Public versus private experimentation*

Mike FELGENHAUER[†] and Petra LOERKE[‡]

May 23, 2013

Abstract

This paper considers an agent who tries to persuade a principal with experimental evidence. We compare sequential private and public experimentation. Under private experimentation, the agent collects hard evidence and selectively reveals the results. Under public experimentation, the principal observes the experimental outcomes. We show that the persuasion probability is lower under private experimentation. We find that the agent tends to prefer public experimentation, whereas the principal tends to favor private experimentation. Under this scheme, the principal can benefit from being sceptical. There can be a deterrence effect that renders the private scheme superior for the principal and the agent.

Keywords: Experimentation, information acquisition, information revelation.

JEL classification: D82, D83

*We thank Rosella Argenziano, Susanne Goldlücke, Johannes Hörner, Andras Niedermayer, Andrew Postlewaite, Elisabeth Schulte, Nicolas Schutz, Dmitry Shapiro, Ran Spiegler, Roland Strausz, Ernst-Ludwig von Thadden and Juuso Välimäki for valuable comments and suggestions. Some of this work was done while Mike Felgenhauer was visiting UPenn. Their hospitality is greatfully acknowledged.

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

[‡] University of Mannheim, Department of Economics, 68131 Mannheim, Germany; email: ploerke@rumms.uni-mannheim.de.

1 Introduction

This paper studies a situation in which an agent tries to persuade a principal by providing experimental evidence. The aim of this paper is to compare public and private experimentation. Under private experimentation, the agent may privately run as many experiments as desired and selectively reveal the results. Under the public scheme instead, the principal observes the experimental outcomes from all experiments that were conducted.[1] Under either scheme, we assume that the outcome of an experiment is hard (i.e., non-manipulable) evidence, that experimentation is costly and that the agent's decision to continue collecting evidence depends on the experimentation history.

Our main contribution is the normative comparison of repeated private and public experimentation from the principal's and the agent's payoff perspective. We find that it is often not in the principal's interest to implement public experimentation. Given that an agent has completed a certain number of studies, it would be best for the principal if all results were revealed. Intuition correctly suggests that unfavorable outcomes may be withheld under the private scheme. This is per se detrimental and cannot occur under the public scheme. Nevertheless, we show that, if handled properly, there are many circumstances where the principal obtains a higher payoff if the agent can hide unfavorable results.

Our focus is on an important aspect in favor of private experimentation: the principal's ability to be sceptical. Under the private scheme, the principal can force the agent to run high quality experiments. The agent's incentive to search for evidence of a required quality depends on his belief about the state of the world. If it is too unlikely that the state is in his favor, then experimentation costs render continuing the search until high quality persuasive evidence is found unattractive. The agent only fails to provide persuasive evidence if his posterior belief that the state is in his favor is weak. Hence, if the agent reveals, e.g., low quality information, then the principal

[1]For example, a law may require that all experiments have to be a priori registered. It is sufficient for our analysis that a regulator knows a priori that the registered experiment will be conducted. In this case, the regulator can force the agent to reveal the outcome by having sceptical beliefs a la Milgrom and Roberts (1986) and, hence, the experiment can be viewed as public. Public experimentation can be implemented in practice by employing whistleblower policies with resulting detection probabilities and appropriate punishments in case there were private experiments.

has good reasons to believe that the posterior that the agent observes is against the agent. Consequently, the decision should be against the agent. The decision rule only to accept high quality evidence for persuasion is, thus, credible and the agent is forced to search for such evidence. Under public experimentation, this is not possible. The principal cannot commit not to be persuaded once the flow of public experimental outcomes yields a posterior belief that passes her threshold of doubt. The agent on the other hand does not have an incentive to continue experimenting once the principal is persuaded.

Private experimentation allows to be sceptical and, therefore, it can dominate public experimentation from the principal's perspective.[2] The agent, however, can suffer from his option to hide unfavorable results. Under the public scheme, he can choose the experiment precision such that there is a higher probability to persuade the principal with only one experiment.

As an application, consider a pharmaceutical company (the agent) that attempts to persuade the U.S. Food and Drug Administration (the principal) to approve a newly developed drug. Given the enormous R&D costs in the pharmaceutical industry, it is plausible that the company prefers that a new drug is approved, even if its merits are doubtful. The FDA instead would like to make the "appropriate" decision, which could be against the company. The FDA mainly has to rely on tests, e.g., clinical studies, provided by the company, which in turn has an incentive to behave strategically. The decision quality can be influenced by the rules under which evidence can be revealed and what evidence is permitted to be considered as decision-relevant. These rules have an important effect on the value of the evidence provided.

Our second contribution is to show that private experimentation limits the extent to which persuasion is possible. Kamenica and Gentzkow (2011) show how an agent

[2] Note that these beliefs differ from sceptical beliefs in Milgrom and Roberts (1986). Sceptical beliefs a la Milgrom and Roberts require that the number of conducted experiments is known. If the number is observed (or deduced in equilibrium), then a principal having sceptical beliefs thinks that all outcomes which are not revealed are against the agent. Under private experimentation in our model, we will see that the principal does not observe the number of conducted experiments. He can deduce the equilibrium experimentation plan, but when the agent stops experimenting depends on the experimentation history. Hence, the actual number of experiments is in general unknown to the principal and sceptical beliefs a la Milgrom and Roberts are not feasible. In contrast to the case with a known number of experiments the beliefs in the present paper, in general, cannot force the agent to reveal all experimental outcomes that he collected.

optimally tries to persuade a principal if the agent can design a hard public signal and they derive the maximum persuasion probability of a Bayesian principal, which is going to correspond to the persuasion probability under public experimentation in our model.[3]

Often, however, information acquisition of hard evidence occurs in private and the maximum persuasion probability a la Kamenica and Gentzkow is not achieved in all equilibria under private experimentation.[4] Even if we restrict attention to sender preferred equilibria, there is no equilibrium with the same persuasion probability as under public experimentation for intermediate experimentation cost. On the one hand if the agent could use the same signal as under public experimentation, then he would excessively search for it in the private case if experimentation costs are not too high. Continuing the search for a favorable result can be profitable even if the agent observes adverse outcomes, as there is a chance that he finds persuasive evidence with the next experiments. Naturally, he would not reveal any adverse experimental outcomes. Excessive private experimentation and selective information revelation dilutes the value of revealed evidence such that a revealed outcome that has the same structure as in Kamenica and Gentzkow may not be persuasive. On the other hand if experimentation costs are not too low, the agent prefers to persuade the principal with one positive result with high precision rather than to search for several positive results. Due to the high precision the persuasion probability in the former case can be lower.

We assume in most of the paper that the quality of the experiments is endogenously chosen. This allows us to highlight the role of beliefs in a parsimonious framework. In an extension, the precision of the experiments is considered as exogenous. For example, in practice it may often not be possible or excessively expensive to run ex-

[3]We will see that if the agent has the option to run multiple public experiments, then an analogous reasoning as in Kamenica and Gentzkow applies and the persuasion probability is the same.

[4]Felgenhauer and Schulte (2013) interpret, e.g., logical arguments as decision relevant hard information. Their reasoning is as follows: If the assumptions underlying a logical argument are revealed, then they cannot be manipulated. Everything that follows from the assumption is logical and logic can also not be manipulated. Logical arguments tend to have persuasive power, therefore, they can be viewed as signals about a decision relevant state of the world. The signals are imperfect, as the underlying assumptions do not cover every real world aspect. A thought experiment (i.e., drawing a set of assumptions and making a logical deduction) yields a signal. For such arguments the assumption that they are acquired in private by running a series of thought experiments and selectively revealed for persuasion is more natural.

periments of very high quality. In this case, assuming an exogenous precision can be a useful approximation. The basic intuition about the influence of beliefs on the equilibrium still holds. The principal can credibly demand many experimental outcomes which tends to increase the informational value of the evidence provided. However, additional effects play a role. For example, there can be a deterrence effect that favors private experimentation even from an agent's perspective. If the experimentation costs are high and the precision of the experiment is low, then it may not pay to start experimenting under the public scheme. The agent knows that an adverse initial outcome cannot be easily reversed. The possibility to hide an unfavorable outcome under the private scheme renders starting more attractive. As valuable information is provided under the private scheme, it can be superior for the principal as well as the agent.

The paper is structured as follows. In the following section we discuss the related literature. Next, we present the model with an endogenous precision and compare private and public experimentation. Then, we analyze a setting with an exogenous precision and study a modified public scheme. Finally, we review our results and conclude.

2 Literature

Our paper is related to two strands of literature: the literature on experimentation (e.g., Rothschild, 1974, Bolton and Harris, 1999, Aghion et al., 1991, and Keller et al., 2005) and the literature on persuasion (e.g., Jovanovich, 1982, Milgrom and Roberts, 1986, Glazer and Rubinstein 2001, 2004, 2006, and Kamenica and Gentzkow, 2011). The present and few other papers (e.g., Henry, 2009, Brocas and Carillo, 2008, Argenziano et al., 2011, and Felgenhauer and Schulte, 2013) combine persuasion and experimentation.

