Eliciting Information from a Committee*

Andriy Zapechelnyuk

Abstract

The paper addresses the mechanism design problem of eliciting truthful information from a committee of informed experts. The experts are engaged in cooperative bargaining about what information to disclose to the principal. The principal strives to achieve full information disclosure by designing the primitives of the bargaining problem. It is shown that, provided the bargaining solution satisfies Pareto efficiency, continuity and dummy axioms, the full information disclosure is possible if and only if the induced outcome is Pareto undominated for the committee members.

Keywords: Communication; multidimensional mechanism design; experts; collusion; axiomatic bargaining; closed rule

JEL classification numbers: D82, C78

*The author thanks Tymofiy Mylovanov for fruitful conversations that inspired this paper, as well as Shintaro Miura, Ronny Razin and two anonymous referees for helpful comments and suggestions.

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 experts, via cheap talk communication [2, 3, 8, 10, 13, 23], through commitment to certain legislative rules [6, 22, 31] or by (constrained) delegation of decisions to the informed party [1, 11, 12, 15, 16, 21, 28, 29]. The underlying assumption in the above literature is that experts communicate information independently or privately. In contrast, this paper focuses on experts who can collude in their actions. It is harder to elicit the truth from experts in this case, since the experts have more freedom in manipulating the policy maker's decisions. An illustrative example is the problem of 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 cutting the benefits is socially optimal, the experts are not likely to give that recommendation, as, for instance, the advice to *do nothing* is Pareto superior. 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, instead the experts might coordinate on communicating false information that would balance their interests better.

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 any advice is given, the policy maker commits to a policy rule that specifies how she will deal with the advice. Then the committee members engage in bargaining over possible outcomes that their advice could induce through influencing the policy maker's decision. This paper uses the axiomatic approach to the bargaining problem. In other words, we are interested in results that hold under a wide range of specifications of bargaining procedures that satisfy a few requirements (axioms).

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 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 (provided the policy maker can commit to a policy rule).

Our main result (Theorem 1) states that for every bargaining solution that satisfies certain axioms, there exists a rule that implements the first best outcome for the policy maker if and only if it is Pareto undominated for the committee. Our axioms are very basic. We demand the solution to be *Pareto efficient*, *continuous*, and to satisfy the *dummy* axiom (i.e., the bargaining solution should not depend on *dummy* experts whose actions have no effect on the policy maker's decision). 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, the policy maker randomizes over the experts' proposed policies with the probabilities that depend on the identities of the experts, but not on their proposals. 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, as illustrated by example in Section 4.

A crucial assumption for the main result is the ability of the policy maker to randomize over policies. Yet one may wonder whether implementation of the first best outcome is possible by *deterministic* policy rules. We illuminate some difficulties of deterministic first best implementation and argue that it cannot be done by many simple deterministic rules. We also show that in the model with two experts, the first best outcome cannot be implemented by deterministic rules under fairly general conditions if the outcome space is unidimensional, but we provide an example of first best implementation on a multidimensional outcome space.

The results presented in this paper are qualitatively different from those in the cheap talk literature. In the case of two experts, the major insight from the literature is that full information revelation is always possible if the policy maker is unconstrained in her out-of-equilibrium beliefs [8], if the state space is multidimensional and the expert's biases are not collinear [8], or if it is unidimensional and the experts' biases are small relative to the size of the state space [13, 22]. This paper contributes to the literature by showing that if experts can collude, full information revelation is not possible whenever the policy maker's first best outcome *is not* Pareto efficient for the experts.[1] In particular, in contrast to Battaglini [8], in our model with two experts and multidimensional state space, the first best outcome is Pareto dominated for generic biases (it is undominated only if the biases are exactly opposing), hence full information extraction is, in general, impossible.

The paper is closely related to Martimort and Semenov [27] (henceforth, MS) who study information effects of collusive behavior among informed parties.[2] 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.[3] 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

[1] Assuming that collusive agreements are binding. See Proposition 4 in Battaglini [8] for a counterexample for the case of non-binding collusion (more discussion in Remark 2 in Section 3).

[2] Another related paper is Wolinsky [36], whose model is different from ours in that (i) it assumes that the experts are not fully informed and hold distinct pieces of information (like in Martimort and Semenov [27]); and (ii) experts who are allowed to communicate are assumed to have identical preferences, so no conflict of interests arises.

[3] Cf. Laffont and Martimort [25], where a side mechanism is a benevolent entity that maximizes the sum of the parties' utilities.

utility functions, and MS study only deterministic policy rules. For this environment, Proposition 4 in our paper shows that a first-best continuous deterministic rule does not exist, whereas MS (Proposition 2) find the second best continuous deterministic rule. MS (Footnote 17) comment that "It is not known whether random mechanisms could help in our context with a continuum of types." Our main result answers this question.

Other related papers are Mylovanov and Zapechelnyuk [31] that characterizes optimal stochastic rules for an uninformed principal in a model with a unidimensional state space and two experts, and Mylovanov and Zapechelnyuk [30] that deals with first-best stochastic rules in a multidimensional environment with two experts. Closed rules in the current paper are similar to stochastic rules in Mylovanov and Zapechelnyuk [30, 31] in that experts' disagreements are punished by randomized policies. As Mylovanov and Zapechelnyuk [30, 31] do not consider collusion of experts, in their models uncertainty after a disagreement is simply a tool to provide incentives to the experts to agree on the truth. In contrast, in this paper the choice of the probability distribution over policies after a disagreement is the policy maker's instrument for manipulation of the experts' collusive behavior.

This paper is also related to the non-cooperative theory of multilateral bargaining that goes back to Rubinstein [35] and has been further developed by Binmore [9], Hart and Mas-Colell [14], Krishna and Serrano [24], Okada [33, 34], with particular focus on bargaining in legislative bodies in Baron and Ferejohn [7], Jackson and Moselle [17], and Banks and Duggan [4, 5]. 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 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 (asymmetric) Nash bargaining solution [32, 18], Kalai-Smorodinsky solution [20], and the egalitarian solution (more generally, the proportional solution) [19].

2 The Model

2.1 Preliminaries. There is a policy maker and $n \geq 2$ 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 $Y = \mathbb{R}^{\mathbf{d}}$, $\mathbf{d} \geq 1$. When a policy $y \in Y$ is chosen, an outcome is $x = y + \theta$, where $\theta \in \Theta = \mathbb{R}^{\mathbf{d}}$ is an unobserved state.[4]

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[5] that stipulates for every tuple of the experts' proposals $m = (m_1, \ldots, m_n)$ to choose a policy (or a lottery over policies) $\mu(m)$.

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 $b_i \in \mathbb{R}^{\mathbf{d}}$ 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 $\{h(a) : a \in A'\}$.

2.2. The Bargaining Problem. The committee members, $N = \{1, \ldots, n\}$, after observing state θ and rule μ but 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 $m = (m_1, \ldots, m_n)$ yields to each expert i expected payoff

$$E_{\mu(m)}[u_i(y + \theta - b_i)],$$

[4]State space Θ and policy set Y need not be unbounded for our results to hold; it is sufficient to require that Θ and Y are rich enough such that $\{y + \theta : y \in Y\}$ contains the policy maker's and experts' most preferred outcomes for any given θ, and so does $\{y + \theta : \theta \in \Theta\}$ for any given y. The outcome set is thus $X = \{y + \theta : y \in Y, \theta \in \Theta\}$.

