A Theorem on Bayesian Updating and Applications to Communication Games*

Navin Kartik[†] Frances Xu Lee[‡] Wing Suen[§]

August 8, 2019

Abstract

We develop a result on Bayesian updating. Roughly, when two agents have different priors, each believes that a (Blackwell) more informative experiment will, on average, bring the other's posterior closer to his own prior. We apply the result to two models of strategic communication: verifiable disclosure and costly falsification. Under some conditions, senders' information revelation are strategic complements when concealing or falsifying information is costly, but strategic substitutes when disclosing verifiable information is costly. In the latter case (but not the others), a receiver can be worse off with additional senders or better external information.

*We thank Nageeb Ali, Ilan Guttman, Keri Hu, Alessandro Lizzeri, Steve Matthews, Ariel Pakes, Andrea Prat, Mike Riordan, Satoru Takahashi, Jianrong Tian, Enrico Zanardo, Jidong Zhou, and various seminar and conference audiences for comments. Sarba Raj Bartaula, Dilip Ravindran, and Teck Yong Tan provided excellent research assistance. Kartik gratefully acknowledges financial support from the NSF (Grant SES-1459877).

[†]Department of Economics, Columbia University. E-mail: nkartik@columbia.edu.

[‡]Quinlan School of Business, Loyola University Chicago. E-mail: francesxu312@gmail.com.

[§]Faculty of Business and Economics, The University of Hong Kong. E-mail: wsuen@econ.hku.hk.

Contents

1. Introduction

Rational agents revise their beliefs upon receiving new information. From an ex-ante point of view, however, one cannot expect new information to systematically alter one's beliefs in any particular direction. More precisely, a fundamental property of Bayesian updating is that beliefs are a martingale: an agent's expectation of his posterior belief is equal to his prior belief. But what about an agent's expectation of *another* agent's posterior belief when their current beliefs are different? Or, intimately related, should agents expect new information to systematically affect any existing disagreement? These questions are not only of intrinsic interest, but tackling them proves useful for analyzing common-prior environments with asymmetric information.

This paper makes two contributions. First, we provide a result concerning the mutual expectations of Bayesian agents who disagree on the distribution of a fundamental. Second, we use this result to derive new insights about some canonical sender(s)-receiver communication games either with multiple senders or when a receiver obtains external information. We identify a unified mechanism that drives senders to reveal more information in some economic settings and less in others; in the latter cases, a receiver can be harmed by having access to better external information or more senders.

A theorem on Bayesian updating. In Section 2, we develop the following result on Bayesian updating. (Throughout this introduction, some technical details are suppressed.) Let $\Omega = \{0, 1\}$ be the possible states of the world. Anne (A) and Bob (B) have mutually-known but possibly-different priors over Ω, $\overline{\beta}_A \in (0, 1)$ and $\overline{\beta}_B \in (0, 1)$ respectively, where $\overline{\beta}_i$ is the probability individual i ascribes to state $\omega = 1$. A signal, s, will be drawn from an information structure or experiment, $\mathcal{E}$, given by a family of probability distributions, $\{p_\omega(\cdot)\}_{\omega\in\Omega}$, where $p_\omega(s)$ is the probability of observing signal s when the true state is ω. Anne and Bob agree on the experiment. Let $\beta_i(s)$ denote i's posterior (on $\omega = 1$) after observing signal s, derived by Bayesian updating. Let $\mathbb{E}_i^{\mathcal{E}}[\beta_j(\cdot)]$ be i's ex-ante expectation of j's posterior under experiment $\mathcal{E}$, where the expectation is taken over signals from i's point of view.

Consider two experiments $\mathcal{E}$ and $\tilde{\mathcal{E}}$ that are comparable in the sense of Blackwell (1951, 1953); specifically, let $\tilde{\mathcal{E}}$ be a garbling of $\mathcal{E}$, or equivalently, $\mathcal{E}$ be more informative than $\tilde{\mathcal{E}}$.[1] We prove (Theorem 1) that for any $i, j \in \{A, B\}$,

$$\overline{\beta}_i \leq (\geq)\, \overline{\beta}_j \implies \mathbb{E}_i^{\mathcal{E}}[\beta_j(\cdot)] \leq (\geq)\, \mathbb{E}_i^{\tilde{\mathcal{E}}}[\beta_j(\cdot)]. \tag{1}$$

[1] Throughout, binary comparisons should be understood in the weak sense unless indicated otherwise, e.g., "greater than" means "at least as large as."

In words: if Anne holds a lower (resp., higher) prior than Bob, then Anne predicts that a more informative experiment will, on average, reduce (resp., raise) Bob's posterior by a larger amount than a less informative experiment. Put differently, Anne expects more information to further validate her prior in the sense of bringing, on average, Bob's posterior closer to her prior. Of course, Bob expects just the reverse. We refer to this result as *information validates the prior*, IVP hereafter.

IVP subsumes the familiar martingale property of Bayesian updating. To see this, consider $j = i$ in Equation 1 and take $\tilde{\mathcal{E}}$ to be an uninformative experiment—which leaves an agent's posterior after any signal equal to his prior—and take $\mathcal{E}$ to be any experiment. IVP also implies that an agent expects another agent's posterior to fall in between their two priors.[2]

IVP implies an interesting point about *disagreement*. When the disagreement between beliefs β_i and β_j is quantified by the canonical metric $|\beta_i - \beta_j|$, a particular signal can lead to higher posterior disagreement than the prior. (By contrast, as shown by Baliga, Hanany, and Klibanoff (2013), no signal can lead to polarization: it cannot be that $(\beta_i - \overline{\beta}_i)(\beta_j - \overline{\beta}_j) < 0$.) Nevertheless, IVP implies (Corollary 1) that both Anne and Bob predict that a more informative experiment will, on average, reduce their posterior disagreement to a greater extent than a less informative experiment; in particular, no matter the experiment, they expect lower posterior disagreement than their prior disagreement. At the extreme, both predict their posterior disagreement to be zero under a fully informative experiment—even though they hold different expectations about what the other's posterior belief will be.

Theorem 2 provides a generalization of IVP to non-binary state spaces using suitable statistics of the agents' beliefs and likelihood-ratio ordering conditions on the agents' priors and on the experiments.

Applications to communication games. IVP is a statistical result. The second part of our paper shows why it is useful in strategic contexts. The idea is that even in common-prior environments, private information can lead to one agent—Anne, an informed agent—holding a different belief about a fundamental than another—Bob, an uninformed agent. Only Anne knows both their beliefs. Anne's strategic incentives may depend on how she expects new information to affect Bob's belief. This new information can be exogenous or endogenous, e.g., owing to the strategic behavior of still other agents. The belief difference between Anne and Bob can occur either *on* an equilibrium path because of Anne's "pooling" behavior, or

[2] To see this, note that (i) letting $\mathcal{U}$ denote an uninformative experiment, $\mathbb{E}_i^{\mathcal{U}}[\beta_j(\cdot)] = \overline{\beta}_j$, and (ii) letting $\mathcal{F}$ denote a fully informative experiment (one in which any signal reveals the state and hence leaves any two agents with the same posterior no matter their priors), $\mathbb{E}_i^{\mathcal{F}}[\beta_j(\cdot)] = \overline{\beta}_i$. The result follows using Equation 1 and the fact that any experiment is more informative than $\mathcal{U}$ but less informative than $\mathcal{F}$.

off path because of a deviation by Anne; as usual, even off-path considerations will generally affect the on-path properties of an equilibrium.

We develop these points in the context of two familiar classes of communication games.

Voluntary disclosure. Section 3 studies a family of voluntary *disclosure games*, also referred to as persuasion games. In these games, biased senders' only instrument of influence on the decision-maker (DM) is the certifiable or verifiable private information they are endowed with; see Milgrom (2008) and Dranove and Jin (2010) for surveys. Senders cannot explicitly lie but can choose what information to disclose and what to withhold. We consider a model in which senders, if informed, draw signals that are independently distributed conditional on an underlying binary state that affects the DM's payoff.[3] We allow for senders to either have opposing or similar biases but assume that a sender's payoff is state-independent and linear (increasing or decreasing) in the DM's belief.

Our focus is on how strategic disclosure interacts with *message costs*: either disclosure or concealment of information can entail direct costs for each sender. Disclosure costs are natural when the process of certifying or publicizing information demands resources such as time, effort, or hiring an outside expert; there can also be other "proprietary costs." A subset of the literature, starting with Jovanovic (1982) and Verrecchia (1983), has modeled such a cost, although primarily only in single-sender problems. On the flip side, there are contexts in which it is the suppression of information that requires costly resources. Besides direct costs, there can also be a psychic disutility to concealing information, or concealment may be discovered ex-post (by auditors, whistleblowers, or mere happenstance) and result in negative consequences for the sender through explicit punishment or reputation loss. A recent example is the $70 million dollar fine imposed by the National Highway Traffic Safety Administration on Honda Motor in January 2015 because "it did not report hundreds of death and injury claims ... for the last 11 years nor did it report certain warranty and other claims in the same period." (The New York Times, 2015)

To determine his own disclosure, it is essential for each sender to predict how other senders' messages will affect the DM's posterior belief, and how this depends on his own message (disclosure or nondisclosure). Our methodological insight is to view senders' messages as endogenously-determined experiments and bring the IVP theorem to bear. This approach allows us to provide a unified treatment—regardless of whether senders have similar

[3] We assume that senders are uninformed with some probability; importantly, an uninformed sender cannot certify that he is uninformed. This modeling device was introduced by Dye (1985) and Shin (1994a,b) to prevent "unraveling" (Grossman and Hart, 1980; Milgrom, 1981), and serves the same purpose here.

or opposing biases, and regardless of whether message costs are disclosure or concealment costs. We thereby uncover new insights into the classic question of when competition promotes disclosure.

Any sender's equilibrium disclosure behavior follows a threshold rule of disclosing all sufficiently favorable signals. We establish a simple but important benchmark: without message costs, any sender in the multi-sender disclosure game uses the same disclosure threshold as he would in a single-sender game. In other words, there is a *strategic irrelevance*. The intuition is that without message costs, a sender's objective is the same regardless of the presence of other senders: he simply wants to induce the most favorable "interim belief" in the DM based on his own message, as this will lead to the most favorable posterior belief based on all senders' messages.

How do message costs alter the irrelevance result? Consider a concealment cost. In a single-sender setting, a sender i's disclosure threshold will be such that the DM's interpretation of nondisclosure is more favorable than i's private belief at the threshold—this wedge is necessary to compensate i for the concealment cost. Nondisclosure thus generates an *interim disagreement* between the DM's belief and the threshold type's private belief. Naturally, disclosure produces no such disagreement. Now add a second sender j to the picture. Our IVP theorem implies that regardless of j's behavior, the threshold type of i predicts that j's message will, on average, make the DM's posterior less favorable to i as compared to the DM's interim belief following i's nondisclosure. Consequently, concealment is now less attractive to sender i. IVP further implies that j's effect on i is stronger when j disclosure more, i.e., when j is more informative. In sum, senders' disclosures are *strategic complements* under a concealment cost.[4]

The logic reverses under a disclosure cost. In a single-sender setting, the DM's interpretation of nondisclosure is now less favorable than i's private belief at the threshold—the gain from disclosing information must compensate i for the direct cost. Reasoning analogously to above, IVP now implies that disclosure becomes less attractive when the other sender is more informative: the threshold type of i expects the other sender's message to make the DM's posterior more favorable to i, reducing the gains from disclosure. Consequently, senders' disclosures are *strategic substitutes* under a disclosure cost.

These results have notable welfare implications. In the case of concealment cost, a DM always benefits from an additional sender not only because of the information this sender

[4] In a different model, Bourjade and Jullien (2011) find an effect related to that we find under concealment cost. Loosely speaking, "reputation loss" in their model plays a similar role to concealment cost in ours.

provides, but also because it improves disclosure from other senders. In the case of disclosure cost, however, the strategic substitution result implies that while a DM gains some direct benefit from consulting an additional sender, the indirect effect on other senders' behavior is deleterious to the DM. In general, the net effect is ambiguous; it is not hard to construct examples in which the DM is made strictly worse off by adding a sender, even if this sender has an opposite bias to that of an existing sender. Thus, competition between senders need not increase information revelation nor benefit the DM.[5] The DM can even be made worse off when disclosure costs become lower or a sender is more likely to be informed, although either change would help the DM in a single-sender setting.

Costly signaling games. Our IVP theorem is also useful in games of asymmetric information even when the information asymmetry is eliminated in equilibrium. We develop such an application in Section 4. For concreteness, we consider information transmission with lying costs (Kartik, 2009), but we also discuss how our themes apply to canonical signaling applications like education signaling (Spence, 1973).

Suppose a sender has imperfect information about a binary state, and can falsify or manipulate his information by incurring costs. A receiver makes inferences about the state based on both the sender's signal and some other exogenous information.[6] The sender has linear preferences over the receiver's belief about the state. Due to a bounded signal space, there is incomplete separation across sender types; we focus on equilibria in which "low" types separate and "high" types pool. What is the effect of better exogenous information (in the Blackwell sense) on the sender's signaling?

Even sender types who separate in equilibrium can induce an incorrect interim belief (i.e., the receiver's belief based only on the sender's signal, before incorporating the exogenous information) by deviating. We use IVP to establish that better exogenous information reduces the sender's expected benefit from falsifying his information to appear more favorable. Intuitively, the sender expects any favorable but incorrect interim belief to get neutralized more

[5] In quite different settings, Milgrom and Roberts (1986), Dewatripont and Tirole (1999), Krishna and Morgan (2001), and Gentzkow and Kamenica (2017) offer formal analyses supporting the viewpoint that competition between senders helps—or at least cannot hurt—a DM. Carlin, Davies, and Iannaccone (2012) present a result in which increased competition leads to less voluntary disclosure. Their model can be viewed as one in which senders bear a concealment cost that is assumed to decrease in the amount of disclosure by other senders. Elliott, Golub, and Kirilenko (2014) show how a DM can be harmed by "information improvements" in a cheap-talk setting, but the essence of their mechanism is not the strategic interaction between senders.

[6] Weiss (1983) is an early paper that introduces exogenous information into a signaling environment; his is not, however, a model in which signaling has direct costs. See also Feltovich, Harbaugh, and To (2002). More similar to our setting are Daley and Green (2014) and Truyts (2015); we elaborate on the connections with these papers at the end of Section 4.

by better exogenous information. This means that better exogenous information makes costly deviations less attractive to the sender, which relaxes his incentive constraints. Consequently, better exogenous information reduces wasteful signaling and enlarges the set of separating types—every type of the sender is better off and so is the receiver because there is more information revealed by the sender. We explain how this result can be also viewed as a strategic complementary between multiple senders.

Other applications. While our paper applies IVP to two models of strategic communication, we hope that the result will also be useful in other contexts. Indeed, the logic of IVP underlies and unifies the mechanisms in a few existing papers that study models with heterogeneous priors under specific information structures. In particular, see the strategic "persuasion motive" that generates bargaining delays in Yildiz (2004), motivational effects of difference of opinion in Che and Kartik (2009) and Van den Steen (2010, Proposition 5), and a rationale for deference in Hirsch (2016, Proposition 8); in a non-strategic setting, see why minorities expect lower levels of intermediate bias in Sethi and Yildiz (2012, Proposition 5). We ourselves have used IVP—specifically invoking Theorem 1 below—to study information acquisition prior to disclosure (Kartik, Lee, and Suen, 2017). We briefly touch on some implications of IVP for Bayesian persuasion with heterogeneous priors in the current paper's conclusion, Section 5.

2. Information Validates the Prior

The backbone of our analysis is a pair of theorems below that relate the informativeness of an experiment to the expectations of individuals with different beliefs. Theorem 1 is a special case of Theorem 2, but for expositional clarity, we introduce them separately.

Throughout, we use the following standard definitions concerning information structures (Blackwell, 1953). Fix any finite state space Ω with generic element ω. An *experiment* is $\mathcal{E} \equiv (S, \mathcal{S}, \{P_\omega\}_{\omega\in\Omega})$, where S is a measurable space of signals, $\mathcal{S}$ is a σ-algebra on S, and each P_ω is a probability measure over the signals in state ω. An experiment $\tilde{\mathcal{E}} \equiv (\tilde{S}, \tilde{\mathcal{S}}, \{\tilde{P}_\omega\}_{\omega\in\Omega})$ is a *garbling* of experiment $\mathcal{E}$ if there is a Markov kernel from $(S, \mathcal{S})$ to $(\tilde{S}, \tilde{\mathcal{S}})$, denoted $Q(\cdot|s)$,[7] such that for each $\omega \in \Omega$ and every set $\Sigma \in \tilde{\mathcal{S}}$, it holds that

$$\tilde{P}_\omega(\Sigma) = \int_S Q(\Sigma|s)\, \mathrm{d}P_\omega(s).$$

This definition captures the statistical notion that, on a state-by-state basis, the distribution

[7] I.e., (i) the map $s \mapsto Q(\Sigma|s)$ is $\mathcal{S}$-measurable for every $\Sigma \in \tilde{\mathcal{S}}$, and (ii) the map $\Sigma \mapsto Q(\Sigma|s)$ is a probability measure on $(\tilde{S}, \tilde{\mathcal{S}})$ for every $s \in S$.

of signals in $\tilde{\mathcal{E}}$ can be generated by taking signals from $\mathcal{E}$ and transforming them through the state-independent kernel $Q(\cdot)$. In a sense, $\tilde{\mathcal{E}}$ does not provide any information that is not contained in $\mathcal{E}$. Indeed, $\mathcal{E}$ is also said to be *more informative* than $\tilde{\mathcal{E}}$ because every expected-utility decision maker prefers $\mathcal{E}$ to $\tilde{\mathcal{E}}$.

2.1. Binary states

In this subsection, let $\Omega = \{0, 1\}$, and take all beliefs to refer to the probability of state $\omega = 1$. There are two individuals, A and B, with respective prior beliefs $\overline{\beta}_A, \overline{\beta}_B \in (0, 1)$. Given any experiment, the individuals' respective priors combine with Bayes rule to determine their respective posteriors after observing a signal s, denoted $\beta_i(s)$ for $i \in \{A, B\}$. Let $\mathbb{E}^{\mathcal{E}}_A[\beta_B(\cdot)]$ denote the ex-ante expectation of individual A over the posterior of individual B under experiment $\mathcal{E}$. If $\overline{\beta}_A = \overline{\beta}_B$, then because the individuals' posteriors always agree, it holds that for any $\mathcal{E}$, $\mathbb{E}^{\mathcal{E}}_A[\beta_B(\cdot)] = \overline{\beta}_A$; this is the martingale or iterated expectations property of Bayesian updating. For individuals with different priors, we have:

Theorem 1. *Let $\Omega = \{0, 1\}$ and experiment $\tilde{\mathcal{E}}$ be a garbling of experiment $\mathcal{E}$. Then, for any $i, j \in \{A, B\}$:*

$$\begin{aligned} \overline{\beta}_i \leq \overline{\beta}_j &\implies \mathbb{E}^{\mathcal{E}}_i[\beta_j(\cdot)] \leq \mathbb{E}^{\tilde{\mathcal{E}}}_i[\beta_j(\cdot)]; \\ \overline{\beta}_i \geq \overline{\beta}_j &\implies \mathbb{E}^{\mathcal{E}}_i[\beta_j(\cdot)] \geq \mathbb{E}^{\tilde{\mathcal{E}}}_i[\beta_j(\cdot)]; \end{aligned}$$

and $\min\{\overline{\beta}_A, \overline{\beta}_B\} \leq \mathbb{E}^{\mathcal{E}}_i[\beta_j(\cdot)] \leq \max\{\overline{\beta}_A, \overline{\beta}_B\}$.

Suppose that individual A is less optimistic than B, i.e., $\overline{\beta}_A < \overline{\beta}_B$. If an experiment $\tilde{\mathcal{E}}$ is uninformative—no signal realization would change any individual's beliefs—then $\mathbb{E}^{\tilde{\mathcal{E}}}_A[\beta_B(\cdot)] = \overline{\beta}_B > \overline{\beta}_A$. On the other hand, if an experiment $\mathcal{E}$ is fully informative—every signal reveals the state—then for any signal s, $\beta_A(s) = \beta_B(s)$, and hence $\overline{\beta}_A = \mathbb{E}^{\mathcal{E}}_A[\beta_A(\cdot)] = \mathbb{E}^{\mathcal{E}}_A[\beta_B(\cdot)] < \overline{\beta}_B$, where the first equality is the previously-noted property of Bayesian updating under any experiment. In other words, individual A believes a fully informative experiment will, on average, bring individual B's posterior perfectly in line with A's own prior, whereas an uninformative experiment will obviously entail no such convergence. (Of course, in turn, B expects A to update to B's prior under a fully informative experiment.) Theorem 1 generalizes this idea to monotonicity among Blackwell-comparable experiments: A anticipates that a more informative experiment will, on average, bring B's posterior closer to A's prior. For short, we will say the theorem shows that *information validates the prior*, or *IVP*.