Henry (2009) and Brocas and Carillo (2008) investigate private versus public experimentation in settings that, in contrast to the present paper, allow for an unraveling argument. Henry assumes that the agent ex ante commits to the number of experiments that he is going to run. Under the public scheme, the principal observes the number of experiments. Under private experimentation, she does not. But since she knows the agent's optimization problem, she can deduce the equilibrium number of experiments under private experimentation and have sceptical beliefs a la Milgrom and Roberts (1986). Henry shows that there are more experiments under private experimentation than under public experimentation and that either scheme can be optimal depending on circumstances. Brocas and Carillo primarily focus on public experimentation. They briefly consider private experimentation and argue that both schemes are equivalent. In this part the principal knows the number of experiments conducted under the private scheme. The equivalence result follows directly also using sceptical beliefs. The comparison of the two schemes in both of these papers crucially depends on the principal knowing or inferring the number of experiments that the agent runs under the private scheme. However, given that experimentation occurs in private and is sequential, we think that it is more natural that the decision to continue experimenting is history dependent and unobservable.[5] Sceptical beliefs a la Milgrom and Roberts are not helpful in such a setting: The principal, in general, cannot deduce the number of experiments that the agent ran, she only knows the equilibrium experimentation plan.[6] As the analysis in both papers depends on an unraveling argument, it is not clear whether the results on the comparison of the schemes hold if the decision to continue experimenting is history dependent.

Argenziano et al. (2012) also compare a covert and an overt scheme, but, in contrast to our work, in a cheap talk framework. The overt and the covert scheme differ with respect to whether the number of experiments can be observed by the principal. They assume that the agent chooses the number of experiments upfront. Furthermore, overt experimentation does not correspond to public experimentation in our model, since under overt experimentation the principal only knows the number of experiments and not the outcomes. Under covert information acquisition, the communication game has additional asymmetric information. Therefore, they find that the principal is

[5]For example, if the agent finds too many adverse results in the first experiments, then he knows that he cannot persuade the principal by conducting the remaining experiments. As experimentation is private, he cannot be forced to continue costly experimentation until the ex ante determined number of experiments is conducted.

[6]In a similar model with private experimentation Celik (2003) shows that there is no equilibrium with experimentation in which the principal can perfectly deduce the agent's information.

better off when experimentation is overt.

Our paper combines elements of Kamenica and Gentzkow (2011) and Felgenhauer and Schulte (2013). Kamenica and Gentzkow show under which conditions the agent can benefit from information acquisition and persuasion. They derive the optimal precision of a signal to persuade another person. We use the same signal generation technology for a particular experiment. We therefore apply their idea to a setup with costly experimentation where experiments are run sequentially and the decision to continue is history dependent. This leads to the same optimal signal precision as in Kamenica and Gentzkow under public experimentation, but not under the private scheme. Felgenhauer and Schulte exclusively study the value of hard evidence under private experimentation with an exogenous precision of the experiments. In contrast to their work, the focus of the present paper is on the comparison of the private and the public scheme with endogenous precision.

3 Model

3.1 Preferences

A principal has to choose $x \in \{0,1\}$. Her payoff depends on a state of the world $s \in \{0,1\}$, where both states are equally likely. The principal's utility is

	$s=1$	$s=0$
$x=1$	1	$1-p_d$
$x=0$	p_d	1

with $p_d \in (\frac{1}{2},1)$. The principal, thus, would like to match the decision with the state of the world if she knew s. At the optimum she only chooses $x=1$ if her posterior belief passes the "threshold of doubt" p_d, i.e., the posterior that $s=1$ must be weakly greater than p_d. The principal's expected utility is denoted by EU_j^p, where $j \in \{pub, priv\}$ indicates the scheme, with *pub* referring to public and *priv* to private. In the absence of information provision the principal at the optimum chooses $x=0$ and her utility is $EU^p = \frac{1}{2}+\frac{1}{2}p_d$. If there exists an equilibrium under private or public experimentation with information provision the principal's utility is[7]

$$EU_j^p = \frac{1}{2} + \frac{1}{2}p_d + \underbrace{(prob\{s=1 \mid x=1\} - p_d)prob\{x=1\}}_{\text{excess utility}}, \tag{1}$$

where $prob\{s = 1 \mid x = 1\}$ is the probability that $s = 1$ given that the principal chooses $x = 1$.[8] The first part of (1) corresponds to the principal's utility in the absence of information provision. The remaining part is the "excess utility" the principal obtains if the agent engages in experimentation. This part is strictly greater than zero if $prob\{s = 1 \mid x = 1\} > p_d$ and $prob\{x = 1\} > 0$. If $prob\{s = 1 \mid x = 1\} = p_d$, then the principal is just persuaded to choose $x = 1$ and the excess utility is equal to zero. In this case the principal does not benefit from experimentation.

There is an agent who prefers $x = 1$ regardless of the state of the world. In case $x = 1$, the agent's gross utility is U and otherwise it is normalized to 0. We assume that the agent holds the same prior belief about s as the principal. Experimentation costs have to be subtracted from the gross utility. The agent's expected utility is EU_j^a, $j \in \{pub, priv\}$. Denote the expected number of experiments until the agent stops experimenting by n. This leads to the agent's expected utility

$$EU_j^a = prob\{x=1\}U - nc. \tag{2}$$

3.2 Experimentation

The agent has access to an experimentation technology that can generate signals about s. The agent can run as many experiments as he likes. Let $y_t \in \{0, 1\}$ be the outcome of the t^{th} experiment. The agent can choose the precision of an experiment t, i.e., $prob\{y_t = 1 \mid s = 1\} = p_t^1$ and $prob\{y_t = 0 \mid s = 0\} = p_t^0$, $p_t^i \in [0, 1]$, $i \in \{0, 1\}$. Define $p_t \equiv (p_t^1, p_t^0)$. We assume that $p_t^1 \geq 1 - p_t^0$, which means that an outcome $y_t = 1$ is more likely if $s = 1$ than if $s = 0$. In the major part of the paper the

[7]The derivation is provided in the appendix.

[8]Let M^* be the set of messages with which the agent can persuade the principal in equilibrium. Then $prob\{s = 1 \mid x = 1\} = \sum_{m^* \in M^*} \frac{prob\{s=1|m^*\}prob\{m^*\}}{prob\{x=1\}}$ where $prob\{m^*\}$ is the probability that evidence m^* is provided in equilibrium. $\frac{prob\{s=1|m^*\}prob\{m^*\}}{prob\{x=1\}}$ corresponds to the principal's posterior probability that $s = 1$, given that he observes evidence that induces him to choose $x = 1$.

agent can make a history dependent choice of the precision. The principal observes the quality of an experiment, once the outcome of this experiment is presented.[9] For simplicity, we assume that all precisions p_t of experiment t are equally costly, i.e., running an experiment costs $c \geq 0$ regardless of the precision.[10] Denote with $h_t = \{(y_1, p_1), ..., (y_t, p_t)\}$ the experimentation history after the first t experiments and the posterior after experiment t by $\overline{p}_t$. The agent cannot manipulate or invent experimental outcomes. In this sense, y_t is "hard" information. After each experiment the agent updates his assessment of the probability distributions regarding s and decides whether to continue or to stop experimenting. Under the public scheme, the experimentation history is common knowledge.

Under the private scheme, the principal cannot observe the experimentation history. The agent can reveal any subset of the experimental outcomes to the principal. Denote an element i of this subset by $(\hat{y}_i, \hat{p}_i^1, \hat{p}_i^0)$, where $\hat{y}_i$ is the outcome and $(\hat{p}_i^1, \hat{p}_i^0)$ characterizes the precision of the revealed experiment i. The principal does not observe the time index t of the experiment, i.e., the index i identifies an element but does not relate to t. The agent cannot prove that he did not conduct a particular experiment.

3.3 Timing

The timing for the private scheme is such that first there is an experimentation phase, which we model as a time interval. At each point of time within the experimentation phase the agent may carry out an experiment. This implies that if the agent runs an experiment at any point of time, then he may still carry out as many experiments as desired before the experimentation phase ends. After the experimentation phase the agent reveals a report and then the principal chooses x.

Under the public scheme, there is again first an experimentation phase, but now

[9] We think that this assumption is natural in many applications. Once a pharmaceutical company reports a clinical study, then experts at the FDA are able to assess its quality. Similarly, if a theoretical argument is considered as evidence, as in Felgenhauer and Schulte (2013), then a scientific audience, e.g., referees, editors or seminar participants, can judge the quality of the argument.

[10] This simplification highlights the effect of asymmetric information and the role of the principal's beliefs under private experimentation. With increasing costs the role of the principal's beliefs in the comparison of public and private experimentation is the same.

there is no announcement stage, since the experimentation history is common knowledge, and the experimentation phase is directly followed by the principal's decision.

3.4 Strategies and equilibrium concept

Under the private scheme, a strategy for the agent consists of an experimentation plan and an announcement plan. For each possible experimentation history h_t the experimentation plan specifies whether to continue or to stop experimenting and which precision to use. The announcement plan states what, if anything, to reveal to the principal after stopping. As the agent cannot manipulate experimental outcomes, the announcement has to be such that he reveals weakly less positive and adverse results than he observes. A strategy for the principal is to choose $x \in \{0, 1\}$ for each possible announcement plan.

Under public experimentation, a strategy for the agent exclusively consists of an experimentation plan. A strategy for the principal is to choose $x \in \{0, 1\}$ for each possible experimentation history h_t.

The equilibrium concept that we use is weak perfect Bayesian equilibrium. In addition, we require that the principal forms beliefs taking into account that the agent cannot signal what he does not know.[11]

4 Equilibrium analysis

This section analyzes the equilibria under both experimentation schemes. Omitted proofs can be found in the appendix.

[11] To illustrate the implications consider public experimentation, where the experimentation history is common knowledge. If the agent deviates from his equilibrium experimentation plan, then this assumption implies that the principal cannot have arbitrary beliefs with respect to the state of the world, since the agent does not have additional information about the state. It follows that the off-the-equilibrium path beliefs with respect to the state of the world have to be Bayesian based exclusively on the publicly observable experimentation history.

