Common agency with Informed Principals: Revelation Principle*

Nicolas Riquelme†

November 1, 2018

Click here for the latest version

Abstract

This paper studies games where a group of privately informed principals design mechanisms to a common agent. The agent has private information (exogenous) and, after observing principals' mechanisms, may have information (endogenous) about feasible allocations and private information from each principal. Thus, each principal may be interested in designing a mechanism to screen all this information, for which a potentially complicated message space to convey this information might be needed. In this paper, we provide sufficient conditions on the agent's payoff such that any equilibrium in this setup has an output-equivalent equilibrium using only mechanisms with simple message spaces (direct mechanisms). Depending on the conditions, we propose two different notions of direct mechanisms and discuss their applicability with some examples.

1 Introduction

Common agency problems suffer the lack of a revelation-principle type of result that allows simplifying their analysis. The problem is originated in the information the agent has at the moment of communication with each principal: Besides agent's exogenous private information (agent's type), the agent observes each principal's mechanism, which

*First draft: August 3, 2017. I am grateful to Srihari Govindan and Paulo Barelli for their guidance and encouragement. I would like to thank Zizhen Ma, Lucas Pahl, Emiel Awad and all participants at Theory seminars at University of Rochester, Midwest Economic Theory Conference 2017 at SMU and 28th Stony Brook Game Theory International Conference for helpful discussions and suggestions. All errors are my own.

†Department of Economics, University of Rochester, email: nriquelm@z.rochester.edu

may give her valuable information that other principals might try elicit using their mechanisms (screen). This information is known in the literature as endogenous information or market information[1]. In the context of privately informed principals, this last kind of information has two components: First, the agent observes which allocations each principal offers to her. Second, the way a principal provides a set of allocations may give information about his private information (principal's type) to the agent (signal). Thus, a message space for the agent to communicate (truthfully) with each principal must be rich enough to encapsulate all this information. In particular, any implemented outcome using direct mechanisms (in which the only information communicated is the agent's type), may have a principal with a strictly profitable deviation to a mechanism that elicits not only agent's type, but also endogenous information. Thus, mechanisms with potentially complicated message spaces are needed. In this paper, we propose sufficient conditions on the agent's payoffs such that we can restrict the analysis without loss of generality to mechanisms with simple message spaces. These conditions are simple to check and, importantly, do not impose any restriction on principals' payoffs. Moreover, these are satisfied in many relevant economic frameworks in the common agency literature. We propose two different notions of mechanisms. In the first one, principals only need the agent's type space as a message space. These mechanisms are called direct mechanism in the mechanism design literature and are the usual restriction when the agent has no endogenous information (one principal case[2]). We refer to these as *strong direct mechanisms*. The sufficient conditions that allow us to restrict to strong direct mechanisms in this context prevent the agent's preference over one principal's allocation to change with other principals' allocations and types. If we relax the conditions related to other principals' types, we have our second notion of mechanisms: Principals can use as a message space the product between agent's type space and other principals' type spaces. We refer to these as *weak direct mechanisms*. The last mechanisms are the natural generalization of the usual direct mechanisms to the present context in the sense that the message space encapsulates all the external information of the model from the perspective of one principal.[3]

The main contribution of the paper is to show that, under our assumptions, the problem of characterizing the implementable outcomes of a common agency model with privately informed principals can be decomposed in two steps: In the first step, we need to

[1] Epstein and Peters (1999)

[2] Myerson (1979)

[3] The agent and other principals' private information.

find incentive compatible direct mechanisms for each principal's type. In these mechanisms, the agent reports truthfully all her information she may have that the principal does not. This step is a **screening** one because the restriction on the mechanisms allows each principal type to screen the information the agent may have. In the second step, each principal type has to choose an incentive compatible direct mechanism. Thus, he may signal his type with his decision. This step is a **signalling** one, because the offered mechanism may signal a principal's private information. Both problems, screening and signaling, have been extensively studied separately in the literature. Because of this decomposition, we can leverage on these two pieces of literature to analyze this new framework where both arise naturally. Finally, we show two relevant economic frameworks that satisfy our assumptions and explicitly study using this new approach. We show that in order to implement some outcomes, it is enough to assume that each principal chooses a price or delegates a policy interval depending on the context considered.

The paper is organized as follows: Section 1.1 presents a brief literature review; Section 2 shows two examples that illustrate our main two results. Section 3 sets the model and the assumptions used in our results. Sections 4 presents our results. Section 5 extends the results to several more general environments. Section 6 discusses the consequences of our results showing how the mechanism design problem is simplified. Section 7 shows two economically relevant environments where our results are useful. We apply our results and show that there are equilibrium outcomes that can be implemented using simple mechanisms as prices or interval policy delegation respectively. Finally, section 8 concludes.

1.1 Literature Review

The lack of a revelation principle in common agency problems has been addressed in two different ways in case principals are not privately informed. The first one, initiated by Peters (2003) and Peters (2007), gives sufficient conditions on the agent and principal's payoffs such that in a common agency model where the agent does not have private information the revelation principle applies. In that context, a direct mechanism corresponds to a take-it-or-leave-it offer from each principal to the agent. Later, Attar et. al (2008) propose sufficient conditions for a setting where the agent has private information. A direct mechanism there corresponds to a map from the agent's type space to the allocations. The second way to address this problem is to focus on what the literature knows as menus. The idea here is that instead of focusing on the message space of the mechanism, we can equivalently look the range of possible outcomes that a mechanism allows the agent to pick. This idea was proposed by Peters (2001) and Martimort and Stole (2002), and it is known The Menu Theorem. In this context, principals offer menus (subsets of the allocation space) from which the agent pick his most preferred allocation. [4] Since principals strategies are now subsets, computing principals' best responses is not an easy task.

In the case principals have private information, Galperti (2015) shows that the same approach using menus can be applied if, besides the menus, each principal is allowed to send cheap-talk signals to the common agent. In contrast, in this paper, we consider a similar framework to Galperti (2015) but we follow the first approach.

2 Illustrative Examples

The following examples give intuition about why the revelation principle fails and illustrate the assumptions that we make to prevent a principal's deviations to mechanisms with complicated message spaces. In the first example, the problem is the agent's willingness to send to a principal a message that conveys information about other principals' feasible allocations. In the second one, the problem is the agent's willingness to send to a principal a message that conveys information about other principals' type. In each case, there is an equilibrium outcome where principals choose mechanisms that are not replicable when we restrict to direct mechanisms that only use as a message space the agent's exogenous type space. Intuitively, each principal may want to screen information about the other principal (either feasible allocations or type). However, since a direct mechanism allows the principal to only screen information about the exogenous type of the agent, that is not possible. For both cases, we show assumptions such that the agent is not willing to transmit other principals' information, and then the usual revelation principle applies. In the second example, we also propose a new set of direct mechanisms that allow capturing the possible information about other principals' types that the agent may have. This allows us to have a broader set of games where a still simple revelation principle holds.

[4] Attar et. al (2011) is an excellent example of how to use this method to analyze the multi-principal competition.

2.1 Example 1:

This example shows how the revelation principle fails when the agent is willing to signal to a principal the allocations that the other principals make available to her. The example is inspired by Attar et. al (2008) and Peters (2003), who showed how the revelation principle fails in an uninformed principal case. Suppose there are two manufacturers (principals) and one retailer (agent). Each manufacturer is looking to invest with the retailer. Suppose that manufacturer 1 privately observes the state of the economy which can take two equally likely values, t_h and t_l. Manufacturer 2 and the retailer do not have private information. Each manufacturer has two actions: Invest (I) or Not invest (NI) with the retailer. In the high state (t_h), both principals prefer jointly invest than jointly not invest. In the low state (t_l), the opposite is true. However, no matter the state, manufacturer 2 prefers not to invest. In the low state, since manufacturer 1 prefers not to invest, there is no problem of not investing jointly. However, in the high state, manufacturer 1 prefers to invest and would like manufacturer 2 to do the same. Depending on the action of each retailer, and the state of the economy, the payoffs of the manufacturers and the agent are the following:

t_h:

	I	NI
I	4,3,3	3,4,0
NI	2,1,2	1,2,1

t_l:

	I	NI
I	1,1,1	1,2,0
NI	2,1,2	2,2,1

Suppose each manufacturer design a mechanism. For that, assume there are message spaces $M_1 = M_2 = \{a, b\}$ that the retailer can use to communicate with manufacturer $i \in \{1, 2\}$ and the manufacturers can commit on how to take action as a function of these messages. More formally, first the manufacturers commit simultaneously to a mechanism $\pi_i : M_i \to \{I, NI\}$. Second, the retailer privately sends a message to each manufacturer $m = (m_1, m_2) \in M_1 \times M_2$. Together with the payoff matrices, these form a game that we call the indirect mechanisms game. A perfect Bayesian equilibrium (PBE) of the indirect mechanism game is the following. Manufacturer 1 selects $(\pi_1(t_h)(a), \pi_1(t_h)(b)) = (I, NI)$ and $\pi_1(t_l)(a) = \pi_1(t_l)(b) = NI$. Manufacturer 2 selects $\pi_2(a) = \pi_2(b) = I$. The outcome of this equilibrium is (I, I) on t_h and (NI, I) on t_l. Thus, even though manufacturer 2 prefers to not invest in both states, in equilibrium he invests. To make that happen, manufacturer 1 uses the messages available in a way that the retailer wants to transmit the information about what actions she will take with manufacturer 2: Whenever manufacturer 2 does not invest, retailer will send message

b and manufacturer 1 will want not to invest; and whenever manufacturer 2 invests, the agent will send message a and manufacturer 1 will invest. Since manufacturer 2 prefers jointly invest than jointly not to invest, and since in the case of not invest the retailer will send message b to manufacturer 1 inducing not to invest, manufacturer 2 will prefer to invest instead. The previous is only possible because the retailer is willing to transmit this information from one manufacturer to the other; another way to put it is that the retailer's preferences over manufacturer 1 actions are reversed depending on manufacturer 2's actions. Additionally, note that in the high state manufacturer 1 makes available the action to not invest to the agent even though in equilibrium manufacturer 1 invests. This is what the literature calls a latent contract.[5] In this example, manufacturer 1 makes available that action to the retailer only to deter manufacturer 2 not to invest.

Let's consider now an alternative game. In this game, the principals are restricted to offer only direct mechanisms with message space the agent's exogenous type space. Since the agent does not have private information, a direct mechanism is a mapping $\pi_i : M_i \to \{I, NI\}$ where M_i is a singleton. Thus, in this new game, denote $M_1 = M_2 = \{c\}$. Because of the previous feature, there is no strategic decision to the agent: Principals cannot screen agent's information and directly select an action. We call this game the direct mechanism game. The only equilibrium of the direct mechanism game is the following: Manufacturer 1 selects $\pi_1(t_h)(c) = I$ and $\pi_1(t_l)(c) = NI$. Manufacturer 2 selects $\pi_2(c) = NI$. The outcome of this equilibrium is (I, NI) on t_h and (NI, NI) on t_l. Thus, the lack of messages prevents manufacturer 1 to get information about manufacturer 2 actions. Equivalently, manufacturer 1 cannot deter manufacturer 2 not to invest. In sum, we have shown an equilibrium allocation on the indirect mechanism game that cannot be replicated in the direct mechanism game: the revelation principle fails.

Now consider the following new payoff matrices:

	I	NI
I	$4,3,\mathbf{2}$	$3,4,0$
NI	$2,1,\mathbf{3}$	$1,2,1$

t_h

	I	NI
I	$1,1,1$	$1,2,0$
NI	$2,1,2$	$2,2,1$

t_l

The only difference to the previous matrices is that on the high state, for every action of manufacturer 2, the agent always prefers manufacturer 1 not to invest. Thus, even

[5] A latent contract is an allocation that belongs to the range of a mechanism offered by a principal that is not selected in equilibrium

though principals' incentives do not change, now the agent is not willing to transmit information about manufacturer 2's actions. Because of this small change, now the previous equilibrium allocation cannot be implemented, and the new equilibrium of the indirect mechanism game is the following: Manufacturer 1 selects $\pi_1(t_h)(a) = \pi_1(t_h)(b) = I$ and $\pi_1(t_l)(a) = \pi_1(t_l)(b) = NI$. Manufacturer 2 selects $\pi_2(a) = \pi_2(b) = NI$. The outcome of this equilibrium is (I, NI) on t_h and (NI, NI) on t_l. Now, if we consider the direct mechanism game, the unique equilibrium is the same than in the previous part: Manufacturer 1 selects $\pi_1(t_h)(c) = I$ and $\pi_1(t_l)(c) = NI$. Manufacturer 2 selects $\pi_2(c) = NI$. The outcome of this equilibrium is (I, NI) on t_h and (NI, NI) on t_l. Thus, we have the same equilibrium outcome if we restrict to direct mechanisms so the revelation principle is now valid. From this example, we learn that it might be enough to assume that the agent does not want to transmit information about other principals' allocations for the revelation principle to be valid. As we will show in the second example, even after this assumption there is a second concern which we need to address to have the revelation principle for general environments.