There is a simple proof for Theorem 1, owing to binary states. The key is to recognize that, because the individuals agree on the experiment and only disagree in their priors over the state, each individual's posterior can be written as a function of both their priors and the other individual's posterior; this point does not rely on binary states. This observation is also used by Gentzkow and Kamenica (2014), Alonso and Câmara (2016), and Galperti (2018) to different ends. For simplicity, consider an experiment with a discrete signal space in which every signal is obtained with positive probability in both states. For any signal realization s, Bayes rule implies that the posterior belief $\beta_i(s)$ for individual $i \in \{A, B\}$ satisfies

$$\frac{\beta_i(s)}{1-\beta_i(s)} = \frac{\overline{\beta}_i}{1-\overline{\beta}_i}\frac{P_1(s)}{P_0(s)},$$

where $P_\omega(s)$ is the probability of observing s in state ω. Eliminating the likelihood ratio $P_1(s)/P_0(s)$ yields $\beta_B(s) = T(\beta_A(s), \overline{\beta}_B, \overline{\beta}_A)$, where

$$T(\beta_A, \overline{\beta}_B, \overline{\beta}_A) := \frac{\beta_A \frac{\overline{\beta}_B}{\overline{\beta}_A}}{\beta_A \frac{\overline{\beta}_B}{\overline{\beta}_A} + (1-\beta_A)\frac{1-\overline{\beta}_B}{1-\overline{\beta}_A}}. \tag{2}$$

It is straightforward to verify that this transformation mapping $T(\cdot, \overline{\beta}_B, \overline{\beta}_A)$ is strictly concave (resp., convex) in A's posterior when $\overline{\beta}_A < \overline{\beta}_B$ (resp., $\overline{\beta}_A > \overline{\beta}_B$). Theorem 1 follows as an application of Blackwell (1953), who showed that a garbling increases (resp., reduces) an individual's expectation of any concave (resp., convex) function of his posterior; see also Rothschild and Stiglitz (1970).

The second part of Theorem 1 is a straightforward consequence of combining the first part with the facts that any experiment is a garbling of a fully informative experiment while an uninformative experiment is a garbling of any experiment (and using the properties of these extreme experiments noted right after the theorem).

Remark 1. The inequalities in the conclusions of Theorem 1 hold strictly if the garbling is "strict." In particular, $\min\{\overline{\beta}_A, \overline{\beta}_B\} < \mathbb{E}_i^{\mathcal{E}}[\beta_j(\cdot)] < \max\{\overline{\beta}_A, \overline{\beta}_B\}$ for any experiment $\mathcal{E}$ that is neither uninformative nor fully informative. □

Theorem 1 has an interesting corollary. A canonical measure of how much agents A and B disagree when they hold beliefs β_A and β_B is $|\beta_A - \beta_B|$. Under this measure, an experiment can generate signals for which posterior disagreement is larger than prior disagreement.[8] Put differently, Bayesian agents who agree on the information-generating process can disagree

[8] For example, let $\overline{\beta}_A = 1/2$, and consider a binary-signal experiment ($S = \{l, h\}$) with $\Pr(l|0) = \Pr(h|1) = p$. For any $p \in (1/2, 1)$, the posterior disagreement after signal h is larger than the prior disagreement if $\overline{\beta}_B < 1-p$. More generally, it can be shown that whenever $\overline{\beta}_A \neq \overline{\beta}_B$, there is an experiment such that posterior disagreement after some signal is larger than prior disagreement.

more after new information.[9] But can they expect, ex-ante, to disagree more? Theorem 1 provides a sharp answer, captured by the following corollary.

Corollary 1. *Let $\Omega = \{0,1\}$ and experiment $\tilde{\mathcal{E}}$ be a garbling of experiment $\mathcal{E}$. Then, for any $i \in \{A, B\}$:*

$$\mathbb{E}_i^{\mathcal{E}}\left[|\beta_A(\cdot) - \beta_B(\cdot)|\right] \leq \mathbb{E}_i^{\tilde{\mathcal{E}}}\left[|\beta_A(\cdot) - \beta_B(\cdot)|\right]. \tag{3}$$

Corollary 1 says that both agents expect a more informative experiment to reduce their disagreement by more. In particular, both predict that any experiment will reduce their disagreement, on average, relative to their prior disagreement. This is the case even though they generally hold different expectations about each other's posteriors; for example, given a fully informative experiment, both anticipate zero posterior disagreement, even though either agent i expects both agents' posteriors to be $\overline{\beta}_i$.

The proof of Corollary 1 is simple. For any priors $\overline{\beta}_A$ and $\overline{\beta}_B$, and any signal s in any experiment, it is a consequence of Bayesian updating (e.g., using (2)) that

$$\text{sign}[\beta_A(s) - \beta_B(s)] = \text{sign}[\overline{\beta}_A - \overline{\beta}_B].$$

In other words, any information preserves the prior ordering of the agents' beliefs. Thus, if (say) $\beta_A \geq \beta_B$, then (3) is equivalent to

$$\mathbb{E}_i^{\mathcal{E}}[\beta_A(\cdot) - \beta_B(\cdot)] \leq \mathbb{E}_i^{\tilde{\mathcal{E}}}[\beta_A - \beta_B],$$

which is implied by Theorem 1 because $\mathbb{E}_i^{\mathcal{E}}[\beta_i(\cdot)] = \mathbb{E}_i^{\tilde{\mathcal{E}}}[\beta_i(\cdot)] = \overline{\beta}_i$.

2.2. Many states

Now consider an arbitrary finite set of states, $\Omega \equiv \{\omega_1, \ldots, \omega_L\} \subset \mathbb{R}$, with $L > 1$ and $\omega_1 < \cdots < \omega_L$. We write $\beta(\omega_l)$ as the probability ascribed to state ω_l by a belief $\boldsymbol{\beta}$, with the notation in bold emphasizing that a belief over Ω is now a vector. We assume that the two individuals' priors, $\overline{\boldsymbol{\beta}}_A$ and $\overline{\boldsymbol{\beta}}_B$, assign strictly positive probability to every state. We say that a belief $\boldsymbol{\beta'}$ *likelihood-ratio dominates* belief $\boldsymbol{\beta}$, written $\boldsymbol{\beta'} \geq_{LR} \boldsymbol{\beta}$ if, for all $\omega' > \omega$,

$$\beta'(\omega')\beta(\omega) \geq \beta(\omega')\beta'(\omega).$$

We denote posterior beliefs given a signal s as $\boldsymbol{\beta}(s) \equiv (\beta(\omega_1|s), \ldots, \beta(\omega_L|s))$.

[9] But agents' posteriors always move in the same direction from their priors, i.e., after any signal s in any experiment, $\text{sign}[\beta_A(s) - \overline{\beta}_A] = \text{sign}[\beta_B(s) - \overline{\beta}_B]$, as can be confirmed using (2). In a sense there cannot be polarization, a point that is established more generally by Baliga et al. (2013, Theorem 1).

Definition 1. An experiment $\mathcal{E} \equiv (S, \mathcal{S}, \{P_\omega\}_{\omega\in\Omega})$ is an *MLRP-experiment* if there is a total order on S, denoted $\succeq$ (with asymmetric relation $\succ$), such that the monotone likelihood ratio property holds: $s' \succ s$ and $\omega' > \omega \implies p(s'|\omega')p(s|\omega) \geq p(s'|\omega)p(s|\omega')$.

As is well known, the monotone likelihood ratio property (in the non-strict version above) is without loss of generality when there are only two states: any experiment is an MLRP-experiment when $L = 2$.

Let $M(\boldsymbol{\beta}) := \sum_{\omega\in\Omega} h(\omega)\beta(\omega)$ represent the expectation of an arbitrary non-decreasing function $h : \Omega \to \mathbb{R}$ under belief $\boldsymbol{\beta}$. The following result generalizes Theorem 1:

Theorem 2. *Consider any two likelihood-ratio ordered priors $\overline{\boldsymbol{\beta}}_A$ and $\overline{\boldsymbol{\beta}}_B$, and any two MLRP-experiments $\mathcal{E}$ and $\tilde{\mathcal{E}}$ with $\tilde{\mathcal{E}}$ a garbling of $\mathcal{E}$. Then, for any $i, j \in \{A, B\}$:*

$$M(\overline{\boldsymbol{\beta}}_i) \leq M(\overline{\boldsymbol{\beta}}_j) \implies \mathbb{E}_i^{\mathcal{E}}\left[M(\boldsymbol{\beta}_j(\cdot))\right] \leq \mathbb{E}_i^{\tilde{\mathcal{E}}}\left[M(\boldsymbol{\beta}_j(\cdot))\right];$$
$$M(\overline{\boldsymbol{\beta}}_i) \geq M(\overline{\boldsymbol{\beta}}_j) \implies \mathbb{E}_i^{\mathcal{E}}\left[M(\boldsymbol{\beta}_j(\cdot))\right] \geq \mathbb{E}_i^{\tilde{\mathcal{E}}}\left[M(\boldsymbol{\beta}_j(\cdot))\right];$$

and $\min\left\{M(\overline{\boldsymbol{\beta}}_A), M(\overline{\boldsymbol{\beta}}_B)\right\} \leq \mathbb{E}_i^{\mathcal{E}}\left[M(\boldsymbol{\beta}_j(\cdot))\right] \leq \max\left\{M(\overline{\boldsymbol{\beta}}_A), M(\overline{\boldsymbol{\beta}}_B)\right\}$.

With many states, we can still represent the posterior belief $\boldsymbol{\beta}_i(s)$ of individual i as a transformation mapping of $\boldsymbol{\beta}_j(s)$, but this mapping is neither concave nor convex even when the prior beliefs of i and j are likelihood-ratio ordered. The proof of Theorem 2 (and of all subsequent formal results) is provided in Appendix A. For an illustration, suppose $M(\cdot)$ is the expected state and Bayesian updating takes the canonical linear form of a convex combination of the prior mean and the signal, as is the case for any exponential family of signals with conjugate prior (e.g., normal-normal). Given a signal $s \in \mathbb{R}$ under experiment $\mathcal{E}$, $\mathbb{E}_j^{\mathcal{E}}[\omega|s] = (1-\alpha^{\mathcal{E}})\mathbb{E}_j[\omega] + \alpha^{\mathcal{E}} s$ for some $\alpha^{\mathcal{E}} \in [0, 1]$. Hence,

$$\mathbb{E}_i^{\mathcal{E}}\left[\mathbb{E}_j^{\mathcal{E}}[\omega|s]\right] = (1-\alpha^{\mathcal{E}})\mathbb{E}_j[\omega] + \alpha^{\mathcal{E}}\mathbb{E}_i[\omega].$$

The weight $\alpha^{\mathcal{E}}$ is larger the more informative is the experiment $\mathcal{E}$, which implies the conclusions of Theorem 2.

Theorem 2 says that either agent i expects the other's posterior to be closer to i's prior under a more informative experiment, when closeness of beliefs is measured by the distance between their $M(\cdot)$ statistics. Theorem 1 obtains when $\Omega = \{0, 1\}$ and $M(\boldsymbol{\beta})$ is simply the expected state under $\boldsymbol{\beta}$ (i.e., $h(\omega) = \omega$), because both likelihood-ratio ordering assumptions in Theorem 2 are without loss of generality with only two states.[10] More generally, since $M(\cdot)$

[10] Indeed, the likelihood-ratio ordering of the priors can always be viewed as without loss of generality—the

is the expectation of any non-decreasing function of the state, Theorem 2 captures the notion that—subject to the ordering assumptions—information validates the prior not only in the sense of the expected state, but also, for instance, in the probability of any $\{\omega_l, \ldots, \omega_L\}$.

The likelihood-ratio ordering assumptions in Theorem 2 are tight in the following respect: (i) there exist priors $\overline{\boldsymbol{\beta}}_B >_{LR} \overline{\boldsymbol{\beta}}_A$ and a non-MLRP-experiment $\mathcal{E}$ such that $\mathbb{E}_A^{\mathcal{E}}[M(\boldsymbol{\beta}_B(\cdot))] > M(\overline{\boldsymbol{\beta}}_B) > M(\overline{\boldsymbol{\beta}}_A)$; and (ii) there exist priors $\overline{\boldsymbol{\beta}}_A \not\geq_{LR} \overline{\boldsymbol{\beta}}_B$ and an MLRP-experiment $\mathcal{E}$ such that $\mathbb{E}_A^{\mathcal{E}}[M(\boldsymbol{\beta}_{\boldsymbol{B}}(\cdot))] > M(\overline{\boldsymbol{\beta}}_A) > M(\overline{\boldsymbol{\beta}}_B)$. See Appendix B for examples demonstrating these points; the examples also establish that the likelihood-ratio ordering hypothesis on the priors cannot be weakened to first-order stochastic dominance.

Theorem 2 has an implication about expected disagreement that generalizes Corollary 1:

Corollary 2. *Consider any two likelihood-ratio ordered priors $\overline{\boldsymbol{\beta}}_{\boldsymbol{A}}$ and $\overline{\boldsymbol{\beta}}_{\boldsymbol{B}}$, and any two MLRP-experiments $\mathcal{E}$ and $\tilde{\mathcal{E}}$ with $\tilde{\mathcal{E}}$ a garbling of $\mathcal{E}$. Then, for any $i \in \{A, B\}$:*

$$\mathbb{E}_i^{\mathcal{E}}\left[|M(\boldsymbol{\beta}_{\boldsymbol{i}}(\cdot)) - M(\boldsymbol{\beta}_{\boldsymbol{j}}(\cdot))|\right] \leq \mathbb{E}_i^{\tilde{\mathcal{E}}}\left[|M(\boldsymbol{\beta}_{\boldsymbol{i}}(\cdot)) - M(\boldsymbol{\beta}_{\boldsymbol{j}}(\cdot))|\right].$$

Corollary 2 says that—subject to the ordering assumptions—more information reduces expected disagreement when the disagreement between beliefs $\boldsymbol{\beta}_A$ and $\boldsymbol{\beta}_B$ is measured by $|M(\boldsymbol{\beta}_A) - M(\boldsymbol{\beta}_B)|$.[11] The corollary follows from Theorem 2 for reasons analogous to our discussion of why Corollary 1 follows from Theorem 1. In particular, information from an MLRP-experiment preserves the prior likelihood-ratio ordering of the agents' beliefs and thus also the ordering in terms of their $M(\cdot)$ statistics; moreover, since $M(\cdot)$ is a linear function, it holds for any $i \in \{A, B\}$ and any experiment $\mathcal{E}$ that $\mathbb{E}_i^{\mathcal{E}}[M(\boldsymbol{\beta}_{\boldsymbol{i}}(\cdot)] = M(\overline{\boldsymbol{\beta}}_{\boldsymbol{i}})$.

In the following sections, we use the logic of information validates the prior to study strategic communication games in which agents begin with a common prior.

3. Multi-Sender Disclosure Games

Our main application is to multi-sender voluntary disclosure of verifiable information. We build on the classic disclosure models mentioned in the Introduction and use IVP to derive new insights on when competition promotes disclosure.

states can be re-labeled so that it holds—so long as (i) the MLRP property of an experiment and (ii) the non-decreasing property of the $h(\cdot)$ function are both understood to be with respect to the ordering of states that ensures likelihood-ratio ordering of the priors.

[11] The distance between expected states—more precisely, the expectations of any non-decreasing function of the states—is an interesting measure of disagreement, but there are of course others. See Zanardo (2017) for an axiomatic approach to measuring disagreement; he provides a result in the spirit of Corollary 2 for a family of disagreement measures that includes the Kullback-Leibler divergence.

3.1. The model

Players. There is an unknown state of the world, $\omega \in \{0, 1\}$. A decision maker (DM) will form a belief β_{DM} that the state is $\omega = 1$. For much of our analysis, all that matters is the belief that the DM holds. For welfare evaluation, however, it is useful to view the DM as taking an action a with von-Neumann Morgenstern utility function $u_{DM}(a, \omega)$ such that the DM is strictly better off when there is strictly more information about the state in the Blackwell sense. There are two senders, indexed by i. (Subsection 3.4 generalizes to many senders.) In a reduced form, each sender i has state-independent preferences over the DM's belief given by the von-Neumann Morgenstern utility function $u(\beta_{DM}, b_i) = b_i\beta_{DM}$, where $b_i \in \{-1, 1\}$ captures a sender's bias. That is, each sender has linear preferences over the DM's expectation of the state; $b_i = 1$ means that sender i is biased upward (i.e., prefers a higher DM's expectation), and conversely for $b_i = -1$. Senders' biases are common knowledge. We say that two senders have similar biases if their biases have the same sign, and opposing biases otherwise.

Information. The DM relies on the senders for information about the state. All players share a common prior π over the state. Each sender may exogenously obtain some private information about the state. Specifically, with independent probability $p_i \in (0, 1)$, a sender i is informed and receives a signal $s_i \in S$; with probability $1 - p_i$, he is uninformed, in which case we denote $s_i = \phi$. If informed, sender i's signal is drawn independently from a distribution that depends upon the true state. Without loss, we equate an informed sender's signal with his *private belief*, i.e., a sender's posterior on state $\omega = 1$ given only his own signal $s \neq \phi$ (as derived by Bayesian updating) is s. For convenience, we assume the cumulative distribution of an informed sender's signals in each state, $F(s|\omega)$ for $\omega \in \{0, 1\}$, have common support $S = [\underline{s}, \overline{s}] \subseteq [0, 1]$ and admit respective densities $f(s|\omega)$. It is straightforward to allow $F(\cdot)$ to be different for different senders, but we abstract from such heterogeneity to reduce notation.

Communication. Signals are "hard evidence"; a sender with signal $s_i \in S \cup \{\phi\}$ can send a message $m_i \in \{s_i, \phi\}$. In other words, an uninformed sender only has one message available, ϕ, while an informed sender can either report his true signal or feign ignorance by sending the message ϕ.[12] We refer to any message $m_i \neq \phi$ as *disclosure* and the message $m_i = \phi$ as *nondisclosure*. When an informed sender chooses nondisclosure, we say he is *concealing* information. That senders must either tell the truth or conceal their information is standard;

[12] Due to the senders' monotonic preferences, standard "skeptical posture" arguments imply that our results would be unaffected if we were to allow for a richer message space, for example if an informed sender could report any subset of the signal space that contains his true signal. Likewise, allowing for cheap talk would not affect our results as cheap talk cannot be influential in equilibrium.

a justification is that signals are verifiable and large penalties will be imposed on a sender if a reported signal is discovered to be untrue. Note that being uninformed is not verifiable.

Message costs. A sender i who sends message $m_i \neq \phi$ bears a known utility cost $c \in \mathbb{R}$. We refer to the case of $c > 0$ as one of *disclosure cost* and $c < 0$ as one of *concealment cost*. A disclosure cost captures the idea that costly resources may be needed to certify or make verifiable the information that one has. A concealment cost captures a resource-related or psychic disutility from concealing information, or the expectation of a penalty from ex-post detection of having withheld information (Daughety and Reinganum (2018) and Dye (2017) consider specific versions). As is well known, a disclosure cost precludes full disclosure (Jovanovic, 1982; Verrecchia, 1983). For this reason, our conclusions under $c > 0$ do not require the assumption that the sender may be uninformed ($p_i < 1$). We make that assumption to provide a unified treatment of both disclosure costs ($c > 0$) and no costs/concealment costs ($c \leq 0$).[13] In the latter cases, there would be full disclosure or "unraveling" were $p_i = 1$.

