NBER WORKING PAPER SERIES

COMPETITION IN PERSUASION

Matthew Gentzkow
Emir Kamenica

Working Paper 17436
http://www.nber.org/papers/w17436

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2011

We are grateful to Eric Budish, Amy Finkelstein, Roger Myerson, Ben Olken, John Quah, Jesse Shapiro, Lars Stole, Xavier Vives, and numerous seminar participants for helpful comments and suggestions. We thank the Neubauer Family Foundation, the William Ladany Faculty Research Fund, and the Initiative on GlobalMarkets at the University of Chicago Booth School of Business for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Competition in Persuasion
Matthew Gentzkow and Emir Kamenica
NBER Working Paper No. 17436
September 2011, Revised April 2015
JEL No. D83,L15,M37

ABSTRACT

We study the impact of competition on information revelation in a class of Bayesian Persuasion games with multiple senders. Senders with no private information choose what information to gather and communicate to a receiver who takes a non-contractible action that affects the welfare of all players. The space of available signals includes all possible conditional distributions and allows any sender to choose a signal that is arbitrarily correlated with signals of others. We characterize the information revealed in pure-strategy equilibria, and show that greater competition tends to increase the amount of information revealed.

Matthew Gentzkow
University of Chicago
Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
gentzkow@chicagobooth.edu

Emir Kamenica
University of Chicago
Booth School of Business
5807 S. Woodlawn Ave.
Chicago, IL 60637
emir.kamenica@chicagobooth.edu

1 Introduction

Kamenica and Gentzkow (2011) study the problem of a single sender choosing what information to gather and communicate to a receiver. In this paper, we extend this analysis to the case of multiple senders. We characterize the set of equilibrium outcomes and examine how competition between the senders affects information revelation.

Several senders, who have no *ex ante* private information, simultaneously conduct costless experiments about an unknown state of the world. The set of possible experiments is rich; it includes all conditional distributions of signal realizations given the state, and it allows arbitrary correlation among the senders' signals. A third party (Receiver) observes the results of these experiments and then takes a non-contractible action that affects the welfare of all players. The state space is finite. Receiver and each of the senders have arbitrary, state-dependent, utility functions over the Receiver's action and the state of the world.. Throughout the paper we focus on pure-strategy equilibria of the game.[1]

The information revealed in an equilibrium of this game can be succinctly summarized by the distribution of Receiver's posterior beliefs (Blackwell 1953). We refer to such a distribution as an *outcome* of the game and order outcomes by informativeness according to the usual Blackwell criterion.

We begin our analysis by establishing a lemma that is the backbone of our main propositions: if the senders other than i together induce some outcome τ', sender i can unilaterally deviate to induce some other outcome τ if and only if τ is more informative than τ'. The "only if" part of this lemma is trivial and captures a basic property of information: an individual sender can unilaterally increase the amount of information revealed, but can never decrease it below the informational content of the other senders' signals. The "if" part of the lemma is more substantive, and draws on the assumption that senders have access to a rich set of possible signals.[2] It implies that no outcome can be a pure-strategy equilibrium if there exists a more informative outcome preferred by any sender. This property is the fundamental reason why competition tends to increase information revelation in our model.

[1] In Section 5, we briefly discuss the complications that arise with mixed strategies.

[2] In the language of Gentzkow and Kamenica (2015), this result means that the rich signal space is "Blackwell-connected." We discuss the relationship of our results to Gentzkow and Kamenica (2015) in greater detail below.

Our main characterization result provides an algorithm for finding the full set of pure-strategy equilibrium outcomes. We consider each sender i's value function over Receiver's beliefs $\hat{v}_i$ and its *concave closure* V_i (the smallest concave function everywhere above $\hat{v}_i$). Kamenica and Gentzkow (2011) show that when there is a single sender $i = 1$, any belief μ that Receiver holds in equilibrium must satisfy $\hat{v}_1(\mu) = V_1(\mu)$. We extend this result and establish that, when there are two or more senders, a distribution of posteriors is an equilibrium outcome *if and only if* every belief μ in its support satisfies $\hat{v}_i(\mu) = V_i(\mu)$ for all i. Identifying the set of these "unimprovable" beliefs for a given sender is typically straightforward. To find the equilibrium outcomes of the game, one then simply takes the intersection of these sets.

We then turn to the impact of competition on information revelation. We consider three ways of increasing competition among senders: (i) moving from collusive to non-cooperative play, (ii) introducing additional senders, and (iii) decreasing the alignment of senders' preferences. Since there are typically many equilibrium outcomes, we state these results in terms of set comparisons based on the strong and the weak set orders introduced by Topkis (1978). We show that, for all three notions of increasing competition, more competition never makes the set of outcomes less informative (under either order).

Competition does not necessarily make the set of outcomes more informative under the usual set orders, however, because the set of outcomes with more competition T can be non-comparable to the set of outcomes with less competition T'. If we restrict attention to comparable outcomes, however, we obtain stronger results. Specifically, given any (maximal) set of comparable elements C, we show that $T \cap C$ is more informative than $T' \cap C$.[3] This relationship holds in the strong set order for the comparison of collusive to non-cooperative play, and in the weak set order for the comparisons based on number of senders and preference alignment. We also show that if the game is zero-sum for any subset of senders, full revelation is the unique equilibrium outcome whenever the value functions are sufficiently nonlinear.

Finally, we discuss an important caveat, namely that when the outcomes under more and less competition are non-comparable, competition can lead to a form of coordination failure that makes Receiver strictly worse off.

[3] By definition, the empty set is not weakly above or below any other set. Hence, we consider only those comparable sets C that intersect T and T'.

The comparative statics results in this paper are special cases of a more general analysis in Gentzkow and Kamenica (2015). In that paper, we allow for arbitrary information environments and derive a necessary and sufficient condition under which competition cannot reduce information revealed regardless of preferences. Here, we specialize to the case of an arbitrarily rich signal space which allows us to use the tools from Kamenica and Gentzkow (2011) to provide a cleaner characterization of equilibria and a stronger version of the comparative statics using set orders. In addition, the focus on the rich signal signal space allows us to extend our results to settings where senders observe the outcomes of their experiments privately and then convey verifiable messages about those outcomes (cf: Gentzkow and Kamenica 2014).

The next section discusses the relationship between our paper and the existing literature. Section 3 develops mathematical preliminaries. Section 4 presents the model. Section 5 develops our main characterization result. Section 6 presents comparative statics. Section 7 presents applications to persuasion in courtrooms and product markets. Section 8 concludes.

2 Relationship to existing literature

Our paper extends the analytical methods developed in Kamenica and Gentzkow (2011) to address a new set of economic questions. The earlier paper establishes conditions under which persuasion is possible and characterizes optimal signals. This paper's primary focus is strategic interaction of multiple senders.

More specifically, our paper relates to four categories of prior literature. First, our model is related to the multiple-senders persuasion game analyzed in Milgrom and Roberts (1986). They identify sufficient conditions on sender preferences under which every equilibrium is fully revealing. They assume exogenous information and do not consider comparative statics.[4]

Second, our work connects to the large literature that examines conflict of interest among senders in cheap talk settings (e.g., Krishna and Morgan 2001; Battaglini 2002). Most papers in this literature focus on establishing conditions under which a fully revealing equilibrium exists.

[4]Bhattacharya and Mukherjee (2013) geometrically characterize equilibrium strategies in a multiple-sender persuasion game with uncertainty about whether each sender is informed. Chen and Olszewski (2014) present a model of debate between two senders with opposed preferences. They take senders' information as exogenous and do not consider comparative statics with respect to the extent of competition.

Since completely uninformative (babbling) equilibria are also always present, and since it is typically infeasible to characterize the full equilibrium set, these papers leave open the question of how much revelation we should expect to see in practice. These papers treat senders' information as exogenous, and focus primarily on the comparison of outcomes with one sender to outcomes with two senders.

Third, our result connect to research on advocacy. Dewatripont and Tirole (1999) show that, when information is costly to gather, hiring advocates with opposed interests is more effective than hiring a single unbiased agent. Shin (1998) present a related result in a setting with exogenous information. He compares a situation where two advocates get independent draws of a signal to a situation where an unbiased investigator gets a single draw. In contrast, in our model competition does not have a mechanical effect on the total information available.

Finally, two other papers examine symmetric information games with multiple senders. Brocas *et al.* (2012) and Gul and Pesendorfer (2012) study models where two senders with opposed interests generate costly signals about a binary state of the world. We assume signals are costless, but consider a more general setting with an arbitrary state space, arbitrary preferences, and arbitrary signals. Neither Brocas *et al.* (2012) nor Gul and Pesendorfer (2012) examine the impact of competition on information revelation.[5]

3 Mathematical preliminaries

3.1 State space and signals

Let Ω be a finite state space. A state of the world is denoted by $\omega \in \Omega$. A belief is denoted by μ. The prior distribution on Ω is denoted by μ_0. Let X be a random variable that is independent of ω and uniformly distributed on $[0,1]$ with typical realization x. We model signals as deterministic functions of ω and x. Formally, a *signal* π is a finite partition of $\Omega \times [0,1]$ s.t. $\pi \subset S$, where S is the set of non-empty Lebesgue measurable subsets of $\Omega \times [0,1]$. We refer to any element $s \in S$ as a *signal realization.*

[5]Ostrovsky and Schwarz (2010) examine how schools provide information about students' abilities in order to maximize the students' job placement. In their setting, each school can only generate information about the quality of its own students, while we assume all senders can generate information about any dimension of the state space. Moreover, Ostrovsky and Schwarz (2010) focus on a different question than we do – they examine whether the amount of information revealed depends on how students' abilities are distributed across schools.

Figure 1: A signal

$\omega = L$ $\omega = R$

X

π

With each signal π we associate an S-valued random variable that takes value $s \in \pi$ when $(\omega, x) \in s$. Let $p(s|\omega) = \lambda(\{x|\,(\omega, x) \in s\})$ and $p(s) = \sum_{\omega \in \Omega} p(s|\omega)\,\mu_0(\omega)$ where $\lambda(\cdot)$ denotes the Lebesgue measure. For any $s \in \pi$, $p(s|\omega)$ is the conditional probability of s given ω and $p(s)$ is the unconditional probability of s.

Our definition of a signal is somewhat non-standard because we model the source of noise in signal realizations (the random variable X) explicitly. This is valuable for studying multiple senders because for any two signals π_1 and π_2, our definition pins down not only their marginal distributions on S but also their joint distribution on $S \times S$. The joint distribution is important as it captures the extent to which observing both π_1 and π_2 reveals more information than observing only π_1 or π_2 alone. The more common definition of a signal, which takes the marginal distribution on S conditional on ω as the primitive, leaves the joint informational content of two or more signals unspecified.

Our definition of a signal is illustrated in Figure 1. In this example, $\Omega = \{L, R\}$ and $\pi = \{l, r\}$ where $l = (L, [0, 0.7]) \cup (R, [0, 0.3])$ and $r = (L, [0.7, 1]) \cup (R, [0.3, 1])$. The signal is a partition of $\Omega \times [0, 1]$ with marginal distribution $\Pr(l|L) = \Pr(r|R) = 0.7$.

3.2 Lattice structure

The formulation of a signal as a partition has the additional benefit of inducing an algebraic structure on the set of signals. This structure allows us to “add” signals together and thus easily examine their joint information content. Let Π be the set of all signals. Let $\trianglerighteq$ denote the refinement order on Π, that is, $\pi_1 \trianglerighteq \pi_2$ if every element of π_1 is a subset of an element of π_2. The pair $(\Pi, \trianglerighteq)$ is a lattice. The join $\pi_1 \vee \pi_2$ of two elements of Π is defined as the supremum of $\{\pi_1, \pi_2\}$. For any

Figure 2: The join of two signals

$\omega = L$ $\omega = R$

X

π_1

π_2

$\pi_1 \vee \pi_2$

finite set (or vector)[6] P we denote the join of all its elements by $\vee P$. We write $\pi \vee P$ for $\pi \vee (\vee P)$.

