

Costly Persuasion by a Partially Informed Sender

SHAOFEI JIANG*

February 12, 2022

I study a model of costly Bayesian persuasion by a privately and partially informed sender who conducts a public experiment. I microfound the cost of an experiment via a Wald's sequential sampling problem and show that it equals the expected reduction in a weighted log-likelihood ratio function evaluated at the sender's belief. I focus on equilibria satisfying the D1 criterion. The equilibrium outcome depends on the relative costs of drawing good and bad news in the experiment. If bad news is more costly, there exists a unique separating equilibrium outcome, and the receiver unambiguously benefits from the sender's private information. If good news is sufficiently more costly, the single-crossing property fails. There exists a continuum of pooling equilibria, and the receiver strictly suffers from sender private information in some equilibria.

KEYWORDS: Bayesian persuasion, Informed principal, Signaling games

JEL CODES: C72, D82, D83

*Department of Economics, the University of Texas at Austin (email: shaofej@utexas.edu). I thank V. Bhaskar, William Fuchs, Vasiliki Skreta, Maxwell Stinchcombe, Caroline Thomas, Mark Whitmeyer, Thomas Wiseman, and seminar participants at UT Austin and 2021 SEA Conference for helpful comments. Errors are my own. For the latest version: https://drive.google.com/file/d/1Ea_MxZ7bFXSUF_pKMR051DWZbvvdwtQe/view?usp=sharing

1. Introduction

Persuasion through public experimentation is prevalent. For example, a pharmaceutical company has to preregister clinical trials and submit complete results to the Food and Drug Administration (FDA) as part of the drug approval process. However, prior to conducting clinical trials, the pharmaceutical company already possess more information about the drug than the FDA. Sources of the pharmaceutical company's private information include internal R&D on related drugs, seeding trials and animal testings. A common belief is that pharmaceutical companies can benefit from having private information, and conversely, making this information public is socially beneficial since the FDA will be better informed when making decisions. This has led to calls for more transparency throughout the pharmaceutical industry beyond just during clinical research.

Another example is startup funding. An entrepreneur uses seed money to develop a prototype in order to convince an investor that a new technology can be reliably deployed. The prototype is developed according to business plans submitted to the investor in advance, and the investor can independently evaluate whether the prototype meets expectations. Despite the transparency in this process, the investor is naturally concerned that the entrepreneur has better knowledge about the technology and may capitalize on her private information when designing the prototype.

However, it is not immediately clear whether these concerns are necessary—that is, whether private information undermines the effectiveness of public experimentation and leads to welfare losses. On the one hand, if the pharmaceutical company designs the clinical trial contingent on the results of its internal research, the FDA is able to infer the pharmaceutical company's private information from the design of the clinical trial and disentangle any effect of the pharmaceutical company's private information. On the other hand, it is also possible that the pharmaceutical company's trial design does not depend on its private information. Moreover, it is unclear whether in equilibrium, the pharmaceutical company chooses a more informative or less informative clinical trial compared to the counterfactual where it does not have private information. Depending on the answers to these questions, private information may be beneficial or detrimental.

We study in this paper a costly persuasion game between a sender (e.g., the pharmaceutical company/the entrepreneur) and a receiver (e.g., the FDA/the investor). The sender persuades the receiver about an unknown, binary state of the world (e.g., whether a drug is effective/whether a technology is reliable) by conducting a public, costly experiment (e.g., a clinical trial/a prototype). The sender is privately and partially informed about the state of the world—at the outset of the game, she privately observes a noisy signal about the state

(e.g., results of internal research/proprietary knowledge about the technology), which is her type. Different types of the sender thus have different prior beliefs about the state of the world and can choose to run different experiments. Hence, this is a signaling game. The receiver can infer the state of the world by observing just the choice of experiment by the sender as well as from the outcome of the experiment.

We first derive a family of cost functions for experiments from a Wald's (1945) sequential sampling problem. That is, the sender sequentially acquires public signals, at a cost, about the state of the world. Two distinct features of this sampling problem are, first, the sender and the receiver have heterogeneous prior beliefs, and second, the cost of acquiring each signal is ex ante a random variable whose value depends on the signal's realization. An experiment in the persuasion game is equivalent to a threshold stopping rule in the Wald sampling problem, and we show that the cost of an experiment equals the expected reduction of a weighted log-likelihood ratio function evaluated at the sender's belief.

Both the pharmaceutical example and the startup funding example can be modeled in this way, but with one key difference. In the pharmaceutical example, a clinical trial consists of tests on individual patients, and the outcome of each patient is a public signal about the effectiveness of the drug. From the point of view of persuading the FDA, if a patient recovers, it is a piece of good news, and if a patient does not recover, it is a piece of bad news. And for the pharmaceutical company, bad news leads to a higher cost of the clinical trial, since the company has to treat the patient using existing drugs if she does not recover. In the startup example, the process of developing a prototype can be viewed as developing a series of features. From the point of view of persuading the investor, successfully developing a feature is good news, and failure to do so is bad news. However, good news leads to a higher cost of developing the prototype because of moral hazard: to incentivize the engineers who work on the prototype, the entrepreneur needs to pay them bonuses atop fixed wages if a feature is successfully developed.

Using this family of cost functions that we derive, our main results show that the equilibrium outcome of the persuasion game depends on the relative costs of drawing good news and bad news in the experiment. As is common for signaling games, multiple equilibria exist, and we focus on equilibria that satisfy the D1 criterion (Banks and Sobel, 1987). If bad news is more costly, there exists a unique separating equilibrium outcome. That is, different types of the sender choose different experiments, hence the sender's choice of experiment perfectly reveals her type. In this case, sender private information increases the equilibrium payoff of the receiver compared to a benchmark model where the sender's type is public. The reason is as follows. The sender type whose prior belief is higher chooses a more Blackwell-informative experiment than she would have chosen if her type is known to the receiver, in

order to deter the sender type whose prior belief is lower from mimicking. This implies that mandating disclosure of pharmaceutical companies' internal research reduces public welfare, as it decentivizes companies to run informative clinical trials.

In contrast, if good news is sufficiently more costly than bad news (in a precise sense), this is a signaling game where the single-crossing property fails. Under some technical conditions, there exists a continuum of pooling equilibria, that is, in addition to the unique separating equilibrium. In at least some pooling equilibria, the receiver's payoff is strictly lower than his equilibrium payoff in the benchmark model where the sender's type is public. This implies that entrepreneurs' private information can hurt investors, and due diligence that aims to reduce information asymmetry is necessary even when experiments are public.

Related literature. Kamenica and Gentzkow (2011) introduce the study of Bayesian persuasion via unrestricted, costless experiments where the sender and the receiver have the same prior belief about the state of the world. Their main result is concavification, that is, the sender's payoff can be expressed as a value function over the (common) posterior belief, and the sender's equilibrium payoff is the concave closure of that value function. Alonso and Câmara (2016) study an extension where the sender and the receiver have heterogeneous priors, but they "agree to disagree." They derive a bijection between the sender's posterior belief and the receiver's posterior belief, hence concavification can be applied after a translation of beliefs. Similarly, Gentzkow and Kamenica (2014) relaxes the assumption that experiments are costless. Using posterior separable (Caplin et al., 2018) costs of experiments, they show that a persuasion game with costly experiments and a common prior is equivalent to one with costless experiments and heterogeneous priors. Hence, the equilibrium can be solved using concavification.

The study of persuasion games with a privately informed sender is more recent. Perez-Richet (2014) studies equilibrium refinement in a persuasion game where the sender is fully informed of a binary state. Koessler and Skreta (2021) study a more general information design problem by a fully informed designer, allowing for many agents and private messages. Alonso and Câmara (2018) study persuasion by a partially informed sender who chooses an experiment from a restricted set of experiments. They provide conditions for a sender to never gain by becoming informed. All three papers assume that information transmission is costless and apply a generalization of the inscrutability principle (Myerson, 1983), which enables them to focus on pooling equilibria without loss of generality. In the current paper, in addition to introducing costs of experiments, we assume that the public experiment is not correlated with the sender's private information conditional on the state of the world. Hence, the inscrutability principal does not apply. Two other papers that feature both

costly experiments and sender private information are by Li and Li (2013) and by Degan and Li (2021). In both papers, a privately informed sender chooses from a restricted class of noisy signals that differ only on their precision, and the cost is increasing in the precision. In contrast, we allow the sender to choose any Blackwell experiment, and the cost of an experiment is increasing in its informativeness.

Within the Bayesian persuasion literature, the closest to the current paper is by Hedlund (2017). Hedlund (2017) studies persuasion by a sender who is partially informed about a binary state, and he also focuses on equilibria satisfying the D1 criterion. However, it is assumed that experiments are costless, and the sender’s payoff is strictly increasing in the receiver’s posterior belief. In equilibrium, the outcomes are either separating or fully disclosing. In a separating equilibrium, the experiments chosen by the sender are more informative than the ones in the benchmark case where the sender’s type is public.

The introduction of costs of experiments adds rich dynamics to persuasion games—the sender’s type is payoff relevant, and the single-crossing property fails if good news is sufficiently more costly than bad news. As a result, pooling equilibria (which are not fully revealing) may exist alongside separating equilibria. In contrast, when experiments are costless, the single-crossing property is always satisfied. The failure of single-crossing in our setting is worth noting beyond the Bayesian persuasion literature, since the majority of studies on signaling games assume single-crossing (see, e.g., the analysis of insurance markets by Rothschild and Stiglitz (1976); Wilson (1977)).

Our paper also contributes to the literature on cost of information, and specifically in the context of persuasion. Posterior-separable cost functions have been popular in modeling attention cost (e.g., Sims (1998, 2003)) and are used by Gentzkow and Kamenica (2014). However, an experiment, as is defined by Blackwell (1953), is a concept independent of beliefs, and with heterogeneous priors, it is not clear which player’s beliefs should be used to measure the cost of an experiment. By studying a Wald’s (1945) sequential sampling problem, we show that the cost of an experiment equals the expected reduction in a weighted log-likelihood ratio function evaluated at the sender’s belief. The setup of the sequential sampling problem is similar to that in Brocas and Carrillo (2007) and Henry and Ottaviani (2019), but we allow for heterogeneous priors and assume that the cost of acquiring each signal is *ex ante* a random variable that depends on its realization. The same class of log-likelihood ratio cost functions is studied also by Pomatto et al. (2020), who provide an axiomatic foundation for the cost function. Our microfoundation results complement their results, and we show that prior dependence of the cost of information can be motivated by the fact that the costs of drawing good and bad signals differ in the Wald sampling problem. Our result also complements other studies which microfound cost of information through sequential information acquisition but

by a single decision maker (e.g., Morris and Strack (2019); Bloedel and Zhong (2020)).

The rest of the paper is organized as follows. Section 2 presents the persuasion game with a privately and partially informed sender and costly experiments. Section 3 studies the symmetric information benchmark. Section 4 and section 5 study the pooling and separating equilibria of the game, respectively, and section 6 compares the receiver’s equilibrium payoff in the persuasion game and the symmetric information benchmark. Section 7 microfound the log-likelihood ratio cost function via a Wald’s sampling problem. Section 8 shows that our main results—failure of the single-crossing property and existence of pooling equilibria—are robust to Shannon entropy cost function. The last section concludes.

2. The Model

There is a sender (she), and a receiver (he). At the outset of the game, Nature determines a binary state of the world $\omega \in \Omega := \{G, B\}$ and a signal $\theta \in \Theta := \{h, l\}$ ¹ according to a commonly known distribution with full support. Let μ_0 be the probability of the good state (i.e., $\omega = G$), and μ_θ the probability of the good state conditional on the signal realization θ ; assume that $0 < \mu_l < \mu_0 < \mu_h < 1$. The sender privately observes the signal realization θ , and neither player observes the state realization ω . Therefore, θ is the sender’s *type*, and her prior belief on the good state is either μ_l or μ_h . On the other hand, the receiver’s prior belief on the good state is μ_0 .

The game proceeds as follows. The sender publicly chooses an experiment π on the state of the world which generates a binary outcome $s \in \{g, b\}$. That is, $\pi : \Omega \rightarrow \Delta(\{g, b\})$. The outcome of the chosen experiment s is determined according to the distribution $\pi(\cdot|\omega)$ and is publicly observed. The receiver takes a binary action $a \in \{0, 1\}$, and payoffs are realized.

2.1. Strategies. Given an experiment π , let $p = \pi(g|G)$ and $q = \pi(g|B)$. Without loss of generality, $p \geq q$. The experiment can thus be identified with (p, q) , and the set of experiments is $\Pi = \{(p, q) : 1 \geq p \geq q \geq 0\}$.

A pure strategy of the sender $\{\pi_\theta\}_{\theta \in \Theta}$ consists of experiments chosen by all types of the sender, where $\pi_\theta \in \Pi$ is the experiment chosen by the type θ sender. A pure strategy of the receiver is denoted $\mathbf{a} : \Pi \times \{g, b\} \rightarrow \{0, 1\}$. It selects an action at every information set of the receiver, which is identified by the sender’s choice of experiment π and its outcome s .

