DEPARTMENT OF ECONOMICS

UNIVERSITY OF OXFORD

Department of Economics Discussion Paper Series

Robustness of Full Revelation in Multisender Cheap Talk

Margaret Meyer, Inés Moreno de Barreda and Julia Nafziger

Number 789
Originally Published: August 2016
Revised: January 2019

Department of Economics
Manor Road Building, Manor Road
Oxford, OX1 3UQ

ISSN 1471-0498

Robustness of Full Revelation in Multisender Cheap Talk*

Margaret Meyer, Inés Moreno de Barreda and Julia Nafziger

First draft: February 2013 This draft: January 12, 2019

Abstract

This paper studies information transmission in a two-sender, multidimensional cheap talk setting where there are exogenous constraints on the (convex) feasible set of policies for the receiver and where the receiver is uncertain about both the directions and the magnitudes of the senders' bias vectors. With the supports of the biases represented by cones, we prove that whenever there exists an equilibrium which fully reveals the state (a FRE), there exists a robust FRE, i.e. one in which small deviations result in only small punishments. We provide a geometric condition, the Local Deterrence Condition, relating the cones of the biases to the frontier of the policy space, that is necessary and sufficient for the existence of a FRE. We also construct a specific policy rule for the receiver, the Min Rule, that supports a robust FRE whenever one exists.

*Meyer: Nuffield College and Dept. of Economics, Oxford University and CEPR, margaret.meyer@nuffield.ox.ac.uk. Moreno de Barreda: Dept. of Economics, Oxford University, ines.morenodebarreda@economics.ox.ac.uk. Nafziger: Dept. of Economics and Business, Aarhus University, jnafziger@econ.au.dk. We would like to thank the Editor, George Mailath, and three referees for extremely thoughtful and detailed comments. We also thank Elizabeth Baldwin, Tom Cunningham, Péter Esö, Navin Kartik, Kohei Kawamura, Paul Klemperer, Patrick Legros, Gilat Levy, John Morgan, Ronny Razin, Larry Samuelson, Joel Sobel, and Peter Norman Sørensen for helpful comments, as well as seminar audiences at Copenhagen, ESSET Gerzensee, LSE, Oxford, the Spanish Economic Association, the Transatlantic Theory Workshop, UAB, and ESEM Toulouse.

1 Introduction

In sender-receiver games with cheap talk, the decision-maker (receiver) has imperfect information about the consequences of a policy (the state of the world) and elicits reports from better-informed experts (the senders), whose preferences are not perfectly aligned with those of the decision-maker (i.e. the experts are "biased"). The advice transmitted by the senders is costless but unverifiable (hence, "cheap talk"), and the receiver cannot commit herself in advance to how she will respond to the senders' advice.[1] Cheap talk games with two biased experts have been used, for example, in organizational economics to analyze the interaction between a CEO and division managers, and in political science to study the transmission of information from legislative committees to the legislature as a whole.[2]

Our objective in this paper is to study the combined impact on information transmission of two frictions that are common in many of the environments which cheap talk models are designed to represent. The first such friction is constraints on the feasible set of policies. In the context of resource allocation by managers or legislatures, these constraints arise from limited budgets. In the context of organizational downsizing or restructuring, they may arise from legal or institutional constraints. More generally, there almost always exist physical constraints on what policies can be implemented within any given time frame. The second friction is uncertainty on the part of the receiver about exactly how and to what extent the preferences of the senders differ from her own. At any given time, the receiver may be unsure about a range of factors that affect how the senders evaluate the different possible decisions, and these privately known factors will affect the senders' incentives when communicating with the receiver.

When the receiver needs to elicit information from biased senders, constraints on feasible policies can generate an informational inefficiency. By limiting the receiver's potential responses to the senders' reports, such constraints can destroy the senders' willingness to truthfully reveal their information. And the more uncertain the receiver is about the senders' preferences, the more likely it is that such constraints will prevent full extraction of the senders' information. We analyze the conditions under which biased senders have incentives to truthfully reveal the information the receiver seeks, despite their privately known biases and despite feasibility constraints in effect tying the receiver's hands with respect to punishments for misreporting.

In our two-sender model of simultaneous cheap talk, both senders (but not the receiver) observe the p-dimensional state of the world. All of the players have quadratic utility functions, and sender i's ideal p-dimensional policy differs from the receiver's by a vector, b_i, sender i's bias vector. Each sender is privately informed about his bias vector, and biases are distributed independently of the state of the world. We assume that the support of each

[1]For the seminal paper in this literature see Crawford and Sobel (1982).

[2]For theoretical models of multisender cheap talk, see Battaglini (2002) and Ambrus and Takahashi (2008). For the former application, see Alonso and Matouschek (2008), and for the latter, Gilligan and Krehbiel (1989) and Krishna and Morgan (2001a;b).

b_i is a *convex cone*; this implies that the receiver is uncertain both about the directions of the biases and about the magnitudes, which may be arbitrarily large. The set of feasible policies for the receiver is a closed, convex, p-dimensional subset of $\mathbb{R}^p$, which may or may not be compact. We study the existence of equilibria in which the receiver fully extracts the senders' information about the state of the world, using a strategy with the desirable property that it is *robust* to small mistakes by the senders, in that small discrepancies between their reports result in only small punishments by the receiver. We call such an equilibrium a robust fully revealing equilibrium (robust FRE).

We prove three main sets of results about the existence of a robust FRE when the receiver is uncertain about the senders' biases and the policy space is exogenously restricted. First, we show that whenever there exists a FRE, there also exists a robust FRE. For any given pair of incompatible reports about the state by the senders, the cones representing the supports of the senders' biases pin down the directions, relative to the reports, in which an effective punishment by the receiver must lie. The convexity of the policy space then guarantees that if there is any feasible punishment in the right directions, there must be a small feasible punishment. Thus, our focus on robust strategies for the receiver does not limit the conditions under which a FRE exists.

Second, we provide geometric characterizations of the necessary and sufficient conditions for existence of a (robust) FRE. These are conditions on the shape of the frontier of the feasible policy space, relative to the cones of the senders' biases. We provide separate characterizations for the cases of deterministic and uncertain bias directions and show that, with uncertain directions, a FRE exists if and only if a FRE would exist for each possible pair of deterministic directions. This result is more subtle than it may at first appear. If the receiver knew the directions of the senders' biases, her strategy in response to incompatible reports could depend on these directions. But with uncertain bias directions, a FRE requires that the receiver's response to incompatible reports *must* constitute a punishment for *all* possible realizations of these directions.

The geometric characterizations can be interpreted as "local deterrence conditions". They guarantee, for appropriate points on the frontier of the policy space, that small deviations from truthful reporting by either sender can be deterred with a local (i.e. robust) punishment. Despite their local nature we prove that these conditions are sufficient for all deviations, including large ones, to be deterred. These geometric conditions are, moreover, easy to check.

Our final set of results shows how to construct a robust FRE, both for deterministic and for uncertain bias directions. We provide a strategy for the receiver, which we call the Min Rule, that implements a robust FRE whenever one exists. The Min Rule takes a very intuitive and appealing form when the bias directions are deterministic or when both the policy space and the supports of the biases are two-dimensional. In both of these cases, there are exactly two dimensions in which the senders' interests conflict with the receiver's. In these dimensions, given any incompatible reports, the policy selected by the Min Rule constitutes the anonymous punishment which is *least severe* for each sender, subject to deterring both of them from misreporting the state, whatever the realizations of their biases.[3]

The rest of the paper is organized as follows. The next subsection briefly reviews the related literature. Section 2 presents the model and some preliminary results. Section 3 proves the existence of a robust FRE whenever there exists a FRE. In Section 4, we provide the geometric characterizations of the necessary and sufficient conditions for existence of a (robust) FRE. Finally, Section 5 shows how to construct a robust FRE, using our Min Rule, whenever one exists.

1.1 Related Literature

The closest papers to ours are Battaglini (2002), Ambrus and Takahashi (2008), and Krishna and Morgan (2001a), in all of which a receiver consults two equally informed senders.

Assuming that the policy space is the whole of $\mathbb{R}^p$, Battaglini constructs a FRE in which each sender has incentives to report truthfully, because his influence over the receiver's policy choice is limited to dimensions orthogonal to his bias vector. In such dimensions, there is no conflict of interest between a sender and the receiver. This construction supports a FRE that is independent of the magnitudes of the biases and also robust to small mistakes. However, this construction breaks down in the presence of the two frictions on which we focus, constraints on the feasible set of policies and uncertainty on the part of the receiver about the directions of the senders' bias vectors. Section 5.1 discusses in more detail the contrast between our construction of a robust FRE, using our Min Rule, and Battaglini's approach.

Ambrus and Takahashi (2008) analyze the implications of restricted policy (or state) spaces for the existence of FRE. They show that when the magnitudes of the biases are sufficiently small, it is always possible to construct a FRE. However, as the magnitudes of the biases increase, the restrictions on the policy space might make it impossible to deter deviations by the senders. For compact policy spaces, they show that there exists a FRE for arbitrarily large magnitudes of the biases if and only if, as the biases become large, the senders have a common least-preferred policy; the receiver could use such a common least-preferred policy as a response to any deviation. Remarking that the use of extreme punishments after even small deviations is unappealing, since such deviations could in practice arise from small mistakes, Ambrus and Takahashi introduce a robustness concept called *continuity on the diagonal*, which is equivalent to our definition of robustness for the case of known biases. For a policy space whose frontier is everywhere smooth, they show that there does not exist such a robust FRE, even when the magnitudes of the biases are small.

Motivated by the same remark, our paper focuses on the existence and construction of robust FRE. Our first main result, Proposition 3, implies, in the special case when the receiver is uncertain about only the magnitudes (not the directions) of the biases and these can be arbitrarily large, that whenever there exists a FRE there also exists a robust FRE. For

[3] An anonymous punishment strategy for the receiver, formally defined in Section 2.2, is one that does not depend on which sender sent which report.

compact policy spaces, this means that when there exists a common least-preferred policy with which deviations can be punished, it is also possible to punish deviations locally. This positive result does not contradict Ambrus and Takahashi (2008)'s negative one for smooth policy spaces, because the condition for existence of a common least-preferred policy *requires* the existence of kinks in the frontier of the policy space.

For convex, closed but unbounded (and hence non-compact) policy spaces, our first characterization result for the existence of (robust) FRE, Proposition 4, which also assumes deterministic bias directions, shows that a robust FRE might exist even if there does not exist a common least-preferred policy and even if the frontier of the policy space is smooth.

In contrast to the analysis in Ambrus and Takahashi (2008), we prove a characterization result for the existence of a (robust) FRE when the receiver is uncertain about the directions of the senders' biases, and we show how to construct robust FRE whenever they exist.

In Krishna and Morgan (2001a), the policy space is one-dimensional. Though their analysis focuses on sequential communication by the senders, they make the following observation about the game with simultaneous reporting. If the senders' ideal points are both larger than the receiver's, then a FRE can be supported by the receiver choosing the smaller of the two reports. Our construction of a robust FRE using the Min Rule can be seen as a multidimensional generalization of Krishna and Morgan's observation.

Finally,[4] Ambrus and Lu (2014) and Rubanov (2015) construct equilibria in a unidimensional policy space that are arbitrarily close to full revelation. Both papers show that their equilibria survive the introduction of a small probability of the senders observing a random state which is independent of the true state. However, the equilibria in these two papers do not satisfy our robustness concept: The senders might observe states that are arbitrarily close to each other, yet the receiver's action in response to the equilibrium messages might be far away from their observations. Though we do not explicitly model the occurrence of mistakes, our robustness concept requires the receiver's response to be close to the senders' reports whenever these reports are themselves close, thus ensuring that small mistakes by the senders do not lead to large responses/punishments by the receiver.

2 The Model

We analyze a game of cheap talk between two senders, S_1, S_2, and a receiver, R. Both senders perfectly observe the *state* $\theta \in \Theta$, where Θ is a convex subset of $\mathbb{R}^p$ with $p \geq 1$. After observing θ, each sender S_i sends a costless and unverifiable message $m_i \in \mathcal{M}_i$ to the receiver, who then chooses a policy y from a closed set Y of feasible policies. We refer to Y as the *policy space* and throughout the paper we assume that $Y = \Theta$.[5]

[4]There is a small recent experimental literature on multi-sender cheap talk. Lai et al. (2015), in particular, discuss robustness, but because their state and policy spaces are discrete, our concept of robustness cannot be applied in their setting.

[5]In our discussion paper, Meyer et al. (2016), we analyzed the case in which $Y \subsetneq \Theta$. Given the utility function for the receiver specified in the next paragraph of the text, Θ is the set of potential ideal policies for the receiver, and $Y \subset \Theta$ means that only a subset of these policies are feasible. We showed in Proposition 2 of Meyer et al. (2016) that given a fixed Y, existence of a (robust) FRE for $\Theta = Y$ is necessary and sufficient for existence of a (robust)

Given the state θ and the chosen policy y, the receiver's utility is $u^R(y,\theta) = -|y-\theta|^2$ and sender S_i's utility is $u^{S_i}(y,\theta,b_i) = -|y-\theta-b_i|^2$, where $|\cdot|$ denotes the Euclidean norm in $\mathbb{R}^p$. Given these utilities, the ideal policy for the receiver is to match the state, whereas the optimal policy for sender S_i is $y = \theta + b_i$. The vector b_i is referred to as the bias vector of sender S_i and is private information to S_i. We denote by F the joint cumulative distribution function of (θ, b_1, b_2), and we assume that (θ, b_1, b_2) are mutually independent with full support on $\Theta \times C_1 \times C_2$, where $C_i \subseteq \mathbb{R}^p$ is a closed and convex cone in $\mathbb{R}^p$, i.e. for any $b_i, b_i' \in C_i$, $tb_i + t'b_i' \in C_i$ for all $t, t' \geq 0$. The assumption that the supports of the privately-observed biases are convex cones implies that the receiver is uncertain not only about the directions of the biases but also about their magnitudes, which may be arbitrarily large.

Given the private information of the senders, we can assume without loss of generality that the message space for S_i is $\mathcal{M}_i = \mathbb{R}^p \times C_i$. A pure strategy for sender S_i is a measurable function $s_i : \Theta \times C_i \longrightarrow \mathcal{M}_i$. A pure strategy for the receiver is a measurable function $y^R : \mathcal{M}_1 \times \mathcal{M}_2 \longrightarrow Y$. Given messages m_1, m_2, $\mu(m_1, m_2)$ denotes the receiver's belief about (θ, b_1, b_2) after receiving messages m_1, m_2.[6] We define

$$Inv^{s_1,s_2}(m_1,m_2) \equiv \{(\theta, b_1, b_2) \in \Theta \times C_1 \times C_2 \mid s_1(\theta, b_1) = m_1,\ s_2(\theta, b_2) = m_2\}.$$

$Inv^{s_1,s_2}(m_1,m_2)$ is the set of triples (θ, b_1, b_2) that lead to messages m_1, m_2 if the senders are using strategies s_1, s_2.

