

Bayesian Persuasion in the Digital Age

Gabriel Martinez-Roa*

For the most recent version click here.

May 18, 2020

Abstract

In digital platforms, agents have vast access to information, but the quality of it is unclear. This paper studies the impact of uncertainty about the quality of information as a Bayesian persuasion game with multiple senders. The model innovates by assuming the receiver samples only a small subset of the information available and whose quality uncertainty is endogenous. The receiver knows the signals chosen by each sender but randomly observes the realization of only one such signal. In equilibrium, senders can pool their signals, so the receiver is uncertain about the informativeness of the message received. Since pooling is prevalent, each sender chooses a signal to affect the average correlation between messages and states of the world. Thus, the strategies of other senders constrain each sender's ability to design information and may incentivize a sender to provide more or less information compared to the single-sender information design benchmark. The model suggests policies to improve the quality of the information on platforms such as social media.

JEL CLASSIFICATION: D82, D83, D91

KEYWORDS: Bayesian Persuasion, Uncertainty, Social Media

1 Introduction

In the current digital age, most people receive their information through digital platforms. These platforms provide people with access to vast amounts of information on practically any topic of interest. Intuitively, the great variety of information and diversity of sources should lead to well-informed actions. Paradoxically, there is growing concern about the spread of low-quality information and its effect on people's behavior (See Lazer et al., 2018). On the one hand, the internet has promoted an increase in the number and variety of sources where people can get information. On the other hand, the internet also enables many of the new sources to mimic reputable sources in form, but not in the editorial rigor and quality of information. According to the Pew Research Center (2018), 68% of Americans report getting news from social media. Of them, however, 57% report that they expect social media news to be mostly inaccurate¹. This paper endogenizes the

*University of Wisconsin-Madison, Department of Economics; E-mail:gamartinez@wisc.edu, Website: <https://gabriel-martinez-roa.github.io>

¹ In their survey, Facebook, Twitter, Youtube, or Reddit were the main pathways to news for 81% of the U.S. adults who get news from social media.

uncertainty about the quality of information and derives policy recommendations to improve the incentives to provide information through communication platforms such as social media.

To study the impact of uncertainty about the quality of information, we develop a Bayesian persuasion game with multiple senders and a representative receiver. The key innovation of the model is that each sender produces information, but the receiver samples only a small subset of the messages produced; this is particularly salient in digital media since people process significantly less information than all of what is available. We also assume that senders have a shared message space, thus allowing pooling. That is, senders can share the support of their signals while correlating messages with states of the world heterogeneously. The receiver knows the signals chosen by each sender but randomly observes the realization of one such signal. If two or more of the signals chosen by senders have the same message in their support, the receiver is uncertain about the precision of such a message. We use this assumption to motivate the uncertainty of messages online. In contrast with traditional news sources like TV, radio, or newspapers, where the consumer knows the reputation of the different sources, digital platforms are less regulated as to who provides information and how (See Pew, 2018 for documentation of stylized facts). This looser regulations create significantly less transparency and allow information of different quality to be lumped together into the same platform and formatting.

Kamenica and Gentzkow’s (2011) Bayesian persuasion model (KG11 henceforth) shows that even senders with extreme bias can benefit from persuasion whenever they can commit to a signal structure. This paper shows how (even with commitment) the receiver’s uncertainty about information quality limits the value of providing information. However, this uncertainty can incentivize senders to produce more precise signals, or obfuscate all valuable communication depending on the types of sources and the likelihood that the receiver will sample them. For example, if the receiver cannot discern respectable information providers (think-tanks, journalists, analysts) from low-quality sources, except by the content of the message received, the receiver faces a double inference problem. From the message, the receiver must infer both the state of the world and the precision of the message. In equilibrium, reputable sources can respond by providing more precise information to remain influential, but if, due to uncertainty, persuading the receiver is impossible, they may instead give up.

To illustrate this mechanism, consider a slight modification to KG11’s model. Suppose there is an election. There are two candidates in the ballot (x and y), and a voter wants to vote for the best candidate. Assume the voter casts a vote for whichever candidate is more likely to be the best, and the common prior belief is that candidate x is more likely to be the best, say with 80% probability. Before the election, the voter will read about the candidates and make up her mind. Suppose there are two kinds of journalists producing information about candidates: a good one (g) who wants to reveal as much information as possible, and a bad one (b) who wants the voter to choose the worse candidate. Each of them can choose a signal which is a function from the state space, $\Theta = \{x, y\}$, into a distribution over recommendations², $\Delta(M)$ where $M = \{x, y\}$. Let us suppose the good journalist chooses a perfectly informative signal, denoted π_g , while the bad

²In section 4 we show how the restriction to direct messages is without loss.

journalist chooses a perfectly uninformative signal, denoted π_b :

$$\begin{aligned} \pi_g(x|x) = 1 \quad \pi_g(y|x) = 0 & & \pi_b(x|x) = \frac{1}{2} \quad \pi_b(y|x) = \frac{1}{2} \\ \pi_g(x|y) = 0 \quad \pi_g(y|y) = 1 & & \pi_b(x|y) = \frac{1}{2} \quad \pi_b(y|y) = \frac{1}{2}. \end{aligned}$$

The voter is aware of these two types of journalists and knows how each of them correlates messages with states of the world. However, the receiver does not read all of the information produced, but instead, consumes whichever information she receives first. Suppose that each journalist's recommendation is equally likely to reach the receiver first. Then, her posterior beliefs after each recommendation are as follows:

$$\mu_1(x|x) = \frac{0.6}{0.6 + 0.05} \approx 0.92 \quad \mu_1(x|y) = \frac{0.2}{0.2 + 0.15} \approx 0.57.$$

The recommendation y is not informative enough, so the receiver's best response after either message is to vote for candidate x . Since this is the same action she would have taken had she not received any information, the receiver obtained no useful information. Further, no sender has a profitable deviation. The probability of choosing the worse candidate is highest: given the prior, there is a 20% probability of choosing the worse candidate. Therefore, b is getting his first-best outcome, while g is obtaining his worst possible outcome. However, the good journalist is already providing as much information as possible, so he has no profitable deviation.

Notice that if the receiver faced no uncertainty, as in KG11, the bad journalist would be indifferent between any uninformative signal, but this is not the case in this model. For example, if b were to choose a signal that always recommends that the receiver choose her default action (to vote for the candidate x), then the receiver's posteriors after each recommendation would be:

$$\mu_1(x|x) = \frac{0.8}{0.8 + 0.1} \approx 0.89 \quad \mu_1(x|y) = \frac{0}{0 + 0.1} = 0.$$

So the receiver will vote for y whenever she receives information recommending y . In this case, the receiver infers that the recommendation y comes from g , thus revealing that the state is y . In this situation, the probability of the receiver voting for the worse candidate is down to 10%. Therefore, b has strict preferences between uninformative signals. This example illustrates how, in contrast with KG11, the choice of signals cannot be reduced to choosing a Bayes-plausible distribution over posteriors.

To study equilibria in this game, we simplify its structure to work in the space of signals. We show how each sender's problem is equivalent to that of a single information designer with an endogenous, linear constraint that depends on the strategies of all other senders in the platform – thus leveraging tools developed by Bergemann and Morris (2019). In this model, a separating equilibrium requires senders' preferences to be at least locally aligned, thus if senders are heterogeneous in their preferences, a semi-pooling equilibrium occurs. Lastly, we use the framework to study large classes of games and derive policy recommendations. We show that a hands-off ap-

proach that allows noisy signals to become more prevalent has a disproportionately large negative effect. We also illustrate the drawbacks of increasing the variety of senders' preferences. Finally, we characterize the value of enabling senders to disclose their biases credibly.

The structure of the remainder of the paper is as follows: the next section reviews the literature; section 3 presents the model, and in section 4 we simplify its analysis; section 5 characterizes equilibrium in general while section 6 specializes to the aforementioned classes of games; finally, section 7 concludes.

2 Literature Review

In this paper, we study a sender-receiver game with many heterogeneous senders and a representative consumer. Senders can commit to a signal (or Blackwell experiment) in the vein of the growing Bayesian persuasion literature. However, the paper departs from previous literature by adopting two key assumptions. The receiver samples only one of the signals produced, and upon observing a signal realization, the receiver is uncertain about the source.

Gentzkow and Kamenica (2017 a,b) explore Bayesian Persuasion models with multiple senders and find sufficient conditions so that more competition weakly increases the informativeness of equilibria. Li and Norman (2018) consider a similar setting but relax three key assumptions to show how each of these assumptions is also necessary for the latter result. In particular, if senders choose their signals sequentially, can play mixed strategies or cannot correlate their signals with that of other senders, increasing the number of senders may decrease the informativeness of equilibria. I relax another assumption in Gentzkow and Kamenica (2017 a,b), namely that the receiver processes all the information available, and find as well that increasing the number of senders can decrease the informativeness of equilibria.

The second assumption pertains to senders being able to pool with other senders, thus allowing uncertainty about the quality of the message to be endogenous. Previous research has studied how the uncertainty about the incentives (or bias) of a sender affects a sender's incentives to disclose information (see Morgan and Stocken, 2003, Dimitrakas and Sarafidis, 2005, Li and Madarasz, 2008 and the appendix in Chakraborty and Harbaug, 2010). That is, the sender knows the state of the world and uses cheap-talk to fully or partially disclose such information. Similarly, Shin (1994) and Wolinsky (2003) study similar setups in communications games with verifiable information, in the vein of Grossman (1981) and Milgrom (1981). This paper fills a gap by examining these questions within the Bayesian persuasion framework. By abstracting from the incentives to disclose information, the model shows how uncertainty can simultaneously limit the value of providing information and incentivize senders to choose more informative signals. Further, since Bayesian persuasion characterizes an upper bound in communication, the negative results in this paper are informative about any more complex model that also considers the incentives to disclose information.

Within the Bayesian persuasion framework, the literature has studied the effect of uncertainty about the source or the quality of the messages in one of three main ways. First, Le Treust and Tomala (2019) and Tsakas and Tsakas (2018) consider a game where recommendations are garbled

with some probability. In these papers, the garbling mechanism is exogenous, while here, the uncertainty is endogenous. Second, Fréchet *et al.* (2019) and Min (2017) consider a sender with partial commitment. In their models, with some exogenous probability, the sender can privately modify their recommendation after the signal has realized. That is, with some probability, the message is cheap talk. In this model, all senders can commit to a signal. Thirdly, several recent papers have studied Bayesian persuasion with a privately informed sender. Perez-Richet (2014), Kosenko (2018), Piermont (2016), and Hedlung (2017) consider the Bayesian persuasion problem with privately informed senders. In this paper, senders are symmetric in terms of their private information but heterogeneous in their preferences.

From the receiver’s perspective, there are two dimensions of uncertainty in this model: the state of the world and the identity of sender providing information. This paper considers senders that can commit to messages that are state-dependent, but not sender-dependent. In a closely related work, Jain (2018) also considers a two-dimensional space of uncertainty, where senders have commitment power in only one dimension. In Jian’s work, both dimensions determine the receiver’s payoff directly. Here, the receiver’s payoff does not directly depend on the sender’s identity.

3 Model

There are S senders and a single receiver. The receiver and each sender has preferences (possibly state-dependent) over the action chosen by the receiver, $a \in A$. Let $u_s(a, \theta)$ be the utility for sender $s \in S$ and $u_r(a, \theta)$ be the utility for the receiver, where $\theta \in \Theta$ is the state of the world. Both Θ and A are assumed finite. Each sender can choose a signal - or Blackwell experiment - $\pi_t : \Theta \rightarrow \Delta(M)$, which specifies a probability distribution over a common message space, M , for each state of the world³. M is assumed finite, but rich enough: $M > \Theta$ ⁴.

The timing is as follows. First, each sender chooses simultaneously their signal, π_s , without knowing the state of the world. The receiver observes the profile of strategies, $\{\pi_s\}_{s \in S}$, and the realization of the signal of one of the senders which is chosen at random. Then the receiver updates her beliefs and takes an action. Finally, all players’ payoffs, $\{u_r, \{u_s\}_{s \in S}\}$, are realized.

Players have a common prior regarding the state of the world, denoted $\mu_0 \in \Delta(\Theta)$. Players also have a common prior belief about the likelihood that the receiver will observe the signal realization of each sender, denoted $\nu_0 \in \Delta(S)$. This probability can be interpreted as the prevalence of each sender in the platform. For simplicity, μ_0 and ν_0 are assumed independent⁵.