2.2. Beliefs. After observing the sender’s choice of experiment but before observing its outcome, the receiver forms a belief about the sender’s type. Let $\gamma(\theta|\pi)$ denote his belief

¹All our results can be generalized to accommodate finite Θ .

that the sender's type is θ after experiment π is chosen. It is more convenient to keep track of $\beta(\pi) := \sum_{\theta \in \Theta} \gamma(\theta|\pi)\mu_\theta$. Notice that $\beta(\pi) \in [\mu_l, \mu_h]$ is the receiver's interim belief on the good state.

After the outcome is observed, both players update their beliefs. Let $\hat{\mu}(\theta, \pi, s)$ and $\hat{\beta}(\pi, s)$ be the posterior beliefs of the type θ sender and the receiver, respectively, that the state is good after observing outcome s from experiment π .

2.3. Cost of experiment and the sender's payoff. The sender strictly prefers the high receiver action over the low action. Her payoff $v(a, \pi|\theta) = a - c(\pi|\mu_\theta)$ consists of two parts: a reward, which is normalized to 1, if the receiver chooses the high action, minus the cost of experiment $c(\pi|\mu_\theta)$, which equals the expected reduction in a weighted log-likelihood ratio function, measured using the sender's beliefs. That is,

$$c(\pi|\mu) = \mathbb{E}[H(\mu) - H(\hat{\mu})],$$

where

$$H(\mu) = C_g \mu \ln \left(\frac{1-\mu}{\mu} \right) + C_b (1-\mu) \ln \left(\frac{\mu}{1-\mu} \right)$$

and $C_g, C_b > 0$, and $\hat{\mu}$ is the sender's posterior belief induced by the experiment π . Fixing the sender's prior and an experiment, $\hat{\mu}$ is a random variable whose value depends on the experiment's outcome and is calculated by the Bayes' rule.

Several remarks are in order. First, in section 6, we microfound this cost function by a Wald sampling problem. The parameters C_g and C_b are affine transformations to the costs of drawing good and bad news in the sampling problem. Hence, the cost function can accommodate both leading examples. The parameterization $C_g < C_b$ models scenarios such as pharmaceutical companies conducting clinical trials, where bad news is more costly than good news, whereas $C_g > C_b$ models scenarios such as entrepreneurs developing prototypes, where good news is more costly than bad news. Second, for almost all experiments and parameters C_g and C_b , the cost of running an experiment depends on the sender's prior belief, and different types of the sender may rank a set of experiments differently in terms of their idiosyncratic costs. Third, the cost is increasing in the Blackwell-informativeness of the experiment. The cost of running an uninformative experiment (i.e., $p = q$) is zero, and the cost of running a fully revealing experiment (i.e., $p = 1 > q$ or $p > q = 0$) is infinity. The intuition is that, to be certain of either state, the sender must in expectation acquire infinite number of signals in the sampling problem. Lastly, in section 7, we show that our result that the single-crossing property fails is robust to Shannon Entropy cost.

2.4. The receiver's payoff. The receiver follows a threshold decision rule and chooses the high action $a = 1$ if and only if his posterior belief is at least some commonly known threshold $\bar{\beta} \in (0, 1)$. We take this threshold as given and fixed throughout the paper.

To analyze the impact of sender private information on the receiver, we specify the payoff of the receiver. The receiver's payoff $u(a, \omega)$ depends on his action a and the state of the world ω . By normalization, let $u(0, G) = u(0, B) = 0$, $u(1, G) = 1$, and $u(1, B) = -\beta^*/(1 - \beta^*)$, where $\beta^* \in (0, 1)$ is the sequentially rational threshold of choosing the high action.

We assume that $\bar{\beta} > \beta^* > \mu_h$. The assumption that $\bar{\beta} > \beta^*$ recognizes that the receiver is able to commit to a higher standard $\bar{\beta}$ so as to elicit more information from the sender. We do not explicitly study the receiver's commitment problem, but we make the following observations. If $\bar{\beta} = \beta^*$, the receiver's equilibrium payoff is zero, as is in a standard persuasion game; if $\bar{\beta} = 1$, persuasion is not possible due to the cost, and the receiver's equilibrium payoff is again zero. Therefore, there is an optimal threshold $\bar{\beta} \in (\beta^*, 1)$ that maximizes the receiver's expected payoff. The assumption $\beta^* > \mu_h$ implies that the receiver is never persuaded at the interim stage. Regardless of the receiver's interim belief, it is not optimal for the receiver to choose the high action if an experiment generates the bad outcome.

2.5. Equilibrium. An *equilibrium* consists of pure strategies of the players, $\{\pi_\theta\}_{\theta \in \Theta}$ and $\mathbf{a} : \Pi \times \{g, b\} \rightarrow \{0, 1\}$, and the receiver's system of beliefs $\beta : \Pi \rightarrow [\mu_l, \mu_h]$, $\hat{\beta} : \Pi \times \{g, b\} \rightarrow [0, 1]$, such that

- (1) Given the receiver's strategy \mathbf{a} , the sender's strategy is optimal, i.e.,

$$\pi_\theta \in \arg \max_{\pi \in \Pi} \mathbb{E}[v(\mathbf{a}(\pi, s), \pi | \theta)]$$

for all $\theta \in \Theta$;

- (2) The receiver follows the threshold rule, i.e., $\mathbf{a}(\pi, s) = 1$ if and only if $\hat{\beta}(\pi, s) \geq \bar{\beta}$;
(3) Beliefs are updated using Bayes' rule whenever possible. That is, $\beta(\pi_l) = \beta(\pi_h) = \mu_0$ if $\pi_l = \pi_h$, $\beta(\pi_\theta) = \mu_\theta$ if $\pi_l \neq \pi_h$, and

$$\hat{\beta}(\pi, s) = \mathbf{B}(\beta(\pi), \pi, s) := \frac{\beta(\pi)\pi(s|G)}{\beta(\pi)\pi(s|G) + (1 - \beta(\pi))\pi(s|B)}$$

if $\pi(s|G) + \pi(s|B) \neq 0$.

We say an equilibrium is a *pooling equilibrium* if $\pi_l = \pi_h$ and a *separating equilibrium* if $\pi_l \neq \pi_h$. Furthermore, an equilibrium is a *persuasion equilibrium* if the high action is taken with positive probability on the equilibrium path, i.e., $\mathbb{E}_{\theta, s}[\mathbf{a}(\pi_\theta, s)] > 0$. Otherwise, it is a *trivial equilibrium*. Given an equilibrium, we call the collection of experiments $\{\pi_\theta\}_{\theta \in \Theta}$ the

equilibrium outcome.

2.6. The D1 criterion. Among all equilibria, we are particularly interested in persuasion equilibria that satisfy the D1 criterion (Banks and Sobel, 1987).

Given the receiver's interim belief $\beta \in [\mu_l, \mu_h]$ and an experiment $\pi \in \Pi$, denote by $\bar{v}(\beta, \pi|\theta)$ the sender's expected payoff if the receiver updates his posterior belief using the Bayes' rule and chooses the sequentially optimal action. That is, $\bar{v}(\beta, \pi|\theta) = \mathbb{P}[\mathbf{B}(\beta, \pi, s) \geq \bar{\beta}] - c(\pi|\mu_\theta)$. Given an equilibrium of the game, denote by $v_\theta^* = \bar{v}(\beta(\pi_\theta), \pi_\theta|\theta)$ the equilibrium payoff of the type θ sender, and for any deviation $\pi \in \Pi \setminus \{\pi_\theta\}_{\theta \in \Theta}$, define

$$\begin{aligned} D_\theta(\pi) &= \{\beta \in [\mu_l, \mu_h] : \bar{v}(\beta, \pi|\theta) > v_\theta^*\}, \\ D_\theta^0(\pi) &= \{\beta \in [\mu_l, \mu_h] : \bar{v}(\beta, \pi|\theta) \geq v_\theta^*\}. \end{aligned}$$

That is, $D_\theta(\pi)$ is the set of receiver interim beliefs that warrant π a profitable deviation for the type θ sender, and $D_\theta^0(\pi)$ is the set of receiver interim beliefs following which the deviation gives the sender at least the same payoff as her equilibrium payoff.

An equilibrium satisfies the D1 criterion if for all $\theta \neq \theta'$ and deviations $\pi \in \Pi \setminus \{\pi_\theta\}_{\theta \in \Theta}$ such that $D_{\theta'}^0(\pi) \subsetneq D_\theta(\pi)$, the receiver's off-path interim belief $\beta(\pi) = \mu_\theta$. Intuitively, if the type θ sender is more keen to deviate to π in the sense that such deviation is profitable for her given a larger set of receiver beliefs than for the other sender type θ' , the receiver should attribute this deviation to the type θ sender.

3. Symmetric Information Benchmarks

To illustrate our settings and introduce some useful notation, consider a benchmark model with symmetric information. Suppose that the sender's prior belief on the good state is μ_s , and the receiver's prior belief is $\mu_r < \bar{\beta}$. That is, we allow the sender and the receiver to have heterogeneous priors, but they "agree to disagree." The sender publicly chooses an experiment. The receiver takes the binary action after observing the outcome of the experiment. Since there is no sender private information, receiver learns only from the outcome of experiment.

Given any experiment $\pi = (p, q) \neq (0, 0)$, we say that it is *persuasive at belief* μ_r if the receiver chooses the high action after seeing the good outcome. That is, the receiver's posterior belief $\mathbf{B}(\mu_r, \pi, g) \geq \bar{\beta}$. Equivalently,

$$\frac{q}{p} \leq \mathbf{Q}(\mu_r) := \frac{\mu_r}{1 - \mu_r} \bigg/ \frac{\bar{\beta}}{1 - \bar{\beta}}.$$

All other experiments, including the uninformative experiment $(0, 0)$, are *unpersuasive at belief* μ_r . That is, the receiver chooses the low action regardless of the outcome of the experiment.

The sender's expected payoff from choosing the experiment π is

$$f(\pi, \mu_s) := \mu_s p + (1 - \mu_s)q - c(\pi|\mu_s)$$

if it is persuasive at belief μ_r . Although the experiment $(0, 0)$ is not persuasive at belief μ_r , it is an accumulation point of the set of persuasive experiments, and the sender's expected payoff is zero from choosing the experiment $(0, 0)$. Hence, the sender can receive a payoff arbitrarily close to zero from choosing persuasive experiments. On the other hand, the sender's payoff is simply $-c(\pi|\mu_r) \leq 0$ if the experiment π is unpersuasive, and the sender can achieve zero payoff if and only if she chooses an uninformative experiment, i.e., $p = q$.

Hence, the sender's equilibrium payoff in the symmetric information benchmark is

$$V(\mu_s, \mu_r) := \sup_{\pi \in \Pi, \frac{q}{p} \leq \mathbf{Q}(\mu_r)} f(\pi, \mu_s) \geq 0.$$

Two scenarios are possible. When the cost of experiment is low (in a precise sense, see appendix A.1), $V(\mu_s, \mu_r) > 0$, and there exists a unique experiment $\hat{\pi}(\mu_s, \mu_r)$ that is persuasive at belief μ_r and obtains the sender's equilibrium payoff, i.e., $f(\hat{\pi}(\mu_s, \mu_r), \mu_s) = V(\mu_s, \mu_r)$. In any equilibrium, the sender chooses the experiment $\hat{\pi}(\mu_s, \mu_r)$, and the receiver's posterior belief after seeing the good outcome is $\bar{\beta}$. That is, $\hat{\pi}(\mu_s, \mu_r)$ is the unique equilibrium outcome. When the cost of experiment is high, $V(\mu_s, \mu_r) = 0$, and persuasion is not possible in equilibrium. In all equilibria, the sender chooses an uninformative experiment which gives her zero payoff. The equilibrium outcome is essentially unique, since all uninformative experiments are equivalent under Blackwell's order of informativeness. Denote by π_u the equivalence class of uninformative experiments, and by abuse of notation, we say $\hat{\pi}(\mu_s, \mu_r) = \pi_u$ is the unique equilibrium outcome in the latter case.

Moreover, fixing the receiver's prior μ_r and the cost parameters C_g and C_b , the equilibrium outcome $\hat{\pi}(\mu_s, \mu_r)$ is ranked by the sender's prior μ_s . If good news is sufficiently more costly than bad news, i.e., C_g is sufficiently greater than C_b (in a precise sense, see appendix A.1), the equilibrium outcome is less Blackwell-informative if the sender's prior is higher, and the equilibrium payoffs of the sender and the receiver are both decreasing in the sender's prior. If C_g is small relative to C_b , the equilibrium outcome is more Blackwell-informative if the sender's prior is higher, and the players' equilibrium payoffs are increasing in the sender's prior. This comparative statics result is closely related to failure of single-crossing in our

main model with sender private information.

A special case where $\mu_s = \mu_r$ is studied in Gentzkow and Kamenica (2014) and pertains to the two welfare benchmarks of our model. In the *no signal benchmark*, the sender does not observe the noisy signal θ , hence both players have prior belief μ_0 . In the *public signal benchmark*, the noisy signal θ is observed by both the sender and the receiver. Hence, after observing the noisy signal, the continuation game is one with common prior μ_θ . The receiver is weakly better off with public signal compared to no signal.