Timing. The game is the following: nature initially determines the state ω and then independently (conditional on the realized state) draws each sender i's private information, $s_i \in S \cup \{\phi\}$; all senders then simultaneously send their respective messages m_i to the DM (whether messages are public or privately observed by the DM is irrelevant); the DM then forms her belief, β_{DM}, according to Bayes rule, whereafter each sender i's payoff is realized as

$$b_i \beta_{DM} - c \cdot \mathbb{1}_{\{m_i \neq \phi\}}. \tag{4}$$

All aspects of the game except the state and senders' signals (or lack thereof) are common knowledge. Our solution concept is the natural adaptation of perfect Bayesian equilibrium, which we will refer to simply as "equilibrium."[14] The notion of welfare for any player is ex-ante expected utility.

[13] When $c = 0$ our setting is related to Jackson and Tan (2013) and Bhattacharya and Mukherjee (2013). Jackson and Tan (2013) assume a binary decision space, which makes their senders' payoffs non-linear in the DM's posterior and shifts the thrust of their analysis. They are ultimately interested in comparing different voting rules, which is effectively like changing the pivotal DM's preferences; we instead highlight how the presence of disclosure or concealment costs affects the strategic interaction between senders holding fixed the DM's preferences. Bhattacharya and Mukherjee (2013) assume that informed senders' signals are perfectly correlated but allow for senders to have non-monotonic preferences over the DM's posterior.

[14] That is: (1) the receiver forms her belief using Bayes rule on path (note that the message $m_i = \phi$ is necessarily on path), and treating any off-path message $m_i \in S$ as proving signal $s_i = m_i$; and (2) each sender chooses his message optimally given his information, the other senders' strategies, and the receiver's belief updating.

3.2. A single-sender benchmark

As a preliminary step and to set a benchmark, begin by considering a (hypothetical) game between a single sender i and the DM. Our analysis in this subsection generalizes some existing results in the literature. For concreteness, suppose the sender is upward biased; straightforward analogs of the discussion below apply if the sender is downward biased.

For any $\beta \in [0,1]$, define $f_\beta(s) := \beta f(s|1) + (1-\beta) f(s|0)$ as the unconditional density of signal s given a belief that puts probability β on state $\omega = 1$. Let F_β be the cumulative distribution of f_β. Since the sender has private belief s upon receiving signal s, disclosure of signal s will lead to the DM also holding belief s. It follows that given any *nondisclosure belief*, i.e., the DM's posterior belief when there is nondisclosure, the optimal strategy for the sender (if informed) is a threshold strategy of disclosing all signals above some *disclosure threshold*, say $\hat{s}$, and concealing all signals below it. If the sender is uninformed, his only available message is ϕ. Suppose the sender uses a disclosure threshold $\hat{s}$. Define the function $\eta : [0,1] \times (0,1) \times [0,1] \to [0,1]$ by

$$\eta(\hat{s}, p, \pi) := \frac{1-p}{1-p+pF_\pi(\hat{s})}\pi + \frac{pF_\pi(\hat{s})}{1-p+pF_\pi(\hat{s})}\mathbb{E}_\pi[s|s<\hat{s}], \tag{5}$$

where $\mathbb{E}_\pi[\cdot]$ refers to an expectation taken with respect to the distribution F_π. The function $\eta(\hat{s}, p, \pi)$ is the posterior implied by Bayes rule in the event of nondisclosure given a conjectured disclosure threshold $\hat{s}$, a probability p of the sender being informed, and prior π.

An increase in the sender's disclosure threshold has two effects on the DM's nondisclosure belief: first, it increases the likelihood that nondisclosure is due to the sender concealing his signal rather than being uninformed; second, conditional on the sender in fact concealing his signal, it causes the DM to expect a higher signal. As the DM's belief conditional on concealment is lower than the prior (since the sender is using a threshold strategy), these two effects work in opposite directions. Moreover, the second effect is stronger than the first if and only if $\hat{s} > \eta(\hat{s}, p, \pi)$. On the other hand, holding the disclosure threshold fixed, an increase in the probability that the sender is informed has an unambiguous effect because it increases the probability that nondisclosure is due to concealed information rather than no information.

Lemma 1. *The nondisclosure belief function, $\eta(\hat{s}, p, \pi)$, has the following properties:*

1. *It is strictly decreasing in $\hat{s}$ when $\hat{s} < \eta(\hat{s}, p, \pi)$ and strictly increasing when $\hat{s} > \eta(\hat{s}, p, \pi)$. Consequently, there is a unique solution to $\hat{s} = \eta(\hat{s}, p, \pi)$, which is the interior point $\arg\min_{\hat{s}}(\hat{s}, p, \pi)$.*
2. *It is weakly decreasing in p, strictly if $\hat{s} \in (\underline{s}, \overline{s})$.*

All but the first sentence of part 1 has appeared previously in Acharya, DeMarzo, and Kremer (2011, Proposition 1 and subsequent discussion); see also Demarzo, Kremer, and Skrzypacz (2018, Proposition 1). The strict quasi-convexity of $\eta(\cdot, p, \pi)$ in part 1 of the lemma will prove crucial for comparative statics.

It follows from the foregoing discussion that any equilibrium is fully characterized by the sender's disclosure threshold. If this threshold is interior, the sender must be indifferent between disclosing the threshold signal and concealing it. As sender i's payoff from disclosing any signal s_i is $s_i - c$, we obtain the following equilibrium characterization.

Proposition 1. *Assume there is only one sender i, and this sender is biased upward.*

1. *Any equilibrium has a disclosure threshold $\hat{s}_i^0$ such that: (i) $\hat{s}_i^0$ is interior and $\eta(\hat{s}_i^0, p_i, \pi) = \hat{s}_i^0 - c$; or (ii) $\hat{s}_i^0 = \underline{s}$ and $\pi \leq \underline{s} - c$; or (iii) $\hat{s}_i^0 = \overline{s}$ and $\pi \geq \overline{s} - c$. Conversely, for any $\hat{s}_i^0$ satisfying (i), (ii), or (iii), there is an equilibrium with disclosure threshold $\hat{s}_i^0$.*
2. *There is a unique equilibrium if there is no message cost or a concealment cost ($c \leq 0$). Moreover, the equilibrium disclosure threshold is interior if there is no message cost ($c = 0$).*
3. *There can be multiple equilibria if there is a disclosure cost ($c > 0$).*

Part 1 of Proposition 1 is straightforward; parts 2 and 3 build on Lemma 1. Multiple equilibria can arise under a disclosure cost because, in the relevant domain (to the right of the fixed point of $\eta(\cdot, p_i, \pi)$), the DM's nondisclosure belief is increasing in the sender's disclosure threshold. In such cases, we will focus on properties of the highest and lowest equilibria in terms of the disclosure threshold. As an equilibrium with a higher disclosure threshold is less (Blackwell) informative, these extremal equilibria are respectively the worst and best equilibria in terms of the DM's welfare. On the other hand, the ranking of these equilibria is reversed for the sender's welfare. To see this, note that because the sender's preferences are linear in the DM's belief and we evaluate welfare at the ex-ante stage, the sender's welfare in an equilibrium with threshold $\hat{s}_i^0$ is $\pi - p_i(1 - F_\pi(\hat{s}_i^0))c$. Thus, when $c > 0$, the sender's welfare is higher when the disclosure threshold is higher: he cannot affect the DM's belief in expectation and prefers a lower probability of incurring the disclosure cost.

When $c = 0$, the sender's belief if he receives the threshold signal is identical to the DM's equilibrium nondisclosure belief. When $c \neq 0$ these two beliefs will differ in any equilibrium: if $c > 0$, the sender's threshold belief, $\hat{s}_i^0$, is higher than the DM's nondisclosure belief, $\eta(\hat{s}_i^0, p_i, \pi)$, and the opposite is true when $c < 0$. This divergence of equilibrium beliefs when the sender withholds information—which, for brevity, we shall refer to as *disagreement*—will prove crucial. Note that when there is disagreement, the sender knows the DM's belief but not vice-versa; consistent with Aumann (1976), the beliefs under disgareement are not common knowledge in this game with a common prior.

Proposition 1 is stated for an upward biased sender. If the sender is instead downward biased, he reveals all signals below some threshold. The DM's nondisclosure belief function changes from (5) to

$$\frac{1-p}{1-p+p(1-F_\pi(\hat{s}))}\pi + \frac{p(1-F_\pi(\hat{s}))}{1-p+p(1-F_\pi(\hat{s}))}\mathbb{E}_\pi[s|s>\hat{s}].$$

This expression is single-peaked in $\hat{s}$. The condition for an interior equilibrium is that it must equal $\hat{s}+c$, since the sender's payoff is now $-\beta_{DM}-c$. As with an upward biased sender, equilibrium is unique when $c \leq 0$ while there can be multiple equilibria when $c > 0$. Note that for any $c \neq 0$, the direction of disagreement between the sender and the DM is now reversed: for a downward biased sender, $c > 0$ implies the sender's threshold belief is lower than the DM's nondisclosure belief, and conversely for $c < 0$. Furthermore, a lower threshold now corresponds to revealing less information.

The following comparative statics hold with an upward biased sender; the modifications for a downward biased sender are straightforward in light of the above discussion.

Proposition 2. *Assume there is only one sender, and this sender is upward biased.*

1. *A higher probability of being informed leads to more disclosure: the highest and lowest equilibrium disclosure thresholds (weakly) decrease.*
2. *An increase in disclosure cost or a reduction in concealment cost leads to less disclosure: the highest and lowest equilibrium disclosure thresholds (weakly) increase.*

The logic for the first part follows from Lemma 1: given any conjectured threshold, a higher p_i leads to a lower nondisclosure belief, which increases the sender's gain from disclosure over nondisclosure of any signal. For the case of $c = 0$, this comparative static has also been noted by other authors, e.g., Jung and Kwon (1988) and Acharya et al. (2011). The second part of Proposition 2 is straightforward, as a higher c makes disclosure less attractive. Since scaling the magnitude of c is equivalent to scaling the agent's bias parameter b_i (cf. expression (4)), an equivalent interpretation is that an agent with a stronger persuasion motive discloses more when $c > 0$ but less when $c < 0$.

Figure 1 summarizes the results of this subsection.[15]

[15] In the figure, $\eta(\cdot)$ has slope less than one when it crosses $s_i - c$ at the highest crossing point. This makes

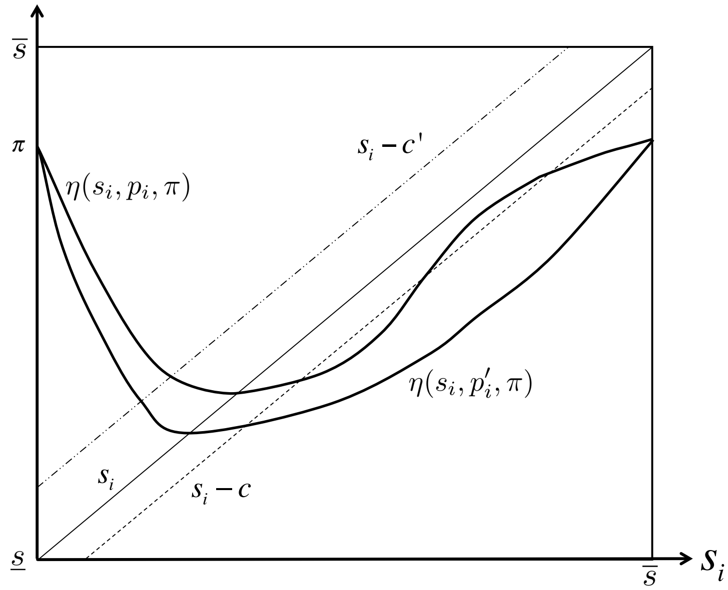

Figure 1. *The single-sender game with an upward biased sender, illustrated with parameters $c > 0 > c'$ and $p_i' > p_i$. Equilibrium disclosure thresholds are given by the intersections of $\eta(s_i, \cdot)$ and $s_i - c$ or $s_i - c'$.*

Although we postpone a formal argument to Subsection 3.3, it is worth observing now that the comparative statics on disclosure have direct welfare implications. Since the DM always prefers more disclosure, a lower message cost and/or a higher probability of the sender being informed (weakly) increases the DM's welfare in a single-sender setting, subject to an appropriate comparison of equilibria, in particular, focussing on the extremal equilibria.

3.3. Strategic substitutes and complements

We are now ready to study the two-sender disclosure game. For concreteness, we will suppose that both senders are upward biased; the modifications needed when one or both senders are downward biased are straightforward.

Lemma 2. *Any equilibrium is a threshold equilibrium, i.e., both senders use threshold strategies.*

In light of Lemma 2, we focus on threshold strategies. A useful simplification afforded by the assumption of conditionally independent signals is that the DM's belief updating is separable in the senders' messages. In other words, we can treat it as though the DM first updates from either sender i's message just as in a single-sender model, and then uses this

transparent that an increase in p_i leads to a reduction in the highest equilibrium threshold. If the slope of $\eta(\cdot)$ were larger than one at the highest crossing point, then the highest equilibrium threshold would be $\overline{s}$, and a small increase in p_i would not alter this threshold. We also note that the $\eta(\cdot)$ depicted in the figure is valid: it can be shown that for any continuously differentiable function, $\psi : [\underline{s}, \overline{s}] \to [0,1]$ with $\psi(\underline{s}) = \psi(\overline{s}) \in (0,1)$ and $\text{sign}[s - \psi(s)] = \text{sign}[\psi'(s)]$ (where ψ' denotes the derivative), there are parameters of the model—viz., $\pi \in (0,1)$, $p_i \in (0,1)$, $f(s|0)$, and $f(s|1)$—such that the nondisclosure belief $\eta(s_i, p, \pi) = \psi(s_i)$ for all s_i.

updated belief as an interim prior to update again from the other sender j's message without any further attention to i's message. Thus, given any conjectured pair of disclosure thresholds, $(\hat{s}_1, \hat{s}_2)$, there are three relevant nondisclosure beliefs for the DM: if only one sender i discloses his signal s_i while sender j sends message ϕ, the DM's belief is $\eta(\hat{s}_j, p_j, s_i)$; if there is nondisclosure from both senders, the DM's belief is $\eta(\hat{s}_j, p_j, \eta(\hat{s}_i, p_i, \pi))$.

Suppose the DM conjectures that sender i is using a disclosure threshold $\hat{s}_i$. As discussed earlier, if i discloses his signal then his expectation of the DM's belief—viewed as a random variable that depends on j's message—is s_i, no matter what strategy j is using.[16] On the other hand, if i conceals his signal, then he views the DM as updating from j's message based on a prior of $\eta(\hat{s}_i, p_i, \pi)$ that may be different from s_i. Denote j's disclosure threshold by $\hat{s}_j$ and any particular message sent by j as $m_j \in S \cup \{\phi\}$. Let $\beta(\mathcal{I}; q)$ denote the posterior belief given an arbitrary information set $\mathcal{I}$ and prior belief q. Then, sender i's posterior belief about the state given j's message and his own signal can be written as $\beta(m_j; s_i)$. The transformation (2) implies that i's expected payoff—equivalently, his expectation of the DM's posterior belief—should he conceal his signal is given by

$$U(s_i, p_i, \hat{s}_j, p_j) := \mathbb{E}_{\hat{s}_j, p_j}[T(\beta(m_j; s_i), \eta(\hat{s}_i, p_i, \pi), s_i)],$$

where $\mathbb{E}_{\hat{s}_j, p_j}$ denotes that the expectation is taken over m_j using the distribution of beliefs that $\hat{s}_j$ and p_j jointly induce in i about m_j (given s_i).

It is useful to study the "best response" of sender i to any disclosure strategy of sender j. More precisely, let $\hat{s}_i^{BR}(\hat{s}_j; p_i, p_j)$ represent the equilibrium disclosure threshold in a (hypothetical) game between sender i and the DM when sender j is conjectured to mechanically adopt disclosure threshold $\hat{s}_j$; we call this sender i's best response. The threshold $\hat{s}_i$ is a best response if and only if it satisfies any one of the following:

$$\begin{cases} U(\hat{s}_i, p_i, \hat{s}_j, p_j) = \hat{s}_i - c & \text{and } \hat{s}_i \in (\underline{s}, \overline{s}); \text{ or} \\ U(\underline{s}, p_i, \hat{s}_j, p_j) \leq \underline{s} - c & \text{and } \hat{s}_i = \underline{s}; \text{ or} \\ U(\overline{s}, p_i, \hat{s}_j, p_j) \geq \overline{s} - c & \text{and } \hat{s}_i = \overline{s}. \end{cases} \tag{6}$$

The necessity of condition (6) is clear; sufficiency follows from the argument given in the proof of Lemma 2. In any equilibrium of the overall game, (s_1^*, s_2^*), condition (6) must hold for each sender i with $\hat{s}_i = s_i^*$ when his opponent uses $\hat{s}_j = s_j^*$.

[16] The distribution of the DM's beliefs as a function of j's message depends both on j's strategy and the DM's conjecture about j's strategy. As the two must coincide in equilibrium, we bundle them to ease exposition.

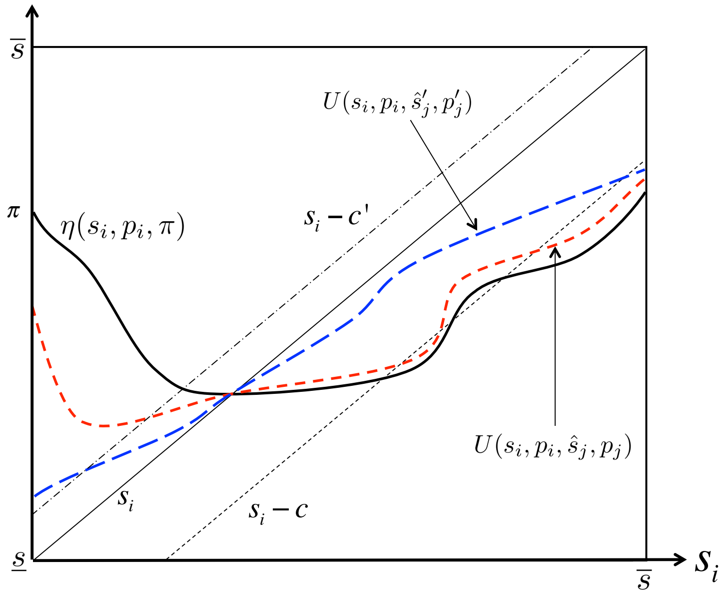

Figure 2. The "best response" of sender i to sender j (both upward biased), illustrated with parameters $c > 0 > c'$, $p'_j > p_j$, and $\hat{s}'_j < \hat{s}_j < \overline{s}$.

Lemma 3. *For any $\hat{s}_j$ and p_j,*

$$
\begin{aligned}
s_i = \eta(s_i, p_i, \pi) &\implies U(s_i, p_i, \hat{s}_j, p_j) = s_i, \\
s_i > \eta(s_i, p_i, \pi) &\implies \eta(s_i, p_i, \pi) \leq U(s_i, p_i, \hat{s}_j, p_j) < s_i, \\
s_i < \eta(s_i, p_i, \pi) &\implies \eta(s_i, p_i, \pi) \geq U(s_i, p_i, \hat{s}_j, p_j) > s_i.
\end{aligned}
$$

Moreover, both weak inequalities above are strict if and only if $\hat{s}_j < \overline{s}$.

Lemma 3 is a direct consequence of Theorem 1 (and Remark 1); it reflects that i expects j's information disclosure to bring the DM's posterior belief closer to s_i, which is i's belief based on his own signal. Graphically, this is seen in Figure 2 by comparing the red (short dashed) curve that depicts $U(\cdot)$ with the black (solid) curve depicting $\eta(\cdot)$. The former is a rotation of the latter curve around its fixed point toward the diagonal.

It is evident from Figure 2 that j's information disclosure has very different consequences for the best response of i depending on whether there is a cost of disclosure or a cost of concealment. If $c > 0$ (disclosure cost) then the smallest and largest solutions to (6) will be respectively larger than the smallest and largest single-sender disclosure thresholds. If $c < 0$ (concealment cost, depicted as c' in Figure 2) the largest solution will be smaller than the unique single-sender threshold.[17] If $c = 0$, the unique solution is the same as the single-sender threshold. These contrasting effects are due to the different nature of disagreement

[17] Although not seen in the figure, there can be multiple solutions to (6) even when $c < 0$.

induced by the sender. When there is a disclosure cost, the threshold type in any single-sender equilibrium has a *higher* belief than the DM upon nondisclosure, and hence an expected shift of the DM's posterior toward the threshold belief makes concealment more attractive. By contrast, when there is a concealment cost, the threshold type has a *lower* belief than the DM upon nondisclosure, and hence an expected shift of the DM's posterior toward the threshold belief makes concealment less attractive.