2.2 Example 2:

Consider the same setup as before, but with the following payoff matrices:

	I	NI
I	4, 4, 3	1, 2, 2
NI	2, 1, 2	1, 1, 0

t_h

	I	NI
I	1, 1, 1	1, 2, 1
NI	2, 1, 1	4, 4, 2

t_l

Now, in each state both manufacturers would like to coordinate: in the high state, they are both better jointly investing, and in the low state they are both better jointly not investing. The problem is that only manufacturer 1 observes the state of the economy, which manufacturer 2 would like to know to coordinate their decisions. Note also that the retailer's payoffs are aligned with the manufacturers: She prefers manufacturers to coordinate in each state. An equilibrium of this indirect mechanism game is the following: Manufacturer 1 selects $\pi_1(t_l)(a) = \pi_1(t_l)(b) = NI$ and $\pi_1(t_h)(a) = \pi_1(t_h)(b) = I$. Manufacturer 2 selects $(\pi_2(a), \pi_2(b)) = (I, NI)$. The equilibrium outcome is (I, I) on t_h and (NI, NI) on t_l. Intuitively, principals use the agent to coordinate their actions. Since the agent's payoffs are aligned with the manufacturers, the agent is willing to transmit the information about the state of the economy from the manufacturer 1 to the manufacturer 2. In the mechanism that manufacturer 2 selects, a message a from the retailer to manufacturer 2 transmits the information that manufacturer 1 will invest, while a message b transmits the information that manufacturer 1 will no invest.

Let's consider now the alternative game where manufacturers are restricted to offer direct mechanisms. The only equilibrium of the direct mechanism game is the following: Manufacturer 1 selects $\pi_1(t_h)(c) = I$ and $\pi_1(t_l)(c) = NI$. Manufacturer 2 selects $\pi_2(c) = NI$. The outcome of this equilibrium is (I, NI) on t_h and (NI, NI) on t_l. Then, in this alternative framework, the retailer has no messages available to transmit information to the manufacturer 2 about manufacturer 1 type. Because of that, the outcome of the mechanism game cannot be replicated in the direct mechanism game, which is a failure of the revelation principle.

We have two ways to proceed: First, if we assume that the agent is not willing to transmit information about others principals types, in this particular example the revelation principle is valid. Second, if we keep the same setup, we can still have a revelation principle if we enlarge the possible messages an agent can send to a principal. In this case, we propose a new set of direct mechanisms such that using those the revelation principle is valid. First, let consider the following payoff matrices:

	I	NI
I	4,4,3	1,2,2
NI	2,1,2	1,1,0

t_h

	I	NI
I	1,1,1	1,2,1
NI	2,1,**2**	4,4,**1**

t_l

The difference with the previous matrices is that at the low state, the agent prefers manufacturer 2 to invest. Then, there are no incentives for the agent to transmit information about the state to manufacturer 2: In both states, the agent prefers manufacturer 2 to invest. Because of this change in agent's payoffs, now the equilibrium outcome where the manufacturers coordinate their actions through the agent cannot be implemented. The new equilibrium of the indirect mechanism game is the following: Manufacturer 1 selects $\pi_1(t_l)(a) = \pi_1(t_l)(b) = NI$ and $\pi_1(t_h)(a) = \pi_1(t_h)(b) = I$. Manufacturer 2 selects $(\pi_2(a), \pi_2(b)) = (NI, NI)$. In the direct mechanism game, the equilibrium is the following: Manufacturer 1 selects $\pi_1(t_l)(c) = NI$ and $\pi_1(t_h)(c) = I$. Manufacturer 2 selects $\pi_2(c) = NI$. In both cases the equilibrium outcome is the same, (I, NI) on t_h and (NI, NI) on t_l, thus the revelation principle is valid. The main message of this result is that if we restrict the incentives of the agent to not be willing to transmit information about other principals' types, the revelation principle may be valid. More generally, as we shall show, assuming the agent is not willing to transmit information about other principals' allocations (example 1) and principals' types (example 2) is sufficient for the revelation principle to be valid.

A second approach is to consider a new set of direct mechanisms with a richer message space. In the second example, the revelation principle fails because the agent does not have a message to send to the principal 2 about the type of the principal 1. Consider a new set of direct mechanisms where now each message space contains one element per each possible type of the other principal. Then, in this case, a direct mechanism for manufacturer 2 will be $\pi_2 : \{t_l, t_h\} \to \{I, NI\}$, while for manufacturer 1 a direct mechanism will be the same as before, with the message being a singleton $\pi_1 : \{c\} \to \{I, NI\}$. Using this new direct mechanisms, the equilibrium is the following: Manufacturer one selects $\pi_1(t_l)(c) = NI$ and $\pi_1(t_h)(c) = I$. Manufacturer two selects $(\pi_2(t_l), \pi_2(t_h)) = (NI, I)$. The equilibrium outcome is (I, I) on t_h and (NI, NI) on t_l. Then, with this new class of mechanisms, we can implement the same equilibrium outcome from the indirect mechanism game; thus the revelation principle is valid (relative to this new class of direct mechanisms). Intuitively, using this class of mechanisms, the agent can transmit information to a principal about the information she can obtain about the other principal's type.

We now formalize all these intuitions for more general environments and show that these assumptions (plus some other, mild assumptions) are enough to have a revelation principle.

3 Model

In this section, we set up a game of a finite number of principals, each of which designs a mechanism for the same agent. There is a finite set of principals N with $|N| \geq 1$. At the beginning of the game principal i privately observes $t_i \in T_i$. Each t_i is independently drawn, with the profile t drawn from a distribution μ_0. For simplicity, assume T_i is a finite set. Principals simultaneously commit to chose an allocation $y_i \in Y_i$ after a private message $m_i \in M_i$ from the agent. We say principal i offers to the agent a mechanism $\pi_i \in \Pi_i$ which consists of a message space[6] M_i and an allocation function[7] $\pi_i : M_i \to Y_i$. We refer to these as *indirect mechanisms*.

There is one agent that privately observes $\theta \in \Theta$ at the start of the game, drawn

[6]As is common in the literature, we assume each M_i is fixed, and each principal chooses only the allocation function. See Epstein and Peters (1999) and Attar et. al (2008).

[7]With a slight abuse of notation, we refer to both, the mechanism and the allocation function, as π_i.

from a distribution λ. After observing the profile of mechanisms, the agent sends a private message to each principal. The message profile is denoted by $m \in M = \times_{i \in N} M_i$. A message profile $m \in M$ induces an allocation profile $\pi(m) \in Y = \times_{i \in N} Y_i$. For convenience, we denote $\pi(m)$ as the allocation profile with element i, $\Big(\pi(m)\Big)_i = \pi_i(m_i)$. Payoffs for principal i and the agent are $V_i\Big(\theta, t, y\Big)$ and $U\Big(\theta, t, y\Big)$ respectively. For a simpler exposition, we focus on pure strategies: For principal i, $\sigma_i : T_i \to \Pi_i$, $\sigma_i \in \Sigma_i$, and for the agent $\eta : \Theta \times \Pi \to M$. All these together define the indirect mechanism game:

$$\Gamma_I = \Big[\Big(T_i\Big)_{i \in N}, \Theta, \Pi, M, U(\ldots), \Big(V_i(\ldots)\Big)_{i \in N}, \mu_0, \lambda\Big]$$

Consider the following alternative game. Now the message space that principal i uses is Θ. Thus, principal i offers to the agent a mechanism $\tilde{\pi}_i : \Theta \to Y_i$. Let $\tilde{\Pi}_i^S$ denote the set of *strong direct mechanisms*. Pure strategies for principal i will be $\tilde{\sigma}_i : T_i \to \tilde{\Pi}_i^S$, $\tilde{\sigma}_i \in \tilde{\Sigma}_i^S$ and for the agent $\tilde{\eta} : \Theta \times \tilde{\Pi}^S \to \Theta^N$. Analogously, these define the strong direct mechanism game:

$$\Gamma_S = \Big[\Big(T_i\Big)_{i \in N}, \Theta, \tilde{\Pi}^S, \Theta^N, U(\ldots), \Big(V_i(\ldots)\Big)_{i \in N}, \mu_0, \lambda\Big]$$

We also consider a second alternative game: Principal i uses message space $\Theta \times T_{-i}$. Thus mechanism i has the form $\tilde{\pi}_i : \Theta \times T_{-i} \to Y_i$. Let $\tilde{\Pi}_i^W$ denote the set of *weak direct mechanisms*. Agent's strategy will be $\tilde{\eta} : \Theta \times \tilde{\Pi}^W \to \times_{i \in N}\Big[\Theta \times T_{-i}\Big]$. This defines the weak direct mechanism game:

$$\Gamma_W = \Big[\Big(T_i\Big)_{i \in N}, \Theta, \tilde{\Pi}^W, \times_{i \in N}\Big[\Theta \times T_{-i}\Big], U(\ldots), \Big(V_i(\ldots)\Big)_{i \in N}, \mu_0, \lambda\Big]$$

The solution concept adopted is pure strategy perfect Bayesian equilibrium. Denote $\mu_i : \Pi_i \to \Delta T_i$ as the belief about principal i's type after the agent observes mechanism $\pi_i \in \Pi_i$. As before, denote $\mu(\pi)$ the profile of beliefs such that element i is $\Big(\mu(\pi)\Big)_i = \mu_i(\pi_i)$. The triple (σ, η, μ) is a *pure strategy perfect Bayesian equilibrium* (PBE) of the game Γ_I if the pair (σ, η) is *sequentially rational* given μ and μ is *Bayesian consistent* with σ (defined below). We denote analogously $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ for a PBE on the games Γ_W and Γ_S; all the definitions are extended analogously.

Definition: (σ, η) is *sequentially rational* given μ if:

(i) $\eta : \Theta \times \Pi \to M$ is a *continuation equilibrium relative to* μ: for every $\theta \in \Theta$ and $\pi \in \Pi$,

$\eta(\theta,\pi) = \arg\max_{m\in M} \mathbb{E}_{\mu(\pi)} \left[U(\theta,t,\pi(m))\right]$

(ii) $\sigma_i : T_i \to Y_i$ is optimal for principal i:

$$\sigma_i = \arg\max_{\pi_i\in\Sigma_i} \mathbb{E}_{\lambda\times\mu_0}\left[V_i\Big(\theta,t,\pi_i\Big(\eta_i(\theta,\pi_i(t_i),\sigma_{-i}(t_{-i}))\Big),\sigma_{-i}(t_{-i})\Big(\eta_{-i}(\theta,\pi_i(t_i),\sigma_{-i}(t_{-i}))\Big)\Big)\right]$$

Definition: μ is *Bayesian consistent* with σ if μ is obtained from Bayes rule whenever is possible using σ and μ_0: for each principal i and for every on-path mechanism $\pi_i \in \sigma_i(T_i)$, $\mu_i(\pi_i)(t_i) = \frac{\mu_0(t_i)}{\sum_{t_i':\sigma_i(t_i')=\pi_i}\mu_0(t_i')}$. There are no restrictions on μ for off-path mechanisms.

3.1 Assumptions

We now introduce some assumptions we use later for different results. Note that all the assumptions are on the agent's payoffs. The first two relates agent's preferences among one principal's allocation with other principals' allocations.

Assumption: U is *weakly outcome separable* if for every $y_i, y_i' \in Y_i$, for every $y_{-i}, y_{-i}' \in Y_{-i}$, and for every $\theta \in \Theta$ and $t \in T$:

$$U\Big(\theta,t,y_i,y_{-i}\Big) > U\Big(\theta,t,y_i',y_{-i}\Big) \Rightarrow U\Big(\theta,t,y_i,y_{-i}'\Big) > U\Big(\theta,t,y_i',y_{-i}'\Big)$$

Intuitively, this assumption requires that agent's preference among Y_i do not depend on the allocation profile y_{-i}. For an extended version of the model, when the agent is allowed to choose a contractible action, this assumption is satisfied for several important games such as provision of public good, public finance, auctions, trade policy, exclusive dealing and lobbying.[8,9] In section 7 we present two additional relevant economic examples that satisfy this assumption. We also discuss the necessity of this assumption for our results. For the case when principals do not have private information, this assumption reduces to the same assumption of Attar et. al (2008) and Peters (2003).
We also consider a stronger version of this assumption:

[8] In the extended model, *weakly outcome separability* is satisfied for each possible agent's action. We discuss more about it in the extensions of our results in Section 5.

[9] Laussel and Le Breton (1998), Grossman and Helpman (1997), Bernheim and Whinston (1986), Grossman and Helpman (1994), Bernheim and Whinston (1998), Le Breton and Salanié (2003)

Assumption: U is *outcome separable* if for every $y_i, y'_i \in Y_i$, for every $y_{-i}, y'_{-i} \in Y_{-i}$, and for every $\theta \in \Theta$ and $t \in T$:

$$U\Big(\theta, t, y_i, y_{-i}\Big) - U\Big(\theta, t, y'_i, y_{-i}\Big) = U\Big(\theta, t, y_i, y'_{-i}\Big) - U\Big(\theta, t, y'_i, y'_{-i}\Big)$$

The next assumption relates agent's preferences over one principal's allocation with other principals' types.

Assumption: U is *type separable* if for every $y_i, y'_i \in Y_i$, for every $y_{-i} \in Y_{-i}$, for all $\theta \in \Theta$, $t_i \in T_i$ and $t_{-i}, t'_{-i} \in T_{-i}$:

$$U\Big(\theta, t_i, t_{-i}, y_i, y_{-i}\Big) - U\Big(\theta, t_i, t_{-i}, y'_i, y_{-i}\Big) = U\Big(\theta, t_i, t'_{-i}, y_i, y_{-i}\Big) - U\Big(\theta, t_i, t'_{-i}, y'_i, y_{-i}\Big)$$

This assumption requires that agent's preference over Y_i not depend on the type profile t_{-i}.These two last assumptions are satisfied in recent privately informed principals' models.[10] The next assumptions are only to simplify the exposition. The next two assume that the message space each principal uses to communicate with the agent has enough elements.

