be published than its rejection. Though we think that this asymmetry is inherent in many problems,[12] there are also problems where either solution is interesting.[13] We finally find that these two types of problems should be treated differently by the scientific community. For a symmetric problem, experimentation should be encouraged rather than deterred.

After discussing the related economics literature, we present our model of strategic experimentation and the equilibrium analysis. Then we discuss several extensions. The final section concludes.

2 Related literature

There is a related literature on strategic experimentation, where the basic underlying trade-off is between benefits from exploration and benefits from exploitation.[14] The ability to observe other players' experimentation yields a rich scope for strategic interaction. Some models have extended the analysis to settings in which outcomes (e.g., Aoyagi, 1998) or outcomes and actions (Reinganum, 1982) cannot be observed.

Brocas and Carillo (2008) study an interested party sending a flow of public experiments

[10] Young researchers have higher stakes and a deviation from standard assumptions tends to lower the value of their arguments. This should also reduce their incentives to engage in fundamentally new research. On the other hand, young researchers may more easily develop fundamentally new ideas because they have not yet been exposed to the established methods as much as the old generation. Whether it is more often the old or the young generation which contributes to fundamentally new research is an empirical question. Kuhn (1970) finds that scientific revolutions are often triggered by young researchers, whereas the old generation tends to be opposed.

[11] Ironically, the present argument is based on a model by not yet tenured researchers.

[12] In this paper, one such hypothesis is that if there is an old experimentation technology and a new technology, where the latter is more precise and cheaper, it may nevertheless be the case that exclusively the old technology is used by the researchers.

[13] Which type of problem is more prevalent depends on the state of the debate. Initially both the rejection and the confirmation of a hypothesis may be equally interesting. However, an initial answer suggests that one of the states is more likely. This influences the expectations of the solution to related problems (Kuhn, 1970). Due to these expectations the rejection or confirmation of related ideas typically are considered as asymmetrically interesting. We do not address whether a researcher prefers to work on a hypothesis where the confirmation and rejection are equally interesting or where there is an asymmetry. We rather think that good ideas are scarce and the scope to choose between these types is limited.

[14] E.g., Rothschild (1974), Bolton and Harris (1999), Bhattacharya et al. (1986), Aghion et al. (1991) and Rosenberg et al. (2007). For a recent survey see Bergemann and Välimäki (2006).

which may be stopped by one of the players.[15] They briefly consider the case in which experimentation is private and find that private and public experimentation yield the same result. In their paper the decision maker knows the number of experiments conducted if experimentation is private, but he does not observe the outcomes. Skeptical beliefs à la Milgrom and Roberts (1986) induce unraveling.[16] Hence the same information is available to the decision maker under public and private experimentation. Henry (2009) studies private experimentation (versus mandatory disclosure) in a framework in which the interested party ex ante chooses how much to invest in experimentation. The agent cannot adjust this decision in the experimentation phase. 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. Henry finds that private experimentation may be socially superior to mandatory disclosure. Our setup is different. The decision maker can only anticipate the optimal experimentation plan but not the actual number of experiments conducted, which depends on the experimentation history. From the decision maker's perspective the latter is uncertain. The decision maker's beliefs are not always degenerate such that skeptical beliefs are not always helpful and there is no unraveling.[17] In a similar setting, Celik (2003) shows that there is no equilibrium with experimentation in which the decision maker can perfectly deduce the interested party's information. Felgenhauer and Loerke (2011) builds on the present framework and compares public versus private experimentation, where the decision to continue experimenting is history dependent. They show that the comparison of the schemes is non-trivial. The private scheme is often better for the decision maker than the public scheme.

Milgrom and Roberts (1986) coined the term "persuasion game". In such games, it is commonly assumed that the individual who tries to persuade is perfectly informed. In our model instead, prior beliefs are symmetric. Glazer and Rubinstein (2001, 2004, 2006) analyze debates in a series of papers in which the debaters have hard evidence. They find that evidence

[15]Gul and Pesendorfer (2009) study competing interested parties who provide a flow of public experiments.

[16]For example, suppose that there are 10 experiments but only 1 favorable result. Since information is hard, the agent cannot lie. However he may publish a report "there was one experiment yielding adverse results and one that yielded an outcome in my favor". The agent can prove this statement. With skeptical beliefs the decision maker thinks that the agent has provided all the favorable information, but all the evidence suppressed is against him. I.e., after obtaining the message the decision maker deduces that of the ten experiments only one was in favor of the agent, which is actually correct. The agent's report is a best response to these beliefs. Hence unraveling occurs.

[17]The argument in the previous footnote requires that the number of experiments conducted is known (or deduced in equilibrium). This is not always the case in our setting. For instance, if the decision maker can be persuaded with a single favorable argument, but the researcher does not give up experimentation after the first failure, the decision maker does not know whether the favorable argument is the result of the first or the second experiment. She can only anticipate the maximum number of failed experiments the agent conducts before he unsuccessfully stops experimenting and hence deduce that the actual number of failed experiments is lower than that.

of the same quality may have a different value. In our paper exactly the same hard evidence (and not only evidence of the same quality) may have a different value. Glazer and Rubinstein assume that a player cannot present all the evidence that he is endowed with. In our paper the evidence acquisition is endogenous. If experimentation is costly, then an agent has no incentive to search for more information than necessary to persuade the decision maker.

Milgrom (2008) allows for the case that the interested party is not perfectly informed and/or cannot perfectly anticipate the result of a test. Tests can be conducted along multiple dimensions, but a test cannot be carried out multiple times. As in our setting, the possibility to hide unfavorable results induces an overinvestment in information. However, unlike in our model, the incentive to invest does not depend on previous investments in Milgrom's paper. In fact, in our model, it is precisely this dependence which renders hard information informative. In the light of uncertainty about the state of the world, Kamenica and Gentzkow (forthcoming) study persuasion from a mechanism design perspective. There, the agent is able to control the conditional probabilities with which certain signals will be issued. In our model, the agent's experimentation technology is exogenously given. The agent can affect the precision of his information by repeated experimentation.

Lewis and Ottaviani (2008) consider a principal delegating research to an agent. The principal does not observe the progress of the agent. He has to device payment schemes such that the agent reveals progress. In our model instead there are no transfers and the focus is different. Olszewski and Sandroni (forthcoming) study the formulation of theories from a contract-theoretical perspective. They show that neither a contract based on falsifiability nor one that is based on verifiability achieves the separation of informed and uniformed theorists. By contrast, a refutation contract can induce separation. In Olszewski and Sandroni's paper, it is the researcher's task to formulate a theory, while the test is carried out by a different party. In our paper instead, the researcher formulates a hypothesis and delivers support for it. Aghion et al. (2005) analyze under which circumstances research by intrinsically motivated researchers should be conducted within firms or in academia. Stern (2004) finds in an empirical analysis that academic researchers derive a value from free research. Intrinsic motivation plays a role in academic research. Yet, if publication-based rewards are helpful to induce effort, i.e., if they affect a trade-off between leisure and hard work in the search for truth, then it appears plausible that there is also a direct trade-off between striving for the truth and striving for rewards.