Persuasion is assumed costless. Senders’ payoffs do not depend on the signal they choose. This assumptions allows us to focus on the role that the uncertainty about which signal realization will be observed by the receiver has on the incentives for senders. I also focus attention on pure strategy equilibria. The latter is without loss, since a mixed strategy is a distribution over Blackwell

³Senders have no restriction on how correlated their messages can be with the state of the world; they can finely choose how much information to produce.

⁴As established by lemmas 1, it is without loss to consider M to be finite.

⁵The qualitative results follow even if these distributions were correlated.

experiments, so it is itself a Blackwell experiment. Thus, any mixed strategy is equivalent to a pure strategy⁶.

An important assumption is that all senders have access to the same messages, M . So in general, after observing a message, the receiver remains uncertain about which signal produced it, hence about the correlation between the message and the states of the world. Since the correlation between a message and each state of the world can be chosen differently by different senders, the receiver has a double inference problem. From the message observed, the receiver must learn both about the state of the world and the quality of the message. By assuming that all senders have access to the same message space we capture the fact that in digital media it is relatively easy for sources to remain anonymous or for sources to misrepresent the quality of their messages. In section ??, I relax this assumption by allowing senders to have access to exclusive messages and discuss its implications.

Let $p_0 = \mu_0 \cdot \nu_0$ denote the joint prior distribution. After receiving a message, m , the receiver learns about the state of the world and about the source of the message. Denote the joint posterior by $p_1(\theta, s|m) \in \Delta(\Theta \times S)$ with marginal distributions over the state of the world denoted by $\mu_1(m) = \sum_S p_1(\cdot|m) \in \Delta(\Theta)$ called the **state-posterior**, and for the source of the message by $\nu_1(m) = \sum_\Theta p_1(\cdot|m) \in \Delta(S)$, called the **source-posterior**. We have that, given a profile of signals $\{\pi_s\}_{s \in S}$, each message in its support⁷ induces the following posterior belief:

$$p_1(\theta, s|m) = \frac{\pi_s(m|\theta)p_0(\theta, s)}{\sum_\Theta \sum_S \pi_{s'}(m|\theta')p_0(\theta', s')}, \quad \forall \theta \in \Theta, s \in S, \text{ given } m \in M; \quad (1)$$

Since the observed message is stochastic, let the distribution over posteriors that is induced by $\{\pi_s\}_{s \in S}$ be written as $\tau(\{\pi_s\}) \in \Delta(\Delta(\Theta \times S))$. This distribution is characterized by two properties:

1. Every point in its support, $p_1(\cdot|m)$, is defined by the previous equation; i.e. using Bayes rule.
2. The probability mass on any posterior p in the support of τ , is given by the total probability of receiving a message that leads to such a posterior:

$$\tau(p) = \sum_{m:p_1(\cdot|m)=p} \sum_\Theta \sum_S \pi_s(m|\theta)\nu_0(s)\mu_0(\theta).$$

The equilibrium concept is Perfect Bayesian Equilibrium. That is, an equilibrium consists of the profile of senders' signals, a receiver's (possibly mixed) strategy as a response function to any belief induced by some profile of senders' signals, and beliefs that are induced by $\{\pi_s\}_{s \in S}$ through Bayesian updating.

Definition 1 (Equilibrium) *An equilibrium is a collection of signals, $\{\pi_s^*\}_{s \in S}$, a response function $a^* : \{\pi_s\}_{s \in S} \rightarrow \Delta(A)$, and beliefs $p_1 : \{\pi_s\}_{s \in S} \rightarrow \Delta(\Theta \times S)$ such that:*

⁶The mechanisms explored by Li and Norman (2018) regarding mixed strategies are mute here since the receiver observes the realization of only one signal.

⁷The support of a collection of signals is the union of the supports of each signal for each state of the world: $\text{support}(\{\pi_s\}_{s \in S}) := \bigcup_{s \in S} \bigcup_{\theta \in \Theta} \text{support}(\pi_s(\theta))$

1. $a^*(p_1(\{\pi_s\}_{s \in S})) \in \Delta(A)$ is the receiver's best response to any belief in the image of p_1 :

$$\text{support } a^*(p_1(\cdot)) \subseteq \underset{a \in A}{\operatorname{argmax}} E_{p_1(\cdot)} [u_r(a, \theta)];$$

2. $p_1(\{\pi_s\}_{s \in S})$ is given by equation (1) for all $m \in \text{support}(\{\pi_s\}_{s \in S})$;

3. π_s^* is a best response to the strategies of all other senders $\{\pi_{s'}^*\}_{s' \neq s}$ and the receiver's response function $a^*(p(\cdot))$:

$$\pi_s^* \in \underset{\Pi}{\operatorname{argmax}} E_{\tau(\{\pi_{s'}^*\})} E_{p_1} u_s(a^*(p_1), \theta) \quad \forall s \in S, \text{ given } \{\pi_{s'}^*\}_{s' \neq s}.$$

Importantly, $a^*(p_1(\cdot))$ is a well defined function everywhere. At each belief where the receiver is indifferent, some mixed strategy is assumed, and this selection rule is commonly known to all senders. Common selections in the literature are sender-optimal strategies, and sender-pessimal strategies. In the case of a single sender, the game can be interpreted as one of an information designer. Therefore, these selections can be motivated in the spirit of partial implementation and robust implementation respectively. In this model, however, different senders may have different optimal and pessimal strategies for beliefs where the receiver is indifferent. In the spirit of partial implementation, I allow the selection to be part of the equilibrium concept.

Though equilibrium is not guaranteed to exist in general, it does exist for a large class of games, as illustrated in section 6. Further, I develop a sufficient condition that guarantees existence and is likely to hold when sender's prevalence is not too concentrated in any single sender, see section 5.

In the next section I develop a geometric interpretation that is useful in characterizing equilibria for specific applications. In general, the problem does not allow for traditional concavification techniques, because the state-posteriors, μ_1 , that a sender can induce not only have to be a mean preserving spread of the prior μ_0 , but also must be "*feasible*" given the strategies of other senders. In the next section I formalize these feasibility constraints for each sender.

4 Model Simplification

In this section I simplify the analysis of the model in four steps. First I show it is reach enough to assume a message space with cardinality $|M| = |A|$. As in previous literature, messages can be assumed to be recommendations, but in contrast with previous literature, recommendations need not always be followed in equilibrium. Second, I justify assuming that the receiver's response function depends only on the state-posterior (as opposed to the joint posterior $p_1 \in \Delta(\Theta \times S)$). Third, I show that the average signal, $\bar{\pi}$, is a sufficient statistic for the strategies of all senders. That is, each individual signal affects the receiver's decision only through its effect on $\bar{\pi}$. Finally, I reformulate each sender's problem to one of a standard information designer (similar to Bergemann and Morris, 2019) with an additional endogenous constraint given by the strategies of all other senders.

We start by noting that senders' signals, π_s , are not restricted to have full support in M . Indeed, whenever sender s is the only one using some message, say m , not only will the receiver identify the sender after receiving message m , i.e. $\nu_1(s|m) = 1$, but the sender has full control over the receiver's posterior beliefs over Θ after this message

$$\mu_1(\theta, |m) = \frac{\pi_s(m|\theta)\mu_0(s)}{\sum_{\Theta} \pi_s(m|\theta')\mu_0(s)} \quad \forall \theta \in \Theta.$$

We could suspect that the cardinality of M affects the set of outcomes that can arise in equilibrium. I show next how this is not the case. In actuality, if the message space is as large as the action space, increasing the cardinality of the M does not affect the distribution of outcomes that can occur in equilibrium⁸.

Proposition 1 (Message space simplification) *If $|A| \leq |M|$, there exists an equilibrium of the game where the message space is M if and only if there is an outcome equivalent equilibrium in the game where the message space is A .*

This result establishes that it is without loss to restrict the cardinality of the message spaces, which greatly simplifies the search for equilibria. Though similar in spirit to proposition 1 in KG11, there is a key distinction between their model and this one. Messages can be assumed *straight-forward*, i.e. each message recommends an action. However, the recommendations are not always followed in equilibrium. Theorem 1, for example, provides a sufficient condition for the receiver to disregard any recommendation that does not agree with her default action, $a^*(\mu_0)$.

Next, we establish that if the receiver's equilibrium response function, $a^*(p_1(\theta, s))$ cannot be written as a function $\hat{a}(\mu_1(\theta))$ that depends only on the state-posterior, there must be two messages, m and m' in the support of $\{\pi_s^*\}_{s \in S}$ that induce the same state-posterior, $\mu_1(m) = \mu_1(m')$, but different source-posteriors, $\nu_1(m) \neq \nu_1(m')$ and $a^*(p_1(\theta, s|m)) \neq a^*(p_1(\theta, s|m'))$. However, corollary 1 stabilizes that if such an equilibrium exists, then all senders that use m and m' are indifferent between $a^*(p_1(\theta, s|m))$ and $a^*(p_1(\theta, s|m'))$. Hence, we can construct an outcome equivalent equilibrium where all senders' strategies are the same, except that the total mass in messages m and m' is redistributed uniformly among these messages. Therefore, $\mu_1(m) = \mu_1(m')$, and $\nu_1(m) = \nu_1(m')$. Finally, the receiver chooses the appropriate randomization between the two actions that keeps the distribution of outcomes constant across equilibria⁹. Therefore, we restrict attention to equilibria where $a^*(\cdot)$ depends only on μ_1 ¹⁰.

The following lemma establishes the previous claim that senders using m and m' with $\mu_1(m) = \mu_1(m')$, but inducing different responses from the receiver are indifferent among these distinct actions.

⁸In section 5 I provide necessary and sufficient conditions for a separating equilibrium. Though some separating equilibria require larger message spaces, I show that these equilibria are outcome equivalent to pooling or semi-pooling equilibria in agreement with the proposition 1.

⁹Given the indifference established, both the senders using m and m' and the receiver are indifferent with this change. The proof constructs the mixed strategy that keeps the rest of the senders indifferent as well.

¹⁰Any ruled out equilibria will be indirectly studied through the multiplicity of outcome equivalent equilibria.

Lemma 1 (Senders' Strategic Incentives Depend only on μ_1) *If an equilibrium has $m, m' \in \text{support}(\{\pi_s^*\}_{s \in S})$ with $\mu_1(m) = \mu_1(m')$, but $a^*(p_1(\theta, s|m)) \neq a^*(p_1(\theta, s|m'))$, then*

$$E_{p_1(\cdot|m)} u_s(a^*(p_1(\cdot|m)), \theta) = E_{p_1(\cdot|m')} u_s(a^*(p_1(\cdot|m')), \theta)$$

for all $s \in S$ with $\{m, m'\} \cap \text{support}(\pi_s^*) \neq \emptyset$.

In light of lemma 1, senders' strict incentives to choose an optimal strategy cannot directly depend on the source-posterior ν_1 . Additionally, the receiver's payoff is constant over ν_1 . Therefore, it is "essentially" without loss to restrict attention to equilibria where $a^*(p_1(\theta, s))$ depends only on the state-posterior so can be written as a function $\hat{a}(\mu_1(\theta))$. Proposition 2 formalizes this statement.

Proposition 2 (Simplification of a^*) *If $(\{\pi_s^*\}_{s \in S}; a^*(p_1); p_1)$ is an equilibrium, then either:*

- $a^*(p_1(\theta, s)(\{\pi_s\}_{s \in S})) \equiv \hat{a}(\mu_1(\theta)(\{\pi_s\}_{s \in S}))$ for every $\{\pi_s\}_{s \in S}$ and $\hat{a} : \Delta(\Theta) \rightarrow \Delta(A)$, or
- There exist an outcome equivalent equilibrium with receiver's response function, a' , such that $a'(p_1(\theta, s)(\{\pi_s\}_{s \in S})) \equiv \hat{a}(\mu_1(\theta)(\{\pi_s\}_{s \in S}))$ for every $\{\pi_s\}_{s \in S}$ and some $\hat{a} : \Delta(\Theta) \rightarrow \Delta(A)$.