Assumption: The game Γ_I satisfies *weak richness* if for every i, $|M_i| \geq |\Theta \times T_{-i}|$

Assumption: The game Γ_I satisfies *strong richness* if for every i, $|M_i| \geq |\Theta|$

For example, in the case of weak richness, this says there is at least one message for each possible information the agent may have about his own type and others principals' types. Thus, we are assuming that the agent has at least one different message to tell the principal i the information that he does not know. The next two assumptions are needed to ensure the agent has a unique preferred allocation with each principal. The first one assumes the agent is never indifferent to a change in the allocation with other principals.

Assumption: U satisfies *no indifference* if for every $y_i \in Y_i$, for all $\theta \in \Theta$, $t \in T$, and $y_{-i} \neq y'_{-i} \in Y_{-i}$, $U\Big(\theta, t, y_i, y_{-i}\Big) \neq U\Big(\theta, t, y_i, y'_{-i}\Big)$

The second one assumes continuity of the agent's payoffs in principals' types.

Assumption: U is *continuous* for every $t \in T$.

[10] Martimort and Moreira (2010), Lima and Moreira (2014)

4 Results

In this section, we present the main results. Note that all our assumptions are either on the agent's payoffs or the cardinality of the message space used by each principal in Γ_I. Thus, all the results are valid for any possible payoffs the principals could have, including cases with allocation and information externalities among the principals. The assumption about principals' message spaces focuses the attention on the interesting case: whenever there is a restriction on how large a message set is, an equilibrium of Γ_I may not be robust to the inclusion of new messages.[11] The interesting cases are indeed the ones we consider here: For any possible large message space, if we focus on robust equilibria, we ask the following question: Is it possible to have an equilibrium in Γ_W or Γ_S where principals choose direct mechanisms, and the same outcome is induced? Putting in another way, is it possible to implement the same outcome using mechanisms with a smaller message set? As we will see, depending on the assumptions on the agent's payoffs, different message spaces are enough to characterize all possible equilibrium outcomes of Γ_I.

The main idea of the revelation principle is to associate each indirect mechanism with a direct mechanism such that the equilibrium on Γ_I gives the same allocation as Γ_W or Γ_S. Implicitly, there is a mapping between indirect and direct mechanism that allows the previous intuition. We generalize these ideas in the following definition:

Definition: Given two sets of mechanisms $\tilde{\Pi}$ and Π, a continuation equilibrium η relative to μ using Π is an *extension* of $\tilde{\eta}$ relative to $\tilde{\mu}$ using $\tilde{\Pi}$ if there exists an one to one mapping $v : \tilde{\Pi} \to \Pi$ such that

$$\tilde{\pi}(\tilde{\eta}(\cdot, \tilde{\pi})) = v(\tilde{\pi})(\eta(\cdot, v(\tilde{\pi}))) \text{ and } \tilde{\mu}(\tilde{\pi}) = \mu(\tau(\tilde{\pi}))$$

Let's first study the implications of our assumptions. The agent's messages to principals maximize his expected payoff $U_\mu\Big(\theta, t, \pi(m)\Big)$, and are denoted by $m(\theta, \pi, \mu)$. In general, the agent's optimal message to principal i is a function $m_i(\theta, \pi, \mu)$. Under our assumptions:

Proposition 4.1: Suppose U satisfies *no indifference condition* and is *continuous*.

[11] Here robust means that the equilibrium strategies remains an equilibrium in the game where principals have available more messages.

i) If U is *outcome separable* and *type separable*, then
$m_i(\theta, \pi, \mu) = m_i(\theta, \pi_i, \pi'_{-i}, \mu_i, \mu'_{-i})$ for every (π'_{-i}, μ'_{-i}).

ii) If U is *outcome separable*, then $m_i(\theta, \pi, \mu) = m_i(\theta, \pi_i, \pi'_{-i}, \mu)$ for every π'_{-i}.

To state our results, it is useful to define an outcome function, which gives an allocation profile for every possible state of the world.

Definition: An *outcome function* is a mapping from states of the world to the allocation profile $F : \Theta \times T \to Y$

We say that a PBE (σ, η, μ) *implements* F if $\sigma(t)\Big(\eta(\theta, \sigma(t))\Big) = F(\theta, t)$ for every $(\theta, t) \in \Theta \times T$, which means that as a result of the interaction of the principals and the agent on equilibrium, the outcome is the same as the one given by the outcome function for every state of the world.[12] Our first result is the more restrictive one. It gives conditions such that instead to studying Γ_I, we can focus to the simpler game Γ_S where principals choose the usual direct mechanism: A mapping from the agent's type space to allocations $\pi_i : \Theta \to Y_i$.

Theorem 4.1: If U is *outcome separable, type separable*, satisfies *no indifference* and Γ_I satisfies *strong richness condition*, then:

i) For every *continuation equilibrium* η relative to μ on Γ_I, there exist a mapping $\tau : \Pi \to \tilde{\Pi}^S$, a belief scheme $\tilde{\mu}$ and a continuation equilibrium $\tilde{\eta}$ relative to $\tilde{\mu}$ on Γ_S such that:

   a) η and $\tilde{\eta}$ induce the same outcome: $\tau(\pi)(\tilde{\eta}(\cdot, \tau(\pi))) = \pi(\eta(\cdot, \pi))$

   b) μ and $\tilde{\mu}$ induce same beliefs when $\tau^{-1}(\tau(\pi))$ is a singleton: $\tilde{\mu}(\tau(\pi)) = \mu(\pi)$

   c) Agent truthfully reports on Γ_S: For every $\tilde{\pi} \in \tau(\Pi)$, $\tilde{\eta}_i(\theta, \tilde{\pi}) = \theta$

ii) If F can be implemented by a PBE (σ, η, μ) on Γ_I, then F can be implemented by a PBE $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ on Γ_S, where the agent truthfully reports.

iii) If $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ is a PBE of Γ_S where the agent truthfully reports, then for any extension of $\tilde{\eta}$ relative to $\tilde{\mu}$, denoted by η relative to μ using mechanisms Π, (σ, η, μ) is a PBE on Γ_I, where $\sigma_i = v(\tilde{\sigma}_i)$.

[12] Denote as before $(\sigma(t))_i = \sigma_i(t_i) \in \Pi_i$

For example, an agent's payoff that is totally separable in principals' allocations and types of the form $U\Big(\theta, t, y\Big) = \sum_i u_i\Big(\theta, t_i, y_i\Big)$ satisfies outcome and type separability.

Fully-revealing PBE

The next result does not assume type separability and uses instead a set of mechanisms with a richer message space: the agent is required to report his exogenous information and also the endogenous information that she may learn about principals' types. For simplicity, we first focus on equilibria of the original game where the principals fully reveal their types to the agent on-path. To state the theorem, denote δ_{t_i} as the Dirac measure on t_i.

Theorem 4.2: If U is *weakly outcome separable*, satisfies *no indifference* and Γ_I satisfies *weak richness condition*, then

i) For every *continuation equilibrium* η relative to μ on Γ_I, there exist a mapping $\tau : \Pi \to \tilde{\Pi}^W$, a belief scheme $\tilde{\mu}$ and a continuation equilibrium $\tilde{\eta}$ relative to $\tilde{\mu}$ on Γ_W such that:

   a) η and $\tilde{\eta}$ induce the same outcome: $\tau(\pi)(\tilde{\eta}(\cdot, \tau(\pi))) = \pi(\eta(\cdot, \pi))$

   b) μ and $\tilde{\mu}$ induce same beliefs when $\tau^{-1}(\tau(\pi))$ is a singleton: $\tilde{\mu}(\tau(\pi)) = \mu(\pi)$

   c) Agent truthfully reports on Γ_W: For every $\tilde{\pi} \in \tau(\Pi)$, $\tilde{\eta}_i(\theta, \tilde{\pi}) = (\theta, t_{-i})$, $t_{-i} \in \text{supp } \tilde{\mu}_{-i}(\tilde{\pi}_{-i})$

ii) If F can be implemented by a PBE (σ, η, μ) on Γ_I with $\mu_i(\sigma_i(t_i)) = \delta_{t_i}$, then F can be implemented by a PBE $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ on Γ_W, where the agent truthfully reports.

iii) If $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ is a PBE of Γ_W with $\tilde{\mu} \in F$ where the agent truthfully reports, then for any extension of $\tilde{\eta}$ relative to $\tilde{\mu}$, denoted by η relative to μ using mechanisms Π, (σ, η, μ) is a PBE on Γ_I, where $\sigma_i = v(\tilde{\sigma}_i)$.

Now, an agent's payoff that is separable only in principals allocation, of the form $U\Big(\theta, t, y\Big) = \sum_i u_i\Big(\theta, t, y_i\Big)$ satisfies outcome separability.

Necessity of weakly outcome separability

At this point, it is important to show the necessity of the weakly outcome separable condition. We will show that if this condition is not satisfied, it is possible to construct principals' payoffs such that there is a PBE of the indirect mechanism game where at least one principal chooses an indirect mechanism that cannot be mapped into any direct mechanism considered in this paper. Moreover, this indirect mechanism will have at least one message associated with an allocation that will never be induced on-path for the agent. In those cases, a direct mechanism may need to have an arbitrarily larger message space to be able to replicate the same equilibrium outcome.[13] Thus, a revelation principle kind of result seems hard to obtain. More formally, note that for any PBE (σ, η, μ) of the game Γ_I, it is always the case that for every $t_i \in T_i$, $\sigma_i(t_i)\Big(\eta_i(\Theta, \sigma_i(t_i), \sigma_{-i}(T_{-i}))\Big) \subseteq \sigma_i(t_i)\Big(M_i\Big)$. The reverse is not always true. Intuitively, there might be equilibrium outcomes supported by some principal including in his mechanism messages associated with allocations that are never induced by the agent on-path. This is known in the literature of uninformed multiple principals as a *latent contract*.[14] The problem with games that have this kind of equilibria is that it is not known how large the message space should be to characterize every possible equilibrium outcome. For that reason, an interesting class of games is the ones where every equilibrium (using indirect mechanisms) do not include *latent contracts*.

Definition: A game Γ_I has *no latent contracts* if for every PBE (σ, η, μ), for every $t_i \in T_i$, $\sigma_i(t_i)\Big(\eta_i(\Theta, \sigma_i(t_i), \sigma_{-i}(T_{-i}))\Big) = \sigma_i(t_i)\Big(M_i\Big)$

Proposition 4.2: If a game Γ_I has no *latent contracts*, then U satisfies *weakly outcome separability*.

Intuitively, whenever the agent's payoff does not satisfy *weakly outcome separability*, it is possible to construct principal's payoffs in a way that a PBE of the constructed game Γ_I involves at least one *latent contract*.

Example of a no fully-revealing equilibrium

Theorem 2 does not assume *weak type separable* condition, but this comes at the cost of enlarging the message space of the direct mechanism considered. Still, this set of direct mechanisms does not allow us to implement outcome functions where the PBE in the indirect mechanism game is implemented by principals' strategies that do not fully

[13] Attar et. al (2011) is an example of equilibrium with this feature.

[14] A *latent contract* is an allocation in the range of the mechanism offered by some principal that the agent never induce on-path.

reveal their types to the agent. We need an extra assumption if we want to implement any PBE. The next example shows why is necessary to assume a new condition.

Let's consider a game between two principals and one agent. Principal 2 is the only with private information $t \in \{t_l, t_h\}$, which only affect agent payoffs. Messages space $M_1 = M_2 = \{a, b\}$. Each principal decides an allocation $Y_1 = Y_2 = \{x, y\}$. Denote p as the probability of the high state $p = \Pr(w_h)$ and assume that $p \in (\frac{1}{2}, \frac{2}{3})$. The payoffs matrices are the following:

	x	y
x	$3,2,\mathbf{0}$	$2,3,\mathbf{2}$
y	$1,0,\mathbf{2}$	$1,1,\mathbf{4}$

t_h

	x	y
x	$3,2,\mathbf{6}$	$2,3,\mathbf{2}$
y	$1,0,\mathbf{2}$	$1,1,\mathbf{0}$

t_l

Note that agent's payoff satisfy *weak outcome separability* for every state of the world, but the expected payoff for the agent (using the prior) does not satisfy *weak outcome separability*. Also, *weak type separability* is not satisfied. The following is a PBE of the indirect mechanism game: $(\pi_1(a), \pi_1(b)) = (x, y)$, $\pi_2(t_h)(a) = \pi_2(t_h)(b) = \pi_2(t_l)(a) = \pi_2(t_l)(b) = x$ and $\mu_2(\pi_2) = \mu_0$ for every π_2. The outcome supported by this equilibrium is (x, x) after both states of the world. Note that this outcome is supported with principal 2 not revealing the state of the world. The intuition of why the revelation principle fails is similar to the first example presented before: principal 1 uses the mechanism to get information about which allocation principal 2 makes available to the agent. Because of the agent's payoff, the agent wants to transmit this information to principal 1: Whenever principal 2 makes available only allocation y, the agent will send message b to principal 1, and whenever it is the allocation x instead, she will send message a. As before, principal 1 does not use the mechanism to screen the principal's 2 types. Instead, he uses it to deter principal 2 to select allocation y (which would be better for principal 2 in both states). Moreover, message b is never selected for the agent on path. Thus, we cannot implement this outcome using any direct mechanisms presented before. We need an extra assumption to be able to implement this kind of outcomes supported by non-revealing principal's strategies. The next result shows that it is enough to assume the strongest version of separability, *outcome separability*.