Theorem 1 implies that these insights are not restricted to comparisons with the single-sender setting. More generally, the same points hold whenever we compare any $(\hat{s}_j, p_j)$ with $(\hat{s}'_j, p'_j)$ such that $\hat{s}'_j \leq \hat{s}_j$ and $p'_j \geq p_j$. The latter experiment is more informative than the former because sender j is more likely to be informed and discloses more conditional on being informed. More precisely, notice that for any message m'_j under $(\hat{s}'_j, p'_j)$, one can garble it to produce message m_j as follows:

$$m_j = \begin{cases} \phi & \text{if } m'_j = \phi \text{ or } m'_j \in [\hat{s}'_j, \hat{s}_j), \\ m'_j \text{ with prob. } p_j/p'_j \text{ or } \phi \text{ with prob. } 1 - p_j/p'_j & \text{if } m'_j \geq \hat{s}_j. \end{cases}$$

In each state of the world, the distribution of m_j as just constructed is the same as the distribution of sender j's message under $(\hat{s}_j, p_j)$. Thus, the message produced under $(\hat{s}_j, p_j)$ is a garbling of that produced under $(\hat{s}'_j, p'_j)$.

The effect of sender j's message becoming more informative is depicted in Figure 2 as a shift from the red (short dashed) curve to the blue (long dashed) curve. Whether sender i's best response is to disclose less or more of his own information turns on whether $c > 0$ or $c < 0$; the logic is the same as discussed earlier. For any given $\hat{s}_j$, there can be multiple solutions to (6), hence $\hat{s}_i^{BR}(\cdot)$ is a best-response correspondence. We say that sender i's best response increases if the largest and the smallest element of $\hat{s}_i^{BR}(\cdot)$ both increase (weakly); we say that the best response is strictly greater than some $\hat{s}$, written $\hat{s}_i{}^{BR}(\cdot) > \hat{s}$, if the smallest element of $\hat{s}_i^{BR}(\cdot)$ is strictly greater than $\hat{s}$.

Proposition 3. *Assume both senders are upward biased. Any sender i's best-response disclosure threshold $\hat{s}_i^{BR}(\hat{s}_j; p_i, p_j)$ is decreasing in p_i. Furthermore, let $\hat{s}_i^0$ denote the unique (resp., smallest) equilibrium threshold in the single-sender game with i when $c \leq 0$ (resp., $c > 0$).*

1. *(Independence). If $c = 0$, then $\hat{s}_i^{BR}(\hat{s}_j; p_i, p_j) = \hat{s}_i^0$ is independent of $\hat{s}_j$ and p_j.*
2. *(Strategic complements). If $c < 0$, then (i) $\hat{s}_i^{BR}(\hat{s}_j; p_i, p_j) \leq \hat{s}_i^0$, with equality if and only if $\hat{s}_i^0 = \underline{s}$ or $\hat{s}_j = \overline{s}$, and (ii) $\hat{s}_i^{BR}(\hat{s}_j; p_i, p_j)$ increases in $\hat{s}_j$ and decreases in p_j.*
3. *(Strategic substitutes). If $c > 0$, then (i) $\hat{s}_i^{BR}(\hat{s}_j; p_i, p_j) \geq \hat{s}_i^0$, with equality if and only if $\hat{s}_i^0 = \overline{s}$*

or $\hat{s}_j = \overline{s}$, and (ii) $\hat{s}_i^{BR}(\hat{s}_j; p_i, p_j)$ decreases in $\hat{s}_j$ and increases in p_j.

Since each sender's best response is monotone, existence of an equilibrium in the two-sender game follows from Tarski's fixed point theorem. When $c < 0$ (concealment cost), the strategic complementarity in disclosure thresholds implies that there is a largest equilibrium, which corresponds to the highest equilibrium thresholds for both senders. Each sender's message in the largest equilibrium is a garbling of his message in any other equilibrium. It follows that the largest equilibrium is the least informative and the *worst* in terms of the DM's welfare. Similarly, the smallest equilibrium—that with the lowest equilibrium thresholds for both senders—is the most informative and the *best* in terms of the DM's welfare. On the other hand, when $c > 0$ (disclosure cost), the two disclosure thresholds are strategic substitutes. There is an i-*maximal* equilibrium that maximizes sender i's threshold and also minimizes sender j's threshold across all equilibria. Likewise, there is a j-*maximal* equilibrium that minimizes sender i's threshold and also maximizes sender j's threshold across all equilibria. These two equilibria are not ranked in terms of (Blackwell) informativeness and in general cannot be welfare ranked for the DM; moreover, neither of these equilibria may correspond to either the best or the worst equilibrium for the DM.[18]

The following result is derived using standard monotone comparative statics arguments. Although the formal statement refers to extremal equilibria, the comparative statics also obtain for any equilibrium that is stable in the sense of adaptive dynamics; see Echenique (2002).

Proposition 4. *Assume both senders are upward biased. For any $i \in \{1, 2\}$:*

1. *If $c \leq 0$, then an increase in p_i or a decrease in c (a higher concealment cost) weakly lowers the disclosure thresholds of both senders in both the worst and the best equilibria.*
2. *If $c > 0$, then an increase in p_i weakly lowers sender i's disclosure threshold and weakly raises sender j's disclosure threshold in both the i-maximal and the j-maximal equilibria. A decrease in c (a lower disclosure cost) has ambiguous effects on the two senders' equilibrium disclosure thresholds in both the i- and j-maximal equilibria.*

One can view the single-sender game with i as a two-sender game where sender j is never informed, i.e., $p_j = 0$. With this in mind, a comparison of the single-sender game and the two-sender game can be obtained as a corollary to Proposition 3 and Proposition 4.

[18] As explained after Proposition 1, the senders' welfare ranking across equilibria just depends on the probability of disclosure. When $c < 0$, both senders' welfare is lowest in the largest equilibrium and highest in the smallest equilibrium. When $c > 0$, sender i's welfare is highest in the i-maximal equilibrium and lowest in the j-maximal equilibrium.

Corollary 3. *Assume both senders are upward biased and let $\hat{s}_i^0$ denote the unique (resp., smallest) equilibrium threshold in the single-sender game with i when $c \leq 0$ (resp., $c > 0$).*

1. *If $c = 0$, equilibrium in the two-sender game is unique and is equal to $(\hat{s}_1^0, \hat{s}_2^0)$. The DM's welfare is strictly higher in the two-sender game than in a single-sender game with either sender.*
2. *If $c < 0$, every equilibrium in the two-sender game is weakly smaller than $(\hat{s}_1^0, \hat{s}_2^0)$, with equality if and only if $\hat{s}_1^0 = \hat{s}_2^0 = \underline{s}$. The DM's welfare is strictly higher in any equilibrium of the two-sender game than in a single-sender game with either sender.*
3. *If $c > 0$, every equilibrium in the two-sender game is weakly larger than $(\hat{s}_1^0, \hat{s}_2^0)$, with equality if and only if $\hat{s}_1^0 = \hat{s}_2^0 = \overline{s}$. The DM's welfare in the best equilibrium of the two-sender game may be higher or lower than in the best equilibrium of the single-sender game with sender i or sender j alone.*

Part 1 of Corollary 3 follows from part 1 of Proposition 3. When $c = 0$, the best response of each sender is to use the same disclosure threshold as in the single-sender setting, regardless of the other sender's strategy. Since the DM receives two messages instead of just one, and the probability distribution of these messages remain the same as in the single-sender game, she is better off when facing both senders than when facing either sender alone.

Part 2 of Corollary 3 can be obtained by considering the worst equilibrium of the two-sender game. For the case of concealment cost ($c < 0$), let $\mathrm{s}^*(p_i, p_j)$ represent the vector of disclosure thresholds in the worst equilibrium. Proposition 4 implies that $s_i^*(p_i, p_j) \leq s_i^*(p_i, 0) = \hat{s}_i^0$ and $s_j^*(p_i, p_j) \leq s_j^*(0, p_j) = \hat{s}_j^0$. Thus $\mathrm{s}^*(p_i, p_j)$ is weakly smaller than $(\hat{s}_i^0, \hat{s}_j^0)$, and hence every equilibrium is weakly smaller than $(\hat{s}_i^0, \hat{s}_j^0)$. It follows that the DM's welfare is higher than in the unique equilibrium of the single-sender game with either sender. This higher welfare is due to both a direct effect of receiving information from an additional sender, and an indirect effect wherein each sender is now disclosing more than in the single-sender setting.

Finally, for the case of disclosure cost ($c > 0$), let $\mathrm{s}^{i*}(p_i, p_j)$ represent the i-maximal equilibrium and $\mathrm{s}^{j*}(p_i, p_j)$ represent the j-maximal equilibrium. In any equilibrium, i's threshold is at least as large as $s_i^{j*}(p_i, p_j) \geq s_i^{j*}(p_i, 0) = \hat{s}_i^0$, where the inequality is by part 2 of Proposition 4. Analogously, sender j's threshold in any equilibrium is at least as large as $s_j^{i*}(p_i, p_j) \geq s_j^{i*}(0, p_j) = \hat{s}_j^0$. Thus, both senders are (weakly) less informative than in the DM's best equilibrium of the single-sender game. The overall welfare comparison between the two-sender game and the single-sender game is generally ambiguous. While adding a second sender has a direct effect of increasing the DM's information, there is an adverse indirect effect due to the strategic substitution in disclosure of the other sender. It is possible the net effect can (strictly) reduce the DM's welfare—even when the two senders have opposite biases, which is often thought to particularly promote information disclosure. An example with a familiar quadratic loss function for the DM is available from the authors on request.

It is appropriate to compare our welfare results with Bhattacharya and Mukherjee (2013). They study a related model to ours, allowing senders to have single-peaked preferences over the DM's posterior, but they assume $c = 0$ (no message costs) and perfectly correlated signals. They show that an increase in the probability of a sender being informed can reduce the DM's welfare. However, a necessary condition for this to happen in their model is that at least one sender must have non-monotonic preferences over the DM's posterior, which in turn implies (because senders have single-peaked preferences) that the senders share the same ranking over decisions on a subset of the decision space. In this sense, their result requires that senders are not in "pure conflict," whereas our setup allows senders to have diametrically opposing preferences. More broadly, our results on equilibrium behavior and welfare for $c \neq 0$ are orthogonal and complementary to their treatment of non-monotonic preferences.

The assumption that senders' signals are conditionally independent is clearly important for our analytical methodology, as without it we cannot apply Theorem 1. If the signals are conditionally correlated, then upon nondisclosure sender i and the DM disagree not only on the probability assessment of the states, but also on the experiment corresponding to sender j's message. Relaxing the conditional independence assumption to obtain a general analysis appears intractable. We illustrate in Supplementary Appendix C.1 how some of our substantive economic conclusions would change under a significantly different information structure: perfectly correlated signals. Another assumption that is important in applying Theorem 1 is that each sender has linear preferences. Supplementary Appendix C.2 discuses how our result under $c = 0$ extends to non-linear preferences and how our results under $c \neq 0$ may or may not hold under non-linear preferences.

3.4. Many senders

Our results readily generalize to any finite number of senders. Suppose in addition to senders i and j, there are K other senders, all of whom simultaneously send messages to the DM. Let $\mathbf{m}$ represent the collection of these K messages. Then, sender i's posterior belief given his own signal s_i, sender j's message m_j, and the K other senders' messages $\mathbf{m}$ is $\beta(m_j, \mathbf{m}; s_i) = \beta(m_j; \beta(\mathbf{m}; s_i))$. The DM's belief given the K senders' messages $\mathbf{m}$ and given nondisclosure by sender i is $\eta(\hat{s}_i, p_i, \beta(\mathbf{m}; \pi))$. Thus, the transformation mapping from (2) and the law of iterated expectations imply that the expected payoff for sender i from concealing his signal is

$$\mathbb{E}\left[\mathbb{E}_{\hat{s}_j,p_j}[T(\beta(m_j;\beta(\mathbf{m};s_i)),\eta(\hat{s}_i,p_i,\beta(\mathbf{m};\pi)),\beta(\mathbf{m};s_i))\mid \mathbf{m}]\right].$$

The inside conditional expectation (given $\mathbf{m}$) is taken over the distribution of m_j, while the outside expectation is taken over the distribution of $\mathbf{m}$ generated from the equilibrium strategies of the K senders. Given any $\mathbf{m}$, the transformation $T(\cdot)$ in the multi-sender case is the same as that in the two-sender case, with the common prior π replaced by $\beta(\mathbf{m};\pi)$. Since our results hold for any π, the logic of strategic substitution or strategic complementarity continues to apply in the multi-sender case. In particular, when $c<0$, $\mathbb{E}_{\hat{s}_j,p_j}[T(\cdot)\mid\mathbf{m}]$ increases in $\hat{s}_j$ and decreases in p_j for any $\mathbf{m}$. Consequently, sender i's expected payoff from nondisclosure, $\mathbb{E}\left[\mathbb{E}_{\hat{s}_j,p_j}[T(\cdot)\mid\mathbf{m}]\right]$ also increases in $\hat{s}_j$ and decreases in p_j. Thus disclosure by any two senders are strategic complements. Similarly, in the case of disclosure cost (i.e., $c>0$), disclosure by any two senders are strategic substitutes.

It follows from these observations that when there is either no message cost or a concealment cost ($c\leq 0$), the DM always benefits from having more senders to supply her with information. When there is disclosure cost ($c>0$), on the other hand, an increase in the number of senders has ambiguous effects on each sender's disclosure threshold, and can lead to either an increase or decrease in the DM's welfare.

3.5. Sequential reporting

The key insight from our analysis of simultaneous disclosure extends to sequential disclosure. For concreteness, consider a two-sender game in which both senders are upward biased but disclosure is sequential: sender 1 reports first and his message m_1 is made public to both the DM and sender 2 before sender 2 submits his report. Sender 2 now effectively faces a single-sender problem where he and the DM share a common prior, say $\beta(m_1;\pi)$, which is a function to be determined in equilibrium. Proposition 2 implies that sender 2 will adopt a disclosure threshold $\hat{s}_2^0$ which depends negatively on p_2.

Consider now the disclosure decision of sender 1 when the DM conjectures that he is using a disclosure threshold $\hat{s}_1$, with corresponding nondisclosure belief $\eta(\hat{s}_1,p_1,\pi)$. If sender 1 discloses his signal s_1, his expectation of the DM's posterior belief is simply s_1. If he chooses nondisclosure, his expectation is $\mathbb{E}_{\hat{s}_2^0,p_2}[T(\beta(m_2;s_1),\eta(\hat{s}_1,p_1,\pi),s_1)]$. Since sender 2 discloses more when he is better informed, a higher p_2 makes the message m_2 more informative, both directly through a higher probability of sender 2 getting a signal and indirectly through a lower disclosure threshold $\hat{s}_2^0$. IVP implies that the DM's belief is expected to move away from $\eta(\hat{s}_1, p_1, \pi)$ toward s_1. The same logic that establishes Proposition 4 therefore gives the following result, whose proof is omitted.

Proposition 5. *Consider sequential disclosure and assume sender 1, the first mover, is upward biased. If $c > 0$ (resp., $c < 0$), a higher p_2 weakly increases (resp., weakly lowers) the equilibrium disclosure threshold of sender 1 in the 1-maximal and 2-maximal equilibria.*

An immediate corollary to Proposition 5 is that, in the case of concealment cost, the first sender discloses more than he does in a single-sender setting. As a result, the DM is always better off in a sequential game than with sender 1 alone. On the other hand, a welfare comparison between the sequential game and the simultaneous move game is generally ambiguous.

We also note that if $c = 0$ the irrelevance result still holds for sender 1: the first sender adopts the same disclosure threshold under sequential reporting as the disclosure threshold in the single-sender problem. The disclosure threshold chosen by the second sender, however, depends on the message sent by sender 1; it may be higher or lower than sender 2's disclosure threshold were he the only sender.

4. Costly Signaling

A key feature of the multi-sender disclosure game of Section 3 is that message costs induce equilibrium disagreement between a sender and the receiver. Since a sender needs to predict how the receiver reacts to the message from another sender, whose informativeness is itself an equilibrium object, our IVP result is particularly pertinent because it provides a useful general tool regarding expectations about beliefs when individuals disagree.

In this section, we illustrate in a costly-signaling application how IVP is useful even when there may be no disagreement *in equilibrium* (because of separation); rather, its usefulness comes from considerations of how disagreement *off the equilibrium path* would be affected by new information. The application also illustrates how IVP is applicable even with just one sender and an exogenous information source.

4.1. The model

We consider a communication game with lying costs, a variation of Kartik (2009). A sender and a receiver share a common prior belief about a state $\omega \in \Omega = \{0, 1\}$. The sender has type $t \in [0, 1]$, which is drawn from a distribution $F(\cdot|\omega)$ with corresponding density $f(\cdot|\omega)$. As in Section 3, we assume without loss of generality that t is the sender's private belief that $\omega = 1$. For simplicity, we assume $f(t|\omega) > 0$ for all $t \in (0, 1)$ and ω. We sometimes refer to t as the "truth." The sender sends a message $m \in [0,1]$ that entails a cost $c(m,t)$, elaborated below. The receiver forms a belief based on both m and an additional signal $s \in [0,1]$ that, conditional on the state ω, is drawn independently of t or m from a distribution $G(\cdot|\omega)$ with density $g(\cdot|\omega)$. Without loss, we assume that the signal s satisfies the weak monotone likelihood ratio property. The signal s can either be the receiver's private information or it can be publicly observed, but only after the sender has chosen his message m. We assume that no realization of s perfectly reveals ω. Denote the receiver's posterior expectation of the state by $\mathbb{E}[\omega|m,s]$. The sender's payoff is linear in the receiver's expectation; specifically, his payoff is

$$\mathbb{E}[\omega|m,s] - c(m,t),$$

That is, the sender is upward biased and prefers a higher posterior belief in the receiver. (Our analysis carries over to a downward bias with straightforward modifications.)

The receiver's belief updating process can be analyzed in two steps: based on the message m from the sender, she forms an interim belief $\pi(m)$ about the state (all beliefs about the state refer to $\Pr[\omega = 1]$) and then uses this interim belief to further update from the signal s to form a posterior belief $\beta(s;\pi(m))$. By Bayes rule,

$$\beta(s;\pi) = \frac{\pi g(s|1)}{\pi g(s|1) + (1-\pi) g(s|0)}.$$

The expected payoff of a type-t sender from sending message m is thus

$$\mathbb{E}_{s|t}[\beta(s;\pi(m))] - c(m,t),$$

where $\mathbb{E}_{s|t}$ denotes the type-t sender's expectation over the signal s. An important observation is that the first term in the above display is strictly increasing in $\pi(m)$: the sender prefers the receiver to hold a higher interim belief.

As in Kartik (2009), assume the cost function $c(\cdot)$ is smooth with $\partial c(t,t)/\partial m = 0$ for all t, i.e., the marginal cost of lying when one is telling the truth is zero. Furthermore, the marginal cost of sending a higher message is increasing in m and decreasing in t, i.e., $\partial^2 c(m,t)/\partial m^2 > 0 > \partial^2 c(m,t)/\partial m \partial t$.

A single-crossing condition. Unlike in a setting without exogenous information for the receiver, the above cost assumptions do not assure a suitable single-crossing property of the sender's indifference curves. The indifference curves of interest are those in the space of the sender's message m and the receiver's interim belief π. These indifference curves are upwardsloping for $m > t$: sending a higher message than the truth requires compensation through a higher interim belief. The single-crossing property we require is that these indifference curves are flatter for higher types, meaning that higher types are more willing to send higher messages (among those above the truth) to induce a given increase in interim belief. It turns out that this property is assured by the following assumption which we will maintain:

$$\text{For } m > t: \quad \frac{\partial^2 c(m,t)/\partial m \partial t}{\partial c(m,t)/\partial m} \leq -\frac{1}{1-t}. \tag{7}$$

Condition (7) can be interpreted as saying that higher types have a sufficiently large marginal lying cost advantage relative to marginal lying costs.[19] An example of a lying cost function that satisfies all our requirements is $c(m,t) = (m-t)^2$.

Lemma 4. $\dfrac{\partial c(m,t)/\partial m}{\partial \mathbb{E}_{s|t}[\beta(s;\pi)]/\partial \pi}$ *strictly decreases in* t *for* $t < m$.