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

where $E_{\mu(m)}[\cdot]$ denotes the expectation with respect to distribution $\mu(m)$ over policies. This is a composition of rule $\mu(m)$ that maps proposal tuple m into a (randomized) outcome and payoff function $u_i(y+\theta-b_i)$ that maps realized outcome y into a payoff of i. We also assume that proposal tuples that lead to very extreme negative payoffs are never supposed to be realized and thus irrelevant for the bargaining. Fix a large constant $L>0$ (much larger than $\max_{i,j}||b_i-b_j||$) and define

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

Thus we have imposed a lower bound on the payoffs. This assumption is needed for compactness of the bargaining problem and it does not affect the results so long as L is sufficiently large.

Let $F^{\mu,\theta}(m)=(F_1^{\mu,\theta}(m)\ldots,F_n^{\mu,\theta}(m))$.

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.

Here we do not specify an underlying strategic-form game. A bargaining solution ϕ is seen as a "black box" that takes state θ and the policy maker's rule μ as inputs and provides a tuple of proposals, m, to be sent to the policy maker as an output, and that satisfies three simple requirements (axioms) described below.

A proposal tuple $m\in Y^n$ is said to be *ex-post Pareto efficient* if there does not exist another proposal tuple, $m'\in Y^n$, $m'\neq m$, 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 ex-post Pareto efficient.*

An expert i is called *dummy* if her proposal does not affect the payoffs, that is, $F(m_i,m_{-i})=F(m'_i,m_{-i})$ for all $m_i,m'_i\in Y$ and all $m_{-i}\in Y^{n-1}$. The next axiom states that if expert i is dummy, the bargaining solution for the others should not depend on whether i is included to the problem or excluded from it.

Denote by $\hat{F}_{-i}$ the restriction of F to the coordinates of the experts other than i, assuming that i is dummy.[6]

Dummy (D) *For every bargaining problem* $F \in \mathcal{F}^n$*, if expert* i *is dummy, then* $\phi_j(F) = \phi_j(\hat{F}_{-i})$ *for every* $j \in N \backslash \{i\}$.

Finally, we assume that a bargaining outcome does not change drastically in response to any small change of a bargaining problem.

Continuity (C) ϕ *is continuous.*[7]

No further constraints on a bargaining solution are necessary for our results.

3 First Best Implementation

We now establish necessary and sufficient conditions under which the first best outcome is implementable, and construct a rule that implements it.

We say that for a given tuple of biases $b = (b_1, \ldots, b_n)$ and 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,

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

Note that this is *unique implementation*: rule μ that satisfies (2) implements the unique outcome, x.[8]

Outcome x is said to be *implementable* if there exists a rule that implements it. Rule μ is called *first best* if it implements the policy maker's first best outcome, $x_0 = 0$.

For every subset of experts S let

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

[6] That is, $\hat{F}_j(m_{-i}) = F_j(m_i, m_{-i})$ for all (m_i, m_{-i}) and all $j \neq i$.

[7] Specifically, $\phi : \mathcal{F}_k \to \mathbb{R}^k$ is continuous w.r.t. the topology of uniform convergence on $\mathcal{F}_k$, $\forall k \leq n$.

[8] By definition, ϕ is a function, i.e., for any (μ, θ) the bargaining outcome $\phi(F^{\mu,\theta})$ is unique. However, in practice, a bargaining outcome is an equilibrium of some non-cooperative bargaining game that may have multiple equilibria (exceptions are rare, such as the unique subgame perfect equilibrium in Rubinstein 35). In this case, bargaining solution ϕ picks a specified equilibrium, and implementation is not unique.

We say that outcome x is *Pareto undominated for experts in S* (or just *Pareto undominated* if $S = N$) if there does not exist another outcome x' with $U_S(x') \gneq U_S(x)$.[9] The set of Pareto undominated outcomes (*Pareto set*) for the experts in S is denoted by P_S.

We now state the main result.

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

This theorem follows from Propositions 1 and 2 below that establish, respectively, *only if* and *if* statements.

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

In other words, there does not exist a rule that implements a Pareto dominated outcome under any bargaining solution ϕ that satisfies Pareto Efficiency. This is a very intuitive result. Assume that outcome x_0 is implementable by rule μ and Pareto dominated by some x', and fix state θ. Then x_0 can be an outcome of ϕ only if x' is not achievable, i.e., $\mu(m) + \theta \neq x'$ for all $m \in Y^n$. But then, at state $\theta' = \theta + x' - x_0$, outcome x_0 is not achievable either, hence it is not implementable at that state, a contradiction.

The proof of the *if* part of Theorem 1 (the possibility of implementation of any Pareto undominated outcome) is constructive. We design a special class of rules for the policy maker, called *closed rules*, that utilize information about the experts' biases and manipulate the bargaining problem among experts in such a way that the bargaining solution yields exactly the desired policy at each state.

A *closed 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, the policy maker randomizes over the experts' proposed policies with the probabilities that depend on the identities of the experts, but not on their proposals.

Definition 3. A rule μ is called *closed* if there exists $p = (p_1, \ldots, p_n) \in \Delta(N)$ such that for all $m = (m_1, \ldots, m_n) \in Y^n$, $\mu(m)$ chooses policy m_i with probability p_i.

[9] We write $z \gneq z'$ if z is greater than z' in every coordinate and strictly greater in some.

A closed rule defined by probability distribution $p \in \Delta(N)$ is denoted by μ_p.

Proposition 2 (Sufficiency). *For every bargaining solution ϕ that satisfies Axioms PE, D and C, and for every Pareto undominated outcome x there exists $p \in \Delta(N)$ such that closed rule μ_p implements x.*

The sketch of the proof is as follows. To begin with, assume that outcome x coincides with the most preferred outcome for expert i (i.e., i has zero bias). Then x is trivially implemented by closed rule μ_p where $p_i = 1$, so expert i is the "dictator" and all the other experts are dummies. By Dummy axiom (applied recursively to each dummy), the bargaining outcome depends entirely on expert i's choice, and by Pareto Efficiency, i will choose the policy that leads i's most preferred outcome, which coincides with x.

Next, suppose that x is located on the Pareto set $P_{\{i,j\}}$ for some pair of experts, i and j. Then x can be implemented by closed rule μ_p with $p_i = 1 - p_j$, and $p_k = 0$ for all $k \notin \{i, j\}$ as follows. Again, by Dummy axiom, the bargaining solution depends only on the choice of $\{i, j\}$ and by Pareto Efficiency it will be in $P_{\{i,j\}}$. The two end-points of the interval $P_{\{i,j\}}$ are implemented by choosing $p_i = 0$ and $p_i = 1$, respectively. By Continuity axiom, assigning probabilities $(p_i, p_j) = (t, 1 - t)$ and moving t from 0 to 1 will span the entire interval $P_{\{i,j\}}$. The proof then proceeds by induction w.r.t. increasing subsets of experts. The complete proof is deferred to the Appendix.

Remark 1. Although Proposition 2 does not tell us how probability vector p should be chosen, the sketch of the proof provides some intuition. One can conjecture that proposals of experts with smaller biases should be assigned higher probabilities. The experts' relative bargaining strength depends on the assigned probabilities, since a relatively high probability p_i means that expert i's opinion carries a large weight, and thus she has less to lose in case of a disagreement. So, experts whose interests are less conflicting with the policy maker (i.e., whose biases are smaller) should be granted greater bargaining power to be able to pull the bargaining outcome closer to their preferred outcomes. This is illustrated by the application of Proposition 2 to Nash bargaining solution in the next section.