Theorem 4.3: If U is *outcome separable*, satisfies *no indifference* and Γ_I satisfies *weak richness condition*, then

i) For every *continuation equilibrium* η relative to μ on Γ_I, there exist a mapping τ :

$\Pi \to \tilde{\Pi}^W$, a belief scheme $\tilde{\mu}$ and a continuation equilibrium $\tilde{\eta}$ relative to $\tilde{\mu}$ on Γ_W such that:

a) η and $\tilde{\eta}$ induce the same outcome: $\tau(\pi)(\tilde{\eta}(\cdot, \tau(\pi))) = \pi(\eta(\cdot, \pi))$

b) μ and $\tilde{\mu}$ induce same beliefs when $\tau^{-1}(\tau(\pi))$ is a singleton: $\tilde{\mu}(\tau(\pi)) = \mu(\pi)$

c) Agent truthfully reports on Γ_W: For every $\tilde{\pi} \in \tau(\Pi)$, $\tilde{\eta}_i(\theta, \tilde{\pi}) = (\theta, t_{-i})$, $t_{-i} \in$ supp $\tilde{\mu}_{-i}(\tilde{\pi}_{-i})$

ii) If F can be implemented by a PBE (σ, η, μ) on Γ_I, then F can be implemented by a PBE $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ on Γ_W, where the agent truthfully reports.

iii) If $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ is a PBE of Γ_W where the agent truthfully reports, then for any extension of $\tilde{\eta}$ relative to $\tilde{\mu}$, denoted by η relative to μ using mechanisms Π, (σ, η, μ) is a PBE on Γ_I, where $\sigma_i = v(\tilde{\sigma}_i)$.

5 Extensions

5.1 No-indifference condition

Regarding the *indifference condition*, we need to assume it only to get rid of equilibria in the indirect mechanism game where the agent, for a given profile of indirect mechanisms, has multiple allocation maximizers with one principal.[15] If that is the case, there could be equilibria where the agent changes his optimal message with one principal (off-path) when others principals change their mechanisms, even though it is not strictly profitable for her. Equilibria like that are not natural because they rely on different agent's responses with a principal even though she is indifferent. An alternative approach is to restrict the set of equilibria that we are interested in studying. Formally, consider the *continuation equilibria* with the following property:

Definition: A *continuation equilibrium* η relative to μ satisfies *separable response* if there is no $\theta \in \Theta$, $\pi \in \Pi$ and $\pi'_i \in \Pi_i$, such that $\mathbb{E}_{\mu(\pi)}[U(\theta, \pi)] = \mathbb{E}_{\mu(\pi'_i, \pi_{-i})}[U(\theta, \pi'_i, \pi_{-i})]$ and $\eta_{-i}(\theta, \pi) \neq \eta_{-i}(\theta, \pi'_i, \pi_{-i})$

[15]This section is based on Attar et. al (2008) which deals with the same issue for the non-privately informed principals' case.

Whenever we restrict to PBE with *continuation equilibria* that satisfies *separable response*, we will not need to assume the *no indifference* and *continuity* conditions. Let's illustrate the equilibria we are ruling out with an example borrowed from Attar et. al (2008):

5.2 Example 3:

Consider a two principal - one agent problem, where $Y_1 = Y_2 = \{x, y, z\}$, $M_1 = M_2 = \{a, b\}$, neither the principals nor the agent have private information. Consider the following payoffs matrices:

	x	y	z
x	1, 1, 0	2, 1, 2	−1, 5, 2
y	1, 2, 0	1, 1, 1	0, 0, 1
z	5, −1, 0	0, 0, 1	0, 0, 1

First, note that the *no indifference* condition is not satisfied. The following constitutes an on-path behavior of an equilibrium from the indirect mechanism game: Both principals select $(\pi_i(a), \pi_i(b)) = (y, z)$, and the agent selects (a, a), inducing the allocation (y, y). This on-path behavior is only an equilibrium if we specify the following off-path behavior: If one principal offers $\pi_i(a) = \pi_i(b) = x$, then the agent instead of choosing a with principal 2, she chooses b (because she is indifferent), inducing the outcome (x, z). Because of that, from a principal perspective offering $\pi_i(a) = \pi_i(b) = x$ does not constitutes a profitable deviation like it would be in the case where the agent chooses the message a with principal 2. Since the agent is changing his message with principal 2 from a to b even though it is not strictly profitable to do it, we focus on equilibria where such off-path behaviour is not present.

5.3 Non-delegation mechanisms

Consider the case when each principal can also participate in the mechanism offered to the agent. Formally, a mechanism $\pi_i \in \Pi_i$ is now a mapping $\pi_i : M_i \times P_i \to Y_i$, where P_i is a message space that the principal can use to induce an allocation. Each principal and the agent send simultaneous messages to each mechanism. In this framework, for a given mechanism profile π, a message profile $m \in M$ from the agent to the mechanisms and a message profile $p \in P$ from the principals to their mechanism will induce the allocation $\pi(m, p) = (\pi_i(m_i, p_i))_{i \in N} \in Y$. A pure strategy for the principal besides the mechanism also includes the message to send to the mechanism, $\sigma_i : T_i \to \Pi_i \times P_i$. The rest of the elements of the game remains equal, defining the indirect mechanism game $\Gamma_I = \Big[\Big(T_i\Big)_{i\in N}, \Theta, \times_{i\in N}\Big[\Pi_i \times P_i\Big], M, U(\ldots), \Big(V_i(\ldots)\Big)_{i\in N}, \mu_0, \lambda\Big]$. Consider now the following alternative class of games. Principal i mechanism $\tilde{\pi}_i \in \tilde{\Pi}_i$ takes the form $\tilde{\pi}_i : \Theta \times T_i \to Y_i$. A pure strategy for the principal is now $\tilde{\sigma}_i : T_i \to \tilde{\Pi}_i^S \times T_i$, with each principal reporting directly his type to his mechanism. That defines the game $\Gamma_S = \Big[\Big(T_i\Big)_{i\in N}, \Theta, \times_{i\in N}\Big[\tilde{\Pi}_i^S \times T_i\Big], \Theta^N, U(\ldots), \Big(V_i(\ldots)\Big)_{i\in N}, \mu_0, \lambda\Big]$. Consider also the second alternative game where principal i use mechanisms of the form $\tilde{\pi}_i : \Theta \times T \to Y_i$, again with each principal reporting his own type to his mechanism. Analogously, that defines the game $\Gamma_W = \Big[\Big(T_i\Big)_{i\in N}, \Theta, \times_{i\in N}\Big[\tilde{\Pi}_i^W \times T_i\Big], \times_{i\in N}\Big[\Theta \times T_{-i}\Big], U(\ldots), \Big(V_i(\ldots)\Big)_{i\in N}, \mu_0, \lambda\Big]$. Our main results (Theorem A and A) holds true in this new framework. The proof is similar to our previews results once we define the mapping τ accordingly and consider that whenever principal and agent report to the mechanism simultaneously, we can treat both as players in a well defined Bayesian game and apply the usual revelation principle for Bayesian games[16] at that stage. For example, to replicate Theorem A in this context, define the mapping $\tau(\pi) \in \tilde{\pi}$ as follows: $\tau_i(\pi_i)(\theta, t) = \pi_i(\eta(\theta, \pi), \sigma_i^r(t_i))$ whenever $t_{-i} \in \text{supp}\ \mu_{-i}(\pi_{-i})$ and $\sigma_i^r(t_i)$ is the projection of $\sigma_i(t_i)$ over P_i . As before, this is well defined because of our assumptions on U. The rest follows from the proof of Theorem A.

5.4 Contractable action

Suppose now the agent can take a public action that each principal can contract on his mechanism. Now, principal i mechanism is a mapping $\pi_i : M_i \times A \to Y_i$ where A is the set of contractible actions of the agent. Assume the agent choose his actions simultaneously with the messages, and that every principal can contract in this action. In this context, an agent's pure strategy is $\eta : \Theta \times \Pi \to M \times A$. Consider the first alternative game where principals set mechanism $\tilde{\pi}_i : \Theta \times A \to Y_i$ and the agent use strategies of the form $\tilde{\eta} : \Theta \times \Pi \to \Theta^N \times A$. In the second alternative game principals set mechanism $\tilde{\pi}_i : \Theta \times T_{-i} \times A \to Y_i$ and the agent uses strategies of the form $\tilde{\eta} : \Theta \times \Pi \to \times_{i\in N}[\Theta \times T_{-i}] \times A$. Our main results still holds true in this new framework if we assume our conditions for every possible action. For example, to replicate Theorem

[16] Myerson (1979)

A the analogue of *outcome separability* for this new framework is the following:

Assumption: U is *outcome separable* if for every $y_i, y'_i \in Y_i$, for every $y_{-i}, y'_{-i} \in Y_{-i}$, for every $\theta \in \Theta$, $t \in T$ and $a \in A$:

$$U\Big(\theta, t, y_i, y_{-i}, a\Big) - U\Big(\theta, t, y'_i, y_{-i}, a\Big) = U\Big(\theta, t, y_i, y'_{-i}, a\Big) - U\Big(\theta, t, y'_i, y'_{-i}, a\Big)$$

For the result, define the mapping $\tau(\pi) \in \tilde{\pi}$ as follows: $\tau_i(\pi_i)(\theta, t_{-i}, a) = \pi_i(\eta_i^M(\theta, \pi), a)$ whenever $t_{-i} \in \text{supp } \mu_{-i}(\pi_{-i})$, where η_i^M is the projections of η over M_i. As before, this is well defined because of our assumptions on U. The rest follows from the proof of Theorem A.

6 Direct mechanism game

In this section, we discuss the consequences of Theorem A. For simplicity assume there are only two principals and conditions of Theorem A are satisfied. We are interested in characterizing all possible outcome functions $F : \Theta \times T \to Y$ implemented by a PBE (σ, η, μ) of the indirect mechanism game Γ_I. Using Theorem A, it is enough to study the game Γ_W and restrict each principal to choose an *incentive compatible* direct mechanism: a direct mechanism where the agent truthfully report his exogenous (his private information) and endogenous information (principals' private information). Fix principal 2 strategy $\tilde{\sigma}_2 = (\tilde{\sigma}_2(t_2))_{t_2 \in T_2}$, where $\tilde{\sigma}_2(t_2) : \Theta \times T_1 \to Y_2$. This determines $\tilde{\mu}_2(\cdot)$ for every $\tilde{\pi}_2 \in \tilde{\sigma}_2(T_2)$. Denote $U_{\tilde{\mu}(\cdot,\cdot)}(\theta, y) = \mathbb{E}_{\tilde{\mu}(\cdot,\cdot)}\Big[U\Big(\theta, t, y\Big)\Big]$ the expected value given $\tilde{\mu}$. Principal 1 solves the following problem:

$$\max_{\tilde{\sigma}_1 = (\tilde{\sigma}_1(t_1))_{t_1 \in T_1}} \int_{\Theta \times T} V_1\Big(t, \theta, \tilde{\sigma}_1(t_1)(\theta, t_2), \tilde{\sigma}_2(t_2)(\theta, t_1)\Big)\mu_0(dt)\lambda(d\theta)$$

subject to *incentive compatible* condition:
For every $(\theta, t) \in \Theta \times T$ and $(\theta', t'_2) \in \Theta \times T_2$,

$$U_{\tilde{\mu}(\tilde{\sigma}_1(t_1), \tilde{\sigma}_2(t_2))}\Big(\theta, \tilde{\sigma}_1(t_1)(\theta, t_2), \tilde{\sigma}_2(t_2)(\theta, t_1)\Big) \geq U_{\tilde{\mu}(\tilde{\sigma}_1(t_1), \tilde{\sigma}_2(t_2))}\Big(\theta, \tilde{\sigma}_1(t_1)(\theta', t'_2), \tilde{\sigma}_2(t_2)(\theta, t_1)\Big)$$

and *Bayesian consistency*:

$$\tilde{\mu}_1 \text{ is a Bayesian consistent with } \tilde{\sigma}_1$$

Suppose we relax the problem and do not consider the first constraint. Thus, the problem is a standard signaling game[17] where principal 1's action, that may signal his private information, is a direct mechanism, and agent's action is the message that she sends to the mechanism, which depend on his updated beliefs. Thus, the only problem is to understand the first constraint. But this is a constraint in the set of direct mechanism available to the principal 1. Specifically, it requires direct mechanisms to be *incentive compatible* for the particular beliefs induced on path. As we show in Section 7, the problem of characterizing equilibrium outcomes for the indirect mechanism game Γ_I simplifies drastically for the two economic environments studied. In any case, *outcome separability* allows us to simplify the problem because the agent reports the same to a principal regardless of any change in allocation with the others' principals. The only effect of one principal's mechanism to the others principals is on the information he may signal to the agent. Because of that, a principal may also screen this information using the mechanism.

As this procedure suggests, the problem can be decomposed in two steps. In the first one, the **screening** step, we need to find *incentive compatible* direct mechanisms for every possible agent's belief. Then, in the second one, the **signaling** step, we need to solve a standard signaling game where the action of the principal are the *incentive compatible* direct mechanisms from the first step.

