Eliciting Information from a Committee*

Andriy Zapechelnyuk[†]
Queen Mary, University of London

June 30, 2011

Abstract

The paper addresses the problem of eliciting truthful information from a committee of experts who collude in their information disclosure strategies. It is shown that under fairly general conditions full information disclosure is possible if and only if the induced outcome is Pareto undominated for the committee members.

Keywords: Information transmission; communication mechanism; cheap talk; experts; collusion; bargaining solution; closed rule

JEL classification numbers: D82, C78, D72

*The author thanks Tymofiy Mylovanov for fruitful conversations that inspired this paper, as well as Ronny Razin for useful comments.

[†]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

Specialized committees are routinely created by political institutions. The purpose of a committee is gathering information that would allow for better policy decisions. Yet its members may have biased interests relative to those of the public (represented by a policy maker or a legislative body), so they may be tempted to manipulate policy decisions by hiding or distorting information.

An important feature of a committee is that its members debate about which information to disclose, or which advice to give to the policy maker. As a result, any information disclosed by the committee is an agreement that equilibrates interests of its members and that, in general, could be different from the truth. This paper addresses the question of how, and under what conditions, truthful information can be elicited from committees.

A considerable body of literature addresses the problem of eliciting truthful information from informed and biased parties (experts), via cheap talk communication (Crawford and Sobel, 1982; Gilligan and Krehbiel, 1989; Austen-Smith, 1990; Krishna and Morgan, 2001b; Battaglini, 2002; Ambrus and Takahashi, 2008), through commitment to certain legislative rules (Baron, 2000; Krishna and Morgan, 2001a; Mylovanov and Zapechelnyuk, 2011) or by (constrained) delegation of decisions to the informed party (Holmström, 1977, 1984; Melumad and Shibano, 1991; Dessein, 2002; Alonso and Matouschek, 2008; Mylovanov, 2008). The underlying assumption in this literature is that experts communicate information independently or privately. Since existence of multiple equilibria is typical in this type of models, coordination of experts on a particular, truthtelling equilibrium (when it exists) need not remain a reasonable prediction if experts can collude on a different outcome.

An illustrative example is consulting with a committee whose members have interests biased in similar directions. Suppose that a policy maker considers cutting social benefits to a certain group of population and consults two experts whose preferences are biased towards higher social benefits. Even when proceeding with that policy is socially optimal, the experts are not likely to give that recommendation, as the advice to do nothing is Pareto superior and easy to coordinate upon. For another example, consider two experts whose interests are biased in the opposite directions with asymmetric magnitude. As disclosing the truth would favor one of them relative to the other, the experts might agree instead to communicate different information that would balance their interests better.

Since experts who collude in their actions have more freedom in manipulating the policy maker's decisions (as compared to those who act independently), eliciting the truth is harder in this case. In fact, if the policy maker trusts the committee whenever its members unanimously agree, then the committee can achieve *any* outcome. So the only way to elicit information from the committee is by using the conflict of interests between its members. The central issue that we investigate here is how it can be done, and what are the necessary and sufficient conditions for that.

We present a model where a policy maker, who acts on behalf of the society, has to choose a (multidimensional) policy whose exact effect is uncertain. In order to make a correct policy choice, the policy maker appoints a committee of n experts, who find out the relationship between policies and their outcomes, and then give advice to the policy maker. The committee members have biased interests relative to the policy maker and relative to one another. Before giving any advice, the committee engages in bargaining over possible outcomes that their advice could induce through influencing the policy maker's decision.

We do not consider any explicit bargaining game, nor we specify how a committee reaches an agreement. Instead, we are interested in general results that hold under a wide range of specifications of bargaining procedures.

The basic insight of the paper is that under a wide range of circumstances there is no loss for a policy maker for dealing with committees, as compared to independent, non-communicating experts. We show that the policy maker can elicit truthful information if and only if the resulting outcome is Pareto undominated for the committee. That is to say, whenever truthtelling is not Pareto inferior to any other outcome, thus being a viable prediction in the case of independent experts, it remains achievable in the case of colluding experts.

Let us briefly describe the results. We begin with the assumption that the policy maker can commit to a rule that chooses a policy as a function of the experts' proposals. The first result (Theorem 1) states that for every bargaining solution that satisfies certain requirements, there exists a rule that implements the desired outcome if and only if it is Pareto undominated for the committee. Our requirements on bargaining solution are very basic. We demand the solution to be *individually rational* (i.e., no expert can guarantee a better payoff by opting out from bargaining), *Pareto efficient*, and *continuous*. The proof is constructive: we show that every Pareto undominated outcome can be implemented by a *closed rule*. Under a closed rule, if the committee unanimously agrees on some policy, then it is adopted with certainty. However, if the committee disagrees and offers a menu of policies, each of which is supported by a subset of committee members, then the policy maker chooses among the policies on the table, with probability of choosing some policy increasing in the number and "quality" of its supporters. By choosing the probabilities of adopting experts' proposals, the decision maker influences the experts relative "bargaining strength", and hence manipulates the bargaining solution to achieve the desired outcome.

Next, we show that a crucial assumption for the above result is the ability of the policy maker to randomize over policies. Theorem 2 asserts that, under an additional mild assumption on a bargaining solution, no deterministic rule can implement the socially optimal outcome.

Lastly, we relax the assumption of commitment to a rule, and consider instead cheap talk communication between the policy maker and the committee. In other words, after receiving proposals from the committee, the policy maker is free to choose any policy. In equilibrium, the chosen policy must maximize the policy maker's payoff given his posterior beliefs formed by the received information. On the one hand, randomizations over policies cannot be rationalized by any beliefs (as payoff functions are assumed to be concave). By the previous result it is immediately clear that achievement of the socially optimal outcome is impossible. On the other hand, if we assume that the policy maker is constrained to use *closed rules*, i.e., to choose a policy *only* among those proposed by the experts, then the positive result is restored. The restriction to closed rules, a type of constraint that in practice could be imposed by institutional regulations, is an advantage for the policy maker, as now lotteries among the proposed set of policies *can be rationalized*, if the policy maker believes to be indifferent among them. We prove that any closed rule can be an equilibrium, once the decision maker's beliefs are chosen appropriately (Theorem 3).

