Who to Listen to?: A Model of Endogenous Delegation*

William Fuchs[†] Satoshi Fukuda[‡] Mahyar Sefidgaran[§]

July 19, 2022

Abstract

Two privately-informed agents must take a joint action without resorting to side-payments. Size and location of the support of each agent's private types (their preferred action) determine the degree of conflict. Under high conflict, it is too costly to elicit agents' information, which leads to an optimal constant allocation. Delegation arises endogenously when there is conflict and asymmetry in the amount of private information. The agent with more private information dictates the allocation within some bounds. When supports overlap information is shared and sometimes ex-post inefficient actions are optimally taken. Welfare relative to the first-best is non-monotone in conflict.

1 Introduction

In many situations, a pair of agents must take a common joint action. Examples include managers of different divisions within a firm deciding over the characteristics of a new product, members of a customs or monetary union deciding over common tariffs or monetary policy or parties in a political coalition deciding on a common political platform. Naturally, the preferred actions for each agent might be their private information. How would they jointly determine the optimal action when utility is non-transferable?

We study this as an ex-ante mechanism design problem where agents only care about the distance of the action taken from their preferred actions (uniformly distributed),

*We would like to thank Manuel Amador, Francesco Nava, Marco Ottaviani, and Thomas Wiseman for their insightful comments. Fuchs gratefully acknowledges the support from ERC Grant 681575, Grant PGC2018-096159-B-I00 financed by MCIN/AEI/10.13039/501100011033, and Comunidad de Madrid (Spain), Grant EPUC3M11 (V PRICIT) and Grant H2019/HUM-5891.

[†]McCombs School of Business, UT Austin, Universidad Carlos III Madrid, and CEPR.

[‡]Department of Decision Sciences and IGIER, Bocconi University.

[§]McCombs School of Business, UT Austin.

there are no transfers, and possibly differing Pareto weights are put on agents. Suppose, as members of Mercosur (a customs union), Argentina and Brazil must decide on a common external tariff for steel. Assume Argentina's preferred tariff (type) is known to lie in $[0, 0.1]$ and Brazil's in $[0.9, 1]$. Each country, conditional on only knowing its own type and the distribution of the other country's type must be induced to reveal truthfully. The distribution of types need not be symmetric. This asymmetry allows us to capture two natural elements of this problem. The first is the ex-ante notion of conflict. For example, if Argentina's type is in $[0, 0.1]$ and Brazil's in $[0.9, 1]$ it is a very different situation than if both countries share the same support. The second is the sense in which there can be more uncertainty about one country's type relative to the other. For example, suppose Argentina's type is in $[0, 0.01]$ while Brazil's is in $[0, 1]$.

With significant ex-ante conflict, the cost of eliciting truthful revelation is so high that the optimal allocation is independent of actual realizations. More interestingly, when there is sufficient asymmetry in the degree of private information and no overlap in the type sets, the optimal allocation endogenously assigns all decision rights to the agent with the larger support. That is, the agent with more private information can take its preferred action up to a cap. Thus, we have a model in which delegation arises endogenously.

If there is some overlap on the supports of both agents' type sets then the optimal allocation will elicit information from both agents. Yet, there will be a limit on how much information can be extracted from certain realized types. Thus, the optimal action may ex-post be constant over a subset of types. Also, similarly to the ex-ante general properties, we may have regions in which the action is only locally responsive to one of the agents' reports. To illustrate, suppose Argentina's preferred tariff is in $[0, 0.75]$ and Brazil's is in $[0.25, 1]$. For all of Argentina's types in $[0, 0.25]$, the country's representative has certainty that Brazil's optimal tariff is above. It is very hard to provide incentives for truthful revelation because they have a strong incentive to lie down. Thus the allocation will simply be unresponsive to Argentina's realized type below 0.25. Instead, for the types in $(0.25, 0.75]$ Argentina's representative must contemplate that Brazil's realized type could be above or below it. This possibility relaxes the incentive-compatibility (IC) constraint and allows the allocation to locally depend on her report.

Lastly, even if stochastic allocations or money burning are not used, the action will still be ex-post inefficient for some type pairs. This implies a second-order cost but it has a first-order effect on the IC constraints and can thus be beneficial.

Related Literature

Our paper is related to the following three strands of literature: (i) communication, (ii) delegation, and (iii) joint decisions without side-payments.

First, within the literature on communication, our paper is related to Alonso, Dessein, and Matouschek (2008a,b). They compare centralized and decentralized coordination when agents have private information and communicate strategically. Goltsman, Hörner, Pavlov, and Squintani (2009) study the value of commitment and stochastic allocations in a principle-agent setting in which the agent has private information. Our contribution is to study how the ex-ante degree of conflict (varying range of private information) affects optimal communication mechanisms. We also show that even without commitment one can improve over the allocation characterized by Alonso, Dessein, and Matouschek (2008a) by allowing multiple rounds of communication (see Proposition 2).

Second, there has been an extensive literature on delegation, pioneered by Holmström (1977, 1984) and subsequently studied by Melumad and Shibano (1991), Dessein (2002), Alonso and Matouschek (2008), Kováč and Mylovanov (2009), Amador and Bagwell (2013, 2020), Ambrus and Egorov (2017), and Kattwinkel (2019) among others.

Gan, Hu, and Weng (2022) consider a similar environment except that the principal can suggest two actions for two agents. They restrict their analysis to deterministic ex-post IC allocations and provide sufficient conditions under which interval delegation is optimal.[1] Unlike our model, both agents have discretion so there is no sense of endogenous delegations. In addition, they do not consider Bayesian incentive compatibility nor conflict.

Garfagnini, Ottaviani, and Sørensen (2014) compare delegation to other organizational structures in a binary decision setting and show that when the principal has access to additional information delegation is dominated by cheap talk communication.

Although the setting of our paper is technically different from the canonical delegation problem in which an uninformed principal delegates control to an informed agent, we show that one agent optimally and endogenously delegates decision rights to the other agent. To the best of our knowledge, this is the first paper to endogenize delegation as an optimal allocation within an organization without a predetermined hierarchy.

Third, and closely related to those above, are the papers studying joint decisions without resorting to transfers. The work on this area has focused mostly on the common support case. Moulin (1980) characterizes ex-post IC allocations for a collective action

[1] Alonso, Brocas, and Carrillo (2014) study a resource allocation in the brain as a mechanism design problem in which the principal can decide multiple actions for multiple agents without transfers.

when agents have single-peaked preferences. Martimort and Semenov (2008) study an optimal continuous ex-post incentive-compatible allocation when a biased principal faces two agents. Myerson (1979) formulates a Bayesian collective-decision problem without transfers and its Lagrangian when agents' type sets are finite. Carrasco and Fuchs (2008) and Fleckinger (2008) study an interesting but suboptimal Bayesian IC allocation for two agents taking a joint action. Our contribution is to study an optimal allocation with a particular focus on the implications of different degrees of conflict on the structure of the allocation and, in particular, which agent's information, if any, is considered. See Section 4.2 for a detailed discussion.

An exception to the common support assumption is the work by Börgers and Postl (2009). They assume that the agents' ordinal rankings over three possible options are known and diametrically opposed. The agents' types are given by how close their value is from the compromise option to each agents' best option. They share our interest in understanding when both agents' private information can be used to determine the optimal allocation. Our environment allows for a richer parametrization of ex-ante conflict and also the possibility of delegation which is not discussed in their paper.

Jackson and Sonnenschein (2007) and Carrasco, Fuchs, and Fukuda (2019) study the repeated version of the case of the symmetric fully-overlapping supports without side-payments. We discuss these papers in more detail in Section 4.2. Although, as they show, repeated interactions can help reduce inefficiency, we show that finding a partner with a lower amount of ex-ante conflict can be more important from a welfare perspective than finding one with whom repeated interactions are likely. Lara (2019) is an exception. It considers an asymmetric case in which one of the agents' type is binary while the other is drawn from a continuum. Furthermore, the two possible realizations of the binary type are symmetric around the average type of the other agent. Instead, we have both agents drawn from a continuum and restrict types to be in an interval.