Henceforth, we restrict attention to response functions that depend on the profile of signals, only through the induced state-posteriors. That is $a^* : \{\pi_s\}_{s \in S}$ is of the form $a^*(\mu_1(\{\pi_s\}_{s \in S})) \in \Delta(A)$. We call such response functions **regular**. Given this assumption, there exists a single element in Π , denoted $\bar{\pi}$ (read: the average signal) that summarizes (as formalized in proposition 3) the whole profile of signals. Let $\bar{\pi} : \{\pi_s\}_{s \in S} \rightarrow \Pi$ be defined as:

$$\bar{\pi}(m|\theta) = \sum_{s \in S} \nu_0(s) \pi_s(m|\theta) \quad \forall m \in M, \quad \forall \theta \in \Theta.$$

Proposition 3 (The average signal, $\bar{\pi}$, as a sufficient statistic for a^*) *If a^* is regular, then it depends on $\{\pi_s\}_{s \in S}$ only through $\bar{\pi}$. That is, it can be written as a function*

$$\hat{a} : \Pi \rightarrow \Delta(A) \text{ where } a^*(\mu_1(\{\pi_s\}_{s \in S})) \equiv \hat{a}(\mu_1(\bar{\pi}(\{\pi_s\}_{s \in S}))).$$

Denote the distribution over state-posteriors induced by an average signal by $\tau^\Theta(\bar{\pi}) \in \Delta(\Delta(\Theta))$.

The previous lemma establishes that any signal chosen by a sender, affects the action of the receiver only through the average signal, $\bar{\pi}$. Hence, whenever a message is used by multiple senders, each of them has partial control over the state-posterior induced by said message. In fact, given the uncertainty on the source, after observing a message, the receiver updates her beliefs as if the message was produced by signal $\bar{\pi}$. That is, the average correlation between messages and states of the world determines the receiver's response, therefore senders choose their communication strategies in order to influence this average correlation.

Remark 1 *We can relax the assumption that the receiver knows the profile of signals chosen by the senders, $\{\pi_s\}_{s \in S}$, and assume instead that the receiver only knows the average communication strategy, $\bar{\pi}$.*

To study each sender's best response I develop a geometrical tool. Notice that any signal, π_s , can be represented by a rectangular row-stochastic matrix of dimension $\Theta \times M$. That is, a matrix with non-negative elements and rows that sum up to 1. Each signal, π_s , is an element of the convex set $\Pi = \prod_{\theta \in \Theta} \Delta(M)$. Since the average signal, $\bar{\pi}$, is the weighted average of elements in Π , it is itself an element in Π :

$$\bar{\pi}(\{\pi_s\}_{s \in S}) = \sum_{s \in S} \nu_0(s) \pi_s \in \Pi. \quad (2)$$

Then, for any sender, s , we can also summarize the strategies of every other sender by a single element in Π : the weighted average of the other senders' signals. Let such signal be called the $-s$ average signal, and be denoted by $\bar{\pi}_{-s}$. It is constructed as follows:

$$\bar{\pi}_{-s}(\{\pi_{s'}\}_{s' \neq s}) := \sum_{s' \neq s} \frac{\nu_0(s')}{1 - \nu_0(s)} \pi_{s'} \in \Pi \quad (3)$$

It is clear that any two profiles of signals $\{\pi_{s'}\}_{s' \neq s}$ and $\{\pi'_{s'}\}_{s' \neq s}$ that define the same $-s$ average signal are equivalent, for strategic purposes, to sender s . This is because $\bar{\pi}$ can always be expressed as the weighted average of π_s and $\bar{\pi}_{-s}$:

$$\bar{\pi}(\pi_s, \bar{\pi}_{-s}(\cdot)) = \nu_0(s) \pi_s + (1 - \nu_0(s)) \bar{\pi}_{-s}(\cdot) \equiv \bar{\pi}(\{\pi_s\}_{s \in S})$$

Since sender s can choose any signal $\pi_s \in \Pi$, the set of average signals, (thus actions of the receiver) that the sender can induce is given by the correspondence $\bar{\Pi}_s : \Pi \times \Delta(S) \rightarrow \mathcal{P}(\Pi)$ defined as follows:

$$\bar{\Pi}_s(\bar{\pi}_{-s} | \nu_0) = \nu_0(s) \Pi + (1 - \nu_0(s)) \bar{\pi}_{-s} \subset \Pi \quad (4)$$

From equation (4) we see that the strategies of other senders constrain the elements in $\bar{\Pi}_s(\bar{\pi}_{-s} | \nu_0)$. Notice that if the prior ν_0 has full support, $\bar{\Pi}_s(\cdot)$ is always a strict subset of Π . The next lemma summarizes the main properties of this correspondence.

Lemma 2 (Properties of $\bar{\Pi}_s(\bar{\pi}_{-s} | \nu_0)$) *The correspondence $\bar{\Pi}_s$ has the following properties:*

1. $\bar{\pi}_{-s} \in \bar{\Pi}_s(\bar{\pi}_{-s} | \nu_0)$ and $\bar{\Pi}_s(\bar{\pi}_{-s} | \nu_0)$ is a contraction of Π (with distances shrunk by $\nu_0(s)$).
2. $\bar{\Pi}_s(\bar{\pi}_{-s} | \nu_0)$ is continuous in $\nu_0(s)$ and $\bar{\pi}_{-s}$, and strictly increasing in $\nu_0(s)$. That is,

$$\nu_0(s) < \nu'_0(s) \text{ implies } \bar{\Pi}_s(\bar{\pi}_{-s} | \nu_0) \subset \bar{\Pi}_s(\bar{\pi}_{-s} | \nu'_0).$$

3. Let \preceq_{bw} denote Blackwell's (1953) partial order over Π . Then, for any $\bar{\pi}_{-s}$, there exists $\bar{\pi}_l, \bar{\pi}_h \in \bar{\Pi}_s(\bar{\pi}_{-s} | \nu_0)$ such that $\bar{\pi}_l \preceq \bar{\pi}_{-s} \preceq \bar{\pi}_h$ with at least one strict inequality.

Figure 1 illustrates these geometric properties for binary Θ , binary A and $M = A$. Since $|\Theta| = |M| = 2$ we can represent $\Pi = \prod_{\Theta} \Delta(M)$ by the cartesian product of two binary simplexes.

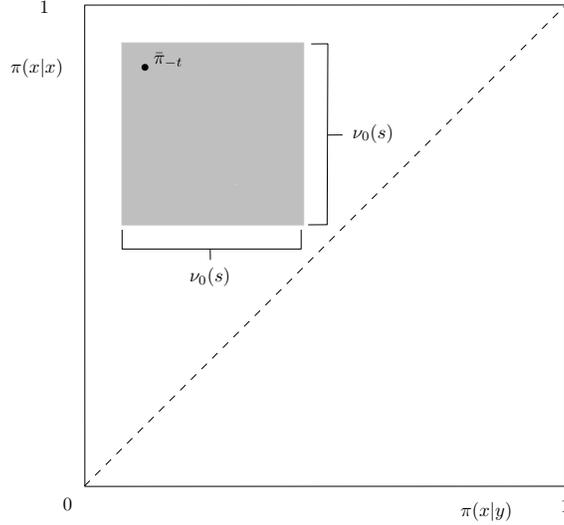


Figure 1: The space of signals Π in the binary benchmark, shaded is the set $\Pi_s(\bar{\pi}_{-s}, \nu_0) \subset \Pi$.

Given $\bar{\pi}_{-s}$, the shaded region represents all the average signals that sender s can induce by choosing some signal $\pi_s \in \Pi$. We can see that for any sender s , choosing $\bar{\pi} = \bar{\pi}_{-s}$ is always feasible, though choosing neighboring signals is also feasible. In fact, both weakly more informative (closer to the top-left or bottom-right corner) and weakly less informative (closer to the 45° line) signals are feasible. Finally, from (4) the set of $\bar{\pi}$'s that sender s can feasible induce depend continuously on the prior ν_0 , and the signal $\bar{\pi}_{-s}$, and increases with the prevalence of the sender, $\nu_0(s)$.

With this, we can write any sender's decision problem as one of choosing $\bar{\pi}$ among their feasible average signals, $\Pi_s(\cdot)$. That is taking as given the signals of other players and their own prevalence. The following result is now immediate, but it will be very useful in the characterization and applications.

Proposition 4 (Game Reformulation) π_s is a best response to $\{\pi_{s'}\}_{s' \neq s}$ given a regular response function $a^*(\cdot)$ if and only if $\bar{\pi} = \nu_0(s)\pi_s + (1 - \nu_0(s))\bar{\pi}_{-s}(\{\pi_{s'}\}_{s' \neq s})$ and

$$\begin{aligned} \bar{\pi} \in \operatorname{argmax} E_{\tau_{\Theta}(\bar{\pi})} E_{\mu_1} u_s(a^*(\mu_1), \theta) \\ \text{s.t. } \bar{\pi} \in \Pi_s(\bar{\pi}_{-s}, \nu_0) \end{aligned} \quad (5)$$

The advantage of such a reformulation is that we can see each sender as a monopolist information designer with endogenous constraints given by the strategies of the other senders. In the latter formulation, each sender chooses $\bar{\pi}$ which is the signal the receiver uses to decide which action to take (hence the ‘‘monopolist’’ interpretation), but the signal is constraint to be in $\Pi_s(\cdot)$, an endogenous set. In equilibrium all senders must choose the same $\bar{\pi}$. However, this is usually achieved by each sender choosing a different signal, π_s , as we will see in the next section.

5 Equilibrium Characterization

In this section I provide necessary and sufficient conditions for separating equilibria to exist. In light of lemma 1 we expect these conditions to be strong insofar they imply that all senders have highly aligned preferences as formalized by propositions 5 and 6. Then, I provide a sufficient condition for equilibria to exist in general (i.e. for semi-pooling equilibria) and discuss why equilibria may fail to exist.

A separating equilibrium is one where $\text{support}(\pi_s) \cap \text{support}(\pi_{s'}) = \emptyset$ for all $s \neq s' \in S$. Given that the supports are pair-wise disjoint, each sender has full control over the receiver's state-posterior beliefs conditional their messages reaching the receiver. Thus, each sender chooses a signal that induces their first-best distribution over posteriors as in KG11. I define the set of such optimal distributions of posteriors as the unrestricted monopolist benchmark and use it to characterize separating equilibria.

Definition 2 (The Unrestricted Monopolist Benchmark) *For each $s \in S$ denote by $\mathcal{T}_s^{KG} \subseteq \Delta(\Delta(\Theta))$ the set of equilibrium posterior distributions in the unrestricted monopolist benchmark: i.e. $S = \{s\}$.¹¹¹²*

Proposition 5 (Necessary Condition for Separating Equilibria) *A game has a separating equilibrium, only if all senders have **locally aligned preferences**. That is if for each $s, s' \in S$ with $s' \neq s$, there exists $\tau_s^{KG} \in \mathcal{T}_s^{KG}$ that is a local maximizer of sender s' monopolist problem:*

$$\begin{aligned} \tau_s^{KG} &\in \operatorname{argmax} E_{\tau} E_{\mu_1} u_{s'}(a^*(\mu_1), \theta) \\ \text{s.t.} \quad &\sum_{\text{support}(\tau)} \mu_1(\theta) d\tau(\mu_1) = \mu_0(\theta) \quad \forall \theta \in \Theta. \end{aligned}$$

The previous lemma, establishes that for a separating equilibrium to exist, senders' preferences have to be perfectly aligned at least locally. In the extreme case, if senders have fully aligned preferences, a separating equilibrium is for all senders to induce the same distribution over posteriors using pair-wise disjoint sets of messages. However, clearly, any such equilibrium is outcome equivalent to a pooling equilibrium where all senders choose the exact same signal. The following lemma establishes a slightly more general result that accounts for multiplicity of equilibria.

Proposition 6 (Sufficient Condition for Separating Equilibria) *If senders have **globally aligned preferences**, i.e. $\bigcap_{s \in S} \mathcal{T}_s^{KG} \neq \emptyset$, there exists a separating equilibrium and an outcome equivalent semi-pooling equilibrium.*

If $\bigcap_{s \in S} \mathcal{T}_s^{KG}$ is a singleton, there exists a separating equilibrium that is outcome equivalent to a pooling equilibria.

¹¹The KG superscript denotes that the unrestricted monopolist benchmark is akin to KG11's sender-receiver game with a single sender. Therefore, concavification techniques or Lipnowski and Mathevet (2017)'s method can be used to characterize these sets.

