Multi-Sender Cheap Talk with Restricted State Spaces

Citation

Published Version

Permanent link

Terms of Use

Share Your Story

Multi-sender cheap talk with restricted state spaces

ATTILA AMBRUS

Department of Economics, Harvard University

SATORU TAKAHASHI

Department of Economics, Princeton University

This paper analyzes multi-sender cheap talk when the state space might be restricted, either because the policy space is restricted or the set of rationalizable policies of the receiver is not the whole space. We provide a necessary and sufficient condition for the existence of a fully-revealing perfect Bayesian equilibrium for any state space. We show that if biases are large enough and are not in similar directions, where the notion of similarity depends on the shape of the state space, then there is no fully-revealing perfect Bayesian equilibrium. The results suggest that boundedness, as opposed to dimensionality, of the state space plays an important role in determining the qualitative implications of a cheap talk model. We also investigate equilibria that satisfy a robustness property, diagonal continuity.

KEYWORDS. Cheap talk, multidimension, multi-sender, full revelation.

JEL CLASSIFICATION. C72, D82, D83.

1. INTRODUCTION

Sender–receiver games with cheap talk have been used extensively in both economics and political science to analyze situations in which an uninformed decision-maker acquires advice from an informed expert whose preferences do not fully coincide with those of the decision-maker. The seminal paper of Crawford and Sobel (1982) has been extended in many directions. In particular, Milgrom and Roberts (1986), Gilligan and Krehbiel (1989), Austen-Smith (1993), and Krishna and Morgan (2001a,b) investigate the case in which the decision-maker can seek advice from multiple experts. More recently, Battaglini (2002) has extended the analysis to multidimensional environments (the decision-maker seeks advice on multiple issues), and has called attention to the importance of equilibrium selection in multi-sender cheap talk games.

In this paper, we further investigate the existence of fully-revealing equilibrium and the existence of informative equilibrium for general state spaces. It might seem that these issues are settled, given that Battaglini (2002) provides a fairly complete analysis of two-sender cheap talk with unidimensional state spaces and shows that if the state space is a multidimensional Euclidean space, then generically a fully-revealing perfect Bayesian equilibrium can be constructed in which there are no out-of-equilibrium messages, so that the equilibrium survives any refinement that puts restrictions on out-of-equilibrium beliefs. The construction Battaglini provides is simple and intuitive: each sender conveys information only in directions along which her interest coincides with that of the receiver (directions that are orthogonal to the bias of the expert). Generically these directions of common interest span the whole state space; therefore, by combining the information obtained from the experts, the decision-maker can perfectly identify the state of the world.

We depart from Battaglini's analysis in allowing for a multidimensional state space that is not the whole Euclidean space, but a closed subset of it. The standard interpretation of states in sender–receiver games is that they represent circumstances under which a given policy action is optimal for the receiver. Given this interpretation, a restricted state space emerges naturally if either the set of available policies is restricted or if the set of rationalizable actions of the receiver is not the whole Euclidean space (that is, some policies would not be chosen by the receiver under any circumstances). In this way, the analysis of multidimensional cheap talk is more comparable to the earlier work on one-dimensional cheap talk, where the state space is standardly assumed to be a compact interval.

To illustrate the difference between bounded and unbounded state spaces, consider the following example. A policymaker needs to allocate a fixed budget between "education," "military spending," and "healthcare," and this decision depends on factors that are unknown to her. Suppose she can ask for advice from two perfectly informed experts, a left-wing analyst and a right-wing analyst. Assume that the left-wing analyst has a bias towards spending more on education, while the right-wing analyst has a bias towards spending more on the military; both of them are unbiased with respect to healthcare. If the state space were unbounded, corresponding to the absence of nonnegativity constraints on spending, a fully-revealing equilibrium could be constructed following Battaglini (2002). In this equilibrium, the amount to be spent on education depends only on the right-wing analyst's report, while the amount to be spent on the military depends only on the left-wing analyst's report (and the remaining budget is allocated to healthcare).

Now suppose that the amount of expenditure on each item must be nonnegative, as in a standard budget allocation problem. The situation may be depicted as in Figure 1: B corresponds to a state in which it is optimal for the policymaker to spend the whole budget on the military; C corresponds to a state in which it is optimal to spend all money on education; and A corresponds to a state in which it is optimal to spend no money on either education or the military. Note that the state space, represented by the triangle ABC, is bounded. The left-wing analyst's bias is orthogonal to AB, in the direction of C.
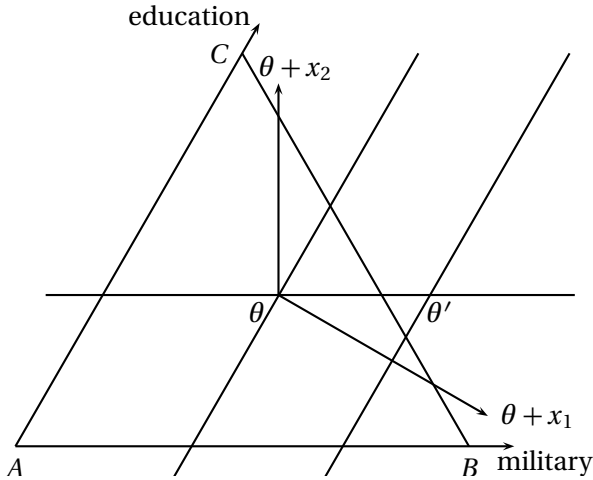
FIGURE 1. Nonnegativity constraints.

The right-wing analyst's bias is orthogonal to AC, in the direction of B. According to the construction that yields a fully-revealing equilibrium for an unbounded state space, the left-wing analyst is expected to report along a line parallel to AC.[1] Similarly, the right-wing analyst is expected to report along a line parallel to AB. Consider state θ in the figure. If the left-wing analyst sends a truthful report, then the right-wing analyst can send reports that are incompatible with the previous message in the sense that the only point compatible with the message pair is outside the state space (like θ' in the figure). Intuitively, these incompatible messages call for a combined expenditure on military and education that exceeds the budget. Such message pairs of course never arise if the experts play according to the candidate equilibrium. Nevertheless, it is important to specify the action the policymaker takes after receiving such a message, in order to make sure that both of the experts have the incentive to tell the truth. We confront this and other issues in our characterization of fully-revealing equilibria.

To extend the analysis to models with restricted state spaces, we first observe that Battaglini's characterization result for the existence of fully-revealing perfect Bayesian equilibrium for one-dimensional compact state spaces can be applied to arbitrary state spaces in any dimension. The result implies that the existence of fully-revealing equilibrium is monotonic in the magnitude of the biases, and that such equilibria always exist if the state space is large enough relative to the biases.

We also characterize the existence of fully-revealing equilibria for a compact state space when the biases are large. The case of senders with large biases is relevant in various applications: for example specialized committees of decision-making bodies frequently consist of preference outliers. We show that a fully-revealing equilibrium exists

---

[1]The most convenient way to think about this is to take a coordinate system in which the horizontal axis is parallel to AB, while the vertical axis is parallel to AC. Then the left-wing analyst reports the true state's horizontal coordinate (which corresponds to a line parallel to the vertical axis, that is parallel to AC).

for arbitrarily large biases if and only if the senders have similar biases. Similarity of biases is defined relative to the shape of the state space: two biases are similar if the intersection of the minimal supporting hyperplanes for the state space that are orthogonal to the biases contains a point of the state space. The intuition is that this point can be used to punish players if they send contradictory messages to the receiver. If the state space has a smooth boundary, then directions are similar if and only if they are exactly the same.

This result reconciles the seeming discontinuity between multi-sender cheap talk when the state space is one-dimensional and when it is multidimensional. In one dimension, there are only two types of biases, the same direction and the opposite direction. Biases of the former type are always similar, and biases of the latter type are never similar. As for multidimensional state spaces, biases with similar directions imply that full revelation is always possible in equilibrium, while non-similar directions imply that if the biases are small enough, then full revelation is possible, and otherwise it is not.

Battaglini (2002) emphasizes that in cheap talk games with multiple senders, perfect Bayesian equilibrium puts only very mild restrictions on out-of-equilibrium beliefs. Hence, not all equilibria are equally plausible: for example, some equilibria might only be supported by beliefs after out-of-equilibrium message pairs that induce the policymaker to choose a policy that is far from states that are compatible with any of the messages sent. Motivated by this concern, we proceed by imposing a robustness property, called diagonal continuity, on beliefs. We demonstrate that imposing this extra restriction on equilibria can reduce the possibility of full revelation in equilibrium drastically. For example, if the state space is a two-dimensional set with a smooth boundary and biases are not in the same direction, then there does not exist a fully-revealing diagonally-continuous equilibrium, no matter how small the biases are. As a counterpart of this result, we show that if the senders' biases are not in completely opposite directions, then under mild conditions information transmission in the most informative diagonally-continuous equilibrium can be bounded away from zero, no matter how large the biases are. This latter result contrasts with the case of only one sender, for which Crawford and Sobel (1982) show that in a unidimensional state space no information can be transmitted if the bias of the sender is large enough, and Levy and Razin (2007) show that in a multidimensional state space there is an open set of environments in which the most informative equilibrium approaches the uninformative equilibrium as the size of the bias goes to infinity.

2. THE MODEL

The model we consider has the same structure as that of Battaglini (2002), with the exception that the state space may be a proper subset of a Euclidean space. There are two senders and one receiver. The senders, labeled 1 and 2, both perfectly observe the state of the world $\theta \in \Theta$. The set Θ is referred to as the state space, and is a closed subset of $\mathbb{R}^d$. The prior distribution of θ is given by F. After observing θ, the senders send messages $m_1 \in M_1$ and $m_2 \in M_2$ to the receiver. The receiver observes these messages and chooses a policy $y \in Y \subseteq \mathbb{R}^d$ that affects the utility of all players. We assume that the policy space Y includes the convex hull of Θ, co(Θ).

For any $x = (x^1, \ldots, x^d)$ and $y = (y^1, \ldots, y^d) \in \mathbb{R}^d$, $x \cdot y = \sum_{j=1}^d x^j y^j$ denotes the inner product and $|x| = \sqrt{x \cdot x}$ denotes the Euclidean norm.