Remark 2. An important assumption behind Theorem 1 is that the bargaining agreement among the experts is *binding*. It exacerbates the problem of first best implementation, as it precludes (unilateral) deviations of the experts and thus allows the experts to sustain any agreement. Indeed, Battaglini [8] shows that without this assumption, almost all outcomes[10] in a multidimensional space can be implemented by consulting with just two experts. Battaglini's equilibrium construction is *collusion proof*, in the sense that collusion on any outcome that is Pareto superior to the first best outcome is unsustainable (i.e., some expert has incentive to deviate).[11] In contrast, in our model collusive agreements are binding, thus the experts are *always* able to sustain a Pareto superior outcome.

Another important assumption is that the policy maker receives no information in addition to a tuple of the experts' reports. For example, she cannot distinguish whether the experts have reached an agreement or, if not, which experts have deviated (or any other details of the underlying bargaining procedure). This assumption also exacerbates the problem of first best implementation, as observing more information would allow to construct more versatile mechanisms that, possibly, could implement some Pareto dominated outcomes.

Remark 3. Suppose that the first best outcome is Pareto dominated, and thus unimplementable by Proposition 1. Then the *second best* outcome (that *is* implementable) is simply the undominated point that is closest (w.r.t. the policy maker's preferences) to the first best outcome.

4 The Nash Bargaining Solution

Let us consider an application of Theorem 1 to a problem with n experts who resolve the bargaining problem according to the Nash bargaining solution [32].[12]

Let the outcome space be $X = \mathbb{R}^{n-1}$. For a given bargaining problem $F^{\mu,\theta}$ define the *disagreement payoff profile* $d^{\mu,\theta} = (d_1^{\mu,\theta}, \ldots, d_n^{\mu,\theta})$ as follows. If the experts fail to agree, we assume that they choose proposals noncooperatively and sequentially, in a random order (all orders are equally likely).[13] For every order a subgame perfect equilibrium

[10] Excluding those on the line that passes through the most preferred outcomes of the two experts.

[11] This result is proven for quadratic payoff functions of the experts.

[12] See Lensberg [26] for the axiomatic extension to multilateral bargaining.

[13] The structure of the order of moves is unimportant for the results of this section (for instance, the experts can move simultaneously), since we will show later that under any closed rule each expert has a unique dominant action. The reason for this assumption is that, in general, subgame perfect equilibrium of the game with a given order of sequential moves typically is unique, so there is no ambiguity about the disagreement payoffs.

is specified (if it is not unique). The *disagreement payoff* $d_i^{\mu,\theta}$ of expert i is her ex ante (before the order is drawn) expected payoff from the described disagreement game.

The *Nash bargaining solution* for bargaining problem $F^{\mu,\theta}$ is the proposal tuple m^* that maximizes the product of the experts' surpluses,[14]

$$m^* = \arg\max_{m \in Y^n} \prod_{i \in N} \left(F_i^{\mu,\theta}(m) - d_i^{\mu,\theta} \right). \tag{3}$$

A noncooperative model that implements the Nash bargaining solution is, for instance, the random proposer model of Okada [33, 34].[15] At every round $t = 1, 2, \ldots$ one player is selected as a proposer with equal probability from the set of *active* players, denoted by N_t (with $N_1 = N$). The selected player i proposes (i) a coalition $S \subset N_t$ with $i \in S$ and (ii) an action tuple m_S for that coalition. All other players in S either accept or reject the proposal sequentially (the order of responders does not affect the results in any crucial way). If all the other players in S accept the proposal, then it is agreed upon, and players in S commit to m_S. The remaining players outside S become active players for the next round, $N_{t+1} = N_t \backslash S$, so long as N_{t+1} is nonempty. If at least one responder rejects the proposal, with probability $1 - \varepsilon$ the negotiations continue in the next round with $N_{t+1} = N_t$. With probability ε the negotiations break down and all players obtain their disagreement payoffs. Okada [34, Theorem 1] shows that as $\varepsilon \to 0$, the payoff profile in any stationary subgame perfect equilibrium of this game approaches the payoff profile of the Nash bargaining solution.

Suppose that the experts' preferences satisfy the following assumptions:

(C1) The payoff functions are quadratic,[16] $u_i(x - b_i) = -||x - b_i||^2$, $i \in N$.

(C2) The experts' biases are vectors whose directions are symmetric w.r.t. the origin (the angle between any pair of vectors is $\arccos[-1/(n-1)]$).

Fig. 1 illustrates the outcome space for $n = 3$. The most preferred outcomes of the experts are labeled by b_1, b_2 and b_3, respectively. The directions of the biases are symmetric so that the angle between each pair is equal to $\alpha = 120^{\mathrm{o}}$. The shaded area depicts the set of Pareto undominated outcomes.

[14]There exists a unique solution, since $F^{\mu,\theta}$ is strictly convex and bounded, and by construction $d^{\mu,\theta} \leq F^{\mu,\theta}(m)$ for some m.

[15]For other noncooperative procedures that lead to the multilateral Nash bargaining solution, see Hart and Mas-Colell [14] and Krishna and Serrano [24].

[16]We write $|| \cdot ||$ for the Euclidean norm on $\mathbb{R}^n$.

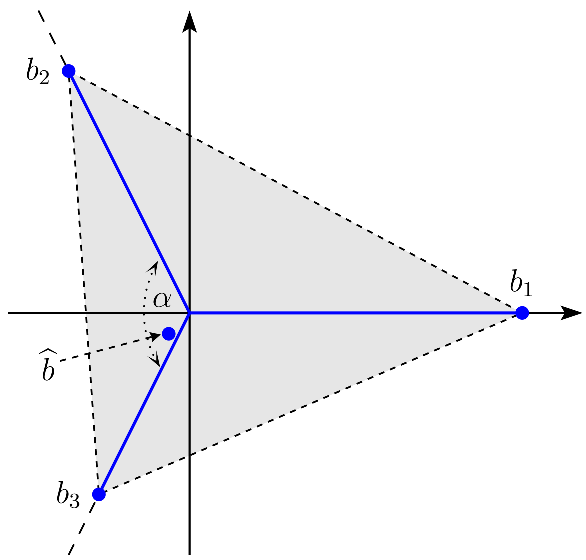

Fig. 1. The problem with three experts.

Let us now assume that the policy maker applies closed rule μ_p (that is, she chooses the proposal of each expert i with fixed probability p_i that is independent of the proposals) and determine probability vector $p = (p_1, \ldots, p_n)$ as a function of experts' biases such that the Nash bargaining solution implements the first best outcome for the policy maker, the origin.

Note that under any closed rule μ_p, the dominant disagreement action of any expert i is to propose her most preferred action, $m_i = b_i - \theta$.[17] Consequently, the disagreement payoff of i is independent of θ and given by

$$d_i^{\mu_p,\theta} = -\sum\nolimits_{j \in N} p_j ||b_i - b_j||^2. \tag{4}$$

The Nash bargaining outcome is a function of the set $F^{\mu_p,\theta}(Y^n)$ of feasible payoff vectors and the disagreement payoff vector d^{μ_p}. Note that under any closed rule, $F^{\mu_p,\theta}(Y^n)$ is independent of p, since the experts can induce any outcome x in X by proposing $m = (x - \theta, \ldots, x - \theta)$. But d^{μ_p} depends on p. The policy maker achieves the desired outcome by adjusting the disagreement payoffs with the choice of p. For example, increasing p_i strengthens the bargaining position of expert i and shifts the bargaining

[17] To see this, assume that i is the last to choose an action. Then, irrespective of what the others have chosen, expert i's payoff is maximized when $m_i = b_i - \theta$. The argument is complete by backward induction.

outcome towards the most preferred outcome of i.