Another interpretation. Although our model is posed as a communication game, we may also interpret it as a variation of the standard Spence (1973) signaling model. In this interpretation, a worker possesses a private trait/characteristic, his type t (e.g., intelligence), which is indicative of whether his job productivity will be high or low, but a potential employer may also observe a signal s about his productivity (e.g., through the job interview process). The wage of a worker depends on the employer's posterior of the worker's productivity, given the schooling level m chosen by the worker and the employer's own signal s. The marginal cost of schooling is decreasing in the characteristic t ($\partial^2 c(m,t)/\partial m \partial t < 0$). The constraint that $m \in [0,1]$ introduces an upper bound on the signaling level m, but this assumption can be relaxed. Further, while we suppose that higher-type workers intrinsically prefer acquiring more education ($\partial c(t,t)/\partial m = 0$ for all t), our analysis also applies under the more traditional assumption that all workers prefer less education ($\partial c(m,t)/\partial m > 0$ for all m, t). See the discussion after Proposition 6. We also note that the model can be reformulated with discrete types with little change in the subsequent analysis, except for possible mixing by some type.

4.2. Equilibrium and comparative statics

A fully separating outcome is not supportable as a (weak perfect Bayesian) equilibrium, for reasons similar to those discussed in Kartik (2009). We focus on *low types separate and high*

[19] Similarly, in Daley and Green (2014), a marginal cost advantage of the higher type does not guarantee single crossing of "belief indifference curves" in the action-interim belief space; rather, single crossing requires the marginal cost advantage to be sufficiently large.

types pool—hereafter LSHP—equilibria. That is, we will look for equilibria in which there is a cutoff $\underline{t} \in [0, 1]$ such that the sender's (pure) strategy, $\mu : [0, 1] \to [0, 1]$, satisfies:

1. Riley condition: $\mu(0) = 0$.
2. Separation at bottom: $\mu(\cdot)$ is strictly increasing on $[0, \underline{t})$.
3. Pooling at top: $\mu(t) = 1$ for all $t > \underline{t}$.

This class of equilibria extends the standard least cost separating equilibrium logic to a bounded signal space. See Kartik (2009) for more detailed discussion and justification of LSHP equilibria, albeit when there is no exogenous information for the receiver.

In an LSHP equilibrium, the receiver can infer the sender's type if his message m belongs to the interval $[0, \mu(\underline{t}))$. For such m, her interim belief upon observing m is simply $\pi(m) = \mu^{-1}(m)$. If $m = 1$, the receiver infers that $t > \underline{t}$; so $\pi(1) = \pi^P(\underline{t}) := \mathbb{E}[\omega | t > \underline{t}]$, where the superscript P is mnemonic for "pooled." We note that $\pi^P(\underline{t}) \in (\underline{t}, 1)$.

The equilibrium strategy $\mu(\cdot)$ for $t < \underline{t}$ is determined by the following differential equation:

$$\frac{\partial c(\mu(t), t)}{\partial m} \mu'(t) = \frac{\partial \mathbb{E}_{s|t}\left[\beta(s; t)\right]}{\partial \pi}, \tag{8}$$

with boundary condition $\mu(0) = 0$. Equation 8 is obtained from the binding local upward incentive compatibility constraints. The left-hand-side is the marginal cost for type t of mimicking a slightly higher type; the right-hand-side is the marginal benefit, which comes from inducing a higher interim belief in the receiver. Note that the benefit is affected by the sender's belief about the exogenous signal s. Since $\mu'(t) > 0$ and the right-hand-side is strictly positive, we must have $\mu(t) > t$ for all $t > 0$. Thus, the solution $\mu(\cdot)$ to the boundary-value problem will hit 1 at some interior $\bar{t}$, i.e., $\mu(\bar{t}) = 1$ for some $\bar{t} \in (0, 1)$. The cutoff $\underline{t} \in [0, \bar{t})$ must satisfy:

$$\underline{t} > 0 \implies \mathbb{E}_{s|\underline{t}}\left[\beta(s; \pi^P(\underline{t}))\right] - c(1, \underline{t}) = \mathbb{E}_{s|\underline{t}}\left[\beta(s; \underline{t})\right] - c(\mu(\underline{t}), \underline{t}), \tag{9}$$

$$\underline{t} = 0 \implies \mathbb{E}_{s|0}\left[\beta(s; \pi^P(0))\right] - c(1, 0) \geq \mathbb{E}_{s|0}\left[\beta(s; 0)\right] - c(0, 0). \tag{10}$$

Condition (9) requires that if $\underline{t}$ is interior, then type $\underline{t}$ must be indifferent between sending message $m = 1$ (pooling with types above him) and sending message $\mu(\underline{t})$ (revealing his true type). In this case $\mu(\cdot)$ is discontinuous at $t = \underline{t}$. If (10) holds, then every type prefers to pool than separate, so there is complete pooling.

We have explained why local incentive compatibility in an LSHP equilibrium implies the differential equation (8) with boundary condition $\mu(0) = 0$, and conditions (9) and (10) for the cutoff type $\underline{t}$. The assumption of (7), which yields the single-crossing property stated in

Lemma 4, guarantees that these conditions are also sufficient for global incentive compatibility. We can thus use these conditions to establish:

Lemma 5. *There is a unique LSHP equilibrium.*

In proving Lemma 5, we assign the off-path belief $\pi(m) = \underline{t}$ for $m \in (\mu(\underline{t}), 1)$, and show that no type t has incentive to deviate from $\mu(t)$. It can be shown that this off-path belief satisfies the natural adaptation of the D1 refinement (Cho and Sobel, 1990) to our setting.

Lemma 5 identifies a unique cutoff type $\underline{t}$ above which pooling occurs. The information the receiver gets about the sender's type is fully characterized by the cutoff: a higher cutoff corresponds to a more (Blackwell) informative equilibrium. The following result will be crucial in understanding how changes in the receiver's exogenous information affects the sender's equilibrium signaling strategy.

Lemma 6. *If $\tilde{s}$ is drawn from a more informative experiment than s, then for any $t < 1$,*

1. $\partial \mathbb{E}_{\tilde{s}|t} \left[\beta(\tilde{s}; t)\right] / \partial \pi \leq \partial \mathbb{E}_{s|t} \left[\beta(s; t)\right] / \partial \pi$.
2. $\mathbb{E}_{\tilde{s}|t} \left[\beta(\tilde{s}; \pi^P(t))\right] \leq \mathbb{E}_{s|t} \left[\beta(s; \pi^P(t))\right]$.

Roughly, Lemma 6 says that the sender's gain from inducing a higher interim belief than the truth is lower when the exogenous signal is more informative. For part 1 of the lemma, note that for any type $t < 1$ and any small $\varepsilon > 0$, inducing an interim belief $t + \varepsilon$ creates disagreement: type t views this interim belief as higher than the truth. IVP implies that t expects a more informative exogenous signal to correct this interim belief more, so that t's expected utility gain from inducing that interim belief is lower under a more informative experiment. More precisely, when $\tilde{s}$ is more informative than s, Theorem 1 implies that

$$\mathbb{E}_{\tilde{s}|t} \left[\beta(\tilde{s}; t + \varepsilon)\right] - \mathbb{E}_{\tilde{s}|t} \left[\beta(\tilde{s}; t)\right] \leq \mathbb{E}_{s|t} \left[\beta(s; t + \varepsilon)\right] - \mathbb{E}_{s|t} \left[\beta(s; t)\right],$$

where we have used $\mathbb{E}_{\tilde{s}|t}[\beta(\tilde{s}; t)] = \mathbb{E}_{s|t}[\beta(s; t)] = t$. Dividing both sides of the above inequality by ε and taking $\varepsilon \to 0$ proves part 1 of Lemma 6. Part 2 of the lemma also follows from Theorem 1, since $\pi^P(t) > t$.

Part 1 of Lemma 6 implies that a more informative exogenous signal reduces the right-hand side of Equation 8. Since each type expects a smaller marginal benefit from inducing a higher belief in the receiver, the solution μ to the differential equation (8) with boundary condution $\mu(0) = 0$ is pointwise lower. Furthermore, part 2 of Lemma 6 implies that a more informative signal reduces the left-hand side of the consequent in (9), while the right-hand side is increased because $\mu(t)$ is lower. Thus, type $\underline{t}$ of the sender now strictly prefers revealing his type to pooling with higher types. Consequently, the equilibrium cutoff type increases when the receiver has access to a more informative signal. In summary:

Proposition 6. *Let $\tilde{s}$ be drawn from a more informative experiment than s. The LSHP equilibrium is more informative under $\tilde{s}$ than s: $\underline{t}_{\tilde{s}} \geq \underline{t}_s$. Furthermore, in the LSHP equilibrium, every type bears a lower signaling cost under $\tilde{s}$ than under s.*

Since Proposition 6 applies to arbitrary experiments, we can readily adapt the foregoing analysis to study a multi-sender signaling game in which each sender gets a conditionally independent signal about the state. For concreteness, suppose there are two upward biased senders, and each adopts an LSHP strategy.[20] Then, from one sender's perspective, the other sender's message is an endogenous experiment. Furthermore, the informativeness of that experiment is increasing in the other sender's cutoff. The logic of Proposition 6 can be used to deduce that the two senders' cutoffs are strategic complements: each sender reveals more and bears a lower signaling cost when the other sender reveals more. If any sender's cost—relative to the persuasion incentive—increases in a suitable sense (for example, if sender i's cost is $k_i c(m,t)$ and $k_i > 0$ increases) then all senders become more informative.

It bears emphasis that the second part of Proposition 6 holds even when there is no pooling at the top. Consider, for example, a canonical model à la Spence (1973) with an unbounded signal space, $m \in [0,\infty)$, and differentiable cost function $kC(m,t)$ where $k > 0$, $\partial C/\partial m > 0$, $\partial^2 C/\partial m^2 > 0$, and $\partial^2 C/\partial m \partial t < 0$. In this setting, (the analog of) the familiar least cost separating equilibrium is given by the solution to Equation 8 with the boundary condition $\mu(0) = 0$. Given condition (7), Lemma 6 (in particular, part 1) still holds: the benefit from mimicking a marginally higher type is lower when the receiver has access to better information. The right-hand side of Equation 8 is thus lower for each type. Consequently, there will still be full separation, but every type bears a lower signaling cost.

Although it is intuitive that more informative exogenous information reduces the sender's gain from misrepresenting his type, and therefore reduces the sender's incentive to incur misrepresentation costs, we emphasize that Proposition 6 relies on the logic of IVP. Supplementary Appendix C.3 contains examples showing that if the sender's payoff is not linear in the receiver's belief, or if the receiver's exogenous information is not from an MLRP-experiment (in a multi-state extension of this section's model), then the sender's marginal benefit from

[20] Any combination of biases can be accommodated, so long as one focus on appropriate equilibria; in particular, a downward biased sender deflates rather than inflates his messages.

mimicking a higher type can be *higher* when the receiver has access to more information, which leads to higher equilibrium signaling costs when the receiver is better informed.

Frank (1985, Section III) also suggests that better exogenous information can reduce dissipative signaling. Weiss (1983) studies when exogenous information allows for separating equilibria even when there is no heterogeneity in the direct costs of signaling; his focus is not comparative statics in the quality of exogenous information. Daley and Green (2014) emphasize the stability of non-separating equilibria when the marginal cost advantage of higher types is insufficiently large relative to the accuracy of exogenous information; this leads to a "double crossing" of appropriate indifference curves, contrary to our single-crossing assumption (cf. fn. 19). Truyts (2015) shows that under noisy signaling—rather than the noiseless case more commonly studied, including here—better exogenous information can exacerbate dissipative signaling.

5. Conclusion

To recap, this paper makes three contributions. First, it provides a fundamental result that "information validates the prior" (IVP, Theorem 2): in an appropriate sense, Anne expects more information to bring Bob's posterior closer to Anne's prior. This result can also be interpreted in terms of their expected disagreement (Corollary 2). Second, it demonstrates how this statistical result concerning agents with heterogenous priors can be fruitfully applied to familiar common-prior environments with asymmetric information—specifically, to disclosure and costly signaling games (Section 3 and Section 4, respectively). Third, it offers new substantive insights into these economic problems. For example, in the disclosure context, the nature of message costs are critical: disclosure costs induce strategic substitutes between senders' disclosure; concealments costs induce strategic complements. In the former scenario, but not the latter, a receiver can be worse off when there are more senders.

While we have focussed on two applications with asymmetric information in this paper, we believe IVP is instructive more broadly. Consider sender-receiver persuasion via information design (Kamenica and Gentzkow, 2011). Alonso and Câmara (2016) develop an elegant general analysis under heterogeneous priors; see also Galperti (2018). When the sender's preferences are state-independent and concave in the receiver's expectation of the state, there is no scope for beneficial persuasion under common priors. Alonso and Câmara (2016, Section 4.3) show that this observation does not hold generically with heterogenous priors when there are more than two states. Our Theorem 2 delivers additional insights. For instance, if the priors are likelihood-ratio ordered and only MLRP-experiments are available, then the sender does not benefit from persuasion if the receiver's prior is favorable in the sense that the receiver's prior dominates (resp., is dominated by) the sender's when the sender's utility is increasing (resp., decreasing) in the receiver's posterior expectation.

We close by highlighting that IVP only speaks to the ex-ante expectation of the posterior—more precisely, certain statistics of the posterior, such as the posterior mean. The ex-ante expectation is, of course, a salient property, and also relevant in various economic problems. Nevertheless, it does restrict the direct applicability of IVP; in particular, our applications assumed certain linearity of utility functions. Future research might explore how more information affects other statistics of one agent's ex-ante belief about another agent's posterior, and how to deploy such results.

Appendices

A. Proofs

Proof of Theorem 1. This is a special case of Theorem 2 below, because with two states, any pair of priors are likelihood-ratio ordered and any experiment is an MLRP-experiment. □

Proof of Theorem 2. For simplicity, we write the proof assuming that experiments have signal distributions with (Radon-Nikodym) densities—subsuming probability mass functions—in each state.

Suppose $M(\overline{\boldsymbol{\beta}}_B) \geq M(\overline{\boldsymbol{\beta}}_A)$. Since the priors are likelihood-ratio ordered, the previous inequality is equivalent to $\overline{\boldsymbol{\beta}}_B$ dominating $\overline{\boldsymbol{\beta}}_A$ in the likelihood-ratio order. Let $p(s|\omega)$ represent the density function of s in state ω under experiment $\mathcal{E}$, and let $\tilde{p}(\tilde{s}|\omega)$ represent the density function under experiment $\tilde{\mathcal{E}}$. By the definition of garbling, there exists a non-negative kernel density $q(\tilde{s}|s)$ with $\int_{\tilde{s}} q(\tilde{s}|s)\,\mathrm{d}\tilde{s} = 1$ for all s, such that for any state ω,

$$\tilde{p}(\tilde{s}|\omega) = \int_s q(\tilde{s}|s)p(s|\omega)\,\mathrm{d}s.$$

In the following, for $k = A, B$, we let $p_k(s) := \sum_\omega p(s|\omega)\overline{\beta}_k(\omega)$ and $\tilde{p}_k(\tilde{s}) := \sum_\omega \tilde{p}(\tilde{s}|\omega)\overline{\beta}_k(\omega)$. The posterior density function over S conditional on $\tilde{s}$ and with prior belief $\overline{\boldsymbol{\beta}}_k$ is given by

$$\hat{q}_k(s|\tilde{s}) = \frac{q(\tilde{s}|s)p_k(s)}{\tilde{p}_k(\tilde{s})}.$$

Therefore,

$$\frac{\hat{q}_B(s|\tilde{s})}{\hat{q}_A(s|\tilde{s})} = \frac{\tilde{p}_A(\tilde{s})}{\tilde{p}_B(\tilde{s})}\frac{p_B(s)}{p_A(s)}.$$

Because $\mathcal{E}$ is an MLRP-experiment and $\overline{\boldsymbol{\beta}}_B$ likelihood-ratio dominates $\overline{\boldsymbol{\beta}}_A$, the ratio $p_B(s)/p_A(s)$ increases in s. Therefore, for any $\tilde{s}$, $\hat{q}_B(\cdot|\tilde{s})$ likelihood-ratio dominates $\hat{q}_A(\cdot|\tilde{s})$.

Since $M(\boldsymbol{\beta}_B(s))$ increases in s, the likelihood-ratio dominance of $\hat{q}_B(\cdot|\tilde{s})$ over $\hat{q}_A(\cdot|\tilde{s})$ implies

$$\mathbb{E}_B^{\mathcal{E}}\left[M\left(\boldsymbol{\beta}_B(s)\right) \mid \tilde{s}\right] \geq \mathbb{E}_A^{\mathcal{E}}\left[M\left(\boldsymbol{\beta}_B(s)\right) \mid \tilde{s}\right]. \tag{11}$$

For any realization $\tilde{s}$ from experiment $\tilde{\mathcal{E}}$,

$$\mathbb{E}_B^{\mathcal{E}}\left[M\left(\boldsymbol{\beta}_B(s)\right) \mid \tilde{s}\right] = M\left(\boldsymbol{\beta}_B(\tilde{s})\right), \tag{12}$$

because, by the law of iterated expectation, B's expectation of any linear function of B's posterior generated by experiment $\mathcal{E}$ is B's belief prior to that experiment. Combining (11) and (12) and taking the expectation over $\tilde{s}$ (using the prior $\overline{\boldsymbol{\beta}}_A$) gives

$$\begin{aligned}
\mathbb{E}_A^{\tilde{\mathcal{E}}}\left[M\left(\boldsymbol{\beta}_B(\tilde{s})\right)\right] &\geq \mathbb{E}_A^{\tilde{\mathcal{E}}}\left[\mathbb{E}_A^{\mathcal{E}}\left[M\left(\boldsymbol{\beta}_B(s)\right) \mid \tilde{s}\right]\right] \\
&= \int_{\tilde{s}}\left(\int_s M\left(\boldsymbol{\beta}_B(s)\right) \hat{q}_A(s|\tilde{s})\, \mathrm{d}s\right) \tilde{p}_A(\tilde{s})\, \mathrm{d}\tilde{s} \\
&= \int_s M\left(\boldsymbol{\beta}_B(s)\right)\left(\int_{\tilde{s}} \frac{q(\tilde{s}|s)}{\tilde{p}_A(\tilde{s})} \tilde{p}_A(\tilde{s})\, \mathrm{d}\tilde{s}\right) p_A(s)\, \mathrm{d}s \\
&= \mathbb{E}_A^{\mathcal{E}}\left[M\left(\boldsymbol{\beta}_B(s)\right)\right].
\end{aligned}$$

Given that the priors are likelihood-ratio ordered, the above inequality would be reversed if and only if $M(\overline{\boldsymbol{\beta}}_A) \geq M(\overline{\boldsymbol{\beta}}_B)$.

Finally, any experiment $\mathcal{E}$ is a garbling of the fully informative experiment $\mathcal{F}$ that reveals the true state ω, whereas the uninformative experiment $\mathcal{U}$ that provides no information is a garbling of $\mathcal{E}$. It holds that

$$\mathbb{E}_A^{\mathcal{F}}\left[M\left(\boldsymbol{\beta}_B(s)\right)\right] = M\left(\overline{\boldsymbol{\beta}}_A\right) \quad \text{and} \quad \mathbb{E}_A^{\mathcal{U}}\left[M\left(\boldsymbol{\beta}_B(s)\right)\right] = M\left(\overline{\boldsymbol{\beta}}_B\right).$$

For $M(\overline{\boldsymbol{\beta}}_B) \geq M(\overline{\boldsymbol{\beta}}_A)$, the experiments $\mathcal{F}$ and $\mathcal{U}$ respectively provide the lower bound and the upper bound in the second part of the theorem; for $M(\overline{\boldsymbol{\beta}}_B) \leq M(\overline{\boldsymbol{\beta}}_A)$, the experiment $\mathcal{F}$ gives the upper bound and $\mathcal{U}$ gives the lower bound. □

Proof of Lemma 1. Partially differentiating (5) with respect to the first argument yields

$$\begin{aligned}
\frac{\partial \eta(\hat{s}, p, \pi)}{\partial \hat{s}} &= \frac{-p(1-p) f_\pi(\hat{s})}{(1-p+pF_\pi(\hat{s}))^2}(\pi - \mathbb{E}_\pi[s|s<\hat{s}]) + \frac{pF_\pi(\hat{s})}{1-p+pF_\pi(\hat{s})} \frac{f_\pi(\hat{s})}{F_\pi(\hat{s})}(\hat{s} - \mathbb{E}_\pi[s|s<\hat{s}]) \\
&= \frac{pf_\pi(\hat{s})}{1-p+pF_\pi(\hat{s})}\left(\frac{-(1-p)}{1-p+pF_\pi(\hat{s})}(\pi - \mathbb{E}_\pi[s|s<\hat{s}]) + (\hat{s} - \mathbb{E}_\pi[s|s<\hat{s}])\right) \\
&= \frac{pf_\pi(\hat{s})}{1-p+pF_\pi(\hat{s})}(\hat{s} - \eta(\hat{s}, p, \pi)).
\end{aligned}$$

Hence, $\operatorname{sign}\left[\partial \eta(\hat{s}, p, \pi)/\partial \hat{s}\right] = \operatorname{sign}\left[\hat{s} - \eta(\hat{s}, p, \pi)\right]$. Part 1 of the lemma follows from the observation that for any p and π, $\eta(\underline{s}, p, \pi) = \eta(\overline{s}, p, \pi) = \pi$.

