



CMS-EMS
Center for Mathematical Studies in Economics
And Management Science

Discussion Paper #1541

Common Agency with Informed Principals:
Menus *and* Signals

Simone Galperti
Northwestern University
June 6, 2012

JEL Classification: Common agency, informed principals, Inscrutability Principals, Extended Taxation Principle, menus, signals, direct mechanisms.

Keywords: C72, D82, D83, D86



NORTHWESTERN
UNIVERSITY



Common Agency with Informed Principals: Menus *and* Signals

Simone Galperti*

Northwestern University

June 6, 2012

Abstract

In this paper I consider games in which multiple informed principals simultaneously compete to influence the decisions of a common agent. I focus on the problem of characterizing the equilibrium outcomes of such games. I first show that, to solve this problem, one can invoke neither Myerson's Inscrutability Principle, which holds in agency games with *one* informed principal, nor the Extended Taxation Principle, which holds in common-agency games with *uninformed* principals. I then provide two characterizations of the equilibrium outcomes: one for games in which the principals delegate the final decisions to the agent, and one for games in which they participate with the agent in making such decisions.

KEYWORDS: Common agency, informed principals, Inscrutability Principle, Extended Taxation Principle, menus, signals, direct mechanisms.

JEL CLASSIFICATION: C72, D82, D83, D86

*Department of Economics, Northwestern University, Evanston, IL 60208 (E-mail: simone-galperti2008@u.northwestern.edu). I am indebted to Eddie Dekel and Alessandro Pavan for long and fruitful discussions that greatly improved the paper. I also thank Jeffrey Ely, Ron Siegel, Balazs Szentes, Asher Wolinsky, Michael Whinston, and seminar participants at Northwestern University. I gratefully acknowledge financial support from the Center of Economic Theory of the Weinberg College of Arts and Sciences of Northwestern University. All remaining errors are mine.

1 Introduction

Common-agency games model settings in which multiple principals non-cooperatively contract with a single agent.¹ These games have proved useful to study many economic problems, such as the possibility of tacit collusion and its implications for antitrust regulation, the consequences of multilateral lobbying on government policies, and the optimality of common vs. exclusive retailers in wholesale trade. Until very recently (Martimort and Moreira (2010)), the literature has focused on common-agency games (hereafter, games for short) in which principals have no private information. Many environments, however, correspond to games with informed principals. For example, Bernheim and Whinston (1986a) consider menu auctions in which the bidders (the principals) submit bidding schemes to the auctioneer (the agent) who then freely chooses a revenue-maximizing allocation. In their model, the bidders are uninformed, although it seems natural to allow for informed bidders as in standard auction settings. More recently, Martimort and Moreira (2010) consider a model of public-good provision in which contributors (the principals) privately know their valuations of such a good and offer its provider (the agent) schemes that tie contributions to production.

The present paper studies the problem of how to characterize the outcomes—defined as maps from players’ types (their information) to allocations—that can arise as equilibria of games with informed principals.² In these games, each principal commits to an arbitrary mechanism that specifies an allocation—e.g., a contribution for the public good—as a function of how she and the agent communicate with the mechanism itself. Since the principals are informed, I consider both games in which each principal can communicate with her mechanisms (as in Myerson (1983)), and games in which only the agent can do so. I refer to the first case as *non-delegation* games, and to the second case as *delegation* games.

Characterizing the outcomes of these games raises several difficulties because standard results in the literature do not apply when principals are informed. From the literature on games with uninformed principals, we know that we cannot invoke the Revelation Principle to characterize the agent’s behavior.³ This is because his behavior with one principal depends not only on his exogenous type, but also on the endogenous information that he learns by interacting with the other principals. Furthermore, in this paper I show that with informed principals two other well-known principles fail: the Inscrutability Principle (Myerson 1983) and the Extended Taxation Principle—also known as Menu Theorem (Peters (2001)) or Delegation Principle (Martimort and Stole (2002)).⁴

The Inscrutability Principle holds in non-delegation games in which there is only one principal and the principal is informed.⁵ It says that, when characterizing the outcomes of such games, there is no loss of generality in assuming that the principal commits to the same mechanism

¹As a convention, throughout the paper I use feminine pronouns for the principals and masculine ones for the agent.