The results presented in this paper are qualitatively different from those in the cheap talk literature. First of all, in cheap talk models the number of experts is at most two. This is because for three and more experts the solution of the truthful communication problem is trivial: let every expert report the truth; if the majority of the reports coincide, ignore any deviating minority. It is clear that no expert can benefit by unilateral deviations, as the majority would still report the truth. However, if the experts can collude, as considered in this paper, the problem looses its triviality.

As concerns two experts in a unidimensional state space, the major insight from the literature (Gilligan and Krehbiel, 1989; Krishna and Morgan, 2001b,a; Battaglini, 2002) is that full information revelation is always possible (provided the biases are small relative to the size of the state space). In the case of experts whose interests are biased in the opposite directions, the equilibrium construction involves extreme punishments, adopting policies that are Pareto inferior to telling the truth for both experts. In contrast, in our model extreme punishments are not feasible. If one expert opts out from bargaining and chooses, say, a policy that is most preferred for her, the other expert always has an option to concede and agree with the deviant. So, the set of policies that may be adopted after a disagreement is limited. We achieve full information revelation by a different means, deterring disagreements by randomizations over proposed policies.

In the case of like-biased experts, construction of truthtelling equilibrium under cheap talk is particularly easy. For every pair of non-identical policy proposals, let the policy maker believe that the truthful one is the lower of the two (with respect to the order of the experts' preferences), thus providing incentives to report truth for both. To the contrary, in our model like-biased experts will never agree to reveal the truth, as it is Pareto dominated.

The above insights carry over to models with two experts and multidimensional state space (Battaglini, 2002; Ambrus and Takahashi, 2008). However, in our model, two experts are typically insufficient for to obtain truthful information disclosure, since, in general, the experts can find a direction of Pareto improvement relative to the socially optimal outcome. The only exception is the case of exactly opposite biases, where the socially optimal outcome is Pareto undominated.

The paper is closely related to Martimort and Semenov (2008) (henceforth, MS) who study information effects of collusive behavior among informed parties. Like in our model, in MS the informed parties are allowed to collude in their information disclosure actions. The bargaining problem of a coalition is resolved by a special type of bargaining solution: a *side mechanism* designed by a mediator who organizes collusion between the parties in exchange for side payments, and whose objective is to maximize his revenue.[1] Though the informed parties are assumed to hold different pieces of information relevant for the legislature (whereas in our model all experts are identically informed), MS consider the type of coalition where all information is observable to its members (called *strong coalition*), in which case the situation identical to our model. The model of MS is limited to unidimensional policies, two informed parties, and special forms of utility functions. Our model extends MS along all these dimensions, though we do not address the same range of questions as MS do.

Another closely related paper is Mylovanov and Zapechelnyuk (2011) who characterize optimal stochastic rules for an uninformed principal in a model with a unidimensional state space and two experts. Closed rules in this paper are similar to *final offer arbitration rules* in Mylovanov and Zapechelnyuk (2011) in that experts' disagreements are punished by randomized policies. The crucial difference is that Mylovanov and Zapechelnyuk (2011) do not consider collusion of experts. Punishment of disagreements is simply a tool to provide incentives to the experts to agree on the truth. Since the experts' payoff functions are concave, randomized policies constitute better punishments than deterministic ones. In contrast, in this paper the policy maker *does not punish* experts. Uncertainty about a policy to be selected is the policy maker's instrument for manipulation of the committee's agreement. Assigning a higher probability on adopting some expert's policy gives that expert more "bargaining strength" and allows her to negotiate the outcome closer to what she favors. The policy maker calibrates the probabilities to induce the committee to agree on a desired outcome.

The paper is also related to the non-cooperative theory of n-player bargaining that goes back to Rubinstein (1982) and has been further developed by Binmore (1987), Krishna and Serrano (1996), with particular focus on bargaining in legislative bodies in Baron and Ferejohn (1989), Jackson and Moselle (2002), and Banks and Duggan (2000, 2006). This paper is different in two aspects. First, the above literature assumes complete information, whereas we assume asymmetric information between the policy maker and the committee. Second, in this paper, instead of specifying a non-cooperative procedure for bargaining (as the above literature does), we simply consider an arbitrary bargaining solution that satisfies a few basic requirements. These requirements are weak and satisfied by all well known bargaining solution concepts, such as Nash bargaining solution (more generally, asymmetric Nash solution) (Nash, 1950;

[1] Cf. Laffont and Martimort (2000), where a side mechanism is a benevolent entity that maximizes the sum of the parties' utilities.

Kalai, 1977a), Kalai-Smorodinsky solution (Kalai and Smorodinsky, 1975), and the egalitarian solution (more generally, the proportional solution) (Kalai, 1977b).

The paper is organized as follows. Section 2 introduces the model. In Section 3 we state our main result, the necessary and sufficient condition for implementation of the socially optimal outcome. Section 4 considers the restriction of the policy maker's rules to deterministic ones. Section 5 addresses the question of truthful information transmission under cheap talk. A few additional issues are discussed in Section 6.

2 The Model

2.1 Preliminaries. There is a policy maker and n experts. The policy maker would like to implement outcome x_0 (normalized to be the origin) in set of outcomes $X = \mathbb{R}^{\mathbf{d}}$ and he chooses a policy from $Y = \mathbb{R}^{\mathbf{d}}$. When a policy $y \in Y$ is chosen, an outcome is $x = y + \theta$, where $\theta \in \Theta \subset \mathbb{R}^{\mathbf{d}}$ is an unobserved state, a random draw represented by some continuous distribution function with support Θ and zero expected value.[2]

Before making any decision, the policy maker appoints a committee of n experts who observe state θ and make proposals about a policy to choose. That is, each expert i proposes policy $m_i = m_i(\theta) \in Y$.

The policy maker follows a rule $\mu : Y^n \to \Delta(Y)$, a measurable function whose image on $\Delta(Y)$ is closed[3], that stipulates for every tuple of the experts' proposals $m = (m_1, \ldots, m_n)$ to choose a policy (or a lottery over policies) $\mu(m)$.[4]

The payoff of the policy maker is given by $u_0(x)$ where $x = y + \theta$ is the implemented outcome. The payoff of each expert $i = 1, \ldots, n$ is given by $u_i(x - b_i)$, where $b_i \in \mathbb{R}^d$ is the most preferred outcome for i, in other words, expert i's *bias*. We assume that $u_i(\cdot)$ is differentiable, strictly concave and attains the maximum at 0 for every $i = 0, 1, \ldots, n$.

