School of Economics and Finance

Competition of E-Commerce Intermediaries

Tymofiy Mylovanov and Andriy Zapechelnyuk

Working Paper No. 674 November 2010 ISSN 1473-0278

Queen Mary
University of London

Decision rules for experts with opposing interests*

Tymofiy Mylovanov
Penn State University†

Andriy Zapechelnyuk
Queen Mary, University of London‡

November 7, 2010

Abstract

This paper studies optimal decision rules for a decision maker who can consult two experts in an environment without monetary payments. This extends the previous work by Holmström (1984) and Alonso and Matouschek (2008) who consider environments with one expert. In order to derive optimal decision rules, we prove a "constant-threat" result that states that any out-of-equilibrium pair of recommendations by the experts are punished with an action that is *independent* of their reports. A particular property of an optimal decision rule is that it is simple and constant for a large set of experts' preferences and distribution of their private information. Hence, it is robust in the sense that it is not affected by errors in specifying these features of the environment. By contrast, the constructions of optimal outcomes absent commitment or with only one expert are sensitive to model details.

JEL classificaton: C72, D82, D83

Keywords: Communication; information; noise; experts; constant threat

---

*We thank Ethem Akyol, Attila Ambrus, Mark Amstrong, Vincent Crawford, Peter Esö, Christian Hellwig, Johannes Hörner, Navin Kartik, Daniel Krähmer, Vijay Krishna, Stephan Lauermann, Gilat Levy, Ming Li, David Martimort, Rony Razin, and Larry Samuelson for helpful comments.

†Department of Economics, Penn State University, Kern Building, University Park, PA 16801, USA *Email*: mylovanov $\alpha\tau$ gmail.com

‡School of Economics and Finance, Queen Mary, University of London, Mile End Road, London E1 4NS, UK. *E-mail:* a.zapechelnyuk $\alpha\tau$ qmul.ac.uk

1 Introduction

The Consolidated Appropriations Act, 2010 of the United States of America grants certain types of auto dealerships that were terminated by the manufacturer on or before April 29, 2009 the right to seek through binding arbitration reinstatement of the dealership agreement.[1] The law mandates that the arbitrator shall balance the economic interest of the dealership, the economic interest of the manufacturer, and the economic interest of the public at large and shall decide whether or not the dealership should be added to the dealer network of the manufacturer.

This is an example of an environment with two parties (a manufacturer and a dealership) who have decision relevant information and whose preferences are not perfectly aligned with each other's and with those of the principal (the government and, more generally, the public). The arbitration is a binding mechanism to which the parties resort if they fail to reach an agreement by themselves.

Two features of the arbitration procedure are noteworthy. First, the arbitrator is restricted in the scope of the award – she can either reinstate the original dealership agreement or not, but the law does not allow the arbitrator to modify the agreement. Second, the arbitration decision is inherently uncertain from the perspective of the dealership and the manufacturer due to their imperfect knowledge about the arbitrator's views on how to balance economic interests at stake. These features are not atypical in arbitration and have been discussed in the law literature.[2]

In this paper, we study optimal decision rules in environments with two informed experts (agents) and an uninformed decision maker (principal). Our results suggest that randomness of arbitration award and its limited scope can be desirable. Furthermore, it can be optimal for the arbitrator's decision, while random, to be independent of the arguments supplied by the parties.

Our model features two strategic experts who are biased in different directions (other things being equal, the dealership is more eager than the manufacturer to get the dealership agreement renewed).[3] The set of feasible actions is a unit interval. A socially optimal

[1]See H.R.3288 - 187, Section 747 of the Act.

[2]The questions of interest in that literature are whether limitations on the remedies that can be awarded by the arbitrator should be enforceable (see, e.g., D. S. Schwartz, "Understanding Remedy-Stripping Arbitration Clauses: Validity, Arbitrability, and Preclusion Principles," 38 U.S.F. L. Rev. 49 (2003-2004), pp. 49-104) and whether randomness in arbitration decisions justifies allowing for ex-post judicial review (see, e.g., S. P. Younger, "Agreements to Expand the Scope of Judicial Review of Arbitration Awards," 63 Alb. L. Rev. 241 (1999-2000), pp. 241-262 and L. Goldman, "Contractually Expanded Review of Arbitration Awards," 8 Harv. Negot. L. Rev. 171 (2003), pp. 171-200).

[3]We discuss the case of similarly biased experts in Section 6.

action is represented by an uncertain state. The experts have unverifiable information about the state. In the first part of the paper, we assume that the experts have identical information and know the state. Later on, we allow the experts to have asymmetric information. There are no monetary payments. A decision rule is a contract that implements an action contingent on the experts' reports about their information.[4] The purpose of the paper is to identify optimal decision rules.

Consider an environment in which the experts have identical information. By the revelation principle, an optimal decision rule can be sought for among direct decision rules in which the experts reports their information truthfully. Thus, the decision rule must provide incentives for the experts to tell the truth through punishment of disagreements. The difficulty here is that when a disagreement is observed, the designer cannot detect which expert deviates from truthful reporting. Furthermore, the experts' opposing interests imply that an action that is a stronger punishment for one expert is often a weaker punishment for the other one. In cheap talk communication models, where the decision maker cannot commit to a decision rule, this issue causes punishments in fully revealing equilibria to depend non-trivially on the experts' reports (Krishna and Morgan, 2001a,b; Battaglini, 2002; Ambrus and Takahashi, 2008)[5] and also makes it difficult to characterize optimal equilibria if full revelation is not feasible.

Proposition 1 is the key insight of the paper. It offers a surprisingly simple way to construct an optimal decision rule for a decision maker with commitment power: we show that one can restrict attention to "constant-threat" rules in which every disagreement between the experts is punished by the same (stochastic) action with a two point support. The proof of this constant-threat result relies in a curious way on a minmax inequality and concavity of the experts' payoff functions.

For our introductory example, the constant-threat result suggests that the optimal decision rule should entail a threat of random arbitration with a restricted scope of the award to provide incentives for the parties to agree on the optimal course of action. Of course, while fully random arbitration might be difficult to implement in practice, at least some randomness and restrictions on the scope of the award can be welfare improving.

An optimal decision rule is constant and very simple across a large set of environments.

[4]This setting is reminiscent of Kalai and Rosenthal (1976) who address the question of implementation through binding arbitration of an efficient outcome in a finite two-player game, where the arbitrator is uninformed about the players' payoff functions. The key difference is that there is no exogenously specified status-quo in our model.

[5]Crawford and Sobel (1982) is the seminal reference on cheap talk communication with one expert. For models of cheap talk communication with two experts see also Krishna and Morgan (2001b); Battaglini (2004); Ambrus and Lu (2009); Li (2008, 2009).

It can be implemented by an indirect mechanism in which both experts suggest an action; if they agree, the suggested recommendation is followed; otherwise, the rule (uniformly) randomizes between the two most extreme actions. The class of environments in which this decision rule is optimal includes those in which the first best outcome can be implemented (Corollary 1) and symmetric environments (Remark 2). This rule is robust to small mistakes in specifications of the experts' preferences and information. By contrast, the constructions of optimal outcomes in cheap talk environments are highly sensitive to these details of the model (Krishna and Morgan, 2001a; Battaglini, 2002; Ambrus and Takahashi, 2008).[6]

The cheap talk literature with two experts has focused on establishing conditions under which the decision maker can achieve the *first best* outcome (Krishna and Morgan, 2001b; Battaglini, 2002; Ambrus and Takahashi, 2008). For comparison, we provide conditions (Proposition 2 and Remark 1) for the first best to be implementable in our environment; the conditions bound the size of bias of each expert. Naturally, these conditions are related to but *weaker* than those in cheap talk environments. For sufficiently small biases, the first best can be attained both with and without commitment. In these circumstances, the role of commitment is to permit a simpler and more robust decision rule. For the intermediate range of biases, the role of commitment is stronger: the first best could be implemented only under commitment. In this case, the ability of the decision maker to commit to a stochastic action out of equilibrium confers an additional benefit.

Propositions 3 – 5 characterize optimal decision rules in the environments in which the first best outcome is not implementable. In particular, adding a second expert always strictly improves the payoff of the decision maker relative to what she would obtain with one expert (Proposition 3). These results are especially useful because a characterization of the best equilibria in cheap talk environments beyond those permitting the first best remains an open question for many environments.

The assumption that the experts have identical information is standard in the related literature.[7] Yet, it is an important assumption in that it allows the decision maker to check the reports of the experts against each other, and inconsistent reports do not occur on the equilibrium path. Therefore, we study noisy environments in Section 5. Our first observation is that the constant-threat result extends to these environments.[8] Further-

[6] Furthermore, as pointed out by Battaglini (2002), these cheap talk equilibria contain implausible out-of-equilibrium beliefs. By contrast, there is no issue of out-of-equilibrium beliefs in our model because the decision maker can commit to her actions.

[7] See references in Section 5.