²In contrast, Martimort and Moreira (2010) focus on a specific game and examine the efficiency properties of a particular class of its equilibria.

³For the Revelation Principle, see Gibbard (1973), Green and Laffont (1977), and Myerson (1979). For the inapplicability of the Revelation Principle to games with multiple principals, see Katz (1991), McAfee (1993), Peck (1997), Epstein and Peters (1999), Peters (2001), and Martimort and Stole (2002).

⁴For the Taxation Principle, see Rochet (1985) and Guesnerie (1995).

⁵As I will explain later, in delegation games the Inscrutability Principle does not apply even if there is only one informed principal.

independently of her type, thereby disregarding the possibility that she may want to choose different mechanisms to convey her information. This is not true in games with multiple informed principals. On the one hand, as noted, with multiple principals the agent’s behavior with, say, principal 1 depends not only on his exogenous but also on his endogenous information about, say, principal 2. On the other hand, with informed principals the agent’s endogenous information includes also what principal 2 conveys about her type with her choices of a mechanism. So, if principal 2 can let these choices depend on her types, she may be able to induce a behavior of the agent with principal 1 that otherwise cannot arise in equilibrium. This means that, in contrast to the games with one informed principal, in games with multiple informed principals signaling through mechanisms represents an essential strategic component.

As far as the Extended Taxation Principle is concerned, it says that one can characterize all the outcomes of a game in which *uninformed* principals compete in arbitrary mechanisms, by studying the equilibria of a simpler game in which the principals can offer only menus of allocations and delegate the final choices to the agent. This Principle is important because it provides a relatively practical framework to study all outcomes of a much more complicated game. Moreover, it represents perhaps the main solution suggested in the literature to the failure of the Revelation Principle, also because it avoids the infinite regress problem that arises if one allows each principal to use mechanisms that depend on the agent’s exogenous type as well as his endogenous information about the other principals.⁶

The Extended Taxation Principle fails for the following reason when principals are informed. As noted, to sustain some outcomes, the principals need to control the information revealed up front by their choices of mechanisms. At the same time, they also need to restrict the options that their communication with the agent can induce, by committing to mechanisms that map to restricted menus of allocations. Clearly, in a delegation game each principal can achieve the second goal by simply offering the agent her desired menu and letting him directly choose from it. However, with general mechanisms each principal can restrict the agent’s choices to the same menu using different communication protocols, which may be essential to signal her desired amount of information. On the other hand, in a non-delegation game, with general mechanisms each principal can directly participate with the agent in choosing an allocation. But this participation is impossible when the principal must delegate the agent to choose from a menu. Thus, the principal may want to restrict such a menu to the allocations that she likes better given her type, but to do so she may have to reveal more information than she wants to.

Building on these remarks about the Inscrutability and the Extended Taxation Principles, I provide two characterizations of the equilibrium outcomes of games with informed principals—one for delegation games and one for non-delegation games. In both cases I focus on (weak) Perfect Bayesian Equilibria in which, conditional on their types, the principals do not randomize over mechanisms. Equilibria with this property represent the focus of (almost) the entire literature on games with uninformed principals—to the best of my knowledge, there is no paper that actually constructs an equilibrium with random choices of mechanisms. Moreover, from an applied perspective it seems hard to imagine, for instance, that bidders in a menu auction randomize over complicated bidding schemes or that competing sellers randomize over schemes of price-quantity pairs.

⁶For more details on this point, see Peters (2001), Martimort and Stole (2002), and Pavan and Calzolari (2010).

In the case of delegation games, I show that we can recover all outcomes that arise when principals compete in arbitrary mechanisms by finding the equilibria of the following ‘signal-menu game.’ In this new game, the principals offer menus of allocations—as suggested by the Extended Taxation Principle—and also privately send a ‘cheap talk’ signal to the agent. Moreover, each principal’s set of signals equals her set of types.