The following notations are in order. For any vector $a = (a_1, \ldots, a_n)$ we write a_S for $(a_i)_{i \in S}$. For a function $h : A \to B$ and a subset $A' \subset A$ we write $h(A')$ for set $\{h(a) : a \in A'\}$.

[2] State space Θ is any measurable subset of $\mathbb{R}^{\mathbf{d}}$, in particular, it may be bounded or finite. Policy set Y need not be unbounded for our results to hold; it is sufficient to require that Y contain policies where the experts' most preferred outcomes are achieved. The outcome set is thus $X = \{y + \theta : y \in Y, \theta \in \Theta\}$.

[3] This assumption is needed for some notions, in particular, Pareto efficiency, to be well defined.

[4] In Section 5 we relax the assumption of commitment to a rule and consider cheap talk communication.

2.2. The Bargaining Problem. The committee members, $N = \{1, \dots, n\}$, before releasing a tuple of proposals to the policy maker, bargain about what they should propose.

A bargaining problem for the committee is defined as a map that associates every tuple of proposals with a payoff vector to all its members. For a given rule μ and state θ, every tuple of proposals $m = (m_1, \dots, m_n)$ yields to each expert i expected payoff

$$F_i^{\mu,\theta}(m) = E_{\mu(m)}[u_i(y + \theta - b_i)], \tag{1}$$

where $E_{\mu(m)}[\cdot]$ denotes the expectation with respect to distribution $\mu(m)$ over policies. Let $F^{\mu,\theta}(m) = (F_1^{\mu,\theta}(m) \dots, F_n^{\mu,\theta}(m))$.

Definition 1. For every rule μ and state θ, function $F^{\mu,\theta} : Y^n \to \mathbb{R}^n$ defined by (1) is called the *bargaining problem for n experts.*

Denote by $\mathcal{F}_k$ the set of all k-player bargaining problems, and let $\mathcal{F}^n = \bigcup_{k=1}^n \mathcal{F}_k$.

Definition 2. A *bargaining solution on $\mathcal{F}^n$* is a mapping ϕ that associates with every bargaining problem F in $\mathcal{F}_k$ a proposal tuple $\phi(F)$ in Y^k.

We allow each expert i to disagree with the other members of the committee and to *opt out* from bargaining. In this event, i sends to the policy maker proposal $\bar{m}_i$ that does not depend on what the rest of the experts agree upon.[5] If a subset S of experts has (simultaneously) opted out with tuple of proposals $\bar{m}_S$, then we assume that the rest of the committee, $N \backslash S$, bargain among themselves. The bargaining problem for $N \backslash S$ is simply a restriction of $F^{\mu,\eta}$ to the coordinates of $N \backslash S$ with a fixed tuple of opt-out proposals $\bar{m}_S$ for the rest of the coordinates, and it is denoted by $F^{\mu,\eta}_{N \backslash S}(\cdot|\bar{m}_S)$.

2.3. The Axioms. We do not specify a bargaining solution. Instead, we consider any solution that satisfies a few basic axioms described below.

An *opt-out payoff* $d_i(F)$ of each expert i is defined as the payoff that i can secure in bargaining problem F, if she anticipates that the rest of the experts, $N \backslash \{i\}$, induce the bargaining outcome for their $(n-1)$-player problem,

$$d_i(F) = \max_{m_i \in Y} F_i\left(m_i, \phi(F_{-i}(\cdot|m_i))\right).$$

[5] We discuss other types of deviations in Section 6.1.

Let $d(F) = (d_1(F), \ldots, d_n(F))$. The opt-out payoff vector $d(F)$ is an important characteristic of the bargaining problem that identifies abilities of experts to enforce certain subsets of outcomes in X, and hence their relative "bargaining strength".

We say that a proposal $m \in Y^n$ is *individually rational* if the resulting payoff to each player is at least as high as her opt-out payoff, $F(m) \geq d(F)$.

Individual Rationality (IR) *For every bargaining problem $F \in \mathcal{F}^n$, $\phi(F)$ is individually rational.*

Proposal tuple $m \in Y^n$ is said to be *Pareto efficient* if there does not exist $m' \in Y^n$, $m' \neq m$, where all experts in N are better off and some are strictly better off.

Pareto Efficiency (PE) *For every bargaining problem $F \in \mathcal{F}^n$, $\phi(F)$ is Pareto efficient.*

The last axiom is technical.

Continuity (C) *ϕ is continuous.*

2.4. The Relation to Nash Bargaining Problem. Let us now relate an n-player bargaining problem F to *Nash bargaining problem* (Nash, 1950). An n-player Nash bargaining problem is a pair (Z, d) where $Z \subset \mathbb{R}^n$ is a set of all payoff vectors achieved under *full agreement* and $d \in \mathbb{R}^n$ is a payoff vector obtained under *disagreement*. Partial agreements, where a subset of players agrees on some outcome, are not allowed – in contrast to our model. Our bargaining problem is a richer object than Nash bargaining problem, so the former can always be translated into the latter (with some loss of information), but not vice versa.

Let $F \in \mathcal{F}_n$. Define the associated Nash bargaining problem (Z, d) as follows. The set of attainable payoffs Z is given by

$$Z = \{F(m) : m \in Y^n\}.$$

A disagreement payoff d_i of each expert i is the opt-out her payoff, $d_i = d_i(F)$.

Axioms IR, PE and C are now trivially translated into Nash bargaining problem setup, where they are satisfied by all well known bargaining solutions, such as Nash bargaining solution (more generally, asymmetric Nash solution) (Nash, 1950; Kalai, 1977a), Kalai-

Smorodinsky solution (Kalai and Smorodinsky, 1975), and the egalitarian solution (more generally, the proportional solution) (Kalai, 1977b).

3 Implementation of the Socially Optimal Outcome

We say that for a given bargaining solution ϕ, *rule* μ *implements outcome* x if in every state θ bargaining outcome $\phi(F^{\mu,\theta})$ leads to policy outcome x,

$$\mu(\phi(F^{\mu,\theta})) + \theta = x \quad \text{for all } \theta \in \Theta.$$

Outcome x is said to be *implementable* if there exists a rule that implements it.

We now establish necessary and sufficient conditions under which the socially optimal outcome, $x_0 = 0$, is implementable, and construct a rule that implements it.