For state θ and policy y, the receiver's utility is $-|y - \theta|^2$, while sender i's utility is $-|y - \theta - x_i|^2$. The number $x_i \in \mathbb{R}^d$ is called sender i's bias. In state θ, the optimal policy of the receiver is θ, while the set of optimal policies of sender i is the set of points in Y that are closest to $\theta + x_i$ in Euclidean distance (which is exactly policy $\theta + x_i$ if the latter is included in the policy space).[2] Note that the magnitude of a sender's bias changes not only his optimal policies; it changes also his preferences over the whole policy space. Intuitively, as the magnitude of the bias increases, the indifference manifolds (curves when $d = 2$) of sender i in any state get closer and closer to hyperplanes (lines) that are orthogonal to x_i. We note that this formulation can be generalized without affecting the main results of the paper. In particular, the quadratic loss function can be changed to any smooth quasiconcave utility function, and some of the results can be extended to state-dependent biases as well.

Let $s_i : \Theta \to M_i$ denote a generic strategy of sender i in the above game, and let $y : M_1 \times M_2 \to Y$ denote a generic strategy of the receiver. Further, let $\mu(m_1, m_2)$ denote the receiver's probabilistic belief of θ given messages m_1, m_2. Strategies s_1, s_2, y constitute a perfect Bayesian equilibrium if there exists a belief function μ such that (i) s_i is optimal given s_{-i} and y for each $i \in \{1, 2\}$, (ii) $y(m_1, m_2)$ is optimal given $\mu(m_1, m_2)$ for each $(m_1, m_2) \in M_1 \times M_2$, and (iii) s_1 and s_2 are measurable and μ is a conditional probability system, given s_1, s_2, and F: if $s_1^{-1}(m_1) \cap s_2^{-1}(m_2)$ has a positive probability with respect to F, then $\mu(m_1, m_2)$ is derived from Bayes' rule. Note that $\mu(m_1, m_2)$ can be any distribution that puts probability 1 on $s_1^{-1}(m_1) \cap s_2^{-1}(m_2)$ if the latter is nonempty. Beliefs μ satisfying (iii) are said to *support* the perfect Bayesian equilibrium (s_1, s_2, y).

Note that the receiver's quadratic utility function implies that condition (ii) is equivalent to requiring that $y(m_1, m_2)$ be equal to the expectation of θ under $\mu(m_1, m_2)$. Let $\bar{\mu}(m_1, m_2)$ denote this expectation. The above implies that in a perfect Bayesian equilibrium the receiver always uses a pure strategy. The senders, however, may use mixed strategies in equilibrium, although the scope of this is rather limited in fully-revealing equilibria, which are at the center of our investigation. In the main part of the paper we ignore this possibility, and focus on pure strategy equilibria. See Section 5.2 for an extension of the results to the case in which the senders use mixed strategies. From now on we refer to a pure strategy perfect Bayesian equilibrium simply as an *equilibrium*.

3. Existence of fully-revealing equilibrium

3.1 *General biases*

By an argument similar to the well-known revelation principle in mechanism design, we do not lose generality by concentrating on truthful equilibria when investigating the existence of fully-revealing equilibria. This makes our task much easier.

---

[2] If $\theta + x_i$ is outside the policy space, then the point $\theta + x_i$ does not have a direct interpretation. In particular it is not the "ideal point" of the sender. Preferences are only defined over Y.

An equilibrium (s_1, s_2, y) is *fully revealing* if $s_1(\theta) = s_1(\theta')$ and $s_2(\theta) = s_2(\theta')$ imply $\theta = \theta'$. In this case, by the definition of a conditional probability system, $\mu(s_1(\theta), s_2(\theta))$ is the point mass on θ. An equilibrium (s_1, s_2, y) is *truthful* if $M_1 = M_2 = \Theta$ and $s_1(\theta) = s_2(\theta) = \theta$ for every $\theta \in \Theta$. A truthful equilibrium is fully revealing. In the next three claims we build heavily on results from Battaglini (2002): Lemma 1 below is essentially the same as Battaglini's Lemma 1, while Propositions 1 and 2 below are straightforward generalizations of Battaglini's Proposition 1 from one-dimensional line-segment state spaces to arbitrary state spaces in any dimension.

LEMMA 1 (Battaglini 2002, Lemma 1). *For any fully-revealing equilibrium, there exists a truthful equilibrium that is outcome-equivalent to the fully-revealing equilibrium.*

In cheap talk games, sequential rationality is a weak requirement. In particular, in truthful equilibria, after incompatible reports $\theta \neq \theta'$, the belief $\mu(\theta, \theta')$ can be an arbitrary distribution on Θ. The only restriction is that no sender has a strict incentive not to report the true state, to change the beliefs of the receiver, given that the other sender reports the truth.

Let $B(x, r) = \{y \in \mathbb{R}^d \mid |y - x| < r\}$ be the open ball with center x and radius r. For each sender i, $B(\theta + x_i, |x_i|)$ is the set of policies that are preferred to θ by sender i at state θ.

PROPOSITION 1. *Belief μ supports a truthful equilibrium if and only if, for every $\theta, \theta' \in \Theta$,*

$$\mu(\theta, \theta) \text{ is a point mass on } \theta \tag{1}$$
$$\overline{\mu}(\theta, \theta') \notin B(\theta' + x_1, |x_1|) \tag{2}$$
$$\overline{\mu}(\theta, \theta') \notin B(\theta + x_2, |x_2|). \tag{3}$$

PROOF.   Under condition (2), sender 1 does not strictly prefer reporting θ to reporting truthfully when the true state is θ'. Condition (3) is similar to (2). □

Figure 2 illustrates this result graphically: in order to keep incentive compatibility in states θ and θ', it is necessary that the policy chosen after the message pair (θ, θ') be a point outside both $B(\theta' + x_1, |x_1|)$ (otherwise, sender 1 would find it profitable to pretend that the state is θ in case the true state is θ') and $B(\theta + x_2, |x_2|)$ (otherwise, sender 2 would find it profitable to pretend that the state is θ' in case the true state is θ).

The above conditions are necessary and sufficient for the existence of a fully-revealing equilibrium, as stated in the next proposition.

PROPOSITION 2. *There exists a fully-revealing equilibrium if and only if $B(\theta' + x_1, |x_1|) \cup B(\theta + x_2, |x_2|) \not\supseteq \text{co}(\Theta)$ for all $\theta, \theta' \in \Theta$.*

PROOF.   By Lemma 1 and Proposition 1, a fully-revealing equilibrium exists if and only if there exists $\overline{\mu}(\theta, \theta')$ satisfying (1)–(3). Since $\overline{\mu}(\theta, \theta')$ is in the convex hull of Θ, if $B(\theta' + x_1, |x_1|) \cup B(\theta + x_2, |x_2|) \supseteq \text{co}(\Theta)$ for some $\theta, \theta' \in \Theta$ then (2)–(3) cannot hold simultaneously for any $\overline{\mu}(\theta, \theta')$. Otherwise, for every $\theta \neq \theta' \in \Theta$, let $\overline{\mu}(\theta, \theta')$ be an arbitrary element of $\text{co}(\Theta) \setminus (B(\theta' + x_1, |x_1|) \cup B(\theta + x_2, |x_2|))$. □
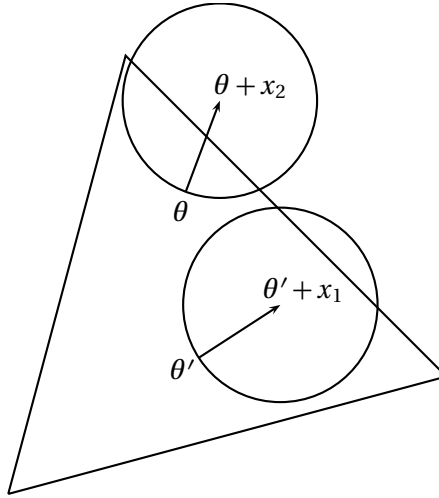
FIGURE 2. Constructing a fully-revealing equilibrium.

There cannot be a fully-revealing equilibrium whenever there exists a pair (θ, θ') of states such that the open balls $B(\theta' + x_1, |x_1|)$ and $B(\theta + x_2, |x_2|)$ cover the convex hull of the state space. Figure 3 depicts such a pair. Note that the existence of fully-revealing equilibrium depends only on the shape of the state space Θ and the biases x_1, x_2, not on the prior distribution F.

In the case of biases in the same direction, Proposition 2 implies that a fully-revealing equilibrium always exists, independently of the state space. The intuition is that $B(\theta' + x_1, |x_1|) \cup B(\theta + x_2, |x_2|)$ in this case does not contain the member of the set $\{\theta, \theta'\}$ that is minimal in the direction of the biases.

DEFINITION 1. The biases x_1 and x_2 are *in the same direction* if $x_1 = \alpha x_2$ for some $\alpha \geq 0$ or $x_2 = 0$.

PROPOSITION 3. *If x_1 and x_2 are in the same direction, then there exists a fully-revealing equilibrium.*

PROOF.   Let μ be the point belief

$$\overline{\mu}(\theta, \theta') = \begin{cases} \theta & \text{if } x_2 \cdot \theta > x_2 \cdot \theta' \\ \theta' & \text{if } x_2 \cdot \theta \leq x_2 \cdot \theta'. \end{cases}$$

Then μ supports a fully-revealing equilibrium.   □

We point out two more general consequences of Proposition 2. Both of them follow from the proposition in a straightforward manner, therefore we omit the proofs. The first one is that the existence of a fully-revealing equilibrium depends monotonically on the magnitudes of the biases: if there exists no fully-revealing equilibrium for biases x_1,
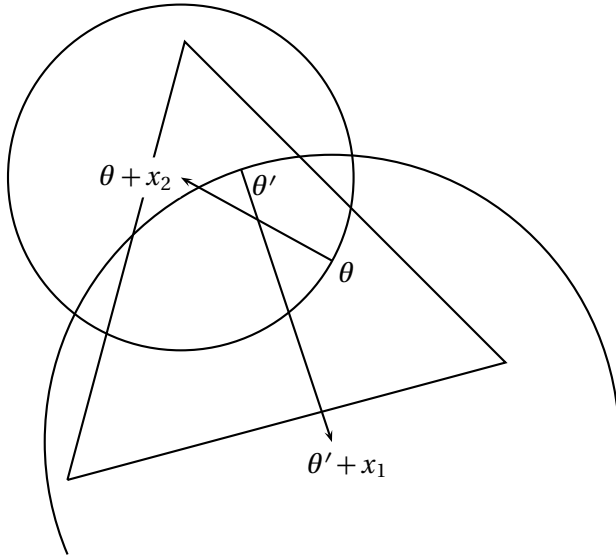
FIGURE 3. Nonexistence of fully-revealing equilibrium.