Partially differentiating with respect to the second argument and simplifying yields

$$\frac{\partial \eta(\hat{s}, p, \pi)}{\partial p} = \frac{F_\pi(\hat{s})\left(\mathbb{E}_\pi[s|s<\hat{s}] - \pi\right)}{(1-p+pF_\pi(\hat{s}))^2},$$

which proves part 2 of the lemma because $\mathbb{E}_\pi[s|s < \hat{s}] < \pi$ if and only if $\hat{s} < \overline{s}$, and $F_\pi(\hat{s}) > 0$ if and only if $\hat{s} > \underline{s}$. $\square$

Proof of Proposition 1. The first part is straightforward and omitted. The second part follows from the fact that for any $p \in (0, 1)$ and π, $\eta(\cdot, p, \pi)$ is strictly decreasing on the domain $[0, \hat{s}]$, where $\hat{s}$ is the fixed point of $\eta(\cdot, p, \pi)$, which is interior (Lemma 1). The third part follows because parameters can be chosen so that $c > 0$ and there are multiple solutions in s to $s - c = \eta(s, p_i, \pi)$, as depicted in Figure 1. This can be seen by fixing all parameters except c and p_i and then considering $p_i \to 1$ with a suitable choice of c; details are available from the authors on request. $\square$

Proof of Proposition 2. Fix any $p_i > \tilde{p}_i$ and let $\hat{s}_i^0$ and $\tilde{s}_i^0$ denote the corresponding highest equilibrium disclosure thresholds. Suppose by way of contradiction that $\hat{s}_i^0 > \tilde{s}_i^0$. Since $\tilde{s}_i^0$ is the highest equilibrium threshold at $\tilde{p}_i$, it follows that for any $\hat{s} > \tilde{s}_i^0$, $\eta(\hat{s}, \tilde{p}_i, \pi) < \hat{s} - c$. But the fact that $\eta(\hat{s}, p, \pi)$ is weakly decreasing in p (Lemma 1) implies that for any $\hat{s} > \tilde{s}_i^0$, we have $\eta(\hat{s}, p_i, \pi) < \hat{s} - c$, which implies that $\hat{s}_i^0 \leq \tilde{s}_i^0$, a contradiction. A similar argument can be used to establish the result for the lowest equilibrium threshold. We omit the proof of the second part as it follows the same logic as the first part. $\square$

Proof of Lemma 2. Fix any equilibrium and any sender i and sender $j \neq i$. It suffices to show that the difference in the expected payoff for i from disclosing versus concealing is strictly increasing in s_i. Denote the expected payoff from concealing as $\mathbb{E}[\beta_{DM}(m_j, m_i = \phi)]$, where $\beta_{DM}(m_j, m_i = \phi)$ denotes the DM's equilibrium belief following any message m_j and nondisclosure by i, and the expectation is taken over m_j given i's beliefs under s_i. Because m_j is uncorrelated with s_i conditional on the state, and i's belief about the state given s_i is just s_i,

$$\mathbb{E}[\beta_{DM}(m_j, m_i = \phi)] = s_i\mathbb{E}[\beta_{DM}(m_j, m_i = \phi)|\omega = 1] + (1 - s_i)\mathbb{E}[\beta_{DM}(m_j, m_i = \phi)|\omega = 0].$$

The derivative of the right-hand-side of the above equation with respect to s_i is strictly less than one because $\mathbb{E}[\beta_{DM}(m_j, m_i = \phi)|\omega = 1] < 1$, as beliefs lie in $[0, 1]$ and m_j cannot perfectly reveal the state. Therefore, the payoff difference, $\mathbb{E}[\beta_{DM}(m_j, m_i = \phi)] - (s_i - c)$ is single-crossing in s_i. $\square$

Proof of Lemma 3. For any p_j, p_i, $\hat{s}_i$, and s_i, $U(s_i, p_i, \hat{s}_j, p_j)$ is i's expectation of the DM's belief (viewed as random variable whose realization depends on j's message) under a prior s_i for i and $\eta(s_i, p_i, \pi)$ for the DM. It follows immediately from Theorem 1 that:

1. $s_i = \eta(s_i, p_i, \pi) \implies U(s_i, p_i, \hat{s}_j, p_j) = s_i$,
2. $s_i > \eta(s_i, p_i, \pi) \implies \eta(s_i, p_i, \pi) \leq U(s_i, p_i, \hat{s}_j, p_j) \leq s_i$,
3. $s_i < \eta(s_i, p_i, \pi) \implies \eta(s_i, p_i, \pi) \geq U(s_i, p_i, \hat{s}_j, p_j) \geq s_i$.

Because $p_j < 1$, the last inequalities in items 2 and 3 above are in fact strict, as an equality in either case requires j's message to be fully informative of the state (cf. Remark 1). Finally, the other inequalities in items 2 and 3 are also strict if and only if $\hat{s}_j < \overline{s}$, as j's message is uninformative if and only if $\hat{s}_j = \overline{s}$. $\square$

Proof of Proposition 3. The transformation $T(\beta_A, \overline{\beta}_B, \overline{\beta}_A)$ defined in (2) is increasing in $\overline{\beta}_B$. Lemma 1 shows that $\eta(\hat{s}_i, p_i, \pi)$ is decreasing in p_i. Hence, sender i's payoff from concealing signal s_i given a correct conjecture by the DM,

$$U(s_i, p_i, \hat{s}_j, p_j) = \mathbb{E}_{\hat{s}_j, p_j}[T(\beta(m_j; s_i), \eta(s_i, p_i, \pi), s_i)],$$

decreases in p_i for any signal s_i, while his payoff from disclosure does not depend on p_i. Following the same argument as in the proof of Proposition 2, the largest and smallest best-response disclosure thresholds must decrease in p_i.

Let $\hat{s}_i^0$ be the smallest equilibrium threshold in the single-sender with i; recall that this is the unique equilibrium threshold if $c \leq 0$.

Consider first $c = 0$. It follows from Part 1 of Lemma 3 that sender i is indifferent between nondisclosure and disclosure when his signal is $\hat{s}_i^0$; hence, $\hat{s}_i^0 \in \hat{s}_i^{BR}(\hat{s}_j; p_i, p_j)$. Next, we claim there exists no other best-response disclosure threshold. Suppose, to contradiction, that $\hat{s}' > \hat{s}_i^0$ and $\hat{s}' \in \hat{s}_i^{BR}(\hat{s}_j; p_i, p_j)$. By Lemma 1, $\eta(\hat{s}', p_i, \pi) < \hat{s}'$. Lemma 3 then implies that $U(\hat{s}', p_i, \hat{s}_j, p_j) < \hat{s}'$. Therefore, sender i strictly prefers disclosure to nondisclosure when his signal is $\hat{s}'$, contradicting $\hat{s}' \in \hat{s}_i^{BR}(\hat{s}_j; p_i, p_j)$. A similar argument establishes that $\hat{s}' < \hat{s}_i^0$ implies $\hat{s}' \notin \hat{s}_i^{BR}(\hat{s}_j; p_i, p_j)$.

Now consider $c < 0$. If $s_i > \hat{s}_i^0$, then $s_i - c > \eta(s_i, p_i, \pi)$. Since Lemma 3 implies that $U(s_i, p_i, \hat{s}_j, p_j) \leq \max\{s_i, \eta(s_i, p_i, \pi)\}$, it follows that $U(s_i, p_i, \hat{s}_j, p_j) < s_i - c$, and hence $s_i \notin \hat{s}_i^{BR}(\hat{s}_j; p_i, p_j)$. Furthermore, if $\hat{s}_j < \overline{s}$, then $\hat{s}_i^0 < \eta(\hat{s}_i^0, p_i, \pi)$ and Lemma 3 together imply $U(\hat{s}_i^0, p_i, \hat{s}_j, p_j) > \eta(\hat{s}_i^0, p_i, \pi)$, and hence $\hat{s}_i^0 \in \hat{s}_i^{BR}(\hat{s}_j; p_i, p_j)$ if and only if $\hat{s}_i^0 = \underline{s}$. Conversely, it is obvious that $\hat{s}_i^0 = \hat{s}_i^{BR}(\hat{s}_j; p_i, p_j)$ if $\hat{s}_j = \overline{s}$, because sender j's message is uninformative. This proves part (i) of the result for $c < 0$. To prove part (ii), we first note from part (i) that if $s_i \in \hat{s}_i^{BR}(\hat{s}_j; p_i, p_j)$, then $s_i \leq \hat{s}_i^0$ and hence $\eta(s_i, p_i, \pi) > s_i$. Theorem 1 then implies that any garbling of sender j's message decreases $U(s_i, p_i, \hat{s}_j, p_j)$. Thus, an increase in $\hat{s}_j$ or a decrease in p_j—both of which represent garblings—lowers sender i's nondisclosure payoff without affecting his disclosure payoff at signal s_i. Following the same argument as in the proof of Proposition 2, the largest and smallest best-response disclosure thresholds must increase.

We omit the proof for $c > 0$ as it follows a symmetric argument to that for $c < 0$; the only point to note is that here the definition of $\hat{s}_i^0$ as the smallest equilibrium threshold in the single-sender game is used to ensure that $s_i < \hat{s}_i^0$ implies $s_i - c < \eta(s_i, p_i, \pi)$. □

Proof of Proposition 4. Consider first the case $c \leq 0$. For each sender i, define the function

$$w_i(\hat{s}_i, \hat{s}_j; p_i, p_j) := \inf\{\hat{s}_i \mid U(\hat{s}_i, p_i, \hat{s}_j, p_j) \leq \hat{s}_i - c\}.$$

That is, $w_i(\cdot)$ gives the smallest element of $\hat{s}_i^{BR}(\hat{s}_j; p_i, p_j)$. Let $w = (w_i, w_j)$, and define

$$\mathrm{s}_*(p_i, p_j) := \inf\{(\hat{s}_i, \hat{s}_j) \mid w(\hat{s}_i, \hat{s}_j; p_i, p_j) \leq (\hat{s}_i, \hat{s}_j)\}.$$

By Proposition 3, $w(\cdot; p_i, p_j)$ is monotone increasing. Hence, $\mathrm{s}_*(p_i, p_j)$ is its smallest fixed point. It remains to be shown that $\mathrm{s}_*(p_i, p_j)$ is the smallest fixed point of the best-response correspondence $(\hat{s}_i^{BR}, \hat{s}_j^{BR})$. Let s be any other fixed point of the correspondence. Since w is monotone, we have $w(\mathrm{s}_* \wedge \mathrm{s}) \leq w(\mathrm{s}) \leq \mathrm{s}$ and $w(\mathrm{s}_* \wedge \mathrm{s}) \leq w(\mathrm{s}_*) \leq \mathrm{s}_*$. (We follow the notation that $\mathrm{s} \wedge \mathrm{s}' \equiv (\min\{s_1, s_1'\}, \min\{s_2, s_2'\})$.) Thus, $w(\mathrm{s}_* \wedge \mathrm{s}) \leq \mathrm{s}_* \wedge \mathrm{s}$. By the definition of s_*, this in turn implies that $\mathrm{s}_* \leq \mathrm{s}_* \wedge \mathrm{s}$, which is possible only if $\mathrm{s}_* \leq \mathrm{s}$, as required. Thus, this argument establishes that the smallest fixed point of the minimal best response is also the smallest fixed point among all best responses. In other words, $\mathrm{s}_*(p_i, p_j)$ is the *smallest* equilibrium. Proposition 3 establishes that w is decreasing in p_i. It follows from standard monotone comparative statics that the smallest fixed point of w decreases in p_i. It is also straightforward to see from the definition of w_i that w is increasing in c. Hence the best equilibrium increases in c. A parallel argument shows that the worst equilibrium also decreases in p_i and increases in c.

For the case $c > 0$, we keep the sign of $\hat{s}_i$ but flip the sign of $\hat{s}_j$ in the definition of w so that it is monotone in $(\hat{s}_i, -\hat{s}_j)$. The smallest fixed point of w then corresponds to the j-maximal equilibrium. By Proposition 3, a higher p_i decreases sender i's best response but increases sender j's best response. Hence, in a j-maximal equilibrium, $(\hat{s}_i, -\hat{s}_j)$ is decreasing in p_i. The same conclusion holds for an i-maximal equilibrium. □

Proof of Lemma 4. The sender's marginal rate of substitution between m and π is given by

$$\frac{\partial c(m,t)/\partial m}{\partial \mathbb{E}_{s|t}[\beta(s;\pi)]/\partial \pi} = \pi(1-\pi)\frac{\partial c(m,t)/\partial m}{\mathbb{E}_{s|t}[\beta(s;\pi)(1-\beta(s;\pi))]}.$$

For $t < m$, the marginal rate of substitution is strictly decreasing in sender's type t if

$$\frac{\partial^2 c(m,t)/\partial m \partial t}{\partial c(m,t)/\partial m} < \frac{\mathbb{E}_{s|1}[\beta(s;\pi)(1-\beta(s;\pi))] - \mathbb{E}_{s|0}[\beta(s;\pi)(1-\beta(s;\pi))]}{t\mathbb{E}_{s|1}[\beta(s;\pi)(1-\beta(s;\pi))] + (1-t)\mathbb{E}_{s|0}[\beta(s;\pi)(1-\beta(s;\pi))]}.$$

Since $\mathbb{E}_{s|1}[\beta(s;\pi)(1-\beta(s;\pi))] > 0$, the right-hand side of the above inequality is strictly greater than

$$\frac{0 - \mathbb{E}_{s|0}[\beta(s;\pi)(1-\beta(s;\pi))]}{0 + (1-t)\mathbb{E}_{s|0}[\beta(s;\pi)(1-\beta(s;\pi))]} = -\frac{1}{1-t}.$$

So, if condition (7) holds, the marginal rate of substitution strictly decreases in t for $t < m$. □

Proof of Lemma 5. For brevity, we denote $\alpha(\pi, t) := \mathbb{E}_{s|t}[\beta(s;\pi)]$.

We first establish uniqueness of the cutoff $\underline{t}$. Suppose, by way of contradiction, that $\underline{t} < \underline{t}'$ are both equilibrium cutoffs. From (9) and (10),

$$\alpha(\pi^P(\underline{t}), \underline{t}) - \alpha(\underline{t}, \underline{t}) \geq c(1, \underline{t}) - c(\mu(\underline{t}), \underline{t}).$$

Since $\pi^P(\underline{t}') > \pi^P(\underline{t})$, we obtain

$$\alpha(\pi^P(\underline{t}'), \underline{t}) - \alpha(\underline{t}, \underline{t}) > c(1, \underline{t}) - c(\mu(\underline{t}), \underline{t}).$$

This implies that type $\underline{t}$ would deviate to $m = 1$ from the $\underline{t}'$ equilibrium, a contradiction.

Next, we show that no type has incentive to deviate from $\mu(\cdot)$.

Case (i). $t \leq \underline{t}$. For any $t' \in (t, \underline{t}]$, Equation 8 and Lemma 4 imply

$$\mu'(t') = \frac{\partial \alpha(t', t')/\partial \pi}{\partial c(\mu(t'), t')/\partial m} > \frac{\partial \alpha(t', t)/\partial \pi}{\partial c(\mu(t'), t)/\partial m}.$$

The inequality is due to Lemma 4, which applies because $\mu(t') > t' > t$. The above inequality can be written as

$$\frac{\partial \alpha(t', t)}{\partial \pi} - \frac{\partial c(\mu(t'), t)}{\partial m}\mu'(t') < 0.$$

This inequality is true for any $t' \in (t, \underline{t}]$. For any $\hat{t}$ in the same interval, integrating the inequality over t' from t to $\hat{t}$ gives

$$\alpha(\hat{t}, t) - c(\mu(\hat{t}), t) < \alpha(t, t) - c(\mu(t), t).$$

Thus, type t has no incentive to deviate upward to mimic type $\hat{t} \in (t, \underline{t}]$.

Further, by (9) and Lemma 4, for any $t < \underline{t}$,

$$\alpha(\pi^P(\underline{t}), \underline{t}) - c(1, \underline{t}) = \alpha(\underline{t}, \underline{t}) - c(\mu(\underline{t}), \underline{t}) \implies \alpha(\pi^P(\underline{t}), t) - c(1, t) < \alpha(\underline{t}, t) - c(\mu(\underline{t}), t).$$

Since we have already shown that type $t < \underline{t}$ has no incentive to mimic type $\underline{t}$, type t has no incentive to mimic types higher than $\underline{t}$ by deviating to $m = 1$ either.

Now, let $t' < t$ be such that $t' \geq \mu^{-1}(t)$. Then, an analogous argument establishes that

$$\frac{\partial \alpha(t', t)}{\partial \pi} - \frac{\partial c(\mu(t'), t)}{\partial m} \mu'(t') > 0$$

for any $t' \in [\mu^{-1}(t), t)$. We can apply Lemma 4 for $t' > \mu^{-1}(t)$ because $\mu(t') > t$. The above also holds for $t' = \mu^{-1}(t)$ because $\frac{\partial c(t,t)}{\partial m} = 0$. For any $\hat{t}$ in the same interval, integrating the inequality over t' from $\hat{t}$ to t gives

$$\alpha(t, t) - c(\mu(t), t) > \alpha(\hat{t}, t) - c(\mu(\hat{t}), t).$$

Thus, type $t \leq \underline{t}$ has no incentive to deviate downward to mimic type $\hat{t} \in [\mu^{-1}(t), t)$.

For any $\hat{t} < \mu^{-1}(t)$, we have $\alpha(\hat{t}, t) < \alpha(\mu^{-1}(t), t)$ and $c(\mu(\hat{t}), t) > c(\mu(\mu^{-1}(t)), t) = 0$. Since type t has no incentive to mimic type $\mu^{-1}(t)$, type t has no incentive to mimic type $\hat{t} < \mu^{-1}(t)$ either.

Case (ii). $t \in (\underline{t}, \mu(\underline{t})]$. By (9) and Lemma 4, for any $t > \underline{t}$,

$$\alpha(\pi^P(\underline{t}), \underline{t}) - c(1, \underline{t}) = \alpha(\underline{t}, \underline{t}) - c(\mu(\underline{t}), \underline{t}) \implies \alpha(\pi^P(\underline{t}), t) - c(1, t) > \alpha(\underline{t}, t) - c(\mu(\underline{t}), t).$$

Thus, type t strictly prefers $m = 1$ to $m = \mu(\underline{t})$.

Take any $\hat{t} \in [\mu^{-1}(t), \underline{t})$. In case (i), we have shown that type $\underline{t}$ has no incentive to deviate downward to mimic $\hat{t}$:

$$\alpha(\underline{t}, \underline{t}) - c(\mu(\underline{t}), \underline{t}) > \alpha(\hat{t}, \underline{t}) - c(\mu(\hat{t}), \underline{t}).$$

By condition (7), this inequality implies

$$\alpha(\underline{t}, t) - c(\mu(\underline{t}), t) > \alpha(\hat{t}, t) - c(\mu(\hat{t}), t)$$

for $t \in (\underline{t}, \mu(\hat{t})]$. Since type t has no incentive to deviate to mimic type $\underline{t}$, type t has no incentive to mimic type $\hat{t} \in [\mu^{-1}(t), \underline{t})$ either.