The equilibrium concept we use is *Perfect Bayesian Equilibrium.*

Definition 1. The strategies (s_1, s_2, y^R) constitute a *Perfect Bayesian Equilibrium* if there exists a belief function μ such that:

(i) $s_i(\theta, b_i)$ is optimal given s_{-i} and y^R, for any $\theta \in \Theta$ and $b_i \in C_i$, for $i \in \{1,2\}$.

(ii) $y^R(m_1,m_2)$ is optimal given $\mu(m_1,m_2)$ for each $(m_1,m_2) \in \mathcal{M}_1 \times \mathcal{M}_2$.

(iii) If $Inv^{s_1,s_2}(m_1,m_2) \neq \emptyset$, $\mu(m_1,m_2)$ puts probability one on $Inv^{s_1,s_2}(m_1,m_2)$. Moreover, if $Inv^{s_1,s_2}(m_1,m_2)$ has positive probability with respect to F, $\mu(m_1,m_2)$ is derived from Bayes' rule.

We focus on Perfect Bayesian Equilibria in which the receiver perfectly learns the state from the senders' messages. We say that the strategies (s_1, s_2) *fully reveal the state* if for any $\theta \in \Theta$ and any $b_1 \in C_1$, $b_2 \in C_2$, the receiver's marginal belief about the state given messages $s_1(\theta, b_1)$ and $s_2(\theta, b_2)$ puts mass one on θ. We call an equilibrium with strategies that fully reveal the state a *fully revealing equilibrium* (FRE).

FRE for $\Theta \supset Y$. For the case $\Theta \subset Y$ and given the receiver's utility function specified below, we can without loss of generality ignore all those policies in Y that are not in Θ, since no such policies could be best responses for the receiver given that Θ is convex.

[6]For clarity of exposition, the players are restricted to using pure strategies throughout the paper. Given the convexity of the policy space, it is always optimal for the receiver to use a pure strategy. Moreover, we can show that our assumption that the supports of the bias vectors are cones (and hence that their magnitudes may be arbitrarily large), coupled with the convexity of Y, implies that if there does not exist a (robust) FRE when the senders are restricted to pure strategies, then there does not exist a (robust) FRE when they are allowed to use mixed strategies. The argument is similar to that used to prove Lemma 1 below.

2.1 Robustness

Our goal is to characterize the conditions under which there exist equilibria that are fully revealing, despite the uncertainty about the biases, and that *also* satisfy an additional desirable property: robustness. Our definition of robustness is motivated by the possibility that the senders might make small mistakes, because they might not perceive the state perfectly accurately. In such situations, it is natural to require the receiver's response to be close to the response that would have resulted in the absence of such mistakes. We now formally define robustness.

Given $x \in \mathbb{R}^p$, $r > 0$, denote by $B(x, r) = \{\theta \in \mathbb{R}^p \mid |\theta - x| < r\}$ the open ball with centre x and radius r.

Definition 2. Given some strategies (s_1, s_2) that fully reveal the state, the receiver's strategy y^R is *robust* if for any $\theta \in \Theta$ and any $\epsilon > 0$, there exists a $\delta > 0$ such that, if $\theta', \theta'' \in B(\theta, \delta) \cap \Theta$, then

$$y^R(s_1(\theta', b_1), s_2(\theta'', b_2)) \in B(\theta, \epsilon) \cap Y \quad \forall b_1 \in C_1, b_2 \in C_2$$

A fully revealing equilibrium (s_1, s_2, y^R) in which y^R is robust is called a *robust* fully revealing equilibrium.[7]

Definition 2 imposes conditions both *on* and *off* the equilibrium path. *On* the equilibrium path, the definition imposes a continuity requirement on the senders' fully revealing strategies. Specifically, when the senders' small mistakes result in a pair of messages that could have been observed jointly on the equilibrium path, we require the receiver's policy choice to be close to the policy that would have resulted in the absence of mistakes.

More importantly, robustness imposes a restriction on the receiver's response *off* the equilibrium path. Consider two incompatible reports m_1, m_2, i.e. reports such that $Inv^{s_1,s_2}(m_1, m_2) = \emptyset$. Suppose that m_1, m_2 are *close* in the following sense: given the fully revealing strategies s_1, s_2, there are two states θ', θ'' that are close to each other and that could have generated those messages, i.e. $s_1(\theta', b_1) = m_1$ for some b_1 and $s_2(\theta'', b_2) = m_2$ for some b_2. The robustness restriction in Definition 2 requires that the receiver's optimal response must also be close to the states θ' and θ''. Note that this definition does not constrain the receiver's response if either of the messages is never sent in equilibrium.

Although our robustness requirement is motivated by the possibility of small mistakes by the senders, our formal analysis assumes that, ex ante, the senders and the receiver are unaware that these mistakes might happen. Definition 2 ensures that as the size of the mistakes goes to zero, any robust equilibrium outcome in the presence of mistakes approaches the outcome when mistakes never occur.

[7] In the Appendix we show that this concept of robustness is equivalent, in the case of known biases, to the concept of diagonal continuity introduced by Ambrus and Takahashi (2008).

2.2 Preliminary Results

In a fully revealing equilibrium, the receiver perfectly learns the state from the pair of messages, and neither sender has an incentive to try to mislead the receiver by sending a different message. Lemma 1 below shows that we can restrict our attention, without loss of generality, to fully revealing equilibria in which (i) each sender S_i truthfully reports (θ, b_i) and (ii) the receiver's action depends only on the states reported by the senders and not on the reported biases. The strategy s_i is *truthful* if for all $(\theta, b_i) \in \Theta \times C_i$, $s_i(\theta, b_i) = (\theta, b_i)$. An equilibrium with truthful strategies is called a *truthful equilibrium*.

Lemma 1. *Given $Y \subseteq \mathbb{R}^p$ and $C_1, C_2 \subseteq \mathbb{R}^p$, if there exists a (robust) fully revealing equilibrium, there exists a (robust) truthful equilibrium in which the receiver's strategy depends only on the states reported.*

Proof: All proofs are in the Appendix.

The logic behind Lemma 1 is as follows. First, an argument similar to the revelation principle implies that for any FRE, there exists a truthful equilibrium that is outcome-equivalent to it. Second, given any truthful equilibrium, consider a new strategy for the receiver constructed by taking the expectation, over b_1 and b_2, of the receiver's strategy in the given truthful equilibrium. This new strategy depends only on the states reported. The assumption that the supports of the bias vectors are cones (and hence that their magnitudes may be arbitrarily large), coupled with the mutual independence of (θ, b_1, b_2), ensures that this new strategy continues to deter both senders from deviating from truthful reporting.

Throughout the paper, $b \cdot x$ denotes the inner product between b and x. Recall that we are assuming $Y = \Theta$, so from now on we will suppress references to Θ. We denote by θ' and θ'' the reports made by S_1 and S_2, respectively.

Proposition 1 articulates the conditions that a strategy and a robust strategy for the receiver must satisfy to deter each sender from misreporting. It allows us to abstract from specifying particular belief functions when proving the existence or nonexistence of robust fully revealing equilibria.

Proposition 1. *Given $Y \subseteq \mathbb{R}^p$, and $C_1, C_2 \subseteq \mathbb{R}^p$,*

(i) There exists a fully revealing equilibrium if and only if for any $\theta', \theta'' \in Y$, there exists $y \in Y$ such that:

$$\begin{aligned} b_1 \cdot y \le b_1 \cdot \theta'' \quad \forall b_1 \in C_1 \\ b_2 \cdot y \le b_2 \cdot \theta' \quad \forall b_2 \in C_2 \end{aligned} \tag{1}$$

(ii) There exists a robust fully revealing equilibrium if and only if condition (i) is satisfied and for any $\theta \in Y$ and any $\epsilon > 0$, there exists $\delta > 0$ such that for any $\theta', \theta'' \in B(\theta, \delta) \cap Y$, there exists $y \in B(\theta, \epsilon) \cap Y$ such that,

$$\begin{aligned} b_1 \cdot y \le b_1 \cdot \theta'' \quad \forall b_1 \in C_1 \\ b_2 \cdot y \le b_2 \cdot \theta' \quad \forall b_2 \in C_2 \end{aligned} \tag{2}$$

Given non-matching reports θ' and θ'', the receiver does not know whether the true state was θ'' and S_1 deviated to a report of θ', or the true state was θ' and S_2 deviated to a report of θ''. If the receiver's response, y, satisfies the first set of inequalities in (1), then $(y-(\theta''+b_1))^2 > (\theta''-(\theta''+b_1))^2$ for any $b_1 \in C_1$, so y would punish S_1 for deviating from truthfully reporting θ'' given any bias $b_1 \in C_1$. Moreover, if $b_1 \cdot y > b_1 \cdot \theta''$ for some $b_1 \in C_1$, then we can find a sufficiently large scalar $t > 0$ such that $(y-(\theta''+tb_1))^2 < (\theta''-(\theta''+tb_1))^2$, so sender S_1 with bias tb_1 would benefit from deviating from truthtelling, and since C_1 is a cone, $tb_1 \in C_1$. Hence the first set of inequalities in (1) is both necessary and sufficient to deter deviations by S_1. Analogously, the second set of inequalities in (1) is necessary and sufficient to deter S_2 from deviating from truthfully reporting θ', given that the support of his bias is the cone C_2. Condition (2) ensures that a feasible punishment for reports $\theta' \neq \theta''$ can be found arbitrarily close to these reports when these are sufficiently close to each other.

Denote by

$$PR_C(\theta) = \{x \in \mathbb{R}^p \mid b \cdot x \leq b \cdot \theta \quad \forall b \in C\}$$

the *Punishment Region* for a sender who observes state θ and whose bias has support C. Any policy in this (individual) punishment region, and only such policies, would deter that sender from deviating from truthfully reporting θ, for any possible realization of the magnitude and direction of the bias in C. Because the support C of the bias is a cone, the (individual) punishment region is also a cone.

Given a pair of non-matching reports θ', θ'' and supports C_1, C_2, for S_1 and S_2, respectively, the *Feasible Punishment Region* for that deviation is $Y \cap PR_{C_1}(\theta'') \cap PR_{C_2}(\theta')$, the set of policies $y \in Y$ that would constitute a punishment for *both senders simultaneously*. Figure 1 illustrates these regions given cones of the biases C_1 and C_2.

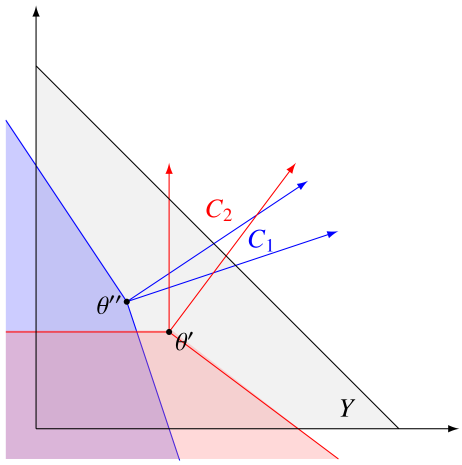

Figure 1: The blue and red shaded areas are $PR_{C_1}(\theta'')$ and $PR_{C_2}(\theta')$, respectively. The joint intersection of these areas with the policy space Y is the feasible punishment region for the incompatible reports θ', θ'', given supports C_1, C_2.

In what follows it will be convenient to focus on *anonymous* FRE. An anonymous FRE is an FRE in which the receiver's strategy satisfies $y^R(\theta', \theta'') = y^R(\theta'', \theta')$ for all θ', θ''. That is, y^R depends only on the reports and not on which sender sent which report.

Incorporating anonymity into Proposition 1 requires replacing the inequalities in Conditions (1) and (2) by

$$b \cdot y \leq \min\{b \cdot \theta', b \cdot \theta''\} \quad \forall b \in co(C_1 \cup C_2), \tag{3}$$

where $co(S)$ denotes the convex hull of the set S. The feasible *anonymous* punishment region for the incompatible reports θ', θ'' then corresponds to $Y \cap PR_C(\theta') \cap PR_C(\theta'')$, where $C = co(C_1 \cup C_2)$.

Proposition 2 shows that in many scenarios, whenever there exists a FRE, there exists an anonymous FRE, and hence restricting attention to anonymous FRE in those cases is without loss of generality. In what follows we define the dimension of a cone C, $dim(C)$, as the dimension of the minimum subspace that contains it.

Proposition 2. *Consider $Y \subseteq \mathbb{R}^p$ and suppose $C_1, C_2 \subseteq \mathbb{R}^p$ satisfy one of the following conditions:*

1. *Deterministic directions of the biases: $C_1 = \{tb_1 \mid t \geq 0\}$, $C_2 = \{tb_2 \mid t \geq 0\}$*
2. *Identical support: $C_1 = C_2$*
3. *Full dimensionality: $dim(C_1) = dim(C_2) = p$*

If there exists a fully revealing equilibrium, there exists an anonymous fully revealing equilibrium.

The conditions listed in Proposition 2 do not exhaust all the cases for which an anonymous FRE exists whenever a FRE does. In fact, the triplets (Y, C_1, C_2) for which a FRE exists but an anonymous FRE does not are non-generic. Example 1 in Appendix A.1 illustrates such non-generic situations.

In light of these observations, we will focus henceforth on anonymous FRE.

3 Robustness of Fully Revealing Equilibrium

This section introduces our first main result. It states that when $co(C_1 \cup C_2)$ has finitely many extreme biases, if there exists a FRE, then there is a FRE that is robust. In other words, requiring robustness does not restrict the conditions for existence of a FRE.

Given a set of vectors $\{b^1, ..., b^m\} \subset \mathbb{R}^p$, we denote by $C(b^1, ..., b^m)$ the convex cone spanned by those vectors: $C(b^1, ..., b^m) = \{t^1 b^1 + ... + t^n b^m \mid t^1, ..., t^n \geq 0\}$. The vectors b^i such that $\{b^1, ..., b^m\}$ is a minimal set spanning $C(b^1, ..., b^m)$ are called the extreme rays.[8] A cone that has a finite number of extreme rays is called a *polyhedral cone*.

Assumption 1. $co(C_1 \cup C_2)$ is a polyhedral cone.

Proposition 3. *Consider $Y \subseteq \mathbb{R}^p$ and $C_1, C_2 \subseteq \mathbb{R}^p$ satisfying Assumption 1. Whenever there exists an anonymous fully revealing equilibrium, there exists a robust anonymous fully revealing equilibrium.*

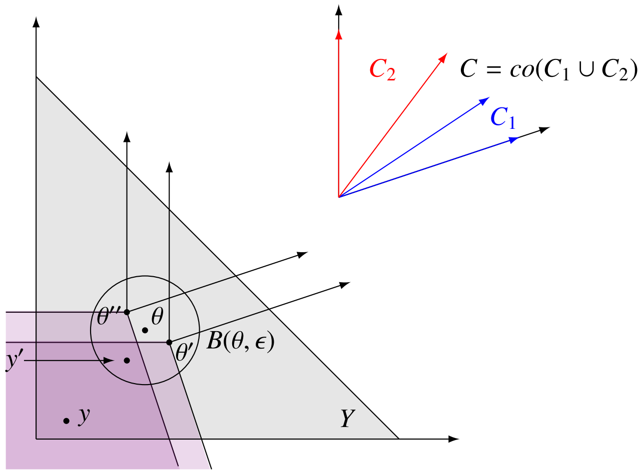

Figure 2: The dark shaded area is the feasible anonymous punishment region: $Y \cap PR_C(\theta') \cap PR_C(\theta'')$. Whenever an anonymous punishment is feasible, a robust anonymous punishment is feasible.