2 Model

Two agents, denoted by $N := \{1, 2\}$, take a common action a without using monetary transfers. Each agent's utility is a quadratic loss function between her favorite action θ_i and the common action, i.e., $-(a - \theta_i)^2$. Each agent's favorite action is her private information. Specifically, suppose that each agent i's favorite action is drawn uniformly from her type set $\Theta_i := [\underline{\theta}_i, \overline{\theta}_i]$ with $0 = \underline{\theta}_1 < \overline{\theta}_1 \leq 1$ and $0 \leq \underline{\theta}_2 < \overline{\theta}_2 = 1$. Define $\Theta := \Theta_1 \times \Theta_2$ as the set of type profiles. The common action a is chosen from a sufficiently large compact interval $\mathcal{A}$ which includes $[0, 1]$.

An *allocation* is a measurable function $a : \Theta \to \mathcal{A}$ which associates, with a profile of announcements $\theta \in \Theta$, the corresponding action $a(\theta) \in \mathcal{A}$. The two agents are committed to an allocation, which maximizes the weighted sum of their (ex-ante) expected utilities subject to the incentive-compatibility and monotonicity constraints to be discussed below.

Our model can capture two types of situations. In the first, the two agents are designing the optimal allocation in the absence of a planner. In this case, the weights $\gamma \in (0, 1)$ and $1 - \gamma$ represent each agent's bargaining power in the design stage. Alternatively, we can capture situations in which there is an actual principal or planner that puts different weights on the payoffs of each agent.[2]

The allocation a satisfies agent i's (Bayesian) incentive-compatibility (IC) constraint if she has an incentive to tell the truth when the opponent is expected to be truthful:

$$-\mathbb{E}_{\theta_{-i}}\left[(a(\theta_i, \theta_{-i}) - \theta_i)^2\right] \geq -\mathbb{E}_{\theta_{-i}}\left[\left(a(\hat{\theta}_i, \theta_{-i}) - \theta_i\right)^2\right] \text{ for all } \theta_i, \hat{\theta}_i \in \Theta_i. \qquad \text{(IC}i\text{)}$$

The allocation is *incentive-compatible* if it satisfies each agent i's IC constraint.

The allocation a satisfies Monotonicity if

$$a \text{ is non-decreasing in each } \theta_i \in \Theta. \qquad \text{(Mon)}$$

Thus, the optimal allocation solves:

$$\max_{a(\cdot)\in\mathcal{A}} -\gamma\mathbb{E}_\theta\left[(a(\theta) - \theta_1)^2\right] - (1-\gamma)\mathbb{E}_\theta\left[(a(\theta) - \theta_2)^2\right] \qquad \text{(OBJ)}$$
$$\text{subject to (IC1), (IC2), and (Mon).}$$

The Bayesian IC constraint (ICi) can be decomposed into the local IC constraint and the expected monotonicity constraint. To see this, define agent i's expected utility from an allocation a when her type is θ_i and she announces $\hat{\theta}_i$:

$$U_i(\theta_i, \hat{\theta}_i) := -\mathbb{E}_{\theta_{-i}}\left[(a(\hat{\theta}_i, \theta_{-i}) - \theta_i)^2\right].$$

Define $U_i(\theta_i) := U_i(\theta_i, \theta_i)$. Now, the IC constraint (ICi) is decomposed into (i) the local IC constraint: for all $\theta_i \in \Theta_i$,

$$U_i(\theta_i) = U_i(\underline{\theta}_i) + 2\int_{\underline{\theta}_i}^{\theta_i}\left(\mathbb{E}_{\theta_{-i}}\left[a(\tau, \theta_{-i})\right] - \tau\right) d\tau; \qquad \text{(LIC}i\text{)}$$

[2]As suggested by Alonso, Dessein, and Matouschek (2008b,a) it could be the CEO of a company that puts different weights based on the profitability of each division.

and (ii) the expected monotonicity constraint:

$$\mathbb{E}_{\theta_{-i}}\left[a(\theta_i, \theta_{-i})\right] \text{ is non-decreasing in } \theta_i. \qquad \text{(Ex-Mon}i\text{)}$$

Clearly, the monotonicity constraint (Mon) is stronger than (Ex-Mon$i$) (for each $i \in N$). Thus, we solve (OBJ) under (LICi) and (Mon).

Using (Mon) simplifies our analyses and has a negligent effect on welfare. Below we discuss this in detail. In Online Appendix, we show that our main result is robust to allowing for stochastic allocations and money burning. Although explicit money burning is not used, Proposition 1 shows that some optimal actions entail ex-post inefficiency.

2.1 Model Discussion

2.1.1 Monotonicity

We impose (Mon) because the use of (Mon) rather than (Ex-Mon$i$) (for each $i \in N$) significantly reduces the difficulty of characterizing certain properties of the optimal allocation. Naturally, the reader might question what happens when one does not impose it. For the commonly studied case of common supports $[0, 1]$, this assumption is actually not binding. For the partially overlapping case we have studied this issue numerically and in our simulations (such as the ones illustrated in Figures 3 and the left-most panel of 5) there are few violations and when we have found them they are of the order of 10^{-3}. For the non-overlapping case one can analytically show that the assumption is binding yet again our simulations show that such violations are also of the order 10^{-3}. Thus, there is no relevant impact on welfare from making this assumption.

It is also important to point out that requiring Bayesian incentive compatibility and (Mon) is not equivalent to assuming ex-post incentive compatibility. This is clearly illustrated in Section 4.1, where we characterize the optimal ex-post IC allocation and show that the welfare differences between the two can be significant. The allocations are also quite different as can be seen by contrasting the left-most panel of Figure 5, the optimal Bayesian IC allocation, which respects (Mon), and Figure 3 in the Online Appendix, the optimal ex-post IC allocation.

2.1.2 Stochastic Allocation and Money Burning

In our model we restrict the allocation to be deterministic and furthermore we rule out the use of money burning. Within the delegation literature, Kováč and Mylovanov (2009), Amador and Bagwell (2013, 2020), and Ambrus and Egorov (2017) among others, have studied the use of stochastic allocations or money burning. In particular, they provide conditions under which money burning or stochastic allocations might be useful and when they are not. Given our use of quadratic preferences (which is common in this literature), agents essentially care about the expected allocation and conditional variance (see, for instance, Goltsman, Hörner, Pavlov, and Squintani, 2009). In particular, the variance term enters just like money burning into the agent's utility function. Thus, if we cannot achieve a better outcome by allowing money burning, stochastic allocations would not help either. When we have a complete characterization of the allocation as in the main results in Theorem 1, we can solve the relaxed problem allowing for money burning and verify analytically that it is not used (see Online Appendix). Thus, our main result about the endogenous delegation is robust to this extension.[3]

For the special case in which the agents' type sets coincide, even though we do not have a complete characterization of the optimal allocation, using the Lagrange multipliers, we can still analytically show that money burning is not used (see Online Appendix). For the case in which the agents' type sets partially overlap, our numerical simulations suggest that money burning is never used.

Finally, note that, even though money burning is not used, Proposition 1 shows that certain actions exhibit ex-post inefficiency.

2.1.3 General Preferences and Distributions

Our analysis is restricted to the uniform-quadratic case, as is common in the literature.[4] Unfortunately, without transferable utility, the problem becomes much harder since the optimization cannot be done type by type as in standard mechanism design. Moreover, an allocation may depend on two agents' information and we employ Bayesian incentive compatibility.[5] That being said, our results are robust to small perturbations of preference or the shape of the distribution. Indeed, the main economic forces behind Theorem 1 suggest that the result is quite robust but unfortunately it is not straightforward to establish this formally given our reliance on the Lagrangian method. What is slightly relevant is knowing that the support is in a given interval and thus being able to rule our certain types for a given agent. We discuss this at length in Remark 1.

[3]For the special case in which one of the agents' type sets is degenerate we can also use the results of Amador and Bagwell (2013) to establish that money burning is not used.

[4]See, for instance, Alonso, Dessein, and Matouschek (2008a,b), Ambrus and Egorov (2017), Goltsman, Hörner, Pavlov, and Squintani (2009), and Martimort and Semenov (2008).