3 Model

Institutional incentives affect the value of a publication for its author, which in turn feeds back on its value for the scientific community. In our model the researcher exclusively cares for publishing his work.[18] Due to the institutional incentives, the game he plays is one of persuasion. We model a researcher's attempt to persuade an editor to publish his research as a game of strategic information acquisition ("experimentation") and strategic information revelation. The editor's aim is to figure out the informational content of the arguments that the researcher brings forward to support his claim. She is the decision maker in our model and decides whether to publish the researcher's work. In the basic version of the model, the editor has no access to the researcher's experimentation technology. We focus on the researcher's activity after coming up with an interesting hypothesis, say "A is a cause for B" that has the potential of being published, i.e., the editor publishes the research if the probability that the hypothesis is true is sufficiently high. The researcher (the "agent" in our model) conducts experiments to support his hypothesis. Such an experiment can be the development of a theoretical model, a regression, or the exploration of his knowledge base as in Aragones et al. (2005). An experimental outcome yields evidence in favor or against the hypothesis in the sense that it is more likely than not that A is a cause for B if A can be shown to imply B in a theoretical model, or if a correlation between A and B can be discovered in a data or knowledge base. Such evidence is imperfect.

3.1 Preferences

There are two equally likely states of the world, $s \in \{0, 1\}$. The decision maker has the choice between two actions, $x \in \{0, 1\}$. She prefers $x = 1$ if the probability that $s = 1$ is sufficiently high, and $x = 0$ else. We denote the decision maker's "threshold of doubt" with p_d and assume that $p_d \in (1/2, 1)$.[19]

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 normalized to 0. For the time being we assume that the agent holds the same prior belief about s as the decision maker. The costs from conducting experiments have to be subtracted from the gross utility. p_d being greater than $\frac{1}{2}$ means that the decision maker is ex ante biased against the agent's preferred action.

[18] Incentives for researchers in academia, such as tenure track decisions, research budgets and salaries, are often based on publications. A strong indicator for opportunistic behavior in academia is the use of single and double blind procedures in the refereeing process. Hiding the referees' identity is meant to protect them in order to induce sincere advice. If such measures are deemed necessary, then scientific progress cannot be the only motivation. Referees who are exclusively interested in unbiased scientific progress do not care to be protected. Similarly, authors exclusively interested in scientific progress do not have an incentive to play games. This paper takes no view on how an "ideal" academic researcher should be. The purpose of this paper is to formulate requirements on academic arguments such that they have a scientific value given that researchers respond to non-scientific incentives.

[19] Journals publish results only if the probability that they are correct is sufficiently high. A similar formulation of preferences is used in the literature on committee decision making (e.g., Feddersen and Pesendorfer, 1998).

3.2 Information acquisition and information transmission

The agent has access to an experimentation technology which can generate signals y_i that are correlated with the state of the world: $prob\{y_i = s'|s = s'\} = p,\ y_i \in \{0,1\},\ p \in (1/2, 1)$. The agent can conduct as many experiments as he wants. Each experiment costs $c \geq 0$. Experimental outcomes are drawn independently conditional on the state s. Let y_t denote the outcome of the t^{th} experiment. Denote with $h_t = \{y_1, ..., y_t\}$ the experimentation history after the first t experiments. Neither the number of conducted experiments nor their outcomes can be observed by the decision maker. In a history h_t, n_t^0 is the number of experiments with outcome 0 and n_t^1 is the number of experiments with outcome 1.

The agent cannot manipulate nor invent experimental outcomes. In that sense, y_i is "hard" information. However, the agent cannot prove that he did not conduct a particular experiment. If $y_i = 1$, experiment i yields an argument (or evidence) in favor of the agent, whereas $y_i = 0$ is an argument (evidence) against him. 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.

In order to capture the idea that the decision maker is not willing to handle an unlimited number of arguments, we assume that the agent can transmit at most $\overline{n}$ arguments.[20]

3.3 Timing

The timing is as follows:

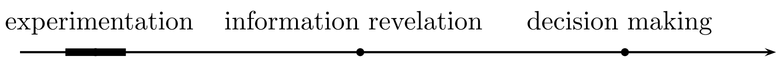

Figure 1: Timing

During the experimentation phase the agent may carry out experiments. The experimentation phase is modeled as a time interval and each experiment occurs at a point in time within this period. This implies that if the agent conducts an experiment at any given point of 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. Alternatively, think of the decision situation appearing to the decision maker only after being approached by the agent (e.g., suggesting a change of the status quo) where she does not know how long he has searched for arguments.[21] After the experimentation

[20] Glazer and Rubinstein (2001, 2004, 2006) motivate this assumption with time constraints. A possible micro-foundation for this assumption is a convex cost of information processing on part of the editor.

[21] It may well be that 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, 2008, in a

phase the agent publishes a report. Finally, the decision maker observes the announcement and chooses x.

3.4 Strategies

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 $\widehat{n} = (\widehat{n}^0, \widehat{n}^1)$, where $\widehat{n}^0$ and $\widehat{n}^1$ are the numbers of the announced unfavorable and favorable results, respectively. The case that the agent does not make an announcement is captured by $\widehat{n} = (0, 0)$. He cannot manipulate experimental outcomes nor invent arguments. Hence, if the agent stops experimenting after t experiments, then the announcement has to satisfy $\widehat{n}^0 \leq n_t^0$ and $\widehat{n}^1 \leq n_t^1$. A strategy for the decision maker is to choose $x \in \{0, 1\}$ for each possible $\widehat{n}$. In equilibrium, players' strategies have to be sequentially rational and players update their beliefs according to Bayes' Law whenever possible.

4 Equilibrium analysis

In subsection 4.1 we introduce our notion of equilibria with persuasive evidence. Next we derive the optimal experimentation plan. We then characterize the set of equilibria with persuasive evidence and state our main results. Omitted proofs 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:

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

The above decision rule maximizes her expected utility, taking into account all the information available to her. She forms beliefs regarding the information hidden from her according to Bayes' Law. Upon observing an event which occurs with probability zero in equilibrium, arbitrary beliefs are allowed. Ex ante, we have $prob\{s = 1\} = 1/2 < p_d$. Without access to information, the decision maker chooses $x = 0$. If she is confronted with an announcement $\widehat{n}$,

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. It may also not be clear when experimentation started. Put differently, the time elapsed contains information, but this information may not be particularly valuable. In this case the major results are not qualitatively affected.

the decision maker takes into account the agent's experimentation and announcement strategy when assessing the informational value of $\widehat{n}$. An announcement $\widehat{n}$ is said to have an informational value (in favor of the agent) if $prob\{s = 1|\widehat{n}\} > \frac{1}{2}$. The informational value of $\widehat{n}$ is the higher, the higher $prob\{s = 1|\widehat{n}\}$.

There exists 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, $prob\{s = 1|\ \widehat{n}\} \leq 1/2$ for all $\widehat{n}$. 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 $\widehat{n} \neq (0, 0)$ 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.[22]

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 called a persuasion equilibrium if $x^* = 1$ for at least one $\widehat{n} = n'$ and n' is announced with positive probability.*

In fact, multiple persuasion equilibria may exist due to the power that out-of-equilibrium-beliefs have in our game.

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

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

is called an equilibrium with persuasive evidence n^*.

In the following we focus on the class of equilibria with persuasive evidence. 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 persuasive evidence.*
(ii) If an equilibrium with persuasive evidence exists, it Pareto-dominates any equilibrium in which the decision maker always chooses $x = 0$.

Proof. Part (i) follows from Lemmata A1–A3 in the appendix. Part (ii): In an equilibrium in which the decision maker always chooses $x = 0$ the agent's expected payoff is zero. In

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

the equilibrium with persuasive evidence, 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 decision maker has access to more information, thus chooses the alternative which maximizes her utility more often, and is strictly better off. Q.E.D.

Focussing on the class of equilibria with persuasive evidence 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 $\widehat{n}^0 = 0$ is harmless because it is available to the agent for any information endowment.