¹²In general, the set \mathcal{T}_s^{KG} depends on μ_0 . I omit this in the notation for simplicity.

From proposition 4 we have that if $\bar{\pi}^*$ is the equilibrium average signal, then it solves the **constrained monopolist problem** for all senders:

$$\begin{aligned} \bar{\pi} \in \operatorname{argmax} \quad & E_{\tau \ominus(\bar{\pi})} E_{\mu_1} u_s(a^*(\mu_1), \theta) \quad \forall s \in S \\ \text{s.t.} \quad & \bar{\pi} \in \Pi_s(\bar{\pi}_{-s}, \nu_0) \end{aligned}$$

The existence of such a signal cannot be guaranteed in general¹³. The following example illustrates that existence fails to exist if two or more senders are highly prevalent and have misaligned preferences.

Example 1 (The binary election game: no equilibrium) *Following with the example in the introduction, let $\Theta = A = M = \{x, y\}$ and a receiver that wants to match the state of the world with her action: $u_r := \mathbf{1}\{a = \theta\}$ so that $a^*(\mu_1) = x$ iff $\mu_1(x) \geq 0.5$.*

Let there be two senders $i, n \in S$ and the priors be $\mu_0(x) = 0.5$ and $\nu_0(i) = 0.5$. Sender i is informative and has preferences convex in μ , in fact $u_r = u_i$, but sender n is noisy and has opposite preferences: $u_n = -u_i$.

It is clear that for sender i choosing a perfectly informative signal is a weakly dominant strategy, e.g. $\pi_i(x|x) = 1, \pi_i(x|y) = 0$. Given such a signal, sender n chooses the least informative signal in $\Pi_n(\pi_i, \nu_0(n))$. However, this is not achieved by choosing π_n to have messages and states uncorrelated. In fact, it is achieved by choosing the other perfectly informative signal! That is $\pi_n(x|x) = 0$ and $\pi_n(x|y) = 1$. By choosing a signal where the correlation of messages and states has the opposite sign, the average signal becomes uninformative (see figure 2). Clearly, this cannot be an equilibrium because sender i has incentives to deviate and imitate n 's signal.

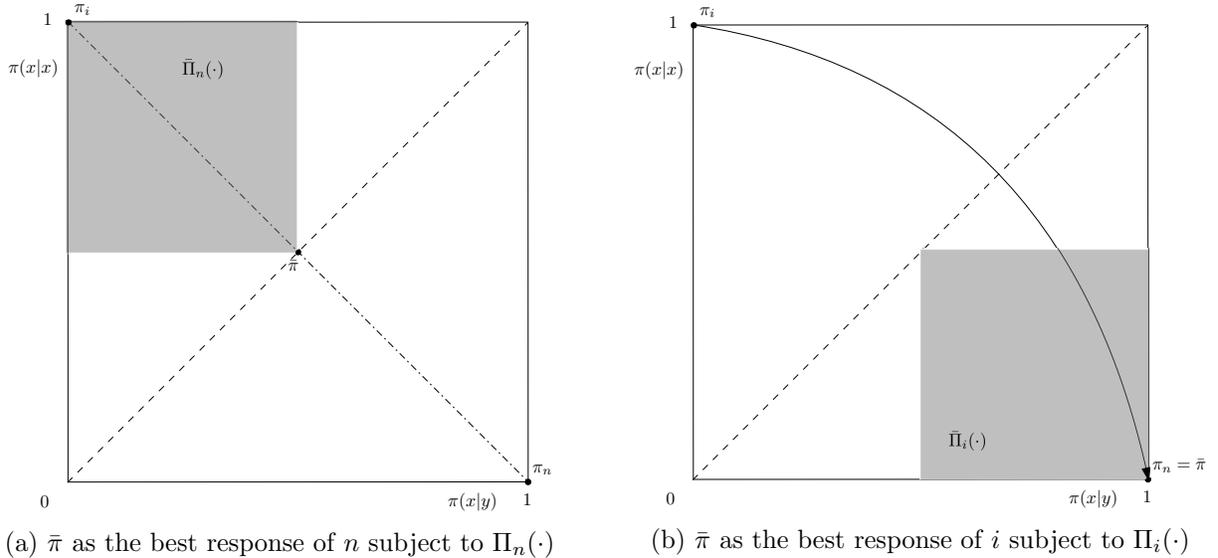


Figure 2: Example of a game with no equilibria: two equally prevalent senders with opposite preferences

¹³See the appendix for a detailed discussion

The previous example illustrates that existence is compromised when senders choose signals with a correlation between messages and states that has the opposite sign. The following proposition gives a sufficient condition for equilibrium to exist in terms of the unrestricted monopolist benchmark as well.

Proposition 7 (Sufficient Condition for the Existence Equilibria) *For each $s \in S$ let Π_s^{KG} be the set of optimal signals in the unrestricted monopolist benchmark. An equilibrium exists if for some selection of signals: $\{\pi_s^{KG}\}_{s \in S}$ where $\pi_s^{KG} \in \Pi_s^{KG}$ for all $s \in S$ we have that*

$$\bar{\pi}_{-s}^{KG}(m|\theta) > \bar{\pi}_{-s}^{KG}(m|\theta') \Rightarrow \exists \pi_s^* \in BR(\bar{\pi}_{-s}^{KG}) \text{ s.t. } \pi_s^*(m|\theta) \geq \pi_s^*(m|\theta') \quad \forall s \in S$$

where $BR(\bar{\pi}_{-s}^{KG})$ is the set of best responses for sender s in the constrained monopolist problem.

In light of this result we can guarantee existence by restricting senders to use signals with the same correlation between messages and states as the rest the other senders. Such signals form a half-space of Π and imposing such a restriction essentially limits the ability for senders to make $\bar{\pi}$ less informative than $\bar{\pi}_{-s}$.

Corollary 1 (Restriction to Guarantee Existence) *Let*

$$\Pi^*(\bar{\pi}_{-s}) = \{\pi \in \Pi : \bar{\pi}_{-s}(m|\theta) > \bar{\pi}_{-s}(m|\theta') \Rightarrow \pi(m|\theta) \geq \pi(m|\theta')\}$$

be the half-space of **message-consistent signals** for s , and let

$$\bar{\Pi}_s^*(\bar{\pi}_{-s}|\nu_0) = \nu_0(s)\Pi^*(\bar{\pi}_{-s}) + (1 - \nu_0(s))\bar{\pi}_{-s}.$$

If senders are restricted to choose $\pi_s \in \Pi^*(\bar{\pi}_{-s})$, or analogously to chose $\bar{\pi} \in \bar{\Pi}_s^*(\bar{\pi}_{-s}|\nu_0)$ an equilibrium exists.

The set $\Pi^*(\bar{\pi}_{-s})$ is not empty. In fact if $\underline{\pi} \in \Pi$ is uninformative then $\underline{\pi}(m|\theta) = \underline{\pi}(m|\theta')$ for all $m \in M$ and all $\theta \in \Theta$, so $\underline{\pi} \in \Pi^*(\bar{\pi}_{-s})$. Informative signals are also in the set. For example, $\bar{\pi}_{-s} \in \Pi^*(\bar{\pi}_{-s})$ and the set contains signals that perfectly reveal one or more states of the world (up to $\min\{|M|, |\Theta|\}$ states). In fact, if a sender was the unrestricted monopolist of information ($S = \{s\}$) but was restricted to choose a signal in $\Pi^*(\pi_0)$ for some $\pi_0 \in \Pi$, the restriction will not affect the sender's optimal payoff regardless of π_0 .

In the examples presented in the next section, all equilibria satisfy the above restriction¹⁴. However, this is not a tight condition in the sense that there are games for which equilibria exist without imposing any restriction on signals, but such equilibria would not use signals that are message-consistent¹⁵. However, any tighter bound would have to still limit the extent to which any sender that wants to make $\bar{\pi}$ less informative than $\bar{\pi}_{-s}$ can.

¹⁴In the general theorems of the applications I do use this restriction in order to be more general with respect to preferences and priors

¹⁵See existence discussion in the appendix

6 Applications

In this section I study equilibria in three important classes of games. I use these applications to derive general insights about the role of the uncertainty faced by receiver regarding the quality of messages¹⁶ on the informativeness of equilibria and on the incentives for senders to provide information. I first explore the consequences of increasing the prevalence of a noisy sender. That is a sender whose first-best is for $\bar{\pi}$ to be uninformative. Then, I study the effect of increasing the competition between senders with state independent preferences and bias towards different actions. Lastly, I relax the assumptions on the message space to enable senders to disclose their bias and study the effectiveness of such voluntary disclosures. I illustrate the results obtained with examples in the binary benchmark defined in example 1.

6.1 Games with a noisy sender

In many situations, a sender has no incentives to provide information to the receiver I call this a **noisy sender** since the signal that maximizes their expected payoff is simply noise. In the unrestricted monopolist benchmark, a sender is noisy if and only if the sender's payoff, $E_{\mu}u_s(a^*(\mu), \theta)$, is concave at μ_0 . This result relies on the fact that the sender can induce any Bayesian-Plausible distribution over posterior. However, when the receiver faces uncertainty about the quality of messages, a sender's first best might be to induce an informative signal, but such a signal might not be in the set of feasible average signals for the a sender. We see that the non-concavity of $E_{\mu}u_s(a^*(\mu), \theta)$ at the prior is necessary, but no longer a sufficient condition for a sender to benefit from persuasion. In the next example, I illustrate how in the presence of a noisy sender – even with relatively low prevalence – all other sender may find optimal to not provide information regardless of the concavity of their payoff function.

Consider the existence of a noisy sender, say $n \in S$ with $E_{\mu}u_n(a^*(\mu), \theta)$ is concave at μ_0 . To guarantee existence we impose the constraint that signals are message-consistent and we explore the effect of increasing the prevalence of n while the relative prevalences of all other senders remains constant. Say that the prevalence of n **increases uniformly** from $\nu_0(n)$ to $\nu'_0(n)$ if $\nu_0(n) < \nu'_0(n)$ and $\frac{\nu_0(s)}{1-\nu_0(n)} = \frac{\nu'_0(s)}{1-\nu'_0(n)}$ for all $s \neq n$.

Theorem 1 (The Effects of Increasing the Prevalence of a Noisy Sender) *Let $n \in S$ be a noisy sender ($E_{\mu}u_n(\cdot)$ is concave at μ_0). Restrict all senders to use message-consistent signals. Then there exists a threshold $\gamma(\mu_0, u_r)$ such that:*

- *If $\nu_0(n) < \gamma(\cdot)$ and $\nu_0(n)$ increases uniformly, then (for any continuous selection of equilibria) π_s^* increases for all $s \neq n$, but $\bar{\pi}$ decreases (both with respect to the order \succ_{bw}).*
- *If $\nu_0(n) \geq \gamma(\cdot)$, then $\bar{\pi}$ provides no valuable information; i.e. $a^*(\mu_1(\bar{\pi})) \equiv a^*(\mu_0)$.*

¹⁶As measured by their correlation with each state of the world

The threshold is given by

$$\gamma(\mu_0, u_r) = \max_i \left[\frac{1}{|\Theta|} \left(\frac{q(\theta_i) - \mu_0(\theta_i)}{(1 - q(\theta_i))\mu_0(\theta_i)} \right) + 1 \right]^{-1}$$

where $q(\theta_i)$ is the infimum belief about the state being θ_i that persuades the receiver to choose an action different than the default. That is

$$q(\theta_i) = \operatorname{arginf}_{\mu(\theta)} \{ \mu(\theta_i) > \mu_0(\theta_i) ; \frac{\mu(\theta_j)}{1 - \mu(\theta_i)} = \frac{\mu_0(\theta_j)}{1 - \mu_0(\theta_i)} \forall \theta_j \neq \theta_i, \text{ and } a^*(\mu) \neq a^*(\mu_0) \}.$$

Note that the threshold does not depend on the number of senders. Rather, it is inversely related to the distance between the prior and the closest belief that induces the receiver to choose a different action. For example, if the receiver requires a lot of information to be persuaded to choose an action different than the default, $\gamma(\cdot)$ is small and even moderate prevalences of the noisy sender preclude the equilibrium from being informative. The following example illustrates this result in the binary benchmark.