[5]The analysis is simpler with ex-post incentive compatibility even if we allow for general distributions. Indeed, Gan, Hu, and Weng (2022) restrict attention to ex-post incentive compatibility and thus can consider general distributions. For the single agent case, there is no distinction between ex-post incentive compatibility and Bayesian incentive compatibility, thus again allowing for general analysis. See, for example, Amador and Bagwell (2013, 2020), Alonso and Matouschek (2008), and Kováč and Mylovanov (2009) among others.

3 Endogenous Delegation

We present our main theorem. It establishes when the planner will forgo the elicitation of information, when it will delegate the decision to one of the agents or when both types will be consulted to determine the action.

Theorem 1. *For each* $(\overline{\theta}_1, \underline{\theta}_2)$*, the optimal allocation* a^* *satisfies the following.*

1. *Suppose* $\underline{\theta}_2 \geq \overline{\theta}_1$. *If* $\underline{\theta}_2 \geq \frac{2-\gamma}{1-\gamma}\overline{\theta}_1 - 1$ *and* $\underline{\theta}_2 \geq \frac{\gamma}{1+\gamma}\overline{\theta}_1 + \frac{1-\gamma}{1+\gamma}$*, then the optimal allocation is the ex-ante optimal constant allocation:* $a^*(\cdot) = \frac{\gamma\overline{\theta}_1 + (1-\gamma)(1+\underline{\theta}_2)}{2}$*. The optimal allocation does not elicit agents' information.*

2. *Suppose* $\underline{\theta}_2 \geq \overline{\theta}_1$. *If* $\underline{\theta}_2 \leq \frac{2-\gamma}{1-\gamma}\overline{\theta}_1 - 1$*, then the optimal allocation is the (constrained) delegation allocation for agent* 1*:*

$$a^*(\theta_1, \theta_2) = \begin{cases} \frac{1-\gamma}{2-\gamma}(1+\underline{\theta}_2) & \text{if } \theta_1 \in \left[\underline{\theta}_1, \frac{1-\gamma}{2-\gamma}(1+\underline{\theta}_2)\right] \\ \theta_1 & \text{if } \theta_1 \in \left[\frac{1-\gamma}{2-\gamma}(1+\underline{\theta}_2), \overline{\theta}_1\right] \end{cases}.$$

   The optimal allocation does not elicit agent 2*'s information.*

3. *Suppose* $\underline{\theta}_2 \geq \overline{\theta}_1$. *If* $\underline{\theta}_2 \leq \frac{\gamma}{1+\gamma}\overline{\theta}_1 + \frac{1-\gamma}{1+\gamma}$*, then the optimal allocation is the (constrained) delegation allocation for agent* 2*:*

$$a^*(\theta_1, \theta_2) = \begin{cases} \theta_2 & \text{if } \theta_2 \in \left[\underline{\theta}_2, \frac{\gamma\overline{\theta}_1 + 1 - \gamma}{1+\gamma}\right] \\ \frac{\gamma\overline{\theta}_1 + 1 - \gamma}{1+\gamma} & \text{if } \theta_2 \in \left[\frac{\gamma\overline{\theta}_1 + 1 - \gamma}{1+\gamma}, \overline{\theta}_2\right] \end{cases}.$$

   The optimal allocation does not elicit agent 1*'s information.*

4. *If* $\underline{\theta}_2 < \overline{\theta}_1$*, then the optimal allocation* a^* *depends on both types* (θ_1, θ_2).

Figure 1 depicts Theorem 1: it illustrates, as $\overline{\theta}_1$ and $\underline{\theta}_2$ vary (i.e., as the agents' ranges of private information vary), when the optimal allocation elicits no information (i.e., the ex-ante optimal constant allocation is optimal); when the (constrained) delegation allocation for each agent is optimal; and when the optimal allocation elicits both agents' information.

The proof relies on the Lagrangian method: for the candidate optimal allocation in each part, we show that the candidate allocation is indeed optimal as we show there exist Lagrange multipliers for the IC constraints under which the Lagrangian is concave in allocations and the candidate allocation is derived from the first-order conditions. As the verification of the Lagrange multipliers is long, computationally intense, and does not add much economic insights, we provide a condensed proof in the Appendix. We direct the interested reader to Online Appendix for detailed derivations.

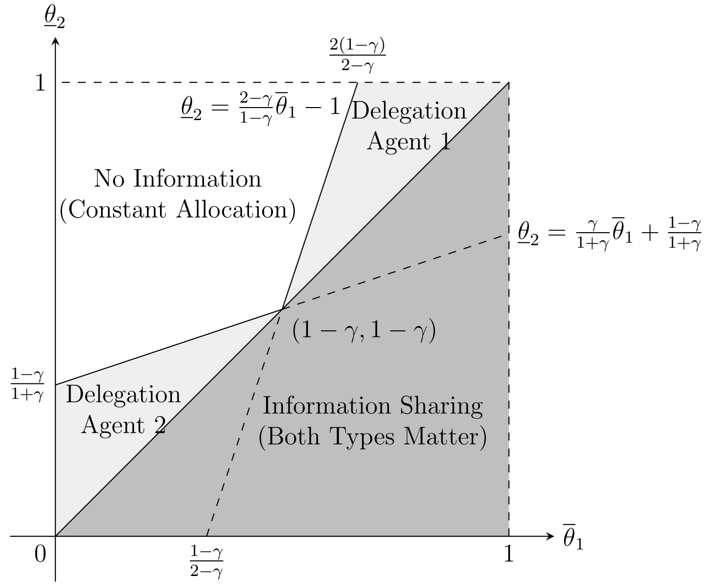

Figure 1: Classification of an Optimal Allocation

Here, we provide a brief overview of its structure. For Part (1), the proof consists of four steps. The first step rewrites the relaxed problem by substituting the local IC constraints into the objective function. The second step formulates the Lagrangian. Denote by Λ_i the Lagrange multiplier (which is a function of bounded variation) associated with agent i's local IC constraint. The third step examines the first-order conditions. The fourth step substitutes $a^*(\cdot) = \gamma\mathbb{E}_{\theta_1}[\theta_1] + (1-\gamma)\mathbb{E}_{\theta_2}[\theta_2]$ and find the Lagrange multipliers (Λ_1, Λ_2) with the properties that the Lagrangian is concave in allocations and that the first-order conditions lead to the constant allocation a^*.

The proof of Part (2) is similar. We consider the following relaxed problem: for each agent i, the local IC constraint is imposed; and the allocation is required to be monotonic in agent 2's types given agent 1's types. We start with the Lagrangian which incorporates agents' IC constraints. Denote by Λ_i the Lagrange multiplier associated with agent i's local IC constraint. Then, the problem is to maximize the Lagrangian subject to the (relaxed) monotonicity constraint. Now, we explicitly incorporate the (relaxed) monotonicity constraint using the Lagrangian approach again. Denote by B the Lagrange multiplier associated with the monotonicity constraint on agent 2's types θ_2 (for any given θ_1). With this set-up, we formulate the first-order conditions, and we find the multipliers $(\Lambda_1, \Lambda_2, B)$ under which the Lagrangian is concave in allocations and under which the first-order conditions are met for the delegation allocation. For the proof of Part (3), exchange the role of agents 1 and 2.

Finally, for Part (4), we note first that the best allocation which depends on at most one agent's information is a constrained delegation allocation with possibly two caps. Since this is a (continuous) ex-post IC allocation, to show that we can do better using both agents' information, it suffices to show that it is not even ex-post optimal, i.e., that the optimal ex-post IC allocation depends on both agents' types.[6] Thus, we characterize the optimal ex-post IC allocation and show it indeed depends on both agents' types when their type sets overlap, establishing the result.