Proposition 3. *Suppose that (C1)–(C2) hold. Let the experts choose their proposals according to the Nash bargaining solution. Then the first best outcome for the policy maker is implementable by closed rule μ_p where $p = (p_1, \dots, p_n)$ is the probability distribution that satisfies $p_i = w_i / \sum_{j \in N} w_j$, where*

$$w_i = \frac{1}{||b_i||^2} - \frac{1}{||b_i||} \frac{n-3}{n-1} \left(\frac{1}{n} \sum\nolimits_{j \in N} ||b_j|| \right)^{-1}, \quad i \in N.$$

In particular, for $n = 3$,

$$w_i = \frac{1}{||b_i||^2}, \quad i = 1, 2, 3,$$

and the dot labeled by $\widehat{b}$ on Fig. 1 shows the weighted average of the depicted biases.

The proof is deferred to the Appendix.

The above proposition shows that probability p_i assigned to a proposal of expert i is decreasing in length of her bias, approaching zero as $||b_i|| \to \infty$. Also, p_i is increasing in the length of bias of any other expert $j \neq i$ and approaches zero as $||b_j|| \to 0$.

As concerns general biases (assumption (C2) dropped), even for $n = 3$ the analytical solution for probabilities p_i is very cumbersome, as it depends not only on lengths of biases, but also on their directions. We omit the solution but note the following properties. Consider Fig. 1 and assume that angle α between biases of experts 2 and 3 increases, while keeping these biases symmetric w.r.t. the horizontal axis. Then the probability p_1 assigned to expert 1 goes down. Eventually, as b_2 and b_3 become vertical and opposite to each other, expert 1 becomes redundant, so $p_1 = 0$. And vice versa, if α decreases, then the probability p_1 assigned to expert 1 goes up. So an expert's opinion matters more not only when her bias is small relative to the others, but also when the conflict of interests with all other experts is large.

5 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 $m \in Y^n$. In this section we discuss some difficulties of deterministic first best implementation and argue that it cannot be done by many simple deterministic rules. We also show that if the outcome space is unidimensional, the first best outcome cannot be implemented by continuous deterministic rules, but we provide an example of first best implementation on a multidimensional outcome space for $n = 2$.

Consider the Nash Bargaining solution.[18] The policy maker who wishes to implement outcome x_0 should be able to choose a rule μ as a function of the experts' biases so that the outcome of the resulting bargaining problem implements exactly x_0. As we have seen in the previous section, under the Nash bargaining solution, the outcome depends only on the experts' disagreement payoffs[19] $d^{\mu,\theta}$, which, in turn, depend on μ.

Let us consider two simple classes of rules and identify some issues that prevent these rules from implementing the first best outcome.

Weighted average rules. A weighted average rule is defined by $\mu(m_1, \ldots, m_n) = \sum_{i \in N} w_i m_i$ for some weights satisfying $0 \leq w_i \leq 1$ for all i and $\sum w_i = 1$.

Suppose that x_0 is in the interior of the set of Pareto undominated outcomes (as on Fig. 1). Then, under the weighted average rule, the disagreement payoffs are the expected payoffs from a fixed lottery over the set $\{b_1, \ldots, b_n\}$ of the most preferred outcomes of the experts. That is, $d^{\mu,\theta}$ does not depend on the choice of weights $(w_1, \ldots, w_n)$, hence the bargaining outcome is constant and, in general, need not coincide with the first best outcome x_0.[20]

Let us explain why the above is true. To implement outcome x_0 in the interior of the set of Pareto undominated outcomes, no expert should be dummy, i.e., $w_i > 0$ for all i. Recall our assumption that the disagreement in the Nash bargaining problem is defined by the noncooperative game among the experts where they move sequentially, according to a random order. Notice that the last expert in that order, called i, is able to induce her preferred outcome by proposing

$$m_i = \frac{1}{w_i}\left(b_i - \theta - \sum\nolimits_{j \neq i} w_j m_j\right)$$

[18] The arguments in this section apply to any bargaining solution that depends entirely on the set of feasible payoff vectors for the experts and the disagreement payoff vector, such as asymmetric Nash solution [18], Kalai-Smorodinsky solution [20], and egalitarian solution (more generally, proportional solution) [19].

[19] It also depends on the set of outcomes achievable for the experts, but it is constant across all rules that implement x_0 as argued in the proof of Proposition 1.

[20] See Remark 4 in Appendix A.1 for a technical explanation why Proposition 2 does not hold if we replace closed rules by weighted average rules.

so that $\mu(m) + \theta = b_i$. As all orders are equally likely, the disagreement outcome is the equal-probability lottery over the set $\{b_i\}_{i \in N}$ of the most preferred outcomes of the experts, which is independent of the weights.

Clearly, the problem described above is not specific to the weighted average rule, it will persist in many other rules that aggregate the proposals. Moreover, it will persist for some other definitions of disagreement. For instance, suppose that an expert can unilaterally disagree with the others, and her disagreement payoff is determined by what she can secure if she anticipates that the rest of the experts induce the bargaining outcome for their $(n-1)$-player problem. In this case, under the weighted average rule, the disagreement payoff vector still depends on the biases of the experts, but not on weights $(w_1, \ldots, w_n)$.

Majority rule. The majority rule stipulates to choose the action that the majority of experts, $k > n/2$, propose. If no action has the majority, then a constant *punishment action* $\hat{y} \in Y$ is chosen.

There are two problems associated with a constant punishment action. First, $\hat{y}$ cannot depend on state θ (which is unknown to the policy maker). At some state θ action $\hat{y}$ may be *extreme*, in the sense that each expert prefers all undominated outcomes (including the most preferred outcomes of the others) to the punishment outcome $\hat{y}+\theta$. But there may be another state θ' where $\hat{y}$ is not extreme, for example, it is the most preferred action for some expert i, i.e., $\hat{y} + \theta' = b_i$. So, at different states the experts may face different bargaining problems and hence agree on different bargaining outcomes, while the policy maker's objective is to make them agree on the same outcome x_0.

Even if we assume that the set of states in bounded, so that there exists action $\hat{y}$ that is *extreme* at all states, there is another problem. As $\hat{y}$ is extreme, in the disagreement game no expert would take any actions that lead to $\hat{y}$ being chosen. That is to say, the disagreement outcome cannot depend on $\hat{y}$. Hence, the bargaining outcome is constant w.r.t. $\hat{y}$ and, in general, different from the first best outcome x_0.