To characterize the outcomes of non-delegation games, I construct a different game in which each principal now offers the agent a menu of direct mechanisms. For each menu, all DMs depend only on the principal’s reports about her type and restrict such reports to a subset of her possible types. Such menus allow the agent to act on his exogenous and endogenous information by selecting different DMs, without causing any infinite regress. At the same time, the DMs allow each principal to retain all her ability to participate with the agent in choosing an allocation. Furthermore, by letting a principal commit to reporting only a subset of her possible types, the restricted DMs allow her to signal that her true type belongs to that set. Finally, I show that we can recover all outcomes of the original non-delegation game by finding the equilibria of the new game in which, after offering a menu of DMs, each principal *truthfully* reports her type to the DM chosen by the agent.

This characterization suggests the following interpretation of non-delegation games with informed principals. In such games, it is as if each principal offers the agent several instruments—the DMs—that he can use to screen her types. After observing all offers, the agent chooses a DM for each principal, in a way that may depend on his type and on how he decides to screen the other principals. Finally, each principal truthfully reveals her type to the chosen DM. In other words, it is as if the principals accept to become ‘agents of the agent’ in a way that they partially control with their initial offers.

Related Literature: As noted, this paper relates to the literature on mechanism design with one informed principal (Myerson (1983)) and on common-agency games with uninformed principals (Peters (2001); Martimort and Stole (2002); Pavan and Calzolari (2010)).

This paper also relates to a recent literature on multilateral contracting games in which multiple, privately informed, parties sign mutual contracts (Peters (2010); Peters and Szentes (2011); Celick and Peters (2011)). This literature has suggested several characterizations of the outcomes of multilateral contracting games. These games, however, differ from common-agency games in one crucial aspect that makes those characterizations inapplicable to the settings considered in the present paper. In multilateral contracting games, each party can communicate with all other parties and design contracts that depend on others’ contracts—examples include international trade agreements and tax treaties. In common-agency games, instead, each principal can neither communicate with the other principals, nor make her mechanisms directly depend on their mechanisms—for example, because of antitrust regulation. This difference significantly changes the strategic nature of the games. In multilateral contracting games, each party can directly punish another party if, for instance, she deviates from a collusive behavior. On the other hand, in common-agency games any principal has to rely on the agent’s collaboration and superior information to punish a deviating opponent.

The paper is organized as follows. Section 2 describes the model. Section 3 is about the failure of the Inscrutability Principle. Section 4 is about the failure of the Extended Taxation Principle. Section 5 contains the main characterization results. The last section concludes. All proofs are in the appendix.

2 The Model

I begin by describing the class of common-agency games on which I focus in this paper. The novelty of this class consists in allowing the principals to have private information.

Players and Information: There are $N > 1$ principals, denoted by $i = 1, \dots, N$, and a common agent, denoted by $i = 0$. Before the game begins, each player i can privately receive some information in the form of a type $\theta_i \in \Theta_i$, where Θ_i is a finite set—as usual, $\Theta = \times_{i=0}^N \Theta_i$ and $\Theta_{-i} = \times_{j \neq i} \Theta_j$ with elements θ and θ_{-i} . I assume that $|\Theta_i| > 1$ for at least one principal i . Finally, the players commonly know that the types are distributed according to the joint probability distribution b over Θ . This distribution need not feature independence across players.

Allocations and Mechanisms: Principals have full commitment and bargaining power. Each principal offers the agent a mechanism to select an allocation for their bilateral relationship. The set of feasible allocations for principal i is the finite set X_i . For example, think of x_i as a quantity together with its price in a relationship between a buyer and a seller, or a piece of regulation together with a kickback in a relationship between a politician (the agent) and a lobbyist (the principal).

To choose an allocation, principal i and the agent communicate using a certain language, which I model with finite sets of messages.⁷ Specifically, the set of messages that principal i can send to the agent is \mathcal{P}_i , and the set of messages that she can receive from the agent is \mathcal{A}_i . I assume that $|\Theta_i| \leq |\mathcal{P}_i|$, so that principal i can communicate what she privately knows; similarly, I assume that $|\Theta| \leq |\mathcal{A}_i|$ for every $i = 1, \dots, N$, so that the agent can communicate to each principal what he privately knows as well as what he may learn about all principals by interacting with them. Formally, an indirect mechanism of principal i is a map $m_i : \mathcal{A}_i \times \mathcal{P}_i \rightarrow \mathcal{D}_i = \Delta(X_i)$ that associates to each pair of messages (a_i, p_i) a *decision* $d_i \in \mathcal{D}_i$, which is a lottery over X_i . For future reference, denote the range of a mechanism m_i by $\text{Ra}(m_i)$. The space of all indirect mechanisms of principal i is M_i .⁸