The intuition behind Proposition 3 for a two-dimensional policy space is illustrated in Figure 2, in which $C = co(C_1 \cup C_2)$. If there exists an anonymous FRE, then for the incompatible reports θ', θ'', the feasible anonymous punishment region $Y \cap PR_C(\theta') \cap PR_C(\theta'')$ must be non-empty. Denote by y a feasible anonymous punishment. Given that Y is convex, and θ', θ'', and y are all feasible policies, any point in the triangle of convex combinations of these three policies is also feasible. As θ' and θ'' converge to θ, $PR_C(\theta') \cap PR_C(\theta'')$ contains policies, such as y', that not only lie in that triangle but also get closer to θ, eventually belonging to the ball $B(\theta, \epsilon)$. Such policies are therefore feasible *local* anonymous punishments.

The intuition for policy spaces of arbitrary dimension is similar, as long as the individual punishment regions $PR_C(\theta')$ and $PR_C(\theta'')$ are polyhedral cones. If so, it remains true that as the two incompatible reports θ' and θ'' converge to θ, there are points, lying in the intersection of $PR_C(\theta') \cap PR_C(\theta'')$ with the feasible triangle formed by θ', θ'', and y, that approach θ. This is not necessarily the case, though, if the cones of the biases have infinitely many extreme rays. In Appendix A.1, we provide a three-dimensional example, Example 2, in which Assumption 1 is violated and, while there exists an anonymous FRE, there is no robust anonymous FRE. The example assumes a very special relationship between Y and C, and we discuss how by enlarging Y marginally, we can ensure the existence of a robust anonymous FRE, even though C does not satisfy Assumption 1. In general, any convex cone can be approximated by a sequence of polyhedral cones, and hence along the sequence, Proposition 3 would hold.

Proposition 3 relies on the assumption that the supports of the biases are cones, which implies that the magnitudes of the biases may be arbitrarily large. If the magnitudes of the biases were known to be sufficiently small, any policy sufficiently far from the reports (in any direction) could serve as a punishment. Such a response would obviously not be robust. When the supports are cones, the cones pin down the directions, relative to the reports, in which any effective punishment must lie. The convexity of the policy space then guarantees

[8] If $C(b^1, ..., b^m)$ is strictly included in a half-space, such a minimal set will be unique.

that if there is any feasible punishment in the right directions, there must be a small feasible punishment.[9]

4 Geometric Characterization of FRE

In this section we provide geometric conditions, on the shape of the policy space and the support of the biases, that characterize when an anonymous FRE exists.

Given a policy space Y, we will denote by $Fr(Y)$ the frontier of Y. For any $\theta \in Fr(Y)$, we will define

$$P_Y(\theta) \equiv \{n \in \mathbb{R}^p \mid |n| = 1,\ n \cdot y \geq n \cdot \theta \quad \forall y \in Y\}.$$

$P_Y(\theta)$ is the set of unit normal vectors of the supporting hyperplanes of Y at θ, pointing in the direction of Y. We define the dimension of $P_Y(\theta)$, $dim(P_Y(\theta))$, as the dimension of the minimal subspace that contains it. We say that θ is a *smooth* point of $Fr(Y)$ if $dim(P_Y(\theta)) = 1$. In such a case, there is only one supporting hyperplane of Y at θ, and we denote by $n_Y(\theta)$ the *inward* normal vector to $Fr(Y)$ at θ, so that $P_Y(\theta) = \{n_Y(\theta)\}$. We say that θ is a *partially smooth* point of $Fr(Y)$ if $1 < dim(P_Y(\theta)) < p$, and that θ is a *kink* point of $Fr(Y)$ if $dim(P_Y(\theta)) = p$. Finally, we will denote by $int(S)$ the interior of set S.

We first analyze the case of deterministic directions of the biases. Our result here will then be used as a building block for the general characterization for uncertain biases.

4.1 Deterministic Directions of the Biases

Consider the case in which the receiver knows the directions of the biases, that is $C_1 = C(b_1) = \{tb_1 \mid t \geq 0\}$ and $C_2 = C(b_2) = \{tb_2 \mid t \geq 0\}$ for some linearly independent $b_1, b_2 \in \mathbb{R}^p$.[10] For this case, Proposition 2 shows that requiring the receiver's strategy to be anonymous does not restrict the conditions for existence of a FRE. Moreover, $co(C_1 \cup C_2) = C(b_1, b_2)$, so Assumption 1 is satisfied.

When the orientations of the bias vectors are known by the receiver, the directions of conflict between the senders and the receiver are limited to those on the plane spanned by these bias vectors. This implies that the senders will have no incentives to deviate by misreporting dimensions of the state orthogonal to this plane. Proposition 4 provides a

[9] In two dimensions, even if the magnitudes of the biases were known to be small, the directions in which an effective *robust* punishment would have to lie would still be determined entirely by the possible orientations of the biases. In our discussion paper, Meyer et al. (2016), we showed that in two-dimensional spaces, a robust FRE exists for known magnitudes of the biases if and only if there exists a FRE for arbitrarily large magnitudes.

[10] If $b_1 = tb_2$, then if $t > 0$ there always exists a robust FRE, independently of the shape of Y. The receiver's policy

$$y^R(\theta', \theta'') = \begin{cases} \theta' & if \quad b_1 \cdot \theta' < b_1 \cdot \theta'' \\ \theta'' & if \quad b_1 \cdot \theta'' < b_1 \cdot \theta' \\ \frac{\theta' + \theta''}{2} & if \quad b_1 \cdot \theta' = b_1 \cdot \theta'' \end{cases}$$

implements an anonymous robust FRE.

If $t < 0$, then there is no FRE unless Y is included in a lower dimensional hyperplane that is orthogonal to b_1, so all policies have the same inner product with b_1, b_2. In such a case there exists a robust FRE as well: a robust FRE is supported by $y^R(\theta', \theta'') = \lambda\theta' + (1 - \lambda)\theta''$, for $\lambda \in [0, 1]$.

geometric characterization for existence of FRE that indeed depends only on the projection of the policy space onto the plane spanned by b_1 and b_2.

Given b_1 and b_2, we denote by Π_{b_1,b_2} the plane spanned by these vectors. For any $\theta \in \mathbb{R}^p$, we denote by θ_{b_1,b_2} the orthogonal projection of θ onto Π_{b_1,b_2} and by Y_{b_1,b_2} the orthogonal projection of Y onto Π_{b_1,b_2}.

Proposition 4. *Consider $Y \subseteq \mathbb{R}^p$ and $C_1 = C(b_1)$, $C_2 = C(b_2)$ with $b_1, b_2 \in \mathbb{R}^p$ linearly independent. If Y_{b_1,b_2} is closed, then there exists a fully revealing equilibrium if and only if for every smooth point $\theta_{b_1,b_2} \in Fr(Y_{b_1,b_2})$,*

$$n_{Y_{b_1,b_2}}(\theta_{b_1,b_2}) \notin int(C(b_1, b_2)) \qquad \textit{(Local Deterrence Condition)}$$

To understand Proposition 4, think first about the two-dimensional case, $p = 2$, for which b_1 and b_2 span the whole space and hence we can abstract from the projections. Here, the Local Deterrence Condition (LDC) can be rewritten as $n_Y(\theta) \notin int(C(b_1, b_2))$ for any smooth point θ on $Fr(Y)$. This condition is satisfied if and only if in state θ, there exists, close to θ, a feasible policy that, no matter the magnitudes of the biases, is worse for both senders than the policy $y = \theta$. In other words, this condition ensures that any small deviation from the smooth point θ by either sender is deterrable with a local (i.e. robust) punishment.

For higher dimensional spaces, $p > 2$, the senders' interests are aligned with those of the receiver in any direction orthogonal to the plane spanned by b_1 and b_2. For a deviation to be profitable for a sender it would have to induce a policy y whose projection y_{b_1,b_2} had a higher inner product with that sender's bias vector than the projection of the true state θ_{b_1,b_2}. Hence, *if* the receiver can ensure that her response to a deviation constitutes a punishment in the projection onto Π_{b_1,b_2}, she can deter all deviations. Furthermore, since the unknown magnitudes of the senders' biases can be arbitrarily large, it is not possible for the receiver to be sure of punishing a deviation *unless* her response constitutes a punishment in the projection. Thus, responding to a deviation by choosing a policy whose projection onto Π_{b_1,b_2} is worse for both senders is both necessary and sufficient for deterrence of misreporting. This is the logic behind the general form of the Local Deterrence Condition stated in Proposition 4.

The "if" part of Proposition 4 demonstrates that the feasibility of deterring small deviations (with local punishments) guarantees the existence of punishments for large deviations as well.[11][12] This proposition shares with Proposition 3 a focus on local punishments for local deviations, but the two messages and their emphases are distinct. Proposition 3 shows that the feasibility of punishing small deviations with *small* punishments is *necessary* for full revelation, while Proposition 4 shows that the feasibility of punishing *small* deviations is *sufficient*.

[11]This result relies on the convexity of Y, which is assumed throughout the paper. In our working paper, Meyer et al. (2016), we showed that for non-convex policy spaces in $\mathbb{R}^2$, the local deterrence condition needs to be supplemented with a *global deterrence condition* in which $Fr(Y)$ is replaced by $Fr(co(Y))$.

[12]For $p > 2$, this result uses the assumption that Y_{b_1,b_2} is closed, as is explained in more detail in the Appendix.

Figure 3 illustrates the applicability of the LDC in a pair of two-dimensional examples. In panel (a), Y is a half-space: $Y = \{y \in \mathbb{R}^2 \mid n \cdot y \geq k\}$. All points on the frontier of Y are smooth, with the same inward normal vector n. It is easy to see that for $C_1 = C(b_1)$ and $C_2 = C(b_2)$, the LDC is satisfied, so all local deviations can be punished locally. For example, the incompatible reports (θ', θ'') close to θ can be punished by y, which is also close to θ. Given that the LDC is satisfied, it then follows from Proposition 4 that there exists a FRE.

In panel (b) of Figure 3, the bias directions are the same as in panel (a), but the shape of Y is different. This panel illustrates the necessity of the LDC for existence of a (robust) FRE. The LDC is violated at θ, which implies that nearby incompatible reports along the frontier, such as (θ', θ''), cannot be punished; this can be confirmed by observing that the punishment region for the deviation (θ', θ'') does not intersect Y. One might argue that violations of the senders' incentives for truthtelling along the frontier are a minor problem if, for instance, the probability of those states arising is close to zero. However, the fact that local deviations along the frontier are not deterrable implies that bigger deviations, such as $(\tilde{\theta}', \tilde{\theta}'')$, are not deterrable either.

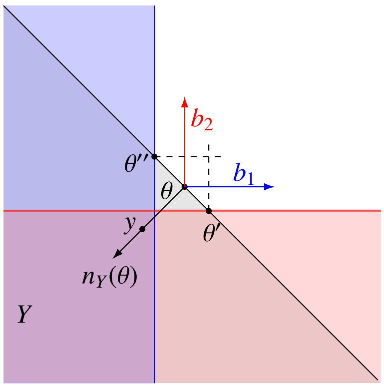

(a) Y is a halfspace. The LDC is satisfied and there exists a (robust) FRE.

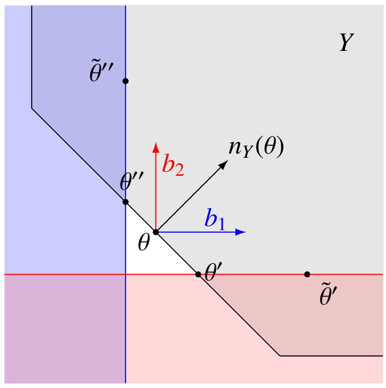

(b) The LDC is violated and there is no FRE.

Figure 3: The LDC for $Y \subset \mathbb{R}^2$.

When Y is compact, the Local Deterrence Condition in Proposition 4 can be shown to be equivalent to the condition that there exists a common "worst point" for the two senders in the policy space.[13] This latter condition, which can equivalently be stated as the existence of a $\theta \in Fr(Y)$ such that $b_1, b_2 \in P_Y(\theta)$, was shown by Ambrus and Takahashi (2008) to characterize existence of FRE for *compact* policy spaces and arbitrarily large, but known, magnitudes of the biases. However, when Y is not bounded, the existence of a common

[13]By a sender's "worst point" in Y, we refer to a point with the lowest inner product with that sender's bias vector. For Y compact and for sufficiently large magnitude of the bias, such a point will indeed minimize the sender's utility over Y, whatever the state. For small biases, it need not do so. However, independent of the magnitude of the bias, such a point will constitute a punishment for that sender for *any* deviation.

worst point is not necessary for the existence of a FRE, as Figure 3(a) illustrates. The LDC in Proposition 4 thus provides a more general characterization for existence of FRE.

4.2 Uncertain Directions of the Biases

We now characterize the existence of robust FRE when the receiver is uncertain about *both* the directions and the magnitudes of the senders' biases, for which the supports are the cones C_1 and C_2. To simplify the characterization we make a second technical assumption about the cones of the biases.

Assumption 2. $dim(co(C_1 \cup C_2)) = p$

Proposition 5. *Consider $Y \subseteq \mathbb{R}^p$ and C_1, C_2 satisfying Assumptions 1 and 2. Denote by $b^1, ..., b^m$ (with superscripts) the extreme rays of $co(C_1 \cup C_2)$. The following statements are equivalent:*

(i) There exists an anonymous FRE given C_1 and C_2.

(ii) For any pair of biases $b_1, b_2 \in co(C_1 \cup C_2)$, there exists a FRE given $C(b_1)$ and $C(b_2)$.

(iii) (a) For any smooth $\theta \in Fr(Y)$, $n_Y(\theta) \notin C(b^1, ..., b^m) \setminus \{b^1, ..., b^m\}$, and
(b) For any partially smooth $\theta \in Fr(Y)$, $P_Y(\theta) \cap int(C(b^1, ..., b^m)) = \emptyset$.

The fact that (*ii*) implies (*i*) in Proposition 5 is more subtle than it may at first appear. Condition (*ii*) allows for the strategy of the receiver supporting a FRE with known bias directions $b_1, b_2 \in co(C_1 \cup C_2)$ to depend on these directions when the senders' reports are incompatible. But condition (*i*) requires that, in response to a deviation, a *single* policy must constitute a punishment for all possible realizations of b_1 in C_1 and b_2 in C_2. Even though condition (*ii*) appears to be weaker than (*i*), our proof demonstrates their equivalence.

Given Assumption 1, one might be tempted to weaken Condition (*ii*) by replacing all possible combinations of $b_1, b_2 \in co(C_1 \cup C_2)$ by all possible pairs of extreme biases $b^i, b^j \in \{b^1, ..., b^m\}$. However this weaker condition is no longer sufficient for existence of an anonymous FRE given C_1 and C_2, as illustrated by Example 3 in Appendix A.1.