For illustration, let us solve the disagreement game under the majority rule for $n = 3$. Fix an order, say, $\{1, 2, 3\}$. We proceed by backward induction. The last in the order, expert 3, will form the majority with either expert 1 or 2, by choosing action $m_3 \in \{m_1, m_2\}$ that gives her the higher payoff among the two. Expert 2 chooses an action that maximizes her payoff among all actions that provide to expert 3 a greater

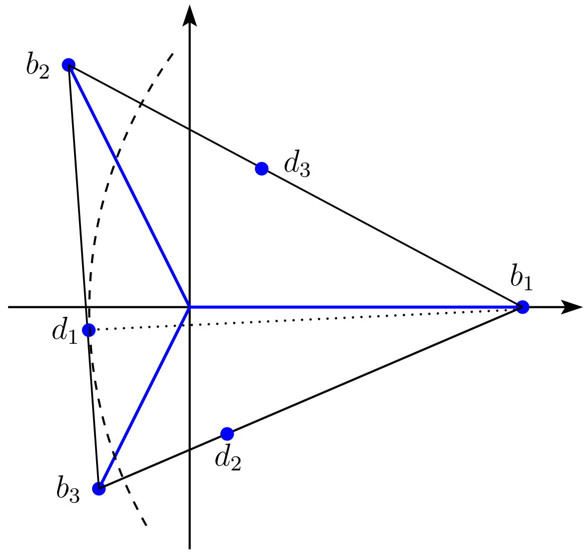

Fig. 2. Disagreement outcomes under the majority rule.

payoff than that at m_1. Expert 1 then chooses m_1 optimally, such that there does not exist action m_2 that provides a greater payoff to 2 and a weakly greater payoff to 3.

This is illustrated by Fig. 2. The optimal action of expert 1 is chosen to induce outcome d_1 (that is, $m_1 = d_1 - \theta$), where d_1 is her preferred outcome among Pareto undominated outcomes for the pair of experts $\{2, 3\}$ (the dashed line depicts expert 1's highest indifference curve that is tangent to segment $[b_2, b_3]$). Then the optimal response of expert 2 is to choose $m_2 = m_1$, since any action that would benefit expert 2 relative to m_1 will make expert 3 worse off, who would then block m_2 by forming the majority with expert 1, $m_3 = m_1$. The disagreement outcomes for order $\{1, 3, 2\}$ coincides with d_1. Similarly we obtain points d_2 and d_3 for orders where 2 and 3 move first, respectively. As each order is equally likely, the resulting disagreement outcome is the equal-probability lottery over outcomes $\{d_1, d_2, d_3\}$. The bargaining outcome is thus a function of the expert's biases, but not of the punishment action $\hat{y}$.

Clearly, the described problems are not specific to the majority rule and generally apply to rules that rely on a constant punishment.

As concerns a broader class of deterministic rules, we consider the case of $n = 2$ and show that, in general, the first best outcome cannot be implemented by continuous deterministic rules if the outcome space is unidimensional.

A bargaining solution ϕ is called *disagreement dependent* if for every bargaining prob-

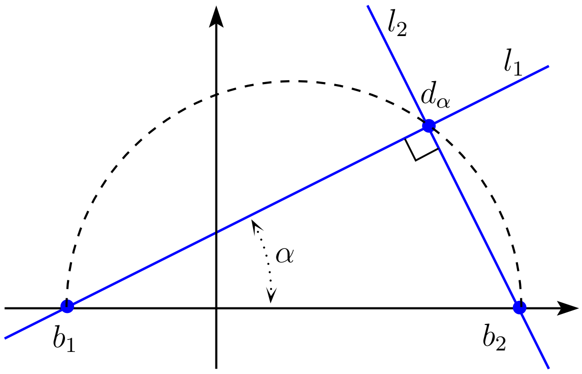

Fig. 3. First best rule for $n = 2$ and $X = \mathbb{R}^2$.

lem $F \in \mathcal{F}_2$, it is a function only of the set of feasible payoff vectors, $F(Y^2)$, and a disagreement payoff vector that the experts obtain in equilibrium by playing noncooperatively, either simultaneously, or sequentially with random order.

Proposition 4. *Let $n = 2$ and $X = \mathbb{R}$. Let the experts choose their proposals according to a disagreement dependent solution that satisfies PE. Then for almost every*[21] *bias tuple $b = (b_1, b_2)$ in X^2, the does not exist a continuous deterministic rule that implements the first best outcome.*

The proof is deferred to the Appendix.

As concerns the case of $X = \mathbb{R}^{\mathbf{d}}$, $\mathbf{d} \geq 2$, we now construct a first best continuous deterministic rule for $X = \mathbb{R}^2$ as follows (it is easily extended to any $\mathbf{d} > 2$). Assume that the biases are opposing (otherwise the first best outcome is not implementable by Proposition 1). W.l.o.g. let $b_1 = (\underline{b}_1, 0)$ and $b_2 = (\underline{b}_2, 0)$ such that $\underline{b}_1 \leq 0 \leq \underline{b}_2$. Fix an angle $\alpha \in [0, \pi/2]$ and suppose that expert 1 chooses a line l_1 in $\mathbb{R}^2$ with slope $\tan\alpha$ and expert 2 chooses a line l_2 with slope $-\cot\alpha$. The implemented policy is the intersection of these lines, as shown on Fig. 3. Formally, for every $(m_1, m_2) \in \mathbb{R}^2$ define

$$\mu_\alpha(m_1, m_2) = \Big(m_1 \sin^2\alpha + m_2 \cos^2\alpha,\ (m_2 - m_1)\sin\alpha\cos\alpha\Big).$$

[21]Excluding a measure zero subset of X^2. For example, with biases $(0, b_2)$ (expert 1 is unbiased) the first best outcome is trivially implemented by $\mu(m_1, m_2) = m_1$.

Then sets $\mu_\alpha(\mathbb{R}, m_2)$ and $\mu_\alpha(m_1, \mathbb{R})$ are the lines in $\mathbb{R}^2$ given by

$$l_1(m_1) = \mu_\alpha(m_1, \mathbb{R}) = \left\{(t, z) \in \mathbb{R}^2 \, : \, z = (t - m_1)\tan\alpha\right\},$$
$$l_2(m_2) = \mu_\alpha(\mathbb{R}, m_2) = \left\{(t, z) \in \mathbb{R}^2 \, : \, z = -(t - m_2)\cot\alpha\right\}$$

Let us now find the disagreement outcomes for two different orders, $\{1, 2\}$ and $\{2, 1\}$. If expert 1 moves first, then expert 2 responds by choosing $m_2 = m_2^*(m_1)$ to minimize the distance between $b_2 - \theta$ and the intersection of $l_2(m_2)$ with $l_1(m_1)$ that is the closest point to $b_2 - \theta$ (i.e., $l_2(m_2)$ and $l_1(m_1)$ are orthogonal). Anticipating that response, expert 1 chooses m_1 to minimize the distance between $b_1 - \theta$ and the intersection of $l_1(m_1)$ and $l_2(m_2^*(m_1))$. It is easy to see that the experts will optimally choose $m_1 = b_1 - \theta$ and $m_2 = b_2 - \theta$ irrespective of the order of moves, as illustrated by Fig. 3 for $\theta = 0$. The disagreement outcome, labeled by d_α on Fig. 3, is the same for both orders $\{1, 2\}$ and $\{2, 1\}$.

The set of disagreement outcomes induced by rules μ_α for various $\alpha \in [0, \pi/2]$ is the arc depicted by the dashed line on Fig. 3. As α increases, d_α moves from b_2 to b_1 along the arc. The disagreement payoff of expert 1 increases and the disagreement payoff of expert 2 decreases monotonically and continuously. As the Nash bargaining outcome depends on d_α only, any outcome between b_1 and b_2 can be implemented by an appropriate choice of α.

Appendix

A.1 Proof of Proposition 2

Consider bargaining problem $F^{\mu_p,\theta}$ defined by closed rule μ_p and state θ. Let ϕ be a bargaining solution that satisfies axioms PE, D and C.