Our proof method has some similarity with Amador and Bagwell (2013, 2020) in that we guess the candidate optimal allocation and verify that it is indeed correct. We do so by showing that there exist Lagrange multipliers under which the candidate allocation is indeed optimal. However, the Lagrange multipliers that we find are different from those of Amador and Bagwell (2013, 2020) especially because the set of feasible allocations, which in principle depend on both agents' types and which has to satisfy the weaker IC notion of Bayesian IC constraints, is larger. To verify that the candidate allocation a^* is optimal, we also need to jointly and simultaneously determine both agents' multipliers. In addition, when a constrained delegation solution is a candidate optimal allocation, the Lagrange multiplier B on (Mon) depends on both agents' types.

Next, we provide an incomplete yet intuitive argument for the structure of an optimal allocation when the agents' type sets do not overlap (i.e., $\overline{\theta}_1 \leq \underline{\theta}_2$), as illustrated in Figure 2. This case highlights the situation in which there exists an ex-ante high degree of conflict, and it is too costly to make an allocation dependent on both agents' information. Section 4 studies the case in which the agents' type sets overlap and the optimal allocation elicits both agents' information.

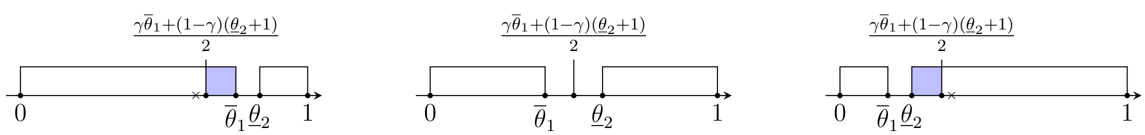

Figure 2: Classification of an Optimal Allocation: Delegation Allocation for Agent 1 (Left), Constant Allocation (Middle), and Delegation Allocation for Agent 2 (Right).

A planner, who has no access to the agents' information, would choose the ex-ante optimal constant allocation, which is the weighted average of the agents' average types: $\gamma\mathbb{E}_{\theta_1}[\theta_1] + (1-\gamma)\mathbb{E}_{\theta_2}[\theta_2]$. The central panel of Figure 2 depicts the case in which this allocation falls into the "gap" of the agents' disjoint type sets (precisely, $\overline{\theta}_1 \leq \gamma\mathbb{E}_{\theta_1}[\theta_1] + (1-\gamma)\mathbb{E}_{\theta_2}[\theta_2] \leq \underline{\theta}_2$). To get an intuition behind why the constant allocation is optimal, start at the ex-ante optimal constant allocation and consider making the allocation dependent on the agents' types. Ideally, when agent 1's type is

[6]Note that, in our problem, an allocation a is ex-post IC if $-(a(\theta_i,\theta_{-i})-\theta_i)^2 \geq -(a(\hat{\theta}_i,\theta_{-i})-\theta_i)^2$ for all $i, \theta_i, \hat{\theta}_i, \theta_{-i}$.

close to 0, the planner would move the action further to the left relative to when agent 1's type is close to $\overline{\theta}_1$. But if so, agent 1 would want to exaggerate downward. Since the allocation is always above agent 1's types, if the allocation depends on agent 1's types then the only incentive-compatible modification under monotonicity is to make the allocation agent 1's favorite actions. However, this is very costly for agent 2, thus inefficient.

When the constant allocation falls into one of the agents' type sets, as depicted in the left and right panels of Figure 2, we can improve upon the constant allocation. To see this, suppose it falls into agent 1's type set. Giving decision rights to type θ_1 between the ex-ante optimal constant and $\overline{\theta}_1$ (the shaded region) is a Pareto improvement. As a result, on average, the allocation moves towards right. Since agent 2 is relatively further away, this change has a larger impact on the payoff. Thus, it is optimal to compensate by granting more discretion to agent 1, i.e., setting the cutoff ("$\times$") below the ex-ante constant allocation. Since the allocation is always below agent 2's types, like in the case above, it is too costly to use agent 2's private information. Thus, the optimal allocation is the constrained delegation for agent 1.

Remark 1. However small a probability of overlap of the agents' type sets is, the allocation will make use of both agents' information. To see this more formally, suppose that with probability $1-\delta$ the agents' type sets are as usual but with probability δ their type sets are the whole interval $[0,1]$. Now, the planner can improve upon the constant allocation using an allocation that elicits information from both agents, for example, by using the median of both agents' reported types and the best constant. This allocation is ex-post IC and strictly better than the original constant however small δ is.[7]

Note that, although the allocation would rely on both agents' information, its effect on welfare would be negligible for small δ. In this sense, this is in contrast to the bargaining literature on the Coase conjecture, where the informed party's type support plays an important role in terms of welfare/outcomes. Thus, our environment is more robust in this sense.

3.1 Effect of Asymmetric Supports and Pareto Weights on Delegation

In many business situations, CEOs delegate certain decisions to unique division managers despite them having implications for other divisions. The model highlights that the guiding principle combines the relative importance for the profitability of each division together with the extent of private information.

[7] Notice that this also implies that Theorem 1 (4) would hold for general distributions on types.

We focus first on the relative extent of private information, and thus we assume $\gamma = \frac{1}{2}$. In this case, the key observation is that the agent needs sufficiently larger private information to be designated as the constrained delegate as reflected in Figure 1. Starting from the 45 degrees line, $\overline{\theta}_1 = \underline{\theta}_2$, which represents a situation of relatively low conflict, we observe that the minimal difference in private information is sufficient to grant some discretion to one of the agents. As conflict increases, i.e., as we move north west in the figure, we need a larger informational asymmetry in order to observe delegation. In all cases, the agent with the largest uncertainty becomes the constrained delegate. Also, the range of discretion given by the cutoff decreases as the degree of conflict increases.

Next, we focus on the effect of Pareto weights. As the importance of profitability of one division increases, the manager of the division will be more likely to have the discretion over joint decisions. Furthermore, the range of discretion increases.

4 Information Sharing

As shown in Part (4) of Theorem 1, the optimal allocation incorporates both agents' information whenever the two agents' type sets overlap. Since we do not have a closed-form characterization of the optimal allocation in this case, it is hard to analytically provide properties of the optimal allocation. Yet, we show that the optimal allocation generally exhibits ex-post inefficiency: ex post, an action associated with high types is even higher than the agents' types; and an action associated with low types is even lower. We establish the result without requiring Monotonicity.

Proposition 1. *Let $\overline{\theta}_1 = 1 - \varepsilon$ and $\underline{\theta}_2 = \varepsilon$ for some $\varepsilon \in [0, \frac{1}{2})$, and let $\gamma = \frac{1}{2}$.*[8] *Let a^* be an optimal allocation. For any $\theta_1 \in [\underline{\theta}_2, \overline{\theta}_1]$, we have:*

$$\begin{cases} a^*(\theta_1, \theta_1) < \theta_1 & \text{if } \theta_1 \in \left(\underline{\theta}_2, \frac{1}{2}\right) \\ a^*(\theta_1, \theta_1) > \theta_1 & \text{if } \theta_1 \in \left(\frac{1}{2}, \overline{\theta}_1\right) \\ a^*(\theta_1, \theta_1) = \theta_1 & \text{if } \theta_1 = \frac{1}{2} \end{cases}.$$

If $\overline{\theta}_1 = 1$ and $\underline{\theta}_2 = 0$, then $a^*(0, 0) = 0$ and $a^*(1, 1) = 1$.

The formal proof is available in Online Appendix, in which we establish the result by expressing the optimal allocation using the Lagrange multipliers. Here, we discuss the outline of the proof, which consists of two steps. In the first step, denoting by Λ_i the Lagrange multiplier associated with agent i's local IC constraint, we can express

[8] The results are qualitatively robust to asymmetric supports and unequal weights.