[8] Though it may become vacuous if in any state the experts' signals have full support.

more, we show that if the state space is finite, then for any sequence of diminishing noise it is possible to construct decision rules that converge to the optimal decision rule in the environment without noise (Remarks 3 and 4). A similar result holds for the environments in which the state space is infinite and the noise is partitional: the experts observe an element of a partition of the state space that contains the realized state (Proposition 6). The proofs of these results are constructive and offer a way to design decision rules that perform close to optimum in environments with small noise. They complement the results in Ambrus and Lu (2009) who show how to construct fully revealing equilibria robust to noise in cheap talk environments.

The results presented here focus on partial implementation of the optimal outcomes and the constructed decision rules permit multiple equilibria. We discuss the issue of full implementation and offer some partial positive results in Section 6.

This paper is a natural continuation of the work on optimal decision rules in environments with one informed expert and no monetary payments that was initiated by Holmström (1977, 1984).[9] We compare our results with the optimal decision rules identified in this literature in Section 6.

Our paper is related to Battaglini (2004) who considers a multidimensional environment with multiple experts and noisy signals. Battaglini shows that minimal commitment power is sufficient to implement an outcome arbitrarily close to the first best as the number of experts becomes sufficiently high.

The problem of optimal decision rules for two experts has been studied in Martimort and Semenov (2008). Our models and approaches are quite different. In particular, they focus on experts who are biased in the same direction and consider dominant strategy implementation. By contrast with our results, Martimort and Semenov (2008) demonstrate impossibility of the first best outcome and show that a sufficiently high bias renders the experts not valuable for the decision maker.

Finally, Esö and Fong (2010) show that the first best outcome can be implemented in a dynamic cheap talk environment in which the decision maker can delay the final decision by choosing an "inaction". In their model, delays are costly: both experts prefer

[9]See Holmström (1977, 1984), Melumad and Shibano (1991), Dessein (2002), Martimort and Semenov (2006), Alonso and Matouschek (2008), Martimort and Semenov (2008), Goltsman et al. (2009), Kovac and Mylovanov (2009), Amador and Bagwell (2010). Armstrong and Vickers (2008), Koessler and Martimort (2009), Li and Li (2009), and Lim (2009) who study optimal decision rules in environments which are related, but not identical to the model of Holmström (1977, 1984). Optimal decision rules for environments in which a decision maker can commit to monetary payments are characterized in Baron (2000), Krishna and Morgan (2008), Bester and Krähmer (2008), Raith (2008), and Ambrus and Egorov (2009).

the decision maker's ideal action to inaction. Our model is different in that we do not assume existence of an outcome with properties of inaction in their model.

The remainder of the paper is organized as follows. Section 2 describes the model. The constant-threat result is derived in Section 3. Applying this result, we characterize the optimal decision rules in Section 4 for the environments with identical information of the experts. Section 5 studies robustness of the optimal decision rules with respect to noise in the experts' information. Section 6 discusses the questions of full implementation, similarly biased experts, and compares the results with those for the environment with one expert.

2 The Model

There are two experts $i = 1, 2$ and a decision maker. The decision maker has to select an action from a set $Y = [0, 1]$ of feasible actions. The most preferred action for the decision maker (the *state*), $x \in X$, is a realized value of a random variable $\tilde{x}$ with support on X. We assume that X is a closed measurable subset of Y that contains the endpoints, 0 and 1. That is, our model includes both the environment in which the set of states is finite and the environment in which the set of states is a compact interval. The restriction to the unit interval is not essential.[10]

The decision maker is uninformed about x and believes that the distribution of $\tilde{x}$ is represented by a c.d.f. F.

We begin with the assumption that the experts know precisely the value of x (it will be relaxed in Section 5). The decision maker can ask for recommendations and commit to take an action that is contingent on their reports.

Let y denote an action. The payoff function of the decision maker is denoted by $u_0(x, y)$ and the payoff functions of expert $i = 1, 2$ by $u_i(x, y)$. We assume that for every $x \in X$ each function $u_i(x, y)$, $i = 0, 1, 2$, is strictly concave in y. The decision maker's payoff function is maximized at the action equal to the state,

$$\arg\max_{y \in Y} u_0(x, y) = \{x\},\ x \in X.$$

For every $x \in X$ we define $\{y_i^*(x)\} = \arg\max_{y \in Y} u_i(x, y)$, $i = 1, 2$. We assume that the

[10] The assumption that the action space is bounded does not drive our results and can be relaxed. In fact, a larger action space makes eliciting information easier for the decision maker. Ultimately, if the action space is unbounded or sufficiently large relative to the state space, the first best outcome can be implemented by threatening experts with extreme actions out of equilibrium.

experts have *opposing interests*:

$$y_1^*(x) < x < y_2^*(x), \quad \text{for every } x \in X. \tag{1}$$

Some of our results are obtained for the environment with quadratic preferences and fixed biases, which is standard in the literature on experts: $u_0(x,y) = -(x-y)^2$, $u_1(x,y) = -((x-b_1)-y)^2$ and $u_2(x,y) = -((x+b_2)-y)^2$, where $b_1, b_2 > 0$.

Let $\mathcal{Y}$ denote the set of distributions on Y (*randomized actions*). Identifying point distributions with points we have $Y \subset \mathcal{Y}$. We extend the definition of u_i to $X \times \mathcal{Y}$ via the statistical expectation:

$$u_i(x,\lambda) = \int u_i(x,y)\lambda(dy), \quad x \in X, \ \lambda \in \mathcal{Y}.$$

A *decision rule* is a measurable function

$$\mu: \quad X^2 \to \mathcal{Y}, \quad (x_1,x_2) \mapsto \mu(x_1,x_2),$$

where $\mu(x_1,x_2)$ is a randomized action that is contingent on the experts' reports (x_1,x_2). A decision rule induces a *game* (a direct mechanism), in which after observing x the experts simultaneously make reports $x_1, x_2 \in X$ and the outcome $\mu(x_1,x_2)$ is implemented.

A decision rule μ is *incentive compatible* if truthtelling, $x_1 = x_2 = x$, is a Nash equilibrium: for all $x, x' \in X$

$$\begin{aligned} u_1(x,\mu(x,x)) &\geq u_1(x,\mu(x',x)), \\ u_2(x,\mu(x,x)) &\geq u_2(x,\mu(x,x')). \end{aligned} \tag{2}$$

By the revelation principle, any equilibrium outcome of the experts' interaction in a game whose space of outcomes is Y or $\mathcal{Y}$ can be represented by the truth-telling equilibrium outcome in some incentive compatible decision rule. In what follows, we will consider only incentive compatible decision rules.

A decision rule μ is *optimal* if it maximizes the expected payoff of the decision maker,

$$v^\mu = \int_X u_0(x,\mu(x,x))dF(x),$$

among all incentive compatible decision rules. Since the set of incentive compatible decision rules is compact in weak topology and v^μ is continuous in μ, an optimal decision rule exists.

3 The Constant-threat Result

In any decision rule, the main incentive issue is to motivate each expert to agree with the other expert who is expected to tell the truth. Therefore, the decision rule must punish disagreements. The difficulty here is that if a disagreement is observed (i) it is unclear which expert, if any, tells the truth, and (ii) since the experts have opposing interests, a punishment that is more severe for one of the experts tends to benefit the other expert. As a result, a punishment after a disagreement may depend non-trivially on the experts' reports. In particular, this is so in the existing constructions of optimal outcomes in cheap talk environments (Krishna and Morgan, 2001b; Battaglini, 2002; Ambrus and Takahashi, 2008).

We now prove our key result, the constant-threat result, which allows us to characterize optimal decision rules. It states that one can restrict attention to decision rules in which the lottery implemented after a disagreement has support on extreme actions 0 and 1 and is *independent* of the reports. This result reduces the problem of finding optimal decision rules to the problem of finding actions that are implemented if the experts report their information truthfully, $\mu(x, x)$, and the probability of implementing $y = 1$ after a disagreement. Thus, it drastically decreases complexity of the design problem, because we avoid the optimization problem in which we search on a continuum of lotteries with support on Y that are implemented after a disagreement (one threat lottery for each pair of reports $x_1, x_2 \in X, x_1 \neq x_2$).

The idea behind the constant-threat result is as follows. First, by concavity of the experts' payoff functions, any lottery over actions implemented after a disagreement can be replaced, using a mean-preserving spread, by a lottery between actions 0 and 1 without affecting the experts' incentives to report the truth.

Now, let μ be a decision rule in which a disagreement always results in a lottery with support $\{0, 1\}$. The crucial step in the proof is to observe that, say, expert 1 (who is left-biased) always prefers action x to the extreme right action 1. Hence, in all states where a disagreement lottery is better than x, his payoff from action 0 must be strictly greater than that from x. It follows that in these states his expected payoff from that lottery must be *decreasing* in the probability assigned on action 1 (similarly, the payoff of expert 2 from a disagreement lottery must be *increasing* in the probability assigned on action 1). Let $\underline{r}$ be the lottery that achieves the highest payoff for expert 1 among the lotteries with support $\{0, 1\}$ that can be achieved by the best deviations of expert 1 in various states $x \in X$. Denote by $\underline{p}$ the probability this lottery assigns to action 1. Define $\overline{p}$ for expert 2 analogously. The result now follows from the observation that $\underline{p} \leq \overline{p}$, which relies on the fact that minmax is larger than or equal to maxmin. Hence, there exists a lottery c that assigns probability p^c to action 1, where $\underline{p} \leq p^c \leq \overline{p}$, such that replacing every threat lottery with c does not violate the incentive constraints of the experts.