Note that $\pi_1 \vee \pi_2$ is a signal that consists of signal realizations s such that $s = s_1 \cap s_2$ for some $s_1 \in \pi_1$ and $s_2 \in \pi_2$. Hence, $\pi_1 \vee \pi_2$ is the signal that yields the same information as observing both signal π_1 and signal π_2. In this sense, the binary operation $\vee$ "adds" signals together. The join of two signals is illustrated in Figure 2.

3.3 Distributions of posteriors

A *distribution of posteriors*, denoted by τ, is an element of $\Delta(\Delta(\Omega))$ that has finite support.[7] A distribution of posteriors τ is *Bayes-plausible* if $\mathrm{E}_\tau[\mu] = \mu_0$.

Observing a signal realization s s.t. $p(s) > 0$ generates a unique posterior belief[8]

$$\mu_s(\omega) = \frac{p(s|\omega)\,\mu_0(\omega)}{p(s)} \text{ for all } \omega.$$

Note that the expression above does not depend on the signal; observing s from any signal π leads to the same posterior μ_s.

Each signal π induces a Bayes-plausible distribution of posteriors. We write $\langle \pi \rangle$ for the distribution of posteriors induced by π. It is easy to see that $\tau = \langle \pi \rangle$ assigns probability $\tau(\mu) =$

[6] In the model we introduce below, a strategy profile will be a vector of signals $\boldsymbol{\pi} = (\pi_1, ..., \pi_n)$ and we will write $\vee \boldsymbol{\pi}$ for $\vee \{\pi_i\}_{i=1}^n$.

[7] The fact that distributions of posteriors have finite support follows from the assumption that each signal has finitely many realizations. The focus on such signals is without loss of generality under the maintained assumption that Ω is finite.

[8] For those s with $p(s) = 0$, set μ_s to be an arbitrary belief.

$\sum_{\{s\in\pi:\mu_s=\mu\}} p(s)$ to each μ. Kamenica and Gentzkow (2011) establish that the image of the mapping $\langle\cdot\rangle$ is the set of all Bayes-plausible τ's:

Lemma 1. *(Kamenica and Gentzkow 2011) For any Bayes-plausible distribution of posteriors τ, there exists a $\pi \in \Pi$ such that $\langle\pi\rangle = \tau$.*

We define a *conditional distribution of posteriors* $\langle\pi|s\rangle$ to be the distribution of posteriors induced by observing signal π after having previously observed some signal realization s with $p(s) > 0$. This distribution assigns probability $\sum_{\{s'\in\pi:\mu_{s\cap s'}=\mu\}} \frac{p(s\cap s')}{p(s)}$ to each belief μ. For any π and s with $p(s) > 0$, we have $\mathrm{E}_{\langle\pi|s\rangle}[\mu] = \mu_s$. Lemma 1 can easily be extended to conditional distributions of posteriors:

Lemma 2. *For any s s.t. $p(s) > 0$ and any distribution of posteriors τ s.t. $\mathrm{E}_\tau[\mu] = \mu_s$, there exists a $\pi \in \Pi$ such that $\tau = \langle\pi|s\rangle$.*

Proof. Given any s s.t. $p(s) > 0$ and any distribution of posteriors τ s.t. $\mathrm{E}_\tau[\mu] = \mu_s$, let S' be a partition of s constructed as follows. For each ω, let $s_\omega = \{x|\,(\omega, x) \in s\}$. Now, partition each s_ω into $\{s_\omega^\mu\}_{\mu\in\mathrm{Supp}(\tau_s)}$ with $\lambda(s_\omega^\mu) = \frac{\mu(\omega)\tau(\mu)}{\mu_s(\omega)}\lambda(s_\omega)$. This is possible because $\mathrm{E}_\tau[\mu] = \mu_s$ implies $\sum_{\mu\in\mathrm{Supp}(\tau)} \mu(\omega)\tau(\mu) = \mu_s(\omega)$; hence, $\sum_{\mu\in\mathrm{Supp}(\tau)} \lambda(s_\omega^\mu) = \lambda(s_\omega)$. For each $\mu \in \mathrm{Supp}(\tau)$, let $s^\mu = \cup_\omega s_\omega^\mu$. Note that $S' = \{s^\mu|\mu \in \mathrm{Supp}(\tau)\}$ is a partition of s. Let $\pi = S' \cup \{\{\Omega \times [0,1] \setminus \{s\}\}\}$. It is easy to check that $\tau = \langle\pi|s\rangle$. □

Note that Lemma 1 is a Corollary of Lemma 2 as we can set s in the statement of Lemma 2 to $\Omega \times [0,1]$ so that $\mu_s = \mu_0$.

3.4 Informativeness

We order distributions of posteriors by informativeness in the sense of Blackwell (1953). We say that τ is *more informative* than τ', denoted $\tau \succsim \tau'$, if for some π and π' s.t. $\tau = \langle\pi\rangle$ and $\tau' = \langle\pi'\rangle$, there exists a *garbling* $g : S \times S \to [0,1]$ such that $\sum_{s'\in\pi'} g(s', s) = 1$ for all $s \in \pi$, and $p(s'|\omega) = \sum_{s\in\pi} g(s', s)\,p(s|\omega)$ for all ω and all $s' \in \pi'$. The relation $\succsim$ is a partial order. The pair $(\Delta(\Delta(\Omega)), \succsim)$ is a bounded lattice. We refer to the minimum element as *no revelation*, denoted

$\underline{\tau}$. Distribution $\underline{\tau}$ places probability one on the prior. We refer to the maximum element as *full revelation*, denoted $\overline{\tau}$. Distribution $\overline{\tau}$ has only degenerate beliefs in its support.[9]

The refinement order on the space of signals implies the informativeness order on the space of distributions of posteriors:

Lemma 3. $\pi \trianglerighteq \pi' \Rightarrow \langle \pi \rangle \succsim \langle \pi' \rangle$.

Proof. Define $g(s', s)$ equal to 1 if $s \subset s'$, and equal to 0 otherwise. Given any π and π' s.t. $\pi \trianglerighteq \pi'$, we know that for each $s \in \pi$, there is exactly one $s' \in \pi'$ s.t. $s \subset s'$. Hence, for all s, $\sum_{s' \in \pi'} g(s', s) = 1$. Moreover, $\pi \trianglerighteq \pi'$ implies that $\cup \{s \in \pi : s \subset s'\} = s'$. Hence, for any ω and any $s' \in \pi'$, $\{x | (\omega, x) \in \cup \{s \in \pi : s \subset s'\}\} = \{x | (\omega, x) \in s'\}$. This in turn implies $p(s'|\omega) = \sum_{s \in \pi} g(s', s) p(s|\omega)$. □

Note that it is not true that $\langle \pi \rangle \succsim \langle \pi' \rangle \Rightarrow \pi \trianglerighteq \pi'$.[10] Note also that Lemma 3 implies $\langle \pi_1 \vee \pi_2 \rangle \succsim \langle \pi_1 \rangle, \langle \pi_2 \rangle$.

We establish one more relationship between $\trianglerighteq$ and $\succsim$.

Lemma 4. *For any τ, τ', and π s.t. $\tau' \succsim \tau$ and $\langle \pi \rangle = \tau$, $\exists \pi'$ s.t. $\pi' \trianglerighteq \pi$ and $\langle \pi' \rangle = \tau'$.*

Proof. Consider any τ, τ', and π s.t. $\tau' \succsim \tau$ and $\langle \pi \rangle = \tau$. By Lemma 1, there is a $\hat{\pi}$ such that $\langle \hat{\pi} \rangle = \tau'$. Hence, by definition of $\succsim$, there is a garbling g such that $p(s|\omega) = \sum_{\hat{s} \in \hat{\pi}} g(s, \hat{s}) p(\hat{s}|\omega)$ for all $s \in \pi$ and ω. Define a new signal $\pi' \trianglerighteq \pi$ as follows. For each $s \in \pi$, for each $\omega \in \Omega$, let $s_\omega = \{x | (\omega, x) \in s\}$. Now, define a partition of each s_ω such that each element of the partition, say $s'(s, \hat{s}, \omega)$, is associated with a distinct $\hat{s} \in \hat{\pi}$ and has Lebesgue measure $g(s, \hat{s}) p(\hat{s}|\omega)$. This is possible since the sum of these measures is $p(s|\omega) = \lambda(s_\omega)$. Let $s'(s, \hat{s}) = \cup_\omega s'(s, \hat{s}, \omega)$. Let $\pi' = \{s'(s, \hat{s}) | \hat{s} \in \hat{\pi}, s \in \pi\}$. For any $s, \hat{s}, \omega_1, \omega_2$, we have

$$\frac{p(s'(s, \hat{s}) | \omega_1)}{p(s'(s, \hat{s}) | \omega_2)} = \frac{g(s, \hat{s}) p(\hat{s}|\omega_1)}{g(s, \hat{s}) p(\hat{s}|\omega_2)} = \frac{p(\hat{s}|\omega_1)}{p(\hat{s}|\omega_2)},$$

which implies $\langle \pi' \rangle = \langle \hat{\pi} \rangle = \tau'$. □

Note that it is not true that for any $\tau' \succsim \tau$ and $\langle \pi' \rangle = \tau'$, $\exists \pi$ s.t. $\pi' \trianglerighteq \pi$ and $\langle \pi \rangle = \tau$.

[9] A belief is degenerate if it places positive probability only on a single state.

[10] For example, suppose that there are two states L and R. π is a perfectly informative signal with two realizations. π' is an uninformative signal with ten realizations, each of which is equally likely in state L and state R. Then $\langle \pi \rangle \succsim \langle \pi' \rangle$, but π cannot be finer than π' because π' has more elements.

3.5 Orders on sets

We will frequently need to compare the informativeness of sets of outcomes. Topkis (1978; 1998) defines two orders on subsets of a lattice. Given two subsets Y and Y' of a lattice $(\mathcal{Y}, \geq)$, consider two properties of a pair $y, y' \in \mathcal{Y}$:

$$(S) \quad y \vee y' \in Y \text{ and } y \wedge y' \in Y'$$

$$(W) \quad \exists \hat{y} \in Y : \hat{y} \geq y' \text{ and } \exists \hat{y}' \in Y' : y \geq \hat{y}'$$

Topkis defines Y to be strongly above Y' (denoted $Y \geq_s Y'$) if property S holds for any $y \in Y$ and $y' \in Y'$, and to be weakly above Y' (denoted $Y \geq_w Y'$) if property W holds for any $y \in Y$ and $y' \in Y'$.

The set of outcomes is not generally a lattice under the Blackwell order (Müller and Scarsini 2006), so the strong set order is not well defined on this set, but this will not be an issue for the modification of the strong set order that we introduce below. The weak set order is of course well defined on any poset, whether or not it is a lattice. Given two sets of outcomes T and T', we say T is *weakly more informative* than T' if $T \succsim_w T'$.

Some of our results will establish that a particular set cannot be strictly less informative than another set. To simplify the statement of those propositions, we say that T is *no less informative* than T' if T is not strictly less informative than T' in the weak order.

Both the strong and the weak order are partial. Broadly speaking, there are two ways that sets Y and Y' can fail to be ordered. The first arises when one set has elements that are ordered both above and below the elements of the other set. For example, suppose that $\max(Y) > \max(Y')$ but $\min(Y) < \min(Y')$. Then, sets Y and Y' are not comparable in either the strong or the weak order, as seems intuitive. The second way that two sets can fail to be comparable arises when individual elements of the two sets are themselves not comparable. For example, suppose that $Y \geq_s Y'$ and $\tilde{y} \in \mathcal{Y}$ is not comparable to any element of $Y \cup Y'$. Then $Y \cup \tilde{y}$ is not comparable to Y' in either the strong or the weak order. The intuitive basis for calling $Y \cup \tilde{y}$ and Y' unordered may seem weaker than in the first case, and in some contexts we might be willing to say that $Y \cup \tilde{y}$ is above Y'.