$x_2 \in \mathbb{R}^d$, then there exists no fully-revealing equilibrium for biases $(t_1 x_1, t_2 x_2)$ for any t_1, $t_2 \geq 1$. The other consequence is that there is a fully-revealing equilibrium if the biases are small enough relative to the size of the state space. Formally, if $|x_1| + |x_2| \leq (\sup_{\theta, \theta' \in \Theta} |\theta - \theta'|)/2$, then there exists a fully-revealing equilibrium. This in particular implies that there always exists a fully-revealing equilibrium if the state space is unbounded.

We close this subsection by showing that the nonexistence part of Proposition 2 can be strengthened, in the sense that if there is no fully-revealing equilibrium then there is an open set of states such that the implemented policy in these states is bounded away from the optimal policy of the receiver.

PROPOSITION 4. *There exists no fully-revealing equilibrium if and only if there exist $\varepsilon > 0$ and open sets U and U' satisfying $U \cap \Theta \neq \emptyset$ and $U' \cap \Theta \neq \emptyset$ such that, for any equilibrium (s_1, s_2, μ), either $|\overline{\mu}(s_1(\theta), s_2(\theta)) - \theta| > \varepsilon$ for all $\theta \in U$ or $|\overline{\mu}(s_1(\theta'), s_2(\theta')) - \theta'| > \varepsilon$ for all $\theta' \in U'$.*

PROOF.   The if part is trivial. For the only if part, suppose that there exists no fully-revealing equilibrium. Then Θ is bounded, and there exist $\widetilde{\theta}$, $\widetilde{\theta}' \in \Theta$ such that

$$B(\widetilde{\theta}' + x_1, |x_1|) \cup B(\widetilde{\theta} + x_2, |x_2|) \supseteq \mathrm{co}(\Theta).$$

Then there exist $\varepsilon > 0$ and neighborhoods U of $\widetilde{\theta}$ and U' of $\widetilde{\theta}'$ such that

$$B(\theta' + x_1, |x_1| - \varepsilon) \cup B(\theta + x_2, |x_2| - \varepsilon) \supseteq \mathrm{co}(\Theta)$$

for any $\theta \in U$ and $\theta' \in U'$.

For any equilibrium (s_1, s_2, μ) and any $\theta \in U$ and $\theta' \in U'$, we must have either $|\overline{\mu}(s_1(\theta), s_2(\theta)) - \theta| > \varepsilon$ or $|\overline{\mu}(s_1(\theta'), s_2(\theta')) - \theta'| > \varepsilon$ because otherwise we have $B(\theta' + x_1, |\theta' + x_1 - \overline{\mu}(s_1(\theta'), s_2(\theta'))|) \cup B(\theta + x_2, |\theta + x_2 - \overline{\mu}(s_1(\theta), s_2(\theta))|) \supseteq \text{co}(\Theta)$, where the first ball is the set of policies sender 1 strictly prefers to $\overline{\mu}(s_1(\theta'), s_2(\theta'))$ in state θ', and the second ball is the set of policies sender 2 strictly prefers to $\overline{\mu}(s_1(\theta), s_2(\theta))$ in state θ. Therefore, as for Proposition 2, no matter what $\overline{\mu}(s_1(\theta), s_2(\theta'))$ is, either sender 1 wants to report θ in state θ' or sender 2 wants to report θ' in state θ, which contradicts the equilibrium condition.

Therefore, if $|\overline{\mu}(s_1(\theta), s_2(\theta)) - \theta| \leq \varepsilon$ for some $\theta \in U$, then $|\overline{\mu}(s_1(\theta'), s_2(\theta')) - \theta'| > \varepsilon$ for all $\theta' \in U'$. Otherwise, $|\overline{\mu}(s_1(\theta), s_2(\theta)) - \theta| \leq \varepsilon$ for all $\theta \in U$. $\qquad \square$

The proof establishes that if there is no fully-revealing equilibrium, then there exist two open balls and a positive constant such that if in an equilibrium the implemented policy for at least one state in one ball is closer than ε to the state itself, then at every state in the other ball, the difference between the implemented policy and the state is at least as much as this constant. Note that the balls are defined independently of the equilibrium at hand; hence the above property applies to all equilibria. This is worth pointing out because typically there are many different types of equilibria, and it is hard to find nontrivial properties that hold for every equilibrium.

3.2 *Examples*

Our primary goal is to characterize conditions for full information revelation for large biases. Before providing the general result, it is useful to look at some concrete examples to develop intuition on how the possibility of full revelation depends on the shape of the state space and the directions and magnitudes of the biases.

We analyze closed balls and hypercubes. In the next subsection, closed balls are generalized to compact spaces with smooth boundaries and hypercubes to compact spaces with kinks.

Let D^d be the d-dimensional unit closed ball $\{\theta \in \mathbb{R}^d \mid |\theta| \leq 1\}$.

PROPOSITION 5. *Suppose* $\Theta = D^d$ *with* $d \geq 2$. *There exists a fully-revealing equilibrium if and only if* x_1 *and* x_2 *are in the same direction or* $\max(|x_1|, |x_2|) \leq 1$.

PROOF. If: By Proposition 3, we can assume that $\max(|x_1|, |x_2|) \leq 1$. For any given (θ, θ'), since $d \geq 2$, there exists a unit vector v perpendicular to $\theta' + x_1$. Let $w = -v$. We have $v, w \in D^d$. Since $|x_1| \leq 1$, (2) is satisfied both by $\overline{\mu}(\theta, \theta') = v$ and by $\overline{\mu}(\theta, \theta') = w$. Since $|v - w| = 2$ and $|x_2| \leq 1$, either v or w satisfies (3).

Only if: Suppose that x_1 and x_2 are in different directions and that $\max(|x_1|, |x_2|) > 1$. Without loss of generality, we can assume $|x_1| > 1$. By rotating the state space, we also have $x_1 = (-a, 0, \ldots, 0)$ with $a > 1$ without loss of generality. Substituting $\theta' = e := (1, 0, \ldots, 0)$ into (2), we have $|\overline{\mu}(\theta, e) - (e + x_1)| \geq a$. By the triangle inequality, $\overline{\mu}(\theta, e) \in D^d$, and $|e + x_1| = a - 1$, we have

$$a \leq |\overline{\mu}(\theta, e) - (e + x_1)| \leq |\overline{\mu}(\theta, e)| + |e + x_1| \leq 1 + (a - 1) = a.$$

Therefore, all the inequalities above hold with equality. Because $|\overline{\mu}(\theta, e)| = 1$, and $\overline{\mu}(\theta, e)$ and $-(e + x_1)$ are in the same direction, we have $\overline{\mu}(\theta, e) = e$. However, this violates (3) when θ is chosen appropriately. Again, without loss of generality, we have $x_2 = (b, c, 0, \ldots, 0)$ with $c \neq 0$, or $b > 0$ and $c = 0$.

For $c > 0$, we choose $\theta = (\sqrt{1 - \varepsilon^2}, -\varepsilon, 0, \ldots, 0)$ for small $\varepsilon > 0$. For $c < 0$, we choose $\theta = (\sqrt{1 - \varepsilon^2}, \varepsilon, 0, \ldots, 0)$ for small $\varepsilon > 0$. For $b > 0$ and $c = 0$, we choose $\theta = (1 - \varepsilon, 0, \ldots, 0)$ for small $\varepsilon > 0$. In each case, we have $e \in B(\theta + x_2, |x_2|)$, which violates (3). $\square$

Therefore, when Θ is a closed ball, as long as x_1 and x_2 are in different directions, whether a fully-revealing equilibrium exists is determined by the size of the biases. If the biases are small enough, then we can construct a fully-revealing equilibrium. If at least one of the biases is large enough, though, there is no such equilibrium.

Consider next $[0, 1]^d$, the unit hypercube in d dimensions. We say that x_1 and x_2 are *in the same orthant* if $x_1^j x_2^j \geq 0$ for every $j \in \{1, \ldots, d\}$.

PROPOSITION 6. *Suppose $\Theta = [0, 1]^d$.*

  (i) *If x_1 and x_2 are in the same orthant, then a fully-revealing equilibrium exists.*

  (ii) *If x_1 and x_2 are in different orthants and $\max_{i \in \{1,2\}} \min_{j \in \{1, \ldots, d\}} |x_i^j| > \frac{1}{2}$, then no fully-revealing equilibrium exists.*

PROOF.   For the first claim, without loss of generality we can assume that $x_i^j \geq 0$ for all $i \in \{1, 2\}$ and $j \in \{1, \ldots, d\}$. Let $\overline{\mu}(\theta, \theta') = (0, \ldots, 0)$ for any $\theta \neq \theta'$. Then (1)–(3) are satisfied.

For the second claim, without loss of generality we can assume that $x_1^j > \frac{1}{2}$ for all $j \in \{1, \ldots, d\}$, and $x_2^1 < 0$. Then, when $\theta' = (0, \ldots, 0)$ in (2), we have $\overline{\mu}(\theta, (0, \ldots, 0)) = (0, \ldots, 0)$ for any $\theta \in [0, 1]^d$. However, this violates (3) when $\theta = (\varepsilon, \ldots, 0)$ for $0 < \varepsilon < \min\{-2x_2^1, 1\}$. $\square$

The second part of the proposition establishes that if one of the biases x_i is large enough that there is a state θ such that $B(\theta + x_i, |x_i|)$ covers the whole hypercube with the exception of θ, then no matter how small is the other bias x_{-i}, as long as it is in a different orthant, there is a state θ' such that $B(\theta' + x_{-i}, |x_{-i}|)$ covers θ (see Figure 4 for an illustration).

For biases in different orthants, the qualitative conditions for the existence of fully-revealing equilibrium are similar to those for the case in which the state space is a d-dimensional unit closed ball. However, for the case of biases in the same orthant, the qualitative conclusion is different. Note that the proof—that, in this case, independent of the magnitudes of the biases, there always exists a fully-revealing equilibrium—uses the fact that for these biases, there is a point in the state space that is minimal among points in the state space in both directions of the biases. This point can serve as a punishment after any incompatible messages, which deters both senders from not revealing the true state.
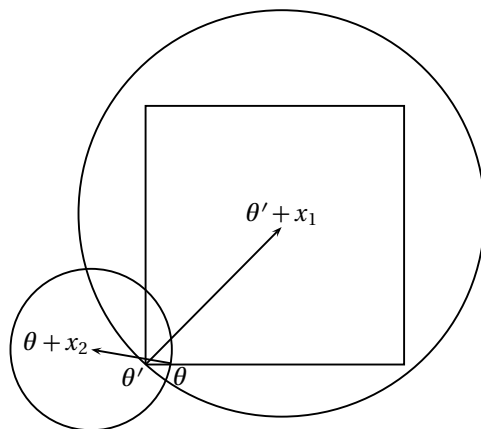
FIGURE 4. Square state space.