Let $\mathcal{Y}^*$ be the set of probability distributions with support on $\{0, 1\}$. We say that a decision rule $\mu = (x_1, x_2)$ is *constant-threat* if it is incentive compatible and

(C) there exists $c \in \mathcal{Y}^*$ such that $\mu(x_1, x_2) = c$ whenever $x_1 \neq x_2$.

We say that two incentive compatible decision rules, μ and μ', are *equivalent* if they implement the same action whenever the reports of the experts coincide, i.e., $\mu(x, x) = \mu'(x, x)$ for all $x \in X$. Thus, two equivalent decision rules implement identical actions *in equilibrium*, but may implement different actions off-equilibrium.

Proposition 1 (Constant-threat result) *For every optimal decision rule there exists an equivalent constant-threat decision rule.*

Note that a constant-threat decision rule which is equivalent to some optimal decision rule must be optimal as well, since it implements the same actions in equilibrium.

Proof. Let μ be an optimal decision rule. Observe that by concavity of $u_i(x, y)$ in y, $i = 1, 2$, for any measure λ,

$$\int u_i(x, y)\lambda(dy) \geq \left(1 - \int y\lambda(dy)\right) u_i(x, 0) + \left(\int y\lambda(dy)\right) u_i(x, 1), \quad x \in X.$$

Hence, replacing $\mu(x_1, x_2)$, $x_1 \neq x_2$, by a lottery that puts probability $\int y\mu(x_1, x_2)(dy)$ on action 1 and the complementary probability on action 0 will not violate the incentive constraints of the experts. Therefore, there exists an equivalent decision rule μ' in which every threat lottery implemented after a disagreement has support on $\{0, 1\}$.

We now show that there exists a constant-threat decision rule μ^c equivalent to μ'. For every pair of different reports, $x_1, x_2 \in X, x_1 \neq x_2$, let $P(x_1, x_2)$ be the probability that $\mu'(x_1, x_2)$ assigns to 1 after a disagreement. We extend the definition of $P(\cdot, \cdot)$ to X^2 by setting $P(x, x) = \int y\mu'(x, x)(dy)$ for all $x \in X$. Define

$$\mathcal{P}_1(x) = \{P(x', x) | x' \in X\} \ \text{ and } \ \mathcal{P}_2(x) = \{P(x, x') | x' \in X\}.$$

For all $x \in X$, $p \in [0, 1]$, and $i = 1, 2$ let

$$D_i(x, p) = \max\{0, pu_i(x, 1) + (1 - p)u_i(x, 0) - u_i(x, \mu(x, x))\}.$$

By construction, a deviation by expert i in state x leading to a lottery in $\mathcal{Y}^*$ that assigns probability $p \in [0,1]$ to action 1 is non-profitable iff $D_i(x,p) = 0$. Furthermore, by definition of $P(x,x)$,

$$D_i(x, P(x,x)) = 0, \quad x \in X, i = 1, 2.$$

Thus, incentive constraints (2) can be written as

$$D_i(x,p) = 0, \quad x \in X, \ p \in \mathcal{P}_i(x), \ i = 1, 2. \tag{IC}$$

We now show that

$$\begin{aligned} &D_1(x,p) \text{ is non-increasing in } p \text{ for every } x \in X; \\ &D_2(x,p) \text{ is non-decreasing in } p \text{ for every } x \in X. \end{aligned} \tag{*}$$

We start by showing that we can restrict attention to decision rules that on the equilibrium path are deterministic and implement actions that are bounded by the experts' most preferred actions,

$$\mu(x,x) \in Y, \quad y_1^*(x) \leq \mu(x,x) \leq y_2^*(x), \quad x \in X. \tag{P_3}$$

To see why this is true, fix some $x' \in X$ and suppose first that $\mu(x',x')$ is a proper lottery. Then, concavity of the payoff functions implies that replacing $\mu(x',x')$ with the expected value of this lottery improves the payoffs of all players without violating any incentive constraints. Next, let $\mu(x',x') = y' \in Y$, $y' > y_2^*(x')$ for some $x' \in X$. Since $y_2^*(x')$ is closer than y' to the most preferred alternatives of all players, concavity of the payoff functions implies that setting $\mu(x',x') = y_2^*(x')$ improves the payoffs of all parties on the equilibrium path without violating incentive constraints. A symmetric argument is valid for $\mu(x',x') = y' < y_1^*(x)$.

Since $u_1(x,y)$ is concave in y and $y_1^*(x) \leq \mu(x,x)$ by (P_3), it follows that $u_1(x,y)$ is decreasing in y on $[\mu(x,x),1]$ for every x, and hence

$$u_1(x, \mu(x,x)) \geq u_1(x,1).$$

If, in addition, $u_1(x,\mu(x,x)) \geq u_1(x,0)$, then, $D_1(x,p) = 0$ for every $p \in [0,1]$. On the other hand, if $u_1(x,\mu(x,x)) < u_1(x,0)$, then $u_1(x,1) < u_1(x,0)$, and hence $D_1(x,p)$ is decreasing in p. This establishes the first statement in (*). The argument for the second statement is analogous.

Next, let

$$\begin{aligned} a_1(x) &= \inf \mathcal{P}_1(x), \quad x \in X; \\ a_2(x) &= \sup \mathcal{P}_2(x), \quad x \in X. \end{aligned}$$

By (IC) and continuity of $D(x,p)$ w.r.t. p, we have $D_i(x, a_i(x)) = 0$ for $x \in X$. By (*),

$$\begin{aligned} D_1(x,p) = &\ 0, \quad p \geq a_1(x), \ x \in X; \\ D_2(x,p) = &\ 0, \quad p \leq a_2(x), \ x \in X. \end{aligned} \tag{3}$$

Define

$$\begin{aligned} \underline{p} &= \sup_{x \in X} a_1(x) = \sup_{x \in X} \inf P_1(x) = \sup_{x \in X} \inf_{x' \in X} P(x', x); \\ \overline{p} &= \inf_{x \in X} a_2(x) = \inf_{x \in X} \sup P_2(x) = \inf_{x' \in X} \sup_{x \in X} P(x', x). \end{aligned}$$

Then, there exists p^c such that $\underline{p} \leq p^c \leq \overline{p}$. By (3),

$$D_i(x, p^c) = 0, \quad x \in X, i = 1, 2.$$

The result now follows from (IC). ■

The result in Proposition 1 can be generalized. We say that an incentive compatible decision rule is *undominated* if there does not exist another incentive compatible decision rule that yields to all players a greater (equilibrium) payoff in every state and a strictly greater payoff in some state. The arguments behind Proposition 1 are not affected if we consider undominated decision rules instead of optimal decision rules.

In the remainder of the paper, we will study optimal decision rules in the set of constant-threat decision rules. Typically, however, there exist decision rules that induce the same equilibrium outcome and are not constant-threat.

Finally, we would like to remark on the multiplicity of equilibria in the constant-threat decision rules. In this paper, we focus on the truthtelling equilibria; this is justified by the revelation principle. At the same time, there are many other equilibria in a given constant-threat decision rule. For instance, there is always a "babbling" equilibrium in which, irrespective of the true state, both experts report some $x' \in X$ such that $\mu(x', x')$

is equal to the expected value of the threat lottery.[11]

4 Optimal Decision Rules

The constant-threat result in Proposition 1 makes characterization of optimal constant-threat decision rules a simple exercise. In this section, we use it to characterize these rules and obtain more specific results under some additional assumptions about the environment.

4.1 First Best Decision Rules

We start our analysis of optimal constant-threat decision rules by identifying conditions under which they implement the most preferred alternative of the decision maker. Let $\mathcal{C}$ be the set of incentive compatible constant-threat decision rules. A decision rule in $\mathcal{C}$ that in each state implements the most preferred action for the decision maker, if it exists, is called *first best*.

We assume that each expert's utility depend only on the distance between her most preferred action and the implemented action: for each $i = 1, 2$

$$u_i(x, y) = -d_i(y - (x + b_i(x))), \tag{4}$$

where $d_i : \mathbb{R} \to \mathbb{R}_+$ is a convex differentiable function which achieves its minimum at zero and which is symmetric around zero, i.e., $d_i(z) = d_i(-z)$ for all $z \in \mathbb{R}$, and $b_i : X \to \mathbb{R}$, $i = 1, 2$, $b_1(x) < 0 < b_2(x)$. The point $x + b_i(x)$ is the most preferred action of i in state x. The values of b_1 and b_2 reflect the conflict of preferences between the experts and the decision maker and are called the experts' *biases*.

The next result provides a sufficient condition for existence of the first best decision rule under these assumptions.

Proposition 2 *Assume that (4) holds. Then there exists the first best decision rule if*

$$\sup_{x \in X, i=1,2} |b_i(x)| \leq 1/2.$$

[11] We discuss this issue in more detail in Section 6, where we show that the decision rules identified in this paper can be modified to ensure uniqueness of the equilibrium outcome if the experts have lexicographic preferences for truthful reporting.