4.1 Public experimentation

Under public experimentation, there is a unique equilibrium. For each experimentation history h_t, there is a common knowledge posterior belief $prob\{s = 1 \mid h_t\}$, which follows from Bayes' law. The principal is persuaded to choose in favor of the agent given history h_t if $prob\{s = 1 \mid h_t\} \geq p_d$. Suppose that experimentation costs are low. Under the optimal experimentation plan, the agent chooses the precision such that one positive result can just persuade the principal, but also such that he does not have an incentive to search further if the first outcome is adverse. This is in line with the results of Kamenica and Gentzkow (2011).[12] Proposition 1 fully characterizes the equilibrium under public experimentation.

Proposition 1 *The agent's equilibrium experimentation strategy is as follows. If he faces a history h_{t-1} where the posterior $\overline{p}_{t-1}$ is such that*

- $\overline{p}_{t-1} < \frac{cp_d}{U}$, *then the agent stops experimenting (unsuccessfully).*
- $\frac{cp_d}{U} \leq \overline{p}_{t-1} < p_d$, *then the agent runs the t'th experiment with $p_t^1 = 1$ and $p_t^0 = \frac{p_d - \overline{p}_{t-1}}{p_d(1-\overline{p}_{t-1})}$.*
- $\overline{p}_{t-1} \geq p_d$, *then the agent stops experimenting (successfully).*

The principal chooses $x = 1$ if $prob\{s = 1 \mid h_t\} \geq p_d$ and $x = 0$ otherwise.

As we have a prior $\overline{p}_0 = \frac{1}{2}$, the experimentation plan implies that the agent runs exactly one experiment if the costs are sufficiently low, i.e., $c \leq \frac{1}{2p_d}U$. If the agent experiments, he chooses the precision such that after a positive result the posterior is equal to the threshold of doubt, which means that he can persuade the principal, and that after an adverse result it is revealed that $s = 0$.

The intuition for the precision of the single experiment relies on the agent's incentive to maximize the persuasion probability. To maximize this probability, he does not want to get an adverse result when the state is good and chooses $p_1^1 = 1$. He also wants

[12] Kamenica and Gentzkow (2011) show that a signal which induces posteriors of 0 or p_d yields the highest probability to persuade that is consistent with Bayes' law. Note that in our setting experimentation is costly and we explicitly consider multiple experiments.

to get a positive result when the state is bad, but he has to choose p_1^0 such that he can still persuade the principal. Therefore, he maximizes $prob\{y_1 = 1 \mid s = 0\}$ under the constraint that $prob\{s = 1 \mid y_1 = 1\} \geq p_d$. This leads to $p_1^0 = \frac{p_d - \overline{p}_0}{p_d(1-\overline{p}_0)}$. Since the prior is $\frac{1}{2}$ the precision of the first experiment is $p_1^1 = 1$, $p_1^0 = \frac{2p_d - 1}{p_d}$. The optimal experimentation strategy for the agent implies that in equilibrium $prob\{s = 1 \mid x = 1\} = p_d$ for all c for which the agent is willing to experiment. Note that the agent stops after the first experiment regardless of the outcome and that the optimal precision if he starts experimenting does not depend on the experimentation costs.

Let us now study the players' utilities under public experimentation. Since the agent chooses the precision of the experiment such that he can just persuade the principal, $prob\{s = 1 \mid x = 1\} = p_d$, the principal's excess utility (as described in (1)) is zero. Therefore, the principal's utility is $EU^p_{pub} = \frac{1}{2}p_d + \frac{1}{2}$. Thus, the principal does not benefit from public experimentation compared to the case without information provision. If the agent chooses the optimal experimentation strategy, he experiments at most once and his utility is $EU^a_{pub} = \frac{1}{2p_d}U - c$ if $c \leq \frac{1}{2p_d}U$, and $EU^a_{pub} = 0$ if $c > \frac{1}{2p_d}U$.

4.2 Private experimentation

Under private experimentation, neither the number of conducted experiments nor their outcomes can be observed by the principal. The agent can reveal any subset of the results to the principal. Denote the reported results by m. A revealed outcome i is characterized by $(\hat{y}_i, \hat{p}_i^1, \hat{p}_i^0)$, where, in order to improve the exposition, we drop the index i if the agent reveals a single outcome. An equilibrium with information provision under private experimentation has to satisfy two conditions. First, the agent's participation constraint $EU^a_{priv} \geq 0$ has to hold. Second, the principal is persuaded and chooses $x = 1$ if $prob\{s = 1 \mid m\} \geq p_d$, otherwise she chooses $x = 0$. In contrast to public experimentation, $prob\{s = 1 \mid m\}$ depends on c under private experimentation: For a given set of messages that are able to persuade the principal, c determines how long the agent is willing to search after an adverse experimentation history. This implies that whether there exists an equilibrium in which an announcement m can indeed persuade the principal depends on the agent's experimentation costs. In general, an m

that can persuade in equilibrium and c have to be such that the agent sometimes stops experimenting unsuccessfully. Otherwise providing m is uninformative with respect to the state of the world.

Let us check whether there is an equilibrium under private experimentation in which the agent can persuade the principal with one positive outcome which has the same precision as under public experimentation. If the first result is adverse, the agent knows that $s = 0$. It is worthwhile to continue experimenting with the same precision, given that the principal can be persuaded with one positive result of this precision, if $c < (1 - p_1^0)U = \frac{1-p_d}{p_d}U$. Thus, in this case upon the provision of one result in favor of the agent, the probability that $s = 1$ must be smaller than the threshold of doubt p_d, since the provided positive result can stem from many experiments. Consequently this cannot be an equilibrium if c is sufficiently low.

In general, there are multiple equilibria for a given c under private experimentation. Each equilibrium is characterized by the set of messages M^* which can persuade the principal. This set results from what the principal believes about the experiments which the agent does not reveal to her. In equilibrium, $prob\{s = 1 \mid m^*\} \geq p_d$ has to be satisfied, where $m^* \in M^*$ denotes a message that persuades the principal in equilibrium.

There can be equilibria in which the principal's posterior that $s = 1$ is strictly greater than her threshold of doubt. If the principal's decision rule is such that only evidence of a high quality induces a favorable decision for the agent, then the agent's experimentation plan is such that he continues searching for such evidence, unless his posterior that $s = 1$ is too low to justify further experimentation.[13] As the agent starts experimenting at the prior $prob\{s = 1\} = 1/2$ in an equilibrium with persuasion, optimal experimentation implies that he only stops searching unsuccessfully given history h_{t-1} if the corresponding posterior $\overline{p}_{t-1}$ is below the prior. Hence, any report by the agent that does not contain persuasive evidence (according to the principal's decision rule) should make the principal think that the experimentation history underlying the report is such that the agent has a corresponding posterior that is below the prior.

[13] We assume that the principal's beliefs are such that the agent can always persuade the principal if he provides at least one positive result with precision $\hat{p}_i^0 = 1$. This positive result implies that $prob\{s = 1 \mid m\} = 1$.

Unsuccessful search should therefore make the principal believe that the probability that $s = 1$ is below 1/2. As the principal's threshold of doubt p_d is greater than 1/2, her decision rule is optimal given these beliefs. In contrast to public experimentation, the principal can thus credibly commit to choose against the agent if the latter fails to provide favorable evidence of a quality that strictly exceeds her threshold of doubt. Thus, her decision rule is sequentially rational and the agent can be forced to provide high quality evidence.[14] This argument in general holds for diverse sets of persuasive messages that can differ with respect to the number and precision of outcomes that are required for persuasion, i.e., there can be multiple equilibria for given parameters.

The following proposition describes the highest $prob\{s = 1 \mid x = 1\}$ that can be obtained in equilibrium, which we need for the comparison of the schemes.

Proposition 2 *Under private experimentation, the maximum probability across all equilibria that $s = 1$ given that the principal chooses $x = 1$ is (i) $prob\{s = 1 \mid x = 1\} = \frac{U}{2c}$ if $\frac{1}{2}U < c \leq \frac{1}{2p_d}U$ and (ii) $prob\{s = 1 \mid x = 1\} = 1$ if $0 \leq c \leq \frac{1}{2}U$.*

Note that for all parameters, for which an equilibrium with information provision exists under public experimentation ($\frac{1}{2p_d}U \geq c$), there also exists at least one equilibrium under the private scheme.

Remark 1 *An equilibrium with information provision exists under public and under private experimentation if $0 \leq c \leq \frac{1}{2p_d}U$.*

Let us now study the players' utilities under private experimentation. In an equilibrium with $prob\{s = 1 \mid x = 1\} = p_d$ the principal's utility is $EU^p_{pri} = \frac{1}{2}p_d + \frac{1}{2}$, since her excess utility (as described in (1)) is zero. She is as well off as in the case without information provision, in which she always optimally chooses $x = 0$. In all equilibria with $prob\{s = 1 \mid x = 1\} > p_d$ the principal's utility is higher than in

[14] A special case is costless experimentation. If $c = 0$, there are only equilibria in which the m^*, which can persuade the principal, include a positive result with $\hat{p}_i^0 = 1$: Assume, in contrast, that m^* does only include positive results with $\hat{p}_i^0 < 1$. The agent can provide this m^* with probability 1, since he never stops experimenting unsuccessfully if experimentation is costless. Thus, there cannot be an equilibrium with this m^*, since it violates the condition that $prob\{s = 1 \mid m^*\} \geq p_d$ in equilibrium. Therefore, in equilibrium $prob\{s = 1 \mid m^*\} = 1$ if $c = 0$.

the case without information provision.[15] In these equilibria the principal's posterior belief that $s = 1$ when she is persuaded is higher than her threshold of doubt. The agent's utility in equilibrium is $EU^a_{priv} \geq 0$, since otherwise the agent would not start experimentation. Furthermore, the agent can always choose to run one experiment with perfect precision, which can persuade the principal if the result is positive. Thus, the agent's utility is $EU^a_{priv} \geq \max\{\frac{1}{2}U - c, 0\}$.