3.3 *Large biases*

A qualitative conclusion of Crawford and Sobel (1982) is that the amount of information that can be transmitted in equilibrium decreases when the sender's preferences diverge from the receiver's. In particular, if the sender's bias is sufficiently large, then no information transmission is possible in equilibrium. Krishna and Morgan (2001b) show that a similar insight holds for two-sender cheap talk games with one-dimensional state spaces, in the sense that the existence of a fully-revealing equilibrium depends on the magnitudes of the biases. However, Battaglini (2002) shows that if the state space is a multidimensional Euclidean space, then generically there exists a fully-revealing equilibrium, no matter how large the biases are. We analyze the case of large biases to revisit this question. Furthermore, large biases are relevant in certain applications. For example, distributive theories of committee formation in political science predict that specialized committees of a decision-making body consist of preference outliers.[3] In general, experts who have specialized knowledge are for many reasons (self-selection in the decision to become an expert, personal financial interests) likely to care in a strongly biased way about policy decisions affecting their fields of expertise.

In our model, if a sender has a large bias, then his or her indifference curves over a bounded policy space are close to hyperplanes orthogonal to the direction of the bias.[4] A natural interpretation is that as the bias of a sender becomes larger, the sender cares more about the direction of conflict with the receiver, and less about directions in which they share a common interest. For different ways of interpreting large biases, see the discussion at the end of this subsection. The formal statements of this subsection are limit results on the existence of fully-revealing equilibrium as the magnitudes of the biases go to infinity (as the preferences of senders approach lexicographic preferences). However,

---

[3]See for example Krishna and Morgan (2001a, p. 448).

[4]This observation, which plays a key role in the results below, was first pointed out by Levy and Razin (2007).

because of the result that the possibility of fully-revealing equilibrium is monotonic in the sizes of biases, the results below apply for all large enough biases.[5]

The next proposition shows that, if the state space is compact, then Proposition 2 for large biases is equivalent to an answer to the question of whether the state space can be covered by the union of two open half spaces with boundaries that are orthogonal to the directions of the biases.

Let S^{d-1} denote the $(d-1)$-dimensional unit sphere $\{x \in \mathbb{R}^d \mid |x| = 1\}$. The sphere S^{d-1} represents the set of possible directions in $\mathbb{R}^d$. For any $\lambda \in S^{d-1}$ and $k \in \mathbb{R}$, let $H^\circ(\lambda, k) = \{x \in \mathbb{R}^d \mid \lambda \cdot x > k\}$. The set $H^\circ(\lambda, \lambda \cdot x)$ is the open half space orthogonal to λ whose boundary goes through x.

PROPOSITION 7. *Fix a compact state space Θ and the directions of biases $z_1, z_2 \in S^{d-1}$. There exists a fully-revealing equilibrium with biases $(x_1, x_2) = (t_1 z_1, t_2 z_2)$ for every t_1, $t_2 \in \mathbb{R}_+$ if and only if $H^\circ(z_1, z_1 \cdot \theta') \cup H^\circ(z_2, z_2 \cdot \theta) \not\supseteq \mathrm{co}(\Theta)$ for all $\theta, \theta' \in \Theta$.*

PROOF.   If: The claim follows from Proposition 2 because $H^\circ(z_1, z_1 \cdot \theta') \cup H^\circ(z_2, z_2 \cdot \theta) \supseteq B(\theta' + t_1 z_1, t_1) \cup B(\theta + t_2 z_2, t_2)$ for every $t_1, t_2 \in \mathbb{R}_+$.

Only if: Suppose that $H^\circ(z_1, z_1 \cdot \theta') \cup H^\circ(z_2, z_2 \cdot \theta) \supseteq \mathrm{co}(\Theta)$ for some $\theta, \theta' \in \Theta$. Then, since $\mathrm{co}(\Theta)$ is compact, there exists $\varepsilon > 0$ such that $H^\circ(z_1, z_1 \cdot \theta' + \varepsilon) \cup H^\circ(z_2, z_2 \cdot \theta + \varepsilon) \supseteq \mathrm{co}(\Theta)$. Since $\mathrm{co}(\Theta)$ is bounded, we have $B(\theta' + t_1 z_1, t_1) \cap \mathrm{co}(\Theta) \supseteq H^\circ(z_1, z_1 \cdot \theta' + \varepsilon) \cap \mathrm{co}(\Theta)$ and $B(\theta + t_2 z_2, t_2) \cap \mathrm{co}(\Theta) \supseteq H^\circ(z_2, z_2 \cdot \theta + \varepsilon) \cap \mathrm{co}(\Theta)$ for sufficiently large t_1 and t_2. Hence the claim follows from Proposition 2.   □

Consider a compact state space Θ. For any $\lambda \in S^{d-1}$, define $k^*(\lambda, \Theta) = \min_{\theta \in \Theta} \lambda \cdot \theta$ and let $H^*(\lambda, \Theta) = \{x \in \mathbb{R}^d \mid \lambda \cdot x \geq k^*(\lambda, \Theta)\}$. Note that the compactness of Θ implies that $k^*(\lambda, \Theta)$ and therefore $H^*(\lambda, \Theta)$ are well-defined. The set $H^*(\lambda, \Theta)$ is the minimal half space that is orthogonal to λ and contains Θ. Let $h^*(\lambda, \Theta)$ denote the boundary of $H^*(\lambda, \Theta)$: $h^*(\lambda, \Theta) = \{x \in \mathbb{R}^d \mid \lambda \cdot x = k^*(\lambda, \Theta)\}$ is the supporting hyperplane to Θ in the direction of λ.

For every $\theta \in \Theta$, let $N_\Theta(\theta) = \{\lambda \in \mathbb{R}^d \mid \lambda \cdot (\theta' - \theta) \leq 0 \ \forall \theta' \in \Theta\}$. The set $N_\Theta(\theta)$ is the set of normal cones to Θ at the point θ. Then z_1 and z_2 are *similar with respect to Θ* if there exists $\theta \in \Theta$ such that $-z_1, -z_2 \in N_\Theta(\theta)$.

PROPOSITION 8. *Fix a compact state space Θ and the directions of biases $z_1, z_2 \in S^{d-1}$. The following conditions are equivalent.*

(i) *There exists a fully-revealing equilibrium with biases $(x_1, x_2) = (t_1 z_1, t_2 z_2)$ for every $t_1, t_2 \in \mathbb{R}_+$.*

(ii) $h^*(z_1, \Theta) \cap h^*(z_2, \Theta) \cap \Theta \neq \varnothing$.

(iii) $h^*(z_1, \Theta) \cap h^*(z_2, \Theta) \cap \mathrm{co}(\Theta) \neq \varnothing$.

(iv) z_1 *and z_2 are similar with respect to Θ.*

---

[5] The concrete meaning of large enough depends on the state space and the directions of the biases. See the examples in Section 3.2 for explicit derivations of threshold magnitudes.

PROOF. $1 \Rightarrow 2$: If not, then we have

$$H^*(z_1,\Theta) \cap H^*(z_2,\Theta) \setminus (h^*(z_1,\Theta) \cap h^*(z_2,\Theta)) \supseteq \Theta.$$

Since the left-hand side of this formula is a convex subset of $H°(z_1,k^*(z_1,\Theta)) \cup H°(z_2,k^*(z_2,\Theta))$, we have

$$H°(z_1,k^*(z_1,\Theta)) \cup H°(z_2,k^*(z_2,\Theta)) \supseteq \mathrm{co}(\Theta),$$

which contradicts Proposition 7.

   $2 \Rightarrow 3$: Trivial.

   $3 \Rightarrow 1$: Pick any $\widetilde{\theta} \in h^*(z_1,\Theta) \cap h^*(z_2,\Theta) \cap \mathrm{co}(\Theta)$. Then the claim follows from Proposition 7 because $\widetilde{\theta} \notin H°(z_i, z_i \cdot \theta)$ for any $i \in \{1,2\}$ and any $\theta \in \Theta$.

   $2 \Leftrightarrow 4$: Straightforward from the definition of $N_\Theta(\theta)$.   □

   This proposition makes it easy to check whether for an arbitrary pair of bias directions full revelation is possible in the limit. If the intersection of the supporting hyperplanes to the state space in the given directions contains a point in the state space, then the answer is no; otherwise, it is yes (as in Figure 5 below, where the intersection of the hyperplanes is a single point, outside the state space). This intersection is a lower dimensional hyperplane, and if it contains a point in the state space and $z_1 \neq z_2$, then that point has to be a kink of the state space. For example, in two dimensions, if $z_1 \neq z_2$ and $h^*(z_1,\Theta) \cap h^*(z_2,\Theta) \cap \Theta \neq \emptyset$, then $h^*(z_1,\Theta) \cap h^*(z_2,\Theta) \cap \Theta$ is a single point, which is such that there are supporting hyperplanes to Θ both in the direction of z_1 and in the direction of z_2. For a concrete example, recall the example of the d-dimensional cube with edges parallel to the axes from the previous subsection and consider $d = 2$. We saw that full revelation in equilibrium is possible even in the limit if biases go to infinity if and only if the directions of the biases are in the same quadrant. Note that for each of these direction pairs, there is a vertex of the square such that there are two lines orthogonal to the biases that are tangential to the square and go through the vertex.

   An immediate consequence of Proposition 8 is that for opposite biases ($z_1 = -z_2$), full revelation is possible in the limit if and only if Θ is included in a lower dimensional hyperspace that is orthogonal to the direction of the biases. To see this, note that in any other case, $h^*(z_1,\Theta) \cap h^*(z_2,\Theta) = \emptyset$; therefore, $h^*(z_1,\Theta) \cap h^*(z_2,\Theta) \cap \Theta = \emptyset$.

   A compact state space Θ has a convex hull with a *smooth boundary* if $\lambda, \lambda' \in N_\Theta(\theta) \cap S^{d-1}$ implies $\lambda = \lambda'$ for any $\theta \in \Theta$. The d-dimensional ball has a smooth boundary, whereas the d-dimensional cube does not. A simple corollary of Proposition 8 is that if the convex hull of Θ has a smooth boundary and $z_1, z_2 \in S^{d-1}$ then there exists a fully-revealing equilibrium with biases $(x_1,x_2) = (t_1 z_1, t_2 z_2)$ for every $t_1, t_2 \in \mathbb{R}_+$ if and only if $z_1 = z_2$.