In the analysis below, we will frequently encounter sets that fail to be ordered only in the latter sense. It will therefore be useful to distinguish these cases from those where sets are unordered even when we restrict attention to their comparable elements. A *chain* is a set in which any two elements are comparable, and a chain is *maximal* if it is not a strict subset of any other chain. We say that Y *is strongly (weakly) above* Y' *along chains* if for any maximal chain $C \subset \mathcal{Y}$ that intersects both Y and Y', $Y \cap C \geq_s Y' \cap C$ ($Y \cap C \geq_w Y' \cap C$).[11] In other words, orders along chains only require that conditions S and W hold for comparable pairs y and y'.

Given two sets of outcomes T and T', we say T is *strongly (weakly) more informative than* T' *along chains* if T is strongly (weakly) above T' along chains under the Blackwell order. Note that the strong (as well as the weak) order along chains is well defined on any poset. This is why it will not matter that the set of outcomes is not generally a lattice.

To gain more intuition about orders along chains, consider again properties S and W. When Y is strongly (weakly) above Y', property S (W) holds for any $y \in Y$ and $y' \in Y'$. When Y is strongly (weakly) above Y' along chains, property S (W) holds for any *comparable* y and y'.

Orders along chains also arise naturally in decision theory. The standard result on monotone comparative statics (Milgrom and Shannon 1994) states that, given a lattice $\mathcal{Y}$ and a poset Z, $\arg\max_{y \in Y} f(y, z)$ is monotone nondecreasing in z in the strong set order if and only if $f(\cdot, \cdot)$ satisfies the single-crossing property[12] and $f(\cdot, z)$ is quasisupermodular[13] for any z. It turns out that if we drop the requirement of quasisupermodularity, we obtain monotone comparative statics in the strong order along chains:[14]

Remark 1. Given a poset $\mathcal{Y}$ and a poset Z, $\arg\max_{y \in Y} f(y, z)$ is monotone nondecreasing in z in the strong set order along chains if and only if $f(\cdot, \cdot)$ satisfies the single-crossing property.

[11]Given any two sets Y and Y', the following three statements are equivalent: (i) for any maximal chain C, $Y \cap C \geq_s Y' \cap C$, (ii) for any chain C, $Y \cap C \geq_s Y' \cap C$, and (iii) for any chain C s.t. $|C| = 2$, $Y \cap C \geq_s Y' \cap C$.

[12]Function $f{:}Y \times Z \to \mathbb{R}$ satisfies the single-crossing property if $y > y'$ and $z > z'$ implies that $f(y, z') \geq f(y', z') \Rightarrow f(y, z) \geq f(y', z)$ and $f(y, z') > f(y', z') \Rightarrow f(y, z) > f(y', z)$.

[13]Function $f : Y \to \mathbb{R}$ is quasisupermodular if $f(y) \geq f(y \wedge y') \Rightarrow f(y \vee y') \geq f(y')$ and $f(y) > f(y \wedge y') \Rightarrow f(y \vee y') > f(y')$.

[14]We thank John Quah for this observation.

4 Bayesian persuasion with multiple senders

4.1 The model

Receiver has a continuous utility function $u(a, \omega)$ that depends on her action $a \in A$ and the state of the world $\omega \in \Omega$. There are $n \geq 1$ senders indexed by i. Each sender i has a continuous utility function $v_i(a, \omega)$ that depends on Receiver's action and the state of the world. All senders and Receiver share the prior μ_0. The action space A is compact.

The game has three stages: Each sender i simultaneously chooses a signal π_i from Π. Next, Receiver observes the signal realizations $\{s_i\}_{i=1}^n$. Finally, Receiver chooses an action.

Receiver forms her posterior using Bayes' rule; hence her belief after observing the signal realizations is μ_s where $s = \cap_{i=1}^n s_i$. She chooses an action that maximizes $\mathrm{E}_{\mu_s} u(a, \omega)$. It is possible for Receiver to have multiple optimal actions at a given belief, but for ease of exposition we suppose that Receiver takes a single action $a^*(\mu)$ at each belief μ. In Section 5 we discuss how our results can be restated to account for the multiplicity of optimal actions.

We denote sender i's expected utility when Receiver's belief is μ by $\hat{v}_i(\mu)$:

$$\hat{v}_i(\mu) \equiv \mathrm{E}_\mu \, v_i(a^*(\mu), \omega).$$

Throughout the paper, we focus exclusively on pure-strategy equilibria. We denote a strategy profile by $\boldsymbol{\pi} = (\pi_1, ..., \pi_n)$ and let $\boldsymbol{\pi}_{-i} = (\pi_1, ...\pi_{i-1}, \pi_{i+1}, ..., \pi_n)$. A profile $\boldsymbol{\pi}$ is an *equilibrium* if

$$\mathrm{E}_{\langle \vee \boldsymbol{\pi} \rangle} \hat{v}_i(\mu) \geq \mathrm{E}_{\langle \pi_i' \vee \boldsymbol{\pi}_{-i} \rangle} \hat{v}_i(\mu) \;\; \forall \pi_i' \in \Pi \; \forall i.$$

We refer to Receiver's equilibrium distribution of posteriors as the *outcome* of the game.[15] We say a belief μ is *induced* in an equilibrium if it is in the support of the equilibrium outcome.

4.2 Discussion of the model

Our model makes several strong assumptions.

[15]It is easy to see that Receiver's distribution of posteriors determines the distribution of Receiver's actions and the payoffs of all the players. The fact that each sender's payoff is entirely determined by the aggregate signal $\vee \boldsymbol{\pi}$ provides a link between our model and the literature on aggregate games (Martimort and Stole 2012).

First, we assume that signals are costless and that each sender can choose any signal whatsoever. This assumption would be violated if different senders had comparative advantage in accessing certain kinds of information, if there were some information that senders could not avoid learning, or if the experimental technology were coarse.

Second, our model implicitly allows each sender to choose a signal whose realizations are arbitrarily correlated, conditional on ω, with the signal realizations of the other senders. This would not be possible if signal realizations were affected by some idiosyncratic noise. One way to motivate our assumption is to consider a setting in which there is an exogenous set of experiments about ω and each sender's strategy is simply a mapping from the outcomes of those experiments to a message space. In that case, each sender can make his messages correlated with those of other senders. Another setting in which senders can choose correlated signals is one where they move sequentially. In that case, each sender can condition his choice of the signal on the realizations of the previous signals. The sequential move version of the game, however, is more cumbersome to analyze as the outcomes depend on the order in which senders move.[16] Moreover, Li and Norman (2015) show that in a sequential move version of our game, adding a sender can reduce the amount of information revealed in equilibrium.

Third, senders do not have any private information that could be inferred from their choice of the signal.

Fourth, our Receiver is a classical Bayesian who can costlessly process all information she receives. The main import of this assumption is that no sender can drown out the information provided by others, say by sending many useless messages. From Receiver's point of view, the worst thing that any sender can do is to provide no information. Hence, unlike in a setting with costly information processing, our model induces an asymmetry whereby each sender can add to but not detract from the information provided by others.

The four assumptions above not only make the model more tractable, but are required for our main results to hold. For example, if senders have access only to a small set of signals, one can construct examples where the unique collusive outcome is more informative than the unique

[16] There is nonetheless a connection between the simultaneous and the sequential move games. If τ is an equilibrium outcome of the sequential move game for all orders of moves by the senders, then τ obeys the characterization from Proposition 1.

equilibrium outcome. We also make several assumptions that are not necessary for the results, but greatly simplify the exposition.

First, we present the model as if there were a single Receiver, but an alternative way to interpret our setting is to suppose there are several receivers $j = 1, .., m$, each with a utility function $u_j(a_j, \omega)$, with receiver j taking action $a_j \in A_j$, and all receivers observing the realizations of all senders' signals. Even if each sender's utility $v_i(a, \omega)$ depends in an arbitrary way on the full vector of receivers' actions $a = (a_1, ..., a_m)$, our analysis still applies directly since, from senders' perspective, the situation is exactly the same as if there were a single Receiver maximizing $u(a, \omega) = \sum_{j=1}^{m} u_j(a_j, \omega)$.

Second, it is easy to extend our results to situations where Receiver has private information. Suppose that, at the outset of the game, Receiver privately observes a realization r from some signal $\xi(\cdot|\omega)$. In that case, Receiver's action, $a^*(s, r)$, depends on the realization of her private signal and is thus stochastic from senders' perspective. However, given a signal realization s, each sender simply assigns the probability $\xi(r|\omega)\mu_s(\omega)$ to the event that Receiver's signal is r and the state is ω. Hence, sender i's expected payoff given s is $\hat{v}_i(\mu_s) = \sum_\omega \sum_r v(a^*(s, r), \omega)\xi(r|\omega)\mu_s(\omega)$. All the results then apply directly with respect to the re-formulated $\hat{v}_i$'s.

Finally, our model assumes that Receiver directly observes the realizations of senders' signals. As noted above, however, the results in Gentzkow and Kamenica (2014) imply that this assumption is not necessary for our results. Gentzkow and Kamenica (2014) study the relationship between games where senders must report their information truthfully and disclosure games where they send verifiable messages. In particular, they consider a disclosure game with the following stages: (i) each sender simultaneously chooses a signal π_i, the choice of which is not observed by the other senders; (ii) each sender privately observes the realization s_i of his own signal; (iii) each sender simultaneously sends a verifiable message $m_i \subset \pi_i$ s.t. $s_i \in m_i$; (iv) Receiver observes the signals chosen by the senders and all of the messages; (v) Receiver chooses an action. Their results imply that the set of pure strategy equilibria of this disclosure game coincides with the set of pure strategy equilibria of the game we study in this paper.[17] Thus, our results are applicable even in settings where senders are able to conceal unfavorable information *ex post*.

[17] Gentzkow and Kamenica (2014) focus on cases where Receiver's optimal action is unique at each μ. When this assumption is not satisfied, the equivalence of the two games can be guaranteed by introducing a small amount of private information for Receiver, so that the distribution of Receiver's optimal actions is single-valued and continuous.

5 Characterizing equilibrium outcomes

In this section, we characterize the set of equilibrium outcomes. As a first step, consider the set of distributions of posteriors that a given sender can induce given the strategies of the other senders. It is immediate that he can only induce a distribution of posteriors that is more informative than the one induced by his opponents' signals alone. The following lemma establishes that he can induce *any* such distribution.

Lemma 5. *Given a strategy profile $\boldsymbol{\pi}$ and a distribution of posteriors τ, for any sender i there exists a $\pi_i' \in \Pi$ such that $\langle \pi_i' \vee \boldsymbol{\pi}_{-i} \rangle = \tau$ if and only if $\tau \succsim \langle \vee \boldsymbol{\pi}_{-i} \rangle$.*

Proof. Suppose $\tau \succsim \langle \vee \boldsymbol{\pi}_{-i} \rangle$. By Lemma 4, there exists a $\pi_i' \trianglerighteq \vee \boldsymbol{\pi}_{-i}$ s.t. $\langle \pi_i' \rangle = \tau$. Since $\pi_i' = \pi_i' \vee \boldsymbol{\pi}_{-i}$, we know $\langle \pi_i' \vee \boldsymbol{\pi}_{-i} \rangle = \langle \pi_i' \rangle = \tau$. The converse follows from Lemma 3. □

Lemma 5 depends on our assumption that each sender can choose a signal whose realizations are arbitrarily correlated, conditional on ω, with the signal realizations of the other senders. As a result, when senders play mixed strategies, the analogue of this lemma does not hold – it is possible to construct an example where senders other than i are playing mixed strategies $\tilde{\boldsymbol{\pi}}_{-i}$, there is a distribution of posteriors $\tau \succsim \langle \vee \tilde{\boldsymbol{\pi}}_{-i} \rangle$, and there is no π_i' such that $\langle \pi_i' \vee \tilde{\boldsymbol{\pi}}_{-i} \rangle = \tau$.[18] Consequently, the approach we develop below cannot be used to characterize mixed strategy equilibria. Li and Norman (2015) provide a more extensive discussion of mixed strategy equilibria in our setting.