4.3 Payoff comparison of public and private experimentation

A welfare comparison of the schemes depends on the weights attached to the players' utilities. Instead of assuming arbitrary weights, this section analyzes separately the principal's and the agent's utilities under both experimentation schemes. The welfare comparison for some given weights is then straightforward. This section then shows that under the private scheme equilibria with $prob\{s = 1 \mid x = 1\} > p_d$, which are particularly relevant for the comparison, can Pareto-dominate equilibria with $prob\{s = 1 \mid x = 1\} = p_d$.

As we have argued above, in an equilibrium under public experimentation $EU^p_{pub} = \frac{1}{2} + \frac{1}{2}p_d$ and $EU^a_{pub} = max\left\{\frac{1}{2p_d}U - c, 0\right\}$. Furthermore, in an equilibrium under private experimentation $EU^p_{priv} \geq \frac{1}{2} + \frac{1}{2}p_d$ and $max\{\frac{1}{2}U - c, 0\} \leq EU^a_{priv} \leq \frac{1}{2p_d}U - c$. The following proposition compares the players' utilities under the two schemes.

Proposition 3 *In all equilibria under the private scheme with $prob\{s = 1 \mid x = 1\} > p_d$ the principal strictly prefers private experimentation, whereas the agent strictly prefers the public scheme. In the equilibria under the private scheme with $prob\{s = 1 \mid x = 1\} = p_d$ the principal is indifferent between the schemes and the agent is weakly better off under public experimentation.*

We have seen that multiple equilibria may exist under private experimentation and in many of them $prob\{s = 1 \mid x = 1\} > p_d$. In these equilibria, the principal strictly prefers private experimentation to the public scheme. However, the question arises whether equilibria with $prob\{s = 1 \mid x = 1\} > p_d$ are plausible. We argue that some

[15]The excess utility is positive in this case, since $prob\{x = 1\}$ is also positive as soon as the agent starts experimenting.

equilibria with $prob\{s = 1 \mid x = 1\} > p_d$ are plausible in the sense that they can Pareto-dominate equilibria with $prob\{s = 1 \mid x = 1\} = p_d$.[16]

Proposition 4 *Equilibria with $prob\{s = 1 \mid x = 1\} = p_d$ are Pareto-dominated by an equilibrium in which $prob\{s = 1 \mid x = 1\} > p_d$ if $\underline{c} < c < \frac{1-p_d}{p_d}U$, with $\underline{c} = max\{\frac{1-p_d}{1+p_d}U, \frac{1-3p_d+\sqrt{p_d^2+6p_d-3}}{4p_d-2}U\}$.*

The idea behind the proof is as follows. In the above utility analysis we have seen that the principal always prefers equilibria with $prob\{s = 1 \mid x = 1\} > p_d$. Thus, to show Pareto dominance, we have to analyze the conditions under which the agent also prefers such an equilibrium.

An equilibrium with $prob\{s = 1 \mid x = 1\} > p_d$ can only Pareto-dominate all equilibria with $prob\{s = 1 \mid x = 1\} = p_d$ if the equilibrium with exactly one positive revealed result with $\hat{p}^1 = 1$ and $\hat{p}^0 = \frac{2p_d-1}{p_d}$ does not exist under the private scheme.[17] This equilibrium does not exist if $c < \frac{1-p_d}{p_d}U$. In the appendix we distinguish three types of equilibria with $prob\{s = 1 \mid x = 1\} = p_d$. A cost range can be identified such that there is an equilibrium with $prob\{s = 1 \mid x = 1\} > p_d$ that dominates all of these three types of equilibria.

First, there may be equilibria with $prob\{s = 1 \mid x = 1\} = p_d$ where the precision of the initial experiment is such that the agent can persuade the principal with one positive result and where the agent stops after the first experiment. In these equilibria it must be that $\hat{p}^1 < 1$ and, therefore, the probability to persuade is relatively low. We show that these equilibria are worse for the agent than some equilibria with $prob\{s = 1 \mid x = 1\} > p_d$ in which $\hat{p}^1 = 1$ and, therefore, the probability to persuade the principal is relatively high.

Second, there may be equilibria with $prob\{s = 1 \mid x = 1\} = p_d$ where the precision of the initial experiment is such that the agent can persuade the principal with one positive result and where the agent does not stop if the first outcome is adverse. Since in these equilibria the agent searches more than once in expected terms, they are

[16] If some equilibria with $prob\{s = 1 \mid x = 1\} > p_d$ Pareto-dominate equilibria with $prob\{s = 1 \mid x = 1\} = p_d$ it follows that an equilibrium with $prob\{s = 1 \mid x = 1\} > p_d$ is also a sender preferred equilibrium, which has received attention in the literature on persuasion (e.g., Kamenica and Gentzkow, 2011).

[17] The precision corresponds to the precision in the equilibrium under public experimentation.

Pareto-dominated by an equilibrium with $prob\{s = 1 \mid x = 1\} > p_d$ in which the agent searches only once if c is high, i.e., if $c > \frac{1-3p_d+\sqrt{p_d^2+6p_d-3}}{4p_d-2}U$.

Third, there may be equilibria with $prob\{s = 1 \mid x = 1\} = p_d$ where the precision of the initial experiment is such that the agent needs at least two outcomes to persuade the principal. As again the expected experimentation costs are relatively high, these equilibria are Pareto-dominated by an equilibrium with $prob\{s = 1 \mid x = 1\} > p_d$ in which the agent searches only once if $c > \frac{1-p_d}{1+p_d}U$.

Corollary 1 *Consider equilibria that are not Pareto-dominated under private experimentation and let $\underline{c} < c < \frac{1-p_d}{p_d}U$. The principal is strictly better off and the agent is strictly worse off under private experimentation than under public experimentation.*

For this parameter range equilibria with $prob\{s = 1 \mid x = 1\} = p_d$ are Pareto-dominated under the private scheme. In all equilibria that are not Pareto-dominated it must be that $prob\{s = 1 \mid x = 1\} > p_d$. Hence, regardless on how players coordinate, the strict preferences result as long as they do not coordinate on a Pareto-dominated equilibrium.

4.4 Comparison of the persuasion probability

This section shows that private experimentation limits the extent to which persuasion is possible compared to the public scheme.

Under sequential public experimentation the agent in equilibrium runs a single experiment with $p_1^1 = 1$ and $p_1^0 = \frac{2p_d-1}{p_d}$. The resulting persuasion probability is $\frac{1}{2p_d}$. As the experiment has the same properties as in Kamenica and Gentzkow (2011), the persuasion probability is the same.

In the private case the persuasion probability can be lower. Consider first $c = 0$. In equilibrium it must be that $prob\{s = 1 \mid x = 1\} = 1$ and $prob\{s = 0 \mid x = 0\} = 1$.[18] Without loss of generality, consider an equilibrium where the agent runs

[18] If there were an equilibrium with $prob\{s = 1 \mid x = 1\} < 1$, then, as experimentation is costless and persuasive evidence can be found by chance even if $s = 0$, the agent would run as many experiments as it takes to find persuasive evidence. He would, thus, find persuasive evidence almost with certainty regardless of the state. But then this evidence is worthless for the principal.

If there were an equilibrium with $prob\{s = 1 \mid x = 1\} = 1$ and $prob\{s = 0 \mid x = 0\} < 1$, then the persuasion probability would be smaller then $1/2$. The agent could improve by running a single experiment which perfectly mirrors the state and yields persuasion probability $1/2$.

a single experiment that perfectly mirrors the state of the world. In this case, the persuasion probability is equal to 1/2 and, hence, strictly lower than under public experimentation.

If $c > 0$, then there can be multiple equilibria that differ with respect to $prob\{s = 1 \mid x = 1\}$. A necessary condition for the persuasion probability to be maximal, however, is that the principal is indifferent between choosing in favor or against the agent upon the observation of persuasive evidence, i.e., $prob\{s = 1 \mid x = 1\} = p_d$, and, thus, the persuasion probability in all equilibria with $prob\{s = 1 \mid x = 1\} > p_d$ is lower than under public experimentation. Following Kamenica and Gentzkow (2011), we focus on sender preferred equilibria. As we have seen in Proposition 4, there are parameter constellations, where equilibria with $prob\{s = 1 \mid x = 1\} = p_d$ are Pareto-dominated by an equilibrium in which a single experiment is run which yields $prob\{s = 1 \mid x = 1\} > p_d$. This implies that for these parameter constellations there is an equilibrium with $prob\{s = 1 \mid x = 1\} > p_d$ that the sender prefers to any equilibrium with $prob\{s = 1 \mid x = 1\} = p_d$. Therefore, the maximum persuasion probability in sender preferred equilibria under the private scheme can be strictly lower than under public experimentation.

5 Extensions

5.1 Deterrence effect

In practice it may not be possible for the agent to choose the precision of an experiment or it may be too expensive to run experiments with high precision. As an approximation we modify our model in this section and assume that the precision p of an experiment is exogenous and symmetric ($p_t^1 = p$, $p_t^0 = p$). This change has noteworthy implications. We are now going to show that there can be a "deterrence effect", which renders private experimentation socially optimal.

Let us start by characterizing the equilibria under public and private experimentation. There is a unique equilibrium under public experimentation. The principal is persuaded to choose in favor of the agent if $prob\{s = 1 \mid h_t\} \geq p_d$. The agent stops at the first t for which $prob\{s = 1 \mid h_t\} \geq p_d$ or if the probability of persuasion becomes too low. However, in contrast to before, the agent cannot adjust the precision. Now, as the precision is exogenous, he needs several positive results to persuade the principal after an initial adverse outcome.