   We can show also from Proposition 8 that we can assume the state space to be convex without loss of generality when we discuss the possibility of full revelation for large biases. This follows because the third condition in Proposition 8 depends only on $\mathrm{co}(\Theta)$.

   Our results imply that the same general result applies for state spaces in any dimension, including one: if the state space is compact, then for biases in similar directions,
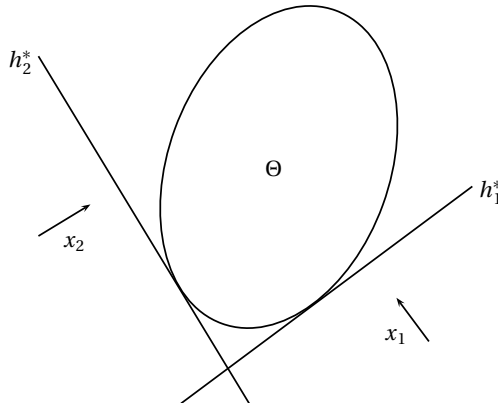
FIGURE 5. Large biases.

full revelation of information is possible for any magnitudes of biases; for biases that are not in similar directions, the magnitudes of biases matter: full revelation of information is possible for small biases, but not possible for large enough biases. In one dimension, there are only two types of direction pairs: the same direction and opposite directions. The former directions are always similar while the latter directions are always nonsimilar as long as the state space is not a singleton. In more than one dimension, the similarity relation depends on the shape of the state space. For state spaces with smooth boundaries, nonsimilar directions are generic, while for other state spaces, neither similar nor nonsimilar direction pairs are generic. In any case, for a two-sender cheap talk model with a compact state space, one can get the same qualitative conclusions with respect to the possibility of fully-revealing equilibrium when using a one-dimensional model (which is typically much easier to analyze) and when using a multidimensional model. There are two caveats, though. The first is that if one considers the one-dimensional model as a simplification of a more realistic multidimensional model, and similar biases are unlikely in that multidimensional model, then the one-dimensional analysis should put more emphasis on the case of opposite biases than on the case of like biases. The second, and more problematic caveat is that the above conclusion is based on the existence of fully-revealing perfect Bayesian equilibria. Cheap talk models typically have a severe multiplicity of equilibria, some of which are supported by implausible out-of-equilibrium beliefs by the receiver. This does not affect the validity of our results concerning the conditions for the nonexistence of fully-revealing equilibrium, since if the game has no fully-revealing perfect Bayesian equilibrium, then also it has no fully-revealing profile that is a refinement of perfect Bayesian Nash equilibrium. The possibility of implausible out-of-equilibrium beliefs does become a concern though for results that establish the existence of fully-revealing perfect Bayesian equilibrium. This is the main motivation for the analysis in the next section.

We conclude this section by briefly discussing alternative ways to model large biases, since, in a compact state space, there is no obvious way to define preferences for

extremely biased senders. When the biases get large, our model has two further qualitative implications besides the property that indifference curves converge to hyperplanes. One is that for large enough biases, a sender's optimal points are always on the boundary of the state space. The other one is that in a strictly convex state space, as the magnitude of bias goes to infinity a sender's optimal points converge to the same point on the boundary, no matter what the true state is.[6] These properties correspond well to the way we intuitively think about "large" or "extreme" biases. One way to generalize our model is to keep the latter two properties, but drop the assumption that the indifference curves converge to hyperplanes as the biases grow larger. In this case, the half spaces $h^*(z_1, \Theta)$ and $h^*(z_2, \Theta)$ in Proposition 8 need to be replaced with the limit upper contour sets in the alternative model.

4. ROBUST EQUILIBRIA

In cheap talk games, perfect Bayesian equilibrium (PBE) does not impose any restriction on the out-of-equilibrium beliefs of the receiver. Given this great flexibility in specifying out-of-equilibrium beliefs—which is made transparent in Proposition 1—the question arises which equilibria can be supported by "plausible" beliefs. This point is made by Battaglini (2002): when analyzing one-dimensional (bounded) state spaces, Battaglini focuses on equilibria that are supported by out-of-equilibrium beliefs satisfying a robustness criterion. The issue does not arise in the multidimensional analysis of Battaglini's paper, since the construction that he gives implies that there are no out-of-equilibrium message pairs when the state space is the whole Euclidean space. However, for restricted state spaces, out-of-equilibrium beliefs become relevant in multidimensional environments, too.

An extensive investigation of the robustness of PBE, and relatedly an investigation of PBE in models with noisy state observation, is difficult for general state spaces and is beyond the scope of this paper.[7] Instead, here we focus on equilibria that satisfy a particular continuity property. The property is motivated by requiring robustness to small mistakes in senders' observations, and is satisfied by the construction provided by Battaglini (2002, 2004) for unrestricted state spaces. We show that a strong definition of consistency of equilibrium beliefs implies this property. We then establish that imposing this property strengthens considerably our nonexistence results for fully-revealing equilibrium for some state spaces. On the other hand, we show that under mild conditions, there exist informative equilibria that satisfy the continuity property, no matter how large the biases are.

4.1 *Diagonal continuity*

The equilibrium construction provided in Battaglini (2002, 2004) satisfies the property that the policy implemented by the receiver is continuous in the observations of the

---

[6]We thank an anonymous referee for pointing this out.

[7]See the analysis of Battaglini (2004) for the case of unrestricted multidimensional state spaces with the improper uniform prior.

senders. In what follows, we investigate a requirement that is weaker than this, in that it requires continuity only at points where the observations of senders are the same.

DEFINITION 2. An equilibrium (s_1, s_2, y) is *continuous on the diagonal* if

$$\lim_{n \to \infty} y(s_1(\theta_1^n), s_2(\theta_2^n)) = y(s_1(\theta), s_2(\theta))$$

for any sequence $\{(\theta_1^n, \theta_2^n)\}_{n \in \mathbb{N}}$ of pairs of states such that $\lim_{n \to \infty} \theta_1^n = \lim_{n \to \infty} \theta_2^n = \theta$.

Our motivation for investigating equilibria that satisfy this property comes from multiple sources. First, we are interested in whether, in a restricted state space, there exist fully-revealing equilibria that can be obtained by a continuous transformation of the Battaglini construction.[8]

Second, this property is equivalent to robustness to *all* small misspecifications of the model. More precisely, suppose that the signals that two senders receive are slightly different from the true state, although all players (incorrectly) believe that both senders know the true state, and that other players believe that both senders know the true state, and so on. In such a situation, if the equilibrium is continuous on the diagonal, then the ex post loss for the receiver that arises from false beliefs is small for any realization of the true state when both senders receive signals close enough to the true state.

Third, as the next proposition shows, diagonal continuity is necessary for nonexistence of incompatible reports. The latter is a convenient property in settings where it is unclear how to specify out-of-equilibrium beliefs.[9]

PROPOSITION 9. *For compact Θ, a fully-revealing equilibrium (s_1, s_2, y) is continuous on the diagonal if*

  (i) *for each sender i, M_i is Hausdorff and $s_i : \Theta \to M_i$ is continuous, and*

  (ii) *for each $(m_1, m_2) \in s_1(\Theta) \times s_2(\Theta)$, there exists a state $\theta \in \Theta$ such that $(s_1(\theta), s_2(\theta)) = (m_1, m_2)$.*

PROOF.   Consider the function g on Θ that maps θ to $(s_1(\theta), s_2(\theta))$. By the assumptions, g is a continuous function onto $s_1(\Theta) \times s_2(\Theta)$. The function g is also one-to-one because (s_1, s_2, y) is fully revealing. Since g is a continuous bijection from the compact space Θ to the Hausdorff space $M_1 \times M_2$, the inverse $\overline{\mu}(m_1, m_2)$ is a continuous function of $(m_1, m_2) \in s_1(\Theta) \times s_2(\Theta)$.[10] Since s_1 and s_2 are continuous, $\overline{\mu}(s_1(\theta_1), s_2(\theta_2))$ is also continuous in (θ_1, θ_2).   □

The last motivation comes from the consistency of beliefs, i.e., the condition that beliefs should be limits of beliefs obtained from noisy models as the noise in the senders'

---

[8]We regard this question as interesting because Battaglini's equilibrium construction is simple and intuitive. A continuous transformation of the equilibrium preserves its basic attractive features, in the sense that the senders still report along different "dimensions" in a generalized sense.

[9]For example, Battaglini's (2002) equilibrium does not have incompatible reports if the state space is an entire Euclidean space.

[10]See Royden (1988, Proposition 5 of Chapter 9).

observations goes to zero. In the Appendix, we show that if we restrict attention to equilibria in which the players' strategies satisfy some regularity conditions, then every PBE that satisfies consistency of beliefs has to satisfy diagonal continuity. The regularity conditions we impose are fairly strong, but they are needed to ensure that the conditional beliefs of the receiver in "nearby" noisy models (which are invoked in the definition of consistent beliefs) are well-defined by Bayes' rule.

4.2 *Nonexistence of diagonally-continuous fully-revealing equilibria*

We now show that requiring diagonal continuity can drastically reduce the possibility of full revelation in equilibrium. First we consider two-dimensional smooth compact sets. Recall the result that if the biases are small enough (positive), then there always exists a fully-revealing equilibrium. In contrast, the next proposition shows that unless the biases are exactly in the same direction, no matter how small they are there does not exist a fully-revealing diagonally-continuous equilibrium.[11]

PROPOSITION 10. *In a two-dimensional smooth compact set Θ, if x_1 and x_2 are not in the same direction, then no diagonally-continuous fully-revealing equilibrium exists.*

PROOF.   Since Θ is a two-dimensional smooth set and x_1 and x_2 are not in the same direction, there exists $\theta \in \Theta$ such that θ is separated from other points in $co(\Theta) \setminus (B(\theta + x_1, |x_1|) \cup B(\theta + x_2, |x_2|))$. Since $y(\theta, \theta')$ is continuous with respect to θ' at $\theta' = \theta$, when we change θ' slightly, $y(\theta, \theta')$ has to move continuously. However, we can change θ' appropriately so that we can cover by $B(\theta' + x_1, |x_1|) \cup B(\theta + x_2, |x_2|)$ the region close enough to θ. This leads to a contradiction.   □

Figure 6 illustrates the argument used in the proof: if the biases are not in the same direction, then there are states θ and θ' arbitrarily close to each other (close to the boundary of the state space) such that the balls $B(\theta' + x_1, |x_1|)$ and $B(\theta + x_2, |x_2|)$ cover an open set that includes both θ and θ'. This means that in order for incentive compatibility to be satisfied for the senders, the policy implemented by the receiver after receiving messages corresponding to θ and θ' has to be "far away" from both θ and θ'. This implies that the equilibrium does not satisfy diagonal continuity at these points.[12]

A similar nonexistence result holds for models in which the state space is the unit d-dimensional cube (note the difference from the result in Proposition 6).