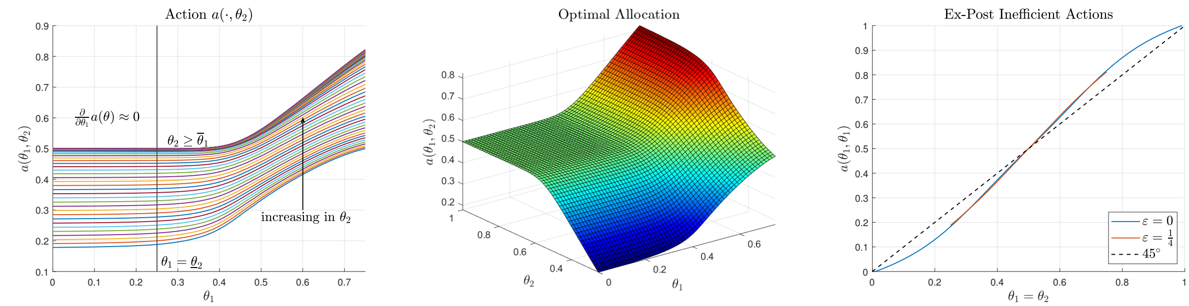

Figure 3: Illustration of the Optimal Allocation, when $\overline{\theta}_1 = 1 - \varepsilon$ and $\underline{\theta}_2 = \varepsilon$ for some $\varepsilon \in [0, \frac{1}{2})$ and $\gamma = \frac{1}{2}$. The left panel illustrates the action $(\cdot, \theta_2)$ for each θ_2 with $\varepsilon = \frac{1}{4}$. The central panel depicts the optimal allocation a^* with $\varepsilon = \frac{1}{4}$. The right panel illustrates $a^*(\theta_1, \theta_1)$ for $\theta_1 \in [\varepsilon, 1 - \varepsilon]$ for $\varepsilon \in \{0, \frac{1}{4}\}$.

the amount of ex-post efficiency as

$$a(\theta_1, \theta_1) - \theta_1 = -\frac{\Lambda_1(\theta_1) - \Lambda_1(1 - \theta_1)}{\frac{1}{1-\varepsilon} - \lambda_1(\theta_1) - \lambda_1(1 - \theta_1)},$$

where λ_1 is the derivative of Λ_1. As in the proof of Theorem 1, this part of the proof involves a lot of algebra. In the second step, although we do not have a closed-form solution for the multiplier, we can still show that the amount of ex-post efficiency is negative when $\theta_1 < \frac{1}{2}$ and is positive when $\theta_1 > \frac{1}{2}$. The right panel of Figure 3 depicts the actions $a(\theta_1, \theta_1)$ for $\varepsilon \in \{0, \frac{1}{4}\}$. When ε increases, while the first-term in the denominator increases, the IC constraints of types close to $\overline{\theta}_1$ and $\underline{\theta}_2$ become stronger (i.e., the Lagrange multipliers are increasing in epsilon). For the latter, this is because with higher probability the agent knows that the other agent lies in the particular direction. Thus, the value/need of distorting actions to achieve the incentive compatibility is increasing in ε. Overall, the amount of ex-post inefficiency appears quite insensitive to the degree of conflict measured by ε, as is observed in the right panel of Figure 3. Note that some actions are ex-post inefficient even when $\theta_1 \neq \theta_2$.

The intuition is best conveyed in a symmetric example with four types: $\{0.1, 0.4, 0.6, 0.9\}$. Since, on average, they expect the other agent to be on the other side of $\frac{1}{2}$, the binding constraints will be the central types wanting to pretend to be extreme. Now consider perturbing the action after both announce 0.1 to a slightly lower value. If both types had been truthfully reporting, this generates an inefficiency but since it is a small distance from their preferred action this only implies a second-order loss. On the other hand, type 0.4 which was very tempted to report 0.1 is now very concerned because if the other agent happens to be type 0.1 the action will be even further away from her preferred action. This has a first-order effect on the IC constraint and can thus be beneficial.

In practice, it may be unlikely to observe ex-post inefficient actions. Thus, it is important to note that Theorem 1 continues to hold if one rules out ex-post inefficient actions.

To gather some additional insight of the properties of the optimal allocation, consider the left and central panels of Figure 3, where we illustrate the optimal allocation for $\varepsilon = \frac{1}{4}$. It is important to highlight the differences between the ex-post types that know that the other agent always lies in a particular direction, i.e., $[0, \underline{\theta}_2]$ for agent 1 and $[\overline{\theta}_1, 1]$ for agent 2, and those types that do not know whether the other agent is above or below. Call a former type a *high-conflict* type. In the illustrated example in Figure 3, this corresponds to θ_1 in $[0, \frac{1}{4}]$ and θ_2 in $[\frac{3}{4}, 1]$. Since it is very hard to provide incentives for the high conflict types to report truthfully, the optimal allocation becomes insensitive to their types when they fall in this region even under Bayesian IC constraints. As a result, $a(\theta)$ is constant in θ_1 for $\theta_1 \leq \frac{1}{4}$ and similarly it is constant in θ_2 for $\theta_2 \geq \frac{3}{4}$.[9] This implies that when $\theta \in [0, \frac{1}{4}] \times [\frac{3}{4}, 1]$, the allocation is approximately constant. Furthermore, by symmetry the constant must equal $\frac{1}{2}$.

As can be observed in Figure 5, even when $\varepsilon = 0$ (and thus there are no high-conflict types), it is still the case that it is very hard to provide incentives for types close to the extremes which understand that the other agent most likely lies in one side of them. As a result, again the optimal allocation responds to this by making the allocation less sensitive to their types. Thus, although not strictly constant, the allocation is observed to be very close to $\frac{1}{2}$ when θ is in the vicinity of $(0, 1)$ and $(1, 0)$.

4.1 Welfare Effect of Increasing Conflict

Here, we ask how the social welfare changes with respect to the first-best as the degree of conflict changes for $\gamma = \frac{1}{2}$. To measure the degree of conflict, we consider the following parametrization of type sets: $\Theta_1 = [0, 1 - \varepsilon]$ and $\Theta_2 = [\varepsilon, 1]$ with $\varepsilon \in [0, \frac{1}{2}]$. When $\varepsilon = 0$, the two agents' type sets coincide. When $\varepsilon = \frac{1}{2}$, they do not overlap. Thus, the degree of conflict is increasing in ε. For reference, the parameter configuration we study corresponds to the anti-diagonal line from $(\frac{1}{2}, \frac{1}{2})$ (i.e., $\varepsilon = \frac{1}{2}$) to $(1, 0)$ (i.e., $\varepsilon = 0$) in Figure 1.

We study the welfare effect on conflict by measuring the relative social welfare loss

$$\frac{V - V_{\mathrm{FB}}}{V_{\mathrm{FB}}},$$

where V denotes the social welfare associated with an optimal allocation and V_{FB} the one associated with the first-best allocation (the weighted sum of the agents' types)

[9] More generally, we observed numerically that even though sometimes derivatives of the allocation are not exactly zero, the optimal allocation is almost insensitive to agents' reports.

for each ε. We normalize the welfare loss $V - V_{\text{FB}}$ by the first-best social welfare V_{FB} because V_{FB} itself changes as the degree of conflict ε changes.

As highlighted before, in order to get information sharing it is important that types do not know whether they are above or below the opponent's. Thus, we compare the social welfare under Bayesian and ex-post IC allocations. Specially, we consider the phantom voter allocation (Moulin, 1980): The joint action is the median of the reported types θ_1 and θ_2 and a "phantom" at $\frac{1}{2}$. As we show in Online Appendix, this is the optimal continuous ex-post IC allocation.