Under private experimentation, there can be multiple equilibria.[19] Any announcement m can be summarized by the number of announced unfavorable m^0 and favorable results m^1, since the precision of each experiment is fixed. We can write $m = (m^0, m^1)$. Each equilibrium where the agent starts experimenting has to satisfy the same two conditions as in the model with an endogenous precision; the agent's participation constraint and that $prob\{s = 1 \mid x = 1\} \geq p_d$. Following Felgenhauer and Schulte (2013) we restrict attention to equilibria in which the principal on the equilibrium path only chooses in favor of the agent if $m^0 = 0$ and $m^1 > 0$. Felgenhauer and Schulte (2013) show that there does not exist an equilibrium in which the agent can persuade the principal if $c = 0$.[20] For $c > 0$, equilibria with $m^1 \in [\underline{m}, \overline{m}]$ can exist, where $\underline{m}$ is the smallest integer such that $prob\{s = 1 \mid x = 1\} \geq p_d$ and $\overline{m}$ is the largest integer such that the agent's participation constraint holds.[21]

We now argue that public experimentation can be worse than private experimentation from a social perspective if the precision p is low and the costs are relatively high. The superiority of private experimentation here stems from the fact that the agent's participation constraint is not satisfied under public experimentation. He is deterred from starting to experiment under public experimentation, whereas he is willing to start under the private scheme, as he may hide an unfavorable result. Since there is no information provision for these parameters under the public scheme, both players are better off under private experimentation.

Proposition 5 *Consider $p < p_d < \frac{p^2+p^2(1-p)}{p^2+p^2(1-p)+(1-p)^2+p(1-p)^2}$. There are $\underline{c}$ and $\overline{c}$ such that for all $c \in [\underline{c}, \overline{c}]$ private experimentation is better than public experimentation for both, the principal and the agent.*

[19] Felgenhauer and Schulte (2013) analyze the setup with private experimentation and an exogenous precision in detail.

[20] The agent could find a given m^1 with probability 1 regardless of the state in a hypothetical equilibrium with persuasion, yielding a contradiction.

[21] Intuitively, for a given c, the equilibrium value of $\underline{m}$ must be sufficiently high such that the agent at least sometimes stops searching unsuccessfully. Only then the principal thinks that $prob\{s = 1 \mid m^1\} > \frac{1}{2}$ upon the provision of m^1. $\underline{m}$ is the lowest number of favorable outcomes such that the principal's threshold of doubt is passed.

We are now going to show that the deterrence effect only plays a role if $p < p_d$. Suppose that $p \geq p_d$. If $c > \frac{1}{2}U$, then the agent does not engage in experimentation under either scheme, as the probability of $\frac{1}{2}$ to obtain a favorable decision does not justify the costs of a single experiment. At $c = \frac{1}{2}U$, the agent is indifferent to experiment once and without loss of generality we suppose he does. Note that this threshold for the costs is the same under both schemes. If $p \geq p_d$, a single experiment can persuade the principal under the public scheme.

Remark 2 *There is no deterrence effect if* $p \geq p_d$.

Note that there is no deterrence effect if the precision of an experiment is endogenous. In this case, equilibria in which the agent starts experimenting exist under both schemes if $c \leq \frac{1}{2p_d}U$.

5.2 Modified public experimentation versus private experimentation

Even if the principal implements a public scheme in the sense that she can only be persuaded with evidence obtained from public experiments, the agent nevertheless may have the option to run experiments privately. Let us return to the setting in which the agent can choose p_t^1 and p_t^0 as in section 4 and investigate a modified public scheme. Under this scheme, the agent can run private as well as public experiments, but he can exclusively persuade the principal with evidence obtained from public experiments. Analogous to our previous analysis of private experimentation, there are equilibria under the modified public scheme in which the agent has to conduct high quality public experiments in order to be able to persuade the principal. Such an equilibrium can be supported with similar beliefs as in the private scheme: If the principal off the equilibrium path observes low quality public experiments, then she should think that the agent also privately collected sufficiently adverse evidence such that it does not pay to run additional public experiments with the required quality. The corresponding beliefs induce a decision against the agent and, therefore, destroy the agent's incentive to run low quality public experiments.

The strategic effects under the modified public scheme are similar to the effects in the setup with pure private experimentation. Both schemes appear to be equally valuable for the principal. However, there are differences with respect to the plausibility of equilibria with $prob\{s = 1|x = 1\} > p_d$. To illustrate this point consider again sender preferred equilibria as in Kamenica and Gentzkow (2011). As we have seen in Proposition 4, the agent prefers an equilibrium with $prob\{s = 1|x = 1\} > p_d$ under private experimentation if $\underline{c} < c < \frac{1-p_d}{p_d}U$. The best equilibrium for the agent under the modified public scheme is where he runs exactly one public experiment with the same quality as under pure public experimentation, which leads to $prob\{s = 1|x = 1\} = p_d$. This equilibrium has the highest probability to persuade, and the agent incurs the cost of running only one experiment. This equilibrium exists for $c \leq \frac{1}{2p_d}U$ under the modified public scheme, i.e., it exists whenever the sender preferred equilibrium under the pure private scheme implies $prob\{s = 1|x = 1\} > p_d$. It follows that, if we use sender preferred equilibrium as the selection criterion, the principal prefers the pure private scheme to modified public experimentation and the agent prefers modified public experimentation to the pure private scheme.

6 Conclusion

There can be severe legal repercussions whenever it is uncovered that an interested party, like a pharmaceutical company, hides unfavorable decision-relevant evidence. In principle, whistle-blower policies and punishments can deter private experimentation. This paper argues that there are reasons for the public to favor little transparency even though each piece of hidden adverse evidence clearly is detrimental.[22] An advantage of a lack of transparency is that the public's representatives can be sceptical. Under public experimentation instead it is not possible to commit not to be persuaded once the threshold of doubt is reached. If the representatives are sufficiently sceptical under the private scheme, then the interested party prefers more transparency. Furthermore, if experimentation costs are high and the quality of an experiment is exogenous and low, there can be a deterrence effect that renders public experimentation inferior for

[22] Private experimentation becomes even more attractive if the detection probabilities cannot be implemented at zero costs.

all parties.

In the literature on persuasion with experimentation it is currently not common to consider additional monetary transfers. In future work it would be interesting to extend this literature by viewing these persuasion situations as contractual problems in which the principal can condition transfers on the signal quality. In our paper in contrast the implementation of a scheme implies a distribution of rents via probabilities. The introduction of transfers allows a redistribution of rents within a scheme. A problem though is that the quality of a test may be observable by experts, but it may not be easily verified by court rendering quality contingent transfers hard to implement.

Appendix

Derivation of the principal's utility

Now we derive the principal's utility $EU^p_{priv} = \frac{1}{2} + \frac{1}{2}p_d + (prob\{s = 1 \mid x = 1\} - p_d)prob\{x = 1\}$. Let M^* be the set of messages with which the principal can be persuaded in equilibrium and which are reached with a positive probability on the equilibrium path. Let m^* be an element of this set.

$$
\begin{aligned}
EU^p_{priv} &= \sum_{m^* \in M^*} (prob\{s = 1 \mid m^*\} + (1 - p_d)(1 - prob\{s = 1 \mid m^*\}))prob\{m^*\} \\
&+ (p_d prob\{s = 1 \mid x = 0\} + prob\{s = 0 \mid x = 0\})prob\{x = 0\} \\
&= \sum_{m^* \in M^*} (prob\{s = 1 \mid m^*\} + (1 - p_d)(1 - prob\{s = 1 \mid m^*\}))prob\{m^*\} \\
&+ \left(p_d \frac{prob\{x = 0 \mid s = 1\}prob\{s = 1\}}{prob\{x = 0\}}\right. \\
&+ \left.\frac{prob\{x = 0 \mid s = 0\}prob\{s = 0\}}{prob\{x = 0\}}\right) prob\{x = 0\} \\
&= \sum_{m^* \in M^*} (prob\{s = 1 \mid m^*\} + (1 - p_d)(1 - prob\{s = 1 \mid m^*\}))prob\{m^*\} \\
&+ \frac{1}{2} p_d prob\{x = 0 \mid s = 1\} + \frac{1}{2} prob\{x = 0 \mid s = 0\} \\
&= \sum_{m^* \in M^*} (prob\{s = 1 \mid m^*\} + (1 - p_d)(1 - prob\{s = 1 \mid m^*\}))prob\{m^*\} \\
&+ \frac{1}{2} p_d (1 - prob\{x = 1 \mid s = 1\}) + \frac{1}{2}(1 - prob\{x = 1 \mid s = 0\}) \\
&= \sum_{m^* \in M^*} (prob\{s = 1 \mid m^*\} + (1 - p_d)(1 - prob\{s = 1 \mid m^*\}))prob\{m^*\} \\
&+ \frac{1}{2} p_d \left(1 - \frac{\sum_{m^* \in M^*} prob\{s = 1 \mid m^*\}prob\{m^*\}}{prob\{s = 1\}}\right) \\
&+ \frac{1}{2} \left(1 - \frac{\sum_{m^* \in M^*} prob\{s = 0 \mid m^*\}prob\{m^*\}}{prob\{s = 0\}}\right) \\
&= \sum_{m^* \in M^*} (prob\{s = 1 \mid m^*\} + (1 - p_d)(1 - prob\{s = 1 \mid m^*\}))prob\{m^*\} \\
&+ \frac{1}{2} p_d (1 - 2 \sum_{m^* \in M^*} prob\{s = 1 \mid m^*\}prob\{m^*\}) \\
&+ \frac{1}{2}(1 - 2 \sum_{m^* \in M^*} (1 - prob\{s = 1 \mid m^*\})prob\{m^*\})
\end{aligned}
$$

Collecting terms yields

$$EU^p_{priv} = \frac{1}{2} + \frac{1}{2}p_d + \sum_{m^* \in M^*} (prob\{s = 1 \mid m^*\} - p_d) prob\{m^*\} \quad .$$

Observing that

$$\sum_{m^* \in M^*} p_d prob\{m^*\} = p_d prob\{x = 1\}$$

and that

$$\begin{aligned}
\sum_{m^* \in M^*} prob\{s = 1 \mid m^*\} prob\{m^*\} &= \sum_{m^* \in M^*} \frac{prob\{s = 1, m^*\}}{prob\{m^*\}} prob\{m^*\} \\
&= \sum_{m^* \in M^*} prob\{s = 1, m^*\} \\
&= prob\{s = 1, x = 1\} \\
&= prob\{s = 1 \mid x = 1\} prob\{x = 1\}
\end{aligned}$$

immediately yields $EU^p_{priv} = \frac{1}{2} + \frac{1}{2}p_d + (prob\{s = 1 \mid x = 1\} - p_d) prob\{x = 1\}$.