The same argument as in case (i) shows that type $t \in (\underline{t}, \mu(\underline{t})]$ has no incentive to deviate to any type lower than $\mu^{-1}(t)$.

Case (iii). $t > \mu(\underline{t})$. Let $\hat{\pi} > \underline{t}$ be the interim belief such that type $\underline{t}$ is indifferent between $m = t$ and $m = 1$:

$$\alpha(\pi^P(\underline{t}), \underline{t}) - c(1, \underline{t}) = \alpha(\hat{\pi}, \underline{t}) - c(t, \underline{t}).$$

Since $t > \mu(\underline{t}) > \underline{t}$,

$$\alpha(\pi^P(\underline{t}), t) - c(1, t) > \alpha(\pi^P(\underline{t}), \underline{t}) - c(1, \underline{t})$$

By condition (7), this implies

$$\alpha(\pi^P(\underline{t}), t) - c(1, t) > \alpha(\hat{\pi}, t) - c(t, t) > \alpha(\underline{t}, t) - c(\mu(\underline{t}), t).$$

Thus, type t strictly prefers $m = 1$ to $m = \mu(\underline{t})$. The same argument as in case (i) shows that type t has no incentive to deviate to types lower than $\underline{t}$ either.

Finally, it remains to be shown that no type has an incentive to deviate to some $m \in (\mu(\underline{t}), 1)$ which is not used in equilibrium. We assign the off-equilibrium belief $\pi(m) = \underline{t}$ for such m. In cases (i) and (ii), we show that type $t \leq \mu(\underline{t})$ has no incentive to deviate to mimic $\underline{t}$. Since $c(\mu(\underline{t}), t) < c(m, t)$ for off-equilibrium m, this type has no incentive to deviate to m either. In case (iii), we show that type $t > \mu(\underline{t})$ prefers belief-message pair $(\pi^P(\underline{t}), 1)$ to $(\hat{\pi}, t)$, where $\hat{\pi} > \underline{t}$. This type must prefer $(\hat{\pi}, t)$ to $(\underline{t}, m)$ because $c(t, t) \leq c(m, t)$. □

Proof of Proposition 6. We first show that the function $\mu(t)$, defined by Equation 8, is pointwise (weakly) decreasing in the informativeness of the experiment.

If we show that $\mu(t)$ decreases pointwise when the right-hand-side of Equation 8 decreases for all t, the result then follows from Lemma 6. Accordingly, let $\tilde{\mu}(t)$ and $\mu(t)$ be two solutions to Equation 8, with $\tilde{\mu}(0) = \mu(0) = 0$, where $\tilde{\mu}$ solves Equation 8 with a pointwise lower right-hand-side. (These are defined over some respective domains $[0, \tilde{\bar{t}}]$ and $[0, \bar{t}]$. The argument establishes that $\tilde{\bar{t}} \geq \bar{t}$.) For any $t > 0$, if $\tilde{\mu}(t) = \mu(t)$ then $\mu'(t) \geq \tilde{\mu}'(t) > 0$. This implies that at any touching point, μ must touch $\tilde{\mu}$ from below. Consequently, by continuity,

$$\mu(t') \geq \tilde{\mu}(t') \text{ for } t' > 0 \implies \mu(t) \geq \tilde{\mu}(t) \text{ for all } t \geq t'.$$

Now suppose, by way of contradiction, that for some $\hat{t} > 0$, $\tilde{\mu}(\hat{t}) > \mu(\hat{t})$. Then, it must be the case that $\tilde{\mu}(t) > \mu(t)$ for all $t \in (0, \hat{t})$. Since $\partial^2 c/\partial m^2 > 0$ and $\tilde{\mu}$ corresponds to a lower right-hand-side of Equation 8, it follows from Equation 8 that $\tilde{\mu}'(t) < \mu'(t)$ for all $t \in (0, \hat{t})$.

But then

$$\tilde{\mu}(\hat{t}) - \mu(\hat{t}) = \int_0^{\hat{t}} [\tilde{\mu}'(t) - \mu'(t)] \, \mathrm{d}t < 0,$$

a contradiction.

Next, we prove that $\underline{t}_{\tilde{s}} \geq \underline{t}_s$. The result is trivial if $\underline{t}_s = 0$, so assume $\underline{t}_s > 0$. Under the more informative experiment $\tilde{s}$ and the corresponding function $\tilde{\mu}$, type $\underline{t}_s$ will (weakly) prefer to separate than pool with the top. There are two reasons, both working in the same direction: (i) separation cost is lower with $\tilde{\mu}$ than with μ (as shown above), i.e., the right-hand-side of the consequent in (9) increases; and (ii) the benefit of pooling is lower (part 2 of Lemma 6), i.e., the left-hand-side of the consequent in (9) decreases. Since $\bar{t}_{\tilde{s}} \geq \bar{t}_s \geq \underline{t}_s$, and $\bar{t}_{\tilde{s}}$ strictly prefers to pool with $[\bar{t}_{\tilde{s}}, 1]$ rather than separate (same message cost, better inference), by continuity there will be some cutoff $\underline{t}_{\tilde{s}} \geq \underline{t}_s$, and this cutoff is unique by Lemma 5.

The second statement of the proposition follows because all types are sending lower signals in the new equilibrium (and each type t's signal is larger than t in both equilibria). □

B. Discussion of Theorem 2

This appendix establishes that the conclusion of Theorem 2—viz., that agent i expects agent j's posterior to be closer to i's prior, in the sense of their $M(\cdot)$ statistics, under a more informative experiment—can fail with a non-MLRP experiment (Example 1) or if the priors are ordered by first-order stochastic dominance (FOSD) rather than by likelihood ratio (Example 2). In both examples, $M(\cdot)$ denotes the simple expectation operator, i.e., we take $h(\omega) = \omega$.

Example 1. Let $\Theta = \{0, 1, 2\}$. Consider the following experiment $\mathcal{E}$ with a binary signal space $\{s_l, s_h\}$:

$$\begin{bmatrix} \Pr(s_l|\theta = 0) & \Pr(s_l|\theta = 1) & \Pr(s_l|\theta = 2) \\ \Pr(s_h|\theta = 0) & \Pr(s_h|\theta = 1) & \Pr(s_h|\theta = 2) \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}.$$

This experiment does not satisfy MLRP. Consider the priors $\overline{\boldsymbol{\beta}}_{\boldsymbol{B}} = (0, 0.2, 0.8)$ and $\overline{\boldsymbol{\beta}}_{\boldsymbol{A}} = (0.9, 0.1, 0)$; the former likelihood-ratio dominates the latter. A direct calculation gives

$$M(\overline{\boldsymbol{\beta}}_{\boldsymbol{A}}) = 0.1 < M(\overline{\boldsymbol{\beta}}_{\boldsymbol{B}}) = 1.8 < \mathbb{E}_A^{\mathcal{E}}[M(\boldsymbol{\beta}_{\boldsymbol{B}}(\cdot))] = 1.9,$$

contrary to Theorem 2's conclusion that $\mathbb{E}_A^{\mathcal{E}}[M(\boldsymbol{\beta}_{\boldsymbol{B}}(\cdot))] \in [M(\overline{\boldsymbol{\beta}}_{\boldsymbol{A}}), M(\overline{\boldsymbol{\beta}}_{\boldsymbol{B}})]$ when $M(\overline{\boldsymbol{\beta}}_{\boldsymbol{A}}) < M(\overline{\boldsymbol{\beta}}_{\boldsymbol{B}})$. □

The priors in Example 1 violate our maintained assumption of full support, but the example would go through (with more elaborate calculations) if we replaced $\overline{\beta}_B(0)$ and $\overline{\beta}_A(2)$ with a small $\varepsilon > 0$ instead of 0. This example is similar to Alonso and Câmara (2016, pp. 674–675), who use their example to illustrate how a "skeptic" can design information to persuade a "believer."

Example 2. Let $\Theta = \{0, 1, 8\}$. Consider the following MLRP experiment $\mathcal{E}$ with signal space $\{s_l, s_m, s_h\}$:

$$\begin{bmatrix} \Pr(s_l|\theta=0) & \Pr(s_l|\theta=1) & \Pr(s_l|\theta=8) \\ \Pr(s_m|\theta=0) & \Pr(s_m|\theta=1) & \Pr(s_m|\theta=8) \\ \Pr(s_h|\theta=0) & \Pr(s_h|\theta=1) & \Pr(s_h|\theta=8) \end{bmatrix} = \begin{bmatrix} \frac{1}{2} & 0 & 0 \\ \frac{1}{4} & \frac{1}{2} & 0 \\ \frac{1}{4} & \frac{1}{2} & 1 \end{bmatrix}.$$

Consider the priors $\overline{\boldsymbol{\beta}}_{\boldsymbol{B}} = (2/3, 0, 1/3)$ and $\overline{\boldsymbol{\beta}}_{\boldsymbol{A}} = (0, 2/3, 1/3)$. The prior $\overline{\boldsymbol{\beta}}_{\boldsymbol{A}}$ dominates $\overline{\boldsymbol{\beta}}_{\boldsymbol{B}}$ in the sense of FOSD, but not in likelihood ratio. Direct calculations establish the following:

1. The posteriors afters signal s_h, $\overline{\boldsymbol{\beta}}_{\boldsymbol{B}}(\cdot|s_h)$ and $\overline{\boldsymbol{\beta}}_{\boldsymbol{A}}(\cdot|s_h)$, are not ordered by FOSD, because $\beta_A(0|s_h) = 0 < \beta_B(0|s_h) = 1/3$ while $\sum_{\theta\in\{0,1\}} \beta_A(\theta|s_h) = 1/2 > \sum_{\theta\in\{0,1\}} \beta_B(\theta|s_h) = 1/3$.
2. $M(\overline{\boldsymbol{\beta}}_{\boldsymbol{B}}) = 8/3 < M(\overline{\boldsymbol{\beta}}_{\boldsymbol{A}}) = 10/3 < \mathbb{E}^{\mathcal{E}}_A[M(\boldsymbol{\beta}_{\boldsymbol{B}}(\cdot))] = 32/9$.

The first point shows that posteriors may not be ranked by FOSD even when priors are and the experiment satisfies MLRP. The second point is contrary to Theorem 2's conclusion that $\mathbb{E}^{\mathcal{E}}_A[M(\boldsymbol{\beta}_{\boldsymbol{B}}(\cdot))] \in [M(\overline{\boldsymbol{\beta}}_{\boldsymbol{B}}), M(\overline{\boldsymbol{\beta}}_{\boldsymbol{A}})]$ when $M(\overline{\boldsymbol{\beta}}_{\boldsymbol{B}}) < M(\overline{\boldsymbol{\beta}}_{\boldsymbol{A}})$. $\square$

The priors in Example 2 violate our maintained assumption of full support, but the example would go through (with more elaborate calculations) if we replaced $\overline{\beta}_B(1)$ and $\overline{\beta}_A(0)$ with a small $\varepsilon > 0$ instead of 0. This example implies that even when priors are first-order-stochastically ordered, a sender with a linear—or, by continuity, even strictly convex—utility function over the receiver's posterior mean could prefer persuading a "skeptic" via a partially-revealing experiment than a fully-revealing one. Theorem 2 implies that this cannot arise under likelihood-ratio ordered priors, as also noted by Alonso and Câmara (2016, Proposition 6.1).

References

ACHARYA, V. V., P. DEMARZO, AND I. KREMER (2011): "Endogenous Information Flows and the Clustering of Announcements," *American Economic Review*, 101, 2955–79.

ALONSO, R. AND O. CÂMARA (2016): "Bayesian Persuasion with Heterogeneous Priors," *Journal of Economic Theory*, 165, 672–706.

AUMANN, R. J. (1976): "Agreeing to Disagree," *Annals of Statistics*, 4, 1236–1239.

BALIGA, S., E. HANANY, AND P. KLIBANOFF (2013): "Polarization and Ambiguity," *American Economic Review*, 103, 3071–3083.

BHATTACHARYA, S. AND A. MUKHERJEE (2013): "Strategic Information Revelation when Experts Compete to Influence," *RAND Journal of Economics*, 522–544.

BLACKWELL, D. (1951): "Comparison of Experiments," in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, University of California Press, vol. 1, 93–102.

——— (1953): "Equivalent Comparisons of Experiments," *Annals of Mathematical Statistics*, 24, 265–272.

BOURJADE, S. AND B. JULLIEN (2011): "The Roles of Reputation and Transparency on the Behavior of Biased Experts," *RAND Journal of Economics*, 42, 575–594.

CARLIN, B. I., S. W. DAVIES, AND A. IANNACCONE (2012): "Competition, Comparative Performance, and Market Transparency," *American Economic Journal: Microeconomics*, 4, 202–37.

CHE, Y.-K. AND N. KARTIK (2009): "Opinions as Incentives," *Journal of Political Economy*, 117, 815–860.

CHO, I.-K. AND J. SOBEL (1990): "Strategic Stability and Uniqueness in Signaling Games," *Journal of Economic Theory*, 50, 381–413.

DALEY, B. AND B. GREEN (2014): "Market Signaling with Grades," *Journal of Economic Theory*, 151, 114–145.

DAUGHETY, A. AND J. REINGANUM (2018): "Evidence Suppression by Prosecutors: Violations of the *Brady* Rule," *Journal of Law, Economics, and Organization*, 34, 475–510.

DEMARZO, P., I. KREMER, AND A. SKRZYPACZ (2018): "Test Design and Disclosure," Forthcoming in the *American Economic Review*.

DEWATRIPONT, M. AND J. TIROLE (1999): "Advocates," *Journal of Political Economy*, 107, 1–39.

DRANOVE, D. AND G. Z. JIN (2010): "Quality Disclosure and Certification: Theory and Practice," *Journal of Economic Literature*, 48, 935–63.

DYE, R. A. (1985): "Disclosure of Nonproprietary Information," *Journal of Accounting research*, 23, 123–145.

——— (2017): "Optimal Disclosure When There are Penalties for Nondisclosure," *RAND Journal of Economics*, 48, 704–732.

ECHENIQUE, F. (2002): "Comparative Statics by Adaptive Dynamics and the Correspondence Principle," *Econometrica*, 70, 833–844.

ELLIOTT, M., B. GOLUB, AND A. KIRILENKO (2014): "How Better Information Can Garble Experts' Advice," *American Economic Review*, 104, 463–468.

FELTOVICH, N., R. HARBAUGH, AND T. TO (2002): "Too Cool for School? Signaling and Countersignaling," *RAND Journal of Economics*, 33, 630–649.

FRANK, R. H. (1985): "The Demand for Unobservable and Other Nonpositional Goods," *American Economic Review*, 75, 101–116.

GALPERTI, S. (2018): "Persuasion: The Art of Changing Worldviews," Forthcoming in the *American Economic Review*.

GENTZKOW, M. AND E. KAMENICA (2014): "Costly Persuasion," *American Economic Review*, 104, 457–462.

——— (2017): "Competition in Persuasion," *Review of Economic Studies*, 84, 300–322.

GROSSMAN, S. J. AND O. D. HART (1980): "Disclosure Laws and Takeover Bids," *Journal of Finance*, 35, 323–334.

HIRSCH, A. V. (2016): "Experimentation and Persuasion in Political Organizations," *American Political Science Review*, 110, 68–84.

JACKSON, M. O. AND X. TAN (2013): "Deliberation, Disclosure of Information, and Voting," *Journal of Economic Theory*, 148, 2–30.

JOVANOVIC, B. (1982): "Truthful Disclosure of Information," *Bell Journal of Economics*, 13, 36–44.

JUNG, W.-O. AND Y. K. KWON (1988): "Disclosure When the Market is Unsure of Information Endowment of Managers," *Journal of Accounting Research*, 26, 146–153.

KAMENICA, E. AND M. GENTZKOW (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

KARTIK, N. (2009): "Strategic Communication with Lying Costs," *Review of Economic Studies*, 76, 1359–1395.

KARTIK, N., F. X. LEE, AND W. SUEN (2017): "Investment in Concealable Information by Biased Experts," *RAND Journal of Economics*, 48, 24–43.

KRISHNA, V. AND J. MORGAN (2001): "A Model of Expertise," *Quarterly Journal of Economics*, 116, 747–775.

MILGROM, P. (2008): "What the Seller Won't Tell You: Persuasion and Disclosure in Markets," *Journal of Economic Perspectives*, 22, 115–131.

MILGROM, P. AND J. ROBERTS (1986): "Relying on the Information of Interested Parties," *RAND Journal of Economics*, 17, 18–32.

MILGROM, P. R. (1981): "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 12, 380–391.

ROTHSCHILD, M. AND J. STIGLITZ (1970): "Increasing Risk: I. A Definition," *Journal of Economic Theory*, 2, 225–243.

SETHI, R. AND M. YILDIZ (2012): "Public Disagreement," *American Economic Journal: Microeconomics*, 4, 57–95.

SHIN, H. S. (1994a): "The Burden of Proof in a Game of Persuasion," *Journal of Economic Theory*, 64, 253–264.

——— (1994b): "News Management and the Value of Firms," *RAND Journal of Economics*, 25, 58–71.

SPENCE, M. (1973): "Job Market Signaling," *Quarterly Journal of Economics*, 87, 355–374.

THE NEW YORK TIMES (2015): "Honda Fined for Violations of Safety Law," http://nyti.ms/1yFNVUa.

TRUYTS, T. (2015): "Noisy Signaling with Exogenous Information," Unpublished.

VAN DEN STEEN, E. (2010): "Culture Clash: The Costs and Benefits of Homogeneity," *Management Science*, 56, 1718–1738.

VERRECCHIA, R. E. (1983): "Discretionary Disclosure," *Journal of Accounting and Economics*, 5, 179–194.

WEISS, A. (1983): "A Sorting-cum-Learning Model of Education," *Journal of Political Economy*, 91, 420–442.

YILDIZ, M. (2004): "Waiting to Persuade," *Quarterly Journal of Economics*, 119, 223–248.

ZANARDO, E. (2017): "How to Measure Disagreement?" Unpublished.

C. Supplementary Appendix (Not for Publication)

The first two parts of this supplementary appendix provide additional results and discussion for the disclosure application in Section 3. The final part deals with the role of linearity and MLRP in the signaling application of Section 4.

C.1. Perfectly correlated signals in the disclosure application

Consider the extreme case where informed senders' signals are perfectly correlated and for simplicity, $c = 0$. In other words, there is a single signal s drawn from a distribution $F(s|\omega)$, and each sender i is independently either informed of s with probability p_i or remains uninformed. This setting is effectively identical to the "extreme agenda" case of Bhattacharya and Mukherjee (2013).[21] If both senders are biased in the same direction then this model can be mapped to a single-sender problem where the sender is informed with probability $p_i+p_j-p_ip_j$, which is larger than $\max\{p_i, p_j\}$.[22] Proposition 2 then implies that each sender discloses more when there is an additional sender; hence, the DM is always better off with two senders than one.

It is instructive to consider why the irrelevance result no longer holds. For simplicity, suppose both senders are upward biased and symmetric ($p_i = p_j = p$). Let $\hat{s}^0$ denote the common single-sender threshold, so that the nondisclosure belief satisfies $\eta(\hat{s}^0, p, \pi) = \hat{s}^0$. The essential observation is that when sender j is added to the picture, say with the hypothesis that he too discloses all signals weakly above $\hat{s}^0$, type $\hat{s}^0$ of sender i no longer expects the DM's belief to be $\hat{s}^0$ should he conceal his signal, in contrast to the case of conditionally independent signals. Rather, he expects the DM's belief to be strictly lower: if j is informed the DM's belief will be $\hat{s}^0$, and if j is uninformed the DM's belief will be strictly lower because of nondisclosure from two senders rather than just one. This makes type $\hat{s}^0$ of sender i strictly prefer disclosure. From the perspective of Theorem 1, the point is that under conditionally correlated signals, when an informed sender i does not disclose his signal, i and the DM do not agree on the experiment generated by j's message; thus, even if the DM's nondisclosure belief agrees with i's belief (over the state), i's expectation of the DM's posterior belief can be different.

Interestingly, welfare conclusions under perfectly correlated signals are very different when the senders are opposite biased. For simplicity continue to consider $c = 0$. The following proposition shows that when senders are opposite biased, each sender discloses strictly less than he does in his single-sender game.