4.2 The optimal experimentation plan

In this subsection, we take a decision rule as defined in (1) with n^* as given. Note that if $n^* > \overline{n}$, then the agent anticipates that the amount of evidence that he is able to transmit will not persuade the decision maker to choose $x = 1$. Hence, it is optimal not to engage in experimentation. Consider in the following $n^* \leq \overline{n}$. The questions which we are going to address are (i) when does the agent engage in experimentation and (ii) when does he stop experimentation. This allows us to determine under which conditions the statement $\widehat{n}^1 = n^*$ has an informational value. The agent updates his beliefs after each experiment. Denote with μ_t the probability that the agent assigns to $s = 1$ after the experimentation history h_t.

Informally, the agent conducts experiments if it is sufficiently likely to find the persuasive evidence n^*. The higher n^*, the lower is his incentive to engage in experimentation. The agent stops experimenting either if he has found persuasive evidence or if, given the experimentation history, it is too unlikely to find the remaining pieces of favorable evidence to justify the expected cost of experimentation.

On the equilibrium path the optimal experimentation plan is characterized by two sets of experimentation histories, H_s and H_f.[23] The set of histories H_s is such that for each $h_t \in H_s$ the agent successfully collects persuasive evidence. For each of these h_t it must be the case that along the way it is always optimal to continue experimenting. $H_s = \emptyset$ if the agent does not engage in experimentation. The set H_f consists of histories for which the collection of persuasive evidence fails and which may be reached by equilibrium play, i.e., the agent stops experimenting before the set of persuasive evidence is complete. For these histories it must also be the case that along the way continuing experimentation is better than stopping. However, after the final experiment it is too unlikely to find persuasive evidence and experimentation ends. In the appendix we provide an algorithm how the unique optimal experimentation plan (on and off the equilibrium path) for a given n^* can be identified. Here, we will present conditions under

[23] As H_s and H_f depend on n^*, a more concise notation would be $H_s(n^*), H_f(n^*)$. We omit the argument for better readability.

which the agent is willing to engage in experimentation (i.e., $H_s \neq \emptyset$), and eventually stops experimenting unsuccessfully (i.e., $H_f \neq \emptyset$).

Suppose experimentation starts. There is an asymmetry regarding the effect of the experimental outcomes on the agent's incentives to continue experimenting: If experiment t yields $y_t = 1$, μ_t is updated favorably, increasing the probability that n^* arguments will be collected. Moreover, having found an argument in his favor, the number of arguments yet to be discovered in order to persuade the decision maker decreases by one. On the other hand, $y_t = 0$ only yields a lower μ_t whereas the number of arguments to be acquired remains unchanged. Consequently, if at any history h_{t-1}, conducting experiment t is optimal, it is optimal to continue the search if $y_t = 1$ (unless the set of persuasive evidence is already complete). With every failed experiment μ_t decreases and converges to zero as the number of failed experiments grows larger and larger. However, this does not necessarily imply that the agent eventually stops experimenting unsuccessfully.

Denote with $\xi(n^*|n_t^1, \mu_t)$ the expected cost of experimenting until the set of persuasive evidence n^* is complete, given the arguments already acquired n_t^1 and according the belief μ_t.[24] The expected cost of this experimentation plan is the expected number of experiments times the cost per experiment, c.

For $\mu_t = 0$, the number of experiments to be conducted to complete the set of persuasive evidence follows a negative binomial distribution with success probability $(1 - p)$. Hence, the expected cost of continuing the search after reaching history h_t until n^* favorable outcomes are obtained is:

$$\xi(n^*|n_t^1, \mu_t = 0) = \frac{(n^* - n_t^1)}{(1 - p)}c.$$

A sufficient condition for $H_f = \emptyset$ is $U \geq \frac{n^*}{1-p}c$, as this implies that the agent continues to search for arguments in his favor even if he is sure that the state is against him and he has not yet acquired a single argument in his favor. In fact, this condition is also necessary for $H_f = \emptyset$. To see this, note that for $n_t^1 = 0$, $\mu_t \to 0$ as $n_t^0 \to \infty$. If $U < \frac{n^*}{1-p}c$, the agent stops experimenting for such an experimentation history in which case $H_f \neq \emptyset$. Denote $\underline{c}(n^*) \equiv \frac{U(1-p)}{n^*}$. Lemma 1 below summarizes the condition under which the agent sometimes gives up experimentation unsuccessfully.

Lemma 1 *Suppose that $x^* = 1$ if and only if $\widehat{n}^1 \geq n^*$ and suppose that the agent starts experimenting. There exists a $\underline{c}(n^*) > 0$ such that the agent stops experimenting with a positive*

[24]Note that μ_t is determined by $n_t^1 - n_t^0$. The expected cost of completing the set of persuasive arguments depends on the stock of valuable arguments that have already been acquired and the probability distribution over the outcomes of experiments yet to be conducted. We choose μ_t as an argument of the cost function instead of n_t^0 in order to make this dependence explicit.

probability before having acquired persuasive evidence if $c > \underline{c}(n^*)$ *but never stops before having acquired persuasive evidence if* $c \le \underline{c}(n^*)$.

If the agent's stakes U are sufficiently high compared to the cost of experimenting c, then he only stops after obtaining persuasive evidence (i.e., $H_f = \emptyset$). If instead U is sufficiently small, then the agent stops searching unsuccessfully after having observed too many adverse experimentation outcomes (i.e., $H_f \neq \emptyset$). The statement $\widehat{n}^1 = n^*$ has an informational value if and only if $H_f \neq \emptyset$ because in this case the decision maker can rule out some experimentation histories upon the observation of $\widehat{n}^1 = n^*$. If $H_f = \emptyset$ instead, then no experimentation history yielding persuasive evidence can be ruled out. The probability to find a set of evidence that allows the statement $\widehat{n}^1 = n^*$ is one in both states and the decision maker's posterior equals her prior.

Consider next the condition for $H_s \neq \emptyset$. Ex ante the agent is willing to engage in experimentation, i.e., $H_s \neq \emptyset$, if and only if the expected experimentation costs are smaller than the probability to acquire n^* arguments, times the gain from persuading the decision maker, U. If experimentation costs are very high, the agent never engages in experimentation. For instance, if $c > U/2$, then it is too costly to engage in information acquisition even if a single argument in favor of the agent would suffice to persuade the decision maker.

Lemma 2 *Suppose that* $x^* = 1$ *if and only if* $\widehat{n}^1 \ge n^*$. *There exists a* $\overline{c}(n^*) > 0$ *such that the agent engages in experimentation if and only if* $c \le \overline{c}(n^*)$.

To derive an explicit condition for $H_s \neq \emptyset$, suppose that the agent conducts experiments until the set of persuasive evidence is complete. If he was sure that the state is in his favor, he would face expected cost $\frac{n^*}{p}c$ for collecting n^* arguments. If, on the other hand, the agent was sure that the state is against him, the expected cost of collecting n^* arguments would be $\frac{n^*}{1-p}c$. Ex ante, the agent does not know whether he draws experiments with a success probability of p or $(1-p)$. He assigns equal probability to both possibilities. Hence, before starting to experiment, the expected cost of experimenting until n^* arguments are found is $\frac{n^*}{2p(1-p)}c$. Thus, a sufficient condition for starting experimentation is $U \ge \frac{n^*}{2p(1-p)}c$, i.e., $c \le \frac{2p(1-p)U}{n^*}$.

From $\frac{2p(1-p)U}{n^*} > \frac{U(1-p)}{n^*}$, it follows immediately that $\overline{c}(n^*) > \underline{c}(n^*)$. Hence, we have the following:

Lemma 3 *Suppose that* $x^* = 1$ *if and only if* $\widehat{n}^1 \ge n^*$. $H_s \neq \emptyset$ *and* $H_f \neq \emptyset$ *for all* $c \in (\underline{c}(n^*), \overline{c}(n^*))$. $\underline{c}(n^*) < \overline{c}(n^*)$ *for all* n^*.