As illustrated in Figure 4, under the ex-post IC constraints, the relative social welfare loss is monotonically increasing in ε. In contrast, under the Bayesian IC constraints, the relative social welfare loss is not monotonic in ε. Intuitively, starting from $\varepsilon = \frac{1}{2}$, as conflict decreases, the first-best welfare rapidly increases. In contrast, due to the tight IC constraints, the social welfare increases at a slower rate under both Bayesian and ex-post IC constraints when ε is close to $\frac{1}{2}$. Thus, both relative measures decrease. When the degree of conflict is small enough, the fact that a large fraction of types are not high-conflict types allows for more efficient information sharing under the Bayesian IC constraints, and thus the social welfare catches up relative to the first-best. This cannot be achieved under the ex-post IC constraints since, in this case, all types are effectively high-conflict types. This results in a significant 20% point difference in welfare loss when $\varepsilon = 0$.[10]

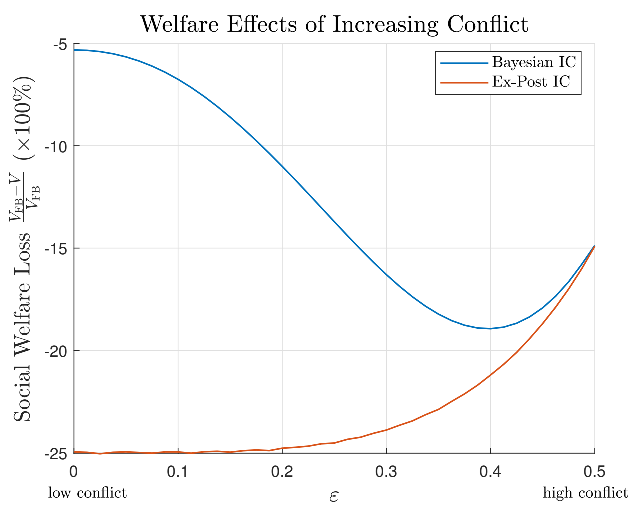

Figure 4: Welfare Effects of Increasing Conflict under Bayesian and Ex-Post IC Constraints

[10]This difference highlights the fact that requiring (Mon) and Bayesian incentive compatibility is quite different from Ex-post incentive compatibility.

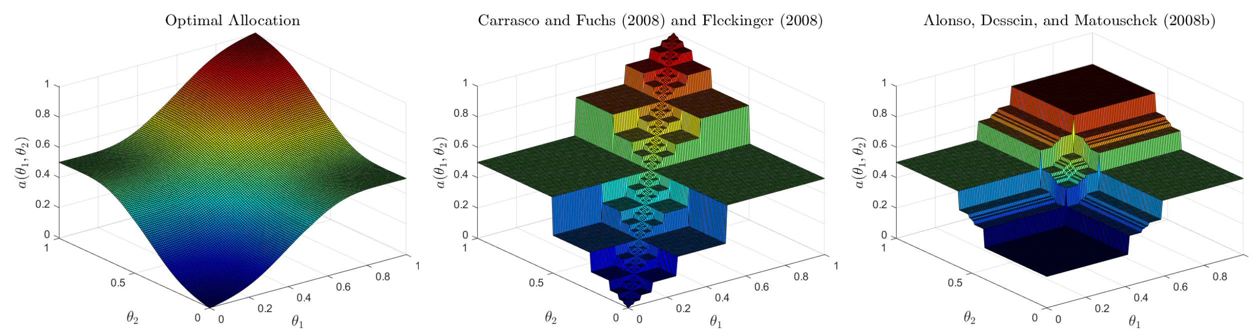

Figure 5: Allocations for the Fully-Overlapping Case

4.2 Fully Overlapping Case

Lastly, to relate to some prior literature, we consider the fully-overlapping case $\Theta = [0,1]^2$ with symmetric weights. The left panel of Figure 5 depicts the optimal allocation. The central panel depicts the allocation suggested by Carrasco and Fuchs (2008) and Fleckinger (2008), and the panel on the right depicts the special case of Alonso, Dessein, and Matouschek (2008b) in which both agents must coordinate on an action. Related to the discussions in Section 3, note that the three allocations share the property that common actions are insensitive to the agents' reports close to extremes and set to $\frac{1}{2}$.

In terms of welfare, the allocations in the center and right panels deliver the exact same value $-\frac{1}{21}$ which is strictly below that of the optimal allocation depicted on the left panel. While the optimal allocation is approximately 5% inefficient, the other two are about 14% inefficient. The allocation by Alonso, Dessein, and Matouschek (2008b) can be obtained without commitment from the planner with just one round of communication. The allocation in the center can be implemented dynamically by a series of simple binary partitions of the type set. The fact that information is conveyed dynamically also allows this allocation to be implemented without commitment. Yet, even within the set of allocations without commitment neither of these are optimal.

A better communication protocol can be obtained by modifying the dynamic process as follows. In the first stage the planner asks each agent which of these four intervals they belong to: $[0, \frac{1}{2} - c]$, $[\frac{1}{2} - c, \frac{1}{2}]$, $[\frac{1}{2}, \frac{1}{2} + c]$, and $[\frac{1}{2} + c, 1]$. If the initial partition reports do not coincide the principal takes the ex-post efficient allocation given this information. If they do coincide, the mechanism then follows the dynamic process described above with just two partitions. One can show that setting c to 0.0321 this communication protocol is IC and more efficient than either of the two described above.[11] Formally:

Proposition 2. *One round of communication is not optimal (even without commit-*

[11]More formally, c is the unique solution to $48c^3 - 36c^2 - 30c + 1 = 0$ in $[0, \frac{1}{2}]$.

ment).

As the proposition establishes, the restriction to one round of communication in Alonso, Dessein, and Matouschek (2008a) is with loss. This is contrary to what was incorrectly claimed by Carrasco and Fuchs (2008).

Finally, the work by Jackson and Sonnenschein (2007) and Carrasco, Fuchs, and Fukuda (2019) shows how, in the fully-overlapping case, repeated interactions can help mitigate the conflict.[12] Both mechanisms are arbitrarily close to efficiency when the discounted number of decisions goes to infinity. Yet, Figure 4 shows that the relative efficiency gains from reducing conflict can be twice as large as those from repeated interactions in the low conflict ($\varepsilon = 0$) case. Thus, it might be more important to find a partner with a low level of conflict than one with which one interacts repeatedly.

5 Conclusion

Although very parsimonious, our model can help us shed light into when a firm or organization might assign exclusive decision rights to different members or divisions. Agents for whom the decision has a higher impact or for whom there is a larger ex-ante uncertainty of what they might prefer are more likely to be granted exclusive decision rights. In this way we provide a theory of endogenous delegation.

Our model also helps understand when a CEO would consult its division managers before making a decision versus act unilaterally. In particular, our model implies that if she needs to make decisions over which it is clear that two divisions would have clear diverging interests, then she would avoid a useless and conflicting meeting in which she would not be able to gather any relevant information and rather dictate the terms herself. Instead, when the situation is less conflictive, it will be possible to extract valuable information and make better decisions. For this, it is important that the agents do not know each other's realized preferences, since, as we illustrated in Figure 4, there is a large gain from only requiring Bayesian IC instead of ex-post IC.

Lastly, although it is natural that repeated interactions can improve welfare, our work suggests that, to the extent possible, it might be more relevant to properly create the teams or decision bodies to lower the degree of conflict. This can potentially, also in relation to the idea of organizational culture, be seen as either molding the agents' preferences or creating a more uniform organizational view of the world which helps reduce conflict.

[12] Jackson and Sonnenschein (2007) propose a clever 'budgeting mechanism' that restricts the agents to report their types in a way that is consistent with the true distribution of types. Instead, Carrasco, Fuchs, and Fukuda (2019) study the optimal mechanism which uses continuation values (i.e., a larger influence on future decisions) as a way to incentivize the agents to report truthfully.

References

Alonso, R., I. Brocas, and J. D. Carrillo (2014): "Resource Allocation in the Brain," *Review of Economic Studies*, 81, 501–534.

Alonso, R., W. Dessein, and N. Matouschek (2008a): "Centralization Versus Decentralization: An Application to Price Setting by a Multi-Market Firm," *Journal of the European Economic Association*, 6, 457–467.

——— (2008b): "When Does Coordination Require Centralization?" *American Economic Review*, 98, 145–179.

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

Amador, M. and K. Bagwell (2013): "The Theory of Optimal Delegation With an Application to Tariff Caps," *Econometrica*, 81, 1541–1599.

——— (2020): "Money Burning in the Theory of Delegation," *Games and Economic Behavior*, 112, 382–412.

Ambrus, A. and G. Egorov (2017): "Delegation and Nonmonetary Incentives," *Journal of Economic Theory*, 171, 101–135.

Börgers, T. and P. Postl (2009): "Efficient Compromising," *Journal of Economic Theory*, 144, 2057–2076.

Carrasco, V. and W. Fuchs (2008): "Dividing and Discarding: A Procedure for Taking Decisions with Non-transferable Utility," in *2008 Meeting Papers 315, Society for Economic Dynamics.*

Carrasco, V., W. Fuchs, and S. Fukuda (2019): "From Equals to Despots: The Dynamics of Repeated Decision Making in Partnerships with Private Information," *Journal of Economic Theory*, 182, 402–432.

