

Bayesian Persuasion with Private Experimentation*

August 29, 2017

BY MIKE FELGENHAUER AND PETRA LOERKE¹

Plymouth University, United Kingdom; NERA Economic Consulting, Germany

Abstract: This paper studies a situation in which a sender tries to persuade a receiver by providing hard evidence that is generated by sequential private experimentation where the sender can design the properties of each experiment contingent on the experimentation history. The sender can selectively reveal as many outcomes as desired. We determine the set of equilibria that are not Pareto dominated. In each of these equilibria under private experimentation the persuasion probability is lower and the receiver obtains access to higher quality information than under public experimentation. The decision quality improves in the sender's stakes.

Keywords: Experimentation, persuasion, information acquisition.

JEL classification: D82, D83

Shortened title: Persuasion with Experimentation

*Manuscript received June 2015; revised December 2015.

¹An early version of this paper circulated under the title "Public versus private experimentation". We thank Johannes Hörner, Andras Niedermayer, Elisabeth Schulte, Nicolas Schutz, Dmitry Shapiro, Ran Spiegler, Roland Strausz, Ernst-Ludwig von Thadden and Juuso Välimäki for valuable comments and suggestions. A substantial part of the work on this paper was done while the authors were at the University of Mannheim. Their support is gratefully acknowledged. The opinions expressed in this article are those of the authors and should not be regarded as those of NERA Economic Consulting. Please address correspondence to: Mike Felgenhauer, Plymouth University, Cookworthy Building, Plymouth, PL4 8AA, UK. Phone: 01752-585541. E-mail: mike.felgenhauer@plymouth.ac.uk.

1 Introduction

There are many situations in which arguments are exchanged, as in lobbying, public discussions of economic policies, the academic publishing process, etc. Arguments have an inherent meaning and are not cheap talk. We think that many arguments, like logical arguments or a regression analysis on a public database, can be viewed as hard (i.e., non-manipulable) and imperfect decision relevant evidence.² Furthermore, arguments often have to be acquired. Kamenica and Gentzkow (2011) describe a trial as a public experiment yielding an argument. The prosecutor can design the error probabilities in order to achieve his objectives. The error probabilities can be affected, e.g., by structuring the examination of witnesses in court. They determine the prosecutor's optimal experiment for persuading a Bayesian judge.

While a public experiment is an interesting case for generating an argument, there is an abundance of situations in which arguments stem from sequential private experimentation. For example, if a sender wants to persuade a receiver with logical arguments, then he runs a series of thought experiments.³ He privately chooses the properties of each experiment, e.g., by choosing the conceptual framework from which to draw a set of specific assumptions. The outcome of each thought experiment is privately observed. Naturally, he may run as many thought experiments with properties of his choice (depending on what he has learnt from previous experimentation) as desired and selectively reveal the results. If a sender instead wants to persuade with an empirical analysis using a public database, then he may run a series of regressions. By choosing the econometric method and the model specification he chooses the properties of each experiment. Again experimentation is sequential and the privately observed results are revealed selectively.

This paper studies a situation in which a sender tries to persuade a Bayesian receiver by providing experimental evidence that stems from sequential experimentation. Information acquisition occurs in private and the experimental evidence can be selectively revealed. We

²Once a regression method is described and the database is public, manipulation is not possible. We make a similar case for logical arguments in footnote 3. Such arguments have persuasive power and, hence, they can be viewed as decision relevant. Naturally, they are also imperfect.

³Felgenhauer and Schulte (2014) interpret, e.g., logical arguments as decision relevant hard information that result from experimentation: If the assumptions underlying a logical argument are revealed, then they cannot be manipulated. The deductions are logical and logic cannot be manipulated. Logical arguments 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 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 natural.

assume that an experimental outcome is hard evidence, that the sender designs the precision of each experiment contingent on the experimentation history and that his decision to continue experimenting also depends on the experimentation history.

If arguments stem from sequential private experimentation and are selectively revealed, then the revealed evidence should not be taken at face value. The value of such arguments depends on the equilibrium experimentation plan, which in turn is influenced by experimentation costs and the sender's benefit from the receiver's decision. An experimentation plan in our model is a complex object, since the sender can make many history dependent choices. The sender also has considerable degrees of freedom regarding the messages that he can send, as he may reveal any subset of the acquired evidence, including "counterarguments".

We derive the set of equilibria that are not Pareto dominated under private experimentation for constant experimentation costs. In each of these equilibria the sender runs one experiment and stops after observing either realization. In the sender preferred equilibrium the precision of the experiment is sufficiently high such that the sender is just deterred from continuing experimentation after observing an adverse outcome. In any other equilibrium that is not Pareto dominated the precision is even higher. The persuasion probability in any such equilibrium is lower than under public experimentation. Private experimentation, thus, limits the extent to which persuasion is possible.

We compare the payoffs under public and private experimentation. Under private experimentation the sender is worse off than under public experimentation due to the lower persuasion probability. The receiver on the other hand is better off in each of these equilibria under private experimentation due to the higher precision of the experiment.

As an application consider a pharmaceutical company (the sender) that attempts to persuade the U.S. Food and Drug Administration (the receiver) 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 acquired and revealed and what evidence is permitted to be considered as decision relevant. The evidence production may be designed as public, by imposing severe penalties if this rule is violated, or as private. Our paper suggests that the FDA would be better off under the private scheme, but the company would benefit more from public experimentation.

We analyze the impact of the sender's stakes on the decision quality. We compare a

situation where a sender does not care much about his favored decision with a situation where he cares substantially. We find that the decision quality in any equilibrium that is not Pareto dominated in the former case is lower than in any equilibrium that is not Pareto dominated in the latter case. In a context where an interested party (like a lobby, student, researcher, etc.) tries to persuade a decision maker (like a politician, teacher, editor, etc.) to choose a favorable action (like a policy, a better mark in the exam, the publication of a paper, etc.) this means that the decision maker is better off the more the interested party benefits from a favorable decision. The interested party then has to provide higher quality information in order to be able to commit not to run additional private experiments after an initial failure.

We further find that Kamenica and Gentzkow's (2011) concavification approach for the derivation of the sender's optimal experiment in general cannot be applied in the case of private experimentation, as the sender may have an incentive to keep an adverse outcome of such an experiment secret and to continue experimenting if costs are low.

2 Literature

Our paper is related to the extensive literature on persuasion (e.g., Jovanovic 1982, Milgrom and Roberts 1986, Glazer and Rubinstein 2001, 2004, 2006). As ours, some papers combine information acquisition with persuasion (e.g., Celik 2003, Brocas and Carillo 2008, Henry 2009, Kamenica and Gentzkow 2011, Gentzkow and Kamenica 2014, 2016, Felgenhauer and Schulte 2014).⁴

Henry (2009) and Brocas and Carillo (2008) investigate private experimentation in settings where the receiver knows or can deduce the number of experiments that the sender ran. Their models allow for an unraveling argument à la Milgrom and Roberts (1986). 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

⁴There is also a literature on strategic experimentation (e.g., Rothschild 1974, Aghion et al.1991, Bolton and Harris 1999, Keller et al. 2005 and Rosenberg et al. 2007). However, in contrast to this literature on bandit problems the purpose of experimentation here is the design of costly signals. For a survey see Bergemann and Välimäki (2008). Experimentation in our sense is also analyzed in papers dealing with the classical problem of sequential analysis (e.g., Moscarini and Smith 2001).

⁵For example, if the sender finds too many unfavorable results in the first experiments, then he knows that he cannot persuade the receiver 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. Baliga and Ely (2016) instead study a repeated receiver - sender interaction where the receiver makes history dependent choices which leads to a commitment problem on the receiver's side.

beliefs à la Milgrom and Roberts are not helpful in such a setting: The receiver, in general, cannot deduce the number of experiments that the sender ran, she only knows the equilibrium experimentation plan.⁶ Nevertheless, in our paper it turns out that the sender runs a single experiment in any equilibrium that is not Pareto dominated and, hence, this experiment is basically public.

The papers that are most closely related to ours are Kamenica and Gentzkow (2011) and Felgenhauer and Schulte (2014), in the following KG and FS, respectively. KG study persuasion via a costless public experiment, where the sender can freely design the experiment. They apply a concavification approach in the context of persuasion that enables the determination of the sender’s optimal experiment for persuading a Bayesian receiver.⁷ Our assumptions regarding the players’ preferences and the experimentation technology resemble KG’s model of public experimentation. Their setup is then extended in order to capture the effects of private experimentation. FS also study private experimentation, but with an exogenously fixed precision of the experiments. The current paper generalizes FS by endogenizing the precision.⁸ We discuss the contribution of our paper relative to FS in detail in section 9.

3 Model

3.1 Preferences

A receiver chooses action $a \in A$, with $A = \{a_1, a_2\}$. Her payoff depends on her action and an unknown state of the world $\omega \in \Omega$, with $\Omega = \{\omega_1, \omega_2\}$ and $prob\{\omega = \omega_1\} = 1/2$. The receiver’s utility is

	$\omega = \omega_1$	$\omega = \omega_2$
$a = a_1$	1	$1 - p_d$
$a = a_2$	p_d	1

⁶In a persuasion setting with a Poisson evidence production technology, Celik (2003) also finds that no productive fully revealing equilibrium exists if the sender is ex ante uninformed about the state.

⁷Aumann and Maschler (1995) use the concavification approach prior to KG in a context with incomplete information in which a fixed zero-sum game is infinitely often repeated. Gentzkow and Kamenica (2014) study KG’s setup with costly experimentation. They show that KG’s concavification approach extends to settings where the costs of a signal are proportional to the expected reduction in uncertainty. Gentzkow and Kamenica (2016) study competition between several senders who try to persuade a receiver.

⁸Felgenhauer (2016) studies endogenous persuasion where the receiver may not pay attention to the persuasion attempt.

with $p_d \in (1/2, 1)$. The receiver, thus, would like to match the decision a with the state of the world ω if she knew ω . At the optimum she only chooses $a = a_1$ if her posterior belief passes the “threshold of doubt” p_d , i.e., the posterior that $\omega = \omega_1$ must be weakly greater than p_d .

There is a sender who prefers $a = a_1$ regardless of ω . His prior belief is also $\text{prob}\{\omega = \omega_1\} = 1/2$. His gross utility is U if $a = a_1$ and 0 otherwise. Experimentation costs have to be subtracted from the gross utility.

3.2 Experimentation

The sender has access to an experimentation technology that can generate signals about ω . The outcome of experiment τ is $\sigma_\tau \in \{s_1, s_2\}$. We call s_1 a “positive outcome” and s_2 an “adverse outcome”. The precision of an experiment τ is $\pi_\tau = (\pi_\tau(s_1 | \omega_1), \pi_\tau(s_2 | \omega_2))$, with $\pi_\tau(s_j | \omega_j) = \text{prob}\{\sigma_\tau = s_j | \omega = \omega_j\}$, $\pi_\tau(s_j | \omega_j) \in [0, 1]$, $j \in \{1, 2\}$. Let $\pi_\tau(s_1 | \omega_1) \geq 1 - \pi_\tau(s_2 | \omega_2)$, i.e., a positive outcome $\sigma_\tau = s_1$ is more likely if $\omega = \omega_1$ than if $\omega = \omega_2$. We say that the precision of an experiment τ increases if ceteris paribus either $\pi_\tau(s_1 | \omega_1)$ increases or $\pi_\tau(s_2 | \omega_2)$ increases or both increase.⁹ The costs of running an experiment are $c \geq 0$.¹⁰ The experimentation history after the first t experiments is denoted by $h = \{(\sigma_j, \pi_j)\}_{j=1, \dots, t}$. The sender chooses the precision for each experiment that he runs contingent on the experimentation history.

3.3 Messages

The sender cannot manipulate or make up experimental outcomes, i.e., each outcome is “hard” information. The sender’s message is denoted by $m = \{(\sigma_i, \pi_i)\}_i$. Note that the receiver does not only observe the outcomes contained in a message, but also the precision of the experiments with which these outcomes were generated.¹¹ Let M^* be the set of messages that are sent with a positive probability on the equilibrium path and let m^* be an element of M^* .

⁹Consider experiments τ and τ' with $\pi_\tau(s_1 | \omega_1) \geq \pi_{\tau'}(s_1 | \omega_1)$ and $\pi_\tau(s_2 | \omega_2) \geq \pi_{\tau'}(s_2 | \omega_2)$ with at least one strict inequality, i.e., experiment τ has the higher precision. According to the Blackwell criterion experiment τ is at least as informative as experiment τ' .

¹⁰In the online appendix we discuss alternative cost structures. We derive a sufficient condition such that the sender runs a single experiment in the sender preferred equilibrium under private experimentation if experimentation costs are not constant.