LEMMA 1. *Let U^{ns} be the receiver's ex ante expected payoff in the no signal benchmark, and U^{ps} that in the public signal benchmark. $U^{ns} \leq U^{ps}$.*

4. Pooling Equilibria

We now turn to the game with incomplete information. Since the receiver updates his belief about the state of the world after seeing the sender's choice of experiment, whether an experiment can persuade the receiver depends on the interim belief held by the receiver. If an experiment π is persuasive at belief $\beta(\pi)$, then the receiver chooses the high action if the experiment π is chosen and yields the good outcome.

Since the receiver's lowest possible interim belief is μ_l , an experiment $\pi = (p, q)$ is *always persuasive* if $\frac{q}{p} \leq \mathbf{Q}(\mu_l)$. That is, regardless of his interim belief, the receiver chooses the high action if an always persuasive experiment yields the good outcome. Conversely, if $\frac{q}{p} > \mathbf{Q}(\mu_h)$, the experiment π is *never persuasive*. For any other experiment, it is persuasive if and only if the receiver's interim belief is high enough, namely, $\beta(\pi) \geq \mathbf{Q}^{-1}(\frac{q}{p})$.

The sender can achieve a *payoff guarantee* $\bar{V}(\mu_\theta) := V(\mu_\theta, \mu_l)$ by choosing the experiment $\hat{\pi}_\theta := \hat{\pi}(\mu_\theta, \mu_l)$. If $\bar{V}(\mu_\theta) > 0$, the experiment $\hat{\pi}_\theta$ is an always persuasive experiment, hence the sender's payoff from choosing $\hat{\pi}_\theta$ is $\bar{V}(\mu_\theta)$ regardless of the receiver's interim belief. If $\bar{V}(\mu_\theta) = 0$, any always persuasive experiment gives her negative payoff, and the sender achieves zero payoff only by choosing an uninformative experiment. Therefore, the type θ sender's ex ante expected payoff must be at least $\bar{V}(\mu_\theta)$ in any equilibrium.

4.1. Pooling equilibrium outcomes. In a pooling persuasion equilibrium, the experiment chosen by both types of the sender, $\pi = (p, q)$, must be persuasive at the receiver's interim belief $\beta(\pi) = \mu_0$. That is, $\frac{q}{p} \leq \mathbf{Q}(\mu_0)$. This condition, along with the the sender's payoff guarantee, fully characterizes the set of pooling persuasion equilibrium outcomes. A pooling trivial equilibrium exists only if both sender types' payoff guarantees are zero. The following proposition characterizes the set of pooling equilibrium outcomes.

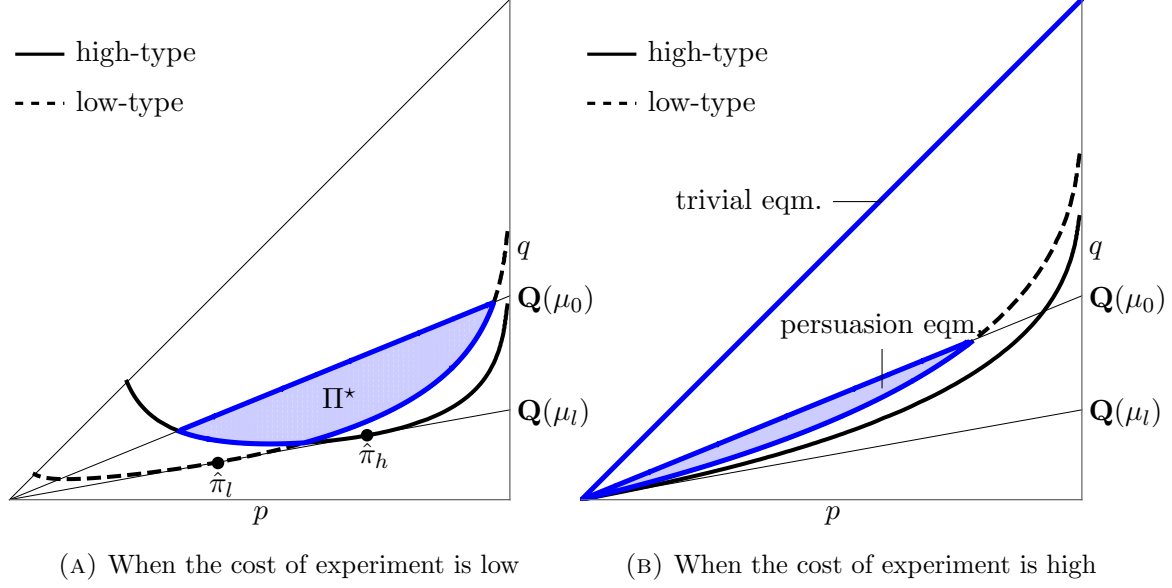


FIGURE 1: The set of pooling equilibria

PROPOSITION 2. *There exists a pooling persuasion equilibrium where both sender types choose experiment $\pi = (p, q)$ if and only if $f(\pi, \mu_\theta) \geq \bar{V}(\mu_\theta)$ for all $\theta \in \Theta$ and $\frac{q}{p} \leq \mathbf{Q}(\mu_0)$. There exists a pooling trivial equilibrium where both sender types choose experiment $\pi = (p, q)$ if and only if $p = q$ and $\bar{V}(\mu_\theta) = 0$ for all $\theta \in \Theta$.*

Figure 1 illustrates the set of pooling equilibrium outcomes. In each subfigure, the right triangle is the set of experiments Π . The horizontal and vertical coordinates of an experiment are the probabilities of the good outcome under the good and bad states, respectively. The set of always persuasive experiments is $\{\pi : \frac{q}{p} \leq \mathbf{Q}(\mu_l)\}$, and the set of experiments that are persuasive at belief μ_0 is $\{\pi : \frac{q}{p} \leq \mathbf{Q}(\mu_0)\}$. The two straight lines show the boundaries of these sets.

When the cost of experiment is low, both sender types' payoff guarantees are positive, as is the case in figure 1(A). By choosing experiment $\hat{\pi}_l$, the low-type sender can achieve payoff $\bar{V}(\mu_l)$ regardless of the interim belief $\beta(\hat{\pi}_l)$ held by the receiver. The dashed curve through $\hat{\pi}_l$ is the low-type sender's indifference curve, i.e., $f(\pi, \mu_l) = \bar{V}(\mu_l)$. An experiment π that is above the indifference curve gives the low-type sender a higher payoff than her payoff guarantee if it is persuasive at belief $\beta(\pi)$. Similarly, the high-type sender can guarantee payoff $\bar{V}(\mu_h)$ by choosing experiment $\hat{\pi}_h$.

The set of pooling equilibrium experiments is the highlighted area Π^* in figure 1(A). That is, the set of all experiments which are persuasive at belief μ_0 and give both sender types at least their respective payoff guarantees. Taking any experiment $\pi^* \in \Pi^*$, we can construct a pooling persuasion equilibrium where the sender chooses the experiment π^* , as

follows. Both sender types choose the experiment π^* , and the receiver's interim belief is such that $\beta(\pi^*) = \mu_0$ and $\beta(\pi) = \mu_l$ for all $\pi' \neq \pi^*$. That is, any deviation is believed to be from the low-type sender. The receiver updates his posterior belief using Bayes' rule and chooses the action according to the threshold rule. Off the equilibrium path, the receiver chooses the high action after seeing the good outcome if the sender deviates to an always persuasive experiment, and he chooses the low action otherwise. It is easy to verify that neither sender type has a profitable deviation. If the sender deviates to an always persuasive experiment π' , her deviation payoff $f(\pi', \mu_\theta) \leq \bar{V}(\mu_\theta) \leq f(\pi^*, \mu_\theta)$ is bounded by her equilibrium payoff; any other deviation is unpersuasive given the receiver's critical off-path interim belief, so the sender's deviation payoff is bounded by zero.

When the cost of experiment is sufficiently high, the payoff guarantees for both types of the sender are zero, as is the case in figure 1(B). The two curves passing through the origin show the zero-payoff curves for the sender, with the dashed curve being for the low-type sender. In this case, there is a continuum of pooling trivial equilibria where both sender types pool on an uninformative experiment and receive zero payoffs, as well as a continuum of pooling persuasion equilibria where both sender types receive nonnegative payoffs.

We discuss in details the existence of pooling equilibria in appendix A.4. Main results are highlighted here. Fixing the sender's possible prior beliefs $\{\mu_\theta\}_{\theta \in \Theta}$, pooling trivial equilibria exist if and only if the cost of experiment is sufficiently high. Pooling persuasion equilibria exist if and only if the cost of experiment is sufficiently low, and the receiver's prior μ_0 is sufficiently high—that is, the probability of the high-type sender is sufficiently high. Intuitively, if the receiver's prior μ_0 is close to μ_l , the low-type sender's gain from mimicking the high-type sender becomes small, since it does not change the receiver's interim belief much. Hence, no pooling persuasion equilibrium exists.

4.2. Single-crossing and the D1 criterion. In the preceding section, pooling equilibria are constructed using critical off-path beliefs of the receiver. Although there are other pooling equilibria with different off-path receiver beliefs, the key to constructing any pooling equilibrium is that, following any deviation π' which the sender may find profitable (i.e., $D_\theta(\pi')$ is nonempty for some sender type θ), the receiver's interim belief $\beta(\pi')$ must be sufficiently low so as to deter the sender from deviating. Applying the D1 criterion, persuasion equilibria can be ruled out generically.

Take, for example, the set of pooling equilibrium outcomes Π^* we solve in figure 1(A) and consider an equilibrium where the sender chooses an experiment $\pi^* \in \Pi^*$. The two indifference curves through π^* in figure 2 show the sender's equilibrium payoffs. That is, any experiment $\pi' = (p', q')$ above the solid curve is a profitable deviation for the high-type




FIGURE 2: The D1 criterion eliminates an pooling equilibrium

sender if it is persuasive at belief $\beta(\pi')$, i.e., if $\beta(\pi') \in D_h(\pi') = [\mathbf{Q}^{-1}(\frac{q'}{p'}), \mu_h]$ (the interval is nonempty if $\frac{q'}{p'} \leq \mathbf{Q}(\mu_h)$). On the other hand, any experiment $\pi' = (p', q')$ below the dashed curve is strictly equilibrium dominated for the low-type sender even if it is persuasive at belief $\beta(\pi')$. That is, $D_l^0(\pi') = \emptyset$.²

Since the two sender types' indifference curves have different slopes at π^* , we can choose a *nearby* experiment π' such that it is above the high-type's indifference curve but below the low-type's indifference curve. Considering π' as a deviation of the sender, $D_l^0(\pi') = \emptyset \subsetneq D_h(\pi') = [\mathbf{Q}^{-1}(\frac{q'}{p'}), \mu_h]$. Hence, the D1 criterion argues that a reasonable receiver should hold off-path belief $\beta(\pi') = \mu_h$. However, given this belief, π' is indeed a strictly profitable deviation for the high-type sender, so this off-path belief cannot be supported by any pooling equilibrium where both sender types choose the experiment π^* .

To formalize this idea, let us define the marginal rate of substitution (of p for q for sender type θ) as

$$MRS(\pi|\mu_\theta) = -\frac{\partial f(\pi, \mu_\theta)/\partial p}{\partial f(\pi, \mu_\theta)/\partial q}.$$

That is, $MRS(\pi|\mu_\theta)$ is the slope of the type θ sender's indifference curve at the experiment $\pi = (p, q)$. Notice that at any uninformative experiment, the marginal rate of substitution

$$MRS(\pi|\mu_\theta) = -\frac{\mu_\theta}{1 - \mu_\theta}$$

is strictly decreasing in μ_θ . We say that *the single-crossing property* is satisfied if this ranking

²Hence, $D_{\theta'}^0(\pi) \subsetneq D_\theta(\pi) \Rightarrow D_{\theta'}^0(\pi) = \emptyset$. That is, the D1 criterion gives the same result as the intuitive criterion (Cho and Kreps, 1987) in our game.

holds true at all experiments (in the interior of Π), that is, if $MRS(\pi|\mu_h) \leq MRS(\pi|\mu_l)$ for all $\pi \in \Pi^\circ$. If the single-crossing property is satisfied, different sender types' indifference curves are never tangent. Hence, we can find a nearby deviation and apply the D1 criterion to rule out all pooling persuasion equilibria.

Our next result characterizes when the single-crossing property holds. We show that, the single-crossing property fails if and only if C_g is sufficiently greater than C_b . That is, if good news is sufficiently more costly than bad news. As a comparison, when experiments are costless, the sender's expected payoff $f(\pi, \mu_\theta) = p\mu_\theta + q(1 - \mu_\theta)$ is linear in π and μ_θ , hence the single-crossing property is always satisfied.

PROPOSITION 3. *There exists an increasing and concave function $\hat{K} : \mathbb{R}_{++} \rightarrow \mathbb{R}_{++}$ such that the single-crossing property is satisfied if and only if $C_g \leq \hat{K}(C_b)$. Moreover, $\hat{K}(C_b) > C_b$ for all $C_b > 0$.*

COROLLARY 4. *If $0 < C_g \leq \hat{K}(C_b)$, there is no pooling persuasion equilibrium satisfying the D1 criterion.*

Specifically, if bad news is more costly, i.e., $0 < C_g < C_b$, no pooling persuasion equilibrium satisfies the D1 criterion. The function \hat{K} is fully characterized in the proof. Although it does not admit a closed form solution, the function is shown in figure 3. The shaded region is the set of parameters where the single-crossing property fails. As a remark, the function \hat{K} does not depend on the sender's prior beliefs, so the above proposition 3 extends to finitely many sender types.