Consider now the problem of implementing an outcome where the PBE is fully revealing. For this case, principal 1 needs to screen two pieces of information: the exogenous agent's private information θ, and the endogenous information about principal 2, t_2. This problem may be not easy to solve because it boils down to a multidimensional screening problem. The same is true if we consider more than two principals. For that reason, in section 7, we show two economic environments where the main concern is the principals' private information and not the agent's one. This simplification allows us to reduce the problem to a standard unidimensional screening problem.

[17] Spence (1973)

7 Examples

In this section, we show how to apply our main results using two examples. In both cases, the new class of direct mechanism where each principal screens the possible endogenous information about others' principals' types is considered.

7.1 Expert Delegation

Suppose two experts observe different pieces of information from the same state of the world. For example, both the environmental and the economic departments from a government privately observe independent information about a public policy project. Suppose each department controls a policy decision about the project and can communicate with the president before deciding about it. In this context, each department would like to know the other department's information and take a decision accordingly, while the president would like to match both decisions with the state of the world. Formally, in this example, each department is a principal, and the president is the agent. Each principal privately observes an independent signal of the state of the world $t_i \in [-1, 1]$, drawn from a distribution with positive density f. Assume the state of the world is simply the average of principals' information $w = \left(\frac{t_1+t_2}{2}\right)$. Each principal controls an action $y_i \in Y_i = [-1, 1]$. Assume there is a message space M_i for every i that satisfies *weak richness*[18] such that the agent can communicate with each principal. We assume that both principals and agent want to match the state of the world. Principal i's payoff is $V_i\left(t, y_i\right) = -\left(y_i - w\right)^2$ and agent's payoff is $U\left(t, y\right) = -\left(y_1 - w\right)^2 - \left(y_2 - w\right)^2$. The timing is the following: First, principals choose mechanisms $\pi_i : M_i \to Y_i$ simultaneously. Then, the agent sends a message to each principal $m = (m_1, m_2)$ and policy decisions are taken accordingly.

In order to show the implications of our results, we focus only on outcomes implementable with principals' fully-revealing strategies. Theorem A allows to restrict without loss of generality to the game Γ_W where principals offer mechanisms $\tilde{\pi}_i : T_{-i} \to Y_i$ and the agent reports truthfully on-path $\tilde{\eta}\left(\tilde{\sigma_1}(t_1), \tilde{\sigma}(t_2)\right) = (t_2, t_1)$. Fix principal 2 strategy: A one-to-one mapping $\tilde{\sigma}_2 : T_2 \to \tilde{\Pi}_2$. Principal 1 optimal strategy is an increasing

[18] Intuitively, we are assuming that the agent, for every type of the other principal $-i$, has one different message to send to principal i.

mapping $\tilde{\sigma}_1 : T_1 \to \tilde{\Pi}_1$ that satisfies for every $t_1 \in T_1$ the following problem:

$$t_1 \in arg \max_{t_1' \in T_1} \int_{T_2} V_1\Big(t_1, t_2, \tilde{\sigma}_1(t_1')(t_2)\Big) f(t_2) dt_2$$

Restricted that all the mechanisms $\tilde{\pi}_1 \in \tilde{\sigma}_1(T_1)$ satisfy the following *incentive compatibility* condition:

$$U\Big(t_1, t_2, \tilde{\pi}_1(t_2), \tilde{\sigma}_2(t_2)(t_1)\Big) \geq U\Big(t_1, t_2, \tilde{\pi}_1(t_2'), \tilde{\sigma}_2(t_2)(t_1)\Big)$$

We first characterize the mechanism that satisfies this restriction. Then, we use this characterization to solve the signaling problem. Finally, we check that the solution is a one-to-one mapping.

The next result is standard in the delegation literature. It shows a necessary condition for a mechanism to be incentive compatible. Note that the result depends on the type the agent perceive is the principal.

Lemma 7.1: If $\tilde{\pi}_1 : T_2 \to Y_1$ satisfies *incentive compatibility* for type t_1, then

(i) $\tilde{\pi}_1$ is weakly increasing and thus almost everywhere differentiable.

(ii) If $\tilde{\pi}_1$ is strictly increasing at t_2, then $\tilde{\pi}_1(t_2) = \frac{t_1+t_2}{2}$

(iii) If $\tilde{\pi}_1$ is discontinuous at ϕ, then:

1. $\tilde{\pi}_1(\phi^+) + \tilde{\pi}_1(\phi^-) = t_1 + \phi$
2. $\tilde{\pi}_1$ is constant on the left and the right of ϕ
3. $\tilde{\pi}_1(\phi) \in \{\tilde{\pi}_1(\phi^+), \tilde{\pi}_1(\phi^-)\}$

A direct corollary is that $\tilde{\pi}_1(t_2) = \frac{t_1+t_2}{2}$ is *incentive compatible* for type t_1. Consider the *outcome function* F that maximizes the sum of the players' payoffs $F(t) \in arg \max_{y \in Y} V_1(t, y_1) + V_2(t, y_2) + U(t, y)$.

Proposition 7.1: There is a fully revealing PBE (σ, η, μ) on Γ_I that implements the *outcome function* F. To implement this outcome, it is without loss of generality to assume that each principal delegates a **closed interval** of policy decisions to the agent.

The proof is included in the appendix. In this PBE, each principal delegates a closed interval to screen the agent's information about the other principal's type. Moreover, different principal's types delegate different intervals, so each principal fully reveals his type to the agent. In this way, signaling and screening are happening at the same time. Since incentives of both principals and the agent are aligned, the outcome function that maximizes the sum of the players' payoffs is implementable despite the private information.

7.2 Manufacturer Competition

Suppose two manufacturers compete in selling their products through the same retailer. Each manufacturer has a perfectly divisible good to sell and has a privately observed opportunity cost to sell one unit of the item $t_i \in T_i = [0,1]$. We assume each t_i is drawn from an independent distribution F with positive pdf f. Assume that the inverse of the hazard rate function $\frac{1-F(t)}{f(t)}$ is decreasing on t. We can think this opportunity cost as a proxy for the demand of the product: the higher the opportunity cost, the higher the buyer's valuation of a unit of the item, and then the higher will be the demand for it. Suppose the retailer wants to buy manufacturers goods and sell it through his stores. Since both goods are going to be sold in the same location, there could be informational externalities: the higher the demand for some of the goods is, the higher the sales of the retailer will be.

Each manufacturer decides a contracts of the type $y_i = (x_i, p_i) \in Y_i = [0,1] \times R_+$ and can communicates with the retailer before deciding about it. For that, we assume there is a message space M_i that satisfies *weak richness*. Intuitively, since from the perspective of a manufacturer the retailer's valuation from a contract depends on the other manufacturer demand, he may want to screen that information. For a given contract $y_i = (x_i, p_i)$ and opportunity cost t_i, manufacturer's payoff is $V_i\Big(t_i, y_i\Big) = p_i - t_i x_i$. On the other side, for a pair of contracts (y_i, y_j) and opportunity cost pair (t_i, t_j), retailer's payoff is $U\Big(t, y\Big) = D(t_1, t_2)(x_1 + x_2) - p_1 - p_2$ where $D(t_1, t_2)$ is the retailer's demand. For simplicity, assume that $D(t_1, t_2) = \left(\frac{t_1+t_2}{2}\right)$. Assume also that the retailer can reject a contract and obtain outside option of zero.[19] The timing is the following. First, manufacturers offer simultaneously mechanisms $\pi_i : M_i \to Y_i$. Then, the retailer sends a message to each principal $m = (m_1, m_2)$ and contracts are enforced accordingly.

As before, for simplicity we focus only on outcomes implementable with principals' fully-revealing strategies.Theorem A allows to restrict without lost of generality to the game Γ_W where principals offer mechanisms $\tilde{\pi}_i : T_{-i} \to Y_i$ and agent reports truthfully

[19]For simplicity we assume the retailer can not contract with only one manufacturer.

on path $\tilde{\sigma}_0(\tilde{\sigma}_1(t_1), \tilde{\sigma}_2(t_2)) = (t_2, t_1)$.
We first characterize the mechanism that satisfies the incentive compatibility condition. Then, we use that characterization to solve the maximization problem.
The first lemma is standard from the mechanism design literature and characterizes the incentive compatible mechanisms.

Lemma 7.2: A mechanism $\tilde{\pi}_1 = (\tilde{x}_1, \tilde{p}_1)$ is *incentive compatible* for type t_1 if and only if

(i) $\tilde{x}_1$ is increasing

(ii) For every $t_2 \in T_2$:

$$\tilde{p}(t_2) = \tilde{p}(\underline{t}_2) + \left(\frac{t_1 + t_2}{2}\right) \tilde{x}_1(t_2) - \left(\frac{t_1 + \underline{t}_2}{2}\right) \tilde{x}_1(\underline{t}_2) - \frac{1}{2} \int_{\underline{t}_2}^{t_2} \tilde{x}_1(y) dy$$

With this lemma in hand, we can show the following result.

Proposition 7.2: There is a unique outcome implemented by a fully revealing PBE (σ, η, μ) on Γ_I. In order to implement this outcome, it is without loss of generality to assume that each principal chooses only a **price**.

The proof is included in the appendix. In this PBE, each principal offers a price to screen the agent's information about the other principal. The information each manufacturer obtains is whether the demand of the other manufacturer is big enough make a deal with the retailer Moreover, different principal's types set different prices, so principal fully reveals his type to the agent. Like the example before, signaling and screening are happening at the same time.

8 Conclusions

In this paper, we study how a group of privately informed principals non-cooperatively design mechanisms to a common agent. We aim to characterize the set of implementable functions. We show two results. First, assuming *outcome* and *type separability* on the agent's payoffs plus some other mild conditions, we can restrict without loss of generality to a simple game where each principal chooses an incentive compatible direct mechanism where the agent reports his exogenous private information truthfully. Thus, we can restrict to the usual set of direct mechanisms used in the mechanism design literature. Second, if we only assume *outcome separability*, we have an analogous result, but now we need to restrict to a different class of mechanisms where the message space is larger than the agent's exogenous private information space. In these mechanisms, the agent needs to truthfully report his exogenous private information and the endogenous information she may learn about others' principals' types. Thus, the result says that it is without loss of generality to assume each principal chooses an incentive compatible direct mechanism of this kind. We discuss extensions of our results to environments where each principal also participates in the mechanism simultaneously with the agent, and where the agent takes a contractible action. Since our assumptions are only on the agent's payoffs, many used economic environments (that do not necessarily assume privately informed principals) fit in our assumptions. It would be interesting to apply the results of this paper to these economic environments when privately informed principals are assumed to understand the role of private information in the principals' side. We leave this for future research.

A Proofs of Section 4

Before the main proofs, we show some useful lemmas. Let $\mu \in \Delta T$ and $U_\mu\Big(\theta, y\Big) = \int_T U\Big(\theta, t, y\Big)\mu(dt)$. Let $\pi \in \Pi$ be a profile of mechanisms.

Definition U is *μ-weakly outcome separable*, if for every $y_i, y_i' \in Y_i$, for every $y_{-i}, y'_{-i} \in Y_{-i}$, for all $\theta \in \Theta$:

$$U_\mu\Big(\theta, y_i, y_{-i}\Big) > U_\mu\Big(\theta, y_i', y_{-i}\Big) \Rightarrow U_\mu\Big(\theta, y_i, y'_{-i}\Big) > U_\mu\Big(\theta, y_i', y'_{-i}\Big)$$

Note that whenever U is *μ-weakly outcome separable* for every degenerate μ, U is *weakly outcome separable*.

Lemma 1 *Suppose U is μ-weakly outcome separable. If $m \in arg\max_{m' \in M} U_\mu\Big(\theta, \pi(m')\Big)$ then for every $y_{-i} \in Y_{-i}$, $m_i \in arg\max_{m_i' \in M_i} U_\mu\Big(\theta, \pi_i(m_i'), y_{-i}\Big)$*

Proof Suppose not. Then there exists $m_i' \neq m_i$ such that $U_\mu\Big(\theta, \pi_i(m_i'), y_{-i}\Big) > U_\mu\Big(\theta, \pi_i(m_i), y_{-i}\Big)$. Because μ-weak outcome separability, $U_\mu\Big(\theta, \pi_i(m_i'), \pi_{-i}(m_{-i})\Big) > U_\mu\Big(\theta, \pi_i(m_i), \pi_{-i}(m_{-i})\Big)$, which contradicts message profile m being optimal.

Lemma 2 *Suppose U is type separable. If $m \in arg\max_{m' \in M} U_\mu\Big(\theta, \pi(m')\Big)$ then for every $\mu'_{-i} \in \Delta T_{-i}$, $m_i \in arg\max_{m_i' \in M_i} U_{\mu_i \times \mu'_{-i}}\Big(\theta, \pi_i(m_i'), \pi_{-i}(m_{-i})\Big)$*

Proof Suppose not. Then there exists $m_i' \neq m_i$ such that $U_{\mu_i \times \mu'_{-i}}\Big(\theta, \pi_i(m_i'), \pi_{-i}(m_{-i})\Big) > U_{\mu_i \times \mu'_{-i}}\Big(\theta, \pi_i(m_i), \pi_{-i}(m_{-i})\Big)$. Because type separability we can replace μ'_{-i} by μ_{-i} and then $U_\mu\Big(\theta, \pi_i(m_i'), \pi_{-i}(m_{-i})\Big) > U_\mu\Big(\theta, \pi_i(m_i), \pi_{-i}(m_{-i})\Big)$, which contradicts message profile m being optimal.