Note that, since the experts' payoff functions are strictly concave, every nondegenerate lottery over actions is Pareto inferior to some deterministic action. Hence by Pareto Efficiency, for every bargaining problem $F^{\mu_p,\theta}$, $\phi(F^{\mu_p,\theta}) = (y, \ldots, y)$ for some action $y \in Y$. With abuse of notation, we will write $\phi(F^{\mu_p,\theta}) = y$.

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

$$\phi(F^{\mu_p,\theta}) + \theta = x \quad \text{for all } \theta \in \Theta. \tag{5}$$

Note that any closed rule μ_p has the following property. For every $(m_1, \dots, m_n) \in Y^n$ and every θ,

$$F^{\mu_p,\theta}(m_1 - \theta, \dots, m_n - \theta) = F^{\mu_p,0}(m_1, \dots, m_n).$$

Hence, the bargaining solutions of $F^{\mu_p,\theta}$ and $F^{\mu_p,0}$ are the same up to translation by θ, so $\phi(F^{\mu_p,\theta}) + \theta = \phi(F^{\mu_p,0})$ for all $\theta \in \Theta$.

Define $\xi : \Delta(N) \to X$ by $\xi(p) = \phi(F^{\mu_p,0})$, which is continuous, since $F^{\mu_p,0}$ is uniformly continuous in p by construction and ϕ is continuous w.r.t. the uniform convergence topology by Continuity axiom. Let N' be the minimal subset of experts whose Pareto set coincides with the Pareto set for N, i.e., $P_{N'} = P_N$, and $P_S \subsetneq P_N$ for all $S \subsetneq N'$. Using these notations, (5) can be equivalently written as $P_{N'} \subset \xi(\Delta(N))$.

For every $S \subset N$ denote by $\Delta_S \subset \mathbb{R}^n$ the set of probability distributions in $\Delta(N)$ with support on the coordinates in S only,

$$\Delta_S = \{p \in \Delta(N) : p_i = 0 \text{ for all } i \notin S\}.$$

The proof of $P_{N'} \subset \xi(\Delta(N))$ is split into two steps.

Step 1. We show that

$$\xi(\Delta_S) \subset P_S \quad \text{for all } S \subset N'. \tag{6}$$

Step 2. By induction in the cardinality of S we show that $\xi(\Delta_S) = P_S$ for all $S \subset N'$.

By Step 2 it follows that $P_{N'} = \xi(\Delta_{N'}) \subset \xi(\Delta(N))$.

Proof of Step 1. Let $S \subset N'$ and consider $p \in \Delta_S$, so that $p_j = 0$ for all $j \in N \backslash S$. Then, under rule μ_p, for every $i \in N$,

$$\begin{aligned} F_i^{\mu_p,\theta}(m) &= \sum\nolimits_{j \in S} p_j u_i(m_j + \theta - b_i) + \sum\nolimits_{j \in N \backslash S} p_j u_i(m_j + \theta - b_i) \\ &= \sum\nolimits_{j \in S} p_j u_i(m_j + \theta - b_i). \end{aligned}$$

Hence, $F^{\mu_p,\theta}$ is independent of $m_{N\backslash S}$. In other words, all experts in $N \backslash S$ are dummies. Applying Dummy axiom repeatedly for all $j \in N \backslash S$ yields that $\phi_i(F^{\mu_p,\theta})$ is equal to the solution of the bargaining problem among experts in S, and by Pareto Efficiency,

$\phi(F^{\mu_p,\theta}) \in P_S$. Consequently, $\xi(\Delta_S) \subset P_S$. ■

Proof of Step 2. Now let us show that $\xi(\Delta'_N) = P_{N'} \equiv P_N$. We proceed by induction in the cardinality of $S \subset N'$. First, let $S = \{i\}$ for some $i \in N'$. Observe that $P_{\{i\}} = \{b_i\}$ (since $u_i(x - b_i)$ is strictly concave and maximized at $x = b_i$). As $\xi(\Delta_{\{i\}})$ is nonempty (it is a singleton) and, by (6), $\xi(\Delta_{\{i\}}) \subset P_{\{i\}} = \{b_i\}$, we have $\xi(\Delta_{\{i\}}) = P_{\{i\}}$.

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

Note that P_S is homeomorphic to set Δ_S, i.e., there exists a continuous bijection $h_S : \Delta_S \to P_S$.[22] Moreover, for every $i \in S$, $P_{S\setminus\{i\}} = h_{S\setminus\{i\}}(\Delta_{S\setminus\{i\}})$. As $\xi(\Delta_S) \subset P_S$ by (6), condition $\xi(\Delta_S) = P_S$ is equivalent to $h_S^{-1}(\xi(\Delta_S)) = \Delta_S$.

Denote by $\partial\Delta_S$ the boundary of the set $\{p \in \Delta(N) : p_i > 0 \text{ for all } i \in S\})$ and let $\partial P_S = h_S(\partial\Delta_S)$. By induction assumption, $P_{S\setminus\{i\}} = \xi(\Delta_{S\setminus\{i\}})$, hence

$$\partial P_S \subset \bigcup\nolimits_{i\in S} P_{S\setminus\{i\}} = \bigcup\nolimits_{i\in S} \xi(\Delta_{S\setminus\{i\}}) = \xi\left(\bigcup\nolimits_{i\in S} \Delta_{S\setminus\{i\}}\right) \subset \xi(\Delta_S).$$

Now, the homotopy group of set $\xi(\Delta_S)$ is trivial (i.e., it has no "holes" inside), since ξ is continuous and we can construct a retraction $r : P_S \times [0,1] \to P_S$ that contracts $\xi(\Delta_S)$ to a point as follows. Fix $i \in S$ and denote by $\delta_{\{i\}}$ the unique point in $\Delta_{\{i\}}$. For every $x \in \xi(\Delta_S)$ define

$$r(x,t) = \xi(t\delta_{\{i\}} + (1-t)h_S^{-1}(x)).$$

In particular, $r(\partial P_S, \cdot)$ contracts ∂P_S (the "boundary" of P_S) to point $b_i \in P_S$. Consequently, $\xi(\Delta_S) = P_S$. ■

Remark 4. This proof does not go through if we consider *weighted average rules* (defined in Section 5) instead of *closed rules*, for the following reason. Let $\{\mu_p\}_{p\in\Delta(N)}$ be the class of weighted average rules. Then function $\xi : \Delta(N) \to X$ defined above

[22] For every $p \in \Delta_S$, one can think of p as a direction in $\mathbb{R}^n_+$. Maximizing payoff vector $U_N(x)$ in direction p yields point $h(p) = \arg\max_{x\in U_N(X)} p \cdot x$ on the Pareto frontier for the experts in S. Function h is well defined, since the payoff functions are strictly convex and bounded from above, and it is continuous by the Maximum theorem. Moreover, since $S \subset N'$, for every point x in the Pareto set P_S there exists a unique direction $p \in \Delta_S$ such that $h(p) = x$, i.e., h is an *onto* mapping. Thus $h_S : \Delta_S \to P_S$ is a homeomorphism.