We next turn to the question of when a given sender would wish to deviate to some more informative τ. For each i, let V_i be the *concave closure* of $\hat{v}_i$:

$$V_i(\mu) \equiv \sup\{z | (\mu, z) \in \mathrm{co}(\hat{v}_i)\},$$

where $\mathrm{co}(\hat{v}_i)$ denotes the convex hull of the graph of $\hat{v}_i$. Note that each V_i is concave by construction. In fact, it is the smallest concave function that is everywhere weakly greater than $\hat{v}_i$.[19] Kamenica and Gentzkow (2011) establish that when there is only a single sender i, $V_i(\mu_0)$ is the greatest payoff that the sender can achieve:

[18] Here, we extend the notation $\langle \cdot \rangle$ to denote the distribution of posteriors induced by a mixed strategy profile.
[19] Aumann and Maschler (1995) refer to V_i as the concavification of $\hat{v}_i$.

Lemma 6. *(Kamenica and Gentzkow 2011) For any belief* μ*,* $\hat{v}_i(\mu) = V_i(\mu)$ *if and only if* $\mathrm{E}_\tau[\hat{v}_i(\mu')] \leq \hat{v}_i(\mu)$ *for all* τ *such that* $\mathrm{E}_\tau[\mu'] = \mu$*.*

In light of this lemma, we refer to a belief μ such that $\hat{v}_i(\mu) = V_i(\mu)$ as *unimprovable* for sender i. Let M_i denote the set of unimprovable beliefs for sender i.

The lemma above establishes that, if there is a single sender, any belief induced in equilibrium has to be unimprovable for that sender. Our main characterization result shows that when $n \geq 2$, any belief induced in equilibrium has to be unimprovable for all senders. Moreover, unlike in the single sender case, this condition is not only necessary but sufficient: for any Bayes-plausible τ whose support lies in the intersection $M = \bigcap_{i=1}^{n} M_i$, there exists an equilibrium that induces τ.

Proposition 1. *Suppose* $n \geq 2$*. A Bayes-plausible distribution of posteriors* τ *is an equilibrium outcome if and only if each belief in its support is unimprovable for each sender.*

Gentzkow and Kamenica (2015) prove a version of this characterization result for more general settings based on the condition that each *distribution* of posteriors must be unimprovable for each sender. Proposition 1 is a stronger result for this more specialized setting that allows us to characterize equilibria based on the characteristics of individual *beliefs* in the support of such a distribution. In practice, this dramatically simplifies the process of computing the equilibrium set.

We provide a sketch of the proof here; a more detailed argument is in the Appendix. Suppose that τ is an equilibrium outcome. If there were some $\mu \in \mathrm{Supp}(\tau)$ such that $\hat{v}_i(\mu) \neq V_i(\mu)$ for some sender i, Lemmas 5 and 6 imply that sender i could profitably deviate by providing additional information when the realization of τ is μ. Conversely, suppose that τ is a Bayes-plausible distribution of beliefs such that for each $\mu \in \mathrm{Supp}(\tau)$, $\hat{v}_i(\mu) = V_i(\mu)$ for all i. Consider the strategy profile where all senders send the same signal π with $\langle \pi \rangle = \tau$. No sender can then deviate to induce any $\tau' \prec \tau$. Moreover, the fact that all beliefs in the support of τ are unimprovable means that no sender would want to deviate to any $\tau' \succ \tau$. Thus, this strategy profile is an equilibrium.

An important feature of Proposition 1 is that it provides a way to solve for the informational content of equilibria simply by inspecting each sender's preferences in turn, without worrying about fixed points or strategic considerations. This is particularly useful because identifying the set of unimprovable beliefs for each sender is typically straightforward. In Section 7, we will use this characterization to develop some applications. For now, Figure 3 illustrates how Proposition 1 can be applied in a simple example with hypothetical value functions. In this example, there are two senders, A and B. Panel (a) displays $\hat{v}_A$ and V_A, while Panel (b) displays $\hat{v}_B$ and V_B. Panel (c) shows the sets of unimprovable beliefs M_A and M_B, as well as their intersection M. Any distribution of beliefs with support in M is an equilibrium outcome. A belief such as μ_1 cannot be induced in equilibrium because sender A would have a profitable deviation. A belief such as μ_2 cannot be induced in equilibrium because sender B would have a profitable deviation.

Recall that, for ease of exposition, we have been taking some optimal $a^*(\cdot)$ as given and focusing on the game between senders. Proposition 1 thus characterizes the set of equilibrium outcomes consistent with this particular strategy by Receiver. To take the multiplicity of Receiver-optimal strategies into account, we could define a separate set of value functions $\hat{v}_i^\alpha(\mu)$ for each Receiver-optimal strategy α. Then, a distribution of posteriors τ is an equilibrium outcome if and only if there is an optimal action strategy α such that the support of τ lies in $\cap_i \{\mu | \hat{v}_i^\alpha(\mu) = V_i^\alpha(\mu)\}$.

Finally, observe that full revelation is an equilibrium in the example of Figure 3 (both $\mu = 0$ and $\mu = 1$ are in M). This is true whenever there are multiple senders, because degenerate beliefs are always unimprovable. This also implies that an equilibrium always exists.[20]

Corollary 1. *If $n \geq 2$, full revelation is an equilibrium outcome.*

As Sobel (2013) discusses, the existence of fully revealing equilibria under weak conditions is a common feature of multi-sender strategic communication models. In many of these models, as in ours, full revelation can be an equilibrium outcome even if all senders have identical preferences and strictly prefer no information disclosure to all other outcomes – a seemingly unappealing prediction.

One response would be to introduce a selection criterion that eliminates such equilibria. Given any two comparable equilibrium outcomes, every sender weakly prefers the less informative one. Hence, while the appropriate selection criterion might depend on the setting, selection criteria that always pick out a minimally informative equilibrium are appealing. We discuss the implications of such a selection criterion in Section 6.4 below. The approach we take in our formal results, however, is to focus on set comparisons of the full range of equilibrium outcomes.

[20]Kamenica and Gentzkow (2011) establish existence for the case $n = 1$. Consider an $a^*(\cdot)$ where Receiver takes a Sender-preferred optimal action at each belief. Such an $a^*(\cdot)$ guarantees that $\hat{v}_i$ is upper semicontinuous and thus that an equilibrium exists.

Figure 3: Characterizing equilibrium outcomes

(a) $\hat{v}$ and V functions for sender A

(b) $\hat{v}$ and V functions for sender B

(c) Sets of unimprovable beliefs $(\mu : \hat{v} = V)$

6 Competition and information revelation

6.1 Comparing competitive and collusive outcomes

One way to vary the extent of competition is to compare the set of non-cooperative equilibria to what senders would choose if they could get together and collude. This might be the relevant counterfactual for analyzing media ownership regulation or the effect of mergers on disclosure.

An outcome τ is *collusive* if $\tau \in \arg\max_{\tau'} \mathrm{E}_{\tau'} \left(\sum \hat{v}_i\left(\mu\right)\right)$. Note that it is without loss of generality to assume that, in choosing the collusive outcome, senders put equal weight on each player's utility; if, say due to differences in bargaining power, the collusive agreement placed weight δ_i on sender i, we could simply redefine each v_i as $\delta_i v_i$.[21]

Proposition 2. *Let T^* be the set of equilibrium outcomes and T^c the set of collusive outcomes. T^* is no less informative than T^c. Moreover, T^* is strongly more informative than T^c along chains.*

If there is a single sender, the proposition holds trivially as $T^* = T^c$, so suppose throughout this subsection that $n \geq 2$. We begin the proof with the following Lemma.

Lemma 7. *If $\tau^* \in T^*$, $\tau^c \in T^c$, and $\tau^c \succsim \tau^*$, then $\tau^c \in T^*$ and $\tau^* \in T^c$.*

Proof. Suppose $\tau^* \in T^*$, $\tau^c \in T^c$, and $\tau^c \succsim \tau^*$. By Lemma 5, we know $\mathrm{E}_{\tau^c}\left[\hat{v}_i\left(\mu\right)\right] \leq \mathrm{E}_{\tau^*}\left[\hat{v}_i\left(\mu\right)\right]$ for all i; otherwise, the sender i for whom $\mathrm{E}_{\tau^c}\left[\hat{v}_i\left(\mu\right)\right] > \mathrm{E}_{\tau^*}\left[\hat{v}_i\left(\mu\right)\right]$ could profitably deviate to τ^c. Since $\tau^c \in T^c$, we know $\mathrm{E}_{\tau^c}\left[\sum \hat{v}_i\left(\mu\right)\right] \geq \mathrm{E}_{\tau^*}\left[\sum \hat{v}_i\left(\mu\right)\right]$. Therefore, $\mathrm{E}_{\tau^c}\left[\hat{v}_i\left(\mu\right)\right] = \mathrm{E}_{\tau^*}\left[\hat{v}_i\left(\mu\right)\right]$ for all i which implies $\tau^* \in T^c$. Now, we know $\tau^c \in T^*$ unless there is a sender i and a distribution of posteriors $\tau' \succsim \tau^c$ s.t. $\mathrm{E}_{\tau'}\left[\hat{v}_i\left(\mu\right)\right] > \mathrm{E}_{\tau^c}\left[\hat{v}_i\left(\mu\right)\right]$. But since $\tau^* \in T^*$, $\mathrm{E}_{\tau^c}\left[\hat{v}_i\left(\mu\right)\right] = \mathrm{E}_{\tau^*}\left[\hat{v}_i\left(\mu\right)\right]$, and $\tau' \succsim \tau^c \succsim \tau^*$, this cannot be. □

Lemma 7 establishes one sense in which competition increases the amount of information revealed: no non-collusive equilibrium outcome is less informative than a collusive outcome, and no equilibrium outcome is less informative than a non-equilibrium collusive outcome. The lemma also plays a central role in the proof of Proposition 2:

[21] Moreover, it is not important that the collusive agreement maximizes the sum rather than the product of senders' payoffs. If we define a collusive outcome as an argmax of $\mathrm{E}_{\tau'}\left(\prod \max\left\{\hat{v}_i\left(\mu\right) - v_i^0, 0\right\}\right)$ where v_i^0 denotes sender i's "disagreement payoff," Proposition 2 would still hold, with a nearly identical proof. The definition of collusion based on the product of payoffs would be appropriate if firms reached a collusive agreement through Nash bargaining rather than through a merger.

Proof. Suppose $T^c \succsim_w T^*$. To establish that T^* is no less informative than T^c, we need to show this implies $T^* \succsim_w T^c$. For any $\tau^c \in T^c$, we know by Corollary 1 there exists $\tau^* \in T^*$ such that $\tau^* \succsim \tau^c$. For any $\tau^* \in T^*$, $T^c \succsim_w T^*$ implies there is a $\tau' \in T^c$ s.t. $\tau' \succsim \tau^*$. By Lemma 7, we must then have $\tau^* \in T^c$. Thus, there is a $\tau^c \in T^c$, namely τ^*, s.t. $\tau^c \precsim \tau^*$. Now, consider any chain C that intersects T and T'. Consider any $\tau^* \in T^* \cap C$ and any $\tau^c \in T^c \cap C$. By Lemma 7, $\tau^* \vee \tau^c \in T^* \cap C$ and $\tau^* \wedge \tau^c \in T^c \cap C$. Therefore, T^* is strongly more informative than T^c along chains. □

Note that the proposition allows for T^* to be non-comparable to T^c. The two sets can indeed be non-comparable in both the strong and the weak order. We will discuss the importance of these caveats below when we analyze whether competition necessarily makes Receiver better off.