Consequently, if the decision maker chooses $x = 1$ upon the provision of n^* favorable arguments, the agent starts experimenting and indeed the presented evidence is informative if $c \in (\underline{c}(n^*), \overline{c}(n^*))$. Histories for which the agent stops experimenting without having acquired persuasive evidence are more likely to occur if the state of the world is against him. Hence in expectation, persuasive evidence indicates that the state is more likely to be in his favor. Thus the statement $\widehat{n}^1 = n^*$ has an informational value.

How do $\underline{c}(n^*)$ and $\overline{c}(n^*)$ depend on n^*? It is easy to see that $\underline{c}(n^*)$ decreases in n^*. Even if the experimentation plan is optimally adjusted, the expected cost of acquiring a larger number of favorable arguments is higher. Hence, $\overline{c}(n^*)$ also decreases in n^*. As noted in Lemma 3, $\underline{c}(n^*) < \overline{c}(n^*)$ for all n^*. For any $n^* \geq 1$ we can find $c > 0$ such that $H_s \neq \emptyset,\ H_f \neq \emptyset$.

4.3 Equilibria with persuasive evidence

The next question to address is under which conditions equilibria with persuasive evidence exist. We already determined the optimal experimentation plan given the decision rule n^*. Remember that the agent does not engage in experimentation if $n^* > \overline{n}$. Hence, the decision maker chooses $x = 0$ according to his prior.

It remains to make sure that the decision rule with $n^* \leq \overline{n}$ is a best response to the experimentation plan that it induces. This means that if $\widehat{n}^1 \geq n^*$, then the informational value of the evidence has to pass the decision maker's threshold of doubt. Not providing sufficient evidence in favor of the agent instead must render $x^* = 0$ optimal.

Formally, a necessary condition for the existence of an equilibrium with persuasive evidence n^* is that the decision maker's assessment of the probability that the state is 1 given n^* is higher than her threshold of doubt, i.e.:

$$prob\{s = 1|\ \widehat{n}^1 = n^*\} \quad \geq p_d \tag{2}$$

$$prob\{s = 1|\ \widehat{n}^1 = n'\} \quad < p_d \text{ for all } n' < n^*. \tag{3}$$

Our first major result is that an equilibrium with persuasive evidence may only exist if experimentation is sufficiently costly.

Proposition 2 *If $c \leq \frac{U(1-p)}{\overline{n}}$, no equilibrium with persuasive evidence exists.*

Proof. Suppose that there is a $n' \leq \overline{n}$ such that the decision maker chooses $x = 1$ if the agent announces $\widehat{n}^1 = n'$. If $U > \frac{n'}{1-p}c$, 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 $c \leq \frac{U(1-p)}{\overline{n}}$, this condition cannot be satisfied. Hence, the announcement n' is equally likely in both states such that $prob\{s = 1|\ \widehat{n}^1 = n'\} = 1/2$ and the decision maker is better off choosing $x = 0$. Q.E.D.

The cheaper experimentation is, the more evidence is required in order to persuade the decision maker. Eventually, more evidence would have to be provided than it is feasible to transmit. At this point, persuasion becomes impossible. Note that if $c = 0$, then the state of the world can be determined almost with certainty at no costs. Nevertheless there is no information provision in equilibrium since the agent cannot commit to eventually stop experimenting. Hence, in order to (potentially) obtain valuable arguments, experimentation must be sufficiently costly. Suppose in the following that $c > \frac{U(1-p)}{\overline{n}}$.

Proposition 3 *Consider U and c such that there is a $n^* \leq \overline{n}$, which induces $H_s \neq \emptyset$ and $H_f \neq \emptyset$. There exist $p_d > \frac{1}{2}$ such that an equilibrium with persuasive evidence exists.*

Proof. We have to show that the decision rule is a best response. Since $H_f \neq \emptyset$, $\widehat{n}^1 = n^*$ has an informational value in the sense that $prob\{s = 1|\widehat{n}\} > \frac{1}{2}$. Hence, there exist $p_d > \frac{1}{2}$ sufficiently close to $\frac{1}{2}$ such that $\widehat{n}^1 = n^*$ passes the threshold of doubt and the optimal decision is $x^* = 1$ if $\widehat{n}^1 = n^*$.

The optimal decision is $x^* = 0$ if $\widehat{n}^1 = n' < n^*$, because for every $h_t \in H_f$, we have $prob\{s = 1|h_t\} < \frac{1}{2}$. This is because at the beginning of the search it is beneficial to search for n^* arguments at a prior $prob\{s = 1\} = \frac{1}{2}$. Hence, it must be beneficial to continue searching for $n \leq n^*$ arguments if $prob\{s = 1|h_t\} \geq \frac{1}{2}$. From the fact that the agent did not continue the search, the decision maker can deduce that the probability that the state is 1 is smaller than $\frac{1}{2}$. Q.E.D.

Suppose an equilibrium with persuasive evidence exists. For a given U and c denote by n' the lowest integer for which $c > \underline{c}(n^* = n')$ and the informational value of $n^* = n'$ passes the threshold of doubt p_d. Further denote by n'' the integer for which $c < \overline{c}(n^* = n'')$ and $c > \overline{c}(n^* = n'' + 1)$. Multiple equilibria exist if $n' \neq n''$.

Proposition 4 *Consider n', n'' as defined above. All $n^* \in \{n', \ldots, n''\}$ can be obtained in an equilibrium with persuasive evidence. There is no equilibrium with persuasive evidence with $n^* \notin \{n', \ldots, n''\}$.*

It is the agent's incentive not to continue the search which renders the evidence persuasive. 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 ex ante probability. 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. Any n^* which has this property and induces experimentation by the agent is attainable in equilibrium. A higher n^* yields a higher probability that the state is in favor of the agent in the case that he presents persuasive evidence, but persuasive evidence is presented with a lower probability.

An interesting question is whether the agent's information provision is biased in his favor. Information provision is called biased in favor of the agent, if $prob\{\widehat{n}^1 = n^*\} > \frac{1}{2}$, i.e., he sends a message in his favor more often than justified by the ex ante likelihood that the state is indeed in his favor. It is easy to construct examples where information provision is biased in favor or against the agent respectively for a given n^*.[25] If the agent's stakes U increase, $\underline{c}(n^*)$ increases and eventually approaches c. The agent only rarely stops experimenting without having acquired persuasive evidence. Consequently, for U sufficiently high, information provision in an equilibrium with persuasive evidence n^* is biased in favor of the agent. Put differently, the higher the stakes of the agent, the more information provision is biased in his favor.

4.4 The value of the same hard evidence

Next we analyze how an exogenous increase of the stakes U affects the informational value of the agent's announcement.