The agent can choose to interact only with some principals by rejecting the mechanisms offered by the other principals. To model this possibility, I assume that each set \mathcal{A}_i contains a ‘rejection’ message \hat{a}_i such that sending \hat{a}_i to any mechanism m_i is equivalent to rejecting it. Note that $m_i(\hat{a}_i, \cdot)$ still depends on the messages of principal i , because she may still be able to affect the final decision in \mathcal{D}_i even without the collaboration of the agent.

As is usual in common-agency games, the principals cannot directly interact among themselves. This means that, for every principal i , m_i cannot directly depend on m_{-i} , nor can x_i depend on x_{-i} ; moreover, principals cannot directly communicate among themselves.⁹

Timing: The game has two periods. In the first period, all principals simultaneously and non-cooperatively offer the agent their mechanisms. The agent observes all mechanisms, but no principal can observe the mechanisms of the other principals—i.e., this is a model of pri-

⁷Of course, the assumption that the sets of messages are finite entails some loss of generality. However, it ensures that continuation equilibria exist after any profile of mechanisms offered by the principals (see Definition 1).

⁸To simplify notation, let $M = \times_{i=1}^N M_i$, $\mathcal{A} = \times_{i=1}^N \mathcal{A}_i$, $\mathcal{P} = \times_{i=1}^N \mathcal{P}_i$, and $\mathcal{D} = \times_{i=1}^N \mathcal{D}_i$, with generic elements $m \equiv (m_1, \dots, m_N)$, $a \equiv (a_1, \dots, a_N)$, $p \equiv (p_1, \dots, p_N)$ and $d \equiv (d_1, \dots, d_N)$.

⁹As I explained in the introduction, these properties are an important difference between my model and those of Peters (2010), Peters and Szentes (2011), and Celik and Peters (2011).

vate contracting. In the second period, both the agent and each principal simultaneously,¹⁰ non-cooperatively, and privately send their messages to the respective mechanisms, which independently deliver an allocation x . The players then enjoy their payoffs and the game is over.

Strategies and Updated Beliefs: Conditional on her type θ_i , in period 1 each principal chooses a mechanism m_i , and in period 2 she chooses a message p_i given m_i . As explained in the Introduction, my analysis focusses on equilibria in which principals choose their mechanisms as a deterministic function of their types. Therefore, a (behavior) strategy of principal i is a pair of (measurable) functions (μ_i, π_i) , where $\mu_i : \Theta_i \rightarrow M_i$ and $\pi_i : \Theta_i \times M_i \rightarrow \Delta(\mathcal{P}_i)$. Hereafter, I refer to μ_i as the *offer strategy* and to π_i as the *communication strategy* of principal i . With regard to the agent, he chooses a profile of messages a as a function of his type θ_0 and the offered mechanisms m . A (behavior) strategy of the agent is then a (measurable) function $\alpha : \Theta_0 \times M \rightarrow \Delta(\mathcal{A})$. Hereafter, I refer to α as the *communication strategy* of the agent.

Finally, after observing a profile of offered mechanisms, the agent will update his belief about the principals' types. For each m and θ_0 , I denote this updated belief with the conditional probability $\beta(\cdot | \theta_0, m)$ on Θ_{-0} .

Preferences and Payoffs: Each player is an expected-utility maximizer. The vN-M utility function of principal i is $u_i : X \times \Theta \rightarrow \mathbb{R}$ and that of the agent is $v : X \times \Theta \rightarrow \mathbb{R}$. Using these functions, let $V(\alpha, \pi; \theta_0, m, \beta)$ be the expected payoff to the agent from the communication strategy profile $(\alpha, \pi) = (\alpha, (\pi_i)_{i=1}^N)$, given that his type is θ_0 , he observes m , and he updates his beliefs according to β . Similarly, let $U_i(\alpha, \pi; \theta_i, m_i, \mu_{-i})$ be the expected payoff to principal i from (α, π) , given that her type is θ_i , she offers m_i , and her opponents play μ_{-i} . Finally, let $\bar{U}_i(\mu, \alpha, \pi; \theta_i)$ be the initial expected payoff to principal i from the strategy profile $\sigma = (\mu, \alpha, \pi)$, given that her type is θ_i .