6.2 Varying the number of senders

A second way to vary the extent of competition is to compare the set of equilibria with many senders to the set of equilibria with fewer senders. This might be the relevant counterfactual for assessing the impact of lowering barriers to entry on equilibrium advertising in an industry.

Proposition 3. *Let T and T' be the sets of equilibrium outcomes when the sets of senders are J and $J' \subset J$, respectively. T is no less informative than T'. Moreover, T is weakly more informative than T' if $|J'| > 1$, and weakly more informative than T' along chains if $|J'| = 1$.*

As suggested by the statement of the proposition, the basic intuition behind this result is somewhat different when we consider a change from many senders to more senders (i.e., when $|J'| > 1$), and when we consider a change from a single sender to many senders (i.e., when $|J'| = 1$).

In the former case, Proposition 1 implies that $T \subset T'$. In other words, adding senders causes the set of equilibrium outcomes to shrink. But, Corollary 1 implies that, even as the set of equilibrium outcomes shrinks, full revelation must remain in the set. Hence, loosely speaking, adding senders causes the set of equilibrium outcomes to shrink "toward" full revelation. We formalize this intuition in the following lemma, which will also be useful in proving Proposition 4 below.

Lemma 8. *Suppose T and T' are sets of outcomes s.t. $T \subset T'$ and $\overline{\tau} \in T$. Then T is weakly more informative than T'.*

Proof. Suppose T and T' are sets of outcomes s.t. $T \subset T'$ and $\overline{\tau} \in T$. For any $\tau' \in T'$ there exists a $\tau \in T$, namely $\bar{\tau}$, s.t. $\tau \succsim \tau'$. For any $\tau \in T$ there exists a $\tau' \in T'$, namely τ, s.t. $\tau \succsim \tau'$. □

In the latter case ($|J'| = 1$), the key observation is that no $\tau \in T \setminus T'$ can be less informative than a $\tau' \in T'$. Otherwise, the single sender in J' would prefer to deviate from τ to τ'. We now turn to the formal proof of Proposition 3.

Proof. If J is a singleton, the proposition holds trivially, so suppose that $|J| \geq 2$. First, consider there case where $|J'| > 1$. By Proposition 1, $T \subset T'$, and by Corollary 1, $\overline{\tau} \in T$. Hence, the proposition follows from Lemma 8. Second, consider the case where $|J'| = 1$. Let i denote the sender in J'. To establish that T is no less informative than T', we need to show that $T' \succsim_w T$ implies $T \succsim_w T'$. Suppose $T' \succsim_w T$. By Corollary 1, for any $\tau' \in T'$, we know there exists $\tau \in T$, namely $\overline{\tau}$, such that $\tau \succsim \tau'$. Given any $\tau \in T$, $T' \succsim_w T$ implies there is a $\tau' \in T'$ s.t. $\tau' \succsim \tau$. But, then it must be the case that τ is also individually optimal for sender i, i.e., $\tau \in T'$; otherwise, by Lemma 5, sender i could profitably deviate to τ' and hence τ would not be an equilibrium. Now, consider any maximal chain C that intersects T'. Since C is maximal, it must include $\overline{\tau}$. Moreover, $\overline{\tau} \in T$. Hence, for any $\tau' \in T' \cap C$ there is a $\tau \in T \cap C$, namely $\overline{\tau}$, s.t. $\tau \succsim \tau'$. It remains to show that for any $\tau \in T \cap C$ there is a $\tau' \in T' \cap C$ s.t. $\tau \succsim \tau'$. Given any $\tau \in T \cap C$, since C is a chain, every element of $T' \cap C$ is comparable to τ. Consider any $\tau' \in T' \cap C$. Since T' intersects C, there must be some such τ'. If $\tau' \precsim \tau$, we are done. Suppose $\tau' \succsim \tau$. Then, it must be the case that τ is also individually optimal for sender i, i.e., $\tau \in T'$; otherwise, by Lemma 5, sender i could profitably deviate to τ' and hence τ would not be an equilibrium. □

6.3 Varying the alignment of senders' preferences

A third way to vary the extent of the competition is to make senders' preferences more or less aligned. This counterfactual sheds lights on the efficacy of adversarial judicial systems and advocacy more broadly (Shin 1998; Dewatripont and Tirole 1999).

Given that senders can have any arbitrary state-dependent utility functions, the extent of preference alignment among senders is not easy to parametrize in general. Hence, we consider a specific form of preference alignment: given any two functions $f, g : A \times \Omega \to \mathbb{R}$ we let $\left\{\boldsymbol{v}^b\right\}_{b \in \mathbb{R}_+}$ denote a collection of preferences where some two senders, say j and k, have preferences of the form

$$\begin{aligned} v_j(a,\omega) &= f(a,\omega) + bg(a,\omega) \\ v_k(a,\omega) &= f(a,\omega) - bg(a,\omega) \end{aligned}$$

while preferences of Receiver and of other senders are independent of b. The parameter b thus captures the extent of preference misalignment between two of the senders.

Proposition 4. *Let T and T' be the sets of equilibrium outcomes when preferences are $\boldsymbol{v}^b$ and $\boldsymbol{v}^{b'}$, respectively, where $b > b'$. T is weakly more informative than T'.*

Proof. For each i, let M_i and M_i' denote the sets of unimprovable beliefs for sender i when preferences are $\boldsymbol{v}^b$ and $\boldsymbol{v}^{b'}$, respectively. Let $M = \cap_i M_i$ and $M' = \cap_i M_i'$. Let $\tilde{M} = M_j \cap M_k$ and $\tilde{M}' = M_j' \cap M_k'$. Let $\hat{f}(\mu) = \mathrm{E}_\mu[f(a^*(\mu),\omega)]$ and $\hat{g}(\mu) = \mathrm{E}_\mu[g(a^*(\mu),\omega)]$. Consider any $\mu \in \tilde{M}$. For any τ s.t. $\mathrm{E}_\tau[\mu'] = \mu$, we know that $\mu \in \tilde{M}_j$ implies $\mathrm{E}_\tau\left[\hat{f}(\mu') + b\hat{g}(\mu')\right] \leq \hat{f}(\mu) + b\hat{g}(\mu)$ and $\mu \in \tilde{M}_k$ implies $\mathrm{E}_\tau\left[\hat{f}(\mu') - b\hat{g}(\mu')\right] \leq \hat{f}(\mu) - b\hat{g}(\mu)$. Combining these two inequalities, we get $\hat{f}(\mu) - \mathrm{E}_\tau\left[\hat{f}(\mu')\right] \geq b\,|\hat{g}(\mu) - \mathrm{E}_\tau[\hat{g}(\mu')]|$,which means $\hat{f}(\mu) - \mathrm{E}_\tau\left[\hat{f}(\mu')\right] \geq b'\,|\hat{g}(\mu) - \mathrm{E}_\tau[\hat{g}(\mu')]|$. This last inequality implies $\mathrm{E}_\tau\left[\hat{f}(\mu') + b'\hat{g}(\mu')\right] \leq \hat{f}(\mu) + b'\hat{g}(\mu)$ and $\mathrm{E}_\tau\left[\hat{f}(\mu') - b\hat{g}(\mu')\right] \leq \hat{f}(\mu) - b\hat{g}(\mu)$. Since these two inequalities hold for any τ s.t. $\mathrm{E}_\tau[\mu'] = \mu$, we know $\mu \in \tilde{M}'$. Hence, $\tilde{M} \subset \tilde{M}'$. Therefore, since $M_i = M_i'$ for all $i \notin \{j,k\}$, we know $M \subset M'$. This in turn implies $T \subset T'$. By Corollary 1, we know $\overline{\tau} \in T$. Hence, the proposition follows directly from Lemma 8. □

Note that proofs of both Proposition 3 and Proposition 4 rely on the fact that, as competition increases (whether through adding senders or increasing misalignment of their preferences), the set of equilibrium outcomes shrinks. This is worth noting since it suggests another sense, not fully captured by the propositions, in which competition increases information revelation. Specifically, $T \subset T'$ implies that the set of unimprovable beliefs is smaller when there is more competition; hence, with more competition there are fewer prior beliefs such that no revelation is an equilibrium outcome.

Proposition 4 establishes that as preference misalignment b grows, the set of equilibrium outcomes shrinks and the extent of information revealed in equilibrium increases. A natural conjecture, therefore, may be that in the limit where two senders have fully opposed preferences, full revelation becomes the only equilibrium.

Specifically, suppose there are two senders j and k s.t. $v_j = -v_k$. Does the presence of two such senders guarantee full revelation? It turns out the answer is no. For example, if $\hat{v}_j$ is linear, and j and k are the only 2 senders, then $M_j = M_k = \Delta(\Omega)$ and any outcome is an equilibrium. Moreover, it will not be enough to simply assume that $\hat{v}_j$ is non-linear; as long as it is linear along some dimension of $\Delta(\Omega)$, it is possible to construct an equilibrium that is not fully revealing along that dimension. Accordingly, we say that $\hat{v}_j$ is *fully non-linear* if it is non-linear along every edge of $\Delta(\Omega)$, i.e., if for any two degenerate beliefs μ_ω and $\mu_{\omega'}$, there exist two beliefs μ_l and μ_h on the segment $[\mu_\omega, \mu_{\omega'}]$ such that for some $\gamma \in [0,1]$, $\hat{v}_j(\gamma\mu_l + (1-\gamma)\mu_h) \neq \gamma\hat{v}_j(\mu_l) + (1-\gamma)\hat{v}_j(\mu_h)$. If $v_j = -v_k$ and $\hat{v}_j$ is fully non-linear, then full revelation is indeed the unique equilibrium outcome. Proposition 5 establishes the analogous result for the more general case where there is some subset of senders for whom the game is zero-sum.

Proposition 5. *Suppose there is a subset of senders $J \subset \{1, ..., n\}$ s.t. (i) for any a and ω, $\sum_{i \in J} v_i(a, \omega) = 0$, and (ii) there exists $i \in J$ s.t. $\hat{v}_i$ is fully non-linear. Then, full revelation is the unique equilibrium outcome.*

6.4 Does competition make Receiver better off?

Propositions 2, 3, and 4 establish a sense in which moving from collusion to non-cooperative play, adding senders, and making senders' preferences less aligned all tend to increase information revelation. Since more information must weakly increase Receiver's utility, increasing competition thus tends to make Receiver better off.