FIGURE 3: The D1 criterion fails in the highlighted area

4.3. Failure of single-crossing. We now turn to the case where the single-crossing property fails. Since the difference between the high- and low-types' marginal rates of substitution change signs, there exist experiments where different sender types' indifference curves are tangent. The following proposition characterizes such experiments.

PROPOSITION 5. *If $C_g > \hat{K}(C_b)$, there exist $\hat{\mathbf{p}}, \check{\mathbf{p}} : (0, 1) \rightarrow (0, 1)$ such that $q < \hat{\mathbf{p}}(q) < \check{\mathbf{p}}(q)$, and $MRS(\pi|\mu_h) = MRS(\pi|\mu_l)$ if and only if $p = \hat{\mathbf{p}}(q)$ or $p = \check{\mathbf{p}}(q)$. Moreover, $\hat{\mathbf{p}}(q)$ is decreasing in C_g and increasing in C_b for all $q \in (0, 1)$.*

As is the case with the function \hat{K} , the functions $\hat{\mathbf{p}}$ and $\check{\mathbf{p}}$ are independent of the sender's prior beliefs μ_l and μ_h . They depend only on the parameters C_g and C_b . Indeed, the marginal rate of substitution at every experiment is weakly monotonic in the sender's prior belief, so we can extend our results to finitely many sender types.

When different sender types' indifference curves are tangent, the aforementioned argument of the D1 criterion may not apply. The highlighted curves in figure 4 plot the set of experiments where the high and low sender types' marginal rates of substitution are equal. Two experiments, $\pi_1 = (p_1, q_1)$ and $\pi_2 = (p_2, q_2)$, are so chosen that $p_1 = \hat{\mathbf{p}}(q_1)$ and $p_2 = \check{\mathbf{p}}(q_2)$, and we plot the indifference curves for the sender at these experiments. At both experiments, the sender's indifference curves are tangent. The difference between them is that, at π_1 , the high-type sender's indifference curve is more convex than the low-type sender's, while at π_2 , the low-type sender's indifference curve is more convex. At experiment π_2 , we can find a nearby experiment $\pi' \neq \pi_2$ that is above the high-type sender's indifference curve but below the low-type sender's indifference curve. In any pooling equilibrium where the sender chooses π_2 , the high-type sender is more keen to deviate to π' . Hence, the pooling equilibrium does not satisfy the D1 criterion.

However, at experiment π_1 , since the high-type sender's indifference curve is more convex, we cannot find a nearby deviation that is always equilibrium dominated for the low-type sender but strictly profitable for the high-type sender when the receiver's interim belief is high. In fact, for any nearby deviation π' that is between the two indifference curves, the low-type sender is more keen to deviate to π' , so the receiver's interim belief should be $\beta(\pi') = \mu_l$ by the D1 criterion, which is consistent with the receiver's critical off-path beliefs. If q_1 is small, the two indifference curves through π_1 intersect again close to $p = 1$, and there exists a *large* deviation $\pi' = (p', q')$ such that $\frac{q'}{p'} \leq \mathbf{Q}(\mu_h)$, and it is above the high-type sender's indifference curve but below the low-type sender's indifference curve. In this case, the high-type sender is more keen to deviate to π' , and we can again apply the D1 criterion to rule out any pooling equilibrium where the sender chooses π_1 . To summarize, in any pooling persuasion equilibrium that satisfies the D1 criterion, the sender's choice of

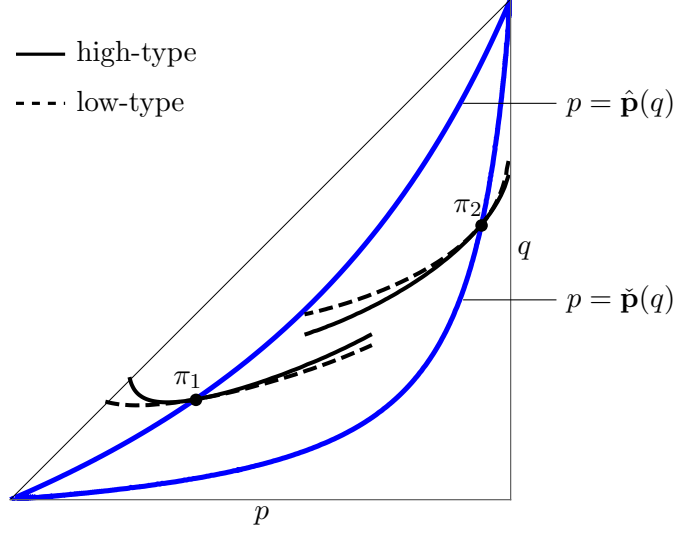


FIGURE 4: Experiments where indifference curves are tangent

experiment $\pi^* = (p^*, q^*)$ must be such that $p^* = \hat{p}(q^*)$ so that no nearby deviation can break the equilibrium, and q^* must be sufficiently large so that no large deviation can break the equilibrium.

PROPOSITION 6. *If $C_g > \hat{K}(C_b)$, there exists \bar{q} such that a pooling persuasion equilibrium outcome $\pi^* = (p^*, q^*)$ can be supported by an equilibrium satisfying the D1 criterion if and only if $p^* = \hat{p}(q^*)$ and $q^* > \bar{q}$.*

5. Separating equilibria

In a separating equilibrium, the experiment chosen by the low-type sender π_l is different from the experiment chosen by the high-type sender π_h . For the low sender type, her equilibrium payoff must be at least her payoff guarantee $\bar{V}(\mu_l)$. However, she cannot achieve a higher payoff than $\bar{V}(\mu_l)$, since at the interim stage, the receiver knows that she is the low-type, i.e., $\beta(\pi_l) = \mu_l$. Therefore, in a separating equilibrium, the low-type sender chooses $\pi_l = \hat{\pi}_l$ if her payoff guarantee $\bar{V}(\mu_l) > 0$, and she chooses an uninformative experiment if $\bar{V}(\mu_l) = 0$. In both cases, the high-type sender's payoff from mimicking the low-type sender is bounded by her payoff guarantee $\bar{V}(\mu_h)$. Hence, the high-type sender does not want to mimic the choice of the low-type sender.

The following proposition characterizes the set of separating equilibrium outcomes.

PROPOSITION 7. *There exists a separating persuasion equilibrium where the type θ sender chooses experiment $\pi_\theta = (p_\theta, q_\theta)$ if and only if $\pi_l = \hat{\pi}_l$, $f(\pi_h, \mu_h) \geq \bar{V}(\mu_h)$, $f(\pi_h, \mu_l) \leq \bar{V}(\mu_l)$,*

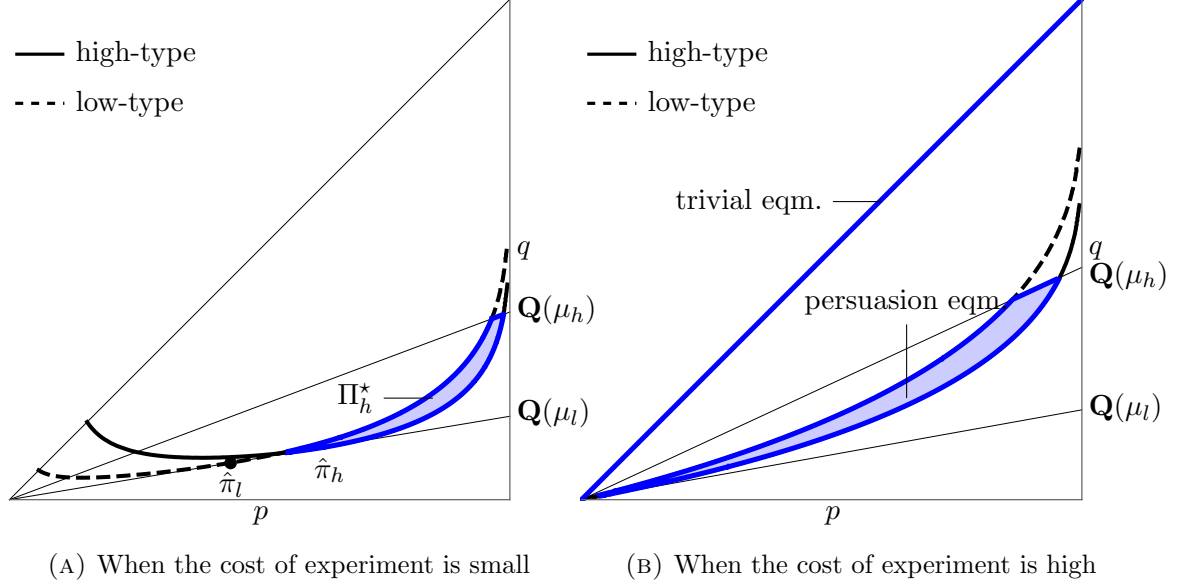


FIGURE 5: The set of separating equilibria

and $\frac{q_h}{p_h} \leq \mathbf{Q}(\mu_h)$. There exists a separating trivial equilibrium if and only if $\bar{V}(\mu_\theta) = 0$ for all $\theta \in \Theta$, and $\pi_l \neq \pi_h$ are both uninformative.

Figure 5 illustrates the result. When the cost of experiment is low, both sender types' payoff guarantees are positive, as is the case in figure 5(A), and there exists a continuum of separating equilibria that differ on the experiment chosen by the high-type sender. The low-type sender chooses the always persuasive experiment $\hat{\pi}_l$ which achieves her payoff guarantee in every separating experiment, and Π_h^* is the set of equilibrium experiments chosen by the high-type sender. When the cost of experiment is high, both sender types' payoff guarantees are zero, as is the case in figure 5(B), and there exist separating trivial equilibria as well as separating persuasion equilibria.

The following proposition characterizes the set of separating persuasion equilibria that satisfy the D1 criterion. Suppose that a separating persuasion equilibrium exists. Then there exists an essentially unique separating persuasion equilibrium that satisfies the D1 criterion which maximizes the sender's expected payoff.

PROPOSITION 8. *A separating persuasion equilibrium outcome can be supported by an equilibrium that satisfies the D1 criterion if and only if it is sender optimal. That is, there exists a unique D1 separating equilibrium outcome $\{\pi_l^*, \pi_h^*\}$ such that the experiment chosen by the low-type sender $\pi_l^* = \hat{\pi}_l$, and the experiment chosen by the high-type sender*

$$\pi_h^* = \operatorname{argmax}_{\pi \in \Pi_h^*} f(\pi, \mu_h).$$

In appendix A.7, we show the condition that trivial equilibria satisfy the D1 criterion. Suppose that $\bar{V}(\mu_\theta) = 0$ for all $\theta \in \Theta$. Then there exist pooling and separating trivial equilibria. If, in addition, there exists a separating persuasion equilibrium, no (pooling or separating) trivial equilibrium satisfies the D1 criterion. If no separating persuasion equilibrium exists, then all (pooling and separating) trivial equilibrium outcomes can be supported by an equilibrium satisfying the D1 criterion. In terms of the cost parameters, trivial equilibria satisfy the D1 criterion if either C_g is large relative to C_b , or if both parameters are sufficiently large. The result does not depend on the single-crossing property.

6. Welfare Comparison

We focus on the case where the cost of experiment is such that persuasion equilibria exist. When the single-crossing property holds, there exists a unique equilibrium outcome $\{\pi_l^*, \pi_h^*\}$ that can be supported by a separating persuasion equilibrium satisfying the D1 criterion. For the low-type sender, the experiment chosen is the same as in the public signal benchmark, i.e., $\pi_l^* = \pi_l^{ps} = \hat{\pi}_l$. For the high-type sender, the unique experiment chosen in D1 separating equilibria is strictly more Blackwell-informative than in the public signal benchmark, i.e., $\pi_h^* \succ_B \pi_h^{ps}$. Hence, the receiver's equilibrium payoff in D1 separating equilibria is strictly higher than that in the public signal benchmark.

When the single-crossing property fails and there exist a continuum of pooling persuasion equilibria that satisfy the D1 criterion, we compare the D1 pooling equilibrium outcome $\pi^* = (p^*, q^*)$ such that $\frac{q^*}{p^*} = \mathbf{Q}(\mu_0)$ with the equilibrium outcome π^{ns} in the no information benchmark. The experiment π^* is strictly less Blackwell-informative than π^{ns} , i.e., $\pi^* \prec_B \pi^{ns}$. Therefore, at least in some D1 pooling equilibria where the sender chooses an experiment close to π^* , the receiver's equilibrium payoff is strictly less than that in the no signal benchmark.