Proof. There exists the first best decision rule if and only if there is $p \in [0,1]$ such that for each expert $i = 1,2$ and for every $x \in X$,

$$u_i(x,x) \geq (1-p)u_i(x,0) + pu_i(x,1). \tag{5}$$

By convexity of d_i, we have for $i = 1,2$ and for every $x \in X$,

$$\begin{aligned} \frac{d_i(x+b_i(x))}{2} + \frac{d_i(1-x-b_i(x))}{2} &\geq d_i\left(\frac{x+b_i(x)}{2} + \frac{1-x-b_i(x)}{2}\right) \\ &= d_i(1/2) \geq d_i(b_i(x)), \end{aligned} \tag{6}$$

where the second inequality follows from the assumption that $\sup\limits_{x\in X, i=1,2} |b_i(x)| \leq 1/2$. Observe that (6) is equivalent to (5) with $p = 1/2$, which implies existence of the first best decision rule with the threat lottery that assigns equal probabilities to 0 and 1. ■

The first best decision rule constructed in the proof of Proposition 2 uses as a threat the lottery that mixes with equal probability between 0 and 1. The logic behind the construction is straightforward: if the experts' biases are not too large, they are better off under the decision maker's most preferred alternative rather than the threat lottery. It is interesting to note that a symmetric threat lottery is optimal even if the experts' biases are not equal.

Under some additional structure of the payoff functions, the sufficient condition in Proposition 2 becomes necessary.

Remark 1 Let[12] $X = [0,1]$ and assume that the biases are constant, $b_i(x) = \tilde{b}_i$, $i = 1,2$, and have the opposite signs, $\tilde{b}_1 < 0 < \tilde{b}_2$. Then there does not exist the first best decision rule whenever $\max |\tilde{b}_i| > 1/2$.

Proof. Assume that $\tilde{b}_2 > 1/2$. First, let $p > 0$. Then,

$$(1-p)d_2(\tilde{b}_2) + pd_2(1-\tilde{b}_2) < (1-p)d_2(\tilde{b}_2) + pd_2(\tilde{b}_2) = d_2(\tilde{b}_2),$$

which contradicts (5) for $x = 0$ and $i = 2$. Next, let $p = 0$. Since $\tilde{b}_1 < 0$, we obtain for x that satisfies $0 < x < |\tilde{b}_1|$

$$u_1(x,0) \equiv -d_1(-x-\tilde{b}_1) > -d_1(\tilde{b}_1) \equiv u_1(x,x),$$

[12]In fact, it suffices to assume that X contains arbitrarily small neighborhoods in $[0,1]$ of the endpoints 0 and 1.

which contradicts (5) for $i = 1$. The argument for $\tilde{b}_1 < -1/2$ is symmetric. ■

The above results are related to Krishna and Morgan (2001a), Battaglini (2002), and Ambrus and Takahashi (2008) who study cheap talk communication with two experts. For the environment considered in Remark 1, Proposition 1 in Battaglini (2002) establishes that a necessary and sufficient condition for a fully revealing cheap talk equilibrium is that the sum of the absolute values of the experts' biases is less than half of the measure of the action space.[13] Proposition 2 and Remark 1 complement this result by providing necessary and sufficient conditions for the first best outcome under commitment. Our condition is weaker and it bounds the size of *each* expert's bias rather than their *sum*; interestingly, the value of the bound is the same in both environments.

The construction of fully revealing equilibria in cheap talk and our construction of a first best decision rule are analogous but not identical. In a cheap talk environment, for any pair of disagreeing reports there is a threat action such that an expert who can induce this pair of reports prefers the first best outcome to the threat action. This threat action is supported by (out-of-equilibrium) beliefs that make it optimal. The proof then verifies that for each pair of states (reports) there exists a threat action that satisfies a number of inequalities that depend on biases of the experts; in equilibrium, the threat action might have to depend non-trivially and discontinuously on the reports of the experts.

By contrast, in our model a decision rule can use lotteries as threat actions that cannot be supported in a cheap talk model, even out of equilibrium.[14] The proof of the possibility of the first best in our environment employs a constant-threat lottery that mixes equally between the extreme actions and makes use of concavity property of payoff functions. Furthermore, the proof of the necessary condition relies on the fact that it is sufficient to consider report-independent threat lotteries.

4.2 Robustness of the First Best Decision Rule

An interesting implication of the above results is that the lottery that mixes between 0 and 1 with equal probability is a sufficient threat for implementing the first best if the biases are not too large. This is so even if the biases are *not* symmetric.

[13] Krishna and Morgan (2001a) prove that a sufficient condition for a fully revealing cheap talk equilibrium in an environment with constant and equal opposing biases is that each expert's bias is less than 1/4.

[14] The concavity of payoff functions implies that a lottery cannot be a best response for the decision maker.

Corollary 1 *Let the conditions in Proposition 2 be satisfied. Then, the first best decision rule is constant in the preferences of the experts.*

The constancy of the optimal decision rule is a useful feature if the decision maker is concerned about robustness of the decision rule with respect to her knowledge of the environment. In particular, if the optimal decision rule is constant, then the decision maker need not possess correct knowledge about the magnitude and the direction of the experts' biases.

4.3 Second Best Decision Rules

What are the properties of an optimal decision rule if the first best outcome cannot be implemented? In what follows, we characterize optimal decision rules that, given the threat lottery, maximize the payoff of the decision maker in each state.[15]

Observe that any decision rule in $\mathcal{C}$ can be identified by a pair

$$(p, g): \quad p \in [0, 1], \quad g : X \to Y,$$

where p is the probability of action 1 after a disagreement and $g(x)$ is the action implemented on the equilibrium path (recall that decision rules in $\mathcal{C}$ have lotteries with support on $\{0, 1\}$).[16]

Assume that $X = [0, 1]$. Let us pick a constant-threat decision rule (p, g) in $\mathcal{C}$. By concavity of payoff functions, both experts prefer $y = x$ in state $x = p$ to the threat lottery,

$$u_i(p, p) \geq p u_i(p, 1) + (1 - p) u_i(p, 0).$$

This implies that an optimal decision rule implements the most preferred alternative for the decision maker, $g(x) = x$, at least in state $x = p$. In addition, since the experts' payoff functions are strictly concave, we obtain $g(x) = x$ for a proper interval containing $p \in (0, 1)$.

Proposition 3 *Let $X = [0, 1]$. Then, an optimal decision rule implements the most preferred alternative of the decision maker on an interval of states.*

This observation highlights the value of two experts for the decision maker. Clearly, the decision maker is weakly better off with two experts than with either of them alone

[15] Trivially, there also exists a continuum of other decision rules that deliver the same expected payoff for the decision maker but do not have this property for a set of states of measure zero.

[16] We ignore decision rules in $\mathcal{C}$ that are stochastic on the equilibrium path as they cannot be optimal.

as she can always implement a decision rule that would be optimal in the environment with one expert ignoring the existence of the other. However, that the decision maker is strictly better off with two experts is particularly clear if their biases are high: Whereas there is no value in employing one expert if she is sufficiently biased,[17] with two experts there exists a decision rule which implements the most preferred action of the decision maker at least in some states.

We now describe the structure of an optimal decision rule in states where the first best outcome is not incentive compatible. For a given probability p of action 1 in an optimal constant-threat decision rule, let $\tilde{X}_i^p$ be the set of states in which expert i strictly prefers the threat lottery to the decision maker's most preferred action,

$$\tilde{X}_i^p = \{x \in [0,1] \, : \, u_i(x,x) < \bar{u}_i(x,p)\},$$

where $\bar{u}_i(x,p)$ is expert i's expected payoff from the threat lottery p,

$$\bar{u}_i(x,p) = (1-p)u_i(x,0) + pu_i(x,1).$$

Hence, $\tilde{X}_1^p \cup \tilde{X}_2^p$ is the set of states where implementing the most preferred action is not incentive compatible.

Lemma 1 *For any state x in $\tilde{X}_1^p \cup \tilde{X}_2^p$, the incentive constraint of only one of the experts is violated, i.e., $\tilde{X}_1^p \cap \tilde{X}_2^p = \varnothing$.*

Proof. By assumption, the experts have opposing interests, i.e., $y_1^*(x) < x < y_2^*(x)$. If $p > x$, then expert 1 prefers action x to action $y = p$ and hence to the threat lottery. Otherwise, expert 2 prefers x to the threat lottery. Hence, at least one expert prefers x to the threat lottery. ■

An optimal decision rule stipulates to choose action $g(x)$ that is the "closest" point to x (from the perspective of the decision maker) subject to the incentive constraints for the experts. Since at every state $x \in \tilde{X}_i^p$ only expert i's incentive constraint is relevant, we obtain

$$g(x) \in \underset{y \, : \, u_i(x,y) \geq \bar{u}_i(x,p)}{\arg\max} \; u_0(x,y).$$

[17]To see why this is true, imagine that the experts are sufficiently biased such that one expert prefers action 0 and the other one prefers action 1 regardless of the state. In this environment, any decision rule with one expert will implement the same action in all states. Hence, with just one expert her information cannot affect the action of the decision maker.