The principal's utility under public experimentation is

$$\begin{aligned}
EU^p_{pub} = & (prob\{s = 1 \mid x = 1\} + (1 - p_d)(1 - prob\{s = 1 \mid x = 1\})) prob\{x = 1\} \\
& + (p_d prob\{s = 1 \mid x = 0\} + prob\{s = 0 \mid x = 0\}) prob\{x = 0\}
\end{aligned}$$

Analogous steps as above yield $EU^p_{pub} = \frac{1}{2} + \frac{1}{2}p_d + (prob\{s = 1 \mid x = 1\} - p_d) prob\{x = 1\}$. Q.E.D.

Proof of Proposition 1

The following probabilities depend on h_t. In order to simplify the exposition, we drop h_t from the formulas for the moment and make the dependence explicit later. We first maximize $prob\{x = 1\}$ subject to $prob\{s = 1 | x = 1\} \in [p_d, 1]$ and $prob\{s = 0 | x = 0\} \in [0, 1]$ and then argue that the precision of a single experiment can be chosen such that the optimal values of $prob\{s = 1 | x = 1\}$ and $prob\{s = 0 | x = 0\}$ can be obtained by an appropriate choice of p_t^0 and p_t^1.

$$\begin{aligned}
prob\{x=1)\} &= prob\{x=1,s=1\}+prob\{x=1,s=0\}\\
&= prob\{s=1\mid x=1\}prob\{x=1\}+prob\{x=1\mid s=0\}prob\{s=0\}\\
&= prob\{s=1\mid x=1\}prob\{x=1\}+(1-prob\{x=0\mid s=0\})prob\{s=0\}\\
&= prob\{s=1\mid x=1\}prob\{x=1\}\\
&+prob\{s=0\}-prob\{s=0\mid x=0\}(1-prob\{x=1\})
\end{aligned}$$

This is equivalent to

$$prob\{x=1\}=\frac{prob\{s=0\mid x=0\}-prob\{s=0\}}{prob\{s=1\mid x=1\}+prob\{s=0\mid x=0\}-1}.$$

The principal chooses $x=1$ if $prob\{s=1\mid x=1\}\in[p_d,1]$ and $x=0$ otherwise. To maximize $prob\{x=1\}$ the agent should set $prob\{s=1\mid x=1\}=p_d$, since $prob\{x=1\}$ is decreasing in $prob\{s=1\mid x=1\}$.

The derivative of $prob\{x=1\}$ with respect to $prob\{s=0\mid x=0\}$ is

$$\frac{\partial prob\{x=1\}}{\partial prob\{s=0\mid x=0\}}>0 \textit{ if} prob\{s=1\mid x=1\}+prob\{s=0\}>1$$

$$\frac{\partial prob\{x=1\}}{\partial prob\{s=0\mid x=0\}}<0 \textit{ if} prob\{s=1\mid x=1\}+prob\{s=0\}<1.$$

Since $prob\{s=1\mid x=1\}\geq p_d$ and $prob\{s=0\}>1-p_d$ (otherwise the agent does not experiment since the principal is already persuaded), it holds that $prob\{s=1\mid x=1\}+prob\{s=0\}>1$. Thus, $prob\{x=1\}$ is maximized if $prob\{s=0\mid x=0\}=1$ or equivalently if $prob\{s=1\mid x=0\}=0$. Consequently,

$$prob\{x=1\mid h_t\}=\frac{1-prob\{s=0\mid h_t\}}{p_d}=\frac{\overline{p}_t}{p_d}.$$

All possible experimentation strategies with the two characteristics that $prob\{s=0\mid x=0\}=1$ and $prob\{s=1\mid x=1\}=p_d$ lead to the maximum probability to persuade the principal. However, the optimal strategy requires only one experiment. Since experimentation is costly, the agent prefers this strategy to all other strategies with the same probability to persuade. He chooses the precision $(p_t^0,\, p_t^1)$ such that after a positive outcome the posterior is equal to p_d and after an adverse outcome the posterior is 0, i.e., he chooses $p_t^1 = 1$ and $p_t^0 = \frac{p_d - \overline{p}_{t-1}}{p_d(1-\overline{p}_{t-1})}$. After this experiment he stops, because either he can already persuade the principal or he cannot persuade her, since the posterior for $s = 1$ is 0. This strategy leads to the maximum probability to persuade the principal $(\frac{\overline{p}}{p_d})$ with exactly one experiment. Note that in Proposition 1 we assume that the agent runs an experiment if he is indifferent.

Given history h_t with $\overline{p}_{t-1} < p_d$, the agent runs exactly one experiment with the above precision iff $prob\{x = 1 \mid h_t\}U \geq c$, i.e., $\frac{\overline{p}_{t-1}}{p_d}U \geq c$. This is equivalent to $\frac{cp_d}{U} \leq \overline{p}_{t-1}$. Otherwise, the agent does not experiment. As the prior is $\frac{1}{2}$, the agent starts experimenting iff $c \leq \frac{1}{2p_d}U$. Q.E.D.

Proof of Proposition 2

If $\frac{1}{2}U < c \leq \frac{1}{2p_d}U$, the maximum $prob\{s = 1 \mid x = 1\}$ in an equilibrium under private experimentation is determined by the agent's participation constraint. The agent is willing to experiment once if $prob\{x = 1\}U - c \geq 0$. If the agent can persuade the principal with one positive result, the condition is $(\frac{1}{2}\hat{p}^1 + \frac{1}{2}(1 - \hat{p}^0))U \geq c$. If $\hat{p}^1 = 1$ and $\hat{p}^0$ is such that $(\frac{1}{2} + \frac{1}{2}(1-\hat{p}^0))U = c$, which is $\hat{p}^0 = \frac{2(U-c)}{U}$, this leads to a maximum $prob\{s = 1 \mid x = 1\}$ of $\frac{U}{2c}$. Since $c > \frac{1}{2}U$, this is less than 1 and for $c = \frac{1}{2p_d}U$ this is equal to the threshold of doubt. If $c \leq \frac{1}{2}U$, the agent is willing to run one experiment with $p_1^1 = 1$, $p_1^0 = 1$. This is an equilibrium if the principal's beliefs are such that he is only persuaded by one positive result that perfectly reveals $s = 1$. Thus, the maximum $prob\{s = 1 \mid x = 1\} = 1$ for $c \leq \frac{1}{2}U$. Q.E.D.

Proof of Remark 1

We show now that for $c \leq \frac{1}{2p_d}U$ there is always an equilibrium under private experimentation in which the agent can persuade the principal with one positive result. Consider an equilibrium in which principal has beliefs such that she is persuaded by exactly one positive result with the precision $\hat{p}^1$ and $\hat{p}^0$. $\hat{p}^1 = 1$, $\hat{p}^0$ and c are such that the agent runs exactly one experiment with this precision ($c \leq (\frac{1}{2}(1-\hat{p}^0) + \frac{1}{2})U$ which is equivalent to $\frac{2(U-c)}{U} \geq \hat{p}^0$) and stops after an adverse outcome ($c \geq (1-\hat{p}^0)U$ which is equivalent to $\hat{p}^0 \geq \frac{U-c}{U}$). Furthermore, $prob\{s = 1 \mid x = 1\} \geq p_d$ which implies that $\hat{p}^0 \geq \frac{2p_d-1}{p_d}$. This leads to $\frac{2(U-c)}{U} \geq \hat{p}^0 \geq max\left\{\frac{U-c}{U}, \frac{2p_d-1}{p_d}\right\}$. In equilibrium, the agent runs one experiment with this precision and the principal chooses $x = 1$ if the agent announces a positive result. Note that $prob\{s = 1 \mid x = 1\} > p_d$, except for $\hat{p}^0 = \frac{2p_d-1}{p_d}$. Q.E.D.

Proof of Proposition 4

We know that an equilibrium under private experimentation in which the agent experiments once with the same precision as under public experimentation ($p_1^1 = 1$, $p_1^0 = \frac{2p_d-1}{p_d}$) does not exist if $c < \frac{1-p_d}{p_d}U$. In this case, equilibria with $prob\{s = 1 \mid x = 1\} = p_d$ can be Pareto-dominated by an equilibrium with $prob\{s = 1 \mid x = 1\} > p_d$.