PROPOSITION 9. *Suppose that persuasion equilibria exist. In the unique D1 separating equilibrium outcome, the receiver's equilibrium payoff is strictly higher than that in the public signal benchmark. If there exist pooling equilibria that satisfy the D1, then the receiver's equilibrium payoff is strictly lower in some D1 pooling equilibria than in the no signal benchmark.*

7. Microfoundation of the Cost Function

We microfound the log-likelihood ratio cost function for experiments used in the persuasion game by a Wald sampling problem. There is a sequence of noisy signals about the

state of the world which is ex ante unobserved to either player. The sender can sequentially acquire signals at a cost, and the cost of acquiring each signal depends on the signal realization. Once acquired, a signal is publicly observed by both players. The sender’s stopping decision is irreversible, and the receiver moves after the sender stops acquiring signals. An experiment in the persuasion game is equivalent to a threshold stopping rule in the Wald sampling problem—that is, the sender stops acquiring additional signals at the first instance the difference between the numbers of good and bad signals reaches some thresholds. To microfound the cost function, we can equivalently model the persuasion game by assuming that the sender commits to such a stopping rule in the sampling problem after she observes her type. Since the receiver learns about the sender’s type only at this interim stage but not during the sampling process, it is sufficient to focus on a complete information version of the sampling problem, where the sender and the receiver hold different prior beliefs but “agree to disagree,” and there is no sender private information.

7.1. The Wald sampling problem. Let us first consider a discrete version of the problem. There is a binary state of the world $\omega \in \{G, B\}$. A sender and a receiver have heterogeneous priors. The sender’s prior belief on the good state is denoted μ , and the receiver’s β , and we assume that $\mu, \beta \in (0, 1)$.

There is a sequence of binary signals $(s_n)_{n=1}^\infty$ such that each $s_n \in \{g, b\}$. It is common knowledge that conditional on the state, signals are distributed iid such that $\mathbb{P}(s_n = g | \omega = G) = \mathbb{P}(s_n = b | \omega = B) = \alpha > \frac{1}{2}$.

Both the state of the world and the signals are realized at the outset of the game and are not observed by either player. We model the sender’s information acquisition as follows. At any public history of the game where the sender has acquired signals s_1, s_2, \dots, s_n and the game has not yet ended, the sender chooses between acquiring an additional signal s_{n+1} and irreversibly stopping signal acquisition. If she chooses to acquire signal s_{n+1} , it is publicly observed by both players. If she chooses to stop, the receiver takes a binary action $a \in \{0, 1\}$, and the game ends.

The receiver updates his belief using Bayes’ rule, and he chooses $a = 1$ if and only if his posterior belief is at least some exogenous threshold $\bar{\beta} \in (\beta_0, 1)$. The sender’s payoff depends on the receiver’s action and the public history. Let $h_n = (s_1, s_2, \dots, s_n)$ denote the public history at which the sender stops information acquisition, which contains the realization of all signals acquired by the sender. Denote by $n_g(h_n)$ and $n_b(h_n)$ the number of good and bad signals in the sequence h_n , respectively. The sender’s payoff

$$v(h_n, a) = a - c_g n_g(h_n) - c_b n_b(h_n),$$

where $c_g, c_b > 0$. That is, each signal is costly for the sender to acquire, and the cost depends on the signal realization. This assumption is suitable in many applications. For example, in the pharmaceutical company example, c_g is the fixed cost of enrolling each patient in the clinical trial, and $c_b - c_g > 0$ is the cost associated with treating a patient who develops side effects; in the entrepreneur example, c_b is the fixed wage paid to the engineers, and $c_g - c_b > 0$ is the bonus the entrepreneur has to pay if the engineers successfully develop a feature. Observe that neither the sender's nor the receiver's prior belief is payoff relevant.

7.2. The sender's strategy and the cost of information. A strategy of the sender is a stopping time adapted to the natural filtration generated by public histories. Specifically, consider the following threshold strategy τ of the sender: she stops at the first history where the difference between the number of good and bad signals equals some threshold values $\underline{n} < 0$ or $\bar{n} > 0$. Notice that the sender stops in finite time with probability one. Hence, interpreting the event $n_g(h_\tau) - n_b(h_\tau) = \bar{n}$ as the good outcome and the event $n_g(h_\tau) - n_b(h_\tau) = \underline{n}$ as the bad outcome, the stopping strategy τ is equivalent to an experiment $\pi = (p, q)$ in our main model, where the probabilities p and q are the probabilities of the good outcome conditional on the good and bad state, respectively.

The sender's posterior belief on the good state induced by the stopping rule τ is a random variable $\hat{\mu}$, which equals either $\mu_{\bar{n}} := \left(1 + \frac{1-\mu}{\mu} x^{\bar{n}}\right)^{-1}$ or $\mu_{\underline{n}} := \left(1 + \frac{1-\mu}{\mu} x^{\underline{n}}\right)^{-1}$, where $x := \frac{1-\alpha}{\alpha} < 1$.

In the following proposition, we calculate the expected cost of implementing the strategy τ . We show that it equals the expected reduction in a weighted log-likelihood ratio measured by the sender's beliefs.

PROPOSITION 10. *The expected cost of implementing the strategy τ is $\mathbb{E}[c_g n_g(h_\tau) + c_b n_b(h_\tau)] = \mathbb{E}[H(\mu_0) - H(\hat{\mu})]$, where*

$$H(\mu) = -\frac{\ln x}{2\alpha - 1} \left[\bar{c}_g \mu \ln \left(\frac{1 - \mu}{\mu} \right) + \bar{c}_b (1 - \mu) \ln \left(\frac{\mu}{1 - \mu} \right) \right],$$

and $\bar{c}_g = \alpha c_g + (1 - \alpha) c_b$, $\bar{c}_b = (1 - \alpha) c_g + \alpha c_b$.

Hence, letting $C_g = -\frac{\ln x}{2\alpha - 1} \bar{c}_g$ and $C_b = -\frac{\ln x}{2\alpha - 1} \bar{c}_b$, the cost of experiment $\pi = (p, q)$ takes the form we assume in section 2.

Notice that for a given experiment π , its cost $c(\pi|\mu)$ is a linear function of the sender's prior belief, and generically (i.e., for all but a measure zero set of parameters $(p, q, C_g, C_b) \in \Pi \times \mathbb{R}_{++}^2$), it is strictly monotonic. That is, the cost of running an experiment depends on the sender's prior belief.

8. Shannon Entropy Cost

Our findings that the single-crossing property may fail in a costly persuasion game with a partially informed sender and that there exist pooling equilibria that satisfy the D1 criterion are not specific to the log-likelihood cost function we use. The same features are present in the model if we assume Shannon Entropy cost of experiments, i.e., letting $H(\mu) = -C [\mu \ln \mu + (1 - \mu) \ln(1 - \mu)]$, where $C > 0$. The Shannon entropy cost is widely used in the literature to model attention cost (cf. Sims (2003)). When using the Shannon Entropy cost function, the single-crossing property fails, and consequently, there may exist a continuum of pooling equilibria that satisfy the D1 criterion.

PROPOSITION 11. *The single-crossing property fails in the persuasion game with Shannon Entropy cost of experiments. There exists $\hat{p} : (0, 1) \rightarrow (0, 1)$ such that $q < \hat{p}(q)$, and $MRS(\pi|\mu_h) < MRS(\pi|\mu_l)$ if $p < \hat{p}(q)$, $MRS(\pi|\mu_h) = MRS(\pi|\mu_l)$ if $p = \hat{p}(q)$, and $MRS(\pi|\mu_h) > MRS(\pi|\mu_l)$ if $p > \hat{p}(q)$. A pooling persuasion equilibrium outcome $\pi^* = (p^*, q^*)$ can be supported by an equilibrium satisfying the D1 criterion if and only if $p^* = \hat{p}(q^*)$.*

9. Conclusions

We study a persuasion game with costly experiments and a partially informed sender. At the outset of the game, the sender privately observes a noisy signal (i.e., her type) about the state of the world. She then chooses a costly experiment which generates public information about the state of the world to the receiver. The receiver can infer the state of the world from the choice of experiment as well as the outcome of the experiment.

From a Wald sampling problem with heterogeneous priors, we show that the cost of experiment equals the expected reduction in a weighted log-likelihood ratio measured by the sender's beliefs. Therefore, the costs of running experiments depend on the sender's type.

We focus on equilibria of the persuasion game that satisfy the D1 criterion. We show that the D1 equilibrium outcome depends on the relative costs of drawing good and bad news in the experiment. If bad news is more costly such as in the pharmaceutical example, there exists a unique separating equilibrium outcome. The experiments chosen by the sender are more informative than the ones chosen by the sender in the public signal benchmark. Therefore, the receiver is unambiguously better off with sender private information. On the other hand, if good news is sufficiently more costly than bad news such as in the entrepreneur example, this is a signaling game where the single-crossing property fails. There may exist a continuum of pooling equilibria in addition to the essentially unique separating equilibrium.

In some pooling equilibria, the receiver's equilibrium payoff is strictly less than that when there is no sender private information.

APPENDIX A. AUXILIARY RESULTS AND PROOFS

A.1. Results relating to symmetric information benchmark. We prove two results relating to the symmetric information benchmark. Lemma A.1 shows that the equilibrium is essentially unique and gives the condition when persuasion is possible. Lemma A.2 shows that fixing the receiver's prior, the equilibrium is ranked by the sender's prior. When bad news is more costly, the equilibrium experiment is more Blackwell-informative if the sender's prior is higher; when good news is sufficiently more costly than bad news, the equilibrium experiment is less Blackwell-informative if the sender's prior is higher.

LEMMA A.1. *Persuasion is possible in the symmetric information benchmark, i.e., $V(\mu_s, \mu_r) > 0$, if and only if $\mathbf{F}(C_g, C_b, \mu_s, \mu_r) > 0$, where*

$$\begin{aligned} \mathbf{F}(C_g, C_b, \mu_s, \mu_r) := & \mu_s + (1 - \mu_s)\mathbf{Q}(\mu_r) + \mu_s C_g (\ln \mathbf{Q}(\mu_r) + 1 - \mathbf{Q}(\mu_r)) \\ & - (1 - \mu_s)C_b (\mathbf{Q}(\mu_r) \ln \mathbf{Q}(\mu_r) + 1 - \mathbf{Q}(\mu_r)). \end{aligned}$$

If $\mathbf{F}(C_g, C_b, \mu_s, \mu_r) > 0$, there exists a unique equilibrium outcome $\hat{\pi}$ which is persuasive at belief μ_r . If $\mathbf{F}(C_g, C_b, \mu_s, \mu_r) \leq 0$, an experiment is an equilibrium outcome if and only if it is uninformative. That is, the equilibrium outcome is essentially unique, and we denote it by $\hat{\pi} = \pi_u$.

PROOF. First, observe that

$$\begin{aligned} c(\pi|\mu) &= H(\mu) - (\mu p + (1 - \mu)q)H(\mu_g) - (\mu(1 - p) + (1 - \mu)(1 - q))H(\mu_b) \\ &= C_g \left[\mu \ln \frac{1 - \mu}{\mu} - \mu p \ln \frac{(1 - \mu)q}{\mu p} - \mu(1 - p) \ln \frac{(1 - \mu)(1 - q)}{\mu(1 - p)} \right] \\ &\quad + C_b \left[(1 - \mu) \ln \frac{\mu}{1 - \mu} - (1 - \mu)q \ln \frac{\mu p}{(1 - \mu)q} - (1 - \mu)(1 - q) \ln \frac{\mu(1 - p)}{(1 - \mu)(1 - q)} \right] \\ &= C_g \left[-\mu p \ln \frac{q}{p} - \mu(1 - p) \ln \frac{1 - q}{1 - p} \right] + C_b \left[-(1 - \mu)q \ln \frac{p}{q} - (1 - \mu)(1 - q) \ln \frac{1 - p}{1 - q} \right] \end{aligned}$$

is convex in p and q , where $\mu_g = \mathbf{B}(\mu, \pi, g)$ and $\mu_b = \mathbf{B}(\mu, \pi, b)$. Hence, $f(\pi, \mu_s)$ is concave in p and q . Therefore, if $V(\mu_s, \mu_r) > 0$, there exists a unique experiment $\hat{\pi} = (\hat{p}, \hat{q})$ that obtains the maximum. Moreover, it must be the case that $\frac{\hat{q}}{\hat{p}} = \mathbf{Q}(\mu_r)$. Assume by way of contradiction that $\frac{\hat{q}}{\hat{p}} < \mathbf{Q}(\mu_r)$. Then there exists a linear combination $\tilde{\pi} = (\tilde{p}, \tilde{q})$ of $\hat{\pi}$ and $(1, 1)$ such that $\frac{\tilde{q}}{\tilde{p}} \leq \mathbf{Q}(\mu_r)$. Since $f((1, 1), \mu_s) = 1 > f(\hat{\pi}, \mu_s)$, concavity of $f(\cdot, \mu_s)$ implies that $f(\tilde{\pi}, \mu_s) > f(\hat{\pi}, \mu_s)$. This is a contradiction to that $\hat{\pi}$ solves the sender's maximization problem.