Equilibrium Concept: I denote the game that I have just described by G . I focus on (weak) PBE of G in which principals play deterministic offer strategies as I defined above. Hereafter, the term 'equilibrium of G ' refers to such a PBE.

Definition 1 (Equilibrium) *The strategy profile $\sigma^* = (\mu^*, \alpha^*, \pi^*)$ together with the belief β^* is an equilibrium of G if and only if:*

1. For every $\theta_0 \in \Theta_0$, $m \in M$, and $a \in \mathcal{A}$,

$$V(\alpha^*(\theta_0, m), \pi^*; \theta_0, m, \beta^*) \geq V(a, \pi^*; \theta_0, m, \beta^*);$$

2. For every $i = 1, \dots, N$, $\theta_i \in \Theta_i$, $m_i \in M_i$, and $p_i \in \mathcal{P}_i$,

$$U_i(\alpha^*, \pi_{-i}^*, \pi_i^*(\theta_i, m_i); \theta_i, m_i, \mu_{-i}^*) \geq U_i(\alpha^*, \pi_{-i}^*, p_i; \theta_i, m_i, \mu_{-i}^*);$$

3. For every $i = 1, \dots, N$, $\theta_i \in \Theta_i$, and $m_i \in M_i$,

$$\bar{U}_i(\mu_{-i}^*, \mu_i^*(\theta_i), \alpha^*, \pi^*; \theta_i) \geq \bar{U}_i(\mu_{-i}^*, m_i, \alpha^*, \pi^*; \theta_i);$$

4. Whenever possible, the belief β^* satisfies Bayes rule given μ^* and b .¹¹

¹⁰Although no principal i can observe m_{-i} at any stage in the game, one could consider the possibility that principals choose their messages after observing the agent's messages. Allowing for this possibility, however, adds no further generality to the analysis, as I will explain after presenting Theorem 2.

¹¹Formally, let $\mu^*(\theta_{-0}) = \times_{i=1}^N \mu_i^*(\theta_i)$ and $\mathbf{1}_{[\cdot]}$ be the indicator function of the event $[\cdot]$. Then, for every θ and

As is common in the literature, I assume that an equilibrium of G exists.¹² Also, to denote a profile of communication strategies (α, π) that satisfies conditions 1 and 2 in Definition 1 given β , I will use the term ‘continuation equilibrium’ of G . Note that, given any profile μ and offered mechanisms m , a continuation equilibrium always exists because the message spaces are finite. So, for any profile of mechanisms that the principals offer in the first period, the players’ behavior in the continuation of the game is always well defined.

Non-Delegation and Delegation Games: In the game G , each principal can send messages to her mechanisms. This possibility allows the principal to affect the final decisions, together with the agent, in the second period of the game. For this reason, hereafter I will refer to the game G as a ‘non-delegation game.’ However, in some economic environments of interest (see, e.g., Martimort and Moreira (2010)), the principals have to completely delegate the final decisions to the agent. To model this situation, I consider a class of games in which the outcomes of each mechanism can depend only on the agent’s messages.

Definition 2 (Delegation Game) *The common-agency game G is a **delegation game**—denoted by G^d —if, for every $i = 1, \dots, N$, every mechanism $m_i \in M_i^d$ satisfies (1) $m_i : \mathcal{A}_i \rightarrow \mathcal{D}_i$, and (2) $m_i(a_i) = \hat{d}_i$ for some $a_i \in \mathcal{A}_i$.*

As in the game G , each principal i must allow the agent to reject her mechanism m_i by sending some message that leads to a ‘default’ decision \hat{d}_i (e.g., no trade). Since in G^d the principals play no role after offering their mechanisms, there is no communication strategy for them. The notion of equilibrium in Definition 1 naturally extends to G^d .