To make this observation more precise, we translate our set comparisons of the informativeness of outcomes into set comparisons of Receiver's utilities. Given two lattices $(\mathcal{Y}, \succeq)$ and $(\mathcal{Z}, \geq)$, a function $f : \mathcal{Y} \to \mathcal{Z}$ is said to be *increasing* if $y \succeq y'$ implies $f(y) \geq f(y')$. Moreover, if the domain of f is a chain, then an increasing f preserves the set order:

Lemma 9. *If $f : (\Delta(\Delta(\Omega)), \succsim) \to (\mathbb{R}, \geq)$ is increasing, then for any chain $C \subset \Delta(\Delta(\Omega))$, $\forall T, T' \subset C$, $T \succsim_{s(w)} T' \Rightarrow f(T) \geq_{s(w)} f(T')$.*

Proof. First consider the strong order. Consider any $y \in f(T)$ and $y' \in f(T')$. If $y \geq y'$, then $y \vee y' \in f(T)$. Suppose $y' > y$. Let τ and τ' be any elements of $f^{-1}(y) \subset T$ and $f^{-1}(y') \subset T'$, respectively. Since f is increasing and $y > y'$, we know $\tau' > \tau$. Hence, since $T \succsim T'$, it must be the case that $\tau' \in T$. Hence, $y \wedge y' = y' = f(\tau') \in f(T)$. Now consider the weak order. Given $y \in f(T)$, consider any $\tau \in f^{-1}(y)$. Since $T \succsim_w T'$ there is a $\tau' \in T'$ s.t. $\tau \succsim \tau'$. Let $y' = f(\tau')$. Since f is increasing, $y \geq y'$. Given $y' \in f(T')$, consider any $\tau' \in f^{-1}(y')$. Since $T \succsim_w T'$ there is a $\tau \in T$ s.t. $\tau \succsim \tau'$. Let $y = f(\tau)$. Since f is increasing, $y \geq y'$. $\square$

By Blackwell's Theorem (1953), the function $f_u : (\Delta(\Delta(\Omega)), \succsim) \to (\mathbb{R}, \geq)$, which maps distributions of posteriors into the expected utility of a decision-maker with a utility function u, is increasing for any u. Hence, Lemma 9 allows us to translate the results of the previous three subsections into results about Receiver's payoff.

Corollary 2. *Let T^* be the set of equilibrium outcomes and T^c be the set of collusive outcomes. Let T and T' be the sets of equilibrium outcomes when the sets of senders are J and $J' \subset J$, respectively. Let T^b and $T^{b'}$ be the sets of equilibrium outcomes when preferences are $\boldsymbol{v}^b$ and $\boldsymbol{v}^{b'}$, respectively, where $b > b'$. For any maximal chain C that intersects T':*

1. *Receiver's payoffs under $T^* \cap C$ are strongly greater than under $T^c \cap C$*
2. *Receiver's payoffs under $T \cap C$ are weakly greater than under $T' \cap C$*
3. *Receiver's payoffs under $T^b \cap C$ are weakly greater than under $T^{b'} \cap C$*

By the definition of Blackwell informativeness, Corollary 2 applies not only to Receiver, whom senders are trying to influence, but also to any third-party who observes the signal realizations and whose optimal behavior depends on ω.[22]

An alternative to comparing sets of Receiver's payoffs is to consider a selection criterion that picks out a particular outcome from the overall set. As mentioned in Section 5, selection criteria that always pick out a minimally informative equilibrium may be appealing. Under any such criterion, there is a strong sense in which competition makes Receiver better off. Proposition 2 implies that any minimally informative equilibrium gives Receiver a weakly higher payoff than any comparable

[22] In the statement of Corollary 2, we do not need to assume that C intersects T^* or T^c because an empty set is strongly above and below any set and we do not need to assume that C intersects T, T^b, or $T^{b'}$ because all these sets contain $\bar{\tau}$ so any maximal chain must intersect them.

collusive outcome. Propositions 3 and 4 imply that any minimally informative equilibrium with more senders or less aligned preferences gives Receiver a weakly higher payoff than any comparable minimally informative equilibrium with fewer senders or more aligned sender preferences.

Whether we consider the entire equilibrium set or a particular selection rule, however, our results apply only to mutually comparable outcomes. This is a substantive caveat. If the outcomes under more and less competition are non-comparable, it is possible that the outcome with more competition makes Receiver worse off.

For example, suppose there are two dimensions of the state space, horizontal and vertical. Senders benefit by providing information only about the vertical dimension but strongly dislike providing information about both dimensions. In this case, competition could lead to a coordination failure; there can exist an equilibrium in which senders provide only horizontal information, even though all senders and Receiver would be strictly better off if only vertical information were provided:

Example 1. The state space is $\Omega = \{l, r\} \times \{u, d\}$. The action space is $A = \{l, m, r\} \times \{u, d\}$. Denote states, beliefs, and actions by ordered pairs (ω_x, ω_y), (μ_x, μ_y), and (a_x, a_y), where the first element refers to the l-r dimension and the second element refers to the u-d dimension. The prior is $\mu_0 = \left(\frac{1}{2}, \frac{1}{2}\right)$. Receiver's preferences are $u(a, \omega) = \frac{1}{100} u_x(a_x, \omega_x) + u_y(a_y, \omega_y)$, where $u_x(a_x, \omega_x) = \frac{2}{3} I_{\{a_x = m\}} + I_{\{a_x = \omega_x\}}$ and $u_y = I_{\{a_y = \omega_y\}}$. There are two senders with identical preferences: $v_1(a, \omega) = v_2(a, \omega) = I_{\{a_x = m\}} I_{\{a_y = \omega_y\}}$. A distribution of posteriors τ^* with support on beliefs $\left(0, \frac{1}{2}\right)$ and $\left(1, \frac{1}{2}\right)$ is an equilibrium outcome. The set of collusive outcomes, T^c, is the same as the set of equilibrium outcomes with a single sender, T'. Each of these sets consists of distributions of posteriors with support on $\left(\left[\frac{1}{3}, \frac{2}{3}\right] \times \{0\}\right) \cup \left(\left[\frac{1}{3}, \frac{2}{3}\right] \times \{1\}\right)$. It is easy to see that Receiver is strictly better off under any outcome in $T^c \cup T'$ than she is under τ^*.

7 Applications

7.1 A criminal trial

In Kamenica and Gentzkow (2011), we introduce the example of a prosecutor trying to persuade a judge that a defendant is guilty. Here, we extend that example to include two senders, a prosecutor

(p) and a defense attorney (d).

There are two states, innocent ($\omega = 0$) and guilty ($\omega = 1$). The prior is $\Pr(\omega = 1) = \mu_0 = 0.3$. Receiver (the judge) can choose to either acquit ($a = 0$) or convict ($a = 1$). Receiver's utility is $u(a, \omega) = I_{\{a=\omega\}}$. The prosecutor's utility is $v_p(a, \omega) = a$. The defense attorney's utility is $v_d(a, \omega) = -a$.

If the prosecutor were playing this game by himself, his optimal strategy would be to choose a signal that induces a distribution of posteriors with support $\left\{0, \frac{1}{2}\right\}$ that leads 60% of defendants to be convicted. If the defense attorney were playing this game alone, his optimal strategy would be to gather no information, which would lead the judge to acquit everyone. Because $v_p + v_d = 0$, all outcomes in this game are collusive outcomes.

When the attorneys compete, the unique equilibrium outcome is full revelation. This follows directly from Proposition 5, since $v_p = -v_d$ and the $\hat{v}_i$'s are fully non-linear. Thus, the equilibrium outcome is more informative than every collusive outcome and more informative than the two outcomes each sender would implement on their own, consistent with Propositions 2 and 3. In this example, competition clearly makes Receiver better off.

To make the analysis more interesting, we can relax the assumption that the two senders' preferences are diametrically opposed. In particular, suppose that the defendant on trial is a confessed terrorist. Suppose that the only uncertainty in the trial is how the CIA extracted the defendant's confession: legally ($\omega = 1$) or through torture ($\omega = 0$). Any information about the CIA's methods released during the trial will be valuable to terrorist organizations; the more certain they are about whether the CIA uses torture or not, the better they will be able to optimize their training methods. Aside from the attorneys' respective incentives to convict or acquit, both prefer to minimize the utility of the terrorists.

Specifically, we assume there is a second receiver, a terrorist organization.[23] The organization must choose a fraction $a_T \in [0, 1]$ of its training to devote to resisting torture. The organization's utility is $u_T(a_T, \omega) = -(1 - a_T - \omega)^2$. The attorneys' utilities are $v_p(a, \omega) = a - cu_T$ and $v_d(a, \omega) = -a - cu_T$. The parameter $c \in [4, 25]$ captures the social cost of terrorism internalized by the attorneys.[24]

[23] As discussed in Section 4.2, our model is easily reinterpreted to allow multiple receivers.

[24] If $c < 4$, the outcome is the same as when $c = 0$; the preferences of the two senders are sufficiently opposed that

Figure 4: Characterizing equilibrium outcomes for the criminal trial example

(a) $\hat{v}$ and V functions for sender p

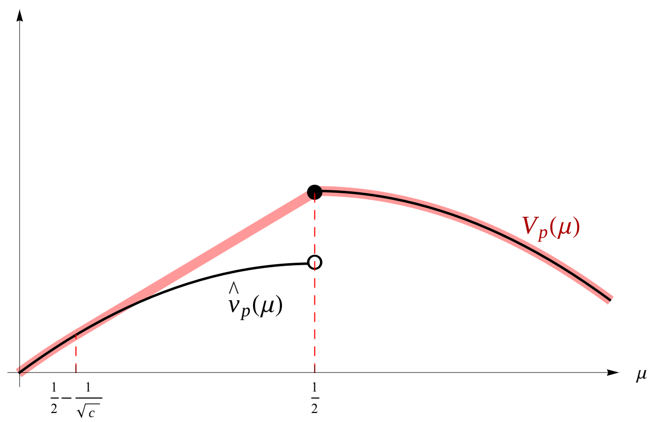

(b) $\hat{v}$ and V functions for sender d

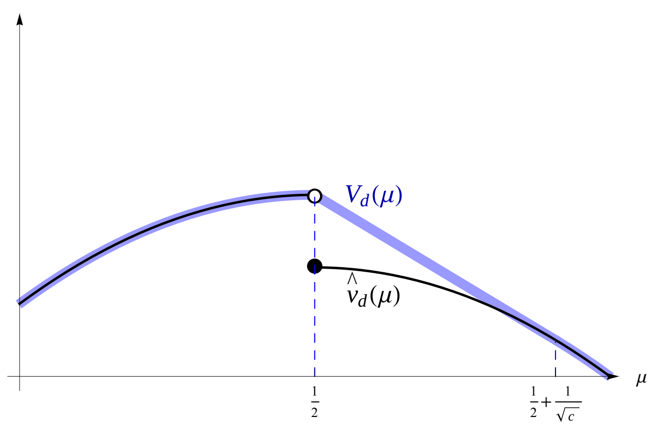

(c) Sets of unimprovable beliefs ($\mu : \hat{v} = V$)

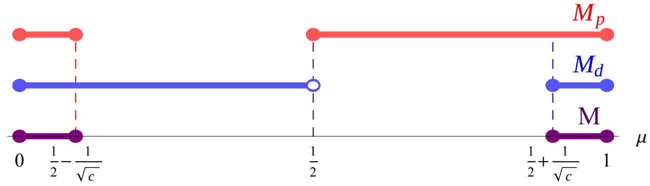

If the prosecutor were playing this game alone, his optimal strategy would be to choose a signal that induces a distribution of posteriors $\left\{\frac{1}{2}-\frac{1}{\sqrt{c}},\frac{1}{2}\right\}$. If the defense attorney were playing this game alone, his optimal strategy would still be to gather no information. The unique collusive outcome is no revelation. To identify the set of equilibrium outcomes, we apply Proposition 1. Panel (a) of Figure 4 plots $\hat{v}_p$ and V_p. We can see that $M_p=\{\mu|\hat{v}_p(\mu)=V_p(\mu)\}=\left[0,\frac{1}{2}-\frac{1}{\sqrt{c}}\right]\cup\left[\frac{1}{2},1\right]$. Panel (b) plots $\hat{v}_d$ and V_d. We can see that $M_d=\{\mu|\hat{v}_d(\mu)=V_d(\mu)\}=[0,\frac{1}{2})\cup\left[\frac{1}{2}+\frac{1}{\sqrt{c}},1\right]$. Hence, as panel (c) shows, $M=M_p\cap M_d=\left[0,\frac{1}{2}-\frac{1}{\sqrt{c}}\right]\cup\left[\frac{1}{2}+\frac{1}{\sqrt{c}},1\right]$. The set of equilibrium outcomes is the set of τ's whose support lies in this M.