¹¹The receiver observes the properties of an experiment, once the outcome of this experiment is presented. This assumption is natural in many applications. E.g., if a theoretical argument is considered as evidence, then a scientific audience, e.g., referees, editors or seminar participants, can assess its quality.

Under private experimentation the receiver observes message m but she cannot observe the experimentation history h at which the sender stops experimenting. A message $m = \{(\sigma_i, \pi_i)\}_i$ is feasible given history h if $(\sigma_i, \pi_i) \in h$ for each i . The assumption that information is hard implies that the sender can only send feasible messages. The sender cannot prove that he did not conduct a particular experiment.

Under public experimentation the receiver observes the experimentation history h at which the sender stops experimenting, i.e., $m = h$.

3.4 Timing and strategies

The sender moves first. His strategy specifies his behavior at each experimentation history h that he may observe. At each h the sender may either continue experimenting with a further experiment with a history dependent precision or he may stop experimenting and send his message.¹² The receiver's strategy specifies an action for each message that she may observe.

4 Equilibrium concept

Our equilibrium concept is weak perfect Bayesian equilibrium. Off-the-equilibrium path beliefs have to satisfy that the sender cannot signal what he does not know.¹³

4.1 Equilibrium conditions under private experimentation

4.1.1 The sender's equilibrium strategy

The sender at each h in equilibrium chooses either the precision of the next experiment or a feasible message such that his continuation payoff at h is maximized given the anticipated

¹²Under private experimentation we exclude the possibility of inferring information from the length of the experimentation phase. Hopenhayn and Squintani (2011) study in a different context the case where the decision maker can deduce something from the time elapsed until he receives information. A longer period may, e.g., suggest many failed experiments or that the sender ran a complex experiment. Our model abstracts from these issues. Often experiments differ regarding the time they require until completed and it is difficult to deduce information from the time elapsed.

¹³To illustrate the implications consider public experimentation, where the experimentation history is common knowledge. If the sender deviates from his equilibrium experimentation plan, then this assumption implies that the receiver cannot have arbitrary beliefs with respect to the state of the world, since the sender 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 under public experimentation.

equilibrium behavior. I.e., at each history h we have the following. Let \tilde{h} be a history where either history $h_1 = (h, (\sigma_\tau = s_1, \pi_\tau))$ or history $h_2 = (h, (\sigma_\tau = s_2, \pi_\tau))$ is a subhistory, with experiment τ being the next experiment at history h . In equilibrium the sender at each h anticipates (i) his equilibrium behavior at any history \tilde{h} for each π_τ and (ii) the receiver's equilibrium response to any message that the sender may send. If the sender in equilibrium runs a further experiment τ at h , then the precision of this experiment π_τ satisfies

$$\pi_\tau \in \arg \max_{\pi_\tau} \text{prob}\{\sigma_\tau = s_1 \mid h, \pi_\tau\} \Upsilon(h_1) + \text{prob}\{\sigma_\tau = s_2 \mid h, \pi_\tau\} \Upsilon(h_2) - c,$$

where $\Upsilon(\cdot)$ denotes the sender's equilibrium continuation utility given the anticipated behavior (i) and (ii). If he sends a feasible message at h that induces $a = a_2$ according to the receiver's equilibrium strategy, then the sender's continuation payoff at h is 0. If he sends a feasible message at h that induces $a = a_1$, then his continuation payoff at h is U . Continuing experimentation is equilibrium behavior if h is such that a message inducing $a = a_1$ is not feasible and $\max_{\pi_\tau} \text{prob}\{\sigma_\tau = s_1 \mid h, \pi_\tau\} \Upsilon(h_1) + \text{prob}\{\sigma_\tau = s_2 \mid h, \pi_\tau\} \Upsilon(h_2) - c > 0$. Otherwise, he sends a best feasible message.

4.1.2 The receiver's equilibrium strategy

In equilibrium the receiver's action is optimal for each message that she may observe given her beliefs, i.e., we have

$$a(m) = \begin{cases} a_1 & \text{if } \text{prob}\{\omega = \omega_1 \mid m\} \geq p_d \\ a_2 & \text{otherwise} \end{cases}$$

for each potential message m , where $\text{prob}\{\omega = \omega_1 \mid m\}$ is the receiver's posterior belief that the state is $\omega = \omega_1$ upon observing message m .

4.1.3 Equilibrium beliefs

The receiver does not observe the sender's experimentation history, but in equilibrium she deduces the sender's equilibrium behavior at each history h that he may face. This allows her to determine a Bayesian posterior belief about the sender's type, i.e., the experimentation history that he faces when sending the message, for each potential on the equilibrium path message. Let $H(m)$ be the set of all histories that are such that the sender sends message m according to his equilibrium strategy. The probability with which each of these histories occurs in each state ω can be deduced from the sender's equilibrium strategy. Hence, if the receiver deduces that the sender follows his equilibrium strategy and she observes some on the

equilibrium path message m^* , then she can determine her on the equilibrium path Bayesian belief

$$prob\{\omega = \omega_1 \mid m^*\} = \frac{\sum_{h \in H(m^*)} prob\{h \mid \omega_1\}}{\sum_{h \in H(m^*)} prob\{h \mid \omega_1\} + \sum_{h \in H(m^*)} prob\{h \mid \omega_2\}}.$$

If the receiver observes some off-the-equilibrium path message $m = \{(\sigma_i, \pi_i)\}_i$, then we assume that she only considers experimentation histories h as possible with $(\sigma_i, \pi_i) \in h$ for each i . This implies $prob\{\omega = \omega_1 \mid m\} = 1$ if m contains at least one positive outcome from an experiment τ with precision $(\pi_\tau(s_1 \mid \omega_1), 1)$ with some $\pi_\tau(s_1 \mid \omega_1) \in (0, 1]$.¹⁴ Otherwise, we do not impose restrictions on off-the-equilibrium path beliefs.

4.2 Equilibrium conditions under public experimentation

The equilibrium conditions for the sender's and the receiver's equilibrium strategy are analogous to the conditions under private experimentation with the modification that $m = h$ at any history h where the sender stops experimenting. As the receiver observes the experimentation history, she observes the sender's type. If the sender stops experimenting at some history h , then her posterior belief that the state is $\omega = \omega_1$ on and off-the-equilibrium path is

$$prob\{\omega = \omega_1 \mid h\} = \frac{prob\{h \mid \omega_1\}}{prob\{h \mid \omega_1\} + prob\{h \mid \omega_2\}}.$$

5 Public experimentation

In this section we derive the equilibrium behavior under public experimentation. It is convenient to first suppose that the sender can run a single public experiment with a precision of his choice. The other assumptions are maintained. KG use a concavification approach to solve the sender's problem to determine the optimal experiment with which he can persuade a Bayesian receiver. The sender anticipates a distribution of the receiver's optimal actions in response to the posterior distributions that result from his experiment. KG show (i) that the sender's payoff can be described as a value function over the posterior generated by the experimental outcome and (ii) that for any distribution of posterior beliefs whose expectation is the prior, there exists an experiment that induces that distribution of posteriors. They find that the

¹⁴Any such message can only stem from experimentation histories that also contain this evidence. Each of these histories implies that the state is $\omega = \omega_1$ with certainty. Regardless of the posterior belief regarding the histories we, thus, have $prob\{\omega = \omega_1 \mid m\} = 1$.

highest payoff that the sender can obtain is equal to the concave closure of the value function evaluated at the prior belief.¹⁵ This determines the precision π_τ of the optimal experiment.

Figure 1 shows the sender's value function that depends on the receiver's posterior probability $\text{prob}\{\omega = \omega_1 \mid m\} \equiv \mu_1$ that the state is ω_1 and the value function's concavification. The highest payoff that the sender can obtain is equal to the concave closure of the value function evaluated at the prior belief $1/2$, i.e., in this case it is equal to $\frac{U}{2p_d} - c$.




Figure 1: The sender's value function (solid curve) and its concavification (dotted curve)

In order to maximize the persuasion probability, the sender designs the experiment yielding posteriors $\mu_1 = 0$ and $\mu_1 = p_d$ such that the expected posterior is $\mu_1 = 1/2$, i.e., he chooses $\pi_\tau(s_1 \mid \omega_1) = 1$ and $\pi_\tau(s_2 \mid \omega_2) = \frac{2p_d - 1}{p_d}$. The intuition for $\pi_\tau(s_1 \mid \omega_1) = 1$ is that the sender does not want to obtain an adverse outcome if $\omega = \omega_1$. He also wants to obtain a positive outcome when the state is bad $\omega = \omega_2$, but he has to choose $\pi_\tau(s_2 \mid \omega_2)$ such that he can still persuade the receiver. Therefore, he maximizes $\text{prob}\{s_1 \mid \omega_2\}$ subject to $\text{prob}\{\omega_1 \mid s_1\} \geq p_d$. The persuasion probability is maximal if $\text{prob}\{\omega_1 \mid s_1\} = p_d$, i.e., such that the receiver is indifferent between a_1 and a_2 upon the observation of a signal realization in the sender's favor, yielding $\pi_\tau(s_2 \mid \omega_2) = \frac{2p_d - 1}{p_d}$.

Using the above analysis we obtain the following lemma if the sender can run multiple public experiments.

Lemma 1 *Consider public experimentation. (i) If $U/c \in [0, 2p_d)$, then there is no equilibrium with experimentation. (ii) If $U/c \in [2p_d, \infty)$, then there is an equilibrium in which the sender on the equilibrium path runs one experiment with precision $\pi_\tau = (1, \frac{2p_d - 1}{p_d})$ and no further experiment. (iii) If $U/c \in (2p_d, \infty)$, then the equilibrium described in (ii) is unique.*

¹⁵The concavification of the value function is the smallest concave function everywhere weakly above the value function.

If U/c is too small, i.e., $U/c \in [0, 2p_d)$, then the probability to obtain U is not high enough for the sender to justify the costs of running an experiment with $\pi_\tau = (1, \frac{2p_d-1}{p_d})$. If $U/c > 2p_d$, then the receiver is better off to run an experiment with precision $\pi_\tau = (1, \frac{2p_d-1}{p_d})$ than to run any other experiment or not to experiment. He does not want to continue experimenting after observing an adverse outcome of the experiment with precision $\pi_\tau = (1, \frac{2p_d-1}{p_d})$, as this outcome implies that the state is ω_2 with certainty and this posterior belief cannot be changed by further experimentation.

6 Private experimentation

Consider costly private experimentation, where the receiver cannot observe the experimentation history and the sender cannot commit not to run additional private experiments. In equilibrium the receiver only chooses the sender's preferred action $a = a_1$ if her posterior belief that $\omega = \omega_1$ upon observing his message exceeds her threshold of doubt p_d . The receiver's posterior belief that the state is $\omega = \omega_1$ upon the observation of some message depends on the sender's strategy.

The following lemma asserts that in any equilibrium with persuasion the sender stops experimenting unsuccessfully at some history h where (i) the posterior at h that the state is $\omega = \omega_1$ is below the prior belief $1/2$ and where (ii) history h does not contain evidence that he can use for persuasion. Only in this case the receiver's threshold of doubt may be passed and, thus, persuasion may be possible. Denote the sender's posterior belief that the state is $\omega = \omega_1$ given that he faces experimentation history h by $\text{prob}\{\omega = \omega_1 \mid h\} \equiv \bar{\mu}_1$.

Lemma 2 *In any equilibrium with persuasion under private experimentation the sender stops experimenting unsuccessfully at some posterior $\bar{\mu}_1 < 1/2$ if he has not yet found an outcome that is part of some $m^* \in M^*$ with which he can persuade the receiver to choose $a = a_1$.*

Consider, e.g., an equilibrium with persuasion in which there is no on the equilibrium path history where the sender never stops experimenting unsuccessfully. In this equilibrium the sender, thus, runs a finite number of experiments. He eventually either stops experimenting successfully or he stops unsuccessfully. If there is no $\bar{\mu}_1$ as described in the lemma, then the sender only stops experimenting unsuccessfully if his posterior is above $1/2$, as shown in the appendix. Bayesian plausibility requires that there is then some persuasive message sent after successful experimentation for which the receiver's posterior belief that the state is $\omega = \omega_1$ is below $1/2$. As this posterior is below the threshold of doubt p_d , the receiver should not be

persuaded by such a message, which violates an equilibrium condition. The other potential equilibria are discussed in the appendix.

In the next subsection we analyze whether a similar equilibrium exists as under public experimentation. Then we derive the sender preferred equilibrium, which enables us to characterize the set of equilibria that are not Pareto dominated in the final subsection.

6.1 Experimentation as under the public scheme

We now determine the circumstances where a similar equilibrium exists under private experimentation as under public experimentation and where this is not the case. Consider a potential equilibrium in which the sender starts experimenting with an experiment, where the precision of this experiment is derived with the concavification approach, then stops after each outcome and sends a message containing the corresponding outcome. The receiver's action after each message equals her action after the corresponding outcome if the experiment is run under public experimentation.