Social-Choice Functions: My aim here is to characterize the outcomes that can arise as equilibria of the interaction between the principals and the agent. I represent such outcomes as social-choice functions $f : \Theta \rightarrow \Delta(X)$.

3 Failure of the Inscrutability Principle

The so called Inscrutability Principle was first stated by Myerson in his 1983 seminal paper and is a fundamental result in the literature on mechanism design with *one* informed principal (see also Maskin and Tirole (1990, 1992)). Using Myerson’s own words, the Inscrutability Principle says: *“there is no loss of generality in assuming that all types of the principal should choose the same mechanism, so that his [...] choice [...] will convey no information.”* (p. 1774) Myerson considered a setting in which the informed principal can send messages to her mechanisms and therefore may not entirely delegate the final decisions to the agent. In this respect Myerson’s setting is similar to the non-delegation game G .

The intuition behind the Inscrutability Principle is as follows. For simplicity, suppose that each type of the principal offers a different mechanism, thereby revealing itself up front. After

m , the updated belief β^* must satisfy

$$\beta^*(\theta_{-0} | \theta_0, m) \left\{ \sum_{\theta'_{-0}} \mathbf{1}_{[\mu^*(\theta'_{-0})=m]} b(\theta'_{-0} | \theta_0) \right\} = \mathbf{1}_{[\mu^*(\theta_{-0})=m]} b(\theta_{-0} | \theta_0).$$

¹²See Peters (2001), Martimort and Stole (2002), Pavan and Calzolari (2010), Celik and Peters (2011), and Peters (2011).

observing a mechanism, the agent replies by sending some messages that, together with the principal’s messages, determine some final decision. The key observation here is that, conditional on the information revealed by the principal’s offer, the agent’s choices of messages depend only on his predetermined exogenous type. Therefore, if the principal knew the agent’s type, she could replicate—on her own and without having to reveal any information up front—the different decisions that her types were originally implementing with that particular type of the agent by offering different mechanisms. To induce the agent to reveal his type, the principal can commit to a single direct mechanism that works as follows: (1) the agent and the principal report one type each, and (2) for every pair of reports, the mechanism delivers the same decision that the reported types were implementing through the original mechanisms. It is easy to see that, under this new mechanism, truthful reporting is optimal for the agent if the principal behaves truthfully (and vice versa). So, the new uninformative mechanism allows the principal to implement the same decisions that she was implementing by revealing her information up front.¹³

The Inscrutability Principle is extremely useful because it allows one to completely disregard that the informed principal may use her choice of a mechanism as a signalling device. It is therefore important to know whether it also applies to games with multiple informed principals. If so, we could safely neglect the signalling role of the mechanisms and assume that each principal follows the same offer strategy independently of her type. The game G would then be essentially the same as a standard game with uninformed principals: As in a standard game, the mechanisms of each principal would respond to the agent’s messages with a potentially random allocation, except that now this randomness may depend on the communication strategy and type of the principal. The analysis of G would then be dramatically simpler.

Unfortunately, when the agent simultaneously interacts with multiple informed principals, the Inscrutability Principle fails. This happens because, conditional on the information revealed by principal i ’s offer, how the agent communicates with i may now depend not only on his exogenous type but also on his endogenous information about i ’s opponents, which includes their offers and how he communicates with them. Furthermore, the agent’s communication with principal i ’s opponents—and hence with i herself—may change depending on the information revealed by principal i ’s offers. Because of this interdependence, if in equilibrium principal i reveals some information up front, then the agent can induce a social-choice function that exhibits correlation between i ’s information and the allocations with *all* principals. However, the agent clearly cannot induce the same correlation if principal i reveals no information up front. The next example illustrates these points.