To show this we distinguish three types of equilibria with $prob\{s = 1 \mid x = 1\} = p_d$ by separating them according to the precision of the initial experiment. We compare these equilibria to an equilibrium with $prob\{s = 1 \mid x = 1\} > p_d$ and show that this equilibrium can Pareto-dominate the equilibria with $prob\{s = 1 \mid x = 1\} = p_d$.

We focus on equilibria with $prob\{s = 1 \mid x = 1\} > p_d$ in which the principal is persuaded by one positive outcome with the precision $\hat{p}^1 = 1$ and $\hat{p}^0 > \frac{2p_d-1}{p_d}$ such that the agent stops if the first result is adverse. This is the case if $(1-\hat{p}^0)U \leq c$, which is equivalent to $\hat{p}^0 \geq \frac{U-c}{U}$. In the equilibrium with $\hat{p}^1 = 1$ and $\hat{p}^0 = \frac{U-c}{U}$, the agent can persuade the principal with probability $0.5(2 - \frac{U-c}{U})$. His expected utility is $EU^a_{priv} = 0.5(2 - \frac{U-c}{U})U - c$, which is $0.5(U - c)$.

Now we compare this utility to the utility in equilibria with $prob\{s = 1 \mid x = 1\} = p_d$. In an equilibrium with $prob\{s = 1 \mid x = 1\} = p_d$, all m^* which can persuade the principal and are revealed with positive probability have to lead to $prob\{s = 1 \mid m^*\} = p_d$. This implies that also the first experiment has to have a precision such that the result may be part of an m^* which leads to $prob\{s = 1 \mid m^*\} = p_d$. Now we distinguish three types of initial experiments.

First, the precision of the initial experiment can be part of an m^* which consists of one positive result and $prob\{s = 1 \mid m^*\} = p_d$. Furthermore, the agent stops after one experiment. A posterior of p_d is reached by one experiment with positive result if

$p_d = \frac{0.5\hat{p}^1}{0.5\hat{p}^1+0.5(1-\hat{p}^0)}$, which is equivalent to $\hat{p}^0 = \frac{p_d-\hat{p}^1(1-p_d)}{p_d}$. The probability to find a positive outcome in the first experiment with this precision is $\frac{1}{2}\hat{p}^1 + \frac{1}{2}(1-\hat{p}^0)$, which is $\frac{\hat{p}^1}{2p_d}$. If the result of the first experiment is adverse, the posterior is $\frac{0.5(1-\hat{p}^1)}{0.5(1-\hat{p}^1)+0.5\hat{p}^0}$, which is $\frac{(1-\hat{p}^1)p_d}{2p_d-\hat{p}^1}$. A necessary condition for the agent to stop experimenting is that he does not want to run a second experiment with the same precision. If he runs such a second experiment, the probability to find a positive outcome in the second experiment is $\frac{(1-\hat{p}^1)p_d}{2p_d-\hat{p}^1}\hat{p}^1 + (1-\frac{(1-\hat{p}^1)p_d}{2p_d-\hat{p}^1})(1-\hat{p}^0)$, which is $\frac{\hat{p}^1(\hat{p}^1-p_d-2\hat{p}^1p_d+2\hat{p}^1p_d^2)}{(\hat{p}^1-2p_d)p_d}$. Since we assume that after an adverse result the agent does not run a second experiment with the same precision, it has to hold that $\frac{\hat{p}^1(\hat{p}^1-p_d-2\hat{p}^1p_d+2\hat{p}^1p_d^2)}{(\hat{p}^1-2p_d)p_d}U < c$. Now we analyze how the probability to find a positive result in the second experiment depends on $\hat{p}^1$. The first derivative of the probability is $\frac{2p_d^2+\hat{p}^1\hat{p}^1(1-2p_d+2p_d^2)-4\hat{p}^1p_d(1-2p_d+2p_d^2)}{(\hat{p}^1-2p_d)^2p_d}$. At $\hat{p}^1 = 1$ it is $-2+\frac{1}{p_d}$, which is negative, and for $\hat{p}^1 = 0$ it is $\frac{2p_d^2}{(-2p_d)^2p_d}$, which is positive. The second derivative of the probability to find a positive result in the second experiment is $\frac{4(1-2p_d)^2p_d}{(\hat{p}^1-2p_d)^3}$, which is negative. This implies that the probability to find a positive result in the second experiment is concave in $\hat{p}^1$. If $\frac{\hat{p}^1(\hat{p}^1-p_d-2\hat{p}^1p_d+2\hat{p}^1p_d^2)}{(\hat{p}^1-2p_d)p_d}U < c$, the agent does not run a second experiment. Even if an equilibrium with $prob\{s=1 \mid m^*\} = p_d$ does not exist for $\hat{p}^1 = 1$, it exists for all $\hat{p}^1$ in the interval $\hat{p}^1 \in (0, \overline{p}^1]$. $\overline{p}^1$ is such that the agent is exactly indifferent to experiment once more if $\hat{p}^1 = \overline{p}^1$, as the probability to find a positive result in the second experiment for $\hat{p}^1 = 0$ is smaller than for $\hat{p}^1 = 1$. Since the probability to find a positive outcome in the first experiment is increasing in $\hat{p}^1$, the best precision for the agent is $\hat{p}^1 = \overline{p}^1$. The upper bound $\overline{p}^1$ follows from $\frac{\overline{p}^1(\overline{p}^1-p_d-2\overline{p}^1p_d+2\overline{p}^1p_d^2)}{(\overline{p}^1-2p_d)p_d}U - c = 0$, which is equivalent to $\overline{p}^1 = \frac{p_d(c+U-\sqrt{c^2-2c(3-8p_d+8p_d^2)U+U^2})}{2(1-2p_d+2p_d^2)U}$ (we show later that the root is in fact positive). The probability to persuade in the first experiment given the precision $\overline{p}^1$ is $\frac{c+U-\sqrt{c^2-2c(3-8p_d+8p_d^2)U+U^2}}{4(1-2p_d+2p_d^2)U}$. Remember that we use a necessary condition such that the agent stops experimenting after the first experiment. Thus, the actual probability to persuade the principal can be smaller. Now we show that an equilibrium in which the agent only experiments once and $prob\{s=1 \mid x=1\} = p_d$ is worse for the agent than the equilibrium in which $prob\{s=1 \mid x=1\} > p_d$. To show this we compare the probability to persuade the principal in both equilibria. The difference between the probability to persuade with the first experiment is at least $0.5(2-\frac{-c+U}{U}) - \frac{c+U-\sqrt{c^2-2c(3-8p_d+8p_d^2)U+U^2}}{4(1-2p_d+2p_d^2)U}$. The first part is independent of p_d and the second one is increasing in p_d. Therefore, $0.5(2-\frac{-c+U}{U})$ is larger than $\frac{c+U-\sqrt{c^2-2c(3-8p_d+8p_d^2)U+U^2}}{4(1-2p_d+2p_d^2)U}$ for all p_d if it holds for $p_d=1$. For $p_d=1$ the difference is $\frac{0.5\sqrt{c^2-6cU+U^2}}{U}$. Since this is positive, the probability to persuade the agent with one positive result is higher in the equilibrium with $prob\{s=1 \mid x=1\} > p_d$ than if $prob\{s=1 \mid x=1\} = p_d$. Therefore, equilibria in which the agent can persuade the principal with one positive result out of one experiment and $prob\{s=1 \mid x=1\} = p_d$ are dominated by an equilibrium with $prob\{s=1 \mid x=1\} > p_d$.

To complete this part we have to show that $\sqrt{c^2-2c(3-8p_d+8p_d^2)U+U^2}$ is positive. It is positive if

$$c \leq 3U - 8p_dU + 8p_d^2U - 2\sqrt{2}\sqrt{U^2-6p_dU^2+14p_d^2U^2-16p_d^3U^2+8p_d^4U^2}. \tag{3}$$

Now we show that it is sufficient for inequality (3) to be true that $c < \frac{1-p_d}{p_d}U$.

$$\frac{1-p_d}{p_d} < 3-8p_d+8p_d^2-2\sqrt{2}\sqrt{1-6p_d+14p_d^2-16p_d^3+8p_d^4}$$

$$\Leftrightarrow 3-8p_d+8p_d^2-2\sqrt{2}\sqrt{1-6p_d+14p_d^2-16p_d^3+8p_d^4}-\frac{1-p_d}{p_d} > 0$$

$$\Leftrightarrow \frac{(2p_d-1)(1+4p_d^2+2p_d(-1+\sqrt{2-4p_d+4p_d^2}))}{p_d} > 0.$$

This is true if $p_d > \frac{1}{2}$. Since we analyze the case in which the equilibrium under private experimentation with the same precision as under public experimentation ($p_1^1=1$, $p_1^0=\frac{2p_d-1}{p_d}$) does not exist, the condition that $c < \frac{1-p_d}{p_d}U$ is satisfied and, therefore, the term under the root is positive.