Lemma 3 *Consider private experimentation. (i) If $U/c \in [0, 2p_d)$, then there is no equilibrium with experimentation. (ii) If $U/c \in [2p_d, \frac{p_d}{1-p_d}]$, then there is an equilibrium in which the sender on the equilibrium path runs one experiment with precision $\pi_\tau = (1, \frac{2p_d-1}{p_d})$ and no further experiment. (iii) If $U/c \in (\frac{p_d}{1-p_d}, \infty)$, then there is no equilibrium as described in (ii).*

Analogous to public experimentation there is no equilibrium with experimentation if U/c is too small, i.e., $U/c \in [0, 2p_d)$. If $U/c \in [2p_d, \frac{p_d}{1-p_d}]$, then there is an equilibrium in which the sender on the equilibrium path runs a single experiment and this experiment has the same precision as under public experimentation. The receiver is persuaded and chooses $a = a_1$ if the sender presents a corresponding positive outcome. $U/c \leq \frac{p_d}{1-p_d}$ implies that the sender does not have an incentive to continue experimenting privately after an adverse outcome. This equilibrium is better for the sender than any other equilibrium under private experimentation, as the sender runs a single experiment with the maximum persuasion probability. If U/c is above $\frac{p_d}{1-p_d}$, then there is no equilibrium under private experimentation where the sender on the equilibrium path runs one experiment with the same precision as under public experimentation and no further experiments. Suppose instead that there is such an equilibrium. If the outcome of the first experiment is s_2 , then the sender knows that $\omega = \omega_2$. It is worthwhile to continue experimenting, given that the receiver can be persuaded with one positive outcome from an experiment with precision $\pi_\tau = (1, \frac{2p_d-1}{p_d})$, if $U/c > \frac{1}{(1-\frac{2p_d-1}{p_d})} = \frac{p_d}{1-p_d}$. Consequently, this cannot be an equilibrium if U/c is sufficiently high. The sender with high stakes cannot

commit not to run a further experiment if he observes an outcome of the first experiment that he does not like and he can hide an adverse outcome of the first experiment.

KG's concavification approach is very useful for determining the equilibrium under public experimentation. However, it is not applicable in many situations where the evidence for persuasion is collected via sequential private experimentation with selective information revelation. Consider the experiment in the unique equilibrium under public experimentation, which is derived with the concavification approach and which has precision $(1, \frac{2p_d-1}{p_d})$. A comparison of Lemma 1 (ii) and Lemma 3 (iii) reveals that an analogous equilibrium does not exist under private experimentation if $U/c > \frac{p_d}{1-p_d}$. Hence, if $U/c > \frac{p_d}{1-p_d}$, the concavification approach can be used under public experimentation, but not under private experimentation.

Lemmas 1 and 3 imply that the persuasion probability, the players' payoffs, etc. are the same under public and private experimentation if $U/c \in [0, \frac{p_d}{1-p_d}]$ given that the players coordinate on the equilibrium described in Lemma 3 (ii) under private experimentation. The graph in Figure 2 illustrates $U/c = \frac{p_d}{1-p_d}$. We have $U/c \in [0, \frac{p_d}{1-p_d}]$ below the curve.




Figure 2: $U/c = \frac{p_d}{1-p_d}$

In the following we focus on the interesting case with $U/c > \frac{p_d}{1-p_d}$, i.e., the region above the curve in Figure 2.

6.2 The sender preferred equilibrium

For the derivation of the set of all equilibria that are not Pareto dominated it is convenient to first determine the sender preferred equilibrium. An equilibrium is sender preferred if there is no other equilibrium in which the sender is strictly better off.

The sender cares about the persuasion probability and expected experimentation costs. If $U/c > \frac{p_d}{1-p_d}$, then there is no equilibrium in which the sender runs a single experiment that leads to the maximum persuasion probability as under public experimentation. The sender in

equilibrium may run multiple experiments and may reveal more than one outcome. However, the sender preferred equilibrium has a surprisingly simple structure.

For the derivation of the sender preferred equilibrium it is useful to characterize the sender's payoff in any equilibrium with persuasion. Define $V_2(\bar{\mu}_1)$ as the sender's expected utility if (i) $\omega = \omega_2$, (ii) he has not yet found an outcome that is an element of some persuasive message, and (iii) he continues experimenting according to his equilibrium experimentation plan given that he holds the posterior belief $prob\{\omega = \omega_1 \mid h\} = \bar{\mu}_1$ and given that he has not yet found an outcome that is an element of some persuasive message. Analogously, define $V_1(\bar{\mu}_1)$ as his expected utility given $\omega = \omega_1$. The sender's utility if he has not yet found an outcome that he can use for persuasion at posterior $\bar{\mu}_1$ is

$$\bar{\mu}_1 V_1(\bar{\mu}_1) + (1 - \bar{\mu}_1) V_2(\bar{\mu}_1).$$

The following lemma characterizes the sender's payoff in each state, given that he does not yet have evidence for persuasion and given the prior belief, as well as his ex ante payoff in any equilibrium with persuasion.

Lemma 4 *In any equilibrium with persuasion we have $V_1(\frac{1}{2}) \geq 0$ and $V_2(\frac{1}{2}) \leq 0$ and $\frac{1}{2}V_1(\frac{1}{2}) + \frac{1}{2}V_2(\frac{1}{2}) \geq 0$.*

$\frac{1}{2}V_1(\frac{1}{2}) + \frac{1}{2}V_2(\frac{1}{2}) \geq 0$ has to be satisfied, as the sender has to have an incentive to start experimenting. We cannot have $V_1(\frac{1}{2}) < 0$ and $V_2(\frac{1}{2}) \leq 0$, as this violates $\frac{1}{2}V_1(\frac{1}{2}) + \frac{1}{2}V_2(\frac{1}{2}) \geq 0$. Similarly, we cannot have $V_1(\frac{1}{2}) \leq 0$ and $V_2(\frac{1}{2}) < 0$. In the appendix we show that we cannot have $V_2(\frac{1}{2}) > 0$ in an equilibrium with persuasion, as the sender in this case would not stop experimenting unsuccessfully for all posteriors $\bar{\mu}_1 < 1/2$, which violates Lemma 2.

Under public experimentation we have determined the experiment that maximizes the sender's payoff subject to a positive outcome having the power to persuade a Bayesian receiver. Lemma 4 asserts that the sender's commitment problem under private experimentation results in the additional constraint that $V_2(\frac{1}{2}) \leq 0$ in any equilibrium with persuasion. The commitment problem limits the sender's payoff in the state where the decision should be against him. If instead $V_2(\frac{1}{2}) > 0$, then he would not stop experimenting unsuccessfully at any posterior belief $\bar{\mu}_1 < 1/2$, which implies that some persuasive message m^* yields a posterior $\mu_1 < 1/2$ for the receiver due to Bayesian plausibility. But then the receiver should not be persuaded by this m^* . In the following, we show that the constraint $V_2(\frac{1}{2}) \leq 0$ affects the nature of the sender preferred equilibrium under private experimentation. We derive an upper bound for the sender's ex ante utility in any equilibrium with persuasion and then show that

there is such an equilibrium in which the sender runs only one experiment that yields this payoff for the sender.

The sender's ex ante utility in any equilibrium with persuasion can be written as $\frac{1}{2}V_1(\frac{1}{2}) + \frac{1}{2}V_2(\frac{1}{2})$. An upper bound for $V_1(\frac{1}{2})$ is $U - c$, i.e., if persuasion occurs in state ω_1 with certainty after the first experiment. An upper bound for $V_2(\frac{1}{2})$ in any equilibrium with persuasion is 0 due to Lemma 4. Hence, an upper bound for the sender's ex ante utility in any equilibrium with persuasion is

$$(1) \quad \frac{1}{2}(U - c).$$

Suppose the sender in a potential equilibrium runs a single experiment τ with precision π_τ and suppose that a positive outcome induces $a = a_1$ and an adverse outcome induces $a = a_2$. We now maximize $\frac{1}{2}V_1(\frac{1}{2}) + \frac{1}{2}V_2(\frac{1}{2})$ with respect to π_τ subject to $V_1(\frac{1}{2}) \geq 0$ and $V_2(\frac{1}{2}) \leq 0$. We have $V_1(\frac{1}{2}) = \pi_\tau(s_1 | \omega_1)U - c$ and $V_2(\frac{1}{2}) = (1 - \pi_\tau(s_2 | \omega_2))U - c$. The solution to the maximization problem is $\pi_\tau(s_1 | \omega_1) = 1$ and $\pi_\tau(s_2 | \omega_2) = \frac{U-c}{U}$. The former implies $V_1(\frac{1}{2}) = U - c$ and the latter implies $V_2(\frac{1}{2}) = 0$. Hence, if such an equilibrium exists, then the sender's ex ante utility in this equilibrium is $\frac{1}{2}(U - c)$ and it is weakly higher than in any other equilibrium with persuasion due to (1), i.e., the equilibrium is sender preferred.

Next, we show that such an equilibrium exists. Consider a potential equilibrium with the following properties:

- (1) The sender's first experiment has precision $\pi_\tau = (1, \frac{U-c}{U})$.
- (2) The sender stops experimenting (successfully) at any history h that contains at least one outcome $\sigma_\tau = s_1$ of an experiment with $\pi_\tau = (1, \frac{U-c}{U})$.
- (3) The sender stops experimenting (unsuccessfully) at any history h with $prob\{\omega = \omega_1 | h\} = 0$ that does not contain at least one outcome $\sigma_\tau = s_1$ of an experiment with $\pi_\tau = (1, \frac{U-c}{U})$.
- (4) The sender sends message $m^* = (\sigma_\tau, \pi_\tau)$ with $\pi_\tau = (1, \frac{U-c}{U})$ for any outcome σ_τ that he observes after the first experiment.
- (5) The sender's remaining behavior off-the-equilibrium path is sequentially rational.
- (6) The receiver chooses $a = a_1$ if the sender sends message $m^* = (\sigma_\tau = s_1, \pi_\tau)$ with $\pi_\tau = (1, \frac{U-c}{U})$ of an experiment τ . The receiver also chooses $a = a_1$ if the sender sends a message m that contains at least one positive outcome from an experiment with precision $(q, 1)$, with $q \in (0, 1]$. Otherwise she chooses a_2 .
- (7) On the equilibrium path, i.e., upon observing message $m^* = (\sigma_\tau, \pi_\tau)$ with $\pi_\tau = (1, \frac{U-c}{U})$, beliefs are formed in accordance with Bayes' Law. Upon observing an off-the-equilibrium path message that does not contain at least one positive outcome from an experiment with precision $(q, 1)$, the receiver forms a probability assessment over experimentation histories such that the

probability that $\omega = \omega_1$ conditional on this assessment is below the threshold of doubt. E.g., she may believe that the sender privately ran a single additional experiment with precision $(1, 1)$ that yielded an adverse outcome. Off-the-equilibrium path beliefs upon observing a message that contains at least one positive outcome from an experiment with precision $(q, 1)$, with $q \in (0, 1]$, are such that the receiver thinks that $\omega = \omega_1$ with certainty, as discussed in section 4.1.3.

An equilibrium condition is that the sender stops experimenting after an adverse outcome from the first experiment. E.g., in this case he must not have an incentive to run the same experiment again. After observing an adverse outcome of an experiment with $\pi_\tau(s_1 | \omega_1) = 1$ the sender knows that the state is $\omega = \omega_2$ with certainty. The payoff from running the same experiment again knowing that $\omega = \omega_2$ is equal to $(1 - \pi_\tau(s_2 | \omega_2))U - c$. This payoff is equal to $V_2(\frac{1}{2})$ and by construction $V_2(\frac{1}{2}) = 0$. Not running this experiment yields a payoff of 0 from stopping unsuccessfully. The sender is just deterred from running the same experiment again after observing an adverse outcome of the first experiment, as it is sufficiently unlikely to obtain a positive outcome if the state is against him. In the appendix we confirm that the remaining equilibrium conditions are also satisfied.¹⁶

Proposition 1 *Consider $U/c > \frac{p_d}{1-p_d}$. (i) There is an equilibrium in which the sender on the equilibrium path runs one experiment with precision $\pi_\tau = (1, \frac{U-c}{U})$ and no further experiment. (ii) The sender strictly prefers the equilibrium in (i) to any other equilibrium with persuasion where the sender runs another first experiment or where he runs multiple experiments on the equilibrium path. (iii) The persuasion probability in the sender preferred equilibrium under private experimentation is strictly lower than under public experimentation.*

Proposition 1 shows that in the sender preferred equilibrium under private experimentation the sender runs a single experiment with $\pi_\tau(s_1 | \omega_1) = 1$ (as under public experimentation), but with $\pi_\tau(s_2 | \omega_2) = \frac{U-c}{U}$ that maximizes the persuasion probability subject to the constraint that further experimentation after an initial adverse outcome is deterred. This limits the extent to which persuasion is possible compared to public experimentation.