Example 2 (The binary election game: a noisy sender) *Consider the binary election game described in the introduction. That is $\Theta = A = M = \{x, y\}$, and the receiver has preferences $u_r = \mathbf{1}\{a = \theta\}$ so that $a^*(\mu_1) = y$ iff $\mu_1(y) \geq 0.5$.*

Assume the prior is $\mu_0(x) = 0.95$. There are two senders $S = \{p, n\}$. One is a partisan journalist (p), and a noisy journalist (n). The partisan journalist wants candidate y to be elected regardless of the state of the world, $u_p = \mathbf{1}\{a = y\}$ while the bad journalist wants the voter to elect the wrong candidate, $u_b = -u_r$. The prevalence of the noisy journalist is $\nu_0(n) = \alpha_n$.

Note first that in this game, the threshold described in theorem 1 is $\gamma(\cdot) = 0.1$; i.e if $\alpha_n \geq 10\%$, then every equilibria provides no valuable information. That is, even if the partisan sender, which has incentives to provide information, is 90% prevalent, the receiver learns nothing in equilibrium. If $\alpha_n < 0.1$, there is an equilibrium where the bad journalist recommends either candidate with equal probability, $(\pi_n^*(x|x), \pi_n^*(x|y)) = (1/2, 1/2)$, and the partisan journalist chooses $(\pi_p^*(x|x), \pi_p^*(x|y)) = \left(\frac{4(2-\alpha_n)}{9(1-\alpha_n)}, 0 \right)$, whose informativeness increases in α_n ¹⁷. Lastly, the equilibrium average signal is $(\bar{\pi}(x|x), \bar{\pi}(x|y)) = \left(\frac{8+\alpha_n}{9}, \frac{\alpha_n}{2} \right)$, and its informativeness decreases in α_n . Figure 3 depicts an equilibrium for $\alpha_n = 5\%$ and Figure 4 for $\alpha_n = 15\%$. The set $\bar{\Pi}_p(\cdot)$ is shaded in grey in panel (a) and $\bar{\Pi}_n(\cdot)$ is shaded grey in panel (b). While the regions shaded in red contain signals that convince the receiver to choose y ¹⁸ with some message, in the white region, both messages induce the default choice, x . Therefore, the partisan journalist wants to induce an average signal that is lays in a red-shaded region and maximizes the probability of sending a message that induces action y . When $\alpha_b = 5\%$, such an average signal is feasible and so is chosen. The bad journalist

¹⁷Note that the perfectly informative signal is $(\pi(x|x), \pi(x|y)) = (1, 0)$ and $\frac{4(2-\alpha_n)}{9(1-\alpha_n)}$ is increasing in α_n .

¹⁸In the top triangle message y induces action y and message x induce action x . In the bottom triangle the reverse is true, message x induces action y and message y induce action x .

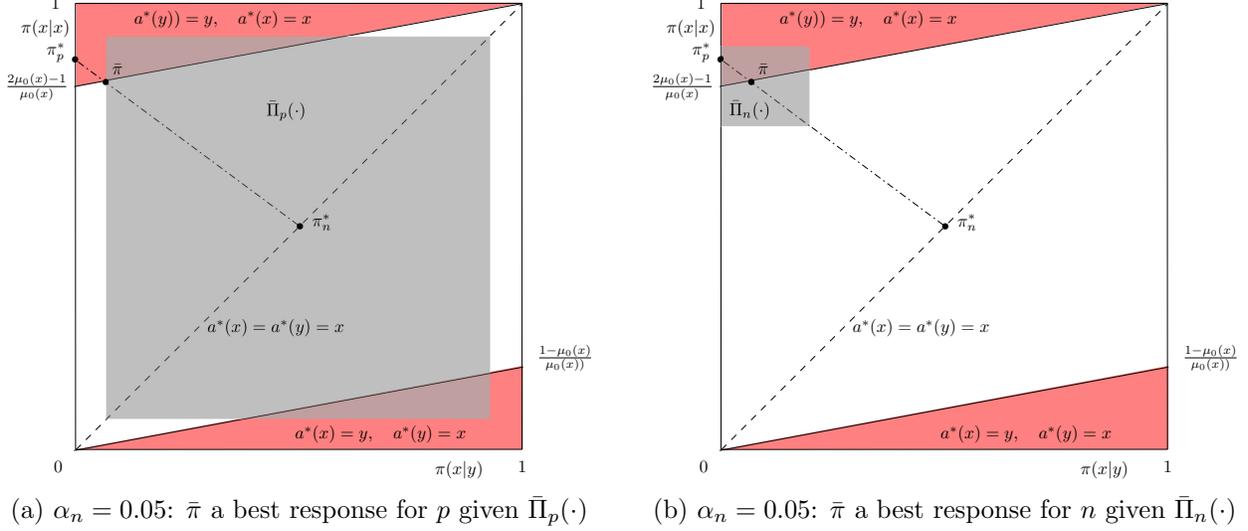


Figure 3: Equilibrium with $\alpha_n = 0.05$, $\bar{\pi}$ is a mutual best response given the constraint each sender faces.

has no profitable deviation since he is indifferent among all the expected signals that are not in the interior of the red-shaded areas.

When $\alpha_n = 15\%$, no average signal that is feasible to p intersects with the red-shaded regions, so p has no incentives to provide information, and the receiver always chooses x .

◇

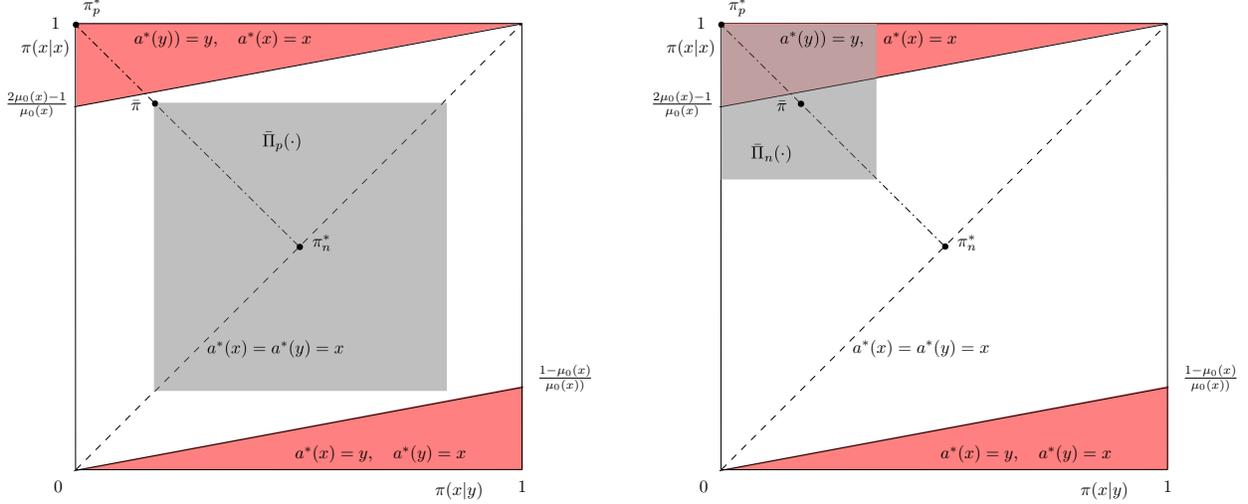
6.2 Multiple senders with heterogeneous biases

A common benchmark in the literature is to assume that sender's preferences do not depend on the true state of the world. For example if senders are sellers, they might want the buyer to buy their good regardless of its true quality. Politicians typically want the electorate to vote for their candidate regardless, etc. In this section, I assume senders with state-independent preferences and heterogeneous biases. In particular I assume that for each sender, s , there is an action, $a_s \in A$, and a state, $\theta_s \in \Theta$, such that $u_s = \mathbf{1}\{a = a_s\}$ and $a^*(\mu) = a_s$ iff $\mu(\theta_s) \geq q_s > \mu_0(\theta_s)$ for some q_s ¹⁹. For games that satisfy these assumptions we say that **senders have biases in different directions**.

By studying games within this class, I show that increasing the number of senders with heterogeneous preferences decreases the informativeness of the average signal. Albeit is true that increasing the number of senders implies that $\bar{\pi}$ is the average over a larger number of signals (i.e. the receiver is more uncertain about the source of the message), the result is that $\bar{\pi}$ is less informative because every sender chooses a less informative signal π_s^* than the equilibrium with less senders.

Theorem 2 (The Effects of Increasing the Variety of Sender-Biases) *Consider a game with senders that have biases in different directions. Let $\{\pi_s^*\}$ be an equilibrium profile, then:*

¹⁹The assumption $q_s > \mu_0(\theta_s)$ guarantees that all senders have incentives to provide information, and so the sufficient condition in proposition 7 is satisfied. The result in theorem 2, however, is robust to imposing instead the restrictions in corollary 1 to have existence and allowing $q_s = \mu_0(\theta_s)$, $\exists s \in S$.



(a) $\alpha_n = 0.15$: $\bar{\pi}$ a best response for p given $\bar{\Pi}_p(\cdot)$ (b) $\alpha_n = 0.15$: $\bar{\pi}$ a best response for n given $\bar{\Pi}_n(\cdot)$

Figure 4: Equilibrium with $\alpha_n = 0.15$, $\bar{\pi}$ is a mutual best response. The receiver always chooses x .

- For every sender, $\pi_s^* \succsim_{bw} \pi_s^{KG}$, $\exists \pi_s^{KG} \in \Pi_s^{KG}$.
- $\bar{\pi}$ and each π_s^* are decreasing (with respect to \succsim_{bw}) in the number of heterogeneous senders. Where if $S' = S \cup \{s'\}$, then $\nu'_0(s') > 0$ and $\nu'_0(s) = (1 - \nu'_0(s))\nu_0(s)$ for all $s \neq s'$.

The intuition behind this result lies in the fact that if the receiver faced no uncertainty about the message's quality, each sender's optimal signal, say π_s^* , is one that induces a posterior that partially obfuscates information. In particular, π_s^* induces a degenerate posterior and a posterior that makes the receiver indifferent between the default action and the action preferred by the sender. The receiver knows that when she receives a recommendation to not choose the sender's preferred action, it is a very informative signal²⁰. When the receiver faces uncertainty about the message's quality, a recommendation to choose a_s can be highly informative if coming from sender $s' \neq s$ that is not biased towards that action, or not very informative if coming from sender s . Taking this into account, each sender can pool with other senders and recommend their preferred action when the state is not favorable even more frequently than in the unrestricted monopolist benchmark (i.e. KG11's game). In equilibrium, all senders choose less informative signals and the average signal is also less informative.

Example 3 (The binary election game: multiple partisan senders) Consider a slight modification to the binary election game described in the introduction. In particular $\Theta = M = \{x, y\}$, but $A = \{x, y, \emptyset\}$.²¹ The receiver's preferences are $u_r = \mathbf{1}\{a = \theta\} + (\frac{9}{10}) \mathbf{1}\{a = \emptyset\}$, so $a^*(\mu) = x$ only if $\mu(\theta = x) \geq 0.9$ similarly $a^*(\mu) = y$ only if $\mu(\theta = y) \geq 0.9$, otherwise $a^*(\mu) = \emptyset$.

²⁰A sender s only recommends an action $a_{s'} \neq a_s$ to relax the Bayes plausibility constraint and so it fully reveals that the state is $\theta_{s'}$

²¹For this game, it is without loss to have $|M| < |A|$, see an argument in the appendix, I use the smaller message set to simplify exposition.

Assume the prior is $\mu_0(\theta = x) = 0.5$ and there are two partisan senders $S = \{x, y\}$ such that $u_s = \mathbf{1}\{a = s\}$ for all $s \in S$. Let both senders to be equally prevalent, i.e. $\nu_0(x) = 0.5$.