Competition between the attorneys increases information revelation. Every equilibrium outcome is more informative than the collusive outcome (cf: Proposition 2) and more informative than what either sender would reveal on his own (cf: Proposition 3). Moreover, when the extent of shared interest by the two attorneys is greater, i.e., when c is greater, the set of equilibrium outcomes becomes weakly less informative (cf: Proposition 4).

7.2 Advertising of quality by differentiated firms

There are two firms $i\in\{1,2\}$ which sell differentiated products. The prices of these products are fixed exogenously and normalized to one, and marginal costs are zero. The uncertain state ω is a two-dimensional vector whose elements are the qualities of firm 1's product and firm 2's product. Receiver is a consumer whose possible actions are to buy neither product ($a=0$), buy firm 1's product ($a=1$), or buy firm 2's product ($a=2$) . We interpret the senders' choice of signals as a choice of verifiable advertisements about quality.[25]

There are three possible states: (i) both products are low quality ($\omega=(-5,-5)$), (ii) firm $1's$ product is low quality and firm 2's product is high quality ($\omega=(-5,5)$), or (iii) both products are high quality ($\omega=(5,5)$). Let $\mu_1=\Pr(\omega=(-5,5))$ and $\mu_2=\Pr(\omega=(5,5))$.

The firms' profits are $v_1=I_{\{a=1\}}$ and $v_2=I_{\{a=2\}}$. Receiver is a consumer whose utility depends

full revelation is the unique equilibrium outcome. If $c>25$, both senders are so concerned about giving information to the terrorists that neither wishes to reveal anything.

[25] Note that in this setting, our model allows for firms' advertisements to provide information about the competitor's product as well as their own. This is a reasonable assumption in certain industries. For example, pharmaceutical companies occasionally advertise clinical trials showing unpleasant side-effects or delayed efficacy of a rival product.

on a, $\omega = (\omega_1, \omega_2)$ and privately observed shocks $\epsilon = (\epsilon_0, \epsilon_1, \epsilon_2)$:[26]

$$\begin{aligned} u(a=0,\omega,\epsilon) &= \epsilon_0 \\ u(a=1,\omega,\epsilon) &= \omega_1 + \epsilon_1 \\ u(a=2,\omega,\epsilon) &= \omega_2 + \epsilon_2 \end{aligned}$$

We assume that the elements of ϵ are distributed i.i.d. type-I extreme value. Senders' expected payoffs at belief μ are thus

$$\begin{aligned} \hat{v}_1(\mu) &= \frac{\exp[\mathrm{E}_\mu(\omega_1)]}{1+\exp[\mathrm{E}_\mu(\omega_1)]+\exp[\mathrm{E}_\mu(\omega_2)]} \\ \hat{v}_2(\mu) &= \frac{\exp[\mathrm{E}_\mu(\omega_2)]}{1+\exp[\mathrm{E}_\mu(\omega_1)]+\exp[\mathrm{E}_\mu(\omega_2)]}. \end{aligned}$$

Figure 5 applies Proposition 1 to solve for the set of equilibrium outcomes. Panel (a) shows $\hat{v}_1$ and $\hat{v}_2$. Panel (b) shows V_1 and V_2. Panel (c) shows the sets of unimprovable beliefs M_1 and M_2 and their intersection M. The set of equilibrium outcomes is the set of τ's with supports in M.

Competition between the firms increases information revelation. The set of equilibrium outcomes is weakly more informative than what either firm would reveal on its own (cf: Proposition 3). Although not immediately apparent from Figure 5, the set of equilibrium outcomes is also weakly more informative than the set of collusive outcomes, and is strongly so along chains (cf: Proposition 2). The functional form of senders' utilities does not allow us to apply Proposition 4.

To understand the set of equilibria in this example, it is useful to consider the following two simpler settings. First, suppose $\mu_1 = 0$, so the only possible states are $\omega = (-5,-5)$ and $\omega = (5,5)$. In this case, the two firms' preferences are aligned: they both want to convince the consumer that $\omega = (5,5)$. The equilibrium outcomes, which one can easily identify by looking at the μ_2-edges in panel (c), involve partial information revelation. Next, suppose $\mu_2 = 0$, so the only possible states are $\omega = (-5,-5)$ or $\omega = (-5,5)$. Here, senders' preferences are opposed: sender 2 would like to convince Receiver that $\omega = (-5,5)$, while sender 1 would like to convince the consumer that $\omega = (-5,-5)$. The unique equilibrium outcome, which one can easily identify by looking at the

[26] As discussed in Section 4.2, our model is easily reinterpreted to allow Receiver to have private information.

Figure 5: Characterizing equilibrium outcomes for the advertising example

(a) $\hat{v}$ functions for senders 1 and 2

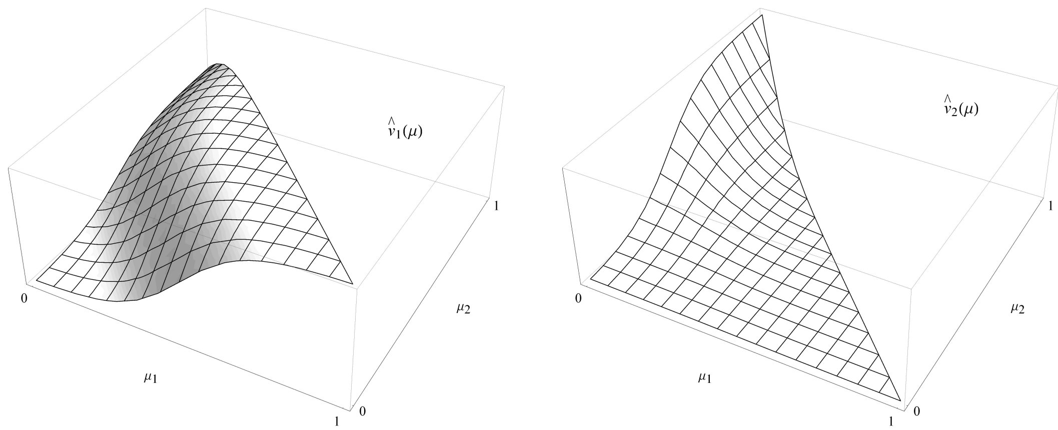

(b) V functions for senders 1 and 2

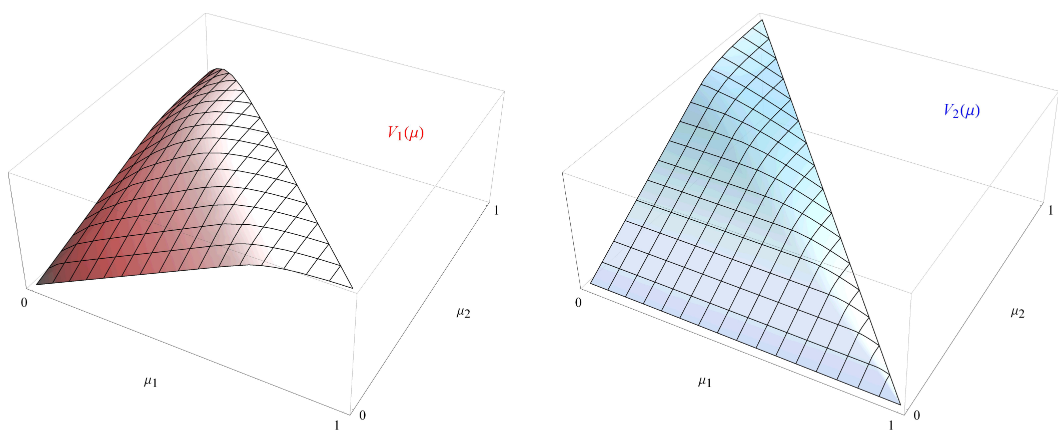

(c) Sets of unimprovable beliefs ($\mu : \hat{v} = V$)

μ_1-edges in panel (c), is full revelation. This is the case even though each firm on its own would prefer a partially revealing signal.[27] Finally, suppose that $\mu_1 + \mu_2 = 1$, so the only possible states are $\omega = (-5, 5)$ or $\omega = (5, 5)$. The firms' preferences are again opposed, and the unique equilibrium outcome, which one can read off the hypotenuses in panel (c), is again full revelation. This is the case despite the fact that firm 1 would strictly prefer no revelation.

In the full three-state example, the equilibrium involves full revelation along the dimensions where senders' preferences are opposed and partial revelation along the dimension where they are aligned. Consequently, the consumer learns for certain whether or not the state is $\omega = (-5, 5)$, but may be left uncertain whether the state is $\omega = (-5, -5)$ or $\omega = (5, 5)$.

8 Conclusion

In his review of the literature on strategic communication, Sobel (2013) points out that the existing work on multiple senders has largely focused on extreme results, such as establishing conditions that guarantee full revelation is an equilibrium outcome in cheap talk games. He mentions that most of this work does not fully characterize the equilibrium set. He also argues that the existing models do not capture the intuition that consulting more than two senders can be helpful even if different senders do not have access to different information.

In this paper, we assume that senders can costlessly choose any signal whatsoever, that their signals can be arbitrarily correlated with those of their competitors, and that Receiver observes all the information that is generated. Under these assumptions, we are able to address some of Sobel's concerns. We provide a simple way to identify the full set of pure-strategy equilibrium outcomes. We show that under quite general conditions competition cannot reduce the amount of information revealed in equilibrium, and in a certain sense tends to increase it. We also discuss the limitations of these results, in particular the fact that when outcomes with more and less competition are informationally non-comparable, an increase in competition can potentially be harmful.

[27] The gain to firm 2 from increasing μ_1 is much larger than the corresponding loss to firm 1; for this reason, at the scale of Figure 5, $\hat{v}_1$ appears flat with respect to μ_1 despite the fact that it is actually decreasing.

9 Appendix

9.1 Proof of Proposition 1

Lemma 10. *For any sender i and any distribution of posteriors τ:*

$$\hat{v}_i(\mu) = V_i(\mu)\,\forall\mu \in \text{Supp}(\tau) \Leftrightarrow \text{E}_{\tau'}[\hat{v}_i(\mu)] \leq \text{E}_\tau[\hat{v}_i(\mu)]\,\forall\tau' \succsim \tau.$$

Proof. Consider any i and any τ s.t. $\hat{v}_i(\mu) = V_i(\mu)\,\forall\mu \in \text{Supp}(\tau)$. Consider any $\tau' \succsim \tau$ and π' such that $\langle\pi'\rangle = \tau'$. For any s s.t. $\mu_s \in \text{Supp}(\tau)$, consider the conditional distribution of posteriors $\langle\pi'|s\rangle$. We know $\text{E}_{\langle\pi'|s\rangle}[\mu] = \mu_s$. Hence, by Lemma 6, $\text{E}_{\langle\pi'|s\rangle}[\hat{v}_i(\mu)] \leq \hat{v}_i(\mu_s)$. Therefore, $\text{E}_{\tau'}[\hat{v}_i(\mu)] = \sum_{s \text{ s.t. } \mu_s \in \text{Supp}(\tau)} p(s)\,\text{E}_{\langle\pi'|s\rangle}[\hat{v}_i(\mu)] \leq \sum_{s \text{ s.t. } \mu_s \in \text{Supp}(\tau)} p(s)\,\hat{v}_i(\mu_s) = \text{E}_\tau[\hat{v}_i(\mu)]$.