¹⁶For the sender the ex ante payoff has to be greater than zero, a deviation to start with another experiment must not be profitable, he must not have an incentive to continue experimentation with an experiment of any precision if he knows that $\omega = \omega_2$ and it has to be optimal to make the announcements as stated above. For the receiver the threshold of doubt has to be passed upon observing message $m^* = (\sigma_\tau = s_1, \pi_\tau)$.

6.3 Equilibria that are not Pareto dominated

The characterization of the sender preferred equilibrium allows us to determine the set of all equilibria that are not Pareto dominated under private experimentation. The following proposition shows that the sender in any equilibrium that is not Pareto dominated runs only one experiment and that the precision of this experiment is $(1, \pi_\tau(s_2 | \omega_2))$. These equilibria differ regarding $\pi_\tau(s_2 | \omega_2)$. The other properties of these equilibria are analogous to properties (1) - (7) of the sender preferred equilibrium.

Proposition 2 *Consider $U/c > \frac{pd}{1-pd}$. (i) There is an equilibrium in which the sender on the equilibrium path runs one experiment with precision $(1, \pi_\tau(s_2 | \omega_2))$ and no further experiment for each $\pi_\tau(s_2 | \omega_2) \in [\frac{U-c}{U}, \min\{1; 2\frac{U-c}{U}\}]$. (ii) All equilibria not described in (i) are Pareto dominated by some equilibrium in (i). (iii) None of the equilibria described in (i) is Pareto dominated.*

According to Proposition 2 (i) we have $\pi_\tau(s_1 | \omega_1) = 1$ in any equilibrium that is not Pareto dominated. The lower bound $\frac{U-c}{U}$ for $\pi_\tau(s_2 | \omega_2)$ in Proposition 2 (i) is the $\pi_\tau(s_2 | \omega_2)$ in the sender preferred equilibrium. The upper bound for $\pi_\tau(s_2 | \omega_2)$ is determined by the equilibrium condition that the sender has to be better off from starting to experiment than from not experimenting. This condition is satisfied if $\frac{1}{2}(U-c) + \frac{1}{2}((1-\pi_\tau(s_2 | \omega_2))U-c) \geq 0$. The lower U/c is, the (weakly) higher the probability of a positive outcome has to be in order to satisfy this condition. A lower $\pi_\tau(s_2 | \omega_2)$ ensures a higher probability that the outcome of the experiment is positive. If $U/c \geq 2$, then $\min\{1; 2\frac{U-c}{U}\} = 1$. The sender is just indifferent to start experimenting with an experiment with precision $(1, 1)$ if $U/c = 2$. He has a strict incentive to start experimenting if $U/c > 2$. If $U/c < 2$ instead, then $\min\{1; 2\frac{U-c}{U}\} = 2\frac{U-c}{U} < 1$. The ex ante probability to obtain a favorable outcome with $\pi_\tau(s_2 | \omega_2) = 2\frac{U-c}{U}$ is higher than with $\pi_\tau(s_2 | \omega_2) = 1$. The sender with low stakes $U/c < 2$ is just indifferent to start experimenting given that $\pi_\tau(s_2 | \omega_2) = 2\frac{U-c}{U}$.

In the following we argue that any equilibrium with persuasion not described in Proposition 2 (i) is Pareto dominated by an equilibrium in the set of equilibria described in Proposition 2 (i). There may be equilibria in which the sender runs only one experiment and there may be equilibria with multiple experiments on the equilibrium path. For expositional convenience we focus on parameters U/c such that the upper bound for $\pi_\tau(s_2 | \omega_2)$ in Proposition 2 (i) is 1 and analyze the remaining parameter constellations in the appendix.

Consider a potential equilibrium in which the sender runs only one experiment on the equilibrium path and where this experiment has some precision (q, p) , with $q < 1$ and $p < \frac{U-c}{U}$.

Such an equilibrium is Pareto dominated by the sender preferred equilibrium: The receiver's benefit increases in $\pi_\tau(s_1 | \omega_1)$ and $\pi_\tau(s_2 | \omega_2)$, given that the sender runs only one experiment τ , as this experiment better predicts the state of the world. The sender is best off in the sender preferred equilibrium by the definition of the sender preferred equilibrium.

Consider a potential equilibrium in which the sender runs only one experiment on the equilibrium path and where this experiment has some precision (q, p) , with $q < 1$ and $p > \frac{U-c}{U}$. Such an equilibrium is Pareto dominated by an equilibrium in which the sender runs one experiment and where this experiment has precision $(1, p)$: Analogous to the previous case, the receiver is better off in the latter equilibrium, as the experiment in this equilibrium is more precise. The sender is better off, as in both equilibria $\pi_\tau(s_2 | \omega_2)$ is the same, i.e., in state $\omega = \omega_2$ the payoff is the same, but $\pi_\tau(s_1 | \omega_1)$ is greater in the latter equilibrium, i.e., in state $\omega = \omega_1$ he obtains a positive outcome with a higher probability.

The arguments for equilibria with multiple experiments on the equilibrium path are similar. Here, an auxiliary (possibly non-equilibrium) situation can be constructed, where the sender runs only one experiment and where the receiver only chooses $a = a_1$ if the outcome of this experiment is positive, that has the same probabilities to induce $a = a_1$ in state ω_1 and $a = a_2$ in state ω_2 as the equilibrium with multiple experiments.¹⁷ The receiver is as well off in the auxiliary situation as in the equilibrium. The sender is better off, as he runs only one costly experiment. Analogous to above and as shown in the appendix, the auxiliary situation can then be compared from a payoff perspective with some equilibrium in the set of equilibria that are not Pareto dominated.

As the sender's payoff decreases in $\pi_\tau(s_2 | \omega_2)$ and the receiver's payoff increases in $\pi_\tau(s_2 | \omega_2)$ given that $\pi_\tau(s_1 | \omega_1) = 1$, none of the equilibria described in Proposition 2 (i) is Pareto dominated. The sender preferred equilibrium has the lowest $\pi_\tau(s_2 | \omega_2)$ in the set of equilibria that are not Pareto dominated. The receiver preferred equilibrium has the highest $\pi_\tau(s_2 | \omega_2)$ in this set of equilibria. If $\pi_\tau(s_2 | \omega_2) = 1$ in the receiver preferred equilibrium, then the receiver learns the state and always makes the appropriate decision.

¹⁷Given the sender's and the receiver's equilibrium strategy, we can determine $\text{prob}\{a = a_1 | \omega = \omega_1\}$ and $\text{prob}\{a = a_2 | \omega = \omega_2\}$ in the equilibrium with multiple experiments. A single experiment with precision (q, p) , with $q = \text{prob}\{a = a_1 | \omega = \omega_1\}$ and $p = \text{prob}\{a = a_2 | \omega = \omega_2\}$ can be designed. In the auxiliary situation the sender (possibly suboptimally) runs exclusively this experiment and reveals the outcome and the receiver mechanically chooses $a = a_1$ only after a positive outcome of this experiment. By construction, the persuasion probability in the auxiliary situation is the same as in the equilibrium with multiple experiments.

7 Payoff comparison of public and private experimentation

We now compare the players' utilities under public and private experimentation.

Proposition 3 *Consider $U/c > \frac{p_d}{1-p_d}$. Suppose the players coordinate on any equilibrium that is not Pareto dominated under private experimentation. The sender strictly prefers public to private experimentation. The receiver strictly prefers private to public experimentation.*

The sender ex ante faces a commitment problem under private experimentation. Choosing a high precision is an endogenous way of committing not to search excessively, which affects the players' utilities.¹⁸

Consider first the sender. The equilibrium under public experimentation is such that the sender designs a public experiment that maximizes his payoff subject to being able to persuade a Bayesian receiver. Under private experimentation the sender's commitment problem creates additional constraints in the maximization problem in the sender preferred equilibrium. Therefore, he is weakly worse off in this equilibrium and, hence, in any other equilibrium than under public experimentation. In the following we determine the sender's payoff in the sender preferred equilibrium under private experimentation and his payoff under the public scheme. This allows us to confirm that he is strictly better off under the latter scheme. The persuasion probability under public experimentation is $\frac{1}{2p_d}$. The persuasion probability in the sender preferred equilibrium under private experimentation is $\frac{U+c}{2U}$ if $U/c > \frac{p_d}{1-p_d}$. Under both schemes the sender runs a single experiment, but the persuasion probability is higher under public experimentation, rendering private experimentation less attractive for the sender.

¹⁸Henry (2009) instead studies private experimentation and mandatory disclosure in a setting where the sender ex ante commits to run a certain number of experiments with an exogenous precision. Under private experimentation the receiver deduces the number of experiments by an unraveling argument. Either of these schemes can be socially optimal.




Figure 3: The sender's expected payoffs under public and private experimentation

Figure 3 illustrates the sender's expected payoffs under public experimentation and in the sender preferred equilibrium under private experimentation. The dotted curve in Figure 3 is the concave closure of the value function, as in Figure 1. Under public experimentation the sender's ex ante payoff is equal to the concave closure evaluated at the prior belief $1/2$. The dashed line below the concave closure corresponds to private experimentation. In contrast to public experimentation the sender designs an experiment yielding posteriors $\mu_1 = 0$ and $\mu_1 = \frac{U}{U+c}$ with expected posterior $\mu_1 = 1/2$. His expected payoff $\frac{1}{2}(U - c)$ is strictly lower than his expected payoff $\frac{U}{2p_d} - c$ under public experimentation, as the dashed line is strictly below the concavification of the value function for all $U/c > \frac{p_d}{1-p_d}$.

Consider next the receiver. The precision of the experiment under public experimentation is such that the receiver is just persuaded upon observing a positive outcome, i.e., the precision is relatively low. The precision of the experiment in the sender preferred equilibrium under the private scheme is higher in order to deter excessive experimentation. We have $\pi_\tau(s_1 | \omega_1) = 1$ in both cases, but $\pi_\tau(s_2 | \omega_2)$ is higher under the private scheme. The receiver benefits from a higher precision, as she makes the correct decision more often. Therefore, she is strictly better off in the sender preferred equilibrium under the private scheme. As established in Proposition 2, $\pi_\tau(s_2 | \omega_2)$ is even higher in any other equilibrium that is not Pareto dominated, where $\pi_\tau(s_1 | \omega_1) = 1$ in each of these equilibria, and, therefore, the receiver strictly prefers private to public experimentation if the players coordinate on any equilibrium that is not Pareto dominated.

8 Comparative statics

We now study the impact of an exogenous change of U/c and p_d on equilibrium behavior and payoffs under public and private experimentation.

8.1 Public experimentation

Under public experimentation the size of U/c affects whether the sender does not experiment, i.e., $U/c < 2p_d$, or whether he runs an experiment in the unique equilibrium, i.e., $U/c > 2p_d$, but it does not have an impact on the experiment's design in the latter case. Hence, for any $U/c > 2p_d$, the receiver's utility is the same. The sender's payoff increases in U/c .

An increase of p_d increases the threshold $2p_d$. The higher the threshold of doubt, the lower is the sender's incentive to run an experiment in equilibrium. An increase of p_d has an impact on the precision of the experiment if $U/c > 2p_d$. The precision in this case is $(1, \frac{2p_d-1}{p_d})$, where $\frac{2p_d-1}{p_d}$ increases in p_d . A higher threshold of doubt can only be passed if the experiment's precision goes up, which is mirrored by the increase of $\frac{2p_d-1}{p_d}$.

8.2 Private experimentation

Under private experimentation the size of U/c again determines whether the sender starts experimenting. Suppose $U/c > \frac{p_d}{1-p_d}$. In contrast to public experimentation, a change of the threshold of doubt p_d does not have an impact on the precision of an experiment that is run in any equilibrium that is not Pareto dominated, but an increase of p_d increases the threshold $\frac{p_d}{1-p_d}$.

The sender's stakes U/c have an impact on the precision of the experiment run in the sender preferred equilibrium and in the receiver preferred equilibrium, as $\frac{U-c}{U}$ and $\min\{1; 2\frac{U-c}{U}\}$ depend on U/c . We now analyze how the precision of the experiment, the persuasion probability and the decision quality in these equilibria are affected if the stakes U/c of the sender change. Our interpretation of the decision quality is motivated by the receiver's preferences. She wants to match the state with the decision. The receiver is better off (and we say that the decision quality increases) if either $\text{prob}\{a = a_2 \mid \omega = \omega_1\}$ decreases without changing $\text{prob}\{a = a_1 \mid \omega = \omega_2\}$ or $\text{prob}\{a = a_1 \mid \omega = \omega_2\}$ decreases keeping $\text{prob}\{a = a_2 \mid \omega = \omega_1\}$ constant or both decrease.