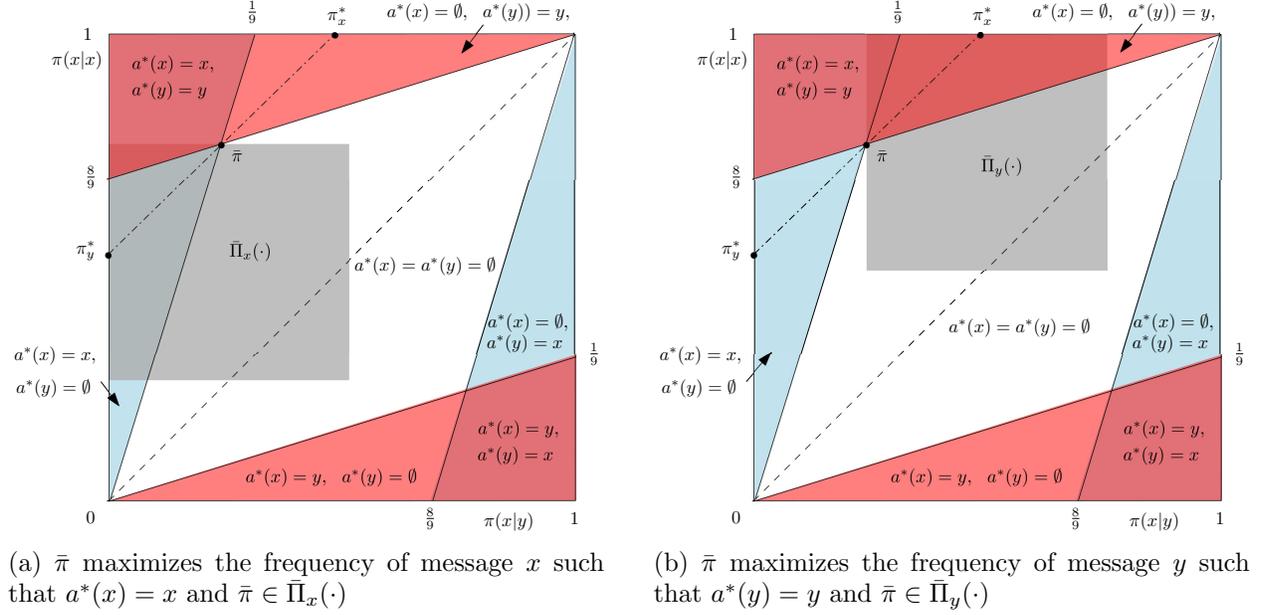


Figure 5: Equilibrium with 2 senders with opposite biases.

Consider first the straight-forward optimal signals in the unrestricted monopoly benchmark for each sender.

$$(\pi_x^{KG}(x|x), \pi_x^{KG}(x|y)) = \left(1, \frac{1}{9}\right), \quad (\pi_y^{KG}(x|x), \pi_y^{KG}(x|y)) = \left(\frac{8}{9}, 0\right),$$

In contrast, the optimal straight-forward signal in the game with uncertainty are:

$$(\pi_x^*(x|x), \pi_x^*(x|y)) = \left(1, \frac{1}{5}\right), \quad (\pi_y^*(x|x), \pi_y^*(x|y)) = \left(\frac{4}{5}, 0\right)$$

note that $\pi_s^* \succ_{bw} \pi_s^{KG}$ for all $s \in \{x, y\}$. Furthermore the average signal is

$$(\bar{\pi}^*(x|x), \bar{\pi}^*(x|y)) = \left(\frac{9}{10}, \frac{1}{10}\right)$$

That is $\bar{\pi}^*$ induces the posteriors $\mu_1(x|x) = 0.9$ and $\mu_1(x|y) = 0.1$, while π_x^{KG} induces posteriors $\mu_1(\theta = x) = 0.9$ and $\mu_1(\theta = y) = 0$. Similarly for y since π_y^{KG} induces posteriors $\mu_1(\theta = x) = 1$ and $\mu_1(\theta = y) = 0.1$. Indeed $\bar{\pi}^* \succ_{bw} \pi_s^{KG}$ for all $s \in \{x, y\}$ illustrating how the average signal also decreases in informativeness when the number of senders increases. Figure 5 illustrates the equilibria. Shaded in red are all the signals where one of the messages induces action y , in blue are the signals where some message induces x , the equilibrium is in the intersection of these two regions. Shaded in grey are the sets $\bar{\pi}_s(\bar{\pi}_s, \nu_0(s))$ for each sender. Panel (a) shows the best response for sender x , while panel (b) shows y 's best response.

6.3 Voluntary Bias Disclosures

In this section I explore the consequences of allowing senders to credibly disclose their preferences. In many platforms online it is common that senders can post information using an account with a verified identity or posting anonymously²². It is key here for identities to be verifiable but, in many cases, seeing a verified identity does not preclude the receiver from being uncertain about the quality of messages. Lazer *et. al* (2018) document how, in major social media platforms, accounts often mimic in appearance reputable sources to gain credibility. If every sender can mimic each others communication strategies, the situation is well approximated by assuming that senders have access to the same message space²³. As I argue in section 5, semi-pooling equilibria is pervasive in such a setting, hence receiver’s uncertainty about quality of messages is as well.

To study the value of voluntary disclosures, I modify slightly the model by enriching each sender’s message space to $M_s = M \cup M_s^*$. For each sender, the message space includes now a common part, M , and an exclusive part, M_s^* , where all $M_s^* \cap M_{s'}^* = \emptyset$ and $M_s^* \cap M = \emptyset$ for all $s \in S$. Again, the cardinality of each of these pair-wise disjoint sets is assumed to be rich²⁴. When these assumptions hold we say that **senders can disclose their biases**.

With this modification, senders face a non-trivial trade-off when deciding whether to use messages in M or in M_s^* . On the one hand, by only using messages in M_s^* the sender can induce any Bayes-plausible distribution over state-posteriors with his exclusive messages. That is, with probability $\nu_0(s)$, s achieves his first-best payoff, but has no control on the payoff realized with complementary probability. On the other hand, by using messages in M and partially pooling with other senders, the sender has partial control over the posteriors induced even when the receiver gets the message from one of these other senders. Though in general pooling gives a lower payoff (since it is a restricted optimization), such payoff is realized with a higher probability.

Theorem 3 (The Value of Voluntary Disclosures) *Consider a game where senders can disclose their biases; i.e. $M_s = M \cup M_s^*$ for all $s \in S$, as previously defined, then:*

- *A separating equilibrium always exists. That is $\text{sup}(\pi_s^*) \cap \text{sup}(\pi_{s'}^*) = \emptyset$ for all $s \neq s'$.*
- *If senders have biases in different directions (as in theorem 2), almost every equilibrium is a separating equilibrium (for generic μ_0).*
- *If senders have biases in different directions, the average signal for any equilibrium where senders can disclose their biases ($\bar{\pi}_{M \cup M_s^*}^*$) is more informative than any equilibrium when*

²²In the late 2000’s and early 2010’s major social media platforms enabled content producers to verify their identity. For example, Twitter adopted this policy in 2009, followed by other platforms like Google+, and Facebook in 2011 and 2012.

²³In recent years, the option to verify your identity has been extended to more than just famous news outlets, public figures or well-known institutions, but essentially to every individual that chooses to verify their account.

²⁴It is a direct result of proposition 1 to assume $|M_s^*| = |A|$ as well.

they cannot ($\bar{\pi}_M^*$):

$$\bar{\pi}_M^* \succsim_{bw} \bar{\pi}_{M \cup M_s^*}^*$$

The first part of the theorem implies that regardless of sender's preferences or prior beliefs, when senders can disclose their biases, a separating equilibrium always exists. In such an equilibrium the receiver faces no uncertainty regarding the quality of messages. The intuition for the second statement in the previous theorem is that when senders have bias in different directions, each sender's payoff is equal to the probability that their favorite action is chosen by the receiver. Since the probabilities with which actions are played are bound to add up to the unity, senders face a zero-sum game. Therefore, for generic priors, if we start assuming that all senders use messages in M , there is always a sender that finds profitable to separate (compared to the equilibrium payoff achieved by pooling). Then, if we assume that the remaining senders are pooling, we find that another sender finds profitable to separate, etc. The uncertainty about S unravels, and only separating equilibria survive. Though in general the receiver does not learn the state θ , the third statement states that the equilibrium average signal is more informative. The following example illustrates.

Example 4 (The binary election game: voluntary disclosures) Consider the game presented in example 3, however let $M_x = \{x, y, x^*, y^*\}$ and $M_y = \{x, y, x_*, y_*\}$ and let $\mu_0(\theta = x) = 0.6^{25}$.

It is immediate to conclude that a separating equilibrium exists. Consider the case where each sender only uses their exclusive messages to induce their first-best distribution of posterior, then each sender will choose:

$$\pi_x^*(x^*|x) = 1; \quad \pi_x^*(x^*|y) = \frac{1}{6}, \quad \pi_x^*(y^*|y) = \frac{5}{6}$$

$$\pi_y^*(x_*|x) = \frac{25}{27}; \quad \pi_y^*(y_*|x) = \frac{2}{27}, \quad \pi_y^*(y_*|y) = 1$$

Note that in this equilibrium, the induced posteriors are $\mu_1(\theta = x) \in \{0, 0.1, 0.9, 1\}$. In contrast, if both senders have only access to messages $M = \{x, y\}$, the equilibrium signals with straightforward messages are $(\pi_{xM}(x|x), \pi_{xM}(x|y)) = (1, \frac{5}{16})$, $(\pi_{yM}(x|x), \pi_{yM}(x|y)) = (\frac{7}{8}, 0)$, the average signal is $(\bar{\pi}_M(x|x), \bar{\pi}_M(x|y)) = (\frac{15}{16}, \frac{5}{32})$, and the induced posteriors are again $\mu_1(x|x) = 0.9$ and $\mu_1(x|y) = 0.1$. Therefore, $\bar{\pi}_M \succsim \bar{\pi}_{M \cup M_s^*}^*$.

Lastly, I illustrate that the semi-pooling equilibria above with the restriction that $M_s = M$ is not an equilibrium when $M_s = M \cup M_s^*$. In the distribution $\tau(\bar{\pi}_M)$ each posterior, $\mu_1(\theta = x) \in \{0.1, 0.9\}$ is induced with equal probability. So the payoff for each sender is 0.5. However, sender x finds profitable to deviate and only pool with y on the signal that induces his favorite action:

$$\pi'_x(x^*|x) = 1; \quad \pi'_x(x^*|y) = \frac{1}{6}, \quad \pi'_x(x|y) = \frac{7}{48}, \quad \pi'_x(y^*|y) = \frac{33}{48}$$

²⁵ $\mu_0(\theta = x) = 0.5$ is the knife-edge case where senders are indifferent between separating or pooling, so I perturb it slightly.

Given π'_x and π_{yM} we have that $\mu_1(x|x) = \mu_1(x|x^*) = 0.9$ i.e. the receiver chooses sender x 's favorite action after any of these messages, while $\mu_1(x|y^*) = 0$ (inducing action y) and $\mu_1(x|y) = \frac{3}{19}$ (inducing action \emptyset). The payoff for x with that strategy is $\frac{5}{8} > 0.5$.

◇

7 Conclusion

This paper explores how the endogenous uncertainty about the precision of signals affects the incentives to provide information. In equilibrium, senders pool their signals, so the receiver updates her beliefs using the average correlation between messages and states of the world. Since each sender chooses only one of the averaged signals, they have a partial influence on the posterior beliefs of the receiver. The paper provides a geometric characterization of this new constraint and reformulates the problem of each sender as one of a standard information designer, but with an additional constraint. This linear constraint is endogenous and determined by the actions of all other senders. Therefore, uncertainty limits the value of providing information, but when providing information remains valuable, it shapes how much information senders provide. In particular, it can incentivize senders to choose more informative signals than the single information designer benchmark.

We apply the framework in three broad classes of games to inform policies aimed to improve the quality of the information in platforms such as social media. In particular, the increase in the prevalence of a noisy sender has a disproportionately negative effect. While such an increase incentivizes other senders to provide more precise messages, the direct effect of the noisy sender dominates, and the receiver learns less in equilibrium. Furthermore, there exists a non-trivial threshold so that if the prevalence of the noisy sender surpasses this threshold, the value of providing information vanishes for all senders, so no relevant information is communicated. As a consequence, if information platforms - such as social media - wish to remain influential, editorial work must be done to limit the prevalence of noisy signals that emulate reputable sources.

The paper also showcases how increasing the variety of senders with biases in different directions can lead to worse outcomes. If senders have state-independent payoffs that are positive whenever the receiver chooses the action they are biased towards, then increasing the variety of senders' biases is detrimental. As the number of competitor senders increases, each sender produces less information, so the receiver learns less as well.