Proposition 5 *Consider U and c such that a class of equilibria with persuasive evidence exists. There is a finite $\overline{U} > U$ such that if the stakes increase above $\overline{U}$, then either the agent has to provide strictly more favorable evidence in order to persuade the decision maker than in any equilibrium in that class, or it is impossible to persuade the decision maker.*

Proof. For U and c consider the equilibrium with the highest evidence requirement and consider an arbitrary history $h_t \in H_f$. After an increase of U, the agent's incentive to search further after observing h_t weakly increases for a given n^*. If $\overline{U}$ is sufficiently larger than U, $h_t \notin H_f$. Since $p < 1$, this weakly increases the probability that n^* favorable pieces of evidence are eventually found, which decreases the informational value of $\widehat{n}^1 = n^*$. Hence, for a given p_d at some point the agent's announcement $\widehat{n}^1 = n^*$ does not induce a posterior on part of the decision maker that passes the threshold of doubt. In order to increase the informational value of the evidence the agent has to stop experimenting earlier after observing unfavorable experimental outcomes. Such an experimentation behavior can only be induced by increasing n^*. By increasing the number of required favorable outcomes, $\underline{c}(.)$ and $\overline{c}(.)$ decrease. Hence, if the stakes increase sufficiently, then a higher number of favorable outcomes is required in order to persuade the decision maker. If the stakes increase, then $\underline{c}(\overline{n})$ increases. Thus for a sufficient increase of U at some point $c < \underline{c}(\overline{n})$ and thus $H_f = \emptyset$ for any $n^* \leq \overline{n}$ and there is no equilibrium with persuasive evidence with n^*. Q.E.D.

Compare an agent with high stakes U^h and an agent with low stakes U^l with $U^h > U^l$. If the difference in the stakes is sufficiently high, then the former has to provide strictly more evidence in any equilibrium with persuasive evidence than the latter.

[25] Note that the higher n^*, the sooner the agent stops searching after observing failed experiments. Hence, the higher n^*, the less often the agent sends a message in his favor. If information provision is biased against the agent in a given equilibrium, then it is also biased against him in any equilibrium with a higher standard of evidence. Similarly, if information provision is biased in favor of the agent, then it is also biased in his favor in any equilibrium with a lower standard of evidence.

4.5 A barrier to new methods

The purpose of this section is to show that a cheap class of experiments may be ignored by the agent given that he also has access to a more expensive one, even if an experiment of the more expensive class yields less accurate information. In the context of our application, this result can be interpreted as a barrier to new methods. We consider two classes of experiments, $i = 1, 2$. The precision and the costs of class i experiments are denoted by p_i and c_i, respectively. Each experiment of class 2 has a higher precision and the experimentation costs are lower, i.e., $p_1 < p_2$ and $c_1 > c_2$.

Proposition 6 *Consider two experimentation technologies 1 and 2, $c1 > c2$. Suppose upon observing an experimental outcome, the decision maker knows with which experimentation technology it was generated. There are U, $\overline{n}$, p_d and $p_1 < p_2$ such that technology 1 can be used to persuade the decision maker, but technology 2 cannot be used to persuade the decision maker.*

Proof. We prove the proposition by construction. Choose U in between $\frac{\overline{n}}{2p_1(1-p_1)}c_1$ and $\frac{\overline{n}}{1-p_1}c_1$. Hence, it is optimal for the agent to start searching for $\overline{n}$ arguments with technology 1, and to stop after too many failures. Choose p_d sufficiently close to 1/2 to make sure that the decision maker can be persuaded by the announcement of $\overline{n}$ favorable arguments acquired with technology 1.

We can find $p_2 > p_1$ and $c_2 < c_1$ such that $\frac{\overline{n}}{1-p_2}c_2 < \frac{\overline{n}}{2p_1(1-p_1)}c_1$ which implies $\frac{\overline{n}}{1-p_2}c_2 < U$. Hence, if $\overline{n}$ arguments acquired with technology 2 would persuade the decision maker, the agent would never stop experimentation without having acquired persuasive evidence. Hence, he cannot persuade the decision maker with arguments exclusively acquired with technology 2. Likewise, this is not possible with any smaller number of arguments.

A composition of arguments acquired with both technologies cannot persuade the decision maker either, if p_1 and p_2 as well as c_1 and c_2 are sufficiently close. To see this, note that conditional on the belief $\mu_t = 0$ the expected cost of acquiring $\overline{n} - n'$ arguments with technology 1 and n' arguments with technology 2 is a convex combination of $\frac{\overline{n}}{1-p_1}c_1$ and $\frac{\overline{n}}{1-p_2}c_2$ which is strictly smaller than $\frac{\overline{n}}{1-p_1}c_1$. Hence, we can find p_1, p_2, c_1 and c_2 with $p_1 < p_2$ and $c_1 > c_2$ for which persuasion is possible for $n' = 0$, but it is impossible for any positive integer n'. Q.E.D.

Figure 2 illustrates the argument. In the figure, $U = 10$, $\overline{n} = 3$, $c_1 = 1$ and $c_2 = c_1/2$. For U below the solid black curve (which represents $\xi(\overline{n}|n_t^1, \mu_t = 0)$ when using technology 1), the agent stops experimenting with technology 1 unsuccessfully for some experimentation histories. If U is above the dotted curve, this is a sufficient condition for starting experimentation with technology 1. For $p_1 \in (p', p'')$ and p_d close enough to 1/2, an equilibrium in which the decision makers is persuaded by three favorable arguments acquired with technology 1 exists.

The gray dashed curve represents $\xi(\overline{n}|n_t^1, \mu_t = 0)$ when using only technology 2. The expected cost of acquiring three favorable arguments using both technologies is at most $\frac{c_2}{1-p_2} +$

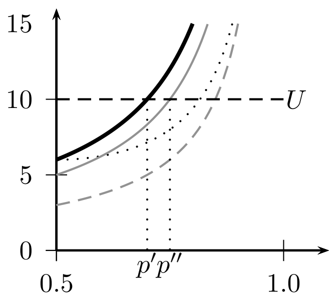

Figure 2: Illustration of Proposition 6. $U = 10$, $\overline{n} = 3$, $c_1 = 1$, $c_2 = c_1/2$. Solid black curve: $\xi(\overline{n}|n_t^1, \mu_t = 0)$ when using technology 1, dotted curve: $\xi(\overline{n}|n_t^1 = 0, \mu_t = 1/2)$ when using technology 1, gray dashed curve: $\xi(\overline{n}|n_t^1, \mu_t = 0)$ when using technology 2, gray solid curve: upper bound for $\xi(.)$ when using both technologies.

$\frac{2c_1}{1-p_1}$. This is less than $\frac{2c_1+c_2}{1-p_2}$, which is depicted by the gray solid curve. Hence, for any p_1 and p_2 in between p' and p'', persuasion is possible only with arguments acquired exclusively with technology 1.

If the superior technology 2 is not much cheaper than the inferior technology 1, then it will also not be used by the agent for learning the state of the world. The reason is that the experimental outcomes of technology 1 have the advantage that in addition to learning about the state of the world, they can be used to persuade the decision maker. If instead experimenting with technology 2 is very cheap, the agent may use the technology to acquire information to get a better estimate of the costs and benefits from investing in the search of persuasive evidence. In this case, the agent is privately informed before the investment in persuasive evidence starts. We discuss the case of an ex ante informed agent in subsection 5.2.

4.6 Symmetric problems

So far, we considered the case that the agent wants to convince the decision maker of a particular state of the world. With reference to our application, the researcher wants to persuade the editor that a surprising effect exists. Evidence against existence is useless, because the editor will not publish a paper on the non-existence of an unanticipated effect, no matter how certain he is that the effect indeed does not exist. Such asymmetries are widespread. Nevertheless, there are applications in which evidence for either state is valuable. In this section, we argue that in such a case experimentation should be encouraged rather than deterred.