4.4 Quadratic Preferences and Constant Biases

We can obtain stronger results if we impose additional structure on the preferences of the experts. Specifically, we make the assumption, which is standard in the literature, that the experts' preferences can be represented by a quadratic payoff function with a constant bias,

$$u_i(x, y) = -(y - (x + b_i))^2, \; i = 1, 2, \tag{7}$$

where $b_1 < 0 < b_2$. Assume also $u_0(x, y) = -(y - x)^2$ and $X = [0, 1]$.

In order to determine the set $\tilde{X}_i^p$ of states where expert i prefers threat lottery p to the most preferred action x for the decision maker, we solve the inequality $u_i(x, x) < \bar{u}_i(x, p)$. Using (7) we obtain

$$(1 - p)(x + b_i)^2 + p(1 - (x + b_i))^2 < b_i^2. \tag{8}$$

In order to state the result, the following definitions are in order. For any $p \in X$, let $D_i = b_i^2 - p(1 - p)$ and let

$$\underline{x}_i^p = p - b_i - \sqrt{D_i}, \;\; \overline{x}_i^p = p - b_i + \sqrt{D_i}. \tag{9}$$

In addition, for $|b_1|$ and $|b_2|$ below 1/2 define

$$\underline{p}^* = \frac{1 - \sqrt{1 - 4b_2^2}}{2}, \;\; \overline{p}^* = \frac{1 + \sqrt{1 - 4b_1^2}}{2}.$$

It is easy to verify that the solution of (8) is the interval $(\underline{x}_i^p, \overline{x}_i^p)$, and hence $\tilde{X}_i^p = (\underline{x}_i^p, \overline{x}_i^p) \cap [0, 1]$.

The next result describes the structure of an optimal decision rule.

Proposition 4 *Let (p, g) be an optimal constant-threat decision rule. Then,*

$$g(x) = \begin{cases} x - |b_1| + \sqrt{-\bar{u}_1(x, p)}, & \text{if } x \in \tilde{X}_1^p; \\ x + |b_2| - \sqrt{-\bar{u}_2(x, p)}, & \text{if } x \in \tilde{X}_2^p, \\ x, & \text{otherwise.} \end{cases}$$

Proof. If $x \notin \tilde{X}_1^p \cup \tilde{X}_2^p$, then the first best action is incentive compatible, $g(x) = x$.

Let $x \in \tilde{X}_1^p$ (the argument for $x \in \tilde{X}_2^p$ will be analogous). In an optimal decision rule the decision maker implements an action $g(x)$ that minimizes the distance to x, subject

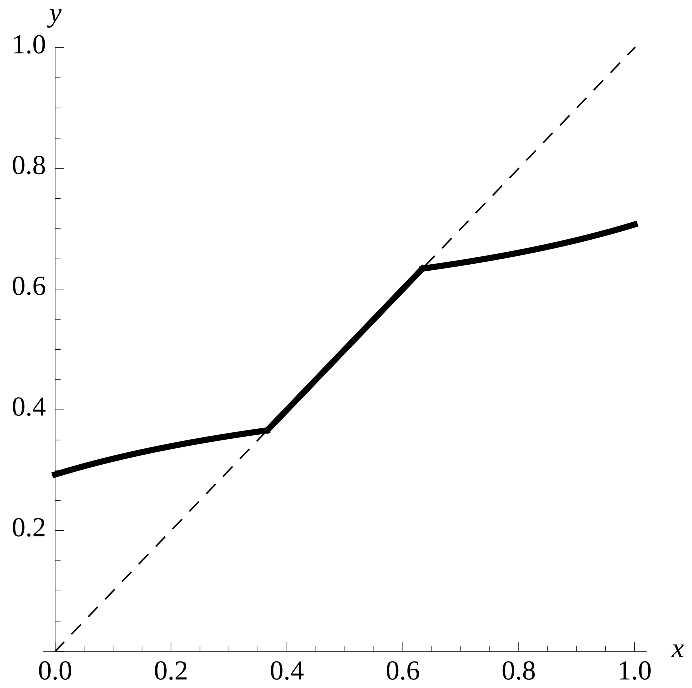

Figure 1: An optimal decision rule with quadratic preferences, $|b_1| = |b_2| = 1$, $p = 1/2$.

to the incentive constraint for expert 1, that is,

$$g(x) \in \arg\min_{y \in [0,1]} (y - x)^2$$

subject to

$$(y - (x + b_1))^2 \leq (1 - p)(x + b_1)^2 + p(1 - (x + b_1))^2.$$

Solving the above inequality for y we obtain

$$y \in [0, 1] \backslash \left(x + b_1 - \sqrt{-\bar{u}_1(x, p)}, x + b_1 + \sqrt{-\bar{u}_1(x, p)} \right).$$

Since $x \in \tilde{X}_1^p$, the above constraint must be binding. As $b_1 < 0$ by assumption, the closest action to x is $g(x) = x + b_1 + \sqrt{-\bar{u}_1(x, p)}$. It is straightforward to verify that in this case $g(x) \in \tilde{X}_1^p$. As $\tilde{X}_1^p \cap \tilde{X}_2^p = \varnothing$, the incentive constraint for expert 2 is satisfied as well. ■

If the absolute value of each of the biases is greater than 1/2, an optimal decision rule looks as follows (Fig. 1). Note that in this case $\tilde{X}_2^p = [0, \overline{x}_2^p)$ and $\tilde{X}_1^p = (\underline{x}_1^p, 1]$. For the "moderate states" in $[\overline{x}_2^p, \underline{x}_1^p]$, both experts prefer the decision maker's most preferred action to the threat lottery, and the first best outcome is achieved (the points along the 45° line on Fig. 1). For the "extreme left" states in $[0, \overline{x}_2^p)$, expert 2 strictly prefers the threat lottery to x, and hence the decision maker implements an action that is closer to expert 2's most preferred action. The distortion for the "extreme right" states is analogous.

The result in Proposition 4 allows us to transfer the problem of finding an optimal decision rule into a one dimensional optimization problem over value of the threat point:

$$\min_{p \in [0,1]} \int_0^1 (g_p(x) - x)^2 dF(x), \tag{10}$$

where g_p, with some abuse of notation, is given by Proposition 4.

In contrast to the first best decision rule, in second-best decision rules the optimal value of the threat point depends on the distribution of the state x. In general, there is no closed form solution for optimal threat points. Nevertheless, under additional assumptions, we obtain the following result.

Proposition 5 *Let the experts' biases be opposing and equal, $-b_1 = b_2 = b$, and distribution of x be symmetric, i.e., $F(1-x) = 1 - F(x)$, $x \in [0,1]$. Then there exists an optimal decision rule with $p = 1/2$.*

Proof. Note that under conditions of Proposition 5, $p = 1/2$ must be an extreme point of the expression in (10) due to full symmetry of the problem w.r.t. 1/2. What remains to prove is that $p = 1/2$ is the minimum of (10). The full proof is deferred to the Appendix. ■

Remark 2 Let the conditions of Proposition 5 hold. Then, the actions implemented on the equilibrium path in the optimal decision rule depend on the value of b. Nevertheless, the outcome of this decision rule can be implemented by the following indirect decision rule that is constant in b: The experts recommend an action. If their recommendations coincide, the action is taken. Otherwise, a threat lottery with $p = 1/2$ is implemented.

5 Continuity and Noise

5.1 Positive Results

The assumption that both experts are perfectly informed is common in the literature that studies cheap talk communication with two experts in payoff environments similar to the one in this paper. It has been made, for example, in Gilligan and Krehbiel (1989), Krishna and Morgan (2001a,b), Battaglini (2002), Levy and Razin (2007), Ambrus and

Takahashi (2008), and Li (2008, 2009).[18]

Nevertheless, this assumption is important because it ensures that any disagreement in experts' reports is evidence of out-of-equilibrium behavior. The optimal constant-threat decision rule might perform poorly if the experts' information about the state is noisy, since any non-identical albeit truthful reports are punished by a threat lottery.

This difficulty created by noise is not specific to the environment with commitment and, moreover, can have significant implications: e.g., cheap talk equilibria attaining the first best outcome are not robust to arbitrarily small amount of noise (Levy and Razin, 2007). The question of robustness to noise of these equilibria has been studied in Battaglini (2004) and Ambrus and Lu (2009). In a model with multiple experts, a multidimensional environment, and noisy signals, Battaglini (2004) shows that minimal commitment power is sufficient for the first best outcome to become feasible in the limit as the number of experts increases. Ambrus and Lu (2009) construct fully revealing equilibria that are robust to a small amount of noise in environments in which the state space is sufficiently large relative to the size of the experts' biases.

In this section, we establish two results. First, if the state space is finite, then either the optimal constant-threat decision rule is robust to small amount of noise or there exists a constant-threat decision rule that implements an outcome arbitrarily close to that in the optimal decision rule as noise vanishes. Second, for environments with continuous state spaces and noisy signals that can be modeled as in Aumann (1976) by elements of a partition of the state space, we can construct a sequence of what we call definite punishment decision rules that are incentive compatible and converge to the optimal decision rule for vanishing noise. These results are valuable as they underscore that commitment can restore continuity of optimal decision rules with respect to noise and suggest how one can construct decision rules that are robust to small noise.