Let

$$U_S(x) = (u_i(x - b_i))_{i \in S}\,.$$

We write $U_S(x') \gneqq U_S(x)$ if $U_S(x')$ is greater than $U_S(x)$ in every coordinate and strictly greater in some.

For every subset of experts $S \subset N$ denote by P_S the *Pareto set* for S, the set of all outcomes x in X such that there does not exist $x' \neq x$ with $U_S(x') \gneqq U_S(x)$. Note that Pareto sets are defined directly on the outcome space and are independent of rule μ or state θ.

Remark 1. Pareto sets have a clear geometric interpretation if we assume that the payoff of each expert is the negative squared Euclidean distance between the realized and most preferred outcomes. That is, for every $i \in N$, let $u_i(x) = -||x - b_i||^2$. Then P_S is the convex hull of the biases of the experts in S.

Any outcome x that is in the Pareto set for the entire committee, $x \in P_N$, is called *Pareto undominated.*

Theorem 1. *For every bargaining solution* ϕ *that satisfies Axioms IR, PE and C, the socially optimal outcome is implementable if and only if it is Pareto undominated.*

The proof proceeds as follows. First, we prove necessity by showing that outcomes outside P_N cannot be implemented by any rule, whenever the bargaining solution satisfies Pareto Efficiency (Proposition 1). Second, to prove sufficiency, we construct a family of rules that implement any outcome in P_N for any bargaining solution that satisfies Axioms IR, PE and C (Proposition 2), thus demonstrating existence of a rule that implements the socially optimal outcome whenever it is in P_N.

Proposition 1 (**Necessity**). *For every bargaining solution ϕ that satisfies PE, outcome x is implementable only if it is Pareto undominated.*

Before proving Proposition 1, let us introduce the notion of *direct rule.*

Definition 3. A rule μ is called *direct* if it satisfies $\mu(y, y, \ldots, y) = y$ for all $y \in Y$.

Direct rules have an important property. Since $\mu(y, y, \ldots, y) = y$ for all $y \in Y$, all outcomes in X can be induced, hence the set of attainable payoffs, Z, in the bargaining problem is equal to

$$Z \equiv F^{\mu,\theta}(Y^n) = \{U_N(x) \,:\, x \in X\}$$

and it does not depend on μ (as long as it is direct) or θ.

By the revelation principle, every rule that implements the socially optimal outcome can be mapped into a direct rule. Thus, to prove that outcomes outside of P_N are not implementable, it it sufficient to show that they are not implementable by direct rules.

Proof of Proposition 1. Consider a direct rule μ. Since the set of attainable payoffs Z is equal to $\{U_N(x) \,:\, x \in X\}$, an agreement on policy y is Pareto efficient if and only if induced outcome $x = y + \theta$ is in P_N. It follows that every bargaining solution that satisfies Pareto Efficiency must lead to selection of an outcome in P_N. ■

To prove sufficiency, let us describe a special class of rules, called *closed rules*, that lead to implementation of any point in P_N.

A closing rule is informally described as follows. The policy maker consults with a committee about what policy to select. If the committee unanimously recommends some policy, then that policy is adopted. If the committee disagrees and offers a menu of policies, each of which is supported by a subset of committee members, then the policy maker chooses among the policies on the table, with probability of choosing some policy increasing in the number and quality of its supporters.

Formally, denote by Δ the set of distributions over set $N = \{1, \ldots, n\}$. For every $\eta = (\eta_1, \ldots, \eta_n) \in \Delta$ and every expert $i \in N$, η_i is the probability assigned to i.

Definition 4. A rule μ is called *closed* if there exists $\eta \in \Delta$ such that for all $m = (m_1, \ldots, m_n) \in Y^n$, $\mu(m)$ chooses policy m_i with probability η_i.

In other words, a rule is called *closed* if it selects a policy y among those proposed by the experts, $y \in \{m_1, m_2, \ldots, m_n\}$, and does so randomly, according to a fixed distribution that does not depend on the proposals.

Every closed rule is fully described by distribution $\eta \in \Delta$ and denoted by μ^η.

It is worth noting that any closed rule is direct, since agreement of the committee on any policy leads to selection of that policy with certainty by the policy maker. Also, notice that any bargaining outcome must be a full agreement among the experts. Indeed, every disagreement results in a lottery, which is inferior for all experts to some deterministic policy choice.

How does this class of rules implement any point in the Pareto set? The channel through which the policy maker manipulates the bargaining outcome is the experts' opt-out payoffs (it cannot be influenced through restriction on the set of outcomes, since under direct rules the committee can enforce any outcome). An increase in probability η_i of adoption of expert i's proposal increases the opt-out payoff of i, and thus her relative "bargaining strength", in turn leading to shifting the bargaining outcome closer to i's most preferred policy. Ultimately, if $\eta_i = 1$ for some i, then expert i is the only one whose opinion matters. In that case IR and PE axioms require the bargaining outcome to be the most preferred policy for i (a "vertex" of the Pareto set). We prove that by an appropriate choice of η the policy maker can manipulate the bargaining solution to induce any outcome in the Pareto set for the committee.

Proposition 2 (**Sufficiency**). *For every bargaining solution ϕ that satisfies Axioms IR, PE and C, and for every $x \in P_N$ there exists $\eta \in \Delta$ such that closed rule μ^η implements x.*

Proof of Proposition 2. Denote by $F^{\eta,\theta}$ the bargaining problem defined by closed rule μ^η and state θ. As noted before, since we consider direct rules, in any Pareto efficient solution experts always agree on some y, i.e., $\phi(F^{\eta,\theta}) = (y, \ldots, y)$. With abuse of notation, we will write $\phi(F^{\eta,\theta}) = y$.

We would like to establish that for every $x \in P_N$ one can find η such that the bargaining solution of $F^{\eta,\theta}$ is exactly the policy that leads to outcome x for every state θ, i.e.,

$$\phi(F^{\eta,\theta}) + \theta = x \quad \text{for all } \theta \in \Theta. \tag{2}$$

Define $\xi : \Delta \to X$ by $\xi(\eta) = \phi(F^{\eta,0})$, which is continuous by Continuity axiom. Observe that for every $(m_1, \ldots, m_n) \in Y^n$ and every θ the following holds,

$$F^{\eta,\theta}(m_1, \ldots, m_n) = F^{\eta,0}(m_1 + \theta, \ldots, m_n + \theta) - \theta.$$

Hence, we have $\phi(F^{\eta(x),\theta}) + \theta = \xi(\eta)$ for all $\theta \in \Theta$. In these notations, (2) can be in short written as $P_N \subset \xi(\Delta)$.