In the sender preferred equilibrium and in the receiver preferred equilibrium we have $\pi_\tau(s_1 \mid \omega_1) = 1$ irrespective of U/c . The stakes of the sender U/c exclusively affect $\pi_\tau(s_2 \mid \omega_2)$. In the sender preferred equilibrium $\pi_\tau(s_2 \mid \omega_2)$ has to increase if U/c increases in order to deter the

sender from running a second experiment if the outcome of the first experiment is adverse. In the receiver preferred equilibrium $\pi_\tau(s_2 | \omega_2)$ strictly increases in U/c for small U/c and it is 1 for all $U/c \geq 2$. In the former case $\pi_\tau(s_2 | \omega_2)$ is such that the sender is just indifferent between starting to experiment or not. If the stakes increase, then he is willing to start experimenting even if there is a lower probability to obtain a positive outcome due to a higher $\pi_\tau(s_2 | \omega_2)$.

Proposition 4 *Consider private experimentation and $U/c > \frac{p_d}{1-p_d}$. (i) The precision of the experiment that is run on the equilibrium path in the sender preferred equilibrium strictly increases in U/c . It weakly increases in U/c in the receiver preferred equilibrium. (ii) The persuasion probability in the sender preferred equilibrium strictly decreases in U/c . It weakly decreases in U/c in the receiver preferred equilibrium. (iii) The decision quality in the sender preferred equilibrium strictly increases in U/c . It weakly increases in U/c in the receiver preferred equilibrium.*

For each $\pi_\tau(s_2 | \omega_2) \in [\frac{U-c}{U}, \min\{1; 2\frac{U-c}{U}\}]$ there is an equilibrium that is not Pareto dominated in which the sender runs one experiment and where this experiment has precision $(1, \pi_\tau(s_2 | \omega_2))$ according to Proposition 2. All other equilibria are Pareto dominated. Hence, an increase of U/c affects the size and “location” of the set of equilibria that are not Pareto dominated. An increase of U/c strictly increases the lower bound $\frac{U-c}{U}$ for $\pi_\tau(s_2 | \omega_2)$ and it weakly increases the upper bound $\min\{1; 2\frac{U-c}{U}\}$. If U/c is sufficiently small, then we have $\pi_\tau(s_2 | \omega_2) \in [\frac{U-c}{U}, 2\frac{U-c}{U}]$ with $2\frac{U-c}{U} < 1$. An increase of the sender’s stakes to U'/c' implies $\pi_\tau(s_2 | \omega_2) \in [\frac{U'-c'}{U'}, \min\{1; 2\frac{U'-c'}{U'}\}]$. If the increase of the stakes is sufficiently high, i.e., if $U'/c' > \frac{1}{1-2\frac{U-c}{U}}$, then these intervals for $\pi_\tau(s_2 | \omega_2)$ are disjunct, since $U'/c' > \frac{1}{1-2\frac{U-c}{U}} \Leftrightarrow 2\frac{U-c}{U} < \frac{U'-c'}{U'}$. The following proposition directly follows.

Proposition 5 *Consider private experimentation. In any equilibrium that is not Pareto dominated a sender with low stakes $U/c \in (\frac{p_d}{1-p_d}, 2)$ runs an experiment with a strictly lower precision than a sender with high stakes $U'/c' > \frac{1}{1-2\frac{U-c}{U}}$.*

If, e.g., the receiver is interpreted as a politician and the sender as a lobby, then Propositions 4 and 5 suggest that the quality of informational lobbying and, hence, the decision quality increase in the stakes of the lobby. A similar point can be made if a researcher with career concerns is viewed as the sender and the receiver is an editor. Young researchers aspiring tenure may care more about a publication in a good journal than researchers with tenure. In order to get published in the same journal, the former may have to write higher quality papers.

9 Comparison with Felgenhauer and Schulte (2014)

The value of any evidence presented depends on the sender's experimentation and information revelation plan. Excessive private experimentation and selective information revelation reduces the value of such evidence. FS and the current paper focus on different aspects of this problem. In their setup with an exogenously fixed design of the experiments, FS show that excessive experimentation can be deterred by requiring a sufficiently high number of positive outcomes in any message that persuades the receiver. The present paper instead focuses on a reduction of excessive experimentation via an appropriate design of the experiments. By granting complete flexibility, and thus generalizing FS, the same channel for deterring too much private experimentation, i.e., requiring a high number of positive outcomes for persuasion, is available here. However, from a Pareto perspective equilibria are more attractive in which excessive experimentation is prevented by an adjustment of the precision. The sender preferred equilibrium here has a simple structure. In this equilibrium the sender runs only one experiment and this experiment is sufficiently precise such that further experimentation is just deterred if its outcome is adverse. In any other equilibrium that is not Pareto dominated the sender also runs exclusively one experiment.

Our paper compares aspects of private and public experimentation à la KG, like the players' payoffs and the persuasion probabilities. This allows us to rank the private and the public scheme from the sender's and the receiver's perspective, which is relevant, e.g., in the FDA example. There is no such comparison in FS. KG's concavification approach also relies on a flexible precision. It is possible to show its limitations under private experimentation in the current paper, but not in FS.

The comparative statics results in the current paper and in FS differ. If the sender's benefit from a favorable decision increases exogenously, then any equilibrium with persuasion in the original situation in FS eventually ceases to exist, as the value of the formerly persuasive evidence is diluted too much by a stronger incentive to experiment excessively. This incentive in FS can only be counteracted by increasing the number of outcomes required for persuasion. With a flexible design an increase of the precision of the only experiment run in the sender preferred equilibrium deters additional private experiments. The different nature of the sender preferred equilibrium in FS and the current paper has different welfare implications. In both papers an equilibrium condition is that a persuasive message is such that the receiver's threshold of doubt is passed. The sender preferred equilibrium in FS is such that the threshold of doubt is *just* passed, regardless of the sender's stakes. I.e., the number of positive outcomes that the sender has to provide for persuasion is as small as possible subject to such a message

being persuasive. Here instead, the threshold of doubt is passed by a margin in the sender preferred equilibrium. This margin increases in the stakes of the sender. Hence, in FS the stakes of the sender do not matter for the receiver's payoff, whereas here they do. In practice, big interest groups, like industrial corporations, often have better access to politicians than small groups. In a context of informational lobbying our paper suggests that the arguments provided by big interest groups are more valuable for politicians than those provided by smaller groups. This explains a higher willingness to grant them access.¹⁹ A bigger lobby in FS may have to provide more evidence for persuasion, but this evidence is not more valuable for the politician.

10 Conclusion

This paper studies a situation in which a sender tries to persuade a Bayesian receiver with evidence that stems from sequential private experimentation and that can be selectively revealed. In each equilibrium that is not Pareto dominated the sender runs a single experiment even when he cares substantially about a favorable decision. To counteract the incentive for excessive private experimentation he designs a single experiment with a sufficiently high precision. If the outcome is unfavorable, then the probability of finding favorable evidence is too low to justify the costs of further experimentation.

Our analysis suggests that the decision quality under private experimentation depends on the sender's stakes. We compare a situation where the sender has low stakes with a situation where he has high stakes and find that the sender in any equilibrium that is not Pareto dominated in the latter case provides higher quality information than in any equilibrium that is not Pareto dominated in the former case. The receiver is strictly better off the more the sender benefits from a favorable decision, as he has to provide higher quality evidence from the single experiment that he runs in order to deter further experimentation.

The sender does not benefit from the option to experiment privately. As under public experimentation he runs a single experiment, but the persuasion probability is lower. The receiver, on the other hand, enjoys an advantage from the sender's commitment problem. Since the high precision of the revealed evidence is not diluted by further private experimentation, the receiver obtains higher quality information than under public experimentation.

¹⁹There is a literature on lobbying where interested parties may buy access to a politician and only then provide information (e.g., Cotton 2012). In a straightforward extension of our paper, a big lobby should find it easier to approach a politician than a small lobby even without contributions and the lobby benefits from getting access.

Appendix

Proof of Lemma 1: (i) As shown in the main text, the optimal experiment under public experimentation has precision $(1, \frac{2p_d-1}{p_d})$. The receiver is persuaded to choose a_1 after observing outcome s_1 . Otherwise she chooses a_2 . The expected utility from running the optimal experiment, anticipating the receiver's behavior, is $\frac{1}{2}U + \frac{1}{2}(1 - \frac{2p_d-1}{p_d})U - c$. It is smaller than zero if $U/c < 2p_d$. In this case the sender does not start to experiment.

(ii) Analogous to the proof of part (i) the sender is better off with equilibrium behavior than not to experiment if $U/c \geq 2p_d$.

A positive outcome from the equilibrium experiment just persuades the receiver (given that he does not observe additional outcomes). After an adverse outcome of this experiment all players know that the state is $\omega = \omega_2$. Running a further costly experiment in this case does not change this posterior belief and is, therefore, not profitable.

Running a different first experiment (and potentially more experiments after some outcome of the first experiment) cannot yield a higher ex ante persuasion probability than with equilibrium behavior, as shown in part (iii) of the proof. Expected experimentation costs are weakly higher with such a deviation. Hence, such a deviation is not profitable.

(iii) An upper bound for the persuasion probability in any equilibrium with persuasion can be determined with the concavification approach. Let μ_1 be the receiver's posterior belief that the state is ω_1 upon observing the sender's message. The persuasion probability depending on μ_1 and its concavification are illustrated in Figure 4.




Figure 4: Upper bound for the persuasion probability

The persuasion probability in any equilibrium cannot be above the concave closure evaluated at the prior belief $1/2$, i.e., in this case the upper bound is equal to $\frac{1}{2p_d}$.

Consider a potential equilibrium in which the sender runs only one experiment on the equilibrium path and where this experiment does not have precision $(1, \frac{2p_d-1}{p_d})$. A deviation for the sender is to run a single experiment with precision $(1, \frac{2p_d-1}{p_d})$. As discussed in section 4.2, the receiver's posterior off-the-equilibrium path belief that the state is ω_1 upon observing a positive outcome s_1 of this experiment is $\frac{1}{1+(1-\frac{2p_d-1}{p_d})} = p_d$ and it is 0 upon observing an adverse outcome s_2 . I.e., she is persuaded to choose a_1 iff she observes a positive outcome of this off-the-equilibrium path experiment. The persuasion probability with such an experiment is $\frac{1}{2} * 1 + \frac{1}{2} * (1 - \frac{2p_d-1}{p_d}) = \frac{1}{2p_d}$. The experiment in the potential equilibrium has a strictly lower persuasion probability.²⁰ Therefore, the deviation is profitable.

Consider a potential equilibrium in which the sender runs multiple experiments on the equilibrium path. Suppose that the persuasion probability in this equilibrium is $\frac{1}{2p_d}$. A deviation for the sender is again to run a single experiment with precision $(1, \frac{2p_d-1}{p_d})$. The deviation also implies the persuasion probability $\frac{1}{2p_d}$. This probability is obtained with the minimum number of (costly) experiments that the sender has to run for persuasion. Therefore, the deviation is profitable. Analogously, there is no equilibrium with multiple experiments on the equilibrium path that implies a persuasion probability that is smaller than $\frac{1}{2p_d}$. Q.E.D.

Proof of Lemma 2: In the following we consider (A) an equilibrium with persuasion in which there is no on the equilibrium path history that contains an infinite number of outcomes and (B) a potential equilibrium with persuasion in which there is an on the equilibrium path history that contains an infinite number of outcomes that are not part of some persuasive message m^* and no further outcome and (C) an equilibrium with persuasion in which there is an on the equilibrium path history that contains an infinite number of outcomes that are not part of some persuasive message m^* and some further outcome(s). In contrast to (B), case (C) allows for “infinite” histories that contain some evidence that is part of some persuasive message m^* .

For the analysis of (A), (B) and (C) it is useful to note the following. The sender is better off at a posterior $\bar{\mu}_1$ if he has a stock of evidence that he can use for persuasion than if he does not have such evidence. Hence, if it is sequentially rational that the sender continues experimenting at posterior $\bar{\mu}_1$ if he does not have evidence that is part of some persuasive message $m^* \in M^*$, then it is also sequentially rational for the sender to continue experimenting at posterior $\bar{\mu}_1$ if he has a stock of evidence that he can use for persuasion. Similarly, if it is

²⁰An experiment with persuasion probability $\frac{1}{2p_d}$ has to yield posteriors $\mu_1 = 0$ and $\mu_1 = p_d$ according to Figure 4. Posterior $\mu_1 = 0$ can only be obtained with a single experiment τ if $\pi_\tau(s_1 | \omega_1) = 1$. Given $\pi_\tau(s_1 | \omega_1) = 1$, the posterior $\mu_1 = p_d$ can only be obtained if $\frac{1}{1+(1-\pi_\tau(s_2|\omega_2))} = p_d$, which is equivalent to $\pi_\tau(s_2 | \omega_2) = \frac{2p_d-1}{p_d}$.

sequentially rational for the sender to stop experimenting unsuccessfully at some posterior $\bar{\mu}_1$ if he has a stock of evidence that is part of some persuasive message $m^* \in M^*$, then it is also sequentially rational for the sender to stop experimenting at posterior $\bar{\mu}_1$ if he does not have a stock of evidence that he can use for persuasion.