Unfortunately, we do not have a continuity result for arbitrary noise structures if the state space is continuous or infinite. Therefore, the question whether the continuity result is specific to finite state space remains open. Nevertheless, since a general continuity result is not the main focus of this paper and it would be valuable to explore it in the environments more general than the one studied here, we leave the question for future research.

Model. We consider a model in which the experts' information is incomplete and not

[18] The experts are imperfectly informed in the models of Austen-Smith (1993), Wolinsky (2002), and Battaglini (2004). See also Li and Suen (2009) for a survey of work on decision making in committees; this literature often assumes that different members of the committee hold distinct pieces of information.

identical. We assume that, instead of observing the state, each expert observes an informative signal $s_i \in X$, $i = 1, 2$, about the state. The definition of the decision rule remains unchanged. A decision rule is incentive compatible if truthful reporting of signals is an equilibrium. The rest of the model is same as before.

The amount of noise is measured by Ky Fan metric

$$\delta_i = \inf_{\epsilon}\{\epsilon > 0 \,:\, \Pr(|s_i - x| > \epsilon | x) \leq \epsilon \text{ for all } x \in X\}.$$

We say that the amount of noise vanishes if $\delta_i \to 0, i = 1, 2$.

Finite X. In a finite state space, an optimal constant-threat decision rule in the environment without noise in which the experts' incentive constraints are satisfied with strict inequality remains incentive compatible for small amount of noise. Furthermore, in each state it implements the same action as in the environment without noise with probability uniformly converging to one as $\max\{\delta_1, \delta_2\} \to 0$. Finally, note that if the optimal constant-threat decision rule satisfies incentive constraints with strict inequality in each state, then this rule must be first best. We collect these observations in the following remark.

Remark 3 Let X be finite and assume that there exists the first best decision rule. If the incentive compatibility constraints are satisfied with strict inequality for each $x \in X$, then the first best decision rule is incentive compatible for a sufficiently small noise and its outcome is state-wise uniformly continuous in the amount of noise.

Consider now an environment in which the incentive constraints in the optimal decision rule hold with equality for some states. This rule may cease to be incentive compatible for an arbitrarily small amount of noise. Nevertheless, it can be shown using Lemma 1 that in the optimal constant-threat decision rule in each state there is a slack in the incentive constraint for at least one expert. Therefore, we can always distort the action in each state in favor of the expert whose incentive constraint is binding in order to make incentives strict for both experts. Clearly, this distortion can be made arbitrarily small. The distorted decision rule is incentive compatible for sufficiently small noise.

Remark 4 Let X be finite. Then, for any $\epsilon > 0$, there exists a constant-threat decision rule, which

(i) is incentive compatible for a sufficiently small amount of noise,

(ii) its outcome is state-wise uniformly continuous in the amount of noise, and

(iii) in the limit without noise, in each state it implements an action that is at most ϵ away from the action that is implemented in the optimal constant-threat decision rule.

Continuous X and partitional noise. Let $X = [0, 1]$. Following Aumann (1976), assume that an expert's information can be represented as an element of a partition of X.

For $i = 1, 2$, let $n_i = \{t_i^1, \ldots, t_i^{K-1}\}$, where $t_i^1 > 0$, $t_i^j < t_i^{j+1}$, $t_i^{K-1} < 1$, be an ordered sequence of real numbers that partition X into K intervals $\Pi_i = \{[0, t_i^1), [t_i^1, t_i^2), \ldots, [t_i^{K-1}, 1]\}$. Expert i observes the element of Π_i which contains the realized state.[19]

Let Π be the *meet* of Π_1 and Π_2, that is, the set of all non-empty intersections of the elements of these sets. Given a partitional information structure, a decision rule is a mapping $\mu : \Pi_1 \times \Pi_2 \to \mathcal{Y}$. The definition of incentive compatibility is standard and the revelation principle applies. A constant-threat decision rule in this environment is a triple $(\hat{\Pi}, g, p)$, where $\hat{\Pi} \subseteq \Pi$, $g : \hat{\Pi} \to Y$, and $p \in [0, 1]$. For every pair of reports $(\pi_1, \pi_2) \in \Pi_1 \times \Pi_2$, we say that experts "agree on $\pi_1 \cap \pi_2$" if $\pi_1 \cap \pi_2 \in \hat{\Pi}$, in which case the decision rule implements $g(\pi_1 \cap \pi_2) \in Y$; otherwise we say that experts "disagree" and the decision rule implements the (constant-threat) lottery in $\mathcal{Y}^*$ that assigns probability p on extreme action 1.

Remark 5 For every optimal decision rule, there exists an equivalent constant-threat decision rule.

Proof. The constant-threat result is verified following the proof of Proposition 1. ■

If $\Pi_1 = \Pi_2 = \Pi$, then the experts have identical information. The model is equivalent to that with discrete states and perfectly informed experts.

Remark 6 If the experts have the same information, $\Pi_1 = \Pi_2$, then the decision maker's payoff in the optimal decision rule converges to her maximal payoff in the environment with perfectly informed experts as the experts' information becomes more refined and $\max_{i,j} |t_i^{j+1} - t_i^j| \to 0$.

Proof. The proof is direct and therefore omitted. ■

For the remainder of the section, we focus on experts with different information, $\Pi_1 \neq \Pi_2$. In a truthtelling equilibrium in this environment, a lie cannot always be

[19] For convenience, we assume that the cardinality of partitions is the same across experts. The arguments in this section can be extended to accommodate different size partitions, but it will significantly increase notational burden.

detected with certainty. As a result, there are additional incentive constraints affecting the structure of optimal decision rules. To keep the presentation focused, we refrain from a full characterization of optimal decision rules for different information structures. Instead, we offer a construction of an incentive compatible decision rule that converges to the optimal decision rule as the difference in the experts' information vanishes.

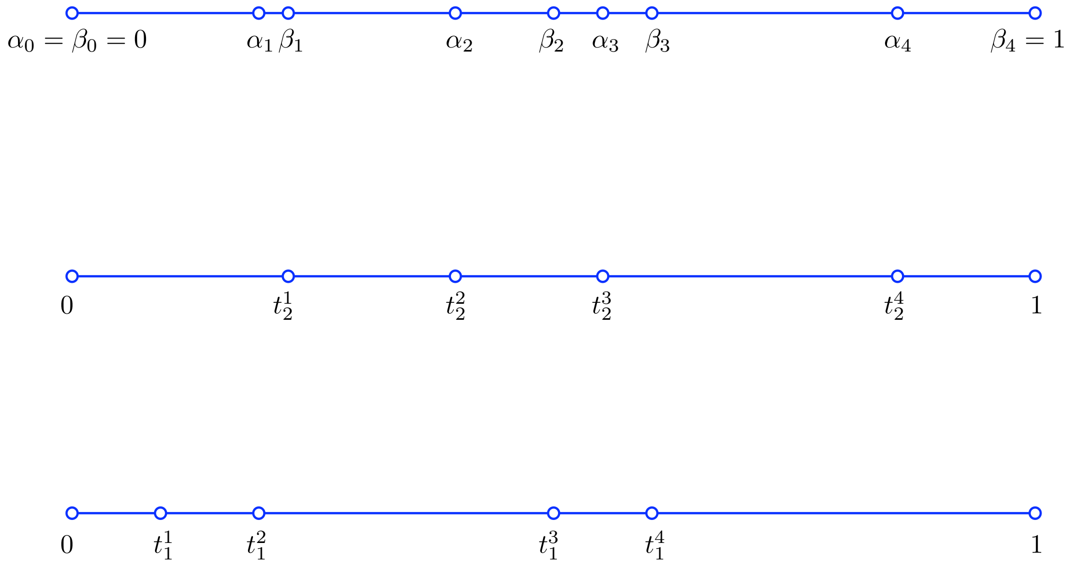

Figure 2: Construction of α_l, β_l sequence.

In order to describe the definite punishment decision rule, we first define a sequence α_l, β_l. Set $\alpha_0 = \beta_0 = 0$. For $l \geq 1$, we define α_l and β_l by induction. Let $\hat{t}_i = \min_k\{t_i^k \in n_i \cup \{1\} | t_i^k > \beta_{l-1}\}$. Without loss of generality assume that $\hat{t}_1 \leq \hat{t}_2$. Then, set $\beta_l = \hat{t}_2$ and $\alpha_l = \max_k\{t_1^k \in n_1 \cup \{1\} | t_1^k \leq \beta_l\}$. The induction stops when $\beta_l = 1$ and α_l is defined as above. Figure 2 depicts an example of a sequence α_l, β_l.