For every $S \subset N$ denote by Δ_S the set of distributions in Δ with support on the coordinates in S only,

$$\Delta_S = \{\eta \in \Delta : \eta_i = 0 \text{ for all } i \notin S\}.$$

We proceed by induction in the cardinality of $S \subset N$. First, let $S = \{i\}$ for some $i \in N$. Observe that $P_{\{i\}} = \{b_i\}$. Consider $\eta \in \Delta_{\{i\}}$, i.e., $\eta_i = 1$ and $\eta_{-i} = 0$. We now show that $\xi(\eta) = b_i$.

Under rule μ^η, for every $j \in N$,

$$\begin{aligned} F_j^{\eta,\theta}(m_i, m_{-i}) &= \eta_i u_j(m_i - b_j + \theta) + \sum_{k \neq i} \eta_k u_j(m_k - b_j + \theta) && (3) \\ &= u_j(m_i - b_j + \theta). && (\text{since } \eta \in \Delta_{\{i\}}) \end{aligned}$$

Hence, $F^{\eta,\theta}$ is independent of m_{-i}. Thus i can attain any outcome by herself, in particular her most preferred outcome $x = b_i$, so her opt-out payoff is equal to the maximum of her utility, $d_i(F^{\eta,\theta}) = u_i(0)$. Thus $y = b_i - \theta$ is the unique policy that satisfies Individual Rationality. Consequently, $\phi(F^{\eta,\theta}) = b_i - \theta$, implementing outcome $x = b_i$. Thus $P_{\{i\}} = \{b_i\} \subset \xi(\Delta_{\{i\}})$.

Next, let $S \subset N$, $|S| \geq 2$. For each $i \in S$, suppose that $P_{S \setminus \{i\}} \subset \xi(\Delta_{S \setminus \{i\}})$. We will now prove that $P_S \subset \xi(\Delta_S)$.

By induction assumption, $P_{S\setminus\{i\}} \subset \xi(\Delta_{S\setminus\{i\}})$, hence

$$\bigcup_{i\in S} P_{S\setminus\{i\}} \subset \bigcup_{i\in S} \xi(\Delta_{S\setminus\{i\}}) = \xi\left(\bigcup_{i\in S} \Delta_{S\setminus\{i\}}\right) \subset \xi(\Delta_S).$$

Denote by ∂P_S the boundary of set P_S. Observe that $\partial P_S \subset \bigcup_{i\in S} P_{S\setminus\{i\}}$. Thus we obtain $\partial P_S \subset \xi(\Delta_S)$.

Now, the homotopy group of set $\xi(\Delta_S)$ is trivial (i.e., it has no "holes" inside), as we can construct a retraction $r : X \times [0,1] \to X$ that contracts $\xi(\Delta_S)$ to a point. Fix $i \in S$. For every $x \in \xi(\Delta_S)$ define

$$r(x,t) = \xi(tb_i + (1-t)x, \theta).$$

In particular, $r(\partial P_S, \cdot)$ contracts boundary of P_S to a point $b_i \in P_S$. Consequently, $P_S \subset \xi(\Delta_S)$. ■

4 Deterministic Rules

A rule μ is said to be *deterministic* if it chooses a non-random policy for every tuple of proposals, $\mu(m) \in Y$ for all $m \in Y^n$.

Let us now discuss why it is very hard, if not impossible, to implement the socially optimal outcome with deterministic rules. Before making a general statement, let us consider a few examples.

For the sake of simplicity, consider a committee of two experts, $N = \{1,2\}$, with exactly opposite, so that the socially optimal outcome $x_0 = 0$ is in the committee's Pareto set.

Consider a *naive rule* that implements a weighted average of the proposals in (m_1, m_2) according to some fixed weights $(w, 1-w)$. Then, if expert 1 opts out to $\bar{m}_1$, expert 2 can enforce her most preferred outcome $x = b_2$ by proposing $m_2^*(m_1) = (b_2 - \theta - w\bar{m}_1)/(1-w)$. Thus the opt-out payoff vector $d(F^{\mu,\theta})$ is independent of chosen weights $(w, 1-w)$, so the bargaining outcome cannot be manipulated.

Next, consider an *extreme punishment rule* rule used in cheap-talk literature (e.g., Battaglini, 2002; Ambrus and Takahashi, 2008) for eliciting truthful information from experts who submit proposals simultaneously and independently. According to this rule, if $m_1 = m_2 = y$, then policy y is adopted; any disagreement ($m_1 \neq m_2$) is punished by policy $\mu(m_1, m_2)$ that makes each expert worse off relative to the agreement, provided the other expert's proposal is socially optimal (i.e., it is the policy that implements x_0). Formally, assuming that i's proposal is socially optimal ($m_i = x_0 - \theta$), consider the set of policies $\hat{Y}(m_i)$ where expert $j \neq i$ is worse off relative to the desired outcome x_0,

$$\hat{Y}_j(m_i) = \{y \in Y : u_j(y + \theta - b_j) \leq u_j(x_0 - b_j)\},$$

where θ is set to $x_0 - m_i$. Then choose $\mu(m_1, m_2) \in \hat{Y}_1(m_2) \cap \hat{Y}_2(m_1)$.

Let us show that an extreme punishment rule that implements the socially optimal outcome does not exist. Suppose otherwise. Under this rule, if expert i opts out to $\bar{m}_i = b_j - \theta$ (the most preferred policy for $j \neq i$), expert j's optimal reaction is to agree. Hence any outcome induced after i opted out optimally must be in Pareto set $P_{\{1,2\}}$. But $P_{\{1,2\}}$ is totally ordered with respect to payoffs of the experts, so that an outcome which is better for one expert is always worse for the other. It is not possible to punish both experts simultaneously just by policies leading to outcomes in $P_{\{1,2\}}$, a contradiction.

We now prove that the above intuition for the impossibility results holds more generally.

Let us introduce an additional axiom, *nondegeneracy with respect to disagreement*, for a bargaining solution to satisfy that states the following. Suppose that two bargaining problems are identical in the set of attainable outcomes, but different in the opt-out payoff vectors, such that some experts (but not all) are strictly better off in one problem relative to the other. Then the bargaining solutions must be different as well.