Hence, the sender's equilibrium value in the symmetric information benchmark can be simplified to

$$V(\mu_s, \mu_r) = \max_{p \in [0,1]} f((p, \mathbf{Q}(\mu_r)p), \mu_s).$$

This single variable convex optimization problem can be solved by the first order condition $\frac{d}{dp} f((p, \mathbf{Q}(\mu_r)p), \mu_s) = 0$. To shorten notation, we will denote the derivative briefly by $f'(p)$, and $\mathbf{Q}(\mu_r)$ by Q . Notice that

$$f'(p) = \mu_s + (1 - \mu_s)Q + \mu_s C_g \left[\ln Q - \ln \frac{1 - Qp}{1 - p} + \frac{1 - Q}{1 - Qp} \right] - (1 - \mu_s)C_b \left[Q \ln Q - Q \ln \frac{1 - Qp}{1 - p} + \frac{1 - Q}{1 - p} \right],$$

and as $p \uparrow 1$, $f'(p)$ tends to negative infinity. Hence, the first order condition admits a unique interior solution $\hat{p} \in (0, 1)$ if and only if

$$f'(0) = \mathbf{F}(C_g, C_b, \mu_s, \mu_r) > 0.$$

If this is satisfied, there exists a unique equilibrium outcome where the sender chooses the experiment $\hat{\pi} = (\hat{p}, \mathbf{Q}(\mu_r)\hat{p})$, and the sender's equilibrium payoff $V(\mu_s, \mu_r) > 0$. If $\mathbf{F}(C_g, C_b, \mu_s, \mu_r) \leq 0$, $f'(p) < 0$ for all $p \in (0, 1)$, hence the maximization problem has a unique corner solution $\hat{p} = 0$, and $V(\mu_s, \mu_r) = f((0, 0), \mu_s) = 0$. In this case, the sender chooses an uninformative experiment in any equilibrium, and any uninformative experiment can be supported by an equilibrium. By abuse of notation, we denote the essentially unique equilibrium outcome in this case by $\hat{\pi} = \pi_u$. \square

Notice that $\ln Q + 1 - Q < 0$, and $Q \ln Q + 1 - Q > 0$ for all $Q \in (0, 1)$. Therefore, $\mathbf{F}(C_g, C_b, \mu_s, \mu_r) > 0$ if and only if the cost parameters C_g and C_b are small. Fixing μ_r , or equivalently fixing $Q = \mathbf{Q}(\mu_r)$, figure A.1 illustrates the result. Given some μ_s , $V(\mu_s, \mu_r) > 0$ if and only if (C_b, C_g) is below the solid line. As the sender's prior belief increases to $\mu'_s > \mu_s$, the boundary of the region rotates around $\left(\frac{Q}{Q \ln Q + 1 - Q}, -\frac{1}{\ln Q + 1 - Q} \right)$, and $V(\mu'_s, \mu_r) > 0$ if and only if (C_b, C_g) is below the dashed line.

Hence, if $C_g \geq -\frac{1}{\ln Q + 1 - Q}$ and $C_b \geq \frac{Q}{Q \ln Q + 1 - Q}$, $\hat{\pi} = \pi_u$ for all $\mu_s \in (0, 1)$; if $C_g \geq -\frac{1}{\ln Q + 1 - Q}$ and $C_b < \frac{Q}{Q \ln Q + 1 - Q}$, $\hat{\pi} = \pi_u$ if and only if μ_s is sufficiently large; if $C_g < -\frac{1}{\ln Q + 1 - Q}$ and $C_b \geq \frac{Q}{Q \ln Q + 1 - Q}$, $\hat{\pi} = \pi_u$ if and only if μ_s is sufficiently small. The following lemma generalizes this finding. Fixing all other parameters, the unique equilibrium outcome in the symmetric information benchmark $\hat{\pi}$ is ranked by the sender's prior μ_s . Whether it is more or less Blackwell-informative as μ_s increases depends on the relative costs of good and bad




FIGURE A.1: Region of persuasion and monotone comparative statics in the symmetric information benchmark

news.

LEMMA A.2. *There exists an increasing and concave function $\hat{K} : \mathbb{R}_{++} \rightarrow \mathbb{R}_{++}$ such that, fixing any μ_r , the unique equilibrium outcome $\hat{\pi}$ is independent of the sender's prior μ_s if and only if $C_g = \hat{K}(C_b)$. Moreover, if $C_g > \hat{K}(C_b)$, the unique equilibrium outcome $\hat{\pi}$ is less Blackwell-informative as the sender's prior μ_s increases; if $C_g < \hat{K}(C_b)$, the unique equilibrium outcome $\hat{\pi}$ is more Blackwell-informative as the sender's prior μ_s increases.*

PROOF. Fix an arbitrary μ_r . Consider first the case $\mathbf{F}(C_g, C_b, \mu_s, \mu_r) > 0$. That is, $\hat{\pi} = (\hat{p}, Q\hat{p})$ where $\hat{p} \in (0, 1)$ is solved by the first order condition $f'(p) = 0$. Notice that

$$f'(p) = \underbrace{\mu_s \left(1 + C_g \left[\ln Q - \ln \frac{1 - Qp}{1 - p} + \frac{1 - Q}{1 - Qp} \right] \right)}_{=:M} + (1 - \mu_s) \underbrace{\left(Q - C_b \left[Q \ln Q - Q \ln \frac{1 - Qp}{1 - p} + \frac{1 - Q}{1 - p} \right] \right)}_{=:N}.$$

Hence, evaluated at $p = \hat{p}$, M and N have different signs. Therefore, \hat{p} is independent of μ_s , i.e., $\frac{\partial}{\partial \mu_s} f'(\hat{p}) = M - N = 0$, if and only if $M = N = 0$.

Given any $C_b > 0$, $M = N = 0$ defines a system of equations of p and C_g , which admits

a unique solution

$$p = 1 - \frac{1-Q}{Q} \frac{1}{\hat{t}(C_b)}, \quad C_g = \hat{K}(C_b) := - \left(\frac{1}{C_b} - \frac{\hat{t}(C_b)^2}{1 + \hat{t}(C_b)} \right)^{-1},$$

where $\hat{t}(C_b) > 0$ is the unique solution to $t - \ln(1+t) = \frac{1}{C_b}$. Notice that the solution of C_g does not depend on μ_r or Q . Thus, it defines a function \hat{K} of C_b . If $C_g > \hat{K}(C_b)$, $M < 0 < N$ at $p = \hat{p}$, so $\frac{\partial}{\partial \mu_s} f'(\hat{p}) < 0$, and \hat{p} is strictly decreasing in μ_s . If $C_g = \hat{K}(C_b)$, $M = N = 0$ at $p = \hat{p}$, so $\frac{\partial}{\partial \mu_s} f'(\hat{p}) = 0$, and \hat{p} is independent of μ_s . If $C_g < \hat{K}(C_b)$, $M > 0 > N$ at $p = \hat{p}$, so $\frac{\partial}{\partial \mu_s} f'(\hat{p}) > 0$, and \hat{p} is strictly increasing in μ_s .

Moreover, $\hat{t}\left(\frac{Q}{Q \ln Q + 1 - Q}\right) = \frac{1}{Q} - 1$, and $\hat{K}\left(\frac{Q}{Q \ln Q + 1 - Q}\right) = -\frac{1}{\ln Q + 1 - Q}$. Hence, we can combine the case $\mathbf{F}(C_g, C_b, \mu_s, \mu_r) \leq 0$ and state that \hat{p} is independent of μ_s if $C_g = \hat{K}(C_b)$, weakly increasing in μ_s if $C_g < \hat{K}(C_b)$, and weakly decreasing in μ_s if $C_g > \hat{K}(C_b)$. Equivalently, the unique equilibrium outcome $\hat{\pi}$ is independent of μ_s if $C_g = \hat{K}(C_b)$, more Blackwell-informative as μ_s increases if $C_g < \hat{K}(C_b)$, and less Blackwell-informative as μ_s increases if $C_g > \hat{K}(C_b)$. Since μ_r is arbitrary, the results hold for all $\mu_r \in (0, 1)$.

The rest of the proof, namely \hat{K} being increasing and concave, is done in the proof of proposition 3. \square

The function \hat{K} reappears in proposition 3, where we show that the single-crossing property in the persuasion game with sender private information fails if and only if $C_g > \hat{K}(C_b)$. The highlighted curve in figure A.1 plots the function.

A.2. Proof of Lemma 1.

PROOF. Let $\bar{U}(\mu)$ be the receiver's equilibrium payoff in the symmetric information benchmark where $\mu_s = \mu_r = \mu$. We are to show that \bar{U} is piecewise linear and weakly convex, hence

$$U^{ni} = \bar{U}(\mu_0) \leq \frac{\mu_0 - \mu_l}{\mu_h - \mu_l} \bar{U}(\mu_h) + \frac{\mu_h - \mu_0}{\mu_h - \mu_l} \bar{U}(\mu_l) = U^{ps}.$$

Given $\mu < \bar{\beta}$, in the symmetric information benchmark with common prior μ , the sender solves

$$\max_{x \leq \mu} \frac{\mu - x}{\bar{\beta} - x} - \left(H(\mu) - \frac{\mu - x}{\bar{\beta} - x} H(\bar{\beta}) - \frac{\bar{\beta} - \mu}{\bar{\beta} - x} H(x) \right),$$

which admits a unique solution x^* . If $x^* < \mu$, there is a unique equilibrium outcome where the sender chooses an experiment that splits the players' belief to x^* and $\bar{\beta}$. If $x^* = \mu$, the sender chooses an uninformative experiment, and persuasion is not possible in equilibrium.

The first order condition of the sender's problem is

$$\frac{\bar{\beta} - \mu}{(\bar{\beta} - x)^2} [1 + H(\bar{\beta}) - H(x)] - \frac{\bar{\beta} - \mu}{\bar{\beta} - x} H'(x) = 0,$$

that is,

$$(A.1) \quad \int_x^{\bar{\beta}} (H'(x) - H'(z)) dz = 1.$$

Notice that the left-hand side of (A.1) is decreasing in x , and as $x \rightarrow 0$, it goes to infinity. Hence, (A.1) admits a unique solution $\underline{\beta} < \bar{\beta}$ which is independent of μ . If $\underline{\beta} < \mu < \bar{\beta}$, $x^* = \underline{\beta}$ solves the sender's problem; if $\mu < \underline{\beta}$, the sender's problem has a corner solution $x^* = \mu$.

To summarize, if $\underline{\beta} < \mu < \bar{\beta}$, the unique equilibrium outcome splits the players' belief to $\underline{\beta}$ and $\bar{\beta}$, and the receiver's expected payoff $\bar{U}(\mu) = \frac{\mu - \underline{\beta}}{\bar{\beta} - \underline{\beta}} \bar{U}(\bar{\beta})$; if $\mu \leq \underline{\beta}$, persuasion is not possible, and the receiver's payoff $\bar{U}(\mu) = 0$. Together with the observation that when $\mu \geq \bar{\beta}$, the receiver chooses $a = 1$ and gets $\bar{U}(\mu) = \mu - \frac{\beta^*}{1 - \beta^*} (1 - \mu)$, these lead to the conclusion that \bar{U} is piecewise linear and weakly convex. \square

A.3. Proof of Proposition 2.

PROOF. *The "only if" part.* Given a pooling equilibrium where the sender chooses experiment $\pi = (p, q)$, the receiver's interim belief $\beta(\pi) = \mu_0$. If it is a persuasion equilibrium, $\frac{q}{p} \leq \mathbf{Q}(\mu_0)$. Otherwise, the receiver's posterior belief $\mathbf{B}(\mu_0, \pi, g) < \bar{\beta}$ even after the good outcome, and the receiver always chooses the low action on the equilibrium path, which contradicts the equilibrium being a persuasion equilibrium.

The type θ sender's ex ante expected payoff in the equilibrium must be at least $\hat{V}(\mu_\theta)$. For a persuasion equilibrium, the sender's equilibrium payoff equals $f(\mu_\theta, \pi)$, hence $f(\pi, \mu_\theta) \geq \hat{V}(\mu_\theta)$. For a trivial equilibrium, the sender's equilibrium payoff is zero, but $\hat{V}(\mu_\theta) \geq 0$, so it must be the case that $\hat{V}(\mu_\theta) = 0$ and π is uninformative.

The "if" part. Let $\pi = (p, q)$ be such that $f(\pi, \mu_\theta) \geq V(\mu_\theta)$ for all $\theta \in \Theta$, and $\frac{q}{p} \leq \mathbf{Q}(\mu_0)$. We show by construction that this is the experiment chosen by the sender in a pooling persuasion equilibrium. Let $\pi_l = \pi_h = \pi$, $\beta(\pi) = \mu_0$, and $\beta(\pi') = \mu_l$ for all $\pi' \neq \pi$. Moreover, let $\hat{\beta}, \mathbf{a}$ be such that $\hat{\beta}(\pi', s) = \mathbf{B}(\beta(\pi'), \pi', s)$ for all $\pi' \in \Pi$ and $s \in \{g, b\}$, and $\mathbf{a}(\pi', s) = 1$ if and only if $\hat{\beta}(\pi', s) \geq \bar{\beta}$. To check that this constitutes an equilibrium, we only need to check that there is no profitable deviation for the sender. For any $\pi' \neq \pi$ that is always persuasive, the sender's payoff by deviating to π' is $f(\pi', \mu_\theta) \leq \hat{V}(\mu_\theta) \leq f(\pi, \mu_\theta)$. For

any $\pi' \neq \pi$ that is not always persuasive, it is not persuasive at belief $\beta(\pi') = \mu_l$. Therefore, the sender's payoff from deviating to π' is at most $0 \leq \bar{V}(\mu_\theta) \leq f(\pi, \mu_\theta)$. Hence, the sender does not have a profitable deviation. Moreover, on the equilibrium path, the receiver chooses the high action if the experiment π yields the good outcome, which happens with probability $\mu_0 p + (1 - \mu_0)q > 0$. Hence, the constructed is a persuasion equilibrium.