Let us adjust the model accordingly. Suppose the decision maker chooses $x = 1$ if it is sufficiently likely that the evidence correctly predicts the state of the world. Formally, the decision maker optimally chooses $x = 1$ if the probability of one state, given the evidence provided, exceeds p_d and he chooses $x = 0$ otherwise. The other assumptions are maintained. In accordance with the arguments underlying Proposition 1 we focus on equilibria in which counterarguments are not required in order to persuade the decision maker and she treats evidence in both directions symmetrically. The decision maker uses the rule

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

Consider a given n^* and suppose that it is optimal for the agent to engage in experimentation. Then, he will only stop experimentation if he has collected n^* arguments in one direction. For any experimentation history, the chances to find persuasive evidence are better than ex ante as the stock of arguments that will be used for persuasion is higher (hence the number to be acquired is lower) and the probability to find a favorable outcome with the next trial is (weakly) higher as well. Thus, the optimal experimentation plan is to either keep on experimenting until persuasive evidence is found 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 $c < c'(n^*)$.[26]

As the agent stops when and only when having 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 $|n_t^1 - n_t^0|$ 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 $|n_t^1 - n_t^0|$, i.e., the higher is the informational value of the agent's announcement. If the informational value of the announcement $\overline{n}$ exceeds p_d and $c < c'(\overline{n})$, an equilibrium with persuasive evidence exists. In particular, if $p > p_d$ and $c < U$, there is an equilibrium in which the agent conducts a single experiment and reports the result to the principal, who is persuaded by the evidence.

In the symmetric case, a higher experimentation cost c deters experimentation if n^* is too high. Unlike in the asymmetric case, a higher experimentation cost cannot increase the informative value of an announcement n^*. With respect to our application, the conclusion is that for symmetric problems, experimentation costs should be as low as possible in order to make experimentation attractive.

[26]Note that this reasoning applies only in a completely symmetric setting. For instance, if the agent has state-dependent preferences, he may stop experimenting after observing (too many) experimental outcomes against the state in which he receives the highest payoff.

5 Extensions

In this section, we extend our analysis of the asymmetric case. First, we discuss the case that the decision maker has access to the same experimentation technology as the agent. Then, we consider the case of an ex ante informational advantage on the agent's side.

5.1 Experimentation by the decision maker

Suppose that the decision maker has the option to conduct experiments on her own and that the experimentation costs are equal for the agent and the decision maker. If experimentation is costless, then the decision maker can determine the state almost with certainty and hence make the correct decision almost with probability one without studying the evidence provided. Our major argument nevertheless qualitatively holds if experimentation is cheap in the sense that $U \gg c$, but higher than the decision maker's stakes. Then, the decision maker does not experiment in equilibrium.

As the agent's utility U from a favorable decision is very high, a large amount of favorable evidence is required in order to persuade the decision maker, eventually exceeding $\overline{n}$. If at the same time information acquisition is too costly for the decision maker, then there is no equilibrium in which the agent engages in research. We find it plausible to assume that an editor's experimentation costs are significant relative to his own stakes, and that a researcher's stakes from a good publication are high. Hence, "cheap science" does not yield scientific progress.

5.2 Ex ante informed agent

If the agent privately conducts experiments, then he is ad interim better informed than the decision maker. In fact, in equilibrium n^* is chosen such that it becomes too costly to conduct further experiments if the (ad interim) probability that the state is in the agent's favor becomes too low. In that sense, n^* induces a separation of (ad interim) types, namely those endowed with unfavorable information from those endowed with information favorable enough to render further experimentation worthwhile.

If the agent has an informational advantage based on soft information to start with, then experimentation may serve as a pure signaling device.[27] Consider the case that the agent is perfectly informed about the state. If the agent engages in experimentation, then he only stops after having collected persuasive evidence. The expected cost of this experimentation plan is $\frac{n^*}{(1-p)}c$ if $s = 0$ and $\frac{n^*}{p}c$ if $s = 1$. In equilibrium, n^* has to be (weakly) higher than $\frac{U(1-p)}{c}$. Otherwise the agent engages in experimentation if the state is against him in which case the presented evidence cannot persuade the decision maker. Any n^* for which $\frac{n^*}{1-p} > \frac{U}{c} \geq$

[27]Experimentation plays a similar role when the agent is fully informed in Celik (2003).

$\frac{n^*}{p}$ induces separation of the types and allows the decision maker to deduce the state of the world. Although experimentation in this case does not yield any new information, still many experiments may be necessary in order to prove one's point.

Remember that $\frac{n^*}{(1-p)}c > U$ is a necessary condition that n^* is informative in the case that the agent has no ex ante informational advantage (which does not yet imply that n^* is also persuasive), whereas $U \geq \frac{n^*}{p}c$ is not sufficient to induce experimentation in that case. Persuasion is easier if the agent has an ex ante informational advantage in the sense that (weakly) less evidence suffices to persuade the decision maker and that persuasion equilibria may exist where no persuasion equilibrium exists if the agent has no informational advantage. Without any informational advantage, the agent first has to incur the cost of learning "his type" before the separation of types is possible.

6 Conclusion

Spectacular cases of academic fraud in which evidence is tempered occur from time to time. These manipulations yield distortions and reduce the faith in science, but they appear to occur rarely. Private experimentation instead is common practice. It rarely happens that a researcher summarily rejects a publishable hypothesis if he fails to find support for it in his first trial. Typically the researcher instead tries a new approach. This yields a fundamental problem: On the one hand, experimentation is necessary to produce information. On the other hand, presented evidence cannot be taken at face value.

Experimentation occurs until the probability of success, which decreases with each failure, is outweighed by the costs. The fewer experiments are conducted in private, the more valuable is the evidence provided. Private experimentation can be deterred by increasing experimentation costs. In academia, increasing the costs of private experiments is possible by requiring that researchers use a set of established assumptions. Confirming a particular hypothesis then becomes harder. These established assumptions may have a value in science even if they are not as well suited to describe reality as an alternative set of assumptions which is "more flexible" and easier to work with. Our theory contributes to explaining why the same hard evidence may be interpreted asymmetrically if it is presented either by a high stake agent (like a non-tenured researcher) or a low stake agent (like an established researcher). If the former has an incentive to conduct more private experiments, the value of the (identical) hard evidence provided is lower.[28] Finally, we find that the standards on arguments should differ depending on whether the confirmation of an idea is either more interesting than its rejection or both findings are

[28] A researcher who can easily generate new publishable ideas reduces the number of experiments he spends on a particular hypothesis. If it is known that a researcher is smart in this sense, then his work is also seen less skeptical.

equally interesting for a journal to publish.

Institutional incentives affect the behavior of academic researchers in the sense that they may excessively experiment.[29] The question arises why such incentives are in place. The analysis in this paper is partial. It is well understood that rewards based on output, such as publications, can be useful in order to induce effort. Output can also signal ability. In general, publication based incentives in conjunction with restrictions on arguments may perform well compared to other schemes in a world with private information.

Technical innovations like the internet, faster computers and better mathematics programs decrease experimentation costs. In the context of our model, this tends to increase experimentation and therefore decreases the informational value of a given piece of evidence. In order to be accepted as convincing evidence, the technical requirements on a model may therefore increase.[30] A technically demanding model may not be better per se, but such a requirement increases the costs of experimentation, which in turn reduces the average number of experiments and thus increases their informational value.

On the other hand it is costly for an author to find a persuasive model in the class of simple models. Researchers tend to be more skeptical in case they are confronted with a model of intermediate complexity, where many auxiliary assumptions are made. The necessity of introducing these auxiliary assumptions suggests excessive private experimentation.

It would be interesting to study strategic experimentation in the light of competition. If multiple competing interested parties simultaneously have to publish their results, then competition should have an effect on the amount of experimentation and the decision quality. If publication is sequential, then the parties can condition their behavior on the past evidence provided. The first experimenter may try to provide a sufficient amount of evidence that deters costly experimentation by his competitors. In academic publishing, editors in general rely on the advice of referees. If referees invest in experimentation, their presence should have a mitigating effect on excessive experimentation. However, we think that the incentive does not disappear.[31]

Finally, the origin of institutional incentives, like publication based rewards, is another interesting issue to be explored.