is not continuous (and hence Step 2 does not hold), since $F^{\mu_p,0}$ is not uniformly continuous in p. To see this, fix expert i and consider a sequence $\{p^k\}$ with $p^k \to \bar{p}$ as $k \to \infty$ such that $p_i^k > 0$ for all k and $\bar{p}_i = 0$. Then at $\bar{p}$ expert i is dummy, so $F_i^{\mu_{\bar{p}},0}(m_i, m_{-i}) = u_i(\sum_{j\neq i} \bar{p}_j m_j - b_i)$ is constant w.r.t. m_i. However, for all k, $F_i^{\mu_{p^k},0}(m_i, m_{-i}) = u_i(\sum_{j\in N} p_j^k m_j - b_i) = u_i(0)$ if

$$m_i = \frac{1}{p_i^k}\left(b_i - \sum\nolimits_{j\neq i} p_j^k m_j\right)$$

Hence for every $\bar{m}_{-i}$ such that $\sum_{j\neq i} \bar{p}_j \bar{m}_j \neq b_i$ and every k we have

$$\max_{m_i\in Y}\left|F_i^{\mu_{p^k},0}(m_i, \bar{m}_{-i}) - F_i^{\mu_{\bar{p}},0}(m_i, \bar{m}_{-i})\right| \geq \left|u_i(0) - u_i\left(\sum\nolimits_{j\neq i} \bar{p}_j m_j - b_i\right)\right| > 0.$$

A.2 Proof of Proposition 3

By (3), (4), and concavity of the experts' payoff functions, the Nash bargaining solution will choose the tuple of proposals $m^* = (x-\theta, x-\theta, \ldots, x-\theta)$, where outcome x maximizes

$$\prod\nolimits_{i\in N}\left(-||x-b_i||^2 + \sum\nolimits_{j\in N} p_j ||b_i - b_j||^2\right).$$

Differentiating the logarithm of the above expression w.r.t. each coordinate of $x = (x_1, \ldots, x_{\mathbf{d}})$ yields the first order conditions

$$\sum\nolimits_{i\in N} \frac{-2(x_k - b_{ik})}{-||x-b_i||^2 + \sum_{j\in N} p_j||b_i-b_j||^2} = 0, \quad k = 1, \ldots, \mathbf{d}.$$

We need to find probabilities p_j such the bargaining outcome x is equal to the first best outcome, $x = 0$. Evaluating the above first order condition at $x = 0$ and dividing the numerator and the denominator of each summand by $||b_i||$ yields

$$\sum\nolimits_{i\in N} \frac{2b_{ik}/||b_i||}{\left(-||b_i||^2 + \sum_{j\in N} p_j||b_i-b_j||^2\right)/||b_i||} = 0, \quad k = 1, \ldots, \mathbf{d}.$$

Since $\sum_{i\in N} b_{ik}/||b_i|| = 0$ by assumption (C2), we need to find probabilities p_j such that the denominator is constant for all i, i.e., there exists a constant K such that

$$\frac{1}{||b_i||}\left(-||b_i||^2 + \sum\nolimits_{j\in N} p_j||b_i - b_j||^2\right) = K \quad \text{for all } i \in N.$$

For every $j \neq i$ we have

$$||b_i - b_j||^2 = ||b_i||^2 + ||b_j||^2 - 2||b_i|| \cdot ||b_j|| \cos\alpha = ||b_i||^2 + ||b_j||^2 + \frac{2||b_i|| \cdot ||b_j||}{n-1},$$

since $\cos\alpha = -1/(n-1)$ by assumption (C2). Hence

$$\begin{aligned} K &= -||b_i|| + \sum\nolimits_{j\neq i} p_j \left(||b_i|| + \frac{||b_j||^2}{||b_i||} + \frac{2||b_j||}{n-1} \right) \\ &= -\left(2 + \frac{2}{n-1}\right) ||b_i|| p_i + \sum\nolimits_{j\in N} p_j \left(\frac{||b_j||^2}{||b_i||} + \frac{2||b_j||}{n-1} \right) \\ &= -\frac{2n}{n-1} ||b_i|| p_i + \frac{1}{||b_i||} \sum\nolimits_{j\in N} ||b_j||^2 p_j + \sum\nolimits_{j\in N} \frac{2||b_j||}{n-1} p_j. \end{aligned}$$

Denote $K' = K - \sum_{j\in N} \frac{2||b_j||}{n-1} p_j$ and $L = \sum_{j\in N} ||b_j||^2 p_j$. Then

$$\frac{2n}{n-1} ||b_i||^2 p_i = -K' ||b_i|| + L. \tag{7}$$

Summing up (7) over $i \in N$ yields

$$\frac{2n}{n-1} L = -K' \sum\nolimits_{i\in N} ||b_i|| + nL.$$

Solving for K' yields

$$K' = \left(n - \frac{2n}{n-1}\right) L \left(\sum\nolimits_{i\in N} ||b_i||\right)^{-1} = \frac{n-3}{n-1} L \left(\frac{1}{n} \sum\nolimits_{i\in N} ||b_i||\right)^{-1}$$

After substitution of K' into (7) and dividing both sides by $||b_i^2||L$ we obtain

$$\frac{2n}{(n-1)L} p_i = -\frac{K'}{||b_i||L} + \frac{1}{||b_i||^2} = -\frac{1}{||b_i||} \frac{n-3}{n-1} \left(\frac{1}{n} \sum\nolimits_{i\in N} ||b_i||\right)^{-1} + \frac{1}{||b_i||^2}.$$

It is now straightforward to verify the above is satisfied by setting $p_i = w_i / \sum_{j\in N} w_j$ with w_i's as defined in Proposition 3.

A.3 Proof of Proposition 4

Let the biases be opposing, w.l.o.g. $b_1 \leq 0 \leq b_2$, and assume that a first best continuous deterministic rule exists. (For like biases a first best rule does not exist by Proposition 1.) Denote that rule by μ and define

$$m_i^*(m_j) = \arg\max_{m_i \in Y} u_i(\mu(m_i, m_j) + \theta - b_i), \quad \text{and}$$
$$\hat{m}_i \in \arg\max_{m_i \in Y} u_i(\mu(m_i, m_j^*(m_i)) + \theta - b_i).$$

Suppose that if the experts disagree, they move *sequentially*, in random order. So, with order $\{i, j\}$, in equilibrium expert i chooses $\hat{m}_i$ and expert j responds by $m_j^*(\hat{m}_i)$. Let us now show that each expert i is always weakly better off to be a second-mover. Specifically, we will prove that

$$b_1 - \theta \leq \mu(m_1^*(\hat{m}_2), \hat{m}_2) \leq \mu(\hat{m}_1, m_2^*(\hat{m}_1)) \leq b_2 - \theta. \tag{8}$$

Fix order $\{1, 2\}$ and let us show that

$$b_1 - \theta \leq \mu(\hat{m}_1, m_2^*(\hat{m}_1)) \leq b_2 - \theta. \tag{9}$$

First, μ is first best by assumption, so $\mu(Y, Y) = Y$ (see the argument in the proof of Proposition 1). Hence, there exists a pair $(\tilde{m}_1, \tilde{m}_2)$ such that $\mu(\tilde{m}_1, \tilde{m}_2) = b_2 - \theta$. That is, expert 1 (who moves first) can achieve the payoff $u_1(b_2 - b_1)$ by proposing $\tilde{m}_1$ (then the optimal response of expert 2 is $m^*(\tilde{m}_1) = \tilde{m}_2$ leading to her most preferred outcome). As expert 1's payoff $u_i(x - b_1)$ is strictly decreasing in x for $x > b_1$, it follows that $\mu(\hat{m}_1, m_2^*(\hat{m}_1)) \leq b_2 - \theta$.