Lemma 3 *Consider any collection of closed sets $(B_i)_{i \in N}$, $B_i \subset Y_i$. If U is μ-weakly outcome separable, satisfies no indifference and is continuous on types, then $arg\max_{y \in B} U_\mu\Big(\theta, y\Big)$ is a singleton.*

Proof Suppose not. Then, there are $\tilde{y} \neq \overline{y}$ such that $\{\tilde{y}, \overline{y}\} \in \text{argmax}_{y \in B} U_\mu\Big(\theta, y\Big)$. Because no indifference condition, it must be that for every i, $\tilde{y}_i \neq \overline{y}_i$. If not, there is j such that $\tilde{y}_j = \overline{y}_j$. Then $U_\mu\Big(\theta, \tilde{y}_j, \tilde{y}_{-j}\Big) = U_\mu\Big(\theta, \tilde{y}_j, \overline{y}_{-j}\Big)$. Using the continuity of U

over t, there must be some t' such that $U\left(\theta, t', \tilde{y}_j, \tilde{y}_{-j}\right) = U\left(\theta, t', \tilde{y}_j, \overline{y}_{-j}\right)$, which contradicts the no indifference condition. Thus, using the optimality conditions, we have that $U_\mu\left(\theta, \overline{y}\right) > U_\mu\left(\theta, \tilde{y}_i, \overline{y}_{-i}\right)$ and $U_\mu\left(\theta, \overline{y}_i, \tilde{y}_{-i}\right) < U_\mu\left(\theta, \tilde{y}\right)$ which contradicts μ-weak outcome separability.

Note that as a corollary of the previous result, whenever μ is degenerate, we only require *weakly outcome separability* and *no indiference* condition to obtains the same result.

Proposition 4.1 Suppose U satisfies *no indifference condition* and is *continuous*.

i) If U is *outcome separable* and *type separable*, then $m_i(\theta, \pi, \mu) = m_i(\theta, \pi_i, \pi'_{-i}, \mu_i, \mu'_{-i})$ for every (π'_{-i}, μ'_{-i}).

ii) If U is *outcome separable*, then $m_i(\theta, \pi, \mu) = m_i(\theta, \pi_i, \pi'_{-i}, \mu)$ for every π'_{-i}.

Proof We will show both propositions separately. In both cases we only assume indifference condition and continuity over t to have a unique optimal message from the agent with each principal.

i) First, we will show that *outcome separability* implies *μ-weakly outcome separability*. Then the result follows from Lemmas 1, 2 and 3. Suppose $U_\mu\left(\theta, y_i, y_{-i}\right) > U_\mu\left(\theta, y'_i, y_{-i}\right)$. Consider the following difference:

$$\begin{aligned} U_\mu\left(\theta, y_i, y'_{-i}\right) - U_\mu\left(\theta, y'_i, y'_{-i}\right) &= \int_T U\left(\theta, t, y_i, y'_{-i}\right) - U\left(\theta, t, y'_i, y'_{-i}\right)\mu(dt) \\ &= \int_T U\left(\theta, t, y_i, y_{-i}\right) - U\left(\theta, t, y'_i, y_{-i}\right)\mu(dt) \\ &= U_\mu\left(\theta, y_i, y_{-i}\right) - U_\mu\left(\theta, y'_i, y_{-i}\right) \\ &> 0 \end{aligned}$$

Where the second equality follows from outcome separability. Thus U is μ-weakly outcome separable.

ii) Using the result showed in i), *outcome separability* implies *μ-weakly outcome separability* and then using Lemma 1 the result follows.

Theorem 4.1: If U is *outcome separable, type separable*, satisfies *no indifference* and Γ_I satisfies *strong richness condition*, then:

i) For every *continuation equilibrium* η relative to μ on Γ_I, there exist a mapping $\tau : \Pi \to \tilde{\Pi}^S$, a belief scheme $\tilde{\mu}$ and a continuation equilibrium $\tilde{\eta}$ relative to $\tilde{\mu}$ on Γ_S such that:

   a) η and $\tilde{\eta}$ induce the same outcome: $\tau(\pi)(\tilde{\eta}(\cdot, \tau(\pi))) = \pi(\eta(\cdot, \pi))$

   b) μ and $\tilde{\mu}$ induce same beliefs when $\tau^{-1}(\tau(\pi))$ is a singleton: $\tilde{\mu}(\tau(\pi)) = \mu(\pi)$

   c) Agent truthfully reports on Γ_S: For every $\tilde{\pi} \in \tau(\Pi)$, $\tilde{\eta}_i(\theta, \tilde{\pi}) = \theta$

ii) If F can be implemented by a PBE (σ, η, μ) on Γ_I, then F can be implemented by a PBE $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ on Γ_S, where the agent truthfully reports.

iii) If $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ is a PBE of Γ_S where the agent truthfully reports, then for any extension of $\tilde{\eta}$ relative to $\tilde{\mu}$, denoted by η relative to μ using mechanisms Π, (σ, η, μ) is a PBE on Γ_I, where $\sigma_i = v(\tilde{\sigma}_i)$.

Proof

i) Consider a continuation equilibrium $\eta : \Theta \times \Pi \to M$ relative to μ on Γ_I. We need to construct a mapping $\tau : \Pi \to \tilde{\Pi}^S$, a belief scheme $\tilde{\mu} : \tilde{\Pi}^s \to \Delta T$ and a continuation equilibrium $\tilde{\eta} : \Theta \times \tilde{\Pi}^S \to \Theta^N$ relative to $\tilde{\mu}$ with the desired properties. Define the mapping $\tau(\pi) \in \tilde{\Pi}^S$ as follows: $\tau(\pi)(\theta) = \pi(\eta(\theta, \pi))$. This is well defined since because our assumptions on U, $\eta_i(\theta, \pi)$ is independent of π_{-i} and any possible belief μ_{-i}. Define belief updating: For every $\tilde{\pi}_i \in \tau_i(\Pi_i)$ such that $\tau_i^{-1}(\tilde{\pi}_i)$ is a singleton, define $\tilde{\mu}_i(\tilde{\pi}_i) = \mu_i(\tau_i^{-1}(\tilde{\pi}_i))$. If it is not singleton, define $\tilde{\mu}_i(\tilde{\pi}_i) = \sum_{j:\tau_i(\pi_i^j)=\tilde{\pi}_i} \alpha_j \mu_i(\pi_i^j)$ with $\alpha_j = \frac{\mu_0^i(\text{supp } (\pi_i^j))}{\sum_{k:\tau_i(\pi_i^k)=\tilde{\pi}_i} \mu_0^i(\text{supp } (\pi_i^k))}$. Define continuation equilibrium: for every $\tilde{\pi}_i \in \tau_i(\Pi_i)$, $\tilde{\eta}_i(\theta, \tilde{\pi}) = \theta$. Note that we have no restriction on continuation equilibrium and updating beliefs for $\tilde{\pi}_i \notin \tau_i(\Pi_i)$. For any $\tilde{\pi}_i \notin \tau_i(\Pi_i)$, there is at least one $\pi_i \in \Pi_i$ such that $\text{Image}(\pi_i) = \text{Image}(\tilde{\pi}_i)$ (this always exists because *strong richness condition*). Then define $\tilde{\eta}_i(\theta, \tilde{\pi}) = \eta_i(\theta, \pi)$ and $\tilde{\mu}_i(\tilde{\pi}_i) = \mu_i(\pi_i)$. Using these definitions, we have the following properties:

   a) η and $\tilde{\eta}$ induce the same outcome since $\tau(\pi)(\tilde{\eta}(\theta, \tau(\pi))) = \tau(\pi)(\theta) = \pi(\eta(\theta, \pi))$.

   b) μ and $\tilde{\mu}$ induce the same beliefs by definition of $\tilde{\mu}$.

   c) Agent truthfully report on Γ_S by definition of $\tilde{\eta}$.

It is only left to show that $\tilde{\eta}$ is a continuation equilibrium relative to $\tilde{\mu}$. We need to show that for every θ and $\tilde{\pi}$, $\tilde{\eta}(\theta, \tilde{\pi}) = arg\max_{\theta' \in \Theta^N} \mathbb{E}_{\tilde{\mu}(\tilde{\pi})}\left[U(\theta, t, \tilde{\pi}(\theta'))\right]$. Consider $\tilde{\pi} \in \tau(\Pi)$. We have that:

$$\begin{aligned}
\mathbb{E}_{\tilde{\mu}(\tilde{\pi})}\left[U(\theta, t, \tilde{\pi}(\theta'))\right] &= \sum_j \alpha_j \mathbb{E}_{\mu(\pi^j)}\left[U(\theta, t, \tau(\pi^j)(\theta'))\right] \\
&= \sum_j \alpha_j \mathbb{E}_{\mu(\pi^j)}\left[U(\theta, t, \pi^j(\eta(\theta', \pi^j)))\right]
\end{aligned}$$

Since η is a continuation equilibrium, $\eta(\theta, \pi^j) = arg\max_{m \in M} \mathbb{E}_{\mu(\pi^j)}\left[U(\theta, t, \pi^j(m))\right]$, and then $\tilde{\eta}_i(\theta, \tilde{\pi}) = \theta$. For $\tilde{\pi} \notin \tau(\Pi)$ the proof is analogue.

ii) Suppose F is implemented by PBE (σ, η, μ). We will construct $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ that implements F on Γ_S. Since η is a continuation equilibrium relative to μ on Γ_I, we apply part i) and obtain updating scheme $\tilde{\mu}$ and a continuation equilibrium $\tilde{\eta}$ relative to $\tilde{\mu}$ on Γ_S. We also have the existence of the mapping τ. For each principal i, define $\tilde{\sigma}_i(t_i) = \tau_i(\sigma_i(t_i))$. Note that by construction, the agent reports truthfully for any on path $\tilde{\sigma}$. Thus, we have already defined $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$. It is left to show that $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ implements F and that it is a PBE. The first is easy since by part i) a) the agent will induce the same allocation in both games. For the second, we first check that $\tilde{\mu}$ is Bayesian consistent with respect to $\tilde{\sigma}$. First, since μ is Bayesian consistent with respect to σ, for any π_i, the updated belief is the following:

$$\mu_i(\pi_i)(t_i) = \frac{\mu_0(t_i)}{\mu_0(\text{supp }(\pi_i))} \mathbb{1}_{\text{supp }(\pi_i)}(t_i)$$

Now, for any $\tilde{\pi}_i \in \Pi_i$ we have by definition:

$$\begin{aligned}
\tilde{\mu}_i(\tilde{\pi}_i)(t_i) &= \sum_{j \in J} \mu_i(\pi_i^j)(t_i) \frac{\mu_0(\text{supp }(\pi_i^j))}{\sum_{j \in J} \mu_0(\text{supp }(\pi_i^j))} \\
&= \sum_{j \in J} \frac{\mu_0(t_i)}{\mu_0(\text{supp }(\pi_i^j))} \mathbb{1}_{\text{supp }(\pi_i^j)}(t_i) \frac{\mu_0(\text{supp }(\pi_i^j))}{\sum_{j \in J} \mu_0(\text{supp }(\pi_i^j))} \\
&= \frac{\mu_0(t_i)}{\sum_{j \in J} \mu_0(\text{supp }(\pi_i^j))} \sum_{j \in J} \mathbb{1}_{\text{supp }(\pi_i^j)}(t_i) \\
&= \frac{\mu_0(t_i)}{\mu_0(\text{supp }(\tilde{\pi}_i))} \mathbb{1}_{\text{supp }(\tilde{\pi}_i)}(t_i)
\end{aligned}$$

Where in the last equality we use that $\sum_{j\in J} \mu_0(\text{supp }(\pi_i^j)) = \mu_0(\text{supp }(\tilde{\pi}_i))$ and $\sum_{j\in J} \mathbb{1}_{\text{supp }(\pi_i^j)}(t_i) = \mathbb{1}_{\text{supp }(\tilde{\pi}_i)}(t_i)$. Then, $\tilde{\mu}$ is Bayesian consistent with respect to $\tilde{\sigma}$. Finally, we need to show that $\tilde{\sigma}$ is optimal for the principals.