[29] In the case that the researcher has state dependent preferences our major results qualitatively hold if the agent's threshold of doubt is sufficiently lower than that of the decision maker.

[30] The appendix of a paper published in Econometrica in 1985 constitutes 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. In our view, this is an indicator that the technical requirements for publishing a paper in Econometrica have increased.

[31] Felgenhauer and Schulte (2010) point out that the refereeing process may be biased and thus may add complications.

APPENDIX

Proof of Proposition 1

In the following three lemmata, we will provide the basis for the statement in Proposition 1 that any persuasion equilibrium is payoff-equivalent to or Pareto-dominated by an equilibrium in which the decision maker uses the following decision rule:

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

Lemma A 1 *Suppose there exists a persuasion equilibrium in which* $x^* = 1$ *if* $\widehat{n}^0 = a,\ \widehat{n}^1 = b$*, with* $a > 0$*. Then there exists a persuasion equilibrium in which* $x^* = 1$ *if* $\widehat{n}^0 = 0,\ \widehat{n}^1 = b$*. 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. 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 1 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 to those where the announcement of counterarguments would be harmful.

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

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 $\widehat{n}^0 = 0$. 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 $\widehat{n}^0 > 0$). 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 $\widehat{n}^0 = 0,\ \widehat{n}^1 \in N^1$, 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 $\widehat{n}^0 = 0,\ \widehat{n}^1 \geq n^*$.*

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 $\hat{n} > n^*$ is not relevant neither for the agent's experimentation and announcement strategy nor for the players' payoffs. Q.E.D.

Proof of Lemma 2

Given the decision rule, the agent engages in experimentation if $c = 0$ and, as U is bounded, does not engage in experimentation if $c = \infty$. It remains to show monotonicity. Suppose that the agent engages in experimentation if $c = c_1$ and consider $c = c_0 < c_1$. If the agent applies the same stopping rule as he does for $c = c_1$, he persuades the decision maker with the same probability, but incurs lower expected experimentation costs. Thus, experimentation with that (possibly not optimal) stopping rule yields a strictly higher expected utility than not engaging in experimentation if experimentation is optimal for $c = c_1$. Suppose that the agent does not engage in experimentation if $c = c_0$ and suppose, contrary to the Lemma, that he engages in experimentation if $c = c_1 > c_0$. The same argument as above applies: If $c = c_0$, the agent can replicate the stopping rule used for $c = c_1$ and obtain a strictly higher expected utility than if $c = c_1$. Hence, it cannot be optimal not to engage in experimentation for $c = c_0$ if experimentation is optimal for $c = c_1$, yielding a contradiction.

Proof of Proposition 4

We distinguish two cases. Case (i) if $n' = n''$, the first statement immediately follows. Consider in the following case (ii): $n' < n''$.

(a) $c > \underline{c}(n') \Rightarrow c > \underline{c}(n)$ for all $n^* > n'$, because $\underline{c}(n^*)$ is decreasing in n^*.

(b) $c < \overline{c}(n'') \Rightarrow c < \overline{c}(n)$ for all $n^* < n''$, because $\overline{c}(n^*)$ is decreasing in n^*.

(a) and (b) imply that $H_s \neq \emptyset$, $H_f \neq \emptyset$ for all $n^* \in \{n', \ldots, n''\}$.

It remains to show that in an equilibrium with $n^* > n'$, the optimal experimentation plan is such that the decision maker is indeed persuaded by the announcement n^*, i.e., his posterior passes his threshold of doubt. We prove this statement by induction, assuming that it is true for n^*, $n' \leq n^* < n''$ and show that this implies that it is true for $n^* + 1$.

Consider an arbitrary history h_t with $\sum_{i=1}^{t} y_i < n^*$ (on or off the equilibrium path). If $n^* + 1$ arguments are required to persuade the decision maker, the expected payoff from continuing experimentation is strictly lower than if only n^* arguments are required (the expected costs are higher the probability to persuade is (weakly) lower). The expected payoff from stopping experimentation is equal in both cases. We conclude that for any h_t with $\sum_{i=1}^{t} y_i < n^*$ for which the agent continues the search if $n^* + 1$ arguments are required, he also continues the search if n^* arguments are required, i.e., every such h_t that occurs with positive probability on the equilibrium path to $n^* + 1$ arguments in the equilibrium with persuasive evidence $n^* + 1$, occurs with positive probability on the equilibrium path to n^* arguments in the equilibrium with persuasive evidence n^*. This implies that the set of experimentation histories h_t which occur with positive probability in the equilibrium $n^* + 1$ and have the properties $\sum_{i=1}^{t} y_i = n^*$ and $y_t = 1$, are a strict subset of $H_s(n^*)$. All histories that are in $H_s(n^*)$, but not in the previously described set of histories are associated with a lower probability that $s = 1$ than any history in that set. As a consequence, conditional on $\sum_{i=1}^{t} y_i = n^*$ and $y_t = 1$, the probability that $s = 1$ is higher in the equilibrium $n^* + 1$ than the decision maker's posterior in the equilibrium with persuasive evidence n^*. After having acquired n^* arguments in the equilibrium with persuasive evidence $n^* + 1$, the agent either (i) experiments until he finds the last remaining piece of favorable evidence, or (ii) stops experimentation for some histories. In case (i), in the equilibrium $n^* + 1$, the decision maker's posterior is equal to the probability that the $s = 1$ conditional on $\sum_{i=1}^{t} y_i = n^*$ and $y_t = 1$, which was already found to be more favorable than the decision maker's posterior in the equilibrium with persuasive n^*. In case (ii), the decision maker's posterior is even higher, as the optimal stopping behavior rules out the worst experimentation histories. Hence, if optimal experimentation aimed at finding n^* arguments leads to a posterior on part of the decision maker that passes his threshold of doubt, than so does optimal experimentation aimed at finding $n^* + 1$ arguments. By construction, the statement is true for n', which completes the proof.

As n' is the lowest integer for which $c > \underline{c}(n^* = n')$ and the informational value of $n^* = n'$ passes the threshold of doubt, $H_s = H_f = \emptyset$ for all $n^* < n'$. Moreover, we have $c > \overline{c}(n^* = n'' + 1)$ and hence $H_s = H_f = \emptyset$ for all $n^* > n''$. Thus, there is no equilibrium with persuasive evidence $n^* \notin \{n', \ldots, n''\}$.

The optimal experimentation plan

In the following, we provide an algorithm for the derivation of the optimal experimentation plan for a given n^* and $c > 0$. We introduce some definitions which improve the exposition.

Definition 3 *h_t contains $h_{t'}$ if $t > t'$ and $h_t = \{h_{t'}, y_{t'+1}, \dots, y_t\}$. Otherwise, h_t does not contain $h_{t'}$. We write $h_{t'} < h_t$ if h_t contains $h_{t'}$, and $h_{t'} \not< h_t$ if h_t does not contain $h_{t'}$.*

Definition 4 *$h_{t'}$ is contained in a set H if there is an $h_t \in H$ that contains $h_{t'}$, and $h_{t'} \notin H$.*

Definition 5 $H_n := \{h_t : n_t^1 = n, y_t = 1\}$.

H_s denotes the set of histories that are reached with positive probability in equilibrium after which the agent stops experimenting, having successfully acquired n^* arguments. H_f denotes the set of histories that are reached with positive probability in equilibrium after which the agent stops experimenting, without having acquired n^* arguments. If the agent does not plan to ever continue the search until the persuasive evidence is complete, then he does not start to search.