Finally, the model characterizes the value of voluntary disclosures of biases. When senders can credibly disclose their biases, a separating equilibrium always exists. Further, if the game between senders is a zero-sum game, the uncertainty about the quality of information unravels, and the receiver is guaranteed to learn more about the state of the world as well. This result implies that Social media platforms can improve the quality of information by enabling sources to credibly disclose their biases without having to directly assess the quality of the information each sender provides. Current efforts of social media to verify the identity of content producers can be enhanced by creating a non-falsifiable record of their biases, such as financial or political sponsorships.

References

- Ambrus, A., & Takahashi, S. (2008). Multi-sender cheap talk with restricted state spaces. *Theoretical Economics*, 3(1), 1-27.
- Battaglini, M. (2002). Multiple referrals and multidimensional cheap talk. *Econometrica*, 70(4), 1379-1401.
- Bergemann, D., & Morris, S. (2019). Information design: A unified perspective. *Journal of Economic Literature*, 57(1), 44-95.
- Chakraborty, A., & Harbaugh, R. (2010). Persuasion by cheap talk. *American Economic Review*, 100(5), 2361-82.
- Crawford, V. P., & Sobel, J. (1982). Strategic information transmission. *Econometrica: Journal of the Econometric Society*, 1431-1451.
- Dimitrakas, V., & Sarafidis, Y. (2005). Advice from an expert with unknown motives. Available at SSRN 797844.
- Gentzkow, M., & Kamenica, E. (2017a). Competition in persuasion. *The Review of Economic Studies*, 84(1), 300-322.
- Gentzkow, M., & Kamenica, E. (2017b). Bayesian persuasion with multiple senders and rich signal spaces. *Games and Economic Behavior*, 104, 411-429.
- Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics*, 24(3), 461-483.
- Hedlund, J. (2017). Bayesian persuasion by a privately informed sender. *Journal of Economic Theory*, 167, 229-268.
- Jain, V. (2018). Bayesian persuasion with cheap talk. *Economics Letters*, 170, 91-95.
- Kamenica, E., & Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6), 2590-2615.
- Kosenko, A. (2018). Bayesian persuasion with private information. Available at SSRN 3236191
- Krishna, V., & Morgan, J. (2001a). Asymmetric Information and Legislative Rules: Some

Amendments. *American Political Science Review*, 95(2).

Krishna, V., & Morgan, J. (2001b). A model of expertise. *The Quarterly Journal of Economics*, 116(2), 747-775.

Lazer, D. M., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., ... & Schudson, M. (2018). The science of fake news. *Science*, 359(6380), 1094-1096.

Le Treust, M., & Tomala, T. (2019). Persuasion with limited communication capacity. *Journal of Economic Theory*, 184, 104940.

Li, M., & Madarasz, K. (2008). When mandatory disclosure hurts: Expert advice and conflicting interests. *Journal of Economic Theory*, 139(1), 47-74.

Li, F., & Norman, P. (2018). On Bayesian persuasion with multiple senders. *Economics Letters*, 66-70

Lipnowski, E., & Mathevet, L. (2017). Simplifying Bayesian Persuasion. mimeo.

Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 380-391.

Milgrom, P., & Roberts, J. (1986). Relying on the information of interested parties. *The RAND Journal of Economics*, 18-32.

Min, D. (2017). Bayesian Persuasion under Partial Commitment. Work. Pap., Univ. Ariz., Tucson

Morgan, J., & Stocken, P. C. (2003). An analysis of stock recommendations. *RAND Journal of economics*, 183-203.

Fréchette, G. R., Lizzeri, A., & Perego, J. (2019). Rules and commitment in communication.

Perez-Richet, E. (2014). Interim bayesian persuasion: First steps. *American Economic Review*, 104(5), 469-74.

Piermont, E. (2016). Distributional uncertainty and persuasion. Working Paper.

Shin, H. S. (1994). The burden of proof in a game of persuasion. *Journal of Economic Theory*, 64(1), 253-264.

Tsakas, E., & Tsakas, N. (2018). Noisy persuasion. Available at SSRN 2940681.

Wolinsky, A. (2003). Information transmission when the sender's preferences are uncertain. *Games and Economic Behavior*, 42(2), 319-326.

8 Appendix

Proof of proposition 1

Claim: For any equilibrium with signals supported in M such that $|A| \leq |M|$ there is an outcome equivalent equilibrium in the game where senders can only use some message space, M' with cardinality $|M'| \leq |A|$.

Proof: Let $\{\pi_s\}_{s \in S}$ be an equilibrium profile of signals in the game with message space M . Define $M(a) = \{m \in M | a^*(\mu_1(m)) = a\}$ for every $a \in A$. Let $M' \subseteq M$ be the collection of messages for which $M(a)$ is not empty. Clearly, $|M'| \leq |A|$. Now, define a set of signals $\{\pi'_s\}_{s \in S}$ with support in M' such that $\pi'_s(m_a | \theta) = \sum_{m \in M(a)} \pi_s(m | \theta)$ for all $m_a \in M'$, all $\theta \in \Theta$ and for each $s \in S$. The posterior induced by, $\{\pi'_s\}_{s \in S}$ after message m_a is in the convex hull of the posteriors induced by $\{\pi_s\}_{s \in S}$ after each of the messages in $M(a)$. Since expected utility is linear in beliefs, action a is also a best response to $\mu_1(m_a)$, so we can assume $a'(\mu_1(m_a)) = a$ for all $a \in A$. Therefore, the distribution over actions is the same under the original and the new set of signals. If we modify the game so that the message space is M' , signals $\{\pi'_s\}_{s \in S}$ along with the response function $a'(\cdot)$ and bayesian beliefs $p_1(\{\pi'_s\}_{s \in S})$ constitute an equilibrium. To see this, note that all messages in M' are used, so any profitable deviation of the game with M' messages would be a profitable deviation in the game with M messages, a contradiction.

Claim: For any equilibrium with signals, supported in a set M' , there is an outcome equivalent equilibrium in the game where senders can use a finite message space M^* with $|M'| \leq |M^*|$.

Proof: Since $|M'| \leq |M^*|$, there exists a partition of M^* with $|M'|$ elements. Denote each element in the partition by $M^*_{m'} \subset M^*$. Since M^* is finite, so is $M^*_{m'}$ for every $m' \in M'$. Construct signals $\{\pi^*_s\}$ supported on M^* as follows: $\pi^*_s(m^* | \theta) = \frac{\pi_s(m' | \theta)}{|M^*_{m'}|}$ for all $m^* \in M^*_{m'}$, $\theta \in \Theta$ and for each $m' \in M'$ and $s \in S$. That is, each sender distributes uniformly the mass chosen for each message $m' \in M'$ among all the messages in $M^*_{m'}$ for each state $\theta \in \Theta$.

The distribution over posteriors induced by $\{\pi^*_s\}$ will be identical than that induced by $\{\pi'_s\}$. Thus the distribution over actions will also be the same. Since there are no un-used messages, any profitable deviation of the game with M^* messages would be a profitable deviation in the game with M' messages, thus there are none.

We conclude that any equilibrium of a game with M messages such that $|M| \geq |A|$ is outcome equivalent to a game with $|M'|$ messages, for some set with $|M'| \leq |A|$. By letting $M^* = A$, the second claim implies that such an equilibrium is outcome equivalent to an equilibrium of a game with $|A|$ messages available, thus completing the "only if" part of the statement. The "if"

part of the statement follows directly by using the second claim again, letting $A = M'$ and $M = M^*$.

Proof of lemma 1

Suppose $\{\pi_s\}_{s \in S}$ is an equilibrium and $m \neq m'$ are in its support with $\mu_1(m) = \mu_1(m')$, but $a^*(p_1(\theta, s|m)) \neq a^*(p_1(\theta, s|m'))$ and $E_{p_1(\cdot|m)} u_s(a^*(p_1(\theta, s|m)), \theta) > E_{p_1(\cdot|m')} u_s(a^*(p_1(\theta, s|m')), \theta)$ for some $s \in S$ with $\{m, m'\} \cap \text{support}(\pi_s) \neq \emptyset$.

Since $a^*(\cdot)$ is a well defined function, $p_1(\theta, s|m) \neq p_1(\theta, s|m')$. However, the receiver's utility does not depend on s , and $\mu_1(m) = \mu_1(m')$, thus the receiver is indifferent between $a_m := a^*(p_1(\theta, s|m))$ and $a_{m'} := a^*(p_1(\theta, s|m'))$. From continuity, sender s can choose a signal π'_s that shifts the posteriors from $p_1(\cdot|m')$ and $p_1(\cdot|m)$, to $p'_1(\cdot|m)$, $p'_1(\cdot|m')$ so that a_m is the unique best response to both messages and $E_{\tau(\{\{\pi_{s'}\}_{s' \neq s} \cup \pi'_s\})} E_{p_1} u_s(a^*(p_1), \theta) > E_{\tau(\pi_s)_{s \in S}} E_{p_1(\cdot|m)} u_s(a^*(p_1(\theta, s|m)), \theta)$.

Proof of Proposition 2

Let $\{\pi_s^*\}_{s \in S}$, $a^*(\cdot)$ and $p_1(\cdot)$ be an equilibrium. If $a^*(p_1(\theta, s|m)) = a^*(p_1(\theta, s|m'))$ for all pair of messages that satisfy that $\mu_1(m) = \mu_1(m')$, then there is a well defined function $\hat{a} : \Delta(\Theta) \rightarrow \Delta(A)$ with $a^*(p_1(\theta, s)(\{\pi_s\}_{s \in S})) \equiv \hat{a}(\mu_1(\theta)(\{\pi_s\}_{s \in S}))$, such that $\{\pi_s^*\}_{s \in S}$, $\hat{a}(\cdot)$ and $p_1(\cdot)$ define an outcome equivalent equilibrium.

If the previous assumption fails, i.e. there is a pair of messages that satisfy $\mu_1(m) = \mu_1(m')$, but $a^*(p_1(\theta, s|m)) \neq a^*(p_1(\theta, s|m'))$ define the mixed strategy $a' = \beta a^*(p_1(\theta, s|m)) + (1-\beta)a^*(p_1(\theta, s|m'))$ with $\beta = \frac{\sum_{\Theta} \sum_S \pi_s(m|\theta)}{\sum_{\Theta} \sum_S (\pi_s(m|\theta) + \pi_s(m'|\theta))}$ and let $\hat{a}(\mu_1(\cdot)) \equiv a^*(p_1(\cdot))$ except at $p_1(\theta, s|m)$ and $p_1(\theta, s|m')$ where $\hat{a}(\mu_1(\cdot)) = a'$. Repeat for all such conflicting pairs. From lemma 1, all the senders with m, m' in the support of their signals are indifferent between these two actions, so they are also indifferent between these actions and the mixed strategy. Thus these senders have no incentive to deviate. For the other senders, the outcome from $\{\pi_s^*\}_{s \in S}$, $\hat{a}(\cdot)$ and $p_1(\cdot)$ is outcome equivalent to the original equilibrium, so any deviation deviation to this triplet would be profitable deviation to $\{\pi_s^*\}_{s \in S}$, $a^*(\cdot)$ and $p_1(\cdot)$.

Proof of proposition 3

If a^* is regular, then $a^*(\mu_1(\theta)(\{\pi_s\}_{s \in S})) \in \Delta(A)$ so it depends on $\{\pi_s\}_{s \in S}$ only through $\mu_1(\cdot)$. Let $\{\pi_s\}_{s \in S}$ be any profile of senders' signals, then:

$$\begin{aligned} \mu_1(\theta|m) &= \sum_{s \in S} \frac{\pi_s(m|\theta) p_0(\theta, s)}{\sum_{\theta' \in \Theta} \sum_{s' \in S} \pi_{s'}(m|\theta') p_0(\theta', s')} \\ &= \frac{(\sum_{s \in S} \pi_s(m|\theta) \nu_0(s)) \mu_0(\theta)}{\sum_{\theta' \in \Theta} [\sum_{s' \in S} \pi_{s'}(m|\theta') \nu_0(s')] \mu_0(\theta')} \\ &= \frac{\bar{\pi}(m|\theta) \mu_0(\theta)}{\sum_{\theta' \in \Theta} \bar{\pi}(m|\theta') \mu_0(\theta')} \end{aligned}$$

Proof of lemma 2

Since $\nu_0(s) \in (0, 1)$ for all s , the contraction follow directly from equation 4. Further, this equation is linear in $\bar{\pi}_{-s}$ and π_s for all $\pi_s \in \Pi$. Therefore, the correspondence $\bar{\Pi}_s(\bar{\pi}_{-s}, \nu_0)$ is also continuous. Therefore, combining that the distances in the contraction are shrunk by $\nu_0(s)$ and continuity, it follows that $\nu_0(s) < \nu'_0(s)$ implies $\bar{\Pi}_s(\bar{\pi}_{-s}, \nu_0) \subset \bar{\Pi}_s(\bar{\pi}_{-s}, \nu'_0)$.