(A) An equilibrium with persuasion requires that the sender stops experimenting unsuccessfully at some history that implies some posterior $\bar{\mu}_1 < 1/2$. If he instead in a hypothetical equilibrium only stops unsuccessfully if $\bar{\mu}_1 \geq 1/2$, then a message indicating unsuccessful experimentation implies that the posterior from the receiver's perspective is above $1/2$. Bayesian plausibility requires that the average posterior from the receiver's perspective after successful experimentation is below $1/2$. The receiver's posterior may depend on the revealed evidence, but at least for some persuasive $m^* \in M^*$ we have $prob\{\omega = \omega_1 \mid m^*\} < 1/2 < p_d$. But then the receiver should not be persuaded by such a message yielding a contradiction.

If the sender stops unsuccessfully at some history that implies some posterior $\bar{\mu}_1 < 1/2$ and that contains some evidence that is part of some persuasive message, then he also stops experimenting unsuccessfully at a history that implies the same posterior $\bar{\mu}_1$ and that does not contain any evidence that is part of some persuasive message, as shown above. Therefore, the sender stops experimenting unsuccessfully at some posterior $\bar{\mu}_1 < 1/2$ if he has not yet found an outcome that is part of some $m^* \in M^*$ with which he can persuade the receiver to choose $a = a_1$ in such an equilibrium.

(B) Consider a potential equilibrium with persuasion in which there is an on the equilibrium path history that contains an infinite number of outcomes that are not part of some persuasive message m^* and no further outcome. If the sender does not stop experimenting unsuccessfully at any history h with $\bar{\mu}_1 < 1/2$ if he has not yet found evidence that can be used for persuasion, then he never stops unsuccessfully at any history with $\bar{\mu}_1 < 1/2$ regardless of the stock of evidence, as shown above. In the next paragraph we show that it is a zero probability event in such a potential equilibrium that the sender runs an infinite number of experiments without finding persuasive evidence. Stopping unsuccessfully may, thus, only occur if the sender faces some history with $\bar{\mu}_1 > 1/2$. Analogous to (A) there cannot be such an equilibrium, as this implies via Bayesian plausibility that the posterior upon observing some persuasive message is below $1/2$, which contradicts that the receiver is persuaded upon observing such a message.

It remains to show that in such a potential equilibrium it is a zero probability event that the sender runs an infinite number of experiments without finding persuasive evidence. At each history h where a persuasive message is not feasible and where he runs an additional experiment according to the equilibrium strategy, sequential rationality requires that the probability to obtain persuasive evidence by future experimentation times U exceeds the costs c of the

next experiment. A necessary condition for sequential rationality is, thus, that this probability at each such history h is weakly above $\frac{c}{\bar{U}}$, with $\frac{c}{\bar{U}} > 0$. As the sender according to the potential equilibrium continues experimenting at any such history h with $\bar{\mu}_1 < 1/2$, the probability that he runs an infinite number of experiments without finding persuasive evidence is zero.

(C) Analogous to (B) it is a zero probability event that the sender runs an infinite number of experiments without finding persuasive evidence. Analogous to (A) the sender stops experimenting unsuccessfully at some posterior $\bar{\mu}_1 < 1/2$ if he has not yet found an outcome that is part of some $m^* \in M^*$ with which he can persuade the receiver to choose $a = a_1$ in such an equilibrium. Q.E.D.

Proof of Lemma 3: (i) There is no equilibrium with persuasion if $U/c < 2p_d$, as it does not pay to start experimenting even if the sender could persuade with the minimum number of experiments required for persuasion and the maximum persuasion probability.

(ii) Consider a potential equilibrium with the following properties:

(1) The sender's first experiment has precision $\pi_\tau = (1, \frac{2p_d-1}{p_d})$.

(2) The sender stops experimenting (successfully) at any history h that contains at least one outcome $\sigma_\tau = s_1$ of an experiment with $\pi_\tau = (1, \frac{2p_d-1}{p_d})$.

(3) The sender stops experimenting (unsuccessfully) at any history h with $\text{prob}\{\omega = \omega_1 \mid h\} = 0$ that does not contain at least one outcome $\sigma_\tau = s_1$ of an experiment with $\pi_\tau = (1, \frac{2p_d-1}{p_d})$.

(4) The sender sends message $m^* = (\sigma_\tau, \pi_\tau)$ with $\pi_\tau = (1, \frac{2p_d-1}{p_d})$ for any outcome σ_τ that he observes after the first experiment.

(5) The sender's remaining behavior off-the-equilibrium path is sequentially rational.

(6) The receiver chooses $a = a_1$ if the sender sends message $m^* = (\sigma_\tau = s_1, \pi_\tau)$ with $\pi_\tau = (1, \frac{2p_d-1}{p_d})$ of an experiment τ . The receiver also chooses $a = a_1$ if the sender sends a message m that contains at least one positive outcome from an experiment with precision $(q, 1)$, with $q \in (0, 1]$. Otherwise she chooses a_2 .

(7) On the equilibrium path, i.e., upon observing message $m^* = (\sigma_\tau, \pi_\tau)$ with $\pi_\tau = (1, \frac{2p_d-1}{p_d})$, beliefs are formed in accordance with Bayes' Law. Upon observing an off-the-equilibrium path message that does not contain at least one positive outcome from an experiment with precision $(q, 1)$, the receiver forms a probability assessment over experimentation histories such that the probability that $\omega = \omega_1$ conditional on this assessment is below the threshold of doubt. E.g., she may believe that the sender privately ran a single additional experiment with precision $(1, 1)$ that yielded an adverse outcome. Off-the-equilibrium path beliefs upon observing a message that contains at least one positive outcome from an experi-

ment with precision $(q, 1)$, with $q \in (0, 1]$, are such that the receiver thinks that $\omega = \omega_1$ with certainty, as discussed in section 4.1.3.

We now show that (1) - (7) are consistent with an equilibrium. (5) and (7) are by construction consistent with perfect Bayesian equilibrium.

Sender behavior: Consider the sender's message. If the sender stops experimenting and the stock of collected experimental outcomes contains an outcome that induces action a_1 if the appropriate message is sent, then he sends such a message as a_1 is the preferred action. If the sender stops experimenting unsuccessfully, then he induces action a_2 with any feasible message. Therefore, sending any feasible message is optimal after unsuccessful experimentation.

We now have to check the sender's equilibrium experimentation behavior described in properties (1) - (3).

(A) Consider a deviation from equilibrium property (1). Analogous to public experimentation, the sender is better off to start experimenting with an experiment with $\pi_\tau = (1, \frac{2p_d-1}{p_d})$ than not to experiment as $U/c \geq 2p_d$. We now have to check whether the sender wants to start experimenting with an experiment that does not have precision $(1, \frac{2p_d-1}{p_d})$. We distinguish the cases (A.a) where he deviates by starting with an experiment with precision $(q, 1)$, with $q \in (0, 1]$, and (A.b) where he deviates by starting with an experiment with a precision that differs from $(1, \frac{2p_d-1}{p_d})$ and $(q, 1)$.

(A.a) The sender does not have an incentive to start experimenting with an experiment with precision $(q, 1)$, with $q \in (0, 1]$: None of these experiments yields a positive outcome in ω_2 . Running one such experiment yields a lower persuasion probability and the same experimentation costs as equilibrium behavior and is, therefore, not profitable. If he runs more than one experiment with a positive probability, then the ex ante upper bound for the persuasion probability in state ω_1 if he exclusively runs these experiments is 1 and it is 0 in state ω_2 . If the sender later (sometimes) runs experiments with $\pi_\tau = (1, \frac{2p_d-1}{p_d})$, then it is sequentially rational to stop after each realization, as implied by equilibrium properties (2) and (3). In state ω_2 , the ex ante upper bound for the persuasion probability can then be increased from 0 to $(1 - \frac{2p_d-1}{p_d})$. Overall the persuasion probability is, therefore, weakly below 1 in state ω_1 and weakly below $(1 - \frac{2p_d-1}{p_d})$ in state ω_2 . Therefore, the persuasion probability is weakly below the persuasion probability if the sender starts with his on the equilibrium path experiment (where it is 1 in state ω_1 and $(1 - \frac{2p_d-1}{p_d})$ in state ω_2), but expected experimentation costs are above c . Hence, this deviation is not profitable.

(A.b) The sender does not have an incentive to start experimenting with an experiment with a precision that differs from $(1, \frac{2p_d-1}{p_d})$ and $(q, 1)$: Due to off-the-equilibrium path beliefs, any outcome from such an experiment cannot be used for persuasion. The sender, therefore,

does not have an incentive to start experimenting with such an experiment for the purpose of inducing action a_1 with this experiment. He may have an incentive to run such an experiment in order to learn about the state. Learning can only be potentially profitable if the sender later runs an experiment with precision $(1, \frac{2p_d-1}{p_d})$ and / or $(q, 1)$ with a positive probability. If he starts experimenting with an off-the-equilibrium path experiment, then the ex ante expected experimentation costs of a potentially profitable deviation are, therefore, greater than c . Ex ante the persuasion probability if the sender at some time after the first experiment runs experiments with precision $(1, \frac{2p_d-1}{p_d})$ and / or $(q, 1)$ with a positive probability and if he behaves sequentially rational is weakly below 1 in state ω_1 and weakly below $(1 - \frac{2p_d-1}{p_d})$ in state ω_2 , analogous to the argument in (A.a). Due to the lower persuasion probability and the higher ex ante costs, the deviation is not profitable.

(B) Consider a deviation from equilibrium property (2). A deviation is not profitable, since sending a message that exclusively contains $\sigma_\tau = s_1$ and π_τ induces his preferred action a_1 .

(C) Consider a deviation from equilibrium property (3). For any history h' for which h is a subhistory we have $prob\{\omega = \omega_1 \mid h'\} = 0$ and according to equilibrium properties (2) and (3) the sender stops experimenting for each such history h' . For sequential rationality we have to check whether the sender can improve with a deviation at history h , without changing the remaining parts of the equilibrium, as discussed in section 4.1.1. We distinguish the cases (C.a) where he deviates by continuing experimentation at h with an experiment that does not have precision $(1, \frac{2p_d-1}{p_d})$ and (C.b) where he deviates by continuing experimentation at h with an experiment that has precision $(1, \frac{2p_d-1}{p_d})$.

(C.a) Consider a deviation at h in which the sender continues experimenting with a next experiment that does not have precision $\pi_\tau = (1, \frac{2p_d-1}{p_d})$. An experiment that does not have precision $(1, \frac{2p_d-1}{p_d})$ or precision $(q, 1)$ cannot be used for persuasion due to off-the-equilibrium-path beliefs (described in equilibrium property (7)) and it also does not contain new information about the state, since the sender knows that $\omega = \omega_2$. Since running it is costly, the deviation is not profitable. An experiment that has precision $(q, 1)$ instead does not yield a positive outcome in state ω_2 and, since it is costly, a deviation where this experiment is run is also not profitable.

(C.b) Consider a deviation at h in which the sender continues experimenting with a next experiment that has precision $\pi_\tau = (1, \frac{2p_d-1}{p_d})$. According to property (2) of this equilibrium he stops experimenting successfully if this experiment yields outcome s_1 . According to property (3) he stops experimenting unsuccessfully if it yields outcome s_2 , since the sender still knows that the state is ω_2 . His continuation utility from running this experiment is $(1 - \frac{2p_d-1}{p_d})U - c$, as $prob\{\omega = \omega_1 \mid h\} = 0$ at h . Continuing experimentation is not optimal if $(1 - \frac{2p_d-1}{p_d})U - c \leq$

$0 \Leftrightarrow U/c \leq \frac{p_d}{1-p_d}$, which is satisfied.

Receiver behavior:

The receiver infers that the sender ran a single experiment if she observes a message $m^* = (\sigma_\tau, \pi_\tau)$, with $\pi_\tau = (1, \frac{2p_d-1}{p_d})$. She is persuaded by a message m^* containing a positive outcome from an experiment with precision $\pi_\tau = (1, \frac{2p_d-1}{p_d})$ if $\frac{1}{1+(1-\frac{2p_d-1}{p_d})} \geq p_d$, which is satisfied.

(iii) Such a potential equilibrium requires that the sender stops experimenting after observing an adverse outcome s_2 of an experiment with precision $\pi_\tau = (1, \frac{2p_d-1}{p_d})$. Analogous to part (ii) of the proof he has an incentive to deviate. Q.E.D.