PROPOSITION 11. *Suppose $\Theta = [0,1]^d$. There exists no diagonally-continuous fully-revealing equilibrium if $x_1^j > 0$ for all $j \in \{1, \ldots, d\}$ and $x_2^k < 0$ for some $k \in \{1, \ldots, d\}$.*

PROOF.   When $\theta = \theta' = (0, \ldots, 0)$, $(0, \ldots, 0)$ is separated from other points in $\Theta \setminus (B(x_1, |x_1|) \cup B(x_2, |x_2|))$. Then, as in the proof of Proposition 10, we can change θ from

---

[11]As for the case of opposite biases, it is easy to see that the equilibrium constructed in Proposition 3 is diagonally continuous, since $\overline{\mu}(\theta, \theta')$ is either θ or θ'.

[12]This argument implicitly assumes, by invoking Lemma 1, that the fully-revealing equilibrium is truthful. This is without loss of generality, though: from the definition it follows that if a fully-revealing equilibrium satisfies diagonal continuity, the outcome-equivalent truthful equilibrium also satisfies diagonal continuity.
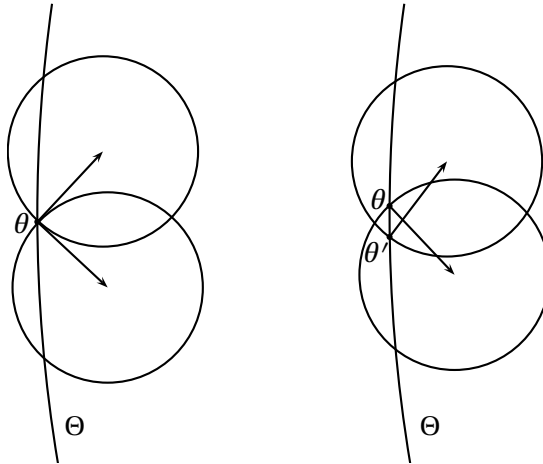
FIGURE 6. Nonexistence of diagonally-continuous fully-revealing equilibrium.

$(0\dots,0)$ toward the positive direction of the kth coordinate so that we can cover by $B(x_1,|x_1|)\cup B(\theta+x_2,|x_2|)$ a neighborhood of $(0,\dots,0)$. This leads to a contradiction. $\quad\square$

4.3 Existence of diagonally-continuous informative equilibria

Here we establish that if the prior distribution is continuous and the expected value of the state space is an interior point of the state space (which holds, for example, if Θ is convex and full dimensional), and biases are not in exactly opposite directions, then information transmission in the most informative equilibrium is bounded away from zero.[13]

PROPOSITION 12. *Assume that $d \geq 2$, the prior mean $E(\theta)$ is in the interior of Θ, and the prior distribution F has density f that is bounded away from 0 in a neighborhood of $E(\theta)$. If x_1 and x_2 are not in opposite directions, then there exists a diagonally-continuous equilibrium in which the receiver's ex ante payoff is strictly larger than that in the babbling equilibrium.*

PROOF.   By rotating and shifting the state space, we can assume $E(\theta)=0$, $x_1^1>0$, and $x_2^1>0$ without loss of generality.

Given positive small numbers a and b, we define the following region for each $t \in [0,1]$:

$$D(t)=\{\theta\in\mathbb{R}^d \mid \underline{\theta}^j(t)\leq\theta^j\leq\overline{\theta}^1(t) \text{ for } j\neq d, \underline{\theta}^d(\theta^{-d},t)\leq\theta^d\leq\overline{\theta}^d(\theta^{-d},t)\},$$

---

[13] Note that the claim is about the most informative equilibrium. As is well known in the literature, there is always a babbling equilibrium in which no information is transmitted.

where $\underline{\theta}^1(t) = a - 2bt/3$, $\overline{\theta}^1(t) = a + bt/3$, and $-\underline{\theta}^j(t) = \overline{\theta}^j(t) = bt/2$ for each $j \neq 1, d$, and, for each θ^{-d}, $\underline{\theta}^d(\theta^{-d}, t)$ and $\overline{\theta}^d(\theta^{-d}, t)$ are such that

$$\int_{\underline{\theta}^d(\theta^{-d}, t)}^{\overline{\theta}^d(\theta^{-d}, t)} f(\theta) d\theta^d = bt \quad \text{and} \quad \int_{\underline{\theta}^d(\theta^{-d}, t)}^{\overline{\theta}^d(\theta^{-d}, t)} \theta^d f(\theta) d\theta^d = 0.$$

Note that if a and b are small enough, then $D(1) \subseteq \Theta$ and, for each $t \in [0, 1]$ and each θ^1 sufficiently close to 0, $\underline{\theta}^d(\theta^{-d}, t)$ and $\overline{\theta}^d(\theta^{-d}, t)$ exist uniquely. Since $\underline{\theta}^d(\theta^{-d}, t)$ and $\overline{\theta}^d(\theta^{-d}, t)$ depend on θ^{-d} continuously, $D(t)$ is a closed set. Let $\partial D(t)$ denote the boundary of $D(t)$ and $D^{-d}(t) = \prod_{j \neq d} [\underline{\theta}^j(t), \overline{\theta}^j(t)]$.

Next, we compute $E(\theta | \theta \in \partial D(t))$ for $t \in [0, 1)$, which is equal to the limit

$$\lim_{t' \searrow t} E(\theta | \theta \in D(t') \setminus D(t)) = \lim_{t' \searrow t} \frac{E(\theta : D(t')) - E(\theta : D(t))}{P(D(t')) - P(D(t))}$$

for almost every $t \in [0, 1)$, where $E(X : A) = P(A)E(X|A)$.[14] For every $j \neq 1, d$, we have

$$E(\theta^j : D(t)) = \int_{D(t)} \theta^j f(\theta) d\theta$$

$$= \int_{D^{-d}(t)} \theta^j \int_{\underline{\theta}^d(\theta^{-d}, t)}^{\overline{\theta}^d(\theta^{-d}, t)} f(\theta) d\theta^d d\theta^{-d}$$

$$= \int_{D^{-d}(t)} \theta^j bt \, d\theta^{-d} = 0$$

and

$$E(\theta^d : D(t)) = \int_{D(t)} \theta^d f(\theta) d\theta$$

$$= \int_{D^{-d}(t)} \int_{\underline{\theta}^d(\theta^{-d}, t)}^{\overline{\theta}^d(\theta^{-d}, t)} \theta^d f(\theta) d\theta^d d\theta^{-d} = 0.$$

Thus $E(\theta^j | \theta \in \partial D(t)) = 0$ for every $j \neq 1$ and almost every $t \in [0, 1)$. For the first component, we have

$$E(\theta^1 : D(t)) = \int_{D(t)} \theta^1 f(\theta) d\theta$$

$$= \int_{D^{-d}(t)} \theta^1 \int_{\underline{\theta}^d(\theta^{-d}, t)}^{\overline{\theta}^d(\theta^{-d}, t)} f(\theta) d\theta^d d\theta^{-d}$$

$$= \int_{D^{-d}(t)} \theta^1 bt \, d\theta^{-d} = \left(a - \tfrac{1}{6} bt \right) \times (bt)^d$$

---

[14] Since $\partial D(t)$ has Lebesgue measure 0 on $\mathbb{R}^d$, $E(\theta | \theta \in \partial D(t))$ is uniquely determined only up to a null set of t.

and

$$P(D(t)) = \int_{D(t)} f(\theta) \, d\theta$$

$$= \int_{D^{-d}(t)} \int_{\underline{\theta}^d(\theta^{-d}, t)}^{\overline{\theta}^d(\theta^{-d}, t)} f(\theta) \, d\theta^d \, d\theta^{-d}$$

$$= \int_{D^{-d}(t)} bt \, d\theta^{-d} = (bt)^d.$$

Thus we have

$$\lim_{t' \searrow t} \frac{E(\theta^1 : D(t')) - E(\theta^1 : D(t))}{P(D(t')) - P(D(t))} = \frac{\frac{d}{dt} E(\theta^1 : D(t))}{\frac{d}{dt} P(D(t))} = a - \frac{d+1}{6d} bt.$$

Here we define

$$\overline{\mu}(t) = \left(a - \frac{d+1}{6d} bt, 0, \dots, 0 \right)$$

for every $t \in [0, 1]$. Note that $\overline{\mu}^1(t)$ is decreasing in t.

Since $E(\theta) = (0, \dots, 0)$, we have

$$E(\theta | \theta \notin D(1)) = \frac{E(\theta) - E(\theta : D(1))}{1 - P(D(1))}$$

$$= \left(-\frac{b^d}{1 - b^d} \left(a - \frac{1}{6} b \right), 0, \dots, 0 \right).$$

We choose $a = [(d+1)b - b^{d+1}]/(6d)$ so that $E(\theta | \theta \notin D(1)) = \overline{\mu}(1)$.

Since each sender has a bias toward the positive direction in the first component of the state and $\overline{\mu}^1(t)$ is decreasing in t, we choose b small enough (hence a is also small) so that, at any state $\theta \in D(1)$, both senders prefer $\overline{\mu}(t)$ to $\overline{\mu}(t')$ whenever $0 \le t < t' \le 1$.

Now we construct the following strategy profile. If the true state θ is outside $D(1)$ or on $\partial D(1)$, each sender sends the message "1." If the true state θ is in the interior of $D(1)$, each sender sends the message "t" such that $\theta \in \partial D(t)$. If two senders send messages t_1 and t_2, then the receiver chooses the policy $\overline{\mu}(\max(t_1, t_2))$.