Lastly, let $\bar{\pi}_{-s} \in \Pi$ be any signal in Π and define $\pi_s(m|\Theta) = 1/|M|$ for all m, θ . Thus the signal $\bar{\pi}_l = \nu_0(s)\pi_s + (1 - \nu_0(s))\bar{\pi}_{-s}$ satisfies that for each message $m \in M : |\bar{\pi}_l(m|\theta) - \bar{\pi}_l(m|\theta')| = (1 - \nu_0(s))|\bar{\pi}_{-s}(m|\theta) - \bar{\pi}_{-s}(m|\theta')| \leq |\bar{\pi}_{-s}(m|\theta) - \bar{\pi}_{-s}(m|\theta')|$ for all pairs of states $\theta \in \Theta$, so $\bar{\pi}_l \preceq \bar{\pi}_{-s}$.

Similarly, define π'_s as follows: For each state θ let $m_\theta = \operatorname{argmax} \mu_1(\theta)(\bar{\pi}_{-s})$. Then $\pi'_s(m|\theta) = \mathbf{1}\{m = m_\theta\}$. Hence the signal $\bar{\pi}_h = \nu_0(s)\pi'_s + (1 - \nu_0(s))\bar{\pi}_{-s}$ satisfies that for each message $m \in M$ let $\alpha = |\bar{\pi}_{-s}(m|\theta)\mu_0(\theta) - \bar{\pi}_{-s}(m|\theta')\mu_0(\theta')| \in [0, 1]$, then $|\bar{\pi}_h(m|\theta)\mu_0(\theta) - \bar{\pi}_h(m|\theta')\mu_0(\theta')| \geq \alpha + \nu_0(s)(1 - \alpha) \geq \alpha = |\bar{\pi}_{-s}(m|\theta)\mu_0(\theta) - \bar{\pi}_{-s}(m|\theta')\mu_0(\theta')|$ Therefore, $\bar{\pi}_{-s} \preceq \bar{\pi}_h$.

Proof of proposition 4 If a^* is regular, then

$$E_{\tau_{\{ \pi_s \}}} E_{p_1} u_s(a^*(p_1), \theta) = E_{\tau_{\Theta}(\bar{\pi})} E_{\mu_1} u_s(\hat{a}(\mu_1), \theta)$$

if and only if $\bar{\pi} = \nu_0(s)\pi_s + (1 - \nu_0(s))\bar{\pi}_{-s}$ and

–

Proof of proposition 5

We will prove the contrapositive of the statement. Suppose there exists a pair of senders $s, s' \in S$ such that $\mathcal{T}_s^{KG} \cap \mathcal{T}_{s'}^{KG} = \emptyset$. In a separating equilibrium, the receiver learns the identity of the receiver upon receiving a message and senders have full control over the receiver's beliefs, conditional on their message being seen by the receiver. Thus, each sender uses signal that induces an element in \mathcal{T}_s^{KG} since it contains all the distributions over posteriors that maximize the sender's objective.

Suppose the senders s, s' induce the distributions τ_s^{KG} and $\tau_{s'}^{KG}$ in equilibrium. Then either $\operatorname{support}(\tau_s^{KG}) \cap \operatorname{support}(\tau_{s'}^{KG}) \neq \emptyset$ or $\operatorname{support}(A^*(\tau_s^{KG})) \cap \operatorname{support}(A^*(\tau_{s'}^{KG})) \neq \emptyset$. Since $\mathcal{T}_s^{KG} \cap \mathcal{T}_{s'}^{KG} = \emptyset$, sender s strictly prefers outcome $A^*(\tau_s^{KG})$ over $A^*(\tau_{s'}^{KG})$ and similarly for sender s' . Consider each case separately.

If $\operatorname{support}(\tau_s^{KG}) \cap \operatorname{support}(\tau_{s'}^{KG}) \neq \emptyset$, then there must be a posterior belief

Proof of Theorem 1

Let t_0 be the sender whose payoff function $v_{t_0}(a, \theta)$ is concave, and consider an uninformative signal $\pi_{t_0}(m|\theta) = 1/|M|$ for all $m \in M$ and all states $\theta \in \Theta$. This sender uses all the messages with positive probability. In the limit when A_{t_0} converges uniformly to the identity matrix I , from the continuity of equation 4, $\bar{\pi} \rightarrow \pi_{t_0}$. Therefore, the receiver's posterior equals the prior, the receiver would then choose its default action, $a^*(\mu_0^\Theta)$. From the continuity of $\bar{\Pi}(\bar{\pi}_{-t}, A_t)$ for every $t \in \mathcal{T}$, for any $\epsilon > 0$ there exists a $\delta > 0$ such that for for any prior distribution, μ_0 satisfying $A_{t_0} \geq (1 - \delta)I$ the expected signal defined by the average of the signal π_{t_0} defined above and identical and perfectly informative signals π_t for all $t \neq t_0$, induces posteriors contained in an open ball centered at the

prior with radius ϵ . The existence of the uniform δ can be guaranteed given the compactness of Θ , \mathcal{T} , and using standard continuity arguments.

For generic preferences of the receiver, the default action is optimal for posteriors in a small open ball around the prior. This concludes the proof. To see it, note that a sender with concave preferences has no incentives to provide information (KG11), since the receiver is choosing the default action with probability one, this sender has no profitable deviation. Any sender with incentives to provide information, would want to induce a mean preserving spread of the posteriors by providing more information. However, from the construction of the δ the receiver is not convinced to change their action even if all other senders were to coordinate to send informative signals. That is, no deviation is profitable.

Now suppose that for the same prior μ_0 satisfying $A_{t_0} \geq (1 - \delta)I$ for some $\delta > 0$ there existed an equilibrium where the receiver does not choose the default action with probability one. Sender t_0 could deviate to the uninformative signal π_{t_0} defined above and guarantee the receiver to choose $a^*(\mu_0^\Theta)$ with probability one. Notice that regardless of the signals of other senders, this uninformative signal uses messages that are congruent across players, and so it is feasible.

Proof of Theorem 2

Let $a_\theta \in A = \Theta$ be the action that is optimal for the receiver when the state is θ , for all $\theta \in \Theta$, and t_θ the sender who would like the receiver to choose a_θ regardless of the state of the world. That is $v_{t_\theta}(a, \theta) = \mathbb{1}\{a = a_\theta\}$. Since $|\mathcal{T}| \leq |\Theta|$, let $\Theta' = \{\theta \in \Theta : t_\theta \in \mathcal{T}\}$. As for the receiver, to simplify calculations we can normalize the payoff of the default action, a_{θ_0} to 1, regardless of the state of the world, and the payoff of action a_θ to $u_\theta > 1$ when the state of the world is θ , and zero otherwise. That is $u(a, \theta) = u_\theta \mathbb{1}\{a = a_\theta\}$ for all $\theta \neq \theta_0$. The normalization allows us to define the decision rule for the receiver to choose any action, $a \neq a_{\theta_0}$, in a simple threshold rule, where $a^*(\mu^\Theta) = a_\theta$ only if $\mu^\Theta(\theta) \geq q_\theta$ for some threshold $q_\theta \in (0, 1)$. This is a necessary, but not a sufficient condition in general. For some beliefs, the receiver could be indifferent between two actions that are different from the default action. In that case, the choice of a_θ might not only depend on the belief that the state of the world is θ , but also on the belief about other states of the world. This is, however, irrelevant for the proof; for simplicity, let's assume that $a^*(\mu^\Theta) = a_\theta$ if only if $\mu^\Theta(\theta) \geq q_\theta$ ²⁶; i.e. $u_\theta \in (1, 2)$ for all $\theta \neq \theta_0$. For generic a marginal prior, the receiver strictly prefers the default action from any other action, thus $q_\theta > \mu_0^\Theta(\theta)$.

Let $\bar{\pi}^\mathcal{T}$ be an equilibrium expected signal with $|\mathcal{T}|$ senders, $\{\mu_1^\Theta(\Theta|m)\}_{m \in M}$ the set of posteriors induced in equilibrium, and $\hat{\Theta} \subseteq \Theta'$ the set of states for which the receiver-optimal action is induced: $\hat{\Theta} = \{a \in \Theta' : \exists m \in M \text{ such that } a^*(\mu_1^\Theta(\Theta|m)) = a\}$.

Claim: If $t_{\theta_0} \notin \mathcal{T}$, then $\hat{\Theta} = \Theta'$ for any $\bar{\pi}^\mathcal{T}$. Proof. Suppose there exists $\theta^* \in \Theta'$ but $\theta^* \notin \hat{\Theta}$. Then sender t_{θ^*} has a payoff of zero in equilibrium. The t_{θ^*} -conditional signal, $\bar{\pi}_{-t_{\theta^*}}$, has all messages in its support; otherwise, sender t_{θ^*} has a profitable deviation by using a message that is not in the support of $\bar{\pi}_{-t_{\theta^*}}$, action a_{θ^*} can be induced with some probability and yield a positive

²⁶For the general case consider the set of beliefs where action a_θ is induced and $q_\theta(\{u_\theta\}_{\theta \neq \theta_0})$ the boundary set of beliefs for which the agent is indifferent between action a_θ and some other action.

payoff to θ^* .

Claim: given the equilibrium $\bar{\pi}$, for every $\theta \in \hat{\Theta}$ there exists an equilibrium posterior, given some message, say m_θ , such that $\mu_1^\Theta(\theta|m_\theta) \geq q_\theta$, and $\mu_1^\Theta(\theta^*|m_{\theta^*}) = q_{\theta^*}$ for at least one $\theta^* \in \hat{\Theta}$. Proof. The first statement of the claim follows from the definition of $\hat{\Theta}$, there must exist a message, m , inducing a posterior $\mu_1^\Theta(\Theta|m)$ that assigns probability to the state of the world being θ , $\mu_1^\Theta(\theta|m_\theta)$, of at least q_θ . For the second claim, suffices to show that $\mu_1^\Theta(\theta^*|m_{\theta^*}) \neq q_{\theta^*}$, since $\theta^* \in \hat{\Theta}$. Proceed by contradiction, assume that for each $\theta \in \hat{\Theta}$ the inequality is strict, $\mu_1^\Theta(\theta|m_\theta) > q_\theta$, given some $m \in M$. Then the equilibrium signal for sender t_θ must send message m_θ with probability 1 in all states of the world, otherwise, there would be a profitable deviation. However, if this is true for all senders, then senders have pair-wise disjoint supports of their signals. In fact, after observing m_θ the receiver knows it came from π_{t_θ} with probability one and it is a perfectly uninformative signal, so $\mu_1^\Theta(\theta|m_\theta) = \mu_0^\Theta(\theta) < q_\theta$, a contradiction.

To see that any equilibrium with 1 sender is more informative. Let θ^* be the state of the world where for equilibrium $\bar{\pi}^\mathcal{T}$ we have that $\mu_1^\Theta(\theta^*|m_{\theta^*}) = q_{\theta^*}$ and consider the equilibrium signal for t_{θ^*} when there is no other sender: $\pi_{t_{\theta^*}}(m_{\theta^*}|\theta^*) = 1$, and $\pi_{t_{\theta^*}}(m_{\theta^*}|\theta) = \left(\frac{1-q_{\theta^*}}{q_{\theta^*}}\right) \left(\frac{\mu_0^\Theta(\theta^*)}{1-\mu_0^\Theta(\theta^*)}\right) < 1$ and send a different message, m_θ , with complementary probability in each state $\theta \neq \theta^*$. Easy algebra shows that this signal maximizes its expected payoff, and since message m_θ is only used in state θ for all $\theta \neq \theta^*$ it perfectly reveals such state of the world. We conclude that this signal must be more informative in the Blackwell order than $\bar{\pi}^\mathcal{T}$.