Proof of Lemma 4: We cannot have $V_1(\frac{1}{2}) \geq 0$ and $V_2(\frac{1}{2}) > 0$. In this case the sender could (possibly suboptimally) at each posterior $\bar{\mu}_1$ replicate the same behavior as at the prior where he has not yet run an experiment, yielding $V_1(\frac{1}{2})$ and $V_2(\frac{1}{2})$ in the respective states with this modified plan at posterior $\bar{\mu}_1$.²¹ As $\bar{\mu}_1$ only allocates probability mass to $V_1(\frac{1}{2})$ and to $V_2(\frac{1}{2})$ in $\bar{\mu}_1 V_1(\frac{1}{2}) + (1 - \bar{\mu}_1) V_2(\frac{1}{2})$ (and both are greater than zero), the sender's benefit from continuing to experiment with the modified plan is greater than zero. With the optimal plan he is weakly better off. Therefore, he would not stop experimenting unsuccessfully for any posterior given that he has not yet acquired an outcome that he can use for persuasion. This implies that he never stops experimenting unsuccessfully for all posteriors $\bar{\mu}_1 < 1/2$, which contradicts the notion of an equilibrium with persuasion due to Lemma 2.

We cannot have $V_1(\frac{1}{2}) \leq 0$ and $V_2(\frac{1}{2}) > 0$. In this case the sender could (possibly suboptimally) at each posterior $\bar{\mu}_1 < 1/2$ replicate the same behavior as at the prior where he has not yet run an experiment. In this case for any $\bar{\mu}_1 < 1/2$, more probability mass would be allocated to $V_2(\frac{1}{2})$ than ex ante and less probability mass would be allocated to $V_1(\frac{1}{2})$ than ex ante. I.e., if $\frac{1}{2} V_1(\frac{1}{2}) + \frac{1}{2} V_2(\frac{1}{2}) \geq 0$, then we have $\bar{\mu}_1 V_1(\frac{1}{2}) + (1 - \bar{\mu}_1) V_2(\frac{1}{2}) > 0$ for all $\bar{\mu}_1 < 1/2$. Optimal behavior implies a weakly greater utility. This implies that he never stops

²¹ "Replicating" behavior means the following. Consider a history $h = \{(\sigma_\tau, \pi_\tau)\}_{\tau=1, \dots, t}$, where the sender has not yet found evidence that is part of some persuasive message and where he faces posterior $\bar{\mu}_1 < 1/2$. Consider further a history $h' = \{(\sigma_\tau, \pi_\tau)\}_{\tau=1, \dots, t, \dots, t'}$, with $t' \geq t$. Based on these histories construct an artificial history h'' as follows: Each (σ_τ, π_τ) in h'' is equal to $(\sigma_{t+\tau}, \pi_{t+\tau})$ in h' , with $\tau = 1, \dots, z$ and $z = t' - t$.

The sender's experimentation plan specifies for each history h (i) whether the sender continues or stops experimenting and (ii) the precision of the next experiment if he continues experimenting. We say that the sender from history h on replicates the same behavior as at the prior where he has not yet found evidence that can be used for persuasion, if the sender facing history h runs the next experiment with precision π_1 and at each history $h' = \{(\sigma_\tau, \pi_\tau)\}_{\tau=1, \dots, t, \dots, t'}$, with $t' > t$, he continues experimenting according to his experimentation plan, as if he faces history h'' instead of history h' .

experimenting unsuccessfully for all posteriors $\bar{\mu}_1 < 1/2$, which contradicts the notion of an equilibrium with persuasion due to Lemma 2.

Therefore, $\frac{1}{2}V_1(\frac{1}{2}) + \frac{1}{2}V_2(\frac{1}{2}) \geq 0$ requires $V_1(\frac{1}{2}) \geq 0$ and $V_2(\frac{1}{2}) \leq 0$. Q.E.D.

Proof of Proposition 1: (i) Consider a potential equilibrium with properties (1) - (7) described in the main text. (1) - (7) are consistent with an equilibrium analogous to the proof of Lemma 3 (ii).²²

(ii) Call the equilibrium in which the sender on the equilibrium path runs exactly one experiment and this experiment has precision $(1, \frac{U-c}{U})$ “equilibrium SP” (where SP stands for “sender preferred”). The sender’s ex ante utility in such an equilibrium is $\frac{1}{2}(U - c) + \frac{1}{2}((1 - \frac{U-c}{U})U - c)$, which is equivalent to

$$(2) \quad \frac{1}{2}(U - c).$$

We show that equilibrium SP is sender preferred if $U/c \in [\frac{p_d}{1-p_d}, \infty)$. In an equilibrium the sender sends a message that contains at least n outcomes on the equilibrium path. The ex ante expected utility in such an equilibrium is at most

$$(3) \quad \frac{1}{2}(U - nc) + \frac{1}{2}V_2\left(\frac{1}{2}\right)$$

because, if $\omega = \omega_1$, then the sender gets his highest utility if he finds an outcome that can be used for persuasion in each experiment and if he can persuade the receiver after n experiments. Note that we have $V_2(\frac{1}{2}) \leq 0$ in (3) due to Lemma 4.

We now show that the sender is strictly better off in equilibrium SP than in any other equilibrium with persuasion.

(A) If all messages $m^* \in M^*$ that are sent on the equilibrium path contain more than one outcome, then the equilibrium cannot be sender preferred as $n \geq 2$ and $V_2(\frac{1}{2}) \leq 0$ according to Lemma 4.

²²Suppose the sender observes s_2 of the first experiment (which has precision $\pi_\tau = (1, \frac{U-c}{U})$). The sender’s continuation utility from running a further experiment with this precision and then to stop after each realization of this experiment is $(1 - \frac{U-c}{U})U - c = 0$. Therefore, it is optimal not to run this experiment. It is also not optimal to run any other experiment. Analogous to the proof of Lemma 3 (ii), he does not have an incentive to start with an experiment with a different precision. As the sender is indifferent to continue experimenting with an experiment with precision $\pi_\tau = (1, \frac{U-c}{U})$ if he knows $\omega = \omega_2$, he has a strict incentive to start experimenting at the prior belief.

The receiver knows that the sender in equilibrium runs only one experiment, where this experiment has precision $\pi_\tau = (1, \frac{U-c}{U})$. Upon observing an on the equilibrium path message containing a positive outcome, she is persuaded to choose a_1 , as $\frac{1}{1+(1-\frac{U-c}{U})} \geq p_d$ due to $U/c > \frac{p_d}{1-p_d}$.

(B) There can be equilibria in which on the equilibrium path messages are sent that contain only one outcome. For $n = 1$ the utility in (3) is weakly smaller than the utility in (2). We show now that equilibrium SP is also strictly preferred by the sender to these equilibria. We distinguish the following cases.

(a) Suppose an outcome of the first experiment that the sender runs is not an element of some m^* with $n = 1$. The sender needs to run a further experiment and $V_1(\frac{1}{2}) < (U - c)$. Thus, the sender prefers equilibrium SP.

(b) Suppose an adverse outcome of the first experiment is an element of a message m^* with $n = 1$, but a positive outcome is not an element of some message m^* with $n = 1$.

Suppose the sender on the equilibrium path runs only one experiment. By assumption $\pi_\tau(s_1 | \omega_1) \geq \pi_\tau(s_1 | \omega_2)$ for all experiments τ , which implies $prob\{\omega = \omega_1 | \sigma_\tau = s_2\} < 1/2 < p_d$. I.e., in such a hypothetical equilibrium the receiver should not be persuaded upon observing a message that contains an adverse outcome and no other outcome. If the receiver cannot be persuaded by such a message, then it cannot be equilibrium behavior for the sender to run (only) the corresponding costly experiment, as not experimenting yields a higher payoff.

Suppose the sender on the equilibrium path runs multiple experiments. In this case $V_1(\frac{1}{2}) < (U - c)$. Thus, the sender prefers equilibrium SP.

(c) Suppose a positive outcome of the first experiment is an element of a message m^* with $n = 1$ that is sent with positive probability on the equilibrium path. We now derive the maximum ex ante utility of the sender in such a potential equilibrium and properties of such a potential equilibrium. We then show that only equilibrium SP yields this utility.

Consider a potential equilibrium in which the precision of the first experiment is such that the receiver is persuaded if the sender presents a positive outcome of this experiment. The maximum $V_1(\frac{1}{2})$ that can potentially be achieved is $(U - c)$, since if $\omega = \omega_1$ payoff $V_1(\frac{1}{2}) = (U - c)$ implies persuasion with probability 1 with the minimum number of experiments required for persuasion. $V_1(\frac{1}{2}) = (U - c)$ can be achieved in a potential equilibrium if and only if $\pi_\tau(s_1 | \omega_1) = 1$. Thus, in such an equilibrium the posterior is $\bar{\mu}_1 = 0$ if the first experimental outcome is adverse.

In such an equilibrium we cannot have that the sender continues running further experiments if he knows that $\bar{\mu}_1 = 0$ after observing an adverse outcome of the first experiment: If the sender instead continues experimenting in equilibrium if the first outcome is adverse knowing $\bar{\mu}_1 = 0$, then he would not stop experimenting unsuccessfully in the following.²³ He

²³Any following experiment does not change the posterior $\bar{\mu}_1 = 0$, but at some later experiment the sender may have found some evidence that he can use for persuasion. Hence, the sender in the following cannot be worse off than in case the first outcome is adverse.

would, thus, run experiments until he finds some persuasive $m^* \in M^*$ regardless of the state ω . But then the receiver should not be persuaded by such an m^* yielding a contradiction.

As the sender stops experimenting if the first experiment yields an adverse outcome in a hypothetical equilibrium with $V_1(\frac{1}{2}) = (U - c)$ the sender runs a single experiment on the equilibrium path. This determines the structure of $V_2(\frac{1}{2})$ in such an equilibrium. As only one experiment is run on the equilibrium path, denoting this experiment τ , we have $V_2(\frac{1}{2}) = (1 - \pi_\tau(s_2 | \omega_2))U - c$. The maximum $V_2(\frac{1}{2})$ that can potentially be achieved in equilibrium is 0, as $V_2(\frac{1}{2}) \leq 0$ according to Lemma 4. Maximizing $V_2(\frac{1}{2}) = (1 - \pi_\tau(s_2 | \omega_2))U - c$ with respect to $\pi_\tau(s_2 | \omega_2)$ subject to the constraint $V_2(\frac{1}{2}) \leq 0$ yields $\pi_\tau(s_2 | \omega_2) = \frac{U-c}{U}$ which implies that $(1 - \pi_\tau(s_2 | \omega_2))U - c = 0$. It follows that in the potential equilibrium which yields $V_1(\frac{1}{2}) = (U - c)$, we may also achieve $V_2(\frac{1}{2}) = 0$. It is, therefore, established that $\pi_\tau(s_1 | \omega_1) = 1$ is the only possibility to have $V_1(\frac{1}{2}) = (U - c)$ and, given that $\pi_\tau(s_1 | \omega_1) = 1$, $\pi_\tau(s_2 | \omega_2) = \frac{U-c}{U}$ is the only precision in state ω_2 that potentially yields $V_2(\frac{1}{2}) = 0$ in equilibrium. It follows that $V_1(\frac{1}{2}) = (U - c)$ and $V_2(\frac{1}{2}) = 0$ can only be achieved if a single experiment is run on the equilibrium path with precision $(1, \frac{U-c}{U})$. An equilibrium in which the sender on the equilibrium path runs a single experiment and this experiment has precision $(1, \frac{U-c}{U})$ corresponds to equilibrium SP.

(iii) The persuasion probability under public experimentation is $\frac{1}{2p_d}$. The persuasion probability under private experimentation is $\frac{U+c}{2U}$, which is $\frac{U+c}{2U} = \frac{1}{2p_d}$ if $U/c = \frac{p_d}{1-p_d}$. As $\frac{U+c}{2U}$ decreases in U/c , the persuasion probability is smaller if $U/c > \frac{p_d}{1-p_d}$. Q.E.D.

Proof of Proposition 2: (i) The proof for existence is analogous to the proof of Lemma 3 (ii).

(ii) We first show that each equilibrium in the class of equilibria that is not described in (i) and where the sender runs only one experiment is Pareto dominated by some equilibrium described in (i). Then we study equilibria in which the sender runs multiple experiments.

(1) Consider a potential equilibrium in which the sender runs only one experiment and where this equilibrium is not described in (i). In such an equilibrium the sender runs an experiment with some precision (q, p) , with $q \in [0, 1]$ and $p \in [0, 1]$. In the following we distinguish the three cases $p \in [\frac{U-c}{U}, \min\{1, 2\frac{U-c}{U}\}]$, $p \in [\min\{1, 2\frac{U-c}{U}\}, 1]$ and $p < \frac{U-c}{U}$.