Dessein, W. (2002): "Authority and Communication in Organizations," *Review of Economic Studies*, 69, 811–838.

Fleckinger, P. (2008): "Bayesian Improvement of the Phantom Voters Rule: An Example of Dichotomic Communication," *Mathematical Social Sciences*, 55, 1–13.

Gan, T., J. Hu, and X. Weng (2022): "Optimal Contingent Delegation," Working paper.

Garfagnini, U., M. Ottaviani, and P. N. Sørensen (2014): "Accept or reject? An Organizational Perspective," *International Journal of Industrial Organization*, 34, 66–74.

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

Holmström, B. (1977): "On Incentives and Control in Organizations," Phd thesis, Stanford University.

——— (1984): "On the Theory of Delegation," in *Bayesian Models in Economic Theory*, ed. by M. Boyer and R. Kihlstrom, North-Holland, 115–141.

Jackson, M. O. and H. F. Sonnenschein (2007): "Overcoming Incentive Constraints by Linking Decisions," *Econometrica*, 75, 241–257.

Kattwinkel, D. (2019): "Allocation with Correlated Information: Too good to be true," Working paper.

Kováč, E. and T. Mylovanov (2009): "Stochastic Mechanisms in Settings without Monetary Transfers: The Regular Case," *Journal of Economic Theory*, 144, 1373–1395.

Lara, L. P. d. (2019): "Delegated Coordination," Master's thesis, Fundação Getulio Vargas, Escola de Pós-Graduação em Economia.

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," *RAND Journal of Economics*, 22, 173–198.

Moulin, H. (1980): "On Strategy-proofness and Single Peakedness," *Public Choice*, 35, 437–455.

Myerson, R. B. (1979): "Incentive Compatibility and the Bargaining Problem," *Econometrica*, 47, 61–73.

A Proof of Theorem 1

This appendix provides a condensed proof of Theorem 1. We direct the interested reader to Online Appendix for detailed derivations.

A.1 Part (1)

Step 1. Consider the relaxed problem in which the monotonicity constraint is ignored. Thus, the problem is to maximize the sum of the agents' ex-ante utilities subject to their local IC constraints. For agent 1, let the "reference" type of the local IC constraint be $\overline{\theta}_1$. For agent 2, let the "reference" type be $\underline{\theta}_2$. After some algebra, the relaxed problem can be rewritten as follows:

$$\max_{a(\cdot)} \gamma \left(U_1(\overline{\theta}_1) - 2\mathbb{E}_\theta \left[(a(\theta) - \theta_1)(\theta_1 - \underline{\theta}_1)\right]\right) + (1-\gamma)\left(U_2(\underline{\theta}_2) + 2\mathbb{E}_\theta \left[(a(\theta) - \theta_2)(\overline{\theta}_2 - \theta_2)\right]\right)$$

subject to $U_1(\theta_1) = U_1(\overline{\theta}_1) - 2\int_{\theta_1}^{\overline{\theta}_1} \mathbb{E}_{\theta_2}\left[a(\tau_1, \theta_2) - \tau_1\right] d\tau_1$ for each $\theta_1 \in [\underline{\theta}_1, \overline{\theta}_1]$ and

$$U_2(\theta_2) = U_2(\underline{\theta}_2) + 2\int_{\underline{\theta}_2}^{\theta_2} \mathbb{E}_{\theta_1}\left[a(\theta_1, \tau_2) - \tau_2\right] d\tau_2 \text{ for each } \theta_2 \in [\underline{\theta}_2, \overline{\theta}_2].$$

Step 2. To formulate the Lagrangian of the problem formulated in Step 1, we denote by Λ_i the Lagrange multiplier associated with agent i's (local) IC constraint. Theoretically, the Lagrange multiplier Λ_i is a function of bounded variation. Without loss of generality, we normalize Λ_i by setting $\Lambda_1(\underline{\theta}_1) = 0$ and $\Lambda_2(\overline{\theta}_2) = 0$.

We conjecture and verify in Step 4 a specific functional form of Λ_i, from the first-order condition to be found in Step 3. In particular, the specific Λ_i is shown to have a density function λ_i on $[\underline{\theta}_i, \overline{\theta}_i]$. With this in mind, we define the Lagrangian as

$$\begin{aligned}\mathcal{L} :=& \gamma \left(U_1(\overline{\theta}_1) - 2\mathbb{E}_\theta \left[(a(\theta) - \theta_1)(\theta_1 - \underline{\theta}_1)\right]\right) + (1-\gamma)\left(U_2(\underline{\theta}_2) + 2\mathbb{E}_\theta \left[(a(\theta) - \theta_2)(\overline{\theta}_2 - \theta_2)\right]\right) \\ &+ \int_{\underline{\theta}_1}^{\overline{\theta}_1} \left(U_1(\theta_1) - U_1(\overline{\theta}_1) + 2\int_{\theta_1}^{\overline{\theta}_1} \mathbb{E}_{\theta_2}\left[a(\tau_1, \theta_2) - \tau_1\right] d\tau_1\right) d\Lambda_1(\theta_1) \\ &+ \int_{\underline{\theta}_2}^{\overline{\theta}_2} \left(U_2(\theta_2) - U_2(\underline{\theta}_2) - 2\int_{\underline{\theta}_2}^{\theta_2} \mathbb{E}_{\theta_1}\left[a(\theta_1, \tau_2) - \tau_2\right] d\tau_2\right) d\Lambda_2(\theta_2).\end{aligned}$$

After some algebra, the Lagrangian can be rewritten as:

$$\begin{aligned}\mathcal{L} =& -\mathbb{E}_{\theta_2}\left[(a(\overline{\theta}_1, \theta_2) - \overline{\theta}_1)^2\right](\gamma - \Lambda_1(\overline{\theta}_1)) - 2\gamma\mathbb{E}_\theta\left[(a(\theta) - \theta_1)(\theta_1 - \underline{\theta}_1)\right] \\ &- \int_{\underline{\theta}_1}^{\overline{\theta}_1} \mathbb{E}_{\theta_2}\left[(a(\theta) - \theta_1)^2\right]\lambda_1(\theta_1) d\theta_1 + 2\int_{\underline{\theta}_1}^{\overline{\theta}_1} \mathbb{E}_{\theta_2}\left[a(\theta) - \theta_1\right]\Lambda_1(\theta_1) d\theta_1 \\ &- \mathbb{E}_{\theta_1}\left[(a(\theta_1, \underline{\theta}_2) - \underline{\theta}_2)^2\right](1 - \gamma + \Lambda_2(\underline{\theta}_2)) + 2(1-\gamma)\mathbb{E}_\theta\left[(a(\theta) - \theta_2)(\overline{\theta}_2 - \theta_2)\right] \\ &- \int_{\underline{\theta}_2}^{\overline{\theta}_2} \mathbb{E}_{\theta_1}\left[(a(\theta) - \theta_2)^2\right]\lambda_2(\theta_2) d\theta_2 + 2\int_{\underline{\theta}_2}^{\overline{\theta}_2} \mathbb{E}_{\theta_1}\left[a(\theta) - \theta_2\right]\Lambda_2(\theta_2) d\theta_2.\end{aligned}$$