Along the equilibrium path, the receiver is sequentially rational. If the true state θ is on $\partial D(t)$, sender i prefers the policy $\overline{\mu}^1(t)$ to any other policy $\overline{\mu}^1(t')$ with $t' > t$, so that, given that sender $j \ne i$ follows the above strategy, it is optimal for sender i to send a message smaller than or equal to t. If the true state θ is outside $D(1)$, then there is no deviation by a single sender that affects the receiver's action. Thus both senders are sequentially rational along the equilibrium path. Thus the above strategy profile is a perfect Bayesian equilibrium. Note that this strategy profile is continuous on the diagonal. □

In the proof, we divide the state space into uncountably many regions such that (i) as θ moves, the region changes continuously, (ii) both senders prefer the conditional

mean of regions with smaller parameters. Then we define the following strategy profile: each sender reports the region parameter, and the receiver believes the higher region parameter. As in Proposition 3, this is an equilibrium due to (ii). Diagonal continuity follows from (i).

Note that the equilibrium constructed above for biases (x_1, x_2) remains an equilibrium for biases $(t_1 x_1, t_2 x_2)$ with any t_1, $t_2 \geq 1$. Therefore even in the limit, as the magnitude of the biases go to infinity, information revelation can be bounded away from zero. This is in contrast to the one-sender case. Crawford and Sobel (1982) show that there is no informative equilibrium for large enough biases if the state space is a compact interval. In multidimensional environments, Levy and Razin (2007) provide a condition for the receiver to play at most k actions with positive probability if the magnitude of the bias is sufficiently large. If this condition holds with $k = 1$, then there is no informative equilibrium for a large enough bias.[15]

For the case of exactly opposite biases, an earlier version of this paper contains the construction of an informative perfect Bayesian equilibrium that is not necessarily diagonally continuous. We do not know if a diagonally-continuous informative equilibrium exists for a general multidimensional state space with opposite biases.[16]

5. Discussion and extensions

5.1 *Long cheap talk*

It is well known that multiple rounds of cheap talk can expand the set of equilibrium payoffs (Aumann and Hart 2003 and Krishna and Morgan 2004). In our model, there might be a fully-revealing equilibrium with multiple rounds of cheap talk, even if there is no such equilibrium with one round of cheap talk.[17] The intuition is that in a game with multiple rounds of cheap talk even if on the equilibrium path players do not mix in any payoff relevant manner (which is necessary for fully-revealing equilibrium), they might do so off the equilibrium path. This means that deviations by the senders can lead to stochastic outcomes, which provides new ways of deterring deviations by senders. Below we show that similar techniques to the ones we used before can be used to derive a necessary condition for the existence of fully-revealing equilibrium in a game with multiple rounds of cheap talk.

Let $D = \text{diam}(\Theta)/2$, where $\text{diam}(\Theta) = \sup_{\theta, \theta' \in \Theta} |\theta - \theta'|$. For $i = 1, 2$, let $r_i = \sqrt{\max(0, |x_i|^2 - D^2)}$.

---

[15]It is not true though that informative equilibria never exist for large enough biases. Chakraborty and Harbaugh (2007) construct an informative equilibrium in symmetric multidimensional environments. They also show that this equilibrium construction is generically robust to small asymmetries of payoff functions and the prior distribution.

[16]We do know, though, that in the limit as the magnitudes of biases go to infinity, all actions taken by the receiver have to be on a lower-dimensional hyperplane through the expectation of the state space that is orthogonal to the biases. A proof of this result is contained in a supplementary file on the journal website, http://econtheory.org/supp/334/supplement.pdf.

[17]A related result in Krishna and Morgan (2001b) is that if two senders send messages sequentially, then introducing a second round of cheap talk in some cases improves the possibility of full revelation (Proposition 5). On the other hand, Krishna and Morgan (2004) prove that with one sender all equilibria with multiple rounds of communication are bounded away from full revelation (Proposition 4).

PROPOSITION 13. *In any game with multiple rounds of cheap talk, there exists no fully-revealing equilibrium if there exist $\theta, \theta' \in \Theta$ such that $B(\theta' + x_1, r_1) \cup B(\theta + x_2, r_2) \supseteq \mathrm{co}(\Theta)$.*

PROOF.   In a fully-revealing equilibrium, for any pair of states θ and θ', player 1 in state θ' cannot gain by deviating to playing what her strategy prescribes in state θ, and in state θ cannot gain by deviating to playing what her strategy prescribes in state θ'. Fix any strategy profile for which in every state the policy outcome is equal to the state. Let $y(\theta, \theta')$ denote the probability distribution of policy outcomes resulting from sender 1 playing the continuation strategy that the strategy profile prescribes for her after observing θ and from sender 2 playing the continuation strategy that the strategy profile prescribes for her after observing θ'. Then since the strategy profile is an equilibrium, we have $-E(y(\theta, \theta') - \theta' - x_1)^2 \le -|x_1|^2$. Note that $-E(y(\theta, \theta') - \theta' - x_1)^2 = -(Ey(\theta, \theta') - \theta' - x_1)^2 - E|y(\theta, \theta') - Ey(\theta, \theta')|^2$. Since $y(\theta, \theta')$ is a distribution over $\mathrm{co}(\Theta)$, $E|y(\theta, \theta') - Ey(\theta, \theta')|^2 \le (\mathrm{diam}(\Theta)/2)^2 = D^2$. This means that a necessary condition for the strategy profile to be an equilibrium is $(Ey(\theta, \theta') - \theta' - x_1)^2 > |x_1|^2 - D^2$. A symmetric argument establishes that another necessary condition is $(Ey(\theta, \theta') - \theta - |x_2|)^2 \ge |x_2|^2 - D^2$. Combining the two conditions yields $Ey(\theta, \theta') \notin B(\theta' + x_1, r_1) \cup B(\theta + x_2, r_2)$. Therefore, $B(\theta' + x_1, r_1) \cup B(\theta + x_2, r_2) \supseteq \mathrm{co}(\Theta)$ for some $\theta, \theta' \in \Theta$ implies that there does not exist a fully-revealing equilibrium.   □

This result is similar in spirit to Proposition 2: if a sender pretends to have observed a state different from the true state, then the resulting probability distribution over outcomes should yield a lower expected utility for her than revealing the true state. For quadratic utilities, this expected utility depends only on the expectation and the variance of the resulting distribution. The variance of the distribution is bounded by a constant that depends on the diameter of the state space. This can be used to show that the expected value of the distribution has to be in the two open balls in the statement, $B(\theta' + x_1, r_1)$ and $B(\theta + x_2, r_2)$ (if player 1 played as if she observed θ and player 2 played as if she observed θ').

We conclude this subsection by showing that in a bounded state space, for any fixed pair of directions of biases, in the limit as the magnitude of biases go to infinity there exists a fully-revealing equilibrium in a game with an arbitrary number of rounds of communication if and only if there exists such an equilibrium in a game with only one round of communication. This means that the results of Section 3.3 on large enough biases hold for games with an arbitrary number of rounds of communication. The key insight is that the open balls in Proposition 13 converge to the ones in Proposition 2.

PROPOSITION 14. *Fix a compact state space Θ and directions of biases $z_1, z_2 \in S^{d-1}$. If there exists $t^* \in \mathbb{R}_+$ such that for every $t_1, t_2 > t^*$ and bias pair $(x_1, x_2) = (t_1 z_1, t_2 z_2)$ there exists no fully-revealing equilibrium in a game with one round of cheap talk, then there exists $t^{**} \in \mathbb{R}_+$ such that for every $t_1, t_2 > t^{**}$ and bias pair $(x_1, x_2) = (t_1 z_1, t_2 z_2)$ there exists no fully-revealing equilibrium in a game with an arbitrary number of rounds of cheap talk.*

PROOF.   Let $r_i(t_i) = \sqrt{\max(0, |t_i z_i|^2 - D^2)}$ for $i = 1, 2$. Note that θ' is not on the boundary of $B(\theta' + t z_1, r_1(t_1))$, but the difference between θ' and $B(\theta' + t z_1, r_1(t_1))$

is $|t_1| - r_1(t_1) = |t_1| - \sqrt{t_1^2 - D^2} = D^2/[t_1 + \sqrt{t_1^2 - D^2}]$ for large enough t_1, which goes to 0 as $t_1 \to \infty$. A symmetric argument shows that $|t_2| - r_2(t_2) \to 0$ as $t_2 \to \infty$. Given this, the same arguments as in Proposition 7 establish that for any θ, $\theta' \in \Theta$, we have $B(\theta' + t z_1, r_1(t_1)) \cup B(\theta + t z_2, r_2(t_2)) \not\supseteq \mathrm{co}(\Theta)$ for all t_1, $t_2 \in \mathbb{R}_+$ if and only if $H^\circ(z_1, z_1 \cdot \theta') \cup H^\circ(z_2, z_2 \cdot \theta) \not\supseteq \mathrm{co}(\Theta)$. The claim then follows from Propositions 7 and 13.   □

5.2 *Mixed strategies in fully-revealing equilibrium*

As mentioned before, in equilibrium the receiver never uses a nondegenerate mixed strategy. Moreover, in a fully-revealing equilibrium, for every $\theta \in \Theta$ and for almost all (m_1, m_2) such that $m_1 \in \mathrm{supp}\, s_1(\theta)$ and $m_2 \in \mathrm{supp}\, s_2(\theta)$ we have $y(m_1, m_2) = \theta$. That is, the outcome of the mixing along the equilibrium path is payoff-irrelevant. Nevertheless, allowing for mixed strategies by the senders can facilitate fully-revealing equilibria in cases when there is no fully-revealing equilibrium in pure strategies. This is for exactly the same reason that multiple rounds of cheap talk can create new equilibria relative to a single round: namely, deviations might lead to randomness in the action chosen by the receiver, which can be an extra deterrent for deviations, given that senders have concave utility functions. To see this, note that although along the equilibrium path the outcome of the randomization of one sender is payoff-irrelevant to the other sender, the same is not necessarily true after deviations.

It is easy to see though that the propositions in the previous subsection hold for the case of a single round of cheap talk when the senders use mixed strategies. The condition for fully-revealing equilibrium in Proposition 13 remains a necessary condition for a fully-revealing equilibrium in this setting, and a result similar to Proposition 14 holds: if the magnitude of the biases goes to infinity, the set of equilibrium payoffs supported by pure strategies and the set of equilibrium payoffs supported by mixed strategies converge to the same limit.

6. Conclusion

This paper argues that in a cheap talk model with multiple senders, the amount of information that can be transmitted in equilibrium depends not on the dimensionality of the state space but on the finer details of the specification of the model. These details include the shape of the boundary of the state space and the similarity of the senders' preferences, where similarity is defined with respect to the state space. It is worth pointing out that the properties of the state space and the senders' preferences cannot be investigated independently once we allow for general (state-dependent) preferences. For example, an open bounded state space with state-independent preferences can be transformed into an unbounded state space with state-dependent preferences in a way that the resulting games are strategically equivalent.