Condition (*ii*) in Proposition 5 could easily be translated to an equivalent geometric condition using Proposition 4. However it might be tedious to check the appropriate Local Deterrence Condition for each possible combination of $b_1, b_2 \in co(C_1 \cup C_2)$. Condition (*iii*) provides a simpler generalization of the LDC for uncertain biases.

To better understand Condition (*iii*), it is helpful to see how it simplifies for several special cases. Consider first the two-dimensional case, $p = 2$. For $Y \subseteq \mathbb{R}^2$, all points on the frontier of Y are either smooth points or kink points; there are no partially smooth points. Hence Condition (*iii*)(*b*) is irrelevant. Furthermore, for $co(C_1 \cup C_2) \subseteq \mathbb{R}^2$, there are at most two extreme biases, which we denote by $\underline{b}$ and $\overline{b}$. Given this, Condition (*iii*)(*a*) reduces to the following condition: For any smooth $\theta \in Fr(Y)$, $n_Y(\theta) \notin int(C(\underline{b}, \overline{b}))$. Thus in two dimensions, the LDC for uncertain bias directions takes the same form as the LDC for deterministic bias directions, except that the *least aligned* possible pair of bias directions, $\underline{b}$ and $\overline{b}$, replaces the known pair b_1 and b_2. Correspondingly, in the two-dimensional case, Condition (*ii*) is equivalent to the condition that there exists a FRE if the bias directions are known to be the extreme biases $\underline{b}$ and $\overline{b}$.

Consider now the case where p is arbitrary and Y is a half-space: $Y = \{y \in \mathbb{R}^p \mid n \cdot y \geq k\}$. In this case, all points on the frontier of Y are smooth, with the same inward normal vector n. Here, therefore, Condition (*iii*) simplifies to $n \notin C(b^1, ..., b^m) \setminus \{b^1, ..., b^m\}$.

Finally, suppose that $Y \subseteq \mathbb{R}^p$ is compact. In this case, it can be shown that Condition (*iii*) reduces to the following single condition: For any $\theta \in Fr(Y)$ such that $dim(P_Y(\theta)) < p$, $P_Y(\theta) \cap int(C(b^1, ..., b^m)) = \emptyset$.

Our analysis so far has assumed that θ, b_1, and b_2 are mutually independent. It is natural to ask to what extent Proposition 5 generalizes if, holding the supports C_1 and C_2 fixed, we allow the (conditional) distributions of the biases to vary with the state. As long as the senders are restricted to making reports about the state only, Condition (*iii*) in Proposition 5 remains sufficient for the existence of an anonymous FRE.[14] The reason is that Condition (*iii*) guarantees that, whatever the state and the realization of b_i in C_i, each sender S_i has incentives for truthful reporting. Moreover, if in every state the conditional distribution of b_i given the state has full support C_i, then Condition (*iii*) also remains necessary for the existence of an anonymous FRE.

5 Construction of Robust FRE

We now study how to construct robust and anonymous FRE. As in the previous section, we start with the case of deterministic directions of the biases. For this case we provide a strategy for the receiver, which we call the Min Rule, that is feasible and implements a robust FRE whenever one exists. We then generalize the construction to the case where the receiver is uncertain about the directions of the biases.

5.1 Deterministic Directions of the Biases

Suppose that the directions of the biases are deterministic, i.e. $C_1 = C(b_1) = \{tb_1 \mid t \geq 0\}$ and $C_2 = C(b_2) = \{tb_2 \mid t \geq 0\}$ for some linearly independent $b_1, b_2 \in \mathbb{R}^p$.[15] Given any pair of reports (θ', θ''), define $M_{b_1,b_2}(\theta', \theta'')$ as follows:

$$M_{b_1,b_2}(\theta', \theta'') = \left\{ x \in \mathbb{R}^p \mid \begin{array}{c} b_1 \cdot x = \min\{b_1 \cdot \theta', b_1 \cdot \theta''\} \\ b_2 \cdot x = \min\{b_2 \cdot \theta', b_2 \cdot \theta''\} \end{array} \right\}$$

$M_{b_1,b_2}(\theta', \theta'')$ is a subspace of $\mathbb{R}^p$ orthogonal to Π_{b_1,b_2}. Its projection onto the plane Π_{b_1,b_2} corresponds to the point in that plane which is the coordinate-wise minimum of the projections of the senders' reports, using the coordinate system formed by the normal

[14] Lemma 1 relied on (θ, b_1, b_2) being mutually independent to show that, given any truthful equilibrium, there is a truthful equilibrium in which the receiver's strategy depends only on the states reported. Since with state-dependent biases the mutual independence condition is violated, we here restrict the senders to reporting the state only.

[15] For deterministic linearly dependent biases, see the discussion in footnote 10.

vectors to b_1 and b_2.

By construction, for any point in $M_{b_1,b_2}(\theta',\theta'')$, its inner product with each bias vector is weakly smaller than the inner products of both θ' and θ'' with that bias. Hence regardless of the magnitudes of the biases, any policy in $M_{b_1,b_2}(\theta',\theta'')$ is weakly worse for both senders than both of the reports. Therefore, any feasible policy in $M_{b_1,b_2}(\theta',\theta'')$ can serve as a punishment for the deviation (θ',θ'').

Figure 4 illustrates the projection of $M_{b_1,b_2}(\theta',\theta'')$ onto Π_{b_1,b_2} for different scenarios. This projection is denoted by $\{M_{b_1,b_2}(\theta',\theta'')\}_{b_1,b_2}$. In panel (a), $b_i \cdot \theta'' < b_i \cdot \theta'$ for $i = 1, 2$, so $\{M_{b_1,b_2}(\theta',\theta'')\}_{b_1,b_2}$ coincides with θ''_{b_1,b_2}. In panel (b), by contrast, $b_1 \cdot \theta'' < b_1 \cdot \theta'$ but $b_2 \cdot \theta'' > b_2 \cdot \theta'$, so $\{M_{b_1,b_2}(\theta',\theta'')\}_{b_1,b_2}$ is distinct from both θ'_{b_1,b_2} and θ''_{b_1,b_2}.

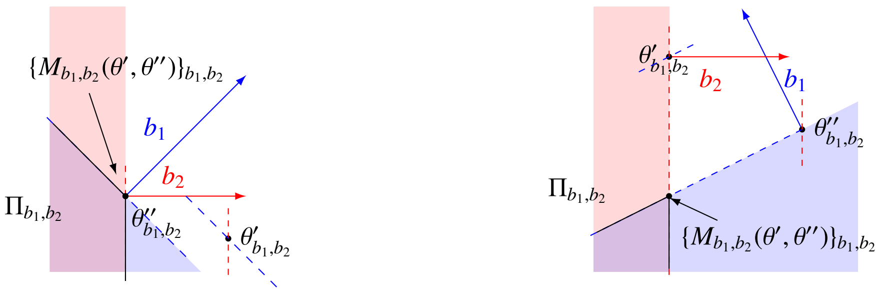

(a) Projection of $M_{b_1,b_2}(\theta',\theta'')$ onto Π_{b_1,b_2} when the senders have the same ranking over the reports θ' and θ''

(b) Projection of $M_{b_1,b_2}(\theta',\theta'')$ onto Π_{b_1,b_2} when the senders have different rankings over the reports θ' and θ''

Figure 4: Projection of $M_{b_1,b_2}(\theta',\theta'')$ onto Π_{b_1,b_2}

Proposition 6 shows that whenever there exists a FRE, the subspace $M_{b_1,b_2}(\theta',\theta'')$ does in fact intersect the policy space.

Proposition 6. *Consider $Y \subseteq \mathbb{R}^p$ and $C_1 = C(b_1), C_2 = C(b_2)$, with $b_1, b_2 \in \mathbb{R}^p$ linearly independent. There exists a fully revealing equilibrium if and only if, for any $\theta', \theta'' \in Y$,*

$$M_{b_1,b_2}(\theta',\theta'') \cap Y \neq \emptyset.$$

In the two-dimensional case, $p = 2$, $M_{b_1,b_2}(\theta',\theta'')$ consists of a single point. Here, we define the Min Rule strategy for the receiver as $y^{b_1\wedge b_2}(\theta',\theta'') \in M_{b_1,b_2}(\theta',\theta'')$. The Min Rule is an anonymous strategy. It is also robust, since as the two reports converge to each other, the policy selected by the Min Rule also converges to the reports. Propositions 3 and 6 together imply that whenever there exists a FRE, the Min Rule is feasible and supports a robust FRE.

The same conclusion holds for higher dimensional spaces, once we define an appropriate generalization of the Min Rule to select a policy in $M_{b_1,b_2}(\theta',\theta'') \cap Y$. For any $p \geq 2$ and any $\theta', \theta'' \in Y$, we define the Min Rule as implementing the policy

$$y^{b_1\wedge b_2}(\theta',\theta'') = \frac{1}{2}\left(\underset{y \in M_{b_1,b_2}(\theta',\theta'') \cap Y}{\arg\min} |y - \theta'| + \underset{y \in M_{b_1,b_2}(\theta',\theta'') \cap Y}{\arg\min} |y - \theta''| \right).$$

Since Y and $M_{b_1,b_2}(\theta', \theta'')$ are both convex, $y^{b_1 \wedge b_2}(\theta', \theta'') \in M_{b_1,b_2}(\theta', \theta'') \cap Y$. This generalized Min Rule is also easily seen to be anonymous and robust.

An appealing feature of the Min Rule is the way it selects a punishment for the senders in the dimensions in which their interests conflict with the receiver's. In these dimensions, policies in $M_{b_1,b_2}(\theta', \theta'')$ constitute the anonymous punishments which are *least severe* for each of the senders, subject to deterring both of them from misreporting, no matter how large their biases. In our working paper Meyer et al. (2016), we showed that this feature of the Min Rule makes it particularly attractive in deterring collusion by the senders. Specifically, for two dimensions, we showed that if the FRE supported by the Min Rule is not collusion-proof, then no other FRE can be collusion-proof when the unknown magnitudes of the biases can be arbitrarily large.[16]

Our multidimensional Min Rule generalizes an observation made by Krishna and Morgan (2001a) for a one-dimensional policy space. They observed that if the senders' ideal points are both larger than the receiver's, then a FRE can be supported by the receiver choosing the smaller of the two reports. In our multidimensional setting, whenever a FRE exists, our Min Rule supports full revelation by exploiting the existence of *some* overlap in the senders' preferences over policies, even when their bias vectors are not perfectly aligned.

The Min Rule may appear reminiscent of Battaglini (2002)'s construction of a FRE, which also uses the coordinate system formed by the normal vectors to the senders' biases. Battaglini's construction provides incentives for each sender to report truthfully, by restricting each sender's influence over the receiver's policy to dimensions orthogonal to his bias vector. However, our Min Rule is distinct from the receiver's strategy in Battaglini's construction in that ours is anonymous. This is important since, in restricted policy spaces, as Proposition 6 shows, a robust FRE exists if and only if the Min Rule is feasible for all pairs of reports. In contrast, in restricted policy spaces, the receiver's strategy in Battaglini's construction may be infeasible and yet a robust FRE could still exist. Such a situation is illustrated in Figure 3(a). There, if S_1 were to report θ'' and S_2 to report θ' (note that this is a reversal from our usual notational convention), Battaglini's construction would dictate the choice of the policy x such that $b_1 \cdot x = b_1 \cdot \theta'$ and $b_2 \cdot x = b_2 \cdot \theta''$, which lies outside the feasible set Y. The point x corresponds to the intersection of the dashed lines in Figure 3.

Another crucial difference between our analysis and Battaglini's is that our construction of robust FRE can be extended to accommodate uncertainty about the directions of the senders' biases, whereas Battaglini's construction cannot.

5.2 Uncertain Directions of the Biases

Now let the receiver be uncertain about *both* the directions and the magnitudes of the senders' biases, and consider first the two-dimensional case, $p = 2$. As noted in Section 4.2, for $p = 2$ there are at most two extreme biases, denoted $\underline{b}$ and $\overline{b}$. For this case, we can generalize the Min Rule for deterministic bias directions b_1 and b_2 by replacing

[16]We used the same concept of collusion-proofness as used by Battaglini (2002).

these known directions with $\underline{b}$ and $\overline{b}$, the least aligned possible pair of directions. Thus, the two-dimensional Min Rule with uncertain biases is the strategy $y^{\underline{b}\wedge\overline{b}}(\theta', \theta'') \in M_{\underline{b},\overline{b}}(\theta', \theta'')$. The following proposition follows from combining Proposition 6 with Proposition 5 when $p = 2$.

Proposition 7. *Suppose that $Y \subseteq \mathbb{R}^2$ and $co(C_1 \cup C_2) = C(\underline{b}, \overline{b})$, with $\underline{b}, \overline{b} \in \mathbb{R}^2$ linearly independent. There exists an anonymous fully revealing equilibrium if and only if, for any $\theta', \theta'' \in Y$,*

$$M_{\underline{b},\overline{b}}(\theta', \theta'') \cap Y \neq \emptyset.$$

The logic for this result is the same as the logic behind the generalization of the LDC to uncertain biases in two dimensions: If deviations can be punished for the least aligned possible pair of bias directions $\underline{b}, \overline{b}$, then they can be punished for any possible bias directions b_1, b_2. See Figure 5.

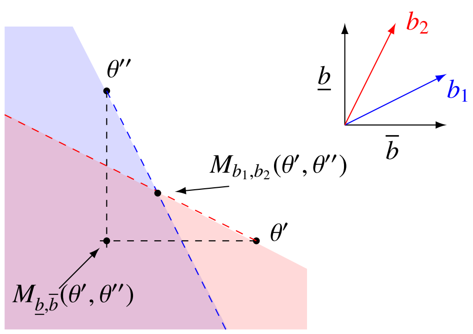

Figure 5: The uncertain biases have support $C(\underline{b}, \overline{b})$. For any realization $b_1, b_2 \in C(\underline{b}, \overline{b})$, $M_{\underline{b},\overline{b}}(\theta', \theta'')$ is an anonymous punishment.

For higher-dimensional spaces, assume that $co(C_1 \cup C_2)$ satisfies Assumption 1 and that there is no b such that $b, -b \in co(C_1 \cup C_2)$. Denote the extreme rays of $co(C_1 \cup C_2)$ by $\{b^1, ..., b^m\}$.

It would be natural to try to generalize the Min Rule for uncertain biases and $p > 2$ by defining the set $M_{b^1,...,b^m}(\theta', \theta'') \equiv \{x \mid b^i \cdot x = \min\{b^i \cdot \theta', b^i \cdot \theta''\}, i = 1, .., m\}$. However, $co(C_1 \cup C_2)$ might have $m > p$ distinct extreme biases, and in such cases it is obviously impossible to satisfy simultaneously all of the equalities defining $M_{b^1,...,b^m}(\theta', \theta'')$. In fact, even if $2 < m \leq p$, it might be that $M_{b^1,...,b^m}(\theta', \theta'') \cap Y = \emptyset$ even though there exists an anonymous FRE. Example 4 in Appendix A.1 demonstrates this possibility.