Consider principal i's payoffs of type t_i of choosing mechanism $\tilde{\pi}_i$. Using the assumptions on U, $\eta_j(\theta, \pi)$ is independent of π_{-j}. Using that allocations are the same after σ_i and $\tilde{\sigma}_i$, we have the following equalities:

$$
\begin{aligned}
&V_i(t_i, \tilde{\pi}_i|\tilde{\sigma}_{-i}) \\
&= \int_{\Theta\times T_{-i}} V_i\Big(t, \theta, \tilde{\pi}_i\Big(\tilde{\eta}_i(\theta, \tilde{\pi}_i, \tilde{\sigma}_{-i}(t_{-i}))\Big), \tilde{\sigma}_{-i}(t_{-i})\Big(\tilde{\eta}_{-i}(\theta, \tilde{\pi}_i, \tilde{\sigma}_{-i}(t_{-i}))\Big)\Big) d\mu_0^{-i}(t_{-i}) d\lambda(\theta) \\
&= \int_{\Theta\times T_{-i}} V_i\Big(t, \theta, \tilde{\pi}_i\Big(\tilde{\eta}_i(\theta, \tilde{\pi}_i)\Big), \tilde{\sigma}_{-i}(t_{-i})\Big(\tilde{\eta}_{-i}(\theta, \tilde{\sigma}_{-i}(t_{-i}))\Big)\Big) d\mu_0^{-i}(t_{-i}) d\lambda(\theta) \\
&= \int_{\Theta\times T_{-i}} V_i\Big(t, \theta, \tilde{\pi}_i\Big(\tilde{\eta}_i(\theta, \tilde{\pi}_i)\Big), \sigma_{-i}(t_{-i})\Big(\eta_{-i}(\theta, \sigma_{-i}(t_{-i}))\Big)\Big) d\mu_0^{-i}(t_{-i}) d\lambda(\theta)
\end{aligned}
$$

Note that $V_i(t_i, \tilde{\pi}_i|\tilde{\sigma}_{-i}) = V_i(t_i, \tilde{\pi}_i|\sigma_{-i})$ and then $V_i(t_i, \tilde{\sigma}_i(t_i)|\tilde{\sigma}_{-i}) = V_i(t_i, \sigma_i(t_i)|\sigma_{-i})$. We need to show that for every i, t_i and $\tilde{\pi}_i$,

$$V_i(t_i, \tilde{\sigma}(t_i)|\tilde{\sigma}_{-i}) \geq V_i(t_i, \tilde{\pi}_i|\tilde{\sigma}_{-i})$$

Suppose not. Then there is $\tilde{\pi}_i' \neq \tilde{\sigma}_i(t_i)$ such that:

$$V_i(t_i, \tilde{\sigma}(t_i)|\tilde{\sigma}_{-i}) < V_i(t_i, \tilde{\pi}_i'|\tilde{\sigma}_{-i})$$

Or equivalently:

$$V_i(t_i, \sigma(t_i)|\sigma_{-i}) < V_i(t_i, \tilde{\pi}_i'|\sigma_{-i})$$

In case $\tilde{\pi}_i' \in \tau_i(\Pi_i)$, we know that there is π_i' such that $\tilde{\pi}_i'(\tilde{\eta}_i(\theta, \tilde{\pi}_i')) = \pi_i'(\eta_i(\theta, \pi_i'))$. Thus $V_i(t_i, \tilde{\pi}_i'|\sigma_{-i}) = V_i(t_i, \pi_i'|\sigma_{-i})$ which implies that there is $\pi_i' \neq \sigma_i(t_i)$ such that:

$$V_i(t_i, \sigma(t_i)|\sigma_{-i}) < V_i(t_i, \pi_i'|\sigma_{-i})$$

This directly contradicts that (σ, η, μ) is a PBE on Γ_I. In the other case $\tilde{\pi}_i' \notin \tau_i(\Pi_i)$ the same argument applies.

iii) Consider a PBE $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ of Γ_S where the agent reports truthfully. Consider any extension η relative to μ on Γ_I. We have to show that (σ, η, μ) is a PBE on Γ_I. First, by definition η is a continuation equilibrium relative to μ. By the assumption that v is a one to one mapping, $\tilde{\mu}$ is Bayesian consistent implies that μ is Bayesian consistent. It is left to show that σ is optimal for the principals. Suppose not. Then there is $\pi'_i \neq \sigma_i(t_i)$ such that:

$$V_i(t_i, \sigma_i(t_i)|\sigma_{-i}) < V_i(t_i, \pi'_i|\sigma_{-i})$$

Following the same logic that part ii), we have that $V_i(t_i, \pi_i|\sigma_{-i}) = V_i(t_i, \pi_i|\tilde{\sigma}_{-i})$ and then $V_i(t_i, \sigma_i(t_i)|\sigma_{-i}) = V_i(t_i, \tilde{\sigma}_i(t_i)|\tilde{\sigma}_{-i})$. Then, equivalently:

$$V_i(t_i, \tilde{\sigma}_i(t_i)|\tilde{\sigma}_{-i}) < V_i(t_i, \pi'_i|\tilde{\sigma}_{-i})$$

In case $\pi'_i \in v_i(\tilde{\Pi}_i)$, we know that there is $\tilde{\pi}'_i$ such that $\tilde{\pi}'_i(\tilde{\eta}_i(\theta, \tilde{\pi}'_i)) = \pi'_i(\eta_i(\theta, \pi'_i))$. Thus $V_i(t_i, \pi'_i|\tilde{\sigma}_{-i}) = V_i(t_i, \tilde{\pi}'_i|\tilde{\sigma}_{-i})$ which implies that there is $\tilde{\pi}'_i \neq \tilde{\sigma}_i(t_i)$ such that:

$$V_i(t_i, \tilde{\sigma}(t_i)|\tilde{\sigma}_{-i}) < V_i(t_i, \tilde{\pi}'_i|\tilde{\sigma}_{-i})$$

This directly contradicts that $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ is a PBE on Γ_S. In the other case $\pi'_i \notin v_i(\tilde{\Pi}_i)$. We have two possibilities. First, $|\pi'_i(\eta_i(\Theta, \pi'_i))| \leq |\Theta|$. In this case define $\tilde{\pi}'_i(\theta) = \pi'_i(\eta_i(\theta, \pi'_i))$. Since $\tilde{\pi}'_i \in \tilde{\Pi}_i$ and the agent is truthful in Γ_S, $\tilde{\pi}'_i(\tilde{\eta}_i(\theta, \tilde{\pi}'_i)) = \pi'_i(\eta_i(\theta, \pi'_i))$. This implies the same previous inequality which is a contradiction with the PBE $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$. The second possibility is that $|\pi'_i(\eta_i(\Theta, \pi'_i))| > |\Theta|$. For this to be true it must be that $|\eta_i(\Theta, \pi'_i)| > |\Theta|$ which is a contradiction.

Theorem 4.2: If U is *weakly outcome separable*, satisfies *no indifference* and Γ_I satisfies *weak richness condition*, then

i) For every *continuation equilibrium* η relative to μ on Γ_I, there exist a mapping $\tau : \Pi \to \tilde{\Pi}^W$, a belief scheme $\tilde{\mu}$ and a continuation equilibrium $\tilde{\eta}$ relative to $\tilde{\mu}$ on Γ_W such that:

  a) η and $\tilde{\eta}$ induce the same outcome: $\tau(\pi)(\tilde{\eta}(\cdot, \tau(\pi))) = \pi(\eta(\cdot, \pi))$

  b) μ and $\tilde{\mu}$ induce same beliefs when $\tau^{-1}(\tau(\pi))$ is a singleton: $\tilde{\mu}(\tau(\pi)) = \mu(\pi)$

  c) Agent truthfully reports on Γ_W: For every $\tilde{\pi} \in \tau(\Pi)$, $\tilde{\eta}_i(\theta, \tilde{\pi}) = (\theta, t_{-i})$, $t_{-i} \in \text{supp } \tilde{\mu}_{-i}(\tilde{\pi}_{-i})$

ii) If F can be implemented by a PBE (σ, η, μ) on Γ_I with $\mu_i(\sigma_i(t_i)) = \delta_{t_i}$, then F can be implemented by a PBE $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ on Γ_W, where the agent truthfully reports.

iii) If $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ is a PBE of Γ_W with $\tilde{\mu} \in F$ where the agent truthfully reports, then for any extension of $\tilde{\eta}$ relative to $\tilde{\mu}$, denoted by η relative to μ using mechanisms Π, (σ, η, μ) is a PBE on Γ_I, where $\sigma_i = v(\tilde{\sigma}_i)$.

Proof The proof is analogue of Theorem 4.1 with the following changes:

i) We need to construct a mapping $\tau : \Pi \to \tilde{\Pi}^W$, a belief scheme $\tilde{\mu} : \tilde{\Pi}^W \to \Delta T$ and a continuation equilibrium $\tilde{\eta} : \Theta \times \tilde{\Pi}^W \to \Theta^N$ relative to $\tilde{\mu}$ with the desired properties. Define the mapping $\tau(\pi) \in \tilde{\pi}$ as following: $\tau(\pi)(\theta, t_{-i}) = \pi(\eta(\theta, \pi))$ whenever supp $\mu_{-i}(\pi_{-i}) = t_{-i}$. This is well defined since because our assumptions on U and the fact that beliefs are degenerate, $\eta_i(\theta, \pi)$ is independent of π_{-i} directly. Define continuation equilibrium: for every $\tilde{\pi}_i \in \tau_i(\Pi_i)$ with $\mu_{-i}(\tilde{\pi}_{-i}) = t_{-i}$, $\tilde{\eta}_i(\theta, \tilde{\pi}) = (\theta, t_{-i})$. Note that we have no restriction on continuation equilibrium and updating beliefs for $\tilde{\pi}_i \notin \tau_i(\Pi_i)$. For any $\tilde{\pi}_i \notin \tau_i(\Pi_i)$, there is at least one $\pi_i \in \Pi_i$ such that $\text{Image}(\pi_i) = \text{Image}(\tilde{\pi}_i)$ (this always exists because *weak richness condition*). Then define $\tilde{\eta}_i(\theta, \tilde{\pi}) = \eta_i(\theta, \pi)$ and $\tilde{\mu}_i(\tilde{\pi}_i) = \mu_i(\pi_i)$. The rest follows the same proof than Theorem 4.1.

ii) Using the new mapping τ, continuation equilibrium and belief scheme defined in i), the construction of PBE on Γ_W follows from the same arguments of Theorem 4.1.

iii) The same proof than Theorem 4.1 applies except where we want to contradicts the existance of $\pi'_i \notin v_i(\tilde{\Pi}_i)$ such that:

$$V_i(t_i, \tilde{\sigma}_i(t_i)|\tilde{\sigma}_{-i}) < V_i(t_i, \pi'_i|\tilde{\sigma}_{-i})$$

We have two possibilities. First, $|\pi'_i(\eta_i(\Theta, \pi'_i, \Pi_{-i}))| \leq |\Theta \times T_{-i}|$. In this case define $\tilde{\pi}'_i(\theta, t_{-i}) = \pi'_i(\eta_i(\theta, \pi'_i, \pi'_{-i}))$ whenever supp $\mu_{-i}(\pi'_{-i}) = t_{-i}$. Since $\tilde{\pi}'_i \in \tilde{\Pi}_i$ and the agent is truthful in Γ_W, $\tilde{\pi}'_i(\tilde{\eta}_i(\theta, \tilde{\pi}')) = \pi'_i(\eta_i(\theta, \pi'))$. This implies the same previous inequality which is a contradiction with the PBE $(\tilde{\sigma}, \tilde{\mu}, \tilde{\mu})$. The second possibility is that $|\pi'_i(\eta_i(\Theta, \pi'_i, \Pi_{-i}))| > |\Theta \times T_{-i}|$. For this to be true it must be that $|\eta_i(\Theta, \pi'_i, \Pi_{-i})| > |\Theta \times T_{-i}|$ which is a contradiction.

Proposition 4.2: If a game Γ_I has no *latent contracts*, then U satisfies *weakly outcome separability*.

Proof Suppose not. Assume we have only two principals, $Y_1 = \{y_1, y_1'\}$, $Y_2 = \{y_2, y_2'\}$ and there is no private information. Then, we can assume WLOG that since U does not satisfy *weakly outcome separability*,

$$U(y_1, y_2) > U(y_1', y_2)$$

and

$$U(y_1, y_2') \leq U(y_1', y_2')$$

By the no indifference condition, we must have that $U(y_1, y_2') < U(y_1', y_2')$. Consider the following preferences. For principal one, $V_1(y_1, y_2^*) > V_1(y_1', y_2^*)$ for every $y_2^* \in Y_2$, and $V_1(y_1, y_2) > V_1(y_1, y_2')$. For principal two, $V_2(y_1^*, y_2') > V_2(y_1^*, y_2)$ for every $y_1^* \in Y_1$ and $V_2(y_1, y_2) > V_2(y_1', y_2')$. Then, $(\pi_1(m_1), \pi_1(m_1')) = (y_1, y_1')$, and $\pi_2 = y_2$ is an equilibrium of the game, where y_1' is a *latent contract*. This is a contradiction.

Theorem 4.3: If U is *outcome separable*, satisfies *no indifference* and Γ_I satisfies *weak richness condition*, then

i) For every *continuation equilibrium* η relative to μ on Γ_I, there exist a mapping $\tau : \Pi \to \tilde{\Pi}^W$, a belief scheme $\tilde{\mu}$ and a continuation equilibrium $\tilde{\eta}$ relative to $\tilde{\mu}$ on Γ_W such that:

   a) η and $\tilde{\eta}$ induce the same outcome: $\tau(\pi)(\tilde{\eta}(\cdot, \tau(\pi))) = \pi(\eta(\cdot, \pi))$

   b) μ and $\tilde{\mu}$ induce same beliefs when $\tau^{-1}(\tau(\pi))$ is a singleton: $\tilde{\mu}(\tau(\pi)) = \mu(\pi)$

   c) Agent truthfully reports on Γ_W: For every $\tilde{\pi} \in \tau(\Pi)$, $\tilde{\eta}_i(\theta, \tilde{\pi}) = (\theta, t_{-i})$, $t_{-i} \in \text{supp } \tilde{\mu}_{-i}(\tilde{\pi}_{-i})$

ii) If F can be implemented by a PBE (σ, η, μ) on Γ_I, then F can be implemented by a PBE $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ on Γ_W, where the agent truthfully reports.

iii) If $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ is a PBE of Γ_W where the agent truthfully reports, then for any extension of $\tilde{\eta}$ relative to $\tilde{\mu}$, denoted by η relative to μ using mechanisms Π, (σ, η, μ) is a PBE on Γ_I, where $\sigma_i = v(\tilde{\sigma}_i)$.