In future work we would like to depart from the assumption made in most of the literature, including this paper, that senders observe the state perfectly. Introducing noise into the senders' information makes the cheap talk model more realistic and potentially

affects the qualitative conclusions.[18] For the latter reason, we think it is an important avenue for future research. It is also a challenging one for general state spaces, since techniques from the existing literature cannot be used, even to investigate the existence of fully-revealing equilibrium.

APPENDIX: CONSISTENCY OF BELIEFS AND DIAGONAL CONTINUITY

In this Appendix, we show that if we restrict attention to strategies that satisfy some regularity conditions, then every equilibrium in which the receivers' beliefs are consistent satisfies diagonal continuity (as defined in Section 4.1).

Consider a PBE (s_1, s_2, y) and conditional beliefs μ of the sender that support this equilibrium. In order to check for consistency of the beliefs, we need to define models in which the observations of senders are noisy. We consider a sequence of noisy models indexed by $k = 1, 2, \ldots$. In the noisy model indexed by k, senders 1 and 2 observe the signals $\theta_1 \in \Theta$ and $\theta_2 \in \Theta$ respectively. For each true state $\theta \in \Theta$, the joint density function of signals (θ_1, θ_2) conditional on θ is given by $g^k(\theta_1, \theta_2 | \theta)$. We assume that the noise disappears in the limit: $g^k(\theta_1, \theta_2 | \theta)$ converges in probability to (θ, θ) as $k \to \infty$.

An example of this construction, which is similar in spirit to the one proposed in Battaglini (2004), is given by

$$\theta_i = \theta + \varepsilon_k u_i,$$

where (u_1, u_2) is a truncated standard normal distribution on $\mathbb{R}^{2d}$ and $\varepsilon_k \to 0$ as $k \to \infty$. Truncation is needed to ensure that θ_i belongs to Θ.[19]

Fixing the senders' strategies in the sequence of noisy models to be $s_i(\theta_i)$, let $\mu^k(m_1, m_2)$ denote the posterior belief of the receiver in the model indexed by k, given the reports (m_1, m_2). Let $\overline{\mu}^k(m_1, m_2)$ be the expectation of θ with respect to $\mu^k(m_1, m_2)$.

DEFINITION 3. We say μ is *consistent* if $\mu^k(m_1, m_2)$ weakly converges to $\mu(m_1, m_2)$ uniformly over $(m_1, m_2) \in s_1(\Theta) \times s_2(\Theta)$, i.e., for any $\varepsilon > 0$ and any continuous and bounded function b on Θ, there exists K such that

$$\left| \int b(\theta) \mu^k(m_1, m_2) \, d\theta - \int b(\theta) \mu(m_1, m_2) \, d\theta \right| < \varepsilon$$

for any $(m_1, m_2) \in s_1(\Theta) \times s_2(\Theta)$ and any $k > K$.[20]

If Θ is bounded, then this definition implies that $\overline{\mu}^k(m_1, m_2)$ uniformly converges to $\overline{\mu}(m_1, m_2)$ as $k \to \infty$.

To show our main result concerning consistent beliefs in the limit model, we first establish a result that applies to beliefs in the noisy models defined above. We show that

---

[18]For existing work along these lines, see Wolinsky (2002) and Battaglini (2004).

[19]On the other hand, noise structures like the one in Battaglini (2002, Section 3) are not compatible with our framework because they do not admit a density function $g^k(\theta_1, \theta_2 | \theta)$.

[20]We note that the requirement of uniform convergence is strong. We do not know whether consistency implies diagonal continuity if we use pointwise convergence.

$\overline{\mu}^k(m_1, m_2)$ is continuous in (m_1, m_2) for any k. Intuitively speaking, in a noisy model, even if the receiver gets two slightly different pairs of messages, she does not drastically change her belief about the true state, for the difference between the message pairs does not necessarily mean a drastic difference in the true state, but means a small change in the noise contained in the senders' signals. Once we establish the continuity of $\overline{\mu}^k$, we show that continuity is inherited by $\overline{\mu}$ in the limit model without noise, which implies diagonal continuity when the reporting functions s_i are continuous.

In order to use Bayes' rule for continuous random variables, we impose several restrictions on the senders' reporting functions. For each i, the message space M_i is a subset of a Euclidean space $\mathbb{R}^{n_i}$ and each inverse image of message m_i with respect to s_i, $s_i^{-1}(m_i) = \{\theta_i \in \Theta \mid s_i(\theta_i) = m_i\}$, is parameterized by $t_i \in T_i \subseteq \mathbb{R}^{d-n_i}$. That is to say, there exists a continuously differentiable bijection

$$h_i \colon M_i \times T_i \to \Theta$$

such that $m_i = s_i(\theta_i)$ if and only if $\theta_i = h_i(m_i, t_i)$ for some $t_i \in T_i$.[21]

Given (h_1, h_2), the density function of (m_1, m_2) with respect to the Lebesgue measure on $M_1 \times M_2$ conditional on the true state θ is

$$\int_{T_2} \int_{T_1} g^k(h_1(m_1, t_1), h_2(m_2, t_2) \mid \theta) \mid J_1(m_1, t_1) J_2(m_2, t_2) \mid d\,t_1 d\,t_2,$$

where $J_i(m_i, t_i)$ is the Jacobian of h_i at (m_i, t_i):

$$J_i(m_i, t_i) = \det \frac{\partial h_i(m_i, t_i)}{\partial(m_i, t_i)}.$$

PROPOSITION 15. *Suppose*

(i) Θ, T_1, and T_2 are compact

(ii) $g^k(\theta_1, \theta_2 \mid \theta)$ is continuous in (θ_1, θ_2), $g^k(\theta_1, \theta_2 \mid \theta) > 0$, and bounded

(iii) for each sender i, $J_i(m_i, t_i)$ is continuous in m_i, $J_i(m_i, t_i) \neq 0$, and $J_i(m_i, t_i)$ is bounded.

Then the expectation $\overline{\mu}^k(m_1, m_2)$ of θ conditional on (m_1, m_2) in the k-th noisy model is continuous in (m_1, m_2).

PROOF.   The expectation $\overline{\mu}^k(m_1, m_2)$ is given by

$$\frac{E\left[\theta \int_{T_1} \int_{T_2} g^k(h_1(m_1, t_1), h_2(m_2, t_2) \mid \theta) \mid J_1(m_1, t_1) J_2(m_2, t_2) \mid d\,t_1 d\,t_2\right]}{E\left[\int_{T_2} \int_{T_1} g^k(h_1(m_1, t_1), h_2(m_2, t_2) \mid \theta) \mid J_1(m_1, t_1) J_2(m_2, t_2) \mid d\,t_1 d\,t_2\right]}.$$

---

[21] For example, in Battaglini's (2002) equilibrium construction, h_i is the identity function on $\mathbb{R}^d$; M_1 and M_2 are subspaces of $\mathbb{R}^d$ that form a coordinate system: every point $\theta \in \mathbb{R}^d$ is uniquely expressed by a linear combination of $m_1 \in M_1$ and $m_2 \in M_2$; M_i contains sender j's bias direction; and $T_i = M_j$. Such a coordinate system exists if $d \geq 2$ and the two senders' biases are not parallel.

The denominator is nonzero. Also, by the Lebesgue Convergence Theorem, both the numerator and the denominator are continuous in (m_1, m_2).[22] Therefore, $\overline{\mu}^k(m_1, m_2)$ is continuous with respect to (m_1, m_2).   □

PROPOSITION 16. *Let (s_1, s_2, y) be an equilibrium in the limit game. In addition to the assumptions in Proposition 15, suppose that $m_i = s_i(\theta_i)$ is continuous in θ_i for $i = 1$, 2. Then every equilibrium that is supported by a consistent belief is continuous on the diagonal.*

PROOF.   By Proposition 15, $\overline{\mu}^k(m_1, m_2)$ is continuous in (m_1, m_2). Since $\overline{\mu}^k(m_1, m_2)$ converges to $\overline{\mu}(m_1, m_2)$ uniformly over (m_1, m_2), $\mu(m_1, m_2)$ is also continuous in (m_1, m_2), and hence $\overline{\mu}(s_1(\theta_1), s_2(\theta_2))$ is continuous in (θ_1, θ_2).   □

REFERENCES

Aumann, Robert J. and Sergiu Hart (2003), "Long cheap talk." *Econometrica*, 71, 1619–1660. [21]

Austen-Smith, David (1993), "Interested experts and policy advice: multiple referrals under open rule." *Games and Economic Behavior*, 5, 3–43. [1]

Battaglini, Marco (2002), "Multiple referrals and multidimensional cheap talk." *Econometrica*, 70, 1379–1401. [1, 2, 4, 6, 11, 15, 16, 24, 25]

Battaglini, Marco (2004), "Policy advice with imperfectly informed experts." *Advances in Theoretical Economics*, 4, Article 1. [15, 24]

Chakraborty, Archishman and Rick Harbaugh (2007), "Comparative cheap talk." *Journal of Economic Theory*, 132, 70–94. [21]

Crawford, Vincent P. and Joel Sobel (1982), "Strategic information transmission." *Econometrica*, 50, 1431–1451. [1, 4, 11, 21]

Gilligan, Thomas W. and Keith Krehbiel (1989), "Asymmetric information and legislative rules with a heterogeneous committee." *American Journal of Political Science*, 33, 459–490. [1]

Krishna, Vijay and John Morgan (2001a), "Asymmetric information and legislative rules: Some amendments." *American Political Science Review*, 95, 435–452. [1, 11]

Krishna, Vijay and John Morgan (2001b), "A model of expertise." *Quarterly Journal of Economics*, 116, 747–775. [1, 11, 21]

Krishna, Vijay and John Morgan (2004), "The art of conversation: eliciting information from experts through multi-stage communication." *Journal of Economic Theory*, 117, 147–179. [21]

---

[22]See Royden (1988, Theorem 16 of Chapter 4).

Levy, Gilat and Ronny Razin (2007), "On the limits of communication in multidimensional cheap talk: a comment." *Econometrica*, 75, 885–893. [4, 11, 21]

Milgrom, Paul and John Roberts (1986), "Relying on the information of interested parties." *RAND Journal of Economics*, 17, 18–32. [1]

Royden, Halsey L. (1988), *Real Analysis*, 3rd edition. Macmillan, New York. [16, 26]

Wolinsky, Asher (2002), "Eliciting information from multiple experts." *Games and Economic Behavior*, 41, 141–160. [24]