To generalize the Min Rule to uncertain biases and $p > 2$, we therefore *must* relax the equality constraints imposed in the definition of $M_{b^1,...,b^m}(\theta', \theta'')$. Define $C = co(C_1 \cup C_2)$. For $p > 2$ and any $\theta', \theta'' \in Y$, define the generalized Min Rule with uncertain biases as implementing the policy

$$y^{\wedge C}(\theta', \theta'') = \frac{1}{2}\left(\underset{y \in PR_C(\theta') \cap PR_C(\theta'') \cap Y}{\arg\min} |y - \theta'| + \underset{y \in PR_C(\theta') \cap PR_C(\theta'') \cap Y}{\arg\min} |y - \theta''| \right).$$

This strategy is anonymous and robust. Moreover, since both Y and $PR_C(\theta') \cap PR_C(\theta'')$ are convex, whenever $PR_C(\theta') \cap PR_C(\theta'') \cap Y \neq \emptyset$, the policy $y^{\wedge C}(\theta', \theta'')$ will be a feasible punishment. Hence this Min Rule supports a robust anonymous FRE whenever one exists.

A Appendix

A.1 Examples

Example 1: Anonymous vs non-anonymous FRE

Example 1 illustrates a non-generic case in which a FRE exists but an anonymous FRE does not.

Example 1. *Consider Y a half-space of $\mathbb{R}^2$, C_1 a one-dimensional cone spanned by the bias b_1 which is orthogonal to the frontier of Y, and C_2 a two-dimensional cone containing b_1 in its interior. See Figure 6. For any pair of incompatible reports (θ', θ''), the feasible (non-anonymous) punishment region, $Y \cap PR_{C_1}(\theta'') \cap PR_{C_2}(\theta')$, is not empty and hence there exists a FRE. Panel (a) depicts a pair of incompatible reports θ', θ'', where θ'' lies on the frontier of Y. The receiver can use a policy in $Y \cap PR_{C_1}(\theta'') \cap PR_{C_2}(\theta')$, for example policy y, to punish this deviation. Panel (b) depicts the same deviation but restricts the receiver to using an anonymous strategy. Since $co(C_1 \cup C_2) = C_2$, the feasible anonymous punishment region is $Y \cap PR_{C_2}(\theta') \cap PR_{C_2}(\theta'')$, which is empty. It is thus impossible to punish this deviation anonymously. However, if the bias b_1 were tilted in either direction even slightly, so that it was no longer orthogonal to the frontier of Y, then a non-anonymous FRE would no longer exist, as illustrated in panels (c) and (d).*

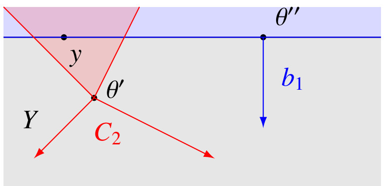

(a) Policy y lies in the feasible punishment region $Y \cap PR_{C_1}(\theta'') \cap PR_{C_2}(\theta')$, so can serve as a non-anonymous punishment for $\theta' \neq \theta''$.

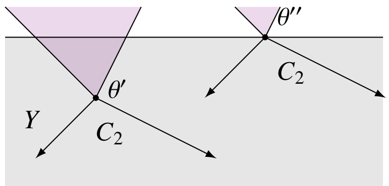

(b) The feasible anonymous punishment region $Y \cap PR_{C_2}(\theta') \cap PR_{C_2}(\theta'')$ is empty, so $\theta' \neq \theta''$ cannot be punished anonymously.

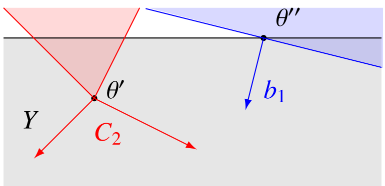

(c) For b_1 tilted slightly relative to panel (a), there is no FRE.

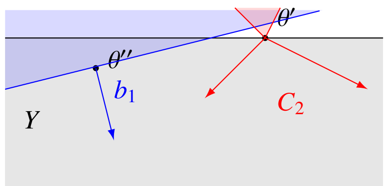

(d) For b_1 tilted slightly relative to panel (a), there is no FRE.

Figure 6: Non-anonymous vs. anonymous FRE

In order for a FRE to exist while an anonymous FRE does not, the support of the bias of one sender must be contained in a subspace orthogonal to the frontier of Y, and the support of the other sender's bias must contain in its interior one bias that is also orthogonal to the frontier of Y.

Example 2: A FRE that is not robust

Example 2 illustrates a situation in which Assumption 1 is not satisfied and, while there exists a FRE, there is no robust FRE.

Example 2. *Let $C = C_1 = C_2$ be a circular convex cone. Such a cone has an infinite number of extreme rays. For $C_1 = C_2$, any policy that punishes the deviation (θ', θ'') also punishes the deviation (θ'', θ'), so is an anonymous punishment.*

Define $Y = \{\theta \in \mathbb{R}^3 \mid b \cdot \theta \geq 0,\ \forall b \in C\}$. See Figure 7. Clearly, given any θ and for any bias $b \in C$, θ is weakly preferred to $\underline{\theta} \equiv (0,0,0)$, and hence $\underline{\theta}$ can be used as an anonymous punishment for any deviation. Therefore an anonymous FRE exists.

However, there is no robust anonymous FRE: Consider θ' on the frontier of Y. Given the definition of Y, there exists $b \in C$ such that $b \cdot \theta' = 0$. Any potential punishment y for a deviation from θ' must be on the ray connecting θ' and $\underline{\theta}$; if it were not, then $b \cdot y > 0 = b \cdot \theta'$, so the deviation would be attractive for a sender with bias tb for sufficiently large t. In other words, for any point θ' on the frontier, the individual feasible punishment region is $PR_C(\theta') \cap Y = [\theta', \underline{\theta}]$. Now consider another point θ'' on the frontier, arbitrarily close to θ', as shown in the figure. Then $PR_C(\theta') \cap PR_C(\theta'') \cap Y = [\theta', \underline{\theta}] \cap [\theta'', \underline{\theta}] = \underline{\theta}$, so the only punishment for the pair of reports θ', θ'' is $\underline{\theta}$. Consequently, condition (ii) in Proposition 1 fails to hold, so a robust anonymous FRE does not exist.

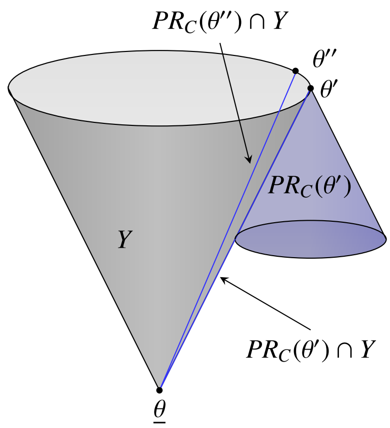

Figure 7: The cone of the biases C has infinitely many extreme rays. There is no local punishment for the deviation (θ', θ''), even though there exists a global punishment $\underline{\theta}$.

The example assumes a very special relationship between Y and C: Y is the dual cone of C, and hence $PR_C(\theta') \cap Y$ has dimension 1. In fact if we slightly enlarge Y, $PR_C(\theta') \cap Y$ will have full dimension, and the conclusion in the example will no longer hold. In particular, if we approximate Y with polyhedral cones Y_n with $Y \subset Y_n$, it can be shown that there would exist a robust anonymous FRE for every Y_n. Analogously, if we fix Y and consider any sequence of polyhedral cones $C_n \subset C$ converging to C, a robust anonymous FRE would exist for any cone in the sequence.

Example 3 for Section 4.2

Example 3 shows that the existence of a FRE for any pair of extreme biases is not sufficient to guarantee the existence of a FRE.

Example 3. *Let $C_1 = C_2 = C(b^1, b^2, b^3)$, where the extreme biases are the canonical vectors in $\mathbb{R}^3$: $b^1 = (1,0,0)$, $b^2 = (0,1,0)$, and $b^3 = (0,0,1)$. Let $Y \subset \mathbb{R}^3$ be the tetrahedron with vertices $A = (1,0,0)$, $B = (0,1,0)$, $C = (0,0,1)$, and $D = (1,1,1)$. See Figure 8. It is easy to check that no FRE exists: To punish the incompatible reports (A, B), the receiver would need to implement a policy y such that $b^i \cdot y \leq 0$ for $i = 1,2,3$, but there is no such policy in Y.*

Yet, for any given pair of extreme biases, any deviation is punishable, as can be confirmed by applying Proposition 4. For example, if the cones were known to be $C(b^1)$ and $C(b^2)$, the deviation (A, B) would be punished by point C; if they were known to be $C(b^2)$ and $C(b^3)$, the same deviation would be punishable by A; and if they were known to be $C(b^1)$ and $C(b^3)$, it would be punishable by B.

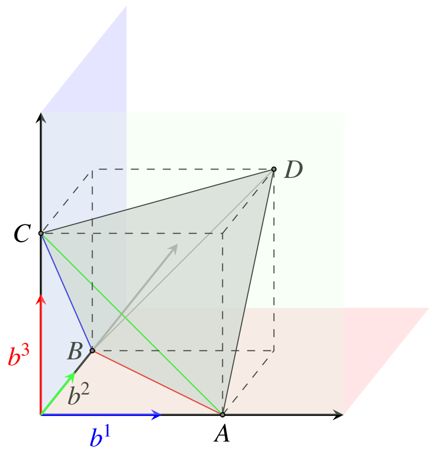

Figure 8: There is no FRE given $C(b^1, b^2, b^3)$, but for any pair of extreme biases b^i, b^j, there exists a FRE given $C(b^i)$ and $C(b^j)$.

Example 4 for Section 5.2

Example 4 shows that even when there exists an anonymous FRE and the set $M_{b^1,\ldots,b^m}(\theta', \theta'')$ is not empty, this set need not contain a feasible policy.

Example 4. *Let $C_1 = C_2 = C(b^1, b^2, b^3)$, where the extreme biases are the canonical vectors in $\mathbb{R}^3$: $b^1 = (1,0,0)$, $b^2 = (0,1,0)$, and $b^3 = (0,0,1)$. For any $\theta' = (x', y', z')$ and $\theta'' = (x'', y'', z'')$, $M_{b^1,b^2,b^3}(\theta', \theta'') = \{(\min\{x', x''\}, \min\{y', y''\}, \min\{z', z''\})\}$.*

Let $Y \subset \mathbb{R}^3$ be the tetrahedron with vertices $O = (0,0,0)$, $A = (0,1,1)$, $B = (1,0,1)$, $C = (1,1,0)$. See Figure 9. It is easy to see that an anonymous FRE exists: For all biases in $C(b^1, b^2, b^3)$, the policy O can serve as a punishment for any deviation. However, for the deviation (A, B), $M_{b^1,b^2,b^3}(A, B) = \{(0,0,1)\}$, but $(0,0,1)$ is not a feasible policy.

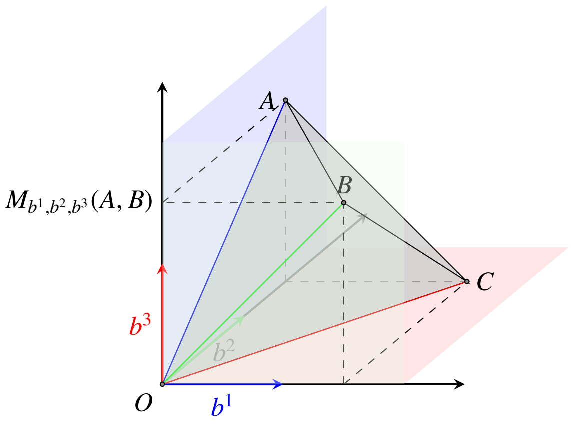

Figure 9: There exists an anonymous FRE given $C(b^1, b^2, b^3)$ but $M_{b^1,b^2,b^3}(A, B) \cap Y = \emptyset$.

A.2 Proofs

Notation used throughout the Appendix

Given a (bias) vector $b \in \mathbb{R}^p$ and a scalar $k \in \mathbb{R}$, we define $H(b, k) \equiv \{x \in \mathbb{R}^p \mid b \cdot x > k\}$ and $h(b, k) \equiv \{x \in \mathbb{R}^p \mid b \cdot x = k\}$, where $b \cdot x$ denotes the inner product between b and x. In words, $H(b, k)$ is the open upper half-space composed of all the points in $\mathbb{R}^p$ whose inner product with b is strictly greater than k, and $h(b, k)$ is the boundary of $H(b, k)$.

Proof of Lemma 1:

Given F, the joint cumulative distribution function of (θ, b_1, b_2), which are mutually independent, we denote by F_θ, F_1, F_2 the corresponding marginal distributions of θ, b_1 and b_2.

Consider a fully revealing equilibrium (s_1, s_2, y^R) supported by the belief function $\mu(\cdot)$. Now consider the strategies $(\tilde{s}_1, \tilde{s}_2, \tilde{y})$, where $\tilde{s}_1$ and $\tilde{s}_2$ are truthful strategies and $\tilde{y} : (\Theta \times C_1) \times (\Theta \times C_2) \longrightarrow Y$ is such that $\tilde{y}((\theta', b_1), (\theta'', b_2)) = y^R(s_1(\theta', b_1), s_2(\theta'', b_2))$. Define the belief function $\tilde{\mu}((\theta', b_1), (\theta'', b_2)) = \mu(s_1(\theta', b_1), s_2(\theta'', b_2))$. Then the strategy profile $(\tilde{s}_1, \tilde{s}_2, \tilde{y})$ supported by the belief $\tilde{\mu}$ is a truthful equilibrium.

We now show that whenever there exists a truthful equilibrium, there exists a truthful equilibrium in which the receiver's strategy depends only on the states reported.