Second, the precision of the first experiment can be part of an m^* which consist of one positive outcome and $prob\{s=1 \mid m^*\} = p_d$. However, the agent continues experimenting if the first result is adverse. The maximum probability to persuade the principal in any equilibrium is $\frac{1}{2p_d}$. Now we identify a lower bound for the agent's expected experimentation costs. Since it is worthwhile to run a second experiment if the first result is adverse, we identify the minimum probability that the agent runs a second experiment. This is at least $(1-\frac{1}{2p_d})+prob\{y_1=0, x=1\}$, where $(1-\frac{1}{2p_d})$ is a lower bound of the probability that the agent does not persuade the principal and $prob\{y_1 = 0, x = 1\}$ is the probability to persuade the principal after the first experiment. It follows that the agent's expected experimentation costs are at least $(1 + (1 - \frac{1}{2p_d}) + prob\{y_1 = 0, x = 1\})c$. Now we identify a lower bound for $prob\{y_1 = 0, x = 1\}$. If the first outcome is adverse, the agent runs a second experiment if $prob\{x = 1 \mid y_1 = 0\}U \geq c$, which is $prob\{y_1 = 0, x = 1\} \geq \frac{c}{U} prob\{y_1 = 0\}$. Furthermore, the probability to find an adverse result in the first experiment has to be at least as high as the probability not to convince the principal, $prob\{y_1 = 0\} \geq (1 - \frac{1}{2p_d})$, since the agent can persuade with one positive result. It follows that $prob\{x = 1, y_1 = 0\} \geq \frac{c}{U} prob\{y_1 = 0\} \geq \frac{c}{U}(1 - \frac{1}{2p_d})$. Thus, the agent's expected utility is $EU^a_{priv} \leq \frac{1}{2p_d}U - (1 + (1 - \frac{1}{2p_d}) + \frac{c}{U}(1 - \frac{1}{2p_d}))c$. We compare this utility with the utility in the equilibrium, in which $prob\{s = 1 \mid x = 1\} > p_d$. $\frac{1}{2p_d}U - (1 + (1 - \frac{1}{2p_d}) + \frac{c}{U}(1 - \frac{1}{2p_d}))c$ is smaller than $0.5(U - c)$ if the difference between the two utilities $(\frac{1}{2p_d}U - (1 + (1 - \frac{1}{2p_d}) + \frac{c}{U}(1 - \frac{1}{2p_d}))c - 0.5(U - c) = \frac{c^2(1-2p_d)+U^2(1-p_d)+cU(1-3p_d)}{2p_dU})$ is negative. By algebraic transformation it can be seen that this is true if $c > \frac{1-3p_d+\sqrt{p_d^2+6p_d-3}}{4p_d-2}U$. Therefore, the equilibrium with $prob\{s = 1 \mid x = 1\} > p_d$ can dominate equilibria with $prob\{s = 1 \mid x = 1\} = p_d$ in which the agent starts experimenting with a precision of which one positive outcome can persuade the principal.

Third, the first precision can be part of an m^* which persuades the principal and consist of at least two outcomes. Now we assume that one outcome of this precision cannot persuade the principal. The agent's expected utility is composed of the probability to persuade the agent and the agent's expected experimentation costs. If the agent can persuade the principal, he has to run at least two experiments, but also if he cannot persuade the principal, the agent has run at least one experiment. Thus, his expected experimentation costs are at least $(prob\{x = 1\} * 2 + (1 - prob\{x = 1\}) * 1)c$, which is $(1 + prob\{x = 1\})c$. The agent's utility $EU^a_{priv} \leq prob\{x = 1\}U - (1 + prob\{x = 1\})c$. $prob\{x = 1\}U - (1 + prob\{x = 1\})c$ is increasing in $prob\{x = 1\}$, since $U > c$. The maximum probability to persuade the principal is $\frac{1}{2p_d}$. Hence, $EU^a_{priv} \leq \frac{1}{2p_d}U - (1 + \frac{1}{2p_d})c$. If we compare this equilibrium again to the equilibrium in which one positive result is enough to persuade the principal and $prob\{s = 1 \mid x = 1\} > p_d$ it follows that $0.5(U - c) > \frac{1}{2p_d}U - (1 + \frac{1}{2p_d})c$ if $c > \frac{(1-p_d)U}{1+p_d}$.

Therefore, if $c > \underline{c} = \max\{\frac{(1-p_d)U}{1+p_d}, \frac{1-3p_d+\sqrt{p_d^2+6p_d-3}}{4p_d-2}U\}$ and $c < \frac{1-p_d}{p_d}U$, equilibria with $prob\{s=1 \mid x=1\} = p_d$ are Pareto-dominated by an equilibrium with $prob\{s = 1 \mid x = 1\} > p_d$. Note that $\underline{c} < \frac{1-p_d}{p_d}U$ if $p_d \in (0.5, 1)$. Q.E.D.

Proof of Proposition 5

Define H_s^j, with $j \in \{priv, pub\}$, as the set of histories such that for each $h_t \in H_s^j$ the agent successfully collects persuasive evidence on the equilibrium path. The set H_f^j consists of histories for which the collection of persuasive evidence fails and which may be reached on the equilibrium path. For each history in both sets it must be the case that along the way it is always optimal to continue experimenting and the agent stops after the final experiment. Since the precision p is given, we can simplify the notation in this proof and do not write explicitly the precision of each experiment. Under this simplification, a history consists of the outcomes of the experiments $h_t = (y_1, y_2, ..., y_t)$. Given the p in the proposition, the principal can be persuaded under private experimentation if the optimal experimentation plan is $H_s^{priv} = \{(1,1),(1,0,1)\}$ and $H_f^{priv} = \{(0),(1,0,0)\}$. We show later that there are parameters such that this plan is optimal. This plan implies that p_d and p are such that the principal can be persuaded if the number of 1s is at least equal to the number of 0s plus two under public experimentation.[23] Consider c such that the above experimentation plan under the private scheme is just optimal for the agent in equilibrium. Under public experimentation, there are two potential equilibrium experimentation plans: (i) $H_s^{pub} = \{(1,1),(1,0,1,1),(1,0,1,0,1,1), ...\}$, $H_f^{pub} = \{(0),(1,0,0),(1,0,1,0,0), ...\}$ and (ii) $H_s^{pub} = H_f^{pub} = \varnothing$, i.e., the agent does not start experimenting.[24] Note that in order to be successful, the agent under experimentation plan (i) has to provide more

[23]If the number of 1s and 0s under public experimentation is equal, then they cancel out in the principal's update concering the state. Two additional 1s under public experimentation yield a higher updated probability that $s = 1$ for the principal than the agent announcing two 1s under private experimentation: The latter can also occur if the agent privately collected an additional 0.

[24]Note that if the plan in (i) is consistent with an equilibrium, then the agent continues experimenting after history $(1,0)$, as the posterior after this history is equal to the prior and the agent started experimenting given the prior.

Note further that there cannot be an equilibrium where a single 1 is persuasive under the public scheme (where the posterior probability that $s = 1$ upon observing a single 1 is p), as $p_d > p$ according to the qualification in the proposition.

1s after history $(1,0)$ than under private experimentation. The probability that $x=1$ is also higher under private experimentation. Hence, the agent's benefit from experimentation in the private case is higher than under public experimentation. Thus, if the costs are such that the above plan under private experimentation is as good as not starting to experiment under private experimentation, then for these costs it is strictly optimal for the agent not to start under public experimentation. Hence, there are costs where the principal's as well as the agent's utility are lower under public experimentation than under private experimentation.

In order to complete the proof it remains to be shown that there are (p, c) combinations, such that the plan $H_s^{priv}=\{(1,1),(1,0,1)\}$, $H_f^{priv}=\{(0),(1,0,0)\}$ is indeed optimal. That is, it pays to start experimenting and it pays to stop after history $(1,0,0)$ (which implies that it pays to stop after the history (0)). The agent starts if $\frac{1}{2}U(p^2+p^2(1-p))+\frac{1}{2}U((1-p)^2+(1-p)^2p)-c(\frac{1}{2}(p^2+p^2(1-p)+(1-p)+(1-p)^2p)+\frac{1}{2}((1-p)^2+(1-p)^2p+p+p^2(1-p)))\geq 0$, which is equivalent to $\frac{c}{U}\leq\frac{1}{2}p^2-\frac{1}{2}p+\frac{1}{2}$.

Let $\frac{c}{U}=\frac{1}{2}p^2-\frac{1}{2}p+\frac{1}{2}$. The agent stops experimenting after history $(1,0,0)$ if the expected utility from running an additional experiment is negative. Given this history it must be true that $prob\{y_4=1\mid(1,0,0)\}=prob\{y_2=1\mid y_1=0\}$ $=\frac{(1-p)p+p(1-p)}{(1-p)p+p(1-p)+(1-p)^2+p^2}$ and the expected utility from experimenting once more is $\frac{(1-p)p+p(1-p)}{(1-p)p+p(1-p)+(1-p)^2+p^2}U-c=\frac{(1-p)p+p(1-p)}{(1-p)p+p(1-p)+(1-p)^2+p^2}U-(\frac{1}{2}p^2-\frac{1}{2}p+\frac{1}{2})U$, which is smaller than zero if $p\geq\frac{1}{10}\sqrt{5}+\frac{1}{2}$. Let $p=\frac{1}{10}\sqrt{5}+\frac{1}{2}$. An equilibrium exists for these parameters if $p_d<\frac{p^2+p^2(1-p)}{p^2+p^2(1-p)+(1-p)^2+p(1-p)^2}=\frac{3}{20}\sqrt{5}+\frac{1}{2}$. Q.E.D.

Literature

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.

ARGENZIANO, R., SQUINTANI, F. and SEVERINOV, S. (2012). "A numerical analysis of strategic information acquisition and transmission", *working paper*.

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*.

FELGENHAUER, M. and SCHULTE, E. (2013). "Strategic private experimentation", *working paper*.

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.

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

JOVANOVIC, B. (1982). "Truthful disclosure of information", *Bell Journal of Economics*, Vol. 13, No. 1, pp. 36-44.

KAMENICA, E. and GENTZKOW, M. (2011). "Bayesian persuasion", *American Economic Review*, Vol. 101, No. 6, pp. 2590-2615.

KELLER, G., RADY, S. and CRIPPS, M. (2005). "Strategic experimentation with exponential bandits", *Econometrica*, Vol. 73, No. 1, pp. 39-68.

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.

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