Fix a threat lottery p. A *definite punishment* decision rule with a threat lottery p is defined as follows. If experts reports $(\pi_1, \pi_2) \in \Pi_1 \times \Pi_2$ are inconsistent, $\pi_1 \cap \pi_2 = \varnothing$, or consistent but $\pi_1 \cap \pi_2 \subset [\alpha_l, \beta_l)$ for some l, then the decision rule implements the threat lottery. Otherwise, if the experts' reports are consistent and $\pi_1 \cap \pi_2 \subset [\beta_l, \alpha_{l+1})$ for some l, the decision rule implements the action that maximizes the expected payoff of the decision maker conditional on the state being in $[\beta_l, \alpha_{l+1})$ subject to the constraint that both experts prefer this action to the threat lottery given their information. By construction, if both experts report their elements of partitions truthfully, any deviation of an expert either does not affect the action or results in the threat lottery. Therefore, the decision rule is incentive compatible.

Proposition 6 *Let* $\max_j |t_1^j - t_2^j| \to 0$*. Then, there exists a sequence of definite punishment decision rules that converges to the optimal constant-threat decision rule in the environment in which the experts have identical information.*

Proof. To construct the sequence it is sufficient to pick the threat lottery in the optimal constant-threat decision rule in the limit environment where the experts have identical information. By definition of the definite punishment decision rule, $\max_j |t_1^j - t_2^j| \to 0$ entails $|\beta_l - \alpha_l| \to 0$ for all l, and hence the distribution of actions implemented on the equilibrium path converges to the distribution of actions in the limit environment. ■

Given the partitional information structure, the out-of-equilibrium lies cannot be detected with certainty. The idea behind our construction in Proposition 6 is to ensure that any deviation from truthful reporting, even if it cannot be detected with certainty, results in a punishment lottery. Clearly, this requires that the decision rule must punish the experts on the equilibrium path in some states. This is costly but as the experts' information becomes more aligned, the probability of punishment on the equilibrium path converges to zero.

Continuous X and replacement noise. A special feature of partitional information structure is that there exist pairs of reports that are incompatible with truthful reporting. If, by contrast, the signals of the experts' have full support in each state, all combinations of experts' reports are consistent with truthful reporting. The main result of this paper – the constant-threat result – will hold but become vacuous in such environments.

Nevertheless, constant-threat rules in which incentives are strict for both experts may be robust to some types of noise with full support: Ambrus and Lu (2009, Section 5.3) consider our constant-threat decision rules in an environment with the replacement noise, in which each expert observes a signal that is equal to the state with some probability and is uninformative with the complementary probability. They show that if the first best is implementable, then the corresponding constant-threat rule is incentive compatible in an environment with small replacement noise and continuous in the amount of noise.

5.2 Continuity of Optimal Decision Rule for $X = [0, 1]$.

Constant-threat decision rules are discontinuous in the experts' reports because an arbitrarily small disagreement between experts results in a punishment lottery. This discontinuity property *per se* may be a source of concern even in environments in which the experts have identical information if experts could make minor mistakes or there is exogenous noise added to their reports as in Blume et al. (2007). However, in these environments constant-threat decision rules could be modified and made continuous at the cost for the decision maker that becomes arbitrarily small as the noise vanishes.

Let $X = [0,1]$. Fix a constant-threat decision rule μ and assume that incentive constraints of both experts hold with strict inequality at all states. For any $\varepsilon > 0$ define

$$\eta(x_1, x_2) = \frac{1}{\varepsilon} \min \{|x_1 - x_2|, \varepsilon\},$$

Construct a new decision rule μ' as follows. If both reports coincide, there is no change. Otherwise, the decision rule is a compounded lottery that implements the threat lottery with probability $\eta(x_1, x_2)$ and the action corresponding to the average of the two reports with the complementary probability. That is,

$$\mu'(x_1, x_2) = \eta(x_1, x_2)\mu(x_1, x_2) + (1 - \eta(x_1, x_2))\mu\left(\frac{x_1 + x_2}{2}, \frac{x_1 + x_2}{2}\right).$$

It is straightforward to verify that this decision rule is continuous in (x_1, x_2) and incentive compatible for a small enough ε, where ε could be interpreted as a measure of mistakes or distortions in the experts' reports during information transmission.

6 Discussion

6.1 Full Implementation

The paper characterizes optimal decision rules that are partially implementable: the constructed decision rules permit equilibria other than truthtelling. Full implementation of the optimal decision rule outcome as a unique equilibrium outcome is impossible in our environment because a necessary condition of Maskin monotonicity (see Dutta and Sen, 1991) is not satisfied.

Nevertheless, if the experts have a lexicographic preference for telling the truth, as in Dutta and Sen (2010),[20] full implementation is feasible. For instance, consider an environment in which the experts know the state and assume that the first best outcome is partially implementable by a constant-threat decision rule with a threat lottery p. Then,

[20] That is, if an expert is asked to report her information, then among all reports that implement the same outcome as the truthful report does she strictly prefers the latter.

it can be fully implemented by the following mechanism:[21] Each expert reports a triple (x_i, y_i, z_i), interpreted as: "The true state is x_i, but I would like y_i to be implemented," and z_i is a positive integer. There are three contigencies to consider.

(i) If $y_1 = y_2 = y'$, then y' is implemented.

(ii) If $y_1 \neq y_2$ and $x_i = y_i$ for at least one of the experts, then the threat lottery p is implemented.

(iii) If $y_1 \neq y_2$ and $x_i \neq y_i$ for both experts, then the "integer game" is played: whoever has a greater integer z_i gets the requested action y_i.

Let true state be x. Then every Nash equilibrium leads to implementation of x and is characterized by $x_i = y_i = x$ and z_i is arbitrary, $i = 1, 2$. To see why this is a Nash equilibrium, observe that any unilateral deviation either does not affect the outcome, or results in the threat lottery p which is inferior to x for each expert. To see why there are no other Nash equilibria, first, observe that miscoordination, $y_1 \neq y_2$, cannot occur in equilibrium. If $x_i \neq y_i$ for at least one of the experts, a deviation that results in the "integer game" is possible, where the deviant chooses a large enough integer and gets the most preferred action. If $x_i = y_i$ for both experts, then one can deviate to $x_i = x$, which does not affect the outcome (the threat lottery is still implemented) but makes the deviant better off since she reports the truth. Finally, a coordination on an action y' different from true state x is not an equilibrium either, since one of the experts can deviate to $x_i = x$, which does not affect the outcome (action y' is still implemented) but makes the deviant better off since she reports the truth.

6.2 Similarly Biased Experts

In our paper, the experts are biased in different directions (c.f., (1)) and the decision maker does not have a choice over the experts available to him. This assumption could be applicable in, e.g., legislative politics where experts represent lobbyists or politicians with different ideologies, political economy of tariffs where experts represent different governments, or organizational economics where experts could be employees of different departments. If the experts are biased in the same direction, e.g., they always prefer an action higher than the decision maker's optimal action, then the decision maker can always implement the first best outcome: The decision rule that threatens the experts

[21] This mechanism is simpler than the one in Dutta and Sen (2010).

to implement $y = 0$ whenever they disagree will achieve the desired outcome. Although we do not provide formal analysis, one disadvantage of the environment with similarly biased experts might be that it is naturally prone to collusion: In the above decision rule, the experts can coordinate to bias their reports upward; this would benefit both experts. In the environment with experts biased in different directions, any agreement to bias reports would benefit one expert at the expense of the other one, and hence collusion is less compelling. These considerations can be important in practice making consulting similarly biased experts less attractive.

6.3 Optimal Decision Rules with One Expert

In this subsection, we comment on the difference between optimal decision rules in our model and in a model with one expert only. Without a second expert, the recommendations to the decision maker by the first expert remain unchecked. Therefore, the relevant incentive constraints are with respect to other actions that can be induced by the expert's reports rather than with respect to the outcome resulting from a disagreement with another expert. Consequently, optimal decision rules have a number of differences: There is bunching of implemented actions across states with one expert (Proposition 3 in Alonso and Matouschek (2008), and Proposition 1 in Kovac and Mylovanov (2009)) and no bunching with two experts (Proposition 4). With one expert, optimal decision rules do not implement first best actions because this cannot be made incentive compatible (see, e.g., Proposition 1 in Kovac and Mylovanov (2009)). This is not so with two experts: there is always a nonempty subset of states where the first best outcome is implemented (Propositions 2–3 in this paper). Furthermore, in the model with one expert the optimal decision rule implements the *expert's* most preferred action for a positive measure of states (Proposition 3 in Alonso and Matouschek (2008) and Proposition 1 in Kovac and Mylovanov (2009)). Again, this is not so with two experts (Propositions 2 and 4).

Appendix

Proof of Proposition 5. For $b \leq 1/2$ the statement holds trivially, since the first best decision rule can be constructed (see Proposition 2 and its proof).