Suppose $p \in [\frac{U-c}{U}, \min\{1, 2\frac{U-c}{U}\}]$, i.e., there is an equilibrium as described in part (i) of the proposition with $\pi_\tau(s_2 | \omega_2) = p$. We compare the latter equilibrium, where $\pi_\tau(s_1 | \omega_1) = 1$, with an equilibrium where $q < \pi_\tau(s_1 | \omega_1)$. Since we assume that a positive outcome is more likely in state $\omega = \omega_1$ than in state $\omega = \omega_2$, the receiver's decision rule in both equilibria is the same. She chooses $a = a_1$ only if she observes a positive outcome. The receiver's utility in

state $\omega = \omega_1$ increases in q , since she then observes the appropriate positive outcome with a higher probability. In state $\omega = \omega_2$ her utility is not affected by q . Hence, she is better off in an equilibrium with a higher q . The sender's utility in state $\omega = \omega_1$ increases in q , since the receiver's decision rule in both equilibria is the same and he obtains a positive outcome with a higher probability if q increases. The sender's utility in state $\omega = \omega_2$ does not depend on q . Hence, he is better off in an equilibrium with a higher q . Hence, any such equilibrium with precision $(q, p = \pi_\tau(s_2 | \omega_2))$, with $q < 1$, is Pareto dominated by an equilibrium in which the sender runs only one experiment and where the experiment has precision $(1, \pi_\tau(s_2 | \omega_2))$.

Suppose $p \in (\min\{1; 2\frac{U-c}{U}\}, 1]$. Consider $\min\{1; 2\frac{U-c}{U}\} < 1$, i.e., parameters for which there is no equilibrium as described in part (i) of the proposition with $\pi_\tau(s_2 | \omega_2) = p$. In this case there is no equilibrium with persuasion in which the only experiment run in equilibrium has precision (q, p) : An equilibrium condition is that the sender starts experimenting with such an experiment. The ex ante payoff from running such an experiment is $\frac{1}{2}(qU - c) + \frac{1}{2}((1-p)U - c)$. The sender by construction is indifferent to start experimenting with an experiment with precision $(1, 2\frac{U-c}{U})$ in this parameter range. Thus, he is strictly better off not to run an experiment if $q \leq 1$ and $p > 2\frac{U-c}{U}$, since such an experiment implies a lower probability of an outcome inducing $a = a_1$.

Suppose $p < \frac{U-c}{U}$, i.e., there is no equilibrium as described in part (i) of the proposition with $\pi_\tau(s_2 | \omega_2) = p$. Analogous to above the receiver's utility increases in q and p . Therefore, the receiver's payoff in any such equilibrium is lower than in the sender preferred equilibrium, where the sender runs only one experiment and this experiment has precision $(1, \frac{U-c}{U})$. The sender preferred equilibrium by definition is better for the sender than any other equilibrium. Therefore, the sender's payoff in the equilibrium where he runs only one experiment and where this experiment has precision $(1, \frac{U-c}{U})$ is also higher. Hence, any such equilibrium with $p < \frac{U-c}{U}$ is Pareto dominated by the sender preferred equilibrium, where the sender runs only one experiment and where this experiment has precision $(1, \frac{U-c}{U})$.

(2) Consider an equilibrium in which the sender runs multiple experiments on the equilibrium path.

Given the sender's and the receiver's equilibrium strategy, we can determine $prob\{a = a_1 | \omega = \omega_1\}$ and $prob\{a = a_2 | \omega = \omega_2\}$. Construct a single experiment with precision (q, p) , with $q = prob\{a = a_1 | \omega = \omega_1\}$ and $p = prob\{a = a_2 | \omega = \omega_2\}$. Suppose (possibly suboptimally) that the sender runs exclusively this experiment and reveals the outcome after each realization and that the receiver chooses $a = a_1$ only after a positive outcome of this experiment. Denote this as (auxiliary) situation A. The receiver by construction is ex ante as well off in situation A as in the equilibrium with the multiple experiments. The persuasion

probability in the equilibrium and in situation A by construction are the same. Expected experimentation costs for the sender if he runs a single experiment are lower. Hence, the sender is better off in situation A with the single experiment than in the equilibrium with multiple experiments. In the following we again distinguish the three cases $p \in [\frac{U-c}{U}, \min\{1; 2\frac{U-c}{U}\}]$, $p \in [\min\{1; 2\frac{U-c}{U}\}, 1]$ and $p < \frac{U-c}{U}$.

Suppose $p \in [\frac{U-c}{U}, \min\{1; 2\frac{U-c}{U}\}]$, i.e., there is an equilibrium as described in part (i) of the proposition with $\pi_\tau(s_2 | \omega_2) = p$. Situation A yields a lower payoff for the sender than the equilibrium in which he runs one experiment and where this experiment has precision $(1, p)$, since the persuasion probability in the latter is higher and the experimentation costs are the same. As the sender's payoff in situation A is higher than in the equilibrium with the multiple experiments, the equilibrium in which he runs one experiment and this experiment has precision $(1, \pi_\tau(s_2 | \omega_2) = p)$ also yields a higher ex ante payoff than the equilibrium with multiple experiments. By construction the receiver is equally well off in situation A as in the equilibrium with multiple experiments. As the receiver's ex ante payoff increases in q , she is better off in the equilibrium in which the sender runs one experiment and this experiment has precision $(1, \pi_\tau(s_2 | \omega_2) = p)$ than in situation A. Hence, she is also better off in the equilibrium in which the sender runs only one experiment and this experiment has precision $(1, \pi_\tau(s_2 | \omega_2) = p)$ than in the equilibrium in which he runs multiple experiments. Hence, the equilibrium in which the sender runs only one experiment and where this experiment has precision $(1, \pi_\tau(s_2 | \omega_2) = p)$ Pareto dominates the equilibrium with multiple experiments.

Suppose $p \in (\min\{1; 2\frac{U-c}{U}\}, 1]$ with $\min\{1; 2\frac{U-c}{U}\} < 1$. Analogous to above, there can be no equilibrium in which the sender runs multiple experiments, as the equilibrium condition that he starts experimenting is violated in this case.

Suppose $p < \frac{U-c}{U}$, i.e., there is no equilibrium as described in part (i) of the proposition with $\pi_\tau(s_2 | \omega_2) = p$. The equilibrium in which the sender runs only one experiment and this experiment has precision $(1, \frac{U-c}{U})$ is better for the sender than the equilibrium in which he runs multiple experiments, as the former is globally sender preferred. The receiver prefers the equilibrium in which the sender runs one experiment that has precision $(1, \frac{U-c}{U})$ to situation A, as $q \leq 1$ and $p < \frac{U-c}{U}$ and her payoff increases in p and q . As situation A is equally good for the receiver as the equilibrium with multiple experiments, the receiver prefers the sender preferred equilibrium to the equilibrium in which the sender runs multiple experiments. Hence, the sender preferred equilibrium in which the sender runs only one experiment and where this experiment has precision $(1, \frac{U-c}{U})$ Pareto dominates the equilibrium with multiple experiments if $p < \frac{U-c}{U}$.

(iii) The receiver's ex ante payoff increases in $\pi_\tau(s_2 | \omega_2)$, but the sender's ex ante payoff

decreases in $\pi_\tau(s_2 | \omega_2)$. Therefore, we cannot rank the equilibria described in part (i) of the proposition according to the Pareto criterion. Q.E.D.

Proof of Proposition 4: (i) Sender preferred equilibrium: The precision of the experiment run in this equilibrium is $(1, \frac{U-c}{U})$, where $\frac{U-c}{U}$ strictly increases in U/c .

Receiver preferred equilibrium: The precision of the experiment run in this equilibrium is $(1, \min\{1; 2\frac{U-c}{U}\})$, where $\min\{1; 2\frac{U-c}{U}\}$ weakly increases in U/c .

(ii) Sender preferred equilibrium: As $\pi_\tau(s_1 | \omega_1) = 1$ does not change and $\pi_\tau(s_2 | \omega_2)$ strictly increases in U/c according to (i), it follows that the persuasion probability strictly decreases in U/c .

Receiver preferred equilibrium: As $\pi_\tau(s_1 | \omega_1) = 1$ does not change and $\pi_\tau(s_2 | \omega_2)$ weakly increases in U/c according to (i), it follows that the persuasion probability weakly decreases in U/c .

(iii) Sender preferred equilibrium: As $\pi_\tau(s_1 | \omega_1) = 1$ does not change and $\pi_\tau(s_2 | \omega_2)$ strictly increases in U/c according to (i), the statement directly follows.

Receiver preferred equilibrium: As $\pi_\tau(s_1 | \omega_1) = 1$ does not change and $\pi_\tau(s_2 | \omega_2)$ weakly increases in U/c according to (i), the statement directly follows. Q.E.D.

REFERENCES

- Aghion, P., P. Bolton, C. Harris and B. Jullien, "Optimal Learning by Experimentation," *The Review of Economic Studies* 58 (1991), 621-654.
- Aumann, R. J. and M. B. Maschler, *Repeated Games with Incomplete Information* (Cambridge, MA: MIT Press, 1995).
- Baliga, S. and J. C. Ely, "Torture and the Commitment Problem," *The Review of Economic Studies* forthcoming (2016).
- Bergemann, D. and J. Välimäki, "Bandit Problems," in S. N. Durlauf and L. E. Blume, eds., *The New Palgrave Dictionary of Economics*, 2nd ed. (Basingstoke: Palgrave Macmillan, 2008).
- Bolton, P. and C. Harris, "Strategic Experimentation," *Econometrica* 67 (1999), 349-374.
- Brocas, I. and J. D. Carillo, "Influence through Ignorance," *The RAND Journal of Economics* 38 (2008), 931-947.
- Celik, G., "Interested Experts: Do They Know More?," Mimeo, University of British Columbia, 2003.
- Cotton, C., "Pay-to-Play Politics: Informational Lobbying and Contribution Limits when Money Buys Access," *Journal of Public Economics* 96 (2012), 369-386.

- Felgenhauer, M. and E. Schulte, "Strategic Private Experimentation," *American Economic Journal: Microeconomics* 6 (2014), 74-105.
- Felgenhauer, M., "Endogenous Persuasion with Rational Inattention," Mimeo, Plymouth University, 2016.
- Gentzkow, M. and E. Kamenica, "Disclosure of Endogenous Information," *Economic Theory Bulletin* forthcoming (2016).
- Gentzkow, M. and E. Kamenica, "Costly Persuasion," *American Economic Review Papers and Proceedings* 104 (2014), 457-462.
- Glazer, J. and A. Rubinstein, "Debates and Decisions: On a Rationale of Argumentation Rules," *Games and Economic Behavior* 36 (2001), 158-173.
- Glazer, J. and A. Rubinstein, "On Optimal Rules of Persuasion," *Econometrica* 72 (2004), 1715-1736.
- Glazer, J. and A. Rubinstein, "A Study in the Pragmatics of Persuasion: A Game Theoretical Approach," *Theoretical Economics* 1 (2006), 395-410.
- Henry, E., "Strategic Disclosure of Research Results: The Cost of Proving Your Honesty," *The Economic Journal* 119 (2009), 1036-1064.
- Hopenhayn, H. and F. Squintani, "Preemption Games with Private Information," *The Review of Economic Studies* 78 (2011), 667-692.
- Jovanovic, B., "Truthful Disclosure of Information," *Bell Journal of Economics* 13 (1982), 36-44.
- Kamenica, E. and M. Gentzkow, "Bayesian Persuasion," *American Economic Review* 101 (2011), 2590-2615.
- Keller, G., S. Rady and M. Cripps, "Strategic Experimentation with Exponential Bandits," *Econometrica* 73 (2005), 39-68.
- Milgrom, P. and J. Roberts, "Relying on the Information of Interested Parties," *The RAND Journal of Economics* 17 (1986), 18-32.
- Moscarini, G. and L. Smith, "The Optimal Level of Experimentation," *Econometrica* 69 (2001), 1629-1644.
- Rosenberg, D., E. Solan and N. Vieille, "Social Learning in One-Arm Bandit Problems," *Econometrica* 75 (2007), 1591-1611.
- Rothschild, M., "A Two-Armed Bandit Theory of Market Pricing," *Journal of Economic Theory* 9 (1974), 185-202.

This is the accepted version of the following article:

“Bayesian Persuasion with Private Experimentation”, Felgenhauer, M. and P. Loerke, *International Economic Review*, Vol. 58, Issue 3, copyright © 2017, pp. 829-856, Wiley.

which has been published in final form at

[<http://onlinelibrary.wiley.com/doi/10.1111/iere.12237/abstract>].

This article may be used for non-commercial purposes in accordance with the Wiley Self-Archiving Policy [<http://olabout.wiley.com/WileyCDA/Section/id-828039.html>].