Step 3. We take the point-wise first-order condition for each $a(\theta)$. For any θ, the first-order condition is

$$\frac{1}{\overline{\theta}_2-\underline{\theta}_2}\left\{\gamma\frac{\theta_1-\underline{\theta}_1}{\overline{\theta}_1-\underline{\theta}_1}+\lambda_1(\theta_1)(a(\theta)-\theta_1)-\Lambda_1(\theta_1)\right\}$$
$$+\frac{1}{\overline{\theta}_1-\underline{\theta}_1}\left\{-(1-\gamma)\frac{\overline{\theta}_2-\theta_2}{\overline{\theta}_2-\underline{\theta}_2}+\lambda_2(\theta_2)(a(\theta)-\theta_2)-\Lambda_2(\theta_2)\right\}$$
$$+\frac{(a(\overline{\theta}_1,\theta_2)-\overline{\theta}_1)(\gamma-\Lambda_1(\overline{\theta}_1))}{(\overline{\theta}_1-\underline{\theta}_1)(\overline{\theta}_2-\underline{\theta}_2)}\mathbb{I}(\theta_1=\overline{\theta}_1)+\frac{(a(\theta_1,\underline{\theta}_2)-\underline{\theta}_2)(1-\gamma+\Lambda_2(\underline{\theta}_2))}{(\overline{\theta}_1-\underline{\theta}_1)(\overline{\theta}_2-\underline{\theta}_2)}\mathbb{I}(\theta_2=\underline{\theta}_2)=0.$$

From now on, we find (Λ_1,Λ_2) such that the first-order conditions are satisfied at

$$a^*=\gamma\mathbb{E}_{\theta_1}[\theta_1]+(1-\gamma)\mathbb{E}_{\theta_2}[\theta_2].$$

For the rest of the proof, we need to consider the following four cases: (i) $\overline{\theta}_1<a^*<\underline{\theta}_2$; (ii) $\overline{\theta}_1=a^*<\underline{\theta}_2$; (iii) $\overline{\theta}_1<a^*=\underline{\theta}_2$; and (iv) $\overline{\theta}_1=a^*=\underline{\theta}_2(=1-\gamma)$. However, since the proof of each case is similar, here we only consider the first case.

Step 4. We conjecture and verify $\Lambda_1(\overline{\theta}_1)=\gamma$ and $\Lambda_2(\underline{\theta}_2)=-(1-\gamma)$. Then, by the first-order conditions, there exist constants α_1 and α_2 such that

$$\alpha_1=\gamma\frac{\theta_1-\underline{\theta}_1}{\overline{\theta}_1-\underline{\theta}_1}+\lambda_1(\theta_1)(a^*-\theta_1)-\Lambda_1(\theta_1), \tag{1}$$
$$\alpha_2=-(1-\gamma)\frac{\overline{\theta}_2-\theta_2}{\overline{\theta}_2-\underline{\theta}_2}+\lambda_2(\theta_2)(a^*-\theta_2)-\Lambda_2(\theta_2), \text{ and} \tag{2}$$
$$0=\frac{\alpha_1}{\overline{\theta}_2-\underline{\theta}_2}+\frac{\alpha_2}{\overline{\theta}_1-\underline{\theta}_1}.$$

Since Expressions (1) and (2) are a linear first-order differential equation, one can show that Λ_1 and Λ_2 are:

$$\Lambda_1(\theta_1)=\gamma\frac{(\theta_1-\underline{\theta}_1)(2a^*-\overline{\theta}_1-\theta_1)}{2(a^*-\theta_1)(\overline{\theta}_1-\underline{\theta}_1)}, \text{ and}$$
$$\Lambda_2(\theta_2)=(1-\gamma)\frac{(\overline{\theta}_2-\theta_2)(\theta_2+\underline{\theta}_2-2a^*)}{2(a^*-\theta_2)(\overline{\theta}_2-\underline{\theta}_2)}.$$

It can be verified that $\Lambda_1(\overline{\theta}_1)=\gamma$ and $\Lambda_2(\underline{\theta}_2)=-(1-\gamma)$.

Finally, it can be seen that, once we substitute Λ_1 and Λ_2, the Lagrangian is a concave function in a. By construction, the first-order conditions are satisfied at a^*. The proof is complete.

A.2 Part (2)

The proof is similar to that of Part (1) except that we consider the relaxed monotonicity constraint which requires the allocation to be monotonic in agent 2's types given agent 1's types. Here, we only report the Lagrange multipliers. Denote by k_1 agent 1's cap:

$$k_1 = \frac{1-\gamma}{2-\gamma}(\underline{\theta}_2 + 1).$$

For agent 1's local IC constraint, the multiplier Λ_1 is:

$$\Lambda_1(\theta_1) = \begin{cases} \dfrac{\gamma}{2}\dfrac{(\theta_1 - \underline{\theta}_1)}{(\overline{\theta}_1 - \underline{\theta}_1)} & \text{if } \theta_1 \in [0, k_1] \\ \dfrac{2\theta_1 - (1-\gamma)(1+\underline{\theta}_2)}{2(\overline{\theta}_1 - \underline{\theta}_1)} & \text{if } \theta_1 \in \left[k_1, \overline{\theta}_1\right] \end{cases}.$$

For agent 2's local IC constraint, the multiplier Λ_2 is:

$$\Lambda_2(\theta_2) = -\frac{(1-\gamma)(\overline{\theta}_2 - \theta_2)}{2(\overline{\theta}_2 - \underline{\theta}_2)}\left(1 + \frac{k_1 - \underline{\theta}_2}{k_1 - \theta_2}\right).$$

For the Monotonicity constraint, denote by B the Lagrange multiplier associated with the monotonicity constraint on agent 1's types θ_1 (for any given θ_2). Then,

$$B(\theta_1, \theta_2) = \begin{cases} 0 & \text{if } \theta_1 \in [0, k_1] \\ (1-\gamma)\dfrac{(\theta_2 - \underline{\theta}_2)(\theta_2 - 1)((1-\gamma)(\underline{\theta}_2 + 1) - (2-\gamma)\theta_1)}{(\overline{\theta}_1 - \underline{\theta}_1)(1 - \underline{\theta}_2)((1-\gamma)(\underline{\theta}_2 + 1) - (2-\gamma)\theta_2)} & \text{if } \theta_1 \in \left[k_1, \overline{\theta}_1\right] \end{cases}.$$

A.3 Part (3)

The proof is similar to that of Part (2), exchanging the role of agents 1 and 2.

A.4 Part (4)

The proof of Theorem 1 (4) is in two steps. First, if an optimal allocation depends on at most one agent's information then the optimal allocation satisfies ex-post IC constraints, including Monotonicity. In particular, the optimal allocation which depends on at most one agent's information is a constrained delegation solution, which is also continuous.

Second, below we find the unique optimal continuous ex-post IC allocation, which can depend on both agents' information. We show that this allocation depends on both agents' information and is better than the best allocation which depends on at most one agent's information.

The proof of this step is in four sub-steps. In the first sub-step, as Martimort and

Semenov (2008) characterize continuous and ex-post IC allocations in the uniform-quadratic setting (see also Moulin, 1980), also in our setting of Θ, an allocation a is ex-post IC and continuous if and only if

$$a(\theta_1, \theta_2) = \min\left(x, \max(\theta_1, y_1), \max(\theta_2, y_2), \max(\theta_1, \theta_2, z)\right)$$

for some (x, y_1, y_2, z) with $z \leq y_1, y_2 \leq x$.

In the second sub-step, for an optimal continuous ex-post IC allocation, x and z can be dropped. Thus, an optimal ex-post incentive-compatible and continuous allocation a satisfies

$$a(\theta_1, \theta_2) = \min\left(\max(\theta_1, y_1), \max(\theta_2, y_2), \max(\theta_1, \theta_2)\right) \text{ for some } (y_1, y_2).$$

Intuitively, if $x < \overline{\theta}_1$, then the allocation a does not use agent 1's information even when both agents' types lie in the common set $[x, \overline{\theta}_1]$, which is inefficient. Likewise, if $z > \underline{\theta}_2$, then the allocation a does not use agent 2's information even when both agents' types lie in the common set $[\underline{\theta}_2, z]$, which is inefficient.

In the third sub-step, after some algebra, we can find the optimal (y_1, y_2). Namely, the optimal ex-post incentive-compatible and continuous allocation a is

$$a(\theta_1, \theta_2) = \min\big(\max(\theta_1, 1-\gamma), \max(\theta_2, \gamma\overline{\theta}_1 + (1-\gamma)\underline{\theta}_2), \max(\theta_1, \theta_2)\big).$$

Now, by construction, the optimal continuous ex-post IC allocation a yields strictly better social welfare than the optimal allocation that incorporates at most one agent's information. The proof of Part (4) is complete.