Assume $b > 1/2$. Let (p, g_p) be a constant-threat decision rule, where g_p is described in Proposition 4. By an argument presented in Section 4.4, if both biases are greater than $1/2$, then $\tilde{X}_2^p = [0, \overline{x}_2^p)$ and $\tilde{X}_1^p = (\underline{x}_1^p, 1]$. We can now write the expected payoff of the decision maker as

$$
\begin{aligned}
v^{(p,g_p)} &\equiv \int_0^1 \left[-(g_p(x)-x)^2\right] dF(x) \\
&= -\int_0^{\overline{x}_2^p} (b-\sqrt{-\overline{u}_2(x,p)})^2 dF(x) - \int_{\underline{x}_1^p}^1 (b-\sqrt{-\overline{u}_1(x,p)})^2 dF(x).
\end{aligned}
$$

Recall that $\overline{u}_i(x,p) = -(1-p)(x+b_i)^2 - p(1-(x+b_i))^2$ and, by (9), $\overline{x}_2^p = p - b_2 + \sqrt{b_2^2 - p(1-p)}$ and $\underline{x}_1^p = p - b_1 - \sqrt{b_1^2 - p(1-p)}$. Using the symmetry assumption $b_1 = -b_2 = b$, we obtain that $\underline{x}_1^p = 1 - \overline{x}_2^{1-p}$ and $\overline{u}_1(x,p) = \overline{u}_2(1-x, 1-p)$, and hence

$$
v^{(p,g_p)} = -\int_0^{\overline{x}_2^p} (b-\sqrt{-\overline{u}_2(x,p)})^2 dF(x) - \int_{1-\overline{x}_2^{(1-p)}}^1 (b-\sqrt{-\overline{u}_2(1-x,1-p)})^2 dF(x).
$$

Next, using the assumption $F(x) = 1 - F(1-x)$ that entails $dF(x) = dF(1-x)$, after the substitution $x' = 1-x$ we obtain

$$
v^{(p,g_p)} = -\int_0^{\overline{x}_2^p} (b-\sqrt{-\overline{u}_2(x,p)})^2 dF(x) - \int_0^{\overline{x}_2^{(1-p)}} (b-\sqrt{-\overline{u}_2(x',1-p)})^2 dF(x').
$$

Let us now differentiate $v^{(p,g_p)}$ with respect to p. Observe that $\bar{u}_2(\bar{x}_2^p, p) = -b^2$, and

$$
\frac{\partial v^{(p,g_p)}}{\partial \bar{x}_2^p} = -(b-\sqrt{-\overline{u}_2(x,p)})^2\Big|_{x=\overline{x}_2^p} = -(b-b)^2 = 0.
$$

Note that $d\bar{x}_2^p/dp$ exists for all $p \in [0,1]$. Hence, the value of the expression $\frac{\partial v^{(p,g_p)}}{\partial \bar{x}_2^p} \cdot \frac{d\bar{x}_2^p}{dp}$ is well defined and equal to zero. An analogous statement holds for $\bar{x}_2^{1-p}$. Thus, derivatives w.r.t. bounds of integration are ignored, and after defining $h(x) = \frac{\partial \overline{u}_2(x,p)}{\partial p} = 2(x+b)-1$ we obtain

$$
\frac{\partial}{\partial p} v^{(p,g_p)} = -\int_0^{\overline{x}_2^p} \left[\frac{b}{\sqrt{-\overline{u}_2(x,p)}} - 1\right] h(x) dF(x) \tag{11}
$$

$$
+ \int_0^{\overline{x}_2^{1-p}} \left[\frac{b}{\sqrt{-\overline{u}_2(x,1-p)}} - 1\right] h(x) dF(x). \tag{12}
$$

It is straightforward to check that $\frac{\partial}{\partial p} v^{(p,g)}|_{p=\frac{1}{2}} = 0$. We now verify that $v^{(p,g_p)}$ is concave in p, thus $p = 1/2$ is a maximum. By $b > 1/2$, we have $h(x) = \frac{\partial \overline{u}_2(x,p)}{\partial p} = 2(x+b)-1 > 0$.

Hence, the expression

$$\left[\frac{b}{\sqrt{-\overline{u}_2(x,p)}} - 1\right] \cdot h(x)$$

is nondecreasing in p. Furthermore, since $\frac{b}{\sqrt{-\overline{u}_2(\overline{x}_2^p,p)}} = 1$, the above expression is non-negative for all $x \leq \overline{x}_2^p$. Thus, the right-hand side term in (11) is nonincreasing in p. A similar argument shows that the term in (12) is nonincreasing in p as well. It follows that $v^{(p,g_p)}$ is concave in p. ■

References

Alonso, R. and N. Matouschek (2008). Optimal delegation. *Review of Economic Studies 75*(1), 259–293.

Amador, M. and K. Bagwell (2010). On the optimality of tariff caps. Mimeo.

Ambrus, A. and G. Egorov (2009). Delegation and nonmonetary incentives. Mimeo.

Ambrus, A. and S.-E. Lu (2009). Robust fully revealing equilibria in multi-sender cheap talk. Mimeo.

Ambrus, A. and S. Takahashi (2008). The multi-sender cheap talk with restricted state spaces. *Theoretical Economics 3*, 1–27.

Armstrong, M. and J. Vickers (2010). A model of delegated project choice. *Econometrica 78*, 213–244.

Aumann, R. J. (1976). Agreeing to disagree. *Annals of Satistics 4*(6), 1236–1239.

Austen-Smith, D. (1993, January). Interested experts and policy advice: Multiple referrals under open rule. *Games and Economic Behavior 5*(1), 3–43.

Baron, D. P. (2000). Legislative organization with informational committees. *American Journal of Political Science 44*(3), 485–505.

Battaglini, M. (2002). Multiple referrals and multidimensional cheap talk. *Econometrica 70*, 1379–1401.

Battaglini, M. (2004). Policy advice with imperfectly informed experts. *Advances in Theoretical Economics 4*. Article 1.

Bester, H. and D. Krähmer (2008). Delegation and incentives. *RAND Journal of Economics 39*(3), 664–682.

Blume, A., O. J. Board, and K. Kawamura (2007). Noisy talk. *Theoretical Economics 2*(4), 395–440.

Crawford, V. P. and J. Sobel (1982). Strategic information transmission. *Econometrica 50*(6), 1431–51.

Dessein, W. (2002, October). Authority and communication in organizations. *Review of Economic Studies 69*(4), 811–38.

Dutta, B. and A. Sen (1991). A necessary and sufficient condition for two-person nash implementation. *Review of Economic Studies 58*(1), 121–128.

Dutta, B. and A. Sen (2010). Nash implementation with partially honest individuals. Mimeo.

Esö, P. and Y. Fong (2010). Wait and see: A theory of communication over time. Mimeo.

Gilligan, T. and K. Krehbiel (1989). Asymetric information and legislative rules with a heterogeneous committee. *American Journal of Political Science 33*, 459–490.

Goltsman, M., J. Hörner, G. Pavlov, and F. Squintani (2009). Mediation, arbitration, and renegotiation. *Journal of Economic Theory 144*, 1397–1420.

Holmström, B. (1977). On incentives and control in organizations. PhD thesis, Stanford University.

Holmström, B. (1984). On the theory of delegation. In M. Boyer and R. E. Kihlstrom (Eds.), *Bayesian Models in Economic Theory*, pp. 115–41. North-Holland.

Kalai, E. and R. W. Rosenthal (1976). Arbitration of two-party disputes under ignorance. *International Journal of Game Theory 7*(2), 65–72.

Koessler, F. and D. Martimort (2009). Optimal delegation with multi-dimensional decisions. Mimeo.

Kovac, E. and T. Mylovanov (2009). Stochastic mechanisms in settings without monetary transfers: The regular case. *Journal of Economic Theory 144*, 1373–1395.

Krishna, V. and J. Morgan (2001a). Asymmetric information and legislative rules: Some amendments. *American Political Science Review 95*(2), 435–52.

Krishna, V. and J. Morgan (2001b). A model of expertise. *Quarterly Journal of Economics 116*(2), 747–75.

Krishna, V. and J. Morgan (2008). Contracting for information under imperfect commitment. *RAND Journal of Economics 39*(4), 905–925.

Levy, G. and R. Razin (2007). On the limits of communication in multidimensional cheap talk: a comment. *Econometrica 75*, 885–893.

Li, H. and W. Li (2009). Optimal delegation with limited commitment. Mimeo.

Li, H. and W. Suen (2009). Viewpoint: Decision-making in committees. *Canadian Journal of Economics 42*, 359–392.

Li, M. (2008). Two (talking) heads are not better than one. *Economics Bulletin 3*(63), 1–8.

Li, M. (2009). Combining expert opinions. Mimeo.

Lim, W. (2009). Selling authority. Mimeo.

Martimort, D. and A. Semenov (2006). Continuity in mechanism design without transfers. *Economics Letters 93*(2), 182–189.

Martimort, D. and A. Semenov (2008). The informational effects of competition and collusion in legislative politics. *Journal of Public Economics 92*(7), 1541–1563.

Melumad, N. D. and T. Shibano (1991). Communication in settings with no transfers. *The RAND Journal of Economics 22*(2), 173–98.

Raith, M. (2008). Specific knowledge and performance measurement. *RAND Journal of Economics 39*(4), 1059–1079.

Wolinsky, A. (2002). Eliciting information from multiple experts. *Games and Economic Behavior 41*(1), 141–160.

Queen Mary
University of London

This working paper has been produced by
the School of Economics and Finance at
Queen Mary, University of London

School of Economics and Finance
Queen Mary, University of London
Mile End Road
London E1 4NS
Tel: +44 (0)20 7882 5096
Fax: +44 (0)20 8983 3580
Web: www.econ.qmul.ac.uk/papers/wp.htm