The intuition behind this requirement is as follows. Opt-out payoffs describe the relative "bargaining strength" of the players. If the opt-out payoffs are increased for some experts relative to others, it must be the case that someone among those experts is better off.

Nondegeneracy (ND) Let F and F' in $\mathcal{F}_n$. If $F(Y^n) = F'(Y^n)$, and there exists subset S, $1 \leq |S| \leq n-1$, such that $d_S(F) > d_S(F')$ and $d_{N\setminus S}(F) \leq d_{N\setminus S}(F')$, then $\phi_i(F) \neq \phi_i(F')$.

The reason for introduction of this axiom is that we have no restrictions on the policy maker's rule μ, nor we specify in any way how bargaining solution ϕ picks outcomes from the set of Pareto efficient and individually rational ones. This joint freedom of choice allows very complicated, intractable relations between the two. Thus, to be able to make a statement about those relations, we need to impose some discipline on ϕ.[6]

We now state the result.

Theorem 2. *For every bargaining solution ϕ that satisfies Axioms IR, PE, C and ND, there does not exist a deterministic rule that implements the socially optimal outcome.*

Proof. Suppose that μ implements the socially optimal outcome x_0, i.e., $\phi(F^{\mu,\theta})+\theta = x_0$ for all $\theta \in \Theta$. W.l.o.g. assume μ is direct.

Fix $\theta \in \Theta$. Consider expert $i \in N$, who has opted out to proposal $\bar{m}_i$. The other experts, $-i$, will propose $m_{-i} = m^*_{-i}(\bar{m}_i, \theta)$ according to their bargaining solution, given that expert i has proposed $\bar{m}_i$, i.e., $m^*_{-i}(\bar{m}_i, \theta) = \phi(F_{-i}(\cdot|\bar{m}_i))$ Expert i thus chooses the proposal $\bar{m}_i = \bar{m}_i(\theta)$ that maximizes her payoff,

$$\max_{m_i \in Y} u_i(\mu(m_i, m^*_{-i}(m_i, \theta)) + \theta - b_i).$$

Denote $\tilde{y}_i(\theta) = \mu(\bar{m}_i(\theta), m^*_{-i}(\bar{m}_i(\theta), \theta))$. That is, $\tilde{y}_i(\theta)$ is the policy implemented after i opts out optimally. Hence, the disagreement payoff of $i = 1, 2$ at state θ is given by

$$d_i(\theta) = u_i(\tilde{y}_i(\theta) + \theta - b_i)$$

and $d(\theta) = (d_1(\theta), \ldots, d_n(\theta))$.

Next, since expert i maximizes her opt-out payoff, at state θ' proposal $\bar{m}_i(\theta')$ will deliver a weakly greater payoff than proposal $\bar{m}_i(\theta)$ for any $\theta \neq \theta'$. That is to say, for every $\theta, \theta' \in \Theta$

$$u_i(\tilde{y}_i(\theta') + \theta' - b_i) \geq u_i(\tilde{y}_i(\theta) + \theta' - b_i) \tag{4}$$

Now, consider any θ, θ' that satisfy $\tilde{y}_i(\theta) + \theta \neq b_i$ and

$$u_i(\tilde{y}_i(\theta) + \theta - b_i) < u_i(\tilde{y}_i(\theta) + \theta' - b_i). \tag{5}$$

[6]Or on μ. In fact, one can prove the result below if the Nondegeneracy requirement on ϕ is replaces by the following requirement of *invariance w.r.t translation* imposed on μ, $\mu(m_1 + \theta, \ldots, m_n + \theta) = \mu(m_1, \ldots, m_n) + \theta$ for all $m = (m_1, \ldots, m_n) \in Y^n$ and all $\theta \in \Theta$.

That is, $y_i(\theta) + \theta'$ is "closer" to b_i than $y_i(\theta) + \theta$. Combining (4) and (5), we obtain

$$d_i(\theta) \equiv u_i(\tilde{y}_i(\theta) + \theta - b_i) < u_i(\tilde{y}_i(\theta) + \theta' - b_i) \leq u_i(\tilde{y}_i(\theta') + \theta' - b_i) \equiv d_i(\theta').$$

That is, for any θ, θ' that satisfy (5) we have $d_i(\theta) < d_i(\theta')$.

Since the experts' most preferred outcomes, $\{b_1, \ldots, b_n\}$, are not identical, one can find θ, θ' and a subset S such that $d_i(\theta) \gneq d_i(\theta')$ for all $i \in S$ and $d_j(\theta) \leq d_j(\theta')$ for all $j \in N \backslash S$. By Nondegeneracy axiom, $\phi(Z^{\mu,\theta}) \neq \phi(Z^{\mu,\theta'})$, a contradiction. ■

5 Cheap Talk

So far we assumed that the policy maker can commit to a rule. We now relax this assumption and show that under certain conditions the socially optimal outcome can be obtained without any commitment, that is, if the communication between the parties is *cheap talk*.

In the cheap talk setting, the policy maker asks the committee to give proposals about a policy to select, but then he is unconstrained in his choice. In this setting, $\mu : Y^n \to \Delta(Y)$ is a *strategy* of the policy maker. The solution concept is sequential equilibrium. In equilibrium, the policy maker will select a policy that maximizes his payoff given his ex-post beliefs about the state.

As the payoff function of the policy maker is strictly concave, randomized policies are not best replies under any beliefs. Thus, in equilibrium the policy maker must choose deterministic policies in every contingency, so $\mu(m) \in Y$ for every $m \in Y^n$. Absent commitment, the policy maker can do even less than with commitment, hence Theorem 2 implies that the socially optimal outcome is not implementable.

Let us now restrict the policy maker's strategies to *closed rules* only. Every closed rule μ is restricts the policy maker to choose a policy (or a lottery over policies) among the experts' proposals. This is a type of constraint that in practice could be imposed by institutional rules.

This restriction to closing rules is an advantage for the policy maker, as now lotteries among policies in $\{m_1, \ldots, m_n\}$ *can be rationalized*, if the policy maker is indifferent among them.

Assumption 1 Strategies of the policy maker are closed rules. That is, strategy μ is characterized by $\eta \in \Delta$ such that for all $m = (m_1, \ldots, m_n) \in Y^n$, $\mu(m)$ chooses policy m_i with probability η_i.