Observation 1 *If $H_s = \emptyset$, then $H_f = (0, 0)$.*

Suppose in the following $H_s \neq \emptyset$. The agent stops experimenting latest when he has collected persuasive evidence.

Observation 2 *H_s is a subset of H_{n^*}.*

If an experimentation history is reached with positive probability in equilibrium, then continuing to experiment is optimal before the last experiment in that history.

Observation 3 *If $h_{t''} \in H_s \cup H_f$, then for all $h_{t'} < h_{t''}$ the following inequality holds:*

$$\sum_{h_t : h_t \in H_s \wedge h_{t'} < h_t} prob\{h_t | h_{t'}\} \cdot U \geq \sum_{h_t \in H_s \cup H_f} prob\{h_t | h_{t'}\} \cdot (t - t') \cdot c.$$

Our equilibrium concept requires that the experimentation plan be optimal, even if an unexpected observation is made, i.e., the agent behaves optimally after any experimentation history, including those that occur with probability zero in equilibrium.[32]

We hypothesize that the agent has already acquired a certain stock of arguments in his favor and identify the histories (from that point on) for which the agent stops experimenting without completing the set of persuasive evidence. Starting with only one missing argument, i.e., the

[32]From the agent's point of view, the experimental outcomes are chosen by a non-strategic player such that there are no degrees of freedom with respect to the determination of his beliefs associated with an out-of-equilibrium-event.

set of histories $h_t : n_t^1 = n^* - 1$, this part of the stopping plan is optimal no matter how the rest of the stopping plan looks like. In the following steps of the derivation we can anticipate subsequently optimal stopping behavior. $\widehat{H}_s$ and $\widehat{H}_f$ denote the hypothesized versions of H_s and H_f that are relevant at each step.

Step 0 *Set* $\hat{H}_s = H_{n^*}, \hat{H}_f = \emptyset$.

Step 1 *Identify* H_{n^*-1,n^*}*, the set of histories which are contained in* H_{n^*} *and which are elements or contain an element of* H_{n^*-1}.

Step 2 *Consider* $h_{t'} \in H_{n^*-1,n^*}$. $h_{t'} \in \hat{H}_f''$ *if and only if* $\sum_{h_t : h_t \in \hat{H}_s \wedge h_{t'} < h_t} prob\{h_t | h_{t'}\} \cdot U < \sum_{h_t \in \widehat{H}_s \cup \widehat{H}_f} prob\{h_t | h_{t'}\} \cdot (t - t') \cdot c$.

Step 3 *Consider* $h_t \in \hat{H}_s$. *If any* $h_{t'} < h_t$ *is an element of* $\hat{H}_f''$*, then* $h_t \notin \hat{H}_s'$. *Else,* $h_t \in \hat{H}_s'$. *Consider* $h_t \in \hat{H}_f \cup \hat{H}_f''$. *If any* $h_{t'} < h_t$ *is an element of* $\hat{H}_f \cup \hat{H}_f''$*, then* $h_t \notin \hat{H}_f'$. *Else* $h_t \in \hat{H}_f'$.

Step 4 *Set* $\hat{H}_s = \hat{H}_s', \hat{H}_f = \hat{H}_f'$.

Repeat (as long as applicable) steps 1-4 subsequently for the (analogously defined) sets of histories $H_{n^*-2,n^*-1}, \ldots, H_{0,1}$. Note that this "backward induction"-procedure yields a unique solution.

REFERENCES

AGHION, P., BOLTON, P., HARRIS, C. and JULLIEN, B. (1991): "Optimal learning by experimentation", *The Review of Economic Studies*, Vol. 58, No. 4, pp. 621-654.

AGHION, P., DEWATRIPONT, M. and STEIN, J. C. (2005): "Academic freedom, private-sector focus, and the process of innovation", *working paper.*

ARAGONES, E., GILBOA, I., POSTLEWAITE, A. and SCHMEIDLER, D. (2005): "Fact-free learning", *The American Economic Review,* Vol. 95, No. 5, pp. 1355-1368.

AOYAGI, M. (1998): "Mutual observability and the convergence of actions in a multi-person two-armed bandit model", *Journal of Economic Theory*, Vol. 82, pp. 405-424.

BHATTACHARYA, S., CHATTERJEE, K. and SAMUELSON, L. (1986): "Sequential research and the adoption of innovations", *Oxford Economic Papers*, New Series, Vol. 38, Supplement: Strategic Behaviour and Industrial Competition, pp. 219-243.

BERGEMANN, D. and VÄLIMÄKI, J. (2006): "Bandit problems", *Cowles Foundation discussion paper* No. 1551.

BOLTON, P. and HARRIS, C. (1999): "Strategic experimentation", *Econometrica*, Vol. 67, No. 2, pp. 349-374.

BROCAS, I. and CARILLO, J. D. (2008): "Influence through ignorance", *The RAND Journal of Economics*, Vol. 38, No. 4, pp. 931-947.

CELIK, G. (2003): "Interested experts: Do they know more?", *working paper.*

CHALMERS, A. F. (1999): "What is this thing called science?", 3rd edition, Hackett Pub Co.

FEDDERSEN, T. and PESENDORFER, W. (1998): "Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting", *The American Political Science Review*, Vol. 92, No. 1, pp. 23-35.

FELGENHAUER, M. and LOERKE, P. (2011): "Public versus private experimentation", mimeo.

FELGENHAUER, M. and SCHULTE, E. (2010): "Preselection and expert advice", *working paper.*

FEYERABEND, P. (1993): "Against method", 3rd edition, London: Verso.

GLAZER, J. and RUBINSTEIN, A. (2001): "Debates and decisions: On a rationale of argumentation rules", *Games and Economic Behavior*, Vol. 36, pp. 158-173.

GLAZER, J. and RUBINSTEIN, A. (2004): "On optimal rules of persuasion", *Econometrica*, Vol. 72, No. 6, pp. 1715-1736.

GLAZER, J. and RUBINSTEIN, A. (2006): "A study in the pragmatics of persuasion: A game theoretical approach", *Theoretical Economics*, Vol. 1, pp. 395-410.

GUL, F. and PESENDORFER, W. (2009): "The war of information", *working paper.*

HENRY, E. (2009): "Strategic disclosure of research results: the cost of proving your honesty", *The Economic Journal*, Vol. 119, pp. 1036-1064.

HOPENHAYN, H. and SQUINTANI, F. (2008): "Preemption games with private information", *working paper*.

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

KAMENICA, E. and GENTZKOW, M. (forthcoming): "Bayesian persuasion", *American Economic Review*.

LAKATOS, I. (1978): "The methodology of scientific research programmes: Philosophical papers volume 1", Cambridge: Cambridge University Press.

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*, Vol. 17, No. 1, pp. 18-32.

MILGROM, P. (2008): "What the seller won't tell you: Persuasion and disclosure in markets", *Journal of Economic Perspectives*, Vol. 22, No. 2, pp. 115-131.

OLSZEWSKI, W. and SANDRONI, A.: "Falsifiability", *The American Economic Review*, forthcoming.

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

REINGANUM J. F. (1982): "Strategic search theory", *International Economic Review*, Vol. 23, No. 1, pp. 1-17.

ROSENBERG, D., SOLAN, E. and VIEILLE, N. (2007): "Social learning in one-arm bandit problems", *Econometrica*, Vol. 75, No. 6, pp. 1591-1611.

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

STERN, S. (2004): "Do scientists pay to be scientists", *Management Science*, Vol. 50, No. 6, pp. 835-853.