Proof The proof is analogue of Theorem 4.1 with the following changes:

i) We need to construct a mapping $\tau : \Pi \to \tilde{\Pi}^W$, a belief scheme $\tilde{\mu} : \tilde{\Pi}^W \to \Delta T$ and a continuation equilibrium $\tilde{\eta} : \Theta \times \tilde{\Pi}^W \to \Theta^N$ relative to $\tilde{\mu}$ with the desired properties. Define the mapping $\tau(\pi) \in \tilde{\pi}$ as following: $\tau(\pi)(\theta, t_{-i}) = \pi(\tau(\theta, \pi))$ whenever $t_{-i} \in \text{supp } \mu_{-i}(\pi_{-i})$. This is well defined since because our assumptions on U and the fact that beliefs are degenerate, $\eta_i(\theta, \pi)$ is independent of π_{-i} directly. Thus the value $\tilde{\pi}_i(\theta, t_{-i})$ only depends on π_i (and the possibly induced belief μ_i), θ and the degenerate belief μ_{-i}. Define continuation equilibrium: for every $\tilde{\pi} \in \tau(\Pi)$ with $\mu_{-i}(\tilde{\pi}_{-i}) = t_{-i}$, $\tilde{\eta}_i(\theta, \tilde{\pi}) = (\theta, t_{-i})$. Note that we have no restriction on continuation equilibrium and updating beliefs for $\tilde{\pi}_i \notin \tau_i(\Pi_i)$. For any $\tilde{\pi}_i \notin \tau_i(\Pi_i)$, there is at least one $\pi_i \in \Pi_i$ such that $\text{Image}(\pi_i) = \text{Image}(\tilde{\pi}_i)$ (this always exists because *weak richness condition*). Then define $\tilde{\eta}(\theta, \tilde{\pi}) = \eta(\theta, \pi)$ and $\tilde{\mu}_i(\tilde{\pi}_i) = \mu_i(\pi_i)$. The rest follows the same proof than Theorem 4.1.

ii) Using the new mapping τ, continuation equilibrium and belief scheme defined in i), the construction of PBE on Γ_W follows from the same arguments of Theorem 4.1.

iii) The same proof from Theorem 4.1 applies except where we want to contradicts the existance of $\pi'_i \notin v_i(\tilde{\Pi}_i)$ such that:

$$V_i(t_i, \tilde{\sigma}_i(t_i)|\tilde{\sigma}_{-i}) < V_i(t_i, \pi'_i|\tilde{\sigma}_{-i})$$

We have two possibilities. First, $|\pi'_i(\eta_i(\Theta, \pi'_i, \Pi_{-i}))| \leq |\Theta \times T_{-i}|$. In this case define $\tilde{\pi}'_i(\theta, t_{-i}) = \pi'_i(\eta_i(\theta, \pi'_i, \pi'_{-i}))$ whenever $\text{supp } \mu_{-i}(\pi'_{-i}) = t_{-i}$. Since $\tilde{\pi}'_i \in \tilde{\Pi}_i$ and the agent is truthful in Γ_W, $\tilde{\pi}'_i(\tilde{\eta}_i(\theta, \tilde{\pi}')) = \pi'_i(\eta_i(\theta, \pi'))$. This implies the same previous inequality which is a contradiction with the PBE $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$. The second possibility is that $|\pi'_i(\eta_i(\Theta, \pi'_i, \Pi_{-i}))| > |\Theta \times T_{-i}|$. For this to be true it must be that $|\eta_i(\Theta, \pi'_i, \Pi_{-i})| > |\Theta \times T_{-i}|$ which is a contradiction.

B Proofs of Section 7

B.1 Expert Delegation

Proposition 7.1: There is a fully revealing PBE (σ, η, μ) on Γ_I that implements the *outcome function F*. To implement this outcome, it is without loss of generality to assume that each principal delegates a **closed interval** of policy decisions to the agent.

Proof Since this game satisfies the assumptions of Theorem 4.3, it is enough to consider the direct mechanism game Γ_W. Consider the following principal i strategy: $\tilde{\sigma}_i(t_i)(t_j) = \frac{t_i+t_j}{2}$. This is a fully revealing strategy for each principal since $\tilde{\sigma}_i(t_i) \neq \tilde{\sigma}_i(t'_i)$ for every $t_i \neq t'_i \in T_i$. Thus $\tilde{\mu}_i(\tilde{\sigma}_i(t_i)) = \delta_{t_i}$. Note that each $\tilde{\sigma}_i(t_i)$ is *incentive compatible* by Lemma 7.1. Thus, $\tilde{\eta}_i(\sigma_i(t_i), \sigma_j(t_j)) = t_j$. For every off path $\tilde{\pi}_i$, define $\tilde{\mu}_i(\tilde{\pi}_i) = \delta_{-1}$ and $\tilde{\eta}_i(\tilde{\pi}_i, \tilde{\pi}_j) = \mathit{arg}\max_{t'_j \in T_j} U_{\mu(\tilde{\pi})}(t, \tilde{\pi}_i(t'_j), \tilde{\pi}_j(t'_i))$. Checking that $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ is a PBE is direct from the fact that each V_i and U is maximized for every $t \in T$. Since the agent reports truthfully to each principal, $y_1 = y_2 = w$, and then this PBE implements F. Note that the direct mechanism $\tilde{\sigma}_i(t_i)(t_j) = \frac{t_i+t_j}{2}$ is equivalent to delegate the decision to the agent restricted to the set $\left[\frac{t_i+\underline{t}_j}{2}, \frac{t_i+\bar{t}_j}{2}\right]$. Thus it is without loss of generality to assume each principal delegates a closed interval of policy decisions.

B.2 Manufacturer Competition

Proposition 7.2: There is a unique outcome implemented by a fully revealing PBE (σ, η, μ) on Γ_I. To implement this outcome, it is without loss of generality to assume that each principal chooses only a **price**.

Proof Consider payoffs of type t_1, which is perceived as type t'_1 and offers *incentive compatible* mechanism $\tilde{\pi}_1 = (\tilde{x}_1, \tilde{p}_1)$:

$$\int_{T_2} \left(\left(\frac{t'_1 + t_2}{2} \right) \tilde{x}_1(t_2) - \frac{1}{2} \int_{\underline{t}_2}^{t_2} \tilde{x}_1(y)dy - t_1 \tilde{x}_1(t_2) \right) dF(t_2) + \left[\tilde{p}(\underline{t}_2) - \left(\frac{t'_1 + \underline{t}_2}{2} \right) \tilde{x}(\underline{t}_2) \right]$$

A fully revealing strategy using *incentive compatible* mechanisms is an injective function $\tilde{\sigma}_1 : T_1 \to \tilde{\Pi}_1$ such that $\tilde{\sigma}^x(t_1)(t_2)$ is increasing as a function of t_2 for every $t_1 \in T_1$ and:

$$t_1 \in arg\max_{t_1'} \int_{T_2} \left(\left(\frac{t_1' + t_2}{2} \right) \tilde{\sigma}_1^x(t_1')(t_2) - \frac{1}{2} \int_{\underline{t}_2}^{t_2} \tilde{\sigma}_1^x(t_1')(y)dy - t_1 \tilde{\sigma}_1^x(t_1')(t_2) \right) dF(t_2) \\ + \left[\tilde{\sigma}_1^p(t_1')(\underline{t}_2) - \left(\frac{t_1' + \underline{t}_2}{2} \right) \tilde{\sigma}_1^x(t_1')(\underline{t}_2) \right]$$

Equivalently, after some manipulations:

$$t_1 \in arg\max_{t_1'} \frac{1}{2} \int_{T_2} \tilde{\sigma}_1^x(t_1')(t_2) \left(\left(t_2 - \frac{1 - F(t_2)}{f(t_2)} \right) + (t_1' - 2t_1) \right) dF(t_2) \\ + \left[\tilde{\sigma}_1^p(t_1')(\underline{t}_2) - \left(\frac{t_1' + \underline{t}_2}{2} \right) \tilde{\sigma}_1^x(t_1')(\underline{t}_2) \right]$$

Denote $\hat{t}_2(t_1)$ an increasing function of t_1. Consider the following particular selling strategy

$$\tilde{\sigma}_1^x(t_1)(t_2) = \begin{cases} 1 & \text{if } t_2 \geq \hat{t}_2(t_1) \\ 0 & \text{if } t_2 < \hat{t}_2(t_1). \end{cases}$$

Together with the induced $\tilde{\sigma}_1^p(t_1)(t_2)$ from Lemma 7.2, it is *incentive compatible* and fully revealing. Using this strategy and considering the participation constraint of the buyer, the maximization problem can be rewritten as follows:

$$t_1 \in arg\max_{t_1'} \frac{1}{2} \int_{\hat{t}_2(t_1')}^{\bar{t}_2} \left(\left(t_2 - \frac{1 - F(t_2)}{f(t_2)} \right) + (t_1' - 2t_1) \right) dF(t_2)$$

Define $J(t_2) = t_2 - \frac{1-F(t_2)}{f(t_2)}$. Using analogue arguments from Cai et. al (2007)[20], there is a unique $\hat{t}_2$ that satisfies the previous maximization problem. The inverse $t_1(t_2)$ is given by the following ODE:

$$\frac{\partial t_1}{\partial t_2} = \frac{\left[J(t_2) - t_1(t_2) \right] f(t_2)}{1 - F(t_2)}$$

[20]They study a single privately informed auctioneer problem which have to set a reserve price. In our case, since there is only one agent, the reserve price corresponds to post a price.

with initial condition $(t_2(\underline{t}_1), \underline{t}_1)$, where:

$$t_2(\underline{t}_1) = \begin{cases} J^{-1}(\underline{t}_1) & \text{if } J(\underline{t}_2) \leq \underline{t}_1 \leq J(\bar{t}_2) \\ \underline{t}_2 & \text{if } \underline{t}_1 < J(\underline{t}_2) \\ \bar{t}_2 & \text{if } J(\bar{t}_2) < \underline{t}_1 \end{cases}$$

Define $\tilde{\sigma}_i(t_i)$ as the previous constructed strategy. Since it is fully revealing, $\tilde{\mu}_i(\tilde{\sigma}_i(t_i)) = \delta_{t_i}$. Also it is *incentive compatible*, thus $\tilde{\eta}_i(\sigma_i(t_i), \sigma_j(t_j)) = t_j$. For every off path $\tilde{\pi}_i$, define $\tilde{\mu}_i(\tilde{\pi}_i) = \delta_{-1}$ and $\tilde{\eta}_i(\tilde{\pi}_i, \tilde{\pi}_j) = arg\max_{t'_j \in T_j} U_{\mu(\tilde{\pi})}(t, \tilde{\pi}_i(t'_j), \tilde{\pi}_j(t'_i))$. Checking that $(\tilde{\sigma}, \tilde{\eta}, \tilde{\mu})$ is a PBE is direct from the previous arguments.

References

Attar, A., Majumdar, D., Piaser, G., & Porteiro, N. (2008). Common agency games: Indifference and separable preferences. *Mathematical Social Sciences*, 56(1), 75-95.

Attar, A., Mariotti, T., & Salanié, F. (2011). Nonexclusive competition in the market for lemons. *Econometrica*, 79(6), 1869-1918.

Bernheim, B. D., & Whinston, M. D. (1986). Menu auctions, resource allocation, and economic influence. *The quarterly journal of economics*, 101(1), 1-31.

Bernheim, B. D., & Whinston, M. D. (1998). Exclusive dealing. *Journal of Political Economy*, 106(1), 64-103.

Cai, H., Riley, J., & Ye, L. (2007). Reserve price signaling. *Journal of Economic Theory*, 135(1), 253-268.

Epstein, L. G., & Peters, M. (1999). A revelation principle for competing mechanisms. *Journal of Economic Theory*, 88(1), 119-160.

Galperti, S. (2015). Common agency with informed principals: Menus and signals. *Journal of Economic Theory*, 157, 648-667.

Grossman, G., & Helpman, E. (1994). Protection for Sale. *The American Economic Review*, 84(4), 833-850.

Dixit, A., Grossman, G. M., & Helpman, E. (1997). Common agency and coordination: General theory and application to government policy making. *Journal of political economy*, 105(4), 752-769.

Laussel, D., & Le Breton, M. (1998). Efficient private production of public goods under common agency. *Games and Economic Behavior*, 25(2), 194-218.

Le Breton, M., & Salanié, F. (2003). Lobbying under political uncertainty. *Journal of Public Economics*, 87(12), 2589-2610.

Lima, R. C., & Moreira, H. (2014). Information transmission and inefficient lobbying. *Games and Economic Behavior*, 86, 282-307.

Martimort, D., & Moreira, H. (2010). Common agency and public good provision under asymmetric information. *Theoretical Economics*, 5(2), 159-213.

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

Martimort, D., & Stole, L. (2002). The revelation and delegation principles in common agency games. *Econometrica*, 70(4), 1659-1673.

Myerson, R. B. (1979). Incentive compatibility and the bargaining problem. *Econometrica* 61-73.

Peters, M. (2001). Common agency and the revelation principle. *Econometrica*, 69(5), 1349-1372.

Peters, M. (2003). Negotiation and take it or leave it in common agency. *Journal of Economic Theory*, 111(1), 88-109.

Peters, M. (2007). Erratum to Negotiation and take it or leave it in common agency:[Journal of Economic Theory 111 (2003) 88109]. *Journal of Economic Theory*, 135(1), 594-595.

Spence, M. (1973). Job Market Signaling. *The Quarterly Journal of Economics*, 87(3), 355-374.