Let us also assume that the payoff of each party is measured as a negative generalized distance between the implemented and the most preferred outcomes.

Assumption 2 Let $\delta : \mathbb{R}^d \to \mathbb{R}_+$ be a differentiable, strictly convex function that satisfies $\delta(0) = 0$ and $\delta(x) = \delta(-x)$. Let

$$\begin{aligned} u_0(x) &= -\delta(x), \\ u_i(x) &= -\delta(x - b_i), \quad i \in N. \end{aligned} \tag{6}$$

We say that socially optimal outcome is *implementable under cheap talk* if there exists a sequential equilibrium where the induced bargaining outcome for the committee is the agreement on policy $y = -\theta$ for every state $\theta \in \Theta$.

Theorem 3. *Let Assumptions 1–2 hold. For every bargaining solution ϕ that satisfies Axioms IR, PE and C, the socially optimal outcome is implementable under cheap talk if and only if it is Pareto undominated.*

Proof. Necessity is immediate by Proposition 1. For sufficiency, we need to show that every closed rule μ^η, $\eta \in \Delta$, defined in Section 3 can be rationalized under cheap talk.

Consider closed rule μ^η for some $\eta \in \Delta$ that implements the socially optimal outcome under commitment. Suppose that the policy maker's strategy in the cheap talk game replicates μ^η. We now show that a lottery η over proposals $(m_1, \ldots, m_n)$ is rational for a specifically selected beliefs of the policy maker.

First, suppose that the condition following holds:

(*) either $n < \mathbf{d} + 1$, or $n = \mathbf{d} + 1$ and the convex hull of P_N has nonempty interior (recall $\mathbf{d}$ is the dimension of outcome space X).

Let $0 \in P_N$. For a tuple $m = (m_1, \ldots, m_n)$ denote by $c(m)$ the point in X which is equally "distant" from all points in $\{m_1, \ldots, m_n\}$,

$$\delta(m_i - c(m)) = \delta(m_j - c(m)) \text{ for all } i, j \in N.$$

Note that $c(m)$ exists if and only if either $n < \mathbf{d} + 1$ or $n \leq \mathbf{d} + 1$ the convex hull of m has nonempty interior.

Now, if all experts agree on some y, the policy maker is constrained to choose y (moreover, in equilibrium his belief that $y = -\theta$ is correct). If the experts disagree, this is out of equilibrium behavior, thus the policy makers' beliefs can be arbitrary.

We set the beliefs as follows. After receiving a proposal tuple m, if $c(m)$ exists, the policy maker believes that $\theta = -c(m)$. Hence, implementing policy $c(m)$ would be optimal. But since he is constrained to choose proposals from m, he is indifferent among those by definition of $c(m)$, and hence he can implement lottery η.

Next, if $n = \mathbf{d} + 1$ and $c(m)$ does not exist, set the beliefs arbitrarily. Note, however, that whenever expert i opts out, her optimal opt-out proposal is $m_i = b_i - \theta$ (see the proof of Proposition 2), while the set of experts who bargain, S, by Pareto Efficiency, agree on a policy y such that $y + \theta \in P_S$. Now, the convex hull of P_S and $\{b_i\}_{i \in N \setminus S}$ equal to P_N that has nonempty interior by assumption, and hence $c(m)$ exists for every set of experts that opt out.

Finally, suppose that condition (*) does not hold. Then one can find a subset of experts, S, such that the convex hull of $\{b_i\}_{i \in N}$ is equal to the convex hull of $\{b_i\}_{i \in N \setminus S}$. By Remark 1, it means that $P_N = P_{N \setminus S}$. Assign $\eta_i = 0$ to every $i \in S$ and let the policy maker ignore their proposals entirely. Further, since any proposal of each $i \in S$ is ignored, let i send a constant, uninformative (independent of the state) proposal. The problem with the remaining experts in $N \setminus S$ satisfies (*). ■

6 Discussion

6.1. Individual Rationality. In our model, the *individually rational* payoff of an expert i is defined as the payoff that i can secure after a unilateral deviation from bargaining. The type of deviation that we consider is *opting out*, that is, announcing the opt-out proposal $\bar{m}_i$ first, and then expecting the rest of the committee to reach the bargaining solution for them when $\bar{m}_i$ is taken as given.

A different type of deviation that one can consider is *defection*, where the committee reaches an agreement first, and then expert i deviates to some proposal m_i, without giving the others a possibility to react. We say that a bargaining solution is *defection-proof* if no expert has incentive to defect.

Defection-proof (DP) *For every bargaining problem $F \in \mathcal{F}^n$, every expert i and every $m_i \in Y$,*

$$u_i(\mu(\phi(F))) \geq u_i(\mu(m_i, \phi_{-i}(F))).$$

Axiom DP turns out to be a very strong requirement, as it is incompatible with Pareto efficiency.

Claim 1. *There does not exist a bargaining solution on $\mathcal{F}^n$ that satisfies PE and DP.*

Proof. Consider a bargaining problem $F^{\eta,\theta}$ defined by any $\theta \in \Theta$ and any closed rule μ^η with full-support distribution, i.e, η is in the interior of Δ. By (3) it is clear that for every expert i and every $m \in Y^n$,

$$\begin{aligned} F_i(\mu^\eta(m)) &= \eta_i u_i(m_i + \theta - b_i) + \sum_{j \neq i} \eta_j u_i(m_j + \theta - b_i) \\ &\leq \eta_i u_i(0) + \sum_{j \neq i} \eta_j u_i(m_j + \theta - b_i) = F_i(\mu^\eta(b_i - \theta, m_{-i})). \end{aligned}$$

That is to say, choosing the policy that induces the most preferred outcome for i, $m_i = b_i - \theta$, is a dominant strategy for i. Hence the only bargaining outcome that satisfies DP is $\phi(\mu^\eta(m)) = (b_1 - \theta, \ldots, b_n - \theta)$, which is Pareto dominated as it induces the lottery over outcomes $\{b_1, \ldots, b_n\}$. ∎

Remark 2. As the strategy of a deviant to propose her most preferred policy is the dominant strategy, the statement of Claim 1 still holds if a deviant's decision and bargaining of the other experts occurs *simultaneously*.

Remark 3. We would like to point out here that our results do not depend on a specific definition of individual rationality. In fact, the Individual Rationality axiom can be replaced by a weaker Dictatorship axiom defined as follows.