Suppose that there exists a truthful equilibrium. Then for any reported states θ', θ'', for any reported biases $\hat{b}_1, \hat{b}_2$, and for any true biases b_1, b_2, we have that

$$
\begin{aligned}
& \int_{C_2} [y^R(\theta', \hat{b}_1, \theta'', b_2) - (\theta'' + b_1)]^2 dF_2(b_2) \geq b_1^2 \\
\Leftrightarrow \quad & \int_{C_2} [y^R(\theta', \hat{b}_1, \theta'', b_2) - \theta'']^2 dF_2(b_2) \geq 2b_1 \cdot (E_{b_2}[y^R(\theta', \hat{b}_1, \theta'', b_2)] - \theta'') \\
\\
\textit{and} \quad & \int_{C_1} [y^R(\theta', b_1, \theta'', \hat{b}_2) - (\theta' + b_2)]^2 dF_1(b_1) \geq b_2^2 \\
\Leftrightarrow \quad & \int_{C_1} [y^R(\theta', b_1, \theta'', \hat{b}_2) - \theta']^2 dF_1(b_1) \geq 2b_2 \cdot (E_{b_1}[y^R(\theta', b_1, \theta'', \hat{b}_2)] - \theta')
\end{aligned}
$$

The first inequality is the incentive compatibility condition for sender S_1, guaranteeing that in state θ'' and for realized bias b_1, and given that S_2 is reporting truthfully, S_1 has no incentive to deviate to reporting $(\theta', \hat{b}_1)$. The second inequality is the analogous incentive compatibility condition for S_2 in state θ' and for realized bias b_2. Since C_1 and C_2 are cones, these two inequalities must be satisfied for any magnitudes of the biases b_1, b_2. This implies that for any θ', θ'',

$$b_1 \cdot (E_{b_2}[y^R(\theta', \hat{b}_1, \theta'', b_2)] - \theta'') \leq 0 \qquad \forall b_1, \hat{b}_1 \in C_1$$
$$b_2 \cdot (E_{b_1}[y^R(\theta', b_1, \theta'', \hat{b}_2)] - \theta') \leq 0 \qquad \forall b_2, \hat{b}_2 \in C_2$$

Now construct a new strategy for the receiver which depends only on the states reported, as follows: $y^*(\theta', \theta'') = E_{b_1,b_2}[y^R(\theta', b_1, \theta'', b_2)]$. Note that $y^*(\theta', \theta'') \in Y$, since Y is convex. Given the mutual independence of (θ, b_1, b_2), $y^*(\theta', \theta'') = E_{b_1}[E_{b_2}[y^R(\theta', b_1, \theta'', b_2)]]$
$= E_{b_2}[E_{b_1}[y^R(\theta', b_1, \theta'', b_2)]]$, and using this, it follows from the preceding inequalities that

$$b_1 \cdot (y^*(\theta', \theta'') - \theta'') \leq 0 \qquad \forall b_1 \in C_1$$
$$b_2 \cdot (y^*(\theta', \theta'') - \theta') \leq 0 \qquad \forall b_2 \in C_2$$

These inequalities ensure that the strategy $y^*(\theta', \theta'')$ serves to deter both S_1 from deviating from truthfully reporting θ'', for any $b_1 \in C_1$, and S_2 from deviating from truthfully reporting θ', for any $b_2 \in C_2$.

Finally, note that if (s_1, s_2, y^R) is a robust FRE, then $y^*(\theta', \theta'') = E_{b_1,b_2}[y^R(\theta', b_1, \theta'', b_2)]$ is robust as well. Hence whenever there exists a robust FRE, there exists a robust truthful equilibrium in which the receiver's strategy depends only on the states reported.

Proof of Proposition 1:

Condition (*i*) in Proposition 1 can be rewritten as

$$Y \not\subseteq \bigcup_{\substack{b_1 \in C_1 \\ b_2 \in C_2}} H(b_1, b_1 \cdot \theta'') \cup H(b_2, b_2 \cdot \theta'),$$

and Condition (*ii*) as

$$B(\theta, \epsilon) \cap Y \not\subseteq \bigcup_{\substack{b_1 \in C_1 \\ b_2 \in C_2}} H(b_1, b_1 \cdot \theta'') \cup H(b_2, b_2 \cdot \theta').$$

Statement (i) (⇒): Suppose there exist $\theta', \theta'' \in Y$ such that
$Y \subseteq \bigcup_{b_1 \in C_1, b_2 \in C_2} H(b_1, b_1 \cdot \theta'') \cup H(b_2, b_2 \cdot \theta')$. We show that there does not exist a truthful equilibrium in which the receiver's response is independent of the biases reported, and therefore by Lemma 1, there does not exist a FRE. Consider any possible response by the receiver given the report (θ', θ'') and denote it by y. Then there exists either a $b_1 \in C_1$ such that $b_1 \cdot (y - \theta'') > 0$ or a $b_2 \in C_2$ such that $b_2 \cdot (y - \theta') > 0$. Suppose that $b_1 \cdot (y - \theta'') > 0$ and consider $t_1 > \frac{|y - \theta''|^2}{2b_1 \cdot (y - \theta'')}$. Then $y \in B(\theta'' + t_1 b_1, t_1|b_1|)$, which implies that sender S_1 with bias $t_1 b_1 \in C_1$ has an incentive to deviate to (θ', b_1) in state θ''. The symmetric argument

could be made if $b_2 \cdot (y - \theta') > 0$ with $t_2 > \frac{|y-\theta'|^2}{2b_2 \cdot (y-\theta')}$.

Statement (i) ($\Leftarrow$): Consider truthful strategies and a belief function $\mu(\cdot)$ such that $\mu(\theta, \theta)$ allocates mass one to θ and, for $\theta' \neq \theta''$, $\mu(\theta', \theta'')$ puts mass one on an element of $Y \setminus \bigcup_{b_1 \in C_1, b_2 \in C_2} H(b_1, b_1 \cdot \theta'') \cup H(b_2, b_2 \cdot \theta')$. Denote by y^R the optimal response by the receiver given those beliefs. Given a report (θ', θ''), $y^R(\theta', \theta'') \notin \bigcup_{b_1 \in C_1, b_2 \in C_2} H(b_1, b_1 \cdot \theta'') \cup H(b_2, b_2 \cdot \theta')$, so in particular $y^R(\theta', \theta'') \notin B(\theta'' + b_1, |b_1|)$ for any $b_1 \in C_1$ and $y^R(\theta', \theta'') \notin B(\theta' + b_2, |b_2|)$ for any $b_2 \in C_2$. So none of the senders has an incentive to deviate.

Statement (ii) ($\Rightarrow$): Suppose there exists a robust fully revealing equilibrium. Then for any $\theta \in Y$ and any $\epsilon > 0$, there exists $\delta > 0$ such that for every $\theta', \theta'' \in B(\theta, \delta)$, $y^R(s_1(\theta'), s_2(\theta'')) \in B(\theta, \epsilon) \cap Y \setminus (B(\theta' + b_1, |b_1|) \cup B(\theta + b_2, |b_2|))$ for any $b_1 \in C_1, b_2 \in C_2$. Hence $B(\theta, \epsilon) \cap Y \nsubseteq \bigcup_{b_1 \in C_1, b_2 \in C_2} H(b_1, b_1 \cdot \theta'') \cup H(b_2, b_2 \cdot \theta')$.

Statement (ii) ($\Leftarrow$): By Lemma 1, we can focus on truthful strategies for the senders and strategies for the receiver that are independent of the biases reported. For any $\theta' \neq \theta'' \in \Theta$, define

$$y^R(\theta', \theta'') \in \arg\min \left\{ |s - \theta'| : s \in Y \setminus \left(\bigcup_{b_1 \in C_1, b_2 \in C_2} H(b_1, b_1 \cdot \theta'') \cup H(b_2, b_2 \cdot \theta') \right) \right\}.$$

By part (i) we know that $Y \setminus \left(\bigcup_{b_1 \in C_1, b_2 \in C_2} H(b_1, b_1 \cdot \theta'') \cup H(b_2, b_2 \cdot \theta') \right)$ is not empty, so such a $y^R(\theta', \theta'')$ exists. To see that this strategy, together with truthful strategies for the senders, constitutes a robust FRE, consider any $\theta \in \Theta$ and any $\epsilon > 0$. By the hypothesis, for $\tilde{\epsilon} = \epsilon/3$ there exists $0 < \delta < \tilde{\epsilon}$ such that for all $\theta', \theta'' \in B(\theta, \delta) \cap Y$, $B(\theta, \tilde{\epsilon}) \cap Y \nsubseteq \bigcup_{b_1 \in C_1, b_2 \in C_2} H(b_1, b_1 \cdot \theta'') \cup H(b_2, b_2 \cdot \theta')$. Consider any $\hat{\theta} \in B(\theta, \tilde{\epsilon}) \cap Y \nsubseteq \bigcup_{b_1 \in C_1, b_2 \in C_2} H(b_1, b_1 \cdot \theta'') \cup H(b_2, b_2 \cdot \theta')$. Then $|y^R(\theta', \theta'') - \theta| \leq |y^R(\theta', \theta'') - \theta'| + |\theta' - \theta| \leq |\hat{\theta} - \theta'| + |\theta' - \theta| \leq |\hat{\theta} - \theta| + 2|\theta' - \theta| < 3\tilde{\epsilon} = \epsilon$, hence $y^R(\theta', \theta'') \in B(\theta, \epsilon) \setminus \left(\bigcup_{b_1 \in C_1, b_2 \in C_2} H(b_1, b_1 \cdot \theta'') \cup H(b_2, b_2 \cdot \theta') \right) \subset B(\theta, \epsilon) \setminus \left(\bigcup_{b_1 \in C_1, b_2 \in C_2} B(\theta'' + b_1, |b_1|) \cup B(\theta' + b_2, |b_2|) \right)$.

□

Proof of Proposition 2:

We prove the assertion for the three scenarios considered:

1. Deterministic directions of the biases:

Consider $\theta' \neq \theta''$. There are two possible scenarios:

(a) Aligned preferences: Either $b_1 \cdot \theta' \leq b_1 \cdot \theta''$ and $b_2 \cdot \theta' \leq b_2 \cdot \theta''$ or $b_1 \cdot \theta'' \leq b_1 \cdot \theta'$ and $b_2 \cdot \theta'' \leq b_2 \cdot \theta'$.

(b) Misaligned preferences: Either $b_1 \cdot \theta'' < b_1 \cdot \theta'$ and $b_2 \cdot \theta' < b_2 \cdot \theta''$ or $b_1 \cdot \theta' < b_1 \cdot \theta''$ and $b_2 \cdot \theta'' < b_2 \cdot \theta'$.

For any θ', θ'' satisfying (a), $y = \arg\min_{\theta \in \{\theta', \theta''\}} b_i \cdot \theta$ is feasible and anonymous.

Suppose θ', θ'' satisfy (b). Suppose $b_1 \cdot \theta'' < b_1 \cdot \theta'$ and $b_2 \cdot \theta' < b_2 \cdot \theta''$ and consider the report (θ', θ''). Since there exists an FRE, there exists $y \in Y$ such that $b_1 \cdot y \leq b_1 \cdot \theta'' < b_1 \cdot \theta'$ and $b_2 \cdot y \leq b_2 \cdot \theta' < b_2 \cdot \theta''$ so $b_i y \leq \min\{b_i \cdot \theta', b_i \cdot \theta''\}$. Analogously, if $b_1 \cdot \theta' < b_1 \cdot \theta''$ and

$b_2 \cdot \theta'' < b_2 \cdot \theta'$, consider the report (θ'', θ'). Since there exists a FRE, there exists a $y \in Y$ such that $b_1 \cdot y \leq b_1 \cdot \theta' < b_1 \cdot \theta''$ and $b_2 \cdot y \leq b_2 \cdot \theta'' < b_2 \cdot \theta'$ so $b_i y \leq \min\{b_i \cdot \theta', b_i \cdot \theta''\}$. And therefore there is always an anonymous FRE.

2. Identical support:

Consider $\theta' \neq \theta'' \in Y$. Given the report (θ', θ''), since there exists a FRE, there is $y \in Y$ such that $b_1 \cdot y \leq b_1 \cdot \theta''$ for all $b_1 \in C$, $b_2 \cdot y \leq b_2 \cdot \theta'$ for all $b_2 \in C$. Therefore, for any $b \in C$, $b \cdot y \leq \min\{b \cdot \theta', b \cdot \theta''\}$ and the same y could be a punishment for the report (θ'', θ'). So there exists an anonymous FRE.

3. Full dimensionality of C_1 and C_2:

We will denote by (C_1, C_2) a situation in which the support of S_1's bias is C_1 and the support of S_2's bias is C_2. Suppose that there exists a FRE for (C_1, C_2) but not an anonymous FRE. Since there exists a FRE, each of the cones C_1 and C_2 must be contained on an open halfspace of $\mathbb{R}^p$. Otherwise it is impossible to punish a deviation for biases b_i and $-b_i$ simultaneously. We proceed in two steps.

Step 1: If there exists a FRE but not an anonymous FRE for (C_1, C_2), then either there does not exist a FRE for (C_1, C_1) or there does not exist a FRE for (C_2, C_2).

Proof of Step 1: Suppose for contradiction that there exists a FRE for (C_i, C_i) $i = 1, 2$, then for any (θ', θ'') there exists $y_i \in Y$ such that $b_i \cdot y_i \leq \min\{b_i \cdot \theta', b_i \cdot \theta''\}$ for any $b_i \in C_i$. Consider the incompatible reports (y_2, y_1). Since there exists a FRE for (C_1, C_2), there exists $\hat{y} \in Y$ such that $b_1 \cdot \hat{y} \leq b_1 \cdot y_1 \leq \min\{b_1 \cdot \theta', b_1 \cdot \theta''\}$ for all $b_1 \in C_1$ and $b_2 \cdot \hat{y} \leq b_2 \cdot y_2 \leq \min\{b_2 \cdot \theta', b_2 \cdot \theta''\}$ for all $b_2 \in C_2$, and hence $\hat{y}$ is a feasible anonymous punishment for (θ', θ''). □

We can, therefore, without loss of generality assume that there is not a FRE for (C_2, C_2).

Step 2: Denote by $Fr(Y)$ the frontier of Y. If there is not a FRE for (C_2, C_2), there exists $\tilde{\theta}' \neq \tilde{\theta}'' \in Fr(Y)$, such that

$$\{\tilde{\theta}'\} = Y \setminus \cap_{b_2 \in C_2} H(b_2, b_2 \cdot \tilde{\theta}') \quad and \quad \{\tilde{\theta}''\} = Y \setminus \cap_{b_2 \in C_2} H(b_2, b_2 \cdot \tilde{\theta}'')$$

Proof of Step 2: If there is not a FRE for (C_2, C_2), there exists $\theta' \neq \theta'' \in Y$ such that there is no $y \in Y$ with $b_2 \cdot y \leq \min\{b_2 \cdot \theta', b_2 \cdot \theta''\}$ for all $b_2 \in C_2$. But given that C_2 is strictly contained in a half space of $\mathbb{R}^p$, there is always a $x \in \mathbb{R}^p$ such that $b_2 x < \min\{b_2 \cdot \theta', b_2 \cdot \theta''\}$ for all $b_2 \in C_2$. Since $x \notin Y$ we can define $\tilde{\theta}', \tilde{\theta}''$ as the unique points such that

$$\{\tilde{\theta}'\} = \arg\min\{|\theta - x| \text{ s.t. } \theta = t\theta' + (1-t)x,\ t \in [0,1], \theta \in Y\}$$
$$\{\tilde{\theta}''\} = \arg\min\{|\theta - x| \text{ s.t. } \theta = t\theta'' + (1-t)x,\ t \in [0,1], \theta \in Y\}$$

Then $\tilde{\theta}' \neq \tilde{\theta}'' \in Fr(Y)$ and there exist $b_2', b_2'' \in C_2$ such that $\tilde{\theta}' \in \arg\min_{y \in Y} b_2' y$, $\tilde{\theta}'' \in \arg\min_{y \in Y} b_2'' y$. Lastly, since $dim(C_2) = p$, this implies that

$$\{\tilde{\theta}'\} = Y \setminus \cap_{b_2 \in C_2} H(b_2, b_2 \cdot \tilde{\theta}') \quad and \quad \{\tilde{\theta}''\} = Y \setminus \cap_{b_2 \in C_2} H(b_2, b_2 \cdot \tilde{\theta}''). \qquad \square$$

Finally, consider $\tilde{\theta}' \neq \tilde{\theta}''$ as in Step 2. Since there exists an FRE for (C_1, C_2),

(a) for the report $(\tilde{\theta}', \tilde{\theta}'')$, $y^R(\tilde{\theta}', \tilde{\theta}'') = \tilde{\theta}'$.