Conversely, suppose $\exists\mu_s \in \text{Supp}(\tau)$ such that $\hat{v}_i(\mu_s) \neq V(\mu_s)$. By Lemma 6, we know there exists a distribution of posteriors τ'_s with $\text{E}_{\tau'_s}[\mu] = \mu_s$ and $\text{E}_{\tau'_s}[\hat{v}_i(\mu)] > \hat{v}_i(\mu_s)$. By Lemma 2, there exists a π' s.t. $\tau'_s = \langle\pi'|s\rangle$. Let π be any signal s.t. $\langle\pi\rangle = \tau$. Let π'' be the union of $\pi \setminus \{s\}$ and $\{s \cap s' : s' \in \pi'\}$. Then $\langle\pi''\rangle \succsim \langle\pi\rangle = \tau$ and $\text{E}_{\langle\pi''\rangle}[\hat{v}_i(\mu)] = p(s)\,\text{E}_{\tau'_s}[\hat{v}_i(\mu)] + \sum_{\tilde{s}\in\pi\setminus\{s\}} p(\tilde{s})\,\hat{v}_i(\mu_{\tilde{s}}) > p(s)\,\hat{v}_i(\mu_s) + \sum_{\tilde{s}\in\pi\setminus\{s\}} p(\tilde{s})\,\hat{v}_i(\mu_{\tilde{s}}) = \text{E}_\tau[\hat{v}_i(\mu)]$ □

With Lemma 10, it is straightforward to establish Proposition 1.

Proof. Suppose $n \geq 2$. Suppose $\hat{v}_i(\mu) = V_i(\mu)\ \forall i\ \forall\mu \in \text{Supp}(\tau)$. By Lemma 1, there is a π such that $\langle\pi\rangle = \tau$. Consider the strategy profile $\boldsymbol{\pi}$ where $\pi_i = \pi\ \forall i$. Since $n \geq 2$, we know that $\vee\boldsymbol{\pi}_{-i} = \vee\boldsymbol{\pi}$. Hence, for any $\pi'_i \in \Pi$ we have $\pi'_i \vee \boldsymbol{\pi}_{-i} = \pi'_i \vee \boldsymbol{\pi} \trianglerighteq \vee\boldsymbol{\pi}$. Hence, by Lemma 3, $\langle\pi'_i \vee \boldsymbol{\pi}_{-i}\rangle \succsim \langle\vee\boldsymbol{\pi}\rangle$. Lemma 10 thus implies $\text{E}_{\langle\vee\boldsymbol{\pi}\rangle}\,\hat{v}_i(\mu) \geq \text{E}_{\langle\pi'_i\vee\boldsymbol{\pi}_{-i}\rangle}\,\hat{v}_i(\mu)$. Hence, $\boldsymbol{\pi}$ is an equilibrium.

Conversely, consider any equilibrium $\boldsymbol{\pi}$. Consider any $\tau' \succsim \langle\vee\boldsymbol{\pi}\rangle$. By Lemma 5, for any sender i there exists $\pi'_i \in \Pi$ such that $\langle\pi'_i \vee \boldsymbol{\pi}_{-i}\rangle = \tau'$. Since $\boldsymbol{\pi}$ is an equilibrium, this means $\text{E}_{\langle\vee\boldsymbol{\pi}\rangle}[\hat{v}_i(\mu)] \geq [\text{E}_{\tau'}\,\hat{v}_i(\mu)]$ for all i. Lemma 10 then implies that $\hat{v}_i(\mu) = V_i(\mu)\ \forall i\ \forall\mu \in \text{Supp}(\langle\vee\boldsymbol{\pi}\rangle)$. □

9.2 Proof of Proposition 5

We build the proof through the following three lemmas.

Lemma 11. *If there is a subset of senders $J \subset \{1, ..., n\}$ s.t. for any a and ω, $\sum_{i\in J} v_i(a,\omega) = 0$, then for any belief μ^* induced in an equilibrium, for any τ s.t. $\mathrm{E}_\tau[\mu] = \mu^*$, we have $\mathrm{E}_\tau[\hat{v}_i(\mu)] = \hat{v}_i(\mu^*)$ for all $j \in J$.*

Proof. Consider J s.t. $\sum_{i\in J} v(a,\omega) = 0$ $\forall a, \omega$ and any μ^* induced in an equilibrium. We must have $\hat{v}_i(\mu^*) = V_i(\mu^*)$ $\forall i$, and thus, by Lemma 6, $\mathrm{E}_\tau[\hat{v}_i(\mu)] - \hat{v}_i(\mu^*) \leq 0$ $\forall i$. We also have $\sum_{i\in J} \hat{v}(\mu) = 0$ $\forall \mu$, which implies $\sum_{i\in J} \mathrm{E}_\tau[\hat{v}_i(\mu)] = 0$. Hence, $\sum_{i\in J} [\mathrm{E}_\tau[\hat{v}_i(\mu)] - \hat{v}_i(\mu^*)] = 0$ $\forall i \in J$. Combining this with the earlier inequality, we obtain that $\mathrm{E}_\tau[\hat{v}_i(\mu)] - \hat{v}_i(\mu^*) = 0$ $\forall i \in J$. □

Lemma 12. *If $\hat{v}_j$ is non-linear, for any $\mu^* \in \mathrm{int}(\Delta(\Omega))$ there exists a τ s.t. $\mathrm{E}_\tau[\mu] = \mu^*$ and $\mathrm{E}_\tau[\hat{v}_j(\mu)] \neq \hat{v}_j(\mu^*)$.*

Proof. If $\hat{v}_j$ is non-linear, there exist $\{\mu_t\}_{t=1}^T$ and weights β_t s.t. $\sum \beta_t \hat{v}_j(\mu_t) \neq \hat{v}_j(\sum_t \beta_t \mu_t)$. Consider any $\mu^* \in \mathrm{int}(\Delta(\Omega))$. There exists some μ_l and $\gamma \in [0,1)$ s.t. $\mu^* = \gamma \mu_l + (1-\gamma)\sum \beta_t \mu_t$. If $\hat{v}_j(\mu^*) \neq \gamma \hat{v}_i(\mu_l) + (1-\gamma)\sum \beta_t \hat{v}_j(\mu_t)$, we are done. So, suppose that $\hat{v}_j(\mu^*) = \gamma \hat{v}_j(\mu_l) + (1-\gamma)\sum \beta_t \hat{v}_i(\mu_t)$. Now, consider the distribution of posteriors τ equal to μ_l with probability γ and equal to belief $\sum \beta_t \mu_t$ with probability $1-\gamma$. We have that $\mathrm{E}_\tau[\mu] = \mu^*$ and $\hat{v}_j(\mu^*) = \gamma \hat{v}_j(\mu_l) + (1-\gamma)\sum \beta_t \hat{v}_j(\mu_t) \neq \gamma \hat{v}_j(\mu_l) + (1-\gamma)\hat{v}_j(\sum \beta_t \mu_t) = \mathrm{E}_\tau[\hat{v}_j(\mu)]$. □

Lemma 13. *If $\hat{v}_j$ is fully non-linear, then the restriction of $\hat{v}_j$ to any n-dimensional face of $\Delta(\Omega)$ is non-linear if $n \geq 1$.*

Proof. The definition of fully non-linear states that the restriction of $\hat{v}_j$ to any 1-dimensional face of $\Delta(\Omega)$ is non-linear. For any $n \geq 1$, every n-dimensional face of $\Delta(\Omega)$ includes some $(n-1)$-dimensional face of $\Delta(\Omega)$ as a subset. Hence, if the restriction of $\hat{v}_j$ to every $(n-1)$-dimensional face is non-linear, so is the restriction of $\hat{v}_j$ to every n-dimensional face. Hence, by induction on n, the restriction of $\hat{v}_j$ to any n-dimensional face of $\Delta(\Omega)$ is non-linear if $n \geq 1$. □

With these lemmas, the proof of Proposition 5 follows easily.

Proof. Suppose there is a subset of senders $J \subset \{1, ..., n\}$ s.t. the conditions of the proposition hold. Let j be the sender in J for whom $\hat{v}_j$ is fully non-linear. Let μ^* be a belief induced in an equilibrium. Lemmas 11 and 12 jointly imply that μ^* must be at the boundary of $\Delta(\Omega)$. Hence,

μ^* is on some n-dimensional face of $\Delta(\Omega)$. But, by Lemma 13, if $n > 0$, the restriction of $\hat{v}_j$ to this n-dimensional face is non-linear. Hence, Lemmas 11 and 12 imply that μ^* must be on the boundary of this n-dimensional face, i.e., it must be on some $(n-1)$-dimensional face. Since this holds for all $n > 0$, we know that μ^* must be on a zero-dimensional face, i.e., it must be an extreme point of $\Delta(\Omega)$. Hence, any belief induced in an equilibrium is degenerate. □

References

Aumann, Robert J. and Michael B. Maschler. 1995. *Repeated Games with Incomplete Information.* Cambridge, MA: MIT Press.

Battaglini, Marco. 2002. Multiple referrals and multidimensional cheap talk. *Econometrica.* Vol. 70, No. 4, 1379-1401.

Bhattacharya, Sourav and Arijit Mukherjee. 2013. Strategic information revelation when experts compete to influence. *RAND Journal of Economics.* Vol. 44, No. 3, 522-544.

Blackwell, David. 1953. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics.* Vol. 24, No. 2, 265-272.

Brocas, Isabelle, Juan D. Carrillo, and Thomas R. Palfrey. 2012. Information gatekeepers: Theory and experimental evidence. *Economic Theory.* Vol. 51, No. 3, 649-676.

Chen, Ying and Wojciech Olszewski. 2014. Effective persuasion. *International Economic Review.* Vol. 55, No. 2, 319-347.

Dewatripont, Mathias and Jean Tirole. 1999. Advocates. *The Journal of Political Economy.* Vol. 107, No. 1, 1-39.

Gentzkow, Matthew and Emir Kamenica. 2014. Disclosure of endogenous information. University of Chicago mimeo.

—–. 2015. Information environments and the impact of competition on information provision. University of Chicago mimeo.

Gul, Faruk and Wolfgang Pesendorfer. 2012. The war of information. *Review of Economic Studies.* Vol. 79, No. 2, 707-734.

Kamenica, Emir and Matthew Gentzkow. 2011. Bayesian persuasion. *American Economic Review.* Vol. 101, No. 6, 2590-2615.

Krishna, Vijay and John Morgan. 2001. A model of expertise. *Quarterly Journal of Economics.* Vol. 116, No. 2, 747-775.

Li, Fei and Peter Norman. 2015. On Bayesian persuasion with multiple senders. University of North Carolina mimeo.

Martimort, David and Lars Stole. 2012. Representing equilibrium aggregates in aggregate games with applications to common agency. *Games and Economic Behavior.* Vol. 76, No. 2, 753-772.

Milgrom, Paul and John Roberts. 1986. Relying on the information of interested parties. *RAND Journal of Economics*. Vol. 17, No. 1, 18-32.

Milgrom, Paul and Chris Shannon. 1994. Monotone comparative statics. *Econometrica.* Vol. 62, No. 1, 157-180.

Müller, Alfred and Marco Scarsini. 2006. Stochastic order relations and lattices of probability measures. *Siam Journal of Optimization.* Vol. 16, No. 4, 1024-1043.

Ostrovsky, Michael and Michael Schwarz. 2010. Information disclosure and unraveling in matching markets. *American Economic Journal: Microeconomics.* Vol. 2, No. 2, 34-63.

Shin, Hyun Song. 1998. Adversarial and inquisitorial procedures in arbitration. *RAND Journal of Economics*. Vol. 29, No. 2, 378-405.

Sobel, Joel. 2013. Giving and receiving advice. In *Advances in Economics and Econometrics*, edited by Daron Acemoglu, Manuel Arellano, and Eddie Dekel. Vol. 1, 305-341. Cambridge: Cambridge University Press.

Topkis, Donald M. 1978. Minimizing a submodular function on a lattice. *Operations Research.* Vol. 26, No.2, 305-321.

——. 1998. *Supermodularity and Complementarity.* Princeton: Princeton University Press.