Expert i is called *dictator* under rule μ if she dictates the policy, $\mu(m_i, m_{-i}) = m_i$ for all $m \in Y^n$. The Dictatorship axiom states that if i is dictator, then the solution for the committee is the same as the solution for i herself.

Dictatorship (D) *For every bargaining problem $F \in \mathcal{F}^n$, if expert i is dictator, then $\phi(F) = \phi(F_i(\cdot|\bar{m}_{-i}))$ for every $\bar{m}_{-i} \in Y^{n-1}$.*

It is straightforward to verify that replacement of IR by Dictatorship does not affect the proof of Proposition 2 (the only place in the paper where IR is used).

6.2. Second Best Outcome. Suppose that the socially optimal outcome (*first best*) is not implementable, i.e., $0 \notin P_N$. Then the outcome that maximizes the policy maker's payoff on the set of implementable outcomes (*second best*) is simply the closest point in P_N to the origin (w.r.t. the policy maker's preferences), $x^* \in \arg\max_{x \in P_N} u_0(x)$. This property follows from Proposition 1 (asserting that only outcomes in P_N can be implemented) and Proposition 2 (asserting that every outcome in P_N is implementable).

6.3. Choosing a Committee. If the policy maker is free to choose a committee from a set of available experts with various biases, what would be his concerns?

Let $\mathcal{N}$ be a finite collection of available experts. Every expert $i \in \mathcal{N}$ is characterized by a strictly convex payoff function u_i with the maximum achieved at b_i (bias of i).

As follows from Theorem 1, any subset $N \subset \mathcal{N}$ of experts whose Pareto set contains the socially optimal outcome will suffice for the policy maker. Thus, it is never optimal to choose experts with *like* biases. The preferred outcomes for experts should be biased in different directions, so that the socially optimal outcome is located inside of their Pareto set.

What is the least number of experts required to guarantee that the socially optimal outcome can be implemented? In the example in Remark 1, Pareto set for experts in N is the convex hull of their biases. So it can be clearly seen that if the dimensionality of the outcome space is $\mathbf{d}$, it suffices to have $\mathbf{d}+1$ experts to implement the socially optimal outcome. For some non-generic payoff functions, this number could be even less, as, for instance, two experts with biases b and $-b$ will have the origin in their Pareto set for any dimensionality of the outcome space.

The next claim shows that the above argument extends to general payoff functions.

Claim 2. *Let $0 \in P_{\mathcal{N}}$. Then the least set N such that $0 \in P_N$ has cardinality at most (and generically equal to) $\mathbf{d}+1$.*

Proof. Let $\bar{C}_i(x)$ be the upper contour set for expert i at point x,

$$\bar{C}_i(x) = \{x' \in X : u_i(x' - b_i) \geq u_i(x - b_i)\}.$$

Since u_i is concave, $\bar{C}_i(x)$ is convex and thus contained in halfspace $H_i(x)$ with boundary passing through x. For any $N \subset \mathcal{N}$ let $H_N(x) = \bigcap_{i \in N} H_i(x)$. Observe that point $x \in X$ is in P_N for some $N \subset \mathcal{N}$ if and only if $H_N(x)$ has empty interior (as payoff functions are differentiable, any point in the interior of $H_N(x)$ must be a Pareto improvement to x).

Since $0 \in P_{\mathcal{N}}$ and X has dimension $\mathbf{d}$, there exists a subset N of at most $\mathbf{d}+1$ experts, such that either $|N| = \mathbf{d} + 1$ and $H_N(0)$ is a singleton, $H_N(0) = \{0\}$, or $|N| < \mathbf{d} + 1$ and $H_N(0)$ is a hyperplane of dimension less than $\mathbf{d}$. ■

References

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

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

Austen-Smith, D. (1990). Information transmission in debate. *American Journal of Political Science 34*, 124–152.

Banks, J. S. and J. Duggan (2000). A bargaining model of collective choice. *American Political Science Review 94*, 73–88.

Banks, J. S. and J. Duggan (2006). A general bargaining model of legislative policy-making. *Quarterly Journal of Political Science 1*, 49–85.

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

Baron, D. P. and J. A. Ferejohn (1989). Bargaining in legislatures. *American Political Science Review 83*, 1181–1206.

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

Binmore, K. (1987). Perfect equilibria in bargaining models. In K. Binmore and P. Dasgupta (Eds.), *The Economics of Bargaining*. Basil Blackwell, Oxford.

Crawford, V. P. and J. Sobel (1982). Strategic information transmission. *Econometrica 50*(6), 1431–1451.

Dessein, W. (2002). Authority and communication in organizations. *Review of Economic Studies 69*, 811–838.

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

Holmström, B. (1977). *On Incentives and Control in Organizations*. Ph. D. 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. 115–141. North-Holland.

Jackson, M. O. and B. Moselle (2002). Coalition and party formation in a legislative voting game. *Journal of Economic Theory 103*, 49–87.

Kalai, E. (1977a). Nonsymmetric Nash solutions and replications of 2-person bargaining. *International Journal of Game Theory 6*, 129–133.

Kalai, E. (1977b). Proportional solutions to bargaining situations: interpersonal utility comparisons. *Econometrica 45*, 1623–1630.

Kalai, E. and M. Smorodinsky (1975). Other solutions to Nash's bargaining problem. *Econometrica 43*, 513–518.

Krishna, V. and J. Morgan (2001a). Asymmetric information and legislative rules: some amendments. *American Political Science Review 95*(2), 435–452.

Krishna, V. and J. Morgan (2001b). A model of expertise. *Quarterly Journal of Economics 116*, 747–775.

Krishna, V. and R. Serrano (1996). Multilateral bargaining. *Review of Economic Studies 63*, 61–80.

Laffont, J.-J. and D. Martimort (2000). Mechanism design with collusion and correlation. *Econometrica 68*, 309–342.

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

Melumad, N. D. and T. Shibano (1991). Communication in settings with no transfers. *The RAND Journal of Economics 22*, 173–198.

Mylovanov, T. (2008). Veto-based delegation. *Journal of Economic Theory 138*, 297–307.

Mylovanov, T. and A. Zapechelnyuk (2011). Optimal arbitration. Mimeo.

Nash, J. (1950). The bargaining problem. *Econometrica 18*, 155–162.

Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica 50*, 97–110.