Let π be any uninformative experiment. We show that it is the experiment chosen by the sender in a pooling trivial equilibrium if $\bar{V}(\mu_\theta) = 0$ for all $\theta \in \Theta$. Using the same construction as above, let $\pi_l = \pi_h = \pi$, $\beta(\pi) = \mu_0$, and $\beta(\pi') = \mu_l$ for all $\pi' \neq \pi$. Moreover, $\hat{\beta}(\pi', s) = \mathbf{B}(\beta(\pi'), \pi', s)$ for all $\pi' \in \Pi$ and $s \in \{g, b\}$, and $\mathbf{a}(\pi', s) = 1$ if and only if $\hat{\beta}(\pi', s) \geq \bar{\beta}$. This gives a trivial equilibrium. For any π' that is always persuasive, the sender's payoff by deviating to π' is strictly negative. For any $\pi' \neq \pi$ that is not always persuasive, it is not persuasive at belief $\beta(\pi') = \mu_l$. Therefore, the sender's payoff from deviating to π' is at most zero, and there is no profitable deviation. \square

A.4. Results relating to existence of pooling equilibria. A direct corollary of proposition 2 and lemma A.1 shows that pooling trivial equilibria exist if and only if C_g and C_b are high.

COROLLARY A.3. *Pooling trivial equilibria exist if and only if $\mathbf{F}(C_g, C_b, \mu_\theta, \mu_l) \leq 0$ for all $\theta \in \Theta$.*

Similarly, a necessary condition for a pooling persuasion equilibrium to exist is that C_g and C_b are small. The following result shows that, in addition, the receiver's prior μ_0 has to be large.

PROPOSITION A.4. *A pooling persuasion equilibrium exists only if $\mathbf{F}(C_g, C_b, \mu_\theta, \mu_h) > 0$ for all $\theta \in \Theta$. Moreover, fixing $\{\mu_\theta\}_{\theta \in \Theta}$ and C_g, C_b such that $\mathbf{F}(C_g, C_b, \mu_\theta, \mu_h) > 0$ for all $\theta \in \Theta$, there exists $\bar{\mu} \geq \mu_l$ such that a pooling persuasion equilibrium exists if $\mu_0 > \bar{\mu}$, and no pooling persuasion equilibrium exists if $\mu_0 < \bar{\mu}$.*

PROOF. Fixing $\{\mu_\theta\}_{\theta \in \Theta}$, pooling persuasion equilibrium does not exist if, for all $\pi = (p, q)$ such that $\frac{q}{p} \leq \mathbf{Q}(\mu_0)$, $f(\pi, \mu_\theta) < 0$ for some $\theta \in \Theta$. Since $\mu_0 < \mu_h$, a necessary condition that a pooling persuasion equilibrium exists is $\mathbf{F}(C_g, C_b, \mu_\theta, \mu_h) > 0$ for all $\theta \in \Theta$. That is, C_g and C_b are small.

Suppose that C_g and C_b are such that $\mathbf{F}(C_g, C_b, \mu_\theta, \mu_h) > 0$ for all $\theta \in \Theta$. We are to show that a pooling persuasion equilibrium exists if and only if μ_0 is sufficiently high. We consider two separate cases.

First, if $\pi_l^* = \pi_h^* = \pi_u$, i.e., if $\mathbf{F}(C_g, C_b, \mu_\theta, \mu_l) \leq 0$ for all $\theta \in \Theta$, letting

$$\bar{\mu} := \inf\{x : \mathbf{F}(C_g, C_b, \mu_\theta, x) > 0, \forall \theta \in \Theta\},$$

a pooling persuasion equilibrium exists if and only if $\mu_0 > \bar{\mu}$. By assumption, $\bar{\mu} \in [\mu_l, \mu_h]$.

Second, if $\hat{\pi}_\theta \neq \pi_u$ for at least some $\theta \in \Theta$, let $U_\theta = \{\pi : f(\pi, \mu_\theta) \geq \bar{V}(\mu_\theta)\}$. It is a closed and convex set, and therefore, so is $U_l \cap U_h$. Moreover, $U_l \cap U_h$ is nonempty, since $(1, 1) \in U_l \cap U_h$. Let

$$\bar{\mu} = \min \left\{ \mu_0 : \exists (p, q) \in U_l \cap U_h, \frac{q}{p} \leq \mathbf{Q}(\mu_0) \right\} \geq \mu_l.$$

Let $\pi^* = (p^*, q^*) \in U_l \cap U_h$ be such that $\frac{q^*}{p^*} \leq \mathbf{Q}(\bar{\mu})$. By convexity of $U_l \cap U_h$, any convex combination of π^* and $(1, 1)$ belongs to $U_l \cap U_h$. Hence, for all $\mu_0 \geq \bar{\mu}$, there exists $(p, q) \in U_l \cap U_h$ such that $\frac{q}{p} \leq \mathbf{Q}(\mu_0)$. Therefore, by proposition 2, a pooling persuasion equilibrium exists if and only if $\mu_0 \geq \bar{\mu}$. Moreover, $\bar{\mu} = \mu_l$ if and only if $\pi_l^* = \pi_h^* \neq \pi_u$, which happens only on a Lebesgue measure zero set of parameters satisfying $C_g = \hat{K}(C_b)$. \square

A.5. Proof of Proposition 3.

PROOF. Given $\mu \in (0, 1)$ and $\pi = (p, q) \in \Pi$ such that $1 > p \geq q > 0$, let $\mu_g = \mathbf{B}(\mu, \pi, g)$, and $\mu_b = \mathbf{B}(\mu, \pi, b)$. That is, μ_g and μ_b are the posterior beliefs induced by the good and bad outcome from experiment π given prior μ , respectively. Observe that

$$f(\pi, \mu) = \mu p + (1 - \mu)q - H(\mu) + (\mu p + (1 - \mu)q)H(\mu_g) + (\mu(1 - p) + (1 - \mu)(1 - q))H(\mu_b).$$

Taking derivatives,

$$\begin{aligned} \frac{\partial f(\pi, \mu)}{\partial p} &= \mu [1 + H(\mu_g) + (1 - \mu_g)H'(\mu_g) - H(\mu_b) - (1 - \mu_b)H'(\mu_b)] \\ &= \mu \left[1 + C_g \left(\ln \frac{1 - \mu_g}{\mu_g} - \ln \frac{1 - \mu_b}{\mu_b} \right) + C_b \left(\frac{1 - \mu_g}{\mu_g} - \frac{1 - \mu_b}{\mu_b} \right) \right] \\ &= \mu - \mu C_g \ln \frac{p(1 - q)}{(1 - p)q} - (1 - \mu)C_b \frac{p - q}{p(1 - p)} \\ &= \underbrace{-C_b \frac{p - q}{p(1 - p)}}_{=: A_1} - \mu \underbrace{\left[-1 + C_g \ln \frac{p(1 - q)}{(1 - p)q} - C_b \frac{p - q}{p(1 - p)} \right]}_{=: A_2}, \end{aligned}$$

and similarly,

$$\begin{aligned}
\frac{\partial f(\pi, \mu)}{\partial q} &= (1 - \mu) [1 + H(\mu_g) - \mu_g H'(\mu_g) - H(\mu_b) + \mu_b H'(\mu_b)] \\
&= (1 - \mu) \left[1 + C_g \left(\frac{\mu_g}{1 - \mu_g} - \frac{\mu_b}{1 - \mu_b} \right) + C_b \left(\ln \frac{\mu_g}{1 - \mu_g} - \ln \frac{\mu_b}{1 - \mu_b} \right) \right] \\
&= 1 - \mu + \mu C_g \frac{p - q}{q(1 - q)} + (1 - \mu) C_b \ln \frac{p(1 - q)}{(1 - p)q} \\
&= \underbrace{1 + C_b \ln \frac{p(1 - q)}{(1 - p)q}}_{=: A_3} + \mu \underbrace{\left[-1 + C_g \frac{p - q}{p(1 - p)} - C_b \ln \frac{p(1 - q)}{(1 - p)q} \right]}_{=: A_4}.
\end{aligned}$$

Therefore,

$$MRS(\pi|\mu) = -\frac{\partial f(\pi, \mu)/\partial p}{\partial f(\pi, \mu)/\partial q} = \frac{A_1 + A_2\mu}{A_3 + A_4\mu}.$$

Noticing that $\frac{\partial f(\pi, \mu)}{\partial q} > 0$, that is, the denominator $A_3 + A_4\mu > 0$. Therefore, for any $\mu_l < \mu_h$, $MRS(\pi|\mu_h) - MRS(\pi|\mu_l)$ has the same sign as

$$\Delta := A_2 A_3 - A_1 A_4 = (C_g - C_b) \ln \frac{p(1 - q)}{(1 - p)q} + C_g C_b \left[\left(\ln \frac{p(1 - q)}{(1 - p)q} \right)^2 - \frac{(p - q)^2}{pq(1 - p)(1 - q)} \right] - 1.$$

Notice that Δ does not depend on μ_h or μ_l . Moreover, as $p \downarrow q$, $\Delta \rightarrow -1 < 0$, and as $p \uparrow 1$, $\Delta \rightarrow -\infty$.

Fix $q \in (0, 1)$. Taking derivative with respect to p ,

$$(A.2) \quad \frac{\partial \Delta}{\partial p} = \frac{C_g - C_b}{p(1 - p)} + \frac{C_g C_b}{p(1 - p)} \left[2 \ln \frac{p(1 - q)}{(1 - p)q} - \frac{p - q}{p(1 - q)} - \frac{p - q}{(1 - p)q} \right].$$

Notice that the bracket in (A.2) is strictly negative for all $p > q$, since

$$\frac{\partial}{\partial p} \left[2 \ln \frac{p(1 - q)}{(1 - p)q} - \frac{p - q}{p(1 - q)} - \frac{p - q}{(1 - p)q} \right] = \frac{2}{p(1 - p)} - \frac{q}{p^2(1 - q)} - \frac{1 - q}{(1 - p)^2 q} < 0,$$

and evaluated at $p = q$, it equals zero. Hence, if $C_g \leq C_b$, Δ is strictly decreasing in p and therefore negative for all $p \geq q$. That is, $MRS(\pi|\mu_h) < MRS(\pi|\mu_l)$ holds for all $\pi \in \Pi^\circ$, and the single-crossing property holds.

Moreover, notice that the bracket in (A.2) goes to negative infinity as $p \uparrow 1$. Hence, if $C_g > C_b$, Δ is first increasing in p when $p \in (q, \mathbf{p}(q))$ and decreasing in p when $p \in (\mathbf{p}(q), 1)$,

where $p = \mathbf{p}(q)$ is the unique solution of

$$(A.3) \quad 2 \ln \frac{p(1-q)}{(1-p)q} - \frac{p-q}{p(1-q)} - \frac{p-q}{(1-p)q} = -\frac{C_g - C_b}{C_g C_b}.$$

Letting $t := \frac{p(1-q)}{(1-p)q} > 1$, we can rewrite (A.3) as

$$(A.4) \quad 2 \ln t + \frac{1}{t} - t + \frac{C_g - C_b}{C_g C_b} = 0.$$

The left-hand side of (A.4) is strictly decreasing in t ; as $t \downarrow 1$, it tends to $\frac{C_g - C_b}{C_g C_b} > 0$, and as $t \rightarrow \infty$, it tends to negative infinity. Hence, (A.4) has a unique solution $t^* > 1$, which is increasing in C_g and decreasing in C_b . Therefore,

$$(A.5) \quad \mathbf{p}(q) = \frac{t^* q}{1 + (t^* - 1)q} \in (q, 1).$$

Using (A.5) to evaluate Δ at $p = \mathbf{p}(q)$, we have

$$\Delta^* := (C_g - C_b) \ln t^* + C_g C_b \left[(\ln t^*)^2 - \frac{(t^* - 1)^2}{t^*} \right] - 1.$$

Notice that Δ^* does not depend on q . That is, fixing any $q \in (0, 1)$ and varying p on $(q, 1)$, Δ first increases, obtains its maximum Δ^* at $p = \mathbf{p}(q)$, and then decreases. Applying envelope theorem,

$$\frac{\partial \Delta^*}{\partial C_g} = \ln t^* + C_b \left[(\ln t^*)^2 - \frac{(t^* - 1)^2}{t^*} \right] \geq \frac{\Delta^* + 1}{C_g} > 0.$$

Notice that $y'(x) = \frac{y(x)+1}{x}$ solves a linear function. Therefore, Δ^* increases in C_g at least as fast as a linear function. Given any $C_b > 0$, as $C_g \downarrow C_b$, $\Delta^* \rightarrow -1$. Therefore, there exists a unique $C_g > C_b$ such that $\Delta^* = 0$. Denote this unique C_g by $\hat{K}(C_b)$ (we later check that \hat{K} is the same function as derived in lemma A.2). Notice that $(\ln t)^2 - \frac{(t-1)^2}{t} < 0$ for all $t > 1$. Therefore, by envelope theorem,

$$\frac{\partial \Delta^*}{\partial C_b} = -\ln t^* + C_g \left[(\ln t^*)^2 - \frac{(t^* - 1)^2}{t^*} \right] < 0$$

Hence, \hat{K} is strictly increasing.