Second, since μ is continuous and its image on Y is closed by assumption, $\mu(m_1, Y)$ is a closed (possibly, unbounded) interval. So $\mu(m_1, m_2^*(m_1))$ is either $b_2 - \theta$ or the endpoint of interval $\mu(m_1, Y)$ that is closest $b_2 - \theta$. Hence, again by continuity of μ, $\mu(Y, m_2^*(Y))$ is a closed interval. If $b_1 - \theta \in \mu(Y, m_2^*(Y))$, then expert 1 can achieve her most preferred point, $\mu(\hat{m}_1, m_2^*(\hat{m}_1)) = b_1 - \theta$. Otherwise, as we have obtained $b_2 - \theta \in \mu(Y, m_2^*(Y))$, it follows that $b_1 - \theta \notin \mu(Y, m_2^*(Y))$ entails $(-\infty, b_1 - \theta] \cap \mu(Y, m_2^*(Y)) = \varnothing$. In either case, $b_1 - \theta \leq \mu(\hat{m}_1, m_2^*(\hat{m}_1))$.

Thus we have shown (9). Then (8) follows immediately by the observation that $u_1(y+$

$\theta - b_1)$ is strictly decreasing and $u_2(y + \theta - b_2)$ is strictly increasing in y for $b_1 - \theta \leq y \leq b_2 - \theta$ and the fact that maxmin never exceeds minmax.

Now, equipped with (8), we consider two cases.

Case 1. For all $\theta \in \Theta$, the second mover obtains her most preferred outcome, i.e., $\mu(m_1, m_2^*(m_1)) = b_2 - \theta$ and $\mu(m_1^*(m_2), m_2) = b_1 - \theta$ for all m_1, m_2. Then the disagreement outcome is the equal probability lottery between outcomes b_1 and b_2 that uniquely defines the bargaining outcome that is, in general, different from the first best outcome.

Case 2. There exists an expert, e.g., expert 1, state θ and action $\hat{m}_1$ such that $\mu(\hat{m}_1, m_2^*(\hat{m}_1)) < b_2 - \theta$. Then at state $\theta' = b_1 - \mu(\hat{m}_1, m_2^*(\hat{m}_1))$, expert 1 can achieve her most preferred outcome when she is first mover, $\mu(\hat{m}_1, m_2^*(\hat{m}_1)) = b_1 - \theta'$. But by (8),

$$b_1 - \theta' \leq \mu(m_1^*(\hat{m}_2), \hat{m}_2) \leq \mu(\hat{m}_1, m_2^*(\hat{m}_1)) = b_1 - \theta'.$$

That is to say, at state θ' expert 1 achieves her most preferred outcome with certainty under disagreement. Consequently, the only implementable outcome is $x = b_1$, which is, in general, different from the first best outcome.

Finally, suppose that the experts move *simultaneously* after a disagreement, so their actions (m_1, m_2) must constitute a Nash equilibrium, $m_1 = m_1^*(m_2)$ and $m_2 = m_2^*(m_1)$.

Then in *Case 2*, the pair $(m_1, m_2) = (\hat{m}_1, m_2^*(\hat{m}_1))$ is a Nash equilibrium at state $\theta' = b_1 - \mu(\hat{m}_1, m_2^*(\hat{m}_1))$, since expert 1 attains his most preferred outcome and expert 2 plays a best reply. So, at that state the only implementable outcome is $x = b_1$.

In *Case 1*, Nash equilibrium in pure strategies does not exist, since for every m_i, best reply $m_j^*(m_i)$ leads to j's most preferred outcome, $\mu(m_i, m_j^*(m_i)) = b_j - \theta$. Nash equilibrium in mixed strategies does not exist either, since the payoff functions of the experts are strictly concave, so the best reply functions are single valued.

References

[1] R. Alonso and N. Matouschek. Optimal delegation. *Rev. Econ. Stud.* 75 (2008), 259–293.

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

[3] D. Austen-Smith. Information transmission in debate. *Am. J. Polit. Sci.* 34 (1990), 124–152.

[4] J. S. Banks and J. Duggan. A bargaining model of collective choice. *Am. Polit. Sci. Rev.* 94 (2000), 73–88.

[5] J. S. Banks and J. Duggan. A general bargaining model of legislative policy-making. *Q. J. Polit. Sci.* 1 (2006), 49–85.

[6] D. P. Baron. Legislative organization with informational committees. *Am. J. Polit. Sci.* 44 (2000), 485–505.

[7] D. P. Baron and J. A. Ferejohn. Bargaining in legislatures. *Am. Polit. Sci. Rev.* 83 (1989), 1181–1206.

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

[9] K. Binmore. Perfect equilibria in bargaining models. In K. Binmore and P. Dasgupta, editors, *The Economics of Bargaining*. Basil Blackwell, Oxford, 1987.

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

[11] W. Dessein. Authority and communication in organizations. *Rev. Econ. Stud.* 69 (2002), 811–838.

[12] A. Frankel. Delegating multiple decisions. Mimeo, 2011.

[13] T. Gilligan and K. Krehbiel. Asymetric information and legislative rules with a heterogeneous committee. *Am. J. Polit. Sci.* 33 (1989), 459–490.

[14] S. Hart and A. Mas-Colell. Bargaining and value. *Econometrica* 64 (1996), 357–380.

[15] B. Holmström. *On Incentives and Control in Organizations*. PhD thesis, Stanford University, 1977.

[16] B. Holmström. On the theory of delegation. In M. Boyer and R. E. Kihlstrom, editors, *Bayesian Models in Economic Theory*, pages 115–141. North-Holland, 1984.

[17] M. O. Jackson and B. Moselle. Coalition and party formation in a legislative voting game. *J. Econ. Theory* 103 (2002), 49–87.

[18] E. Kalai. Nonsymmetric Nash solutions and replications of 2-person bargaining. *Int. J. Game Theory* 6 (1977), 129–133.

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

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

[21] F. Koessler and D. Martimort. Optimal delegation with multi-dimensional decisions. *J. Econ. Theory* 147 (2012), 18501881.

[22] V. Krishna and J. Morgan. Asymmetric information and legislative rules: some amendments. *Am. Polit. Sci. Rev.* 95 (2001), 435–452.

[23] V. Krishna and J. Morgan. A model of expertise. *Q. J. Econ.* 116 (2001), 747–775.

[24] V. Krishna and R. Serrano. Multilateral bargaining. *Rev. Econ. Stud.* 63 (1996), 61–80.

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

[26] T. Lensberg. Stability and the Nash solution. *J. Econ. Theory* 45 (1988), 330–341.

[27] D. Martimort and A. Semenov. The informational effects of competition and collusion in legislative politics. *J. Pub. Econ.* 92 (2008), 1541–1563.

[28] N. D. Melumad and T. Shibano. Communication in settings with no transfers. *RAND J. Econ.* 22 (1991), 173–198.

[29] T. Mylovanov. Veto-based delegation. *J. Econ. Theory* 138 (2008), 297–307.

[30] T. Mylovanov and A. Zapechelnyuk. Decision rules revealing commonly known events. *Econ. Letters* 119 (2013), 8–10.

[31] T. Mylovanov and A. Zapechelnyuk. Optimal arbitration. *Int. Econ. Rev.*, forthcoming.

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

[33] A. Okada. A noncooperative coalitional bargaining game with random proposers. *Games Econ. Behav.* 16 (1996), 97–108.

[34] A. Okada. The Nash bargaining solution in general n-person cooperative games. *J. Econ. Theory* 145 (2010), 2356–2379.

[35] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica* 50 (1982), 97–109.

[36] A. Wolinsky. Eliciting information from multiple experts. *Games Econ. Behav.* 41 (2002), 141–160.