Proposition 7. *Assume perfectly correlated signals, $c = 0$, and that the two senders are opposite biased. Then, each sender discloses strictly less than what he does in his single-sender game.*

[21] Note that they allow for the senders' utility functions to be non-linear; the ensuing discussion does not depend on linearity either because our single-sender analysis does not require linearity.

[22] Perfect correlation implies that there is only one relevant nondisclosure belief, viz., when both senders don't disclose. So senders who are biased in the same direction must use the same equilibrium disclosure threshold. Given any such threshold, the nondisclosure belief is then computed just as in (5) (assuming the bias is upward), but with $1 - p$ replaced by the probability that both senders are uninformed, i.e., $(1 - p_i)(1 - p_j)$.

Proof of Proposition 7. We prove it for the upward biased sender; the argument is symmetric for the other sender. Let sender 1 be upward biased and sender 2 be downward biased. Let $\hat{s}_1^0$ denote the single-sender threshold, i.e., $\eta(\hat{s}_1^0) = \hat{s}_1^0$. Write $\eta(\hat{s}_1, \hat{s}_2)$ as the nondisclosure belief (in the event both senders do not disclose) in the two-sender game when the respective thresholds are $\hat{s}_1$ and $\hat{s}_2$. Even though the DM's updating is not separable as in our baseline model, it is clear that $\eta(\hat{s}_1, \hat{s}_2) \geq \eta(\hat{s}_1)$, with equality if and only if $\hat{s}_2 = \underline{s}$. This follows from the simple observation that the nondisclosure event can be viewed as the union of two events: (i) $m_1 = m_2 = \phi$ and $s_2 = \phi$; and (ii) $m_1 = m_2 = \phi$ and $s_2 > \hat{s}_2$. Conditional on the first event, the DM's posterior is $\eta(\hat{s}_1)$, whereas conditional on the second event, the posterior is larger than $\eta(\hat{s}_1)$ (strictly if and only if $\hat{s}_2 > \underline{s}$). It follows that for any $\hat{s}_2 > \underline{s}$, because $\eta(\hat{s}_1) \geq \hat{s}_1$ for all $\hat{s}_1 \leq s_1^0$, if the DM conjectures thresholds $(\hat{s}_1, \hat{s}_2)$, sender 1 with signal $s_1 = \hat{s}_1$ will strictly prefer nondisclosure to disclosure, given that if sender 2 discloses, he necessarily discloses s_1. Therefore, since sender 2 will use a threshold strictly larger than $\underline{s}$ in any equilibrium, any equilibrium involves sender 1's threshold being strictly larger than $\hat{s}_1^0$. ☐

Thus, despite the increased availability of information, the overall disclosure of information in the two-sender setting is not more informative than under either single-sender problem. Consequently, for either sender i, there exist preferences of the DM such that she would strictly prefer to face sender i alone rather than the two senders simultaneously. An implication is that the welfare conclusion in Corollary 2 of Bhattacharya and Mukherjee (2013) can be reversed under alternative DM preferences.

In general, for an arbitrary c, an interior equilibrium $(\hat{s}_1, \hat{s}_2)$ requires

$$\Pr[m_2 \neq \phi | s_1 = \hat{s}_1, \hat{s}_2]\hat{s}_1 + \Pr[m_2 = \phi | s_1 = \hat{s}_1, \hat{s}_2]\eta(\hat{s}_1, \hat{s}_2) = \hat{s}_1 - c,$$

or

$$(\hat{s}_1 - \eta(\hat{s}_1, \hat{s}_2)) = \frac{c}{\Pr[m_2 = \phi | s_1 = \hat{s}_1, \hat{s}_2]}.$$

Thus, when $c \geq 0$, using $\eta(\hat{s}_1, \hat{s}_2) > \eta(\hat{s}_1)$ (strictness by interiority of $\hat{s}_2$), it follows that $\hat{s}_1 > \eta(\hat{s}_1)$. Furthermore, for any $c \geq 0$, there is an equilibrium in which $\hat{s}_1$ is weakly larger than the largest single-sender equilibrium. When $c < 0$, the comparison is ambiguous.

C.2. Non-linear utility functions in the disclosure application

Another assumption that is important in applying Theorem 1 is that each sender has linear preferences. Suppose, more generally, that a sender i's utility is given by some function, $V_i(\beta_{DM})$. The comparative statics of sender i's disclosure depends on the comparative statics of

$$\mathbb{E}[V_i(T(\beta_i, \overline{\beta}_{DM}, \overline{\beta}_i))] - \mathbb{E}[V_i(\beta_i)] \tag{13}$$

across Blackwell-comparable experiments, where the expectation is taken over the posterior β_i using i's beliefs and $T(\cdot)$ is the transformation of (2). When $V_i(\cdot)$ is linear, one can ignore the second term in expression (13) as it is constant across experiments and Theorem 1 tells us that the sign of the change in the first term is determined by the sign of $\overline{\beta}_i - \overline{\beta}_{DM}$. Unambiguous comparative statics of expression (13) cannot be obtained for arbitrary specifications of $V_i(\cdot)$. However, because $T(\beta_i, x, x) = \beta_i$ for any β_i and x, the logic behind our irrelevance result extends generally. In particular:

Proposition 8. *If $c = 0$ and $V_i(\cdot)$ is strictly monotone, then no matter j's disclosure strategy, the best response disclosure threshold for i is the same as when he is a single sender.*

Proof of Proposition 8. Denote $r \equiv \frac{1-\overline{\beta}_{DM}}{\overline{\beta}_{DM}} \frac{\overline{\beta}_i}{1-\overline{\beta}_i}$, and define $W(\beta, r) := V_i(T(\beta, r)) - V_i(\beta)$, where $T(\beta, r) = \frac{\beta}{\beta+(1-\beta)r}$ is a shorthand for the $T(\beta, \overline{\beta}_{DM}, \overline{\beta}_i)$ transformation defined in (2). When $V_i(\cdot)$ is strictly monotone and $c = 0$, i's best response threshold must be such that $\mathbb{E}[W(\cdot, r)] = 0$ when r is determined by i's threshold type and the DM's nondisclosure belief.

Observe that when $r = 1$, then for any β, $T(\beta, 1) = \beta$ and hence $W(\beta, 1) = 0$. Furthermore, because $T(\beta, r)$ is strictly decreasing in r for all interior β, it follows that for any non-perfectly-informative experiment, $\mathbb{E}[W(\cdot, r)] = 0$ if and only if $r = 1$. Thus, no matter j's disclosure strategy (so long as it is not perfectly informative of the state, which it cannot be since $p_j < 1$), i's best response threshold is such that $r = 1$, i.e., the DM's nondisclosure belief is the same as i's threshold type. But this is the same condition as in the single-sender game. $\square$

When $c \neq 0$, there are non-linear specifications for $V_i(\cdot)$ under which our themes about strategic complementarity under concealment cost or substitutability under disclosure cost do extend, and there are other specifications which make conclusions ambiguous or even reversed. Below, we show through a family of exponential utility functions how our conclusions are affected by departures from linearity of $V_i(\cdot)$. To this end, define $W(\beta, r) := V_i(T(\beta, r)) - V_i(\beta)$, as in the proof of Proposition 8. Thus, under a disclosure cost $(c > 0)$ the relevant case is $r > 1$ if the sender is upward biased and $r < 1$ if the sender is downward biased; under a concealment cost $(c < 0)$ the relevant case is $r < 1$ if the sender is upward biased and $r > 1$ if the sender is downward biased.

In the following proposition, we say that an upward biased sender i's disclosure is a strategic substitute (resp., complement) to j's if whenever j's message is more Blackwell-informative, i's largest and smallest best response disclosure thresholds increase (resp., decrease).

Proposition 9. *Assume $V_i(\beta) = \gamma\beta^\alpha$, where either $\gamma, \alpha > 0$ or $\gamma, \alpha < 0$, so that sender i is upward biased. Then, i's disclosure is:*

1. *a strategic substitute to j's under disclosure cost if $0 < \alpha \leq 1$ and $\gamma > 0$.*
2. *a strategic substitute to j's under concealment cost if $\alpha < 0$ and $\gamma < 0$.*
3. *a strategic complement to j's under disclosure cost if $\alpha \leq -1$ and $\gamma < 0$.*

Part 1 of Proposition 9 is a generalization of Part 2 of Proposition 3 to some non-linear preferences; on the other hand, Parts 2 and 3 of Proposition 9 show how our main findings of strategic complementary under concealment cost and strategic substitutability under disclosure cost can actually be reversed for other non-linear preferences. Note that one has to be careful with the analog of Proposition 9 for the case of a downward biased sender, because the direction of disagreement between i and the DM reverses. Thus, if $V_i(\beta) = -\gamma\beta^\alpha$, then in each part of Proposition 9 one should replace "disclosure cost" with "concealment cost" and vice-versa.

Given the discussion preceding Proposition 9, and invoking Blackwell's results as in discussion in the text after Theorem 1, Proposition 9 is a straightforward consequence of the following lemma.

Lemma 7. *If $V_i(\beta) = \beta^\alpha$ then $G(\beta, r)$ is:*

1. *convex in β if $0 < \alpha \leq 1$ and $r > 1$;*
2. *concave in β if $\alpha < 0$ and $r < 1$;*
3. *convex in β if $\alpha < -1$ and $r > 1$;*

Proof of Lemma 7. Denoting partial derivatives with subscripts as usual, we obtain that $W_{\beta\beta}(\cdot)$ is equal to

$$V_i''\left(\frac{\beta}{\beta + (1-\beta)r}\right)\left(\frac{r^2}{(\beta + (1-\beta)r)^4}\right) + V_i'\left(\frac{\beta}{\beta + (1-\beta)r}\right)\left(\frac{2r(r-1)}{(\beta + (1-\beta)r)^3}\right) - V_i''(\beta).$$

Plugging in $V_i(\beta) = \beta^\alpha$ and doing some algebra shows that $W_{\beta\beta}(\cdot)$ has the same sign as:

$$\alpha\left[(1-\alpha) + \frac{r(2\beta(r-1) - r(1-\alpha))}{(\beta + (1-\beta)r)^{\alpha+2}}\right] =: H(\beta, \alpha, r).$$

Observe that $H(0, \alpha, r) = \alpha(1-\alpha)(1-r^{-\alpha})$, and hence if $\alpha < 1$ and $\alpha \neq 0$ then $\text{sign}[H(0, \alpha, r)] = \text{sign}[r-1]$. Differentiating yields

$$H_\beta(\cdot) = \frac{\alpha(\alpha+1)(r-1)r(\alpha r + 2\beta(r-1))}{(\beta + (1-\beta)r)^{\alpha+3}}.$$

We now consider four cases:

1. Suppose $0 < \alpha \leq 1$ and $r > 1$. Then $H(0, \alpha, r) \geq 0$ and $H_\beta(\cdot) > 0$, and hence $H(\beta, \alpha, r) > 0$ for all $\beta \in (0, 1)$.
2. Suppose $-1 \leq \alpha < 0$ and $0 \leq r < 1$. Then $H(0, \alpha, r) < 0$ and $H_\beta(\cdot) \leq 0$, and hence $H(\beta, \alpha, r) < 0$ for all $\beta \in (0, 1)$.
3. Suppose $\alpha < -1$ and $r > 1$. Then $H(0, \alpha, r) > 0$ and $H(1, \alpha, r) = \alpha(r-1)(\alpha - 1 + r(1+\alpha)) > 0$. We will show that $H_\beta(\beta, \alpha, r) = 0$ implies $H(\beta, \alpha, r) > 0$, which combines with the previous two inequalities to imply that $H(\cdot) > 0$. Accordingly, assume $H_\beta(\beta, \alpha, r) = 0$, which occurs when $\beta = \frac{\alpha r}{2(1-r)}$, which implies $\alpha \in (-2, -1)$ and $r \geq \frac{2}{2+\alpha}$ (because $\beta \leq 1$ and $\alpha < -1$) . Furthermore,

$$H\left(\frac{\alpha r}{2(1-r)}, \alpha, r\right) = \alpha\left[1 - \alpha - r^2\left(\frac{2}{r(\alpha+2)}\right)^{\alpha+2}\right].$$

The derivative of the above expression with respect to r is $\alpha\left(\frac{2}{\alpha+2}\right)^{\alpha+2} r^{-\alpha-1}$, which is strictly positive given $\alpha \in (-2,-1)$ and $r > \frac{2}{\alpha+2} > 2$. Moreover, when evaluated with $r = 2$, the expression reduces to $1-\alpha-\frac{4}{(\alpha+2)^2}$, which is strictly positive given $\alpha \in (-2,-1)$. Therefore, $H\left(\frac{\alpha r}{2(1-r)}, \alpha, r\right) > 0$, as was to be shown.

4. Suppose $\alpha < -1$ and $0 \leq r < 1$. Then $H(0,\alpha,r) < 0$ and $H(1,\alpha,r) = \alpha(r-1)(\alpha-1+r(1+\alpha)) < 0$. As argued in the previous case, $H_\beta(\beta,\alpha,r) = 0$ requires $\alpha \in (-2,-1)$ and $r \geq \frac{2}{2+\alpha} > 2$, which is not possible given that we have assumed $r < 1$. Thus, $H_\beta(\cdot)$ has a constant sign in the relevant domain, which implies that $H(\cdot) < 0$ in the relevant domain. □

C.3. The role of linearity and MLRP in the signaling application

Recall that in an LSHP equilibrium, the strategy $\mu(\cdot)$ for types that separate is determined by the initial condition $\mu(0) = 0$ and the differential equation (8):

$$\frac{\partial c(\mu(t),t)}{\partial m}\mu'(t) = \frac{\partial \mathbb{E}_{s|t}\left[\beta(s;t)\right]}{\partial \pi}.$$

Part 1 of Lemma 6 established that

$$\frac{\partial \mathbb{E}_{\tilde{s}|t}\left[\beta(\tilde{s};t)\right]}{\partial \pi} \leq \frac{\partial \mathbb{E}_{s|t}\left[\beta(s;t)\right]}{\partial \pi} \tag{14}$$

when signal $\tilde{s}$ is more informative (i.e., drawn from a more informative experiment) than signal s. Inequality (14) implies that the solution to the aforementioned initial-value problem is pointwise lower under the more informative experiment, and hence the equilibrium signaling level $\mu(t)$ is lower for every type when the receiver has access to $\tilde{s}$ rather than s.

We show below how the conclusion can be altered by dropping either linearity of the sender's payoff in the receiver's posterior (Example 3) or the MLRP property of the receiver's experiments (Example 4).

Example 3. Letting $V(\beta) \equiv \beta/(1-\beta)$, suppose the sender's payoff is

$$V(\beta) - c(m,t),$$

which is convex in β. Assumption (7) in Section 4 continues to imply the relevant single-crossing condition for this modified objective. Using Bayes rule, we compute

$$\mathbb{E}_{s|t}[V(\beta(s;\pi))] = \frac{\pi}{1-\pi}\mathbb{E}_{s|t}\left[\frac{g(s|1)}{g(s|0)}\right].$$

Differentiating and evaluating at $\pi = t$,

$$\frac{\partial \mathbb{E}_{s|t}[V(\beta(s;t))]}{\partial \pi} = \frac{1}{t(1-t)}\mathbb{E}_{s|t}\left[\frac{\beta(s;t)}{1-\beta(s;t)}\right].$$

The term inside the expectation operator on the right-hand side above is a convex function of $\beta(\cdot)$. It follows that

$$\frac{\partial \mathbb{E}_{\tilde{s}|t}[V(\beta(\tilde{s};t))]}{\partial \pi} \geq \frac{\partial \mathbb{E}_{s|t}[V(\beta(s;t))]}{\partial \pi},$$

by contrast to inequality (14). That is, the convexity in $V(\cdot)$ is strong enough to ensure that the marginal benefit from inducing a higher interim belief π (locally, at $\pi = t$) is higher when the exogenous signal is more informative. It follows that in an LSHP equilibrium, all types bear a *higher* (at least weakly) signaling cost when the exogenous signal is more informative.[23] □

Example 4. To see the role of MLRP-experiments, we have to modify the signaling model of Section 4 by introducing more states, because any experiment in a two-state model is an MLRP-experiment.

Assume a full-support common prior about the state $\omega \in \Omega = \{0, 1, 2\}$. The sender receives some private information, indexed by $t \in [0, 1]$, which updates his belief about the state to $(z, 1 - z(1 + t), zt)$, where each element of this vector is the probability assigned to the corresponding state. The parameter $z \in (0, 1/2)$ is a commonly known constant. We refer to t as the sender's type. Letting $M(\boldsymbol{\beta}) \equiv \sum_{\omega} \omega\beta(\omega)$ be the receiver's expectation of the state when she holds belief $\boldsymbol{\beta}$, the sender's payoff is

$$M(\boldsymbol{\beta}) - c(m, t).$$

Let s represent the outcome of an uninformative experiment, and let $\boldsymbol{\beta}(s; \hat{t})$ represent the posterior of the receiver after observing s when she puts probability one on the sender's type $\hat{t}$. It clearly holds that

$$\mathbb{E}_{s|t}\left[M(\boldsymbol{\beta}(s;\hat{t}))\right] = M(\boldsymbol{\beta}(s;\hat{t})) = 1 - z + z\hat{t}.$$

Now, consider an experiment $\tilde{\mathcal{E}}$ with a binary signal space, $\tilde{s} \in \{s_l, s_h\}$. Let $g(\tilde{s}|\omega)$ denote the probability of the signal realization in each state, specified as follows:

$$\begin{bmatrix} g(s_l|0) & g(s_l|1) & g(s_l|2) \\ g(s_h|0) & g(s_h|1) & g(s_h|2) \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}.$$

This experiment is the same as that used in Example 1 of Appendix B; it is not an MLRP-experiment. Note also that $\tilde{s}$ is more informative than s.

Suppose the receiver ascribes probability one to the sender's type $\hat{t}$. By Bayes rule, if the signal realization is s_l, the receiver's posterior belief about the state is $\boldsymbol{\beta}(s_l; \hat{t}) = (0, 1, 0)$, and

[23] On the other hand, if $V(\beta) \equiv \log[\beta/(1-\beta)]$, then the local marginal benefit of inducing a higher interim belief is independent of the exogenous experiment. The reason is that now $V(\beta(s,\pi)) = \log\left(\frac{\pi}{1-\pi}\right) + \log\left(\frac{g(s|1)}{g(s|0)}\right)$ and hence $\partial \mathbb{E}_{s|t}[V(\beta(s;t))]/\partial \pi$ does not vary with the experiment. Note that in this example, $V(\cdot)$ is neither convex nor concave.

therefore $M(\boldsymbol{\beta}(s_l;\hat{t}))=1$. For signal realization s_h,

$$\boldsymbol{\beta}(s_h;\hat{t})=\left(\frac{1}{1+\hat{t}},0,\frac{\hat{t}}{1+\hat{t}}\right),$$

and therefore

$$M(\boldsymbol{\beta}(s_h;\hat{t}))=\frac{2\hat{t}}{1+\hat{t}}.$$

The sender of type t's expectation is

$$\mathbb{E}_{\tilde{s}|t}\left[M(\boldsymbol{\beta}(\tilde{s};\hat{t}))\right]=(1-z(1+t))M(\boldsymbol{\beta}(s_l;\hat{t}))+z(1+t)M(\boldsymbol{\beta}(s_h;\hat{t})).$$

The derivative of the above expression with respect to $\hat{t}$, evaluated at $\hat{t}=t$, is

$$\frac{\partial\mathbb{E}_{\tilde{s}|t}\left[M(\boldsymbol{\beta}(\tilde{s};t))\right]}{\partial\hat{t}}=\frac{2z}{1+t}\geq z=\frac{\partial\mathbb{E}_{s|t}\left[M(\boldsymbol{\beta}(s;t))\right]}{\partial\hat{t}}. \tag{15}$$

This inequality is opposite to inequality (14), even though $\tilde{s}$ is more informative than s.

In this example, $\mathbb{E}_{s|t}[M(\boldsymbol{\beta}(s;\hat{t}))]$ and $\mathbb{E}_{\tilde{s}|t}[M(\boldsymbol{\beta}(\tilde{s};\hat{t}))]$ are both (weakly) supermodular in the sender's type t. The assumption that $c_{mt}(m,t)<0$ ensures that indifference curves in the space of $(m,\hat{t})$ for different types are single-crossing. As local incentive compatibility then implies global incentive compatibility, inequality (15) implies that all types incur higher (at least weakly) signaling costs when the receiver has access to more information. □