(b) for the report $(\tilde{\theta}'', \tilde{\theta}')$, $y^R(\tilde{\theta}'', \tilde{\theta}') = \tilde{\theta}''$.

In particular (a) implies $b_1 \cdot \tilde{\theta}' \le b_1 \cdot \tilde{\theta}''$ for all $b_1 \in C_1$, and (b) implies $b_1 \cdot \tilde{\theta}'' \le b_1 \cdot \tilde{\theta}'$ for all $b_1 \in C_1$. Hence $b_1 \cdot \tilde{\theta}' = b_1 \cdot \tilde{\theta}''$ for all $b_1 \in C_1$. But $\tilde{\theta}' \neq \tilde{\theta}''$ implies $dim(C_1) < p$ which is a contradiction.

Proof of Proposition 3:

Define $C \equiv Co(C_1 \cup C_2)$ and suppose that C satisfies Assumption 1, so $C = C(b^1, ..., b^m)$. Consider $\theta \in \Theta$ and $\epsilon > 0$. Recall that the (individual) punishment region $PR_C(\theta) = \{x \in \mathbb{R}^p \mid b \cdot x \le b \cdot \theta \text{ for all } b \in C\}$ is the cone consisting of all points that are weakly worse than θ for all realizations of the bias $b \in C$. Note that $\theta \in PR_C(\theta) \cap Y$ and hence $PR_C(\theta) \cap Y \neq \emptyset$. Define

$$N(\theta) = \min_{\theta' \in PR_C(\theta) \cap Y} \#\{b^k \in \{b^1, ..., b^m\} \mid b^k \cdot \theta' = b^k \cdot \theta\}.$$

Note that $N(\theta) \in \{0, ..., p\}$.

Suppose that $N(\theta) = p$. Then $PR_C(\theta) \cap Y \equiv \{\theta\}$. Note that since there exists a FRE, for any $\theta' \in \Theta$, $Y \cap PR_C(\theta) \cap PR_C(\theta') \neq \emptyset$. Therefore, since $PR_C(\theta) \cap Y \equiv \{\theta\}$, it means that for any $\theta' \in \Theta$, $\theta \in PR_C(\theta')$, or in other words, $b \cdot \theta \le b \cdot \theta'$ for all $\theta' \in Y$, and θ is itself a local punishment.

Suppose now that $N(\theta) < p$ and consider $\bar{\theta} \in PR_C(\theta) \cap Y$ such that $\#\{b^k \in \{b^1, ..., b^m\} \mid b^k \cdot \bar{\theta} = b^k \cdot \theta\} = N(\theta)$. By convexity of Y, we can pick up $\bar{\theta} \in B(\theta, \epsilon)$. Denote by $B(\bar{\theta}) = \{b^k \in \{b^1, ..., b^m\} \mid b^k \cdot \bar{\theta} = b^k \cdot \theta\}$. Define

$$\delta = \min_{b \in \{b^1, ..., b^m\} \setminus B(\bar{\theta})} \frac{b}{|b|} \cdot (\theta - \bar{\theta}) > 0$$

We now show that for all $\theta' \in B(\theta, \delta) \cap Y$, $b \cdot \bar{\theta} \le b \cdot \theta'$ for all $b \in C$, and hence $\bar{\theta}$ is a local punishment:

- By construction of δ, if $b \in \{b^1, ..., b^m\} \setminus B(\bar{\theta})$ and $\theta' \in B(\theta, \delta)$,

$$b \cdot \theta' > b \cdot \theta - \delta b \cdot \frac{b}{|b|} \ge b \cdot \theta - \left(\frac{b}{|b|} \cdot (\theta - \bar{\theta})\right) b \cdot \frac{b}{|b|} = b \cdot \bar{\theta}.$$

- Consider $b^k \in B(\bar{\theta})$. If there exists a $\theta' \in B(\theta, \delta) \cap Y$ such that $b^k \cdot \theta' < b^k \cdot \bar{\theta}$, then by the existence of a FRE, there exists $\theta^* \in Y$ such that $b \cdot \theta^* \le b \cdot \theta'$ and $b \cdot \theta^* \le b \cdot \bar{\theta} \le b \cdot \theta$ for all $b \in C$. In particular $\theta^* \in PR_C(\theta) \cap Y$ and for all $b \in \{b^1, ..., b^m\} \setminus B(\bar{\theta})$, $b \cdot \theta^* \le b \cdot \bar{\theta} < b \cdot \theta$, and $b^k \cdot \theta^* \le b^k \cdot \theta' < b^k \cdot \bar{\theta} = b^k \cdot \theta$. Therefore $\#\{b^k \in \{b^1, ..., b^m\} \mid b^k \cdot \theta^* = b \cdot \theta\} < N(\theta)$ which contradicts the definition of $N(\theta)$. Therefore $b^k \cdot \bar{\theta} \le b^k \cdot \theta'$ for any $\theta' \in B(\theta, \delta)$.

- Finally, any $b \in C$ can be written as a linear combination of $b^1, ..., b^m$ and hence, for any $b \in C$ and any $\theta' \in B(\theta, \delta)$, $b \cdot \bar{\theta} \le b \cdot \theta'$.

Therefore $\bar{\theta}$ is a local punishment for θ.

Proof of Proposition 4:

We will first look at the case of two-dimensional spaces:

Proposition 8. *Given $Y \subseteq \mathbb{R}^2$ and $C_1 = C(b_1)$, $C_2 = C(b_2)$ with $b_1, b_2 \in \mathbb{R}^2$ linearly independent, there exists a fully revealing equilibrium if and only if for every smooth point $\theta \in Fr(Y)$,*

$$n_Y(\theta) \notin int(C(b_1, b_2))$$

Proof of Proposition 8:

$\Rightarrow$) Suppose there exists $\theta \in Fr(Y)$ smooth such that $n_Y(\theta) \in C(b_1, b_2)$. Then $Y \subset \bar{H}(b_1, b_1\theta) \cup \bar{H}(b_2, b_2\theta)$. Moreover, since $n_Y(\theta) \neq b_1$ and $n_Y(\theta) \neq b_2$, for any $\delta > 0$ there exists $\theta' \in Y \cap B(\theta, \delta)$ and $\theta'' \in Y \cap B(\theta, \delta)$ such that $b_2 \cdot \theta' < b_2 \cdot \theta$ and $b_1 \cdot \theta'' < b_1 \cdot \theta$. But then $Y \subset \bar{H}(b_1, b_1\theta) \cup \bar{H}(b_2, b_2\theta) \subset H(b_1, b_1 \cdot \theta'') \cup H(b_2, b_2 \cdot \theta')$ and local deviations from θ cannot be deterred, which contradicts the existence of a FRE.

$\Leftarrow$) Suppose that there exist $\theta', \theta'' \in Y$ such that $PR_{C(b_1)}(\theta'') \cap PR_{C(b_2)}(\theta') \cap Y = \emptyset$ and consider x, such that $b_1 \cdot x = b_1 \cdot \theta''$ and $b_2 \cdot x = b_2 \cdot \theta'$. The point $x \in PR_{C(b_1)}(\theta'') \cap PR_{C(b_2)}(\theta')$ and hence $x \notin Y$. Consider any $\theta \in Fr(Y)$ smooth that lies in the triangle formed by θ', θ'' and x.[17] See Figure 10. In particular, since Y is convex, $h(n_Y(\theta), n_Y(\theta) \cdot \theta)$ is a hyperplane separating Y from x, and

$$n_Y(\theta)(\theta' - \theta) \geq 0 \tag{4}$$

$$n_Y(\theta)(\theta'' - \theta) \geq 0 \tag{5}$$

$$n_Y(\theta)(x - \theta) < 0. \tag{6}$$

Note that the inequality in (6) is strict. This is because if, for all smooth points in the triangle formed by $\{\theta', \theta'', x\}$, $n_Y(\theta) \cdot (x - \theta) = 0$, then x would be a kink point in the frontier of Y. In particular, it would be feasible (since Y is closed), and it would constitute a punishment for the deviation.

Since b_1, b_2 span $\mathbb{R}^2$, there exist $\alpha, \beta \in \mathbb{R}$ such that $n_Y(\theta) = \alpha b_1 + \beta b_2$. Substituting this into equations (4), (5), (6), and then substracting (6) from (4) and (5), we obtain

$$0 < \alpha b_1 \cdot (\theta' - x) - \beta b_2 \cdot (\theta' - x) = \alpha b_1 \cdot (\theta' - \theta'') \tag{7}$$

$$0 < \alpha b_1 \cdot (\theta'' - x) - \beta b_2 \cdot (\theta'' - x) = \beta b_2 \cdot (\theta'' - \theta'), \tag{8}$$

where the equalities follow by the definition of x. And given that $b_1 \cdot \theta' > b_1 \cdot \theta''$ and $b_2 \cdot \theta' < b_2 \cdot \theta''$, (7) and (8) imply $\alpha > 0$ and $\beta > 0$, respectively. Hence $n_Y(\theta) \in int(C(b_1, b_2))$. $\square$

In order to prove Proposition 4 for higher dimensions, we show the following result:

Proposition 9. *Given $Y \subseteq \mathbb{R}^p$ and $C_1 = C(b_1)$, $C_2 = C(b_2)$ with $b_1, b_2 \in \mathbb{R}^p$ linearly independent, the following statements are equivalent:*

[17] Note that $Fr(Y)$ has at most a countable number of kinks. Since Y is convex, $Fr(Y)$ is locally the graph of a concave (convex) function and hence the derivative of this function is monotonic, and it has at most a countable number of jumps.

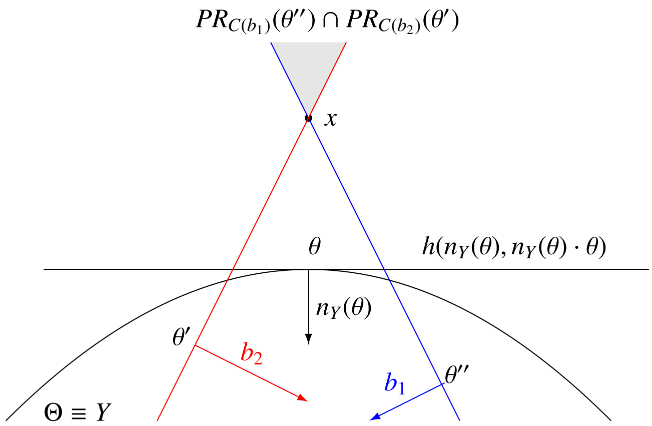

Figure 10: If there is no FRE then the LDC is violated.

(i) *There exists a fully revealing equilibrium in Y.*

(ii) *There exists a fully revealing equilibrium in Y_{b_1,b_2}.*

Proof of Proposition 9: Given $b_1, b_2 \in \mathbb{R}^p$, Π_{b_1,b_2} the plane spanned by b_1, b_2 and $b \in \Pi_{b_1,b_2}$, denote by $H_{b_1,b_2}(b,k) = \{x_{b_1,b_2} \in \Pi_{b_1,b_2} \mid b \cdot x_{b_1,b_2} > k\}$. Then, for any $\theta \in \mathbb{R}^p$ and $b \in \Pi_{b_1,b_2}$,

$$x \in H(b, b\cdot\theta) \iff b\cdot x > b\cdot\theta \iff b\cdot x_{b_1,b_2} > b\cdot\theta_{b_1,b_2} \iff x_{b_1,b_2} \in H_{b_1,b_2}(b, b\cdot\theta_{b_1,b_2}). \quad (9)$$

Given Proposition 1-(i), there exists a FRE in Y given $C(b_1)$ and $C(b_2)$ if and only if for any $\theta', \theta'' \in Y$, $Y \not\subseteq H(b_1, b_1 \cdot \theta'') \cup H(b_2, b_2 \cdot \theta')$. By (9) this is equivalent to the statement that for any $\theta'_{b_1,b_2}, \theta''_{b_1,b_2} \in Y_{b_1,b_2}$, $Y_{b_1,b_2} \not\subseteq H_{b_1,b_2}(b_1, b_1 \cdot \theta''_{b_1,b_2}) \cup H_{b_1,b_2}(b_2, b_2 \cdot \theta'_{b_1,b_2})$ which, given Proposition 1-(i), is equivalent to the existence of a FRE in Y_{b_1,b_2} given $C(b_1)$ and $C(b_2)$. □

The proof of Proposition 4 follows immediately from Propositions 8 and 9 given that we have assumed that Y_{b_1,b_2} is closed. Throughout the paper, Y is assumed to be closed, but since the projection of a closed set onto a plane is not necessarily closed, Y_{b_1,b_2} is not necessarily closed. Proposition 9 holds for any $Y \in \mathbb{R}^p$ (closed or not) but the characterization in Proposition 8 requires Y to be closed. □

Proof of Proposition 5:

In order to prove this result, we introduce some new concepts and intermediate results. For any convex cone C, we denote by C^* the dual cone of C: $C^* = \{v \in \mathbb{R}^p \mid b \cdot v \geq 0 \,\forall b \in C\}$. Note that $PR_C(\theta) = \{x \in \mathbb{R}^p \mid x - \theta \in -C^*\}$.

Lemma 2. *For any convex cone $C \subset \mathbb{R}^p$ and $\theta \in \mathbb{R}^p$, if $x \in PR_C(\theta)$ then $PR_C(x) \subseteq PR_C(\theta)$.*

Proof. For any $z \in PR_C(x)$, $z - x \in -C^*$. Similarly, since $x \in PR_C(\theta)$, $x - \theta \in -C^*$. But then $z - \theta = (z - x) + (x - \theta) \in -C^*$ since $-C^*$ is a cone. Therefore $z \in PR_C(\theta)$. □

Lemma 3. *Given $Y \subseteq \mathbb{R}^p$, $C \subset \mathbb{R}^p$ a convex cone, and $\theta' \in Y$. If $PR_C(\theta') \cap Y$ is unbounded and $dim(PR_C(\theta') \cap Y) = dim(Y)$, then $PR_C(\theta') \cap PR_C(\theta'') \cap Y \neq \emptyset$ for any $\theta'' \in Y$.*

Proof. If $PR_C(\theta') \cap Y$ is unbounded and $dim(PR_C(\theta') \cap Y) = dim(Y)$, there exists $r \in -int(C^*)$ such that $\theta' + \lambda r \in PR_C(\theta') \cap Y$ for any $\lambda \geq 0$. Since $r \in -int(C^*)$, for any vector $s \in \mathbb{R}^p$ there exists $\lambda_s \geq 0$ sufficiently large such that $s + \lambda_s r \in -C^*$. Therefore, for $s = \theta' - \theta''$, there exists λ_s such that $\theta' - \theta'' + \lambda_s r \in -C^*$. But then $\theta' + \lambda_s r \in PR_C(\theta') \cap PR_C(\theta'') \cap Y$. □

Lemma 4. *Consider C_1 and C_2 satisfying Assumption 1 with $C = co(C_1 \cup C_2) = C(b^1, ..., b^m)$. Let Y be a half-space: $Y = \{y \in \mathbb{R}^p \mid n \cdot y \geq k\}$. If $n \notin C \setminus \{b^1, ..., b^m\}$ then there exists an anonymous FRE in Y given C_1 and C_2.*

Proof. If $n \notin C \setminus \{b^1, ..., b^m\}$ then either $n \notin C$ or $n = b^j$ for some $j \in \{1, ..., m\}$.