If $C_g \leq \hat{K}(C_b)$, $MRS(\pi|\mu_h) \leq MRS(\pi|\mu_l)$ for all π in the interior of Π , and the single-crossing property is satisfied. If $C_g > \hat{K}(C_b)$, $\Delta^* > 0$. Hence, for any $\pi = (p, q)$ such that $p = \mathbf{p}(q)$, $MRS(\pi|\mu_h) > MRS(\pi|\mu_l)$, hence the single-crossing property does not hold.

We are left to check that \hat{K} is the same function as derived in appendix A.1. We do so by showing that for all $Q \in (0, 1)$, $\hat{K}\left(\frac{Q}{Q \ln Q + 1 - Q}\right) = -\frac{1}{\ln Q + 1 - Q}$. Substituting $C_g = -\frac{1}{\ln Q + 1 - Q}$ and $C_b = \frac{Q}{Q \ln Q + 1 - Q}$ in (A.4) yields the unique solution $t^* = \frac{1}{Q}$, hence

$$\Delta^* = \frac{(2Q \ln Q + 1 - Q^2) \ln Q - Q \left[(\ln Q)^2 - \frac{(1-Q)^2}{Q} \right]}{(\ln Q + 1 - Q)(Q \ln Q + 1 - Q)} - 1 = 0.$$

That is, $\hat{K}\left(\frac{Q}{Q \ln Q + 1 - Q}\right) = -\frac{1}{\ln Q + 1 - Q}$. \square

A.6. Proof of Proposition 5.

PROOF. If $C_g > \hat{K}(C_b)$, $\Delta^* > 0$. Hence, fixing $q \in (0, 1)$, Δ as a function of p has two zeros, denoted $\hat{\mathbf{p}}(q)$ and $\check{\mathbf{p}}(q)$, such that $q < \hat{\mathbf{p}}(q) < \mathbf{p}(q) < \check{\mathbf{p}}(q) < 1$. At $p = \hat{\mathbf{p}}(q)$ or $p = \check{\mathbf{p}}(q)$, $MRS(\pi|\mu_h) = MRS(\pi|\mu_l)$. If $p < \hat{\mathbf{p}}(q)$ or $p > \check{\mathbf{p}}(q)$, $MRS(\pi|\mu_h) < MRS(\pi|\mu_l)$; if $\hat{\mathbf{p}}(q) < p < \check{\mathbf{p}}(q)$, $MRS(\pi|\mu_h) > MRS(\pi|\mu_l)$. Therefore, at $p = \hat{\mathbf{p}}(q)$, the high-type sender's indifference curve is more convex than the low-type sender's indifference curve, while at $p = \check{\mathbf{p}}(q)$, the high-type sender's indifference is more convex than the high-type sender's.

We are left to show that that $\hat{\mathbf{p}}(q)$ is strictly decreasing in C_g and increasing in C_b . By definition, $p = \hat{\mathbf{p}}(q)$ is a solution of $\Delta = 0$, and at $p = \hat{\mathbf{p}}(q)$, $\frac{\partial \Delta}{\partial p} > 0$. Notice that Δ is strictly decreasing in C_b . Therefore, $p = \hat{\mathbf{p}}(q)$ is strictly increasing in C_b . On the other hand,

$$(A.6) \quad \frac{\partial \Delta}{\partial C_g} = \ln \frac{p(1-q)}{(1-p)q} + C_b \left[\left(\ln \frac{p(1-q)}{(1-p)q} \right)^2 - \frac{(p-q)^2}{pq(1-p)(1-q)} \right].$$

Using the fact that $\Delta = 0$ at $p = \hat{\mathbf{p}}(q)$ to replace the bracket in (A.6), we have

$$\frac{\partial \Delta}{\partial C_g} = \frac{C_b}{C_g} \ln \frac{\hat{\mathbf{p}}(q)(1-q)}{(1-\hat{\mathbf{p}}(q))q} > 0$$

at $p = \hat{\mathbf{p}}(q)$. Therefore, $\hat{\mathbf{p}}(q)$ is decreasing in C_g . \square

A.7. Results relating to existence of pooling trivial equilibria that satisfy the D1 criterion. Suppose that $\bar{V}(\mu_\theta) = 0$ for all $\theta \in \Theta$. That is, $\mathbf{F}(C_g, C_b, \mu_\theta, \mu_l) \leq 0$ for all $\theta \in \Theta$. Then there exists pooling and separating trivial equilibria.

If there also exists a separating persuasion equilibrium where the high-type sender chooses experiment $\pi_h = (p, q)$. By proposition 7, $\frac{q}{p} \leq \mathbf{Q}(\mu_h)$, and $f(\pi_h, \mu_h) > 0 > f(\pi_h, \mu_l)$. Consider any trivial equilibrium. Both types of the sender receive zero payoff in equilibrium. Consider π_h as a deviation of the sender. Since $f(\pi_h, \mu_h) > 0$, π_h is a profitable deviation for the high-type sender if it is persuasive at belief $\beta(\pi_h)$. That is, $D_h(\pi_h) = [\mathbf{Q}^{-1}(\frac{q}{p}), \mu_h] \neq \emptyset$.

On the other hand, since $f(\pi_h, \mu_l) < 0$, π_h is strictly equilibrium dominated for the low-type sender regardless of the receiver's interim belief, i.e., $D_l^0(\pi_h) = \emptyset$. Hence, the D1 criterion requires that $\beta(\pi_h) = \mu_h$, which cannot be supported by any trivial equilibrium. Since $\mu_h \in D_h(\pi_h)$, π_h is a profitable deviation for the high-type sender given such off-path belief.

If there does not exist a separating persuasion equilibrium, then for all $\pi = (p, q)$ such that $f(\pi, \mu_h) > 0$, either $\frac{q}{p} > \mathbf{Q}(\mu_h)$, or $f(\pi, \mu_l) > 0$. Hence, fixing any trivial equilibrium where the receiver uses the critical off-path belief and any deviation π' that is not uninformative, the D1 criterion either requires that $\beta(\pi') = \mu_l$, which is consistent with the receiver's critical off-path belief, or the D1 criterion is muted about the off-path belief $\beta(\pi')$. Hence, the trivial equilibrium satisfies the D1 criterion.

By lemma A.1 and proposition 7, when $\bar{V}(\mu_\theta) = 0$ for all $\theta \in \Theta$, a separating persuasion equilibrium exists if and only if $C_g \leq \hat{K}(C_b)$ and $\mathbf{F}(C_g, C_g, \mu_l, \mu_h) > 0$.

A.8. Proof of Proposition 10.

PROOF. For $\underline{n} \leq n \leq \bar{n}$, let z_n be the expected cost of acquiring additional signals before the difference between the number of g 's and b 's reaches either threshold, conditional on the state being good and the current difference being n . $\{z_n\}$ satisfies the following recurrence relation:

$$(A.7) \quad z_n = [\alpha c_g + (1 - \alpha)c_b] + \alpha z_{n+1} + (1 - \alpha)z_{n-1}$$

for all $\underline{n} < n < \bar{n}$ and the boundary conditions $z_{\underline{n}} = z_{\bar{n}} = 0$. Letting $\bar{c}_g = \alpha c_g + (1 - \alpha)c_b$, (A.7) can be rewritten as

$$\alpha \left(z_{n+1} + \frac{n+1}{2\alpha-1} \bar{c}_g \right) - \left(z_n + \frac{n}{2\alpha-1} \bar{c}_g \right) + (1-\alpha) \left(z_{n-1} + \frac{n-1}{2\alpha-1} \bar{c}_g \right) = 0.$$

Therefore,

$$z_n = C_1 x^n + C_2 - \frac{n}{2\alpha-1} \bar{c}_g,$$

where

$$C_1 = -\frac{\bar{c}_g}{2\alpha-1} \frac{\bar{n} - \underline{n}}{x^{\underline{n}} - x^{\bar{n}}}, \quad C_2 = \frac{\bar{c}_g}{2\alpha-1} \frac{\bar{n}x^{\underline{n}} - \underline{n}x^{\bar{n}}}{x^{\underline{n}} - x^{\bar{n}}}$$

are solved using the boundary conditions. Hence, the expected cost of implementing the threshold stopping rule τ conditional on the good state, i.e., $\mathbb{E}[c_g n_g(h_\tau) + c_b n_b(h_\tau) | \omega = G]$, is

$$z_0 = \frac{\bar{c}_g}{2\alpha-1} \frac{\bar{n}(x^{\underline{n}} - 1) - \underline{n}(x^{\bar{n}} - 1)}{x^{\underline{n}} - x^{\bar{n}}}.$$

Similarly, conditional on the bad state, the expected cost of implementing τ , i.e., $\mathbb{E}[c_g n_g(h_\tau) +$

$c_b n_b(h_\tau) | \omega = B$], is

$$\frac{\bar{c}_b}{2\alpha - 1} \frac{\bar{n}(x^{\bar{n}} - x^{\bar{n}+n}) - \underline{n}(x^n - x^{\bar{n}+n})}{x^n - x^{\bar{n}}},$$

where $\bar{c}_b = \alpha c_b + (1 - \alpha)c_g$. Hence, the expected cost of implementing the strategy τ , $\mathbb{E}[c_g n_g(h_\tau) + c_b n_b(h_\tau)]$, is

$$\frac{1}{2\alpha - 1} \left[\mu \frac{\bar{n}(x^n - 1) - \underline{n}(x^{\bar{n}} - 1)}{x^n - x^{\bar{n}}} \bar{c}_g + (1 - \mu) \frac{\bar{n}(x^{\bar{n}} - x^{\bar{n}+n}) - \underline{n}(x^n - x^{\bar{n}+n})}{x^n - x^{\bar{n}}} \bar{c}_b \right].$$

It can be equivalently written as $\mathbb{E}[H(\mu_0) - H(\hat{\mu})]$, where

$$H(\mu) = -\frac{\ln x}{2\alpha - 1} \left[\bar{c}_g \mu \ln \left(\frac{1 - \mu}{\mu} \right) + \bar{c}_b (1 - \mu) \ln \left(\frac{\mu}{1 - \mu} \right) \right],$$

and $\hat{\mu}$ is the sender's posterior belief at stopping calculated by Bayes' rule, which is a random variable supported on $\{\mu_{\underline{n}}, \mu_{\bar{n}}\}$. \square

REFERENCES

- Alonso, R. and Câmara, O. (2016). Bayesian persuasion with heterogeneous priors. *Journal of Economic Theory*, 165:672–706.
- Alonso, R. and Câmara, O. (2018). On the value of persuasion by experts. *Journal of Economic Theory*, 174:103–123.
- Banks, J. S. and Sobel, J. (1987). Equilibrium selection in signaling games. *Econometrica*, 55(3):647.
- Blackwell, D. (1953). Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272.
- Bloedel, A. W. and Zhong, W. (2020). The cost of optimally acquired information. *Working Paper*.
- Brocas, I. and Carrillo, J. D. (2007). Influence through ignorance. *The RAND Journal of Economics*, 38(4):931–947.
- Caplin, A., Dean, M., and Leahy, J. (2018). Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Working Paper*.
- Cho, I.-K. and Kreps, D. M. (1987). Signaling games and stable equilibria. *The Quarterly Journal of Economics*, 102(2):179.
- Degan, A. and Li, M. (2021). Persuasion with costly precision. *Economic Theory*.
- Gentzkow, M. and Kamenica, E. (2014). Costly persuasion. *American Economic Review*, 104(5):457–462.
- Hedlund, J. (2017). Bayesian persuasion by a privately informed sender. *Journal of Economic Theory*, 167:229–268.
- Henry, E. and Ottaviani, M. (2019). Research and the approval process: The organization of persuasion. *American Economic Review*, 109(3):911–955.

- Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.
- Koessler, F. and Skreta, V. (2021). Information design by an informed designer. *Working Paper*.
- Li, H. and Li, W. (2013). Misinformation. *International Economic Review*, 54(1):253–277.
- Morris, S. and Strack, P. (2019). The wald problem and the relation of sequential sampling and ex-ante information costs. *Working Paper*.
- Myerson, R. B. (1983). Mechanism design by an informed principal. *Econometrica*, 51(6):1767.
- Perez-Richet, E. (2014). Interim bayesian persuasion: First steps. *American Economic Review*, 104(5):469–474.
- Pomatto, L., Strack, P., and Tamuz, O. (2020). The cost of information. *Working Paper*.
- Rothschild, M. and Stiglitz, J. (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, 90(4):629.
- Sims, C. A. (1998). Stickiness. *Carnegie-Rochester Conference Series on Public Policy*, 49:317–356.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.
- Wald, A. (1945). Sequential method of sampling for deciding between two courses of action. *Journal of the American Statistical Association*, 40(231):277–306.
- Wilson, C. (1977). A model of insurance markets with incomplete information. *Journal of Economic Theory*, 16(2):167–207.