Case 1: $n \notin C$. Since C is a closed convex cone, $(C^*)^* = C$ by the Bipolar Theorem, and hence $C = \{b \in \mathbb{R}^p \mid b \cdot v \geq 0 \quad \forall v \in C^*\}$. In particular, if $n \notin C$, there exists $v \in C^*$ such that $n \cdot v < 0$, or analogously, there exists a $r \in -C^*$ such that $n \cdot r > 0$. Consider any pair of reports $\theta', \theta'' \in Y$. Since $n \cdot r > 0$ and $r \in -C^*$, $PR_C(\theta') \cap Y$ is unbounded and $dim(PR_C(\theta') \cap Y) = dim(Y)$. Hence by Lemma 3, $PR_C(\theta') \cap PR_C(\theta'') \cap Y \neq \emptyset$, and there exists an anonymous FRE.

Case 2: $n = b^j$ for some $j \in \{1, ..., m\}$. Since C is a convex cone, then there exists $r \in \mathbb{R}^p$ orthogonal to n (i.e., $n \cdot r = 0$) such that $C \subset \{v \in \mathbb{R}^p \mid v \cdot r \geq 0\}$. In particular, for any $\theta', \theta'' \in Y$, there exists $\theta'_n \in Fr(Y) \cap PR_C(\theta')$, $\theta''_n \in Fr(Y) \cap PR_C(\theta'')$ and $dim(PR_C(\theta'_n) \cap Fr(Y)) = dim(Fr(Y)) = p - 1$ and unbounded. Hence by Lemma 3, $PR_C(\theta'_n) \cap PR_C(\theta''_n) \cap Fr(Y) \neq \emptyset$ and by Lemma 2, $PR_C(\theta'_n) \cap PR_C(\theta''_n) \cap Fr(Y) \subset PR_C(\theta') \cap PR_C(\theta'') \cap Y$, so there exists an anonymous FRE. □

We are now ready to prove Proposition 5.

Proof of Proposition 5

$(i) \Rightarrow (ii)$ is trivial.

$(ii) \Rightarrow (iii)$: Suppose first that there exists a smooth $\theta \in Fr(Y)$ such that $n_Y(\theta) \in C(b^1, ..., b^m) \setminus \{b^1, ..., b^m\}$, then there exists $b, b' \in C$ such that $n \in int(C(b, b'))$. Moreover, as θ is smooth, $\theta_{b,b'}$ is a smooth point of $Fr(Y_{b,b'})$ and $n_{Y_{b,b'}}(\theta_{b,b'}) \in int(C(b, b'))$. Hence, by Proposition 4 there is no FRE given $C(b)$ and $C(b')$.

Suppose instead that there exists a partially smooth $\theta \in Fr(Y)$, i.e., $1 < dim(P_Y(\theta)) < p$, such that there exists $n \in P_Y(\theta) \cap int(C(b^1, ..., b^m))$. Since θ is partially smooth, there exists a plane Π such that the projection of θ onto Π is a smooth point of $Fr(Y_\Pi)$, the frontier of the projection of Y onto Π. Moreover, since $n \in int(C(b^1, ..., b^m))$ and $dim(C(b^1, ..., b^m)) = p$, we can choose $b, b' \in C(b^1, ..., b^m)$ such that $n \in int(C(b, b'))$, and $\Pi_{b,b'} = \Pi$. In particular, $n = n_{Y_{b,b'}}(\theta_{b,b'}) \in int(C(b, b'))$, so by Proposition 4 there is no FRE given $C(b)$ and $C(b')$.

$(iii) \Rightarrow (i)$: Denote by $C = co(C_1 \cup C_2) = C(b^1, ..., b^m)$. Suppose for contradiction that there is not an anonymous FRE given C_1 and C_2. If so, there exist $\theta', \theta'' \in Y$ such that

$PR_C(\theta') \cap PR_C(\theta'') \cap Y = \emptyset$. Since both Y and $PR_C(\theta') \cap PR_C(\theta'')$ are convex sets, there exists a supporting hyperplane to Y that strictly separates Y from $PR_C(\theta') \cap PR_C(\theta'')$. In other words, there exists $n \in \mathbb{R}^p$ and $\theta \in Fr(Y)$ such that $Y \subseteq \{x \in \mathbb{R}^p \mid n \cdot x \geq n \cdot \theta\}$ and $PR_C(\theta') \cap PR_C(\theta'') \subset \{x \in \mathbb{R}^p \mid n \cdot x < n \cdot \theta\}$. Note that for any normal vector n to a separating hyperplane supported at θ, $n \in P_Y(\theta)$. We consider in turn the two possible cases:

Case 1: There exists a smooth $\theta \in Fr(Y)$ that supports a hyperplane strictly separating Y from $PR_C(\theta') \cap PR_C(\theta'')$. Denoting by $H = \{x \in \mathbb{R}^p \mid n_Y(\theta) \cdot x \geq n_Y(\theta) \cdot \theta\}$ the supporting half-space that contains Y, we have $PR_C(\theta') \cap PR_C(\theta'') \cap H = \emptyset$, so there is not an anonymous FRE in H given C. But then, by Lemma 4, $n_Y(\theta) \in C \setminus \{b^1, ..., b^m\}$ which contradicts (iii) (a).

Case 2: There are no smooth points in $Fr(Y)$ that support a hyperplane that strictly separates Y from $PR_C(\theta') \cap PR_C(\theta'')$. We show below that there must exist a partially smooth $\theta \in Fr(Y)$ supporting a strictly separating hyperplane whose normal vector n satisfies $n \in int(C)$.

1. If all the strictly separating hyperplanes are supported on a kink point $\theta \in Fr(Y)$, then $C \subseteq P_Y(\theta)$ which implies that $b \cdot \theta \leq b \cdot y$ for all $y \in Y$ and hence θ is a punishment for any deviation and $\theta \in PR_C(\theta') \cap PR_C(\theta'')$, which contradicts the fact that there was a strictly separating hyperplane supported at θ.
2. If a partially smooth point θ supports a strictly separating hyperplane, denoting by $H = \{x \in \mathbb{R}^p \mid n \cdot x \geq n \cdot \theta\}$ the corresponding half-space that containing Y, we have that by Lemma 4, $n \in C \setminus \{b_1, ..., b^m\}$.
3. If all the normal vectors of the supporting strictly separating half-spaces at a partially smooth point θ satisfy that $n \notin int(C)$, then there exists b^i, b^j such that all those normal vectors n lie in the face of C spanned by b^i and b^j. This is because if there is $n, n' \in P_Y(\theta)$ such that n is in the face spanned by b^i, b^j and n' is in the face spanned by b^k, b^l then any convex combination of n and n' will also by in $P_Y(\theta)$, so it will support a half-space strictly separating Y from $PR_C(\theta') \cap PR_C(\theta'')$, and it will be in the interior of C. Moreover, given b^i, b^j, there exists a direction r orthogonal to b^i, b^j, such that θ is smooth in that direction, i.e., r belongs to all the supporting hyperplanes to Y at θ. Finally, since C is strictly contained in a half-space, either $C \subset \{v \in \mathbb{R}^p \mid v \cdot r \geq 0\}$ or $C \subset \{v \in \mathbb{R}^p \mid v \cdot r \leq 0\}$. Without loss of generality we can assume that $b \cdot r \geq 0$ for all $b \in C$.
4. Suppose that for all the partially smooth points supporting strictly separating hyperplanes, and for all the normal vectors to those separating hyperplanes, $n \notin int(C)$.
5. Consider any of those partially smooth point θ and the direction of smoothness r defined in 3. By 4, $\{\theta - \lambda r \mid \lambda \in [0, \Lambda_r]\} \subset Fr(Y)$. If not we would be able to find another partially smooth point supporting a strictly separating hyperplane with $n \in int(C)$. If all the Λ_r are finite, then there exists $\hat{\theta}$ such that $C \subset P_C(\hat{\theta})$, in which case as in 1, $\hat{\theta} \in PR_C(\theta') \cap PR_C(\theta'')$, which contradicts the existence of a strictly separating hyperplane.

If there exists a Λ_r that is not finite, i.e, Y is unbounded in the direction $-r$, we can use such a direction to punish any deviation. in other words, $PR_C(\theta') \cap PR_C(\theta'') \cap FR(Y) \neq \emptyset$, which again contradicts the existence of a strictly separating hyperplane.

Therefore, there must exists a partially smooth point supporting a strictly separating hyperplane with normal vector $n \in int(C)$. But $n \in P_Y(\theta)$ so $n \in P_Y(\theta) \cap int(C)$ which contradicts (iii) (b). □

Proof of Proposition 6:

$\Rightarrow$): Define x_{b_1,b_2} to be the projection of $M_{b_1,b_2}(\theta', \theta'')$ onto Π_{b_1,b_2}, i.e, x_{b_1,b_2} is the point in Π_{b_1,b_2} such that $b_1 \cdot x_{b_1,b_2} = \min\{b_1 \cdot \theta', b_1 \cdot \theta''\}$ and $b_2 \cdot x_{b_1,b_2} = \min\{b_2 \cdot \theta', b_2 \cdot \theta''\}$. If $M_{b_1,b_2}(\theta', \theta'') \cap Y = \emptyset$, then $x_{b_1,b_2} \notin Y_{b_1,b_2}$. Replicating the argument in the proof of Proposition 4, this implies the LDC is violated and hence there is not a FRE.
$\Leftarrow$): Suppose that for all $\theta', \theta'' \in Y$, there exists $y \in M_{b_1,b_2}(\theta', \theta'') \cap Y$. Then $b_1 \cdot y \leq b_1 \cdot \theta''$ and $b_2 \cdot y \leq b_2 \cdot \theta'$ and hence y is a punishment for the deviation (θ', θ''). □

Proof of Proposition 7:

The proposition is a direct consequence of Propositions 5 and 6. □

A.3 Equivalence of Robustness and Continuity on the Diagonal

Ambrus and Takahashi (2008) define a notion of continuity for the receiver's strategy. We state their definition for fully revealing equilibrium:

Definition 3 (Ambrus and Takahashi (2008))**.** A fully revealing equilibrium (s_1, s_2, y^R) is *continuous on the diagonal* if

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

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

We now show that this notion of continuity is equivalent to our definition of robustness in the case of known biases b_1, b_2.

Lemma 5. *A fully revealing equilibrium (s_1, s_2, y^R) is robust if and only if it is continuous on the diagonal.*

Proof. $\Rightarrow$) Consider any pair of sequences $\{(\theta_1^n, \theta_2^n)\}_{n\in\mathbb{N}} \subset \Theta$ such that $\lim_{n\to\infty} \theta_1^n = \lim_{n\to\infty} \theta_2^n = \theta$. Since y^R is robust, for every $\epsilon > 0$ there exists a $\delta > 0$ such that for all $\theta', \theta'' \in B(\theta, \delta) \cap Y$, $y^R(s_1(\theta'), s_2(\theta'')) \in B(\theta, \epsilon)$. Now, $\lim_{n\to\infty} \theta_1^n = \lim_{n\to\infty} \theta_2^n = \theta$ implies that for that $\delta > 0$, there exists $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$, $\theta_1^n, \theta_2^n \in B(\theta, \delta) \cap Y$, which implies that $y^R(s_1(\theta_1^n), s_2(\theta_2^n)) \in B(\theta, \epsilon)$ and hence the equilibrium is continuous on the diagonal.

$\Leftarrow$) We argue by contradiction. Suppose that y^R is not robust. Then there exists $\theta \in \Theta$ and $\epsilon > 0$ such that for all $n_0 \in \mathbb{N}$ there exists $n > n_0$ with θ_1^n, θ_2^n such that $\theta_1^n, \theta_2^n \in B(\theta, \frac{1}{n}) \cap Y$ and

$$y^R(s_1(\theta_1^n), s_2(\theta_2^n)) \notin B(\theta, \epsilon) \setminus \left(B(\theta_1^n + b_2, |b_2|) \cup B(\theta_2^n + b_1, |b_1|)\right).$$

Note that for any n such that $\frac{1}{n} < \epsilon$, $\theta_1^n \neq \theta_2^n$, because if $\theta_1^n = \theta_2^n$, $y^R(s_1(\theta_1^n), s_2(\theta_2^n)) = \theta_1^n \in B(\theta, \epsilon) \setminus \left(B(\theta_1^n + b_2, |b_2|) \cup B(\theta_2^n + b_1, |b_1|)\right)$. Since (s_1, s_2, y^R) is an equilibrium, $y^R(s_1(\theta_1^n), s_2(\theta_2^n)) \notin B(\theta_1^n + b_2, |b_2|) \cup B(\theta_2^n + b_1, |b_1|)$, otherwise either sender S_1 would have an incentive to deviate to $s_1(\theta_1^n)$ when θ_2^n is realized, or sender S_2 would have an incentive to deviate to $s_2(\theta_2^n)$ when θ_1^n is realized. Hence $y^R(s_1(\theta_1^n), s_2(\theta_2^n)) \notin B(\theta, \epsilon)$, which contradicts the diagonal continuity of the equilibrium. $\square$

References

Alonso, R., W. D. and Matouschek, N. (2008). When Does Coordination Require Centralization?, *American Economic Review* **98**: 145–179.

Ambrus, A. and Lu, S. E. (2014). Almost Fully Revealing Cheap Talk with Imperfectly Informed Senders, *Games and Economic Behavior* **88**: 174–189.

Ambrus, A. and Takahashi, S. (2008). Multi-Sender Cheap Talk with Restricted State Spaces, *Theoretical Economics* **3**: 1–27.

Battaglini, M. (2002). Multiple Referrals and Multidimensional Cheap Talk, *Econometrica* **70**: 1379–1401.

Crawford, V. P. and Sobel, J. (1982). Strategic Information Transmission, *Econometrica* **50**: 1431–1451.

Gilligan, T. W. and Krehbiel, K. (1989). Asymmetric Information and Legislative Rules with a Heterogeneous Committee, *American Journal of Political Science* **33**: 459–490.

Krishna, V. and Morgan, J. (2001a). A Model of Expertise, *Quarterly Journal of Economics* **116**: 747–775.

Krishna, V. and Morgan, J. (2001b). Asymmetric Information and Legislative Rules: Some Amendments, *American Political Science Review* **95**: 435–452.

Lai, E. K., Lim, W. and Wang, J. T. (2015). An Experimental Analysis of Multidimensional Cheap Talk, *Games and Economic Behavior* **91**: 114–144.

Meyer, M., de Barreda, I. M. and Nafziger, J. (2016). Robustness of Full Revelation in Multisender Cheap Talk, *CEPR Discussion Paper 11224* .

Rubanov, O. (2015). Asymptotic Full Revelation in Cheap Talk with a Large Number of Senders. Working Paper.