Example 1 (Failure of the Inscrutability Principle) *There are two principals. Principal 1 has two equally-likely types: $\Theta_1 = \{\bar{\theta}_1, \underline{\theta}_1\}$ with $b(\bar{\theta}_1) = 1/2$. Principal 2 and the agent are uninformed: $|\Theta_0| = |\Theta_2| = 1$. For each principal there are two possible allocations: $X_1 = \{\bar{x}_1, \underline{x}_1\}$ and $X_2 = \{\bar{x}_2, \underline{x}_2\}$. The sets of messages are $\mathcal{P}_1 = \{\underline{p}_1, \bar{p}_1\}$, $\mathcal{A}_1 = \{\underline{a}_1, \bar{a}_1\}$, $\mathcal{A}_2 = \{\underline{a}_2, \bar{a}_2\}$; for simplicity, I neglect \mathcal{P}_2 . Table 1 reports the players’ payoffs; in each cell, the first entry is principal 1’s payoff, the second is principal 2’s, and the third is the agent’s.*

¹³The Inscrutability Principle clearly relies on the principal’s ability to send her messages to her mechanism, so as to tailor the final decision to her type for any report of the agent. Therefore, the Principle does not apply to settings with only one principal who is informed but has to delegate the final decisions to the agent. A fortiori, this inapplicability extends to the delegation game G^d .

$\bar{\theta}_1$	\bar{x}_2	\underline{x}_2
\bar{x}_1	4, 3, 3	1, -1, 0
\underline{x}_1	1, 0, 0	-1, 2, 1

$\underline{\theta}_1$	\bar{x}_2	\underline{x}_2
\bar{x}_1	-1, 3, -1	1, -1, 0
\underline{x}_1	1, 0, 0	4, 2, 3

Table 1: Payoffs of Example 1

Note that the agent's preference over X_1 depends on x_2 . So, from the point of view of principal 1, the agent can have endogenous information to the extent that the final x_2 depends on the mechanism of principal 2 and on how the agent communicates with it. Moreover, the agent's preference over $X_1 \times X_2$ depends on θ_1 . Therefore, if the agent learns θ_1 by observing principal 1's offers, his endogenous information about x_2 can change.

I will construct an equilibrium that sustains the social-choice function $f^*(\bar{\theta}_1) = (\bar{x}_1, \bar{x}_2)$ and $f^*(\underline{\theta}_1) = (\underline{x}_1, \underline{x}_2)$. Note that f^* 'matches' both x_1 and x_2 with principal 1's type.

In this equilibrium, principal 2 offers a mechanism m_2^* such that $m_2^*(\bar{a}_2) = \bar{x}_2$ and $m_2^*(\underline{a}_2) = \underline{x}_2$. Principal 1 offers the mechanism \bar{m}_1 if her type is $\bar{\theta}_1$, and \underline{m}_1 if her type is $\underline{\theta}_1$, where

$$\bar{m}_1(a, p) = \begin{cases} \bar{x}_1 & \text{if } (a, p) = (\bar{a}_1, \bar{p}_1) \\ \underline{x}_1 & \text{if } (a, p) = (\underline{a}_1, \bar{p}_1) \\ \bar{x}_1 & \text{otherwise} \end{cases} \quad \text{and} \quad \underline{m}_1(a, p) = \begin{cases} \bar{x}_1 & \text{if } (a, p) = (\bar{a}_1, \underline{p}_1) \\ \underline{x}_1 & \text{if } (a, p) = (\underline{a}_1, \underline{p}_1) \\ \bar{x}_1 & \text{otherwise} \end{cases} .$$

Since \bar{m}_1 and \underline{m}_1 differ, principal 1 reveals her type up front. For $m_1 \in \{\bar{m}_1, \underline{m}_1\}$, her communication strategy is $\pi_1(\bar{\theta}_1, m_1) = \bar{p}_1$ and $\pi_1(\underline{\theta}_1, m_1) = \underline{p}_1$. Finally, when the agent is offered m_2^* and either \bar{m}_1 or \underline{m}_1 , his communication strategy is

$$\alpha(m_1, m_2^*) = \begin{cases} (\bar{a}_1, \bar{a}_2) & \text{if } m_1 = \bar{m}_1 \\ (\underline{a}_1, \underline{a}_2) & \text{if } m_1 = \underline{m}_1 \end{cases} . \quad (1)$$

Note that according to (1), the agent selects different allocations by communicating to m_2^* depending on whether principal 1 offers \bar{m}_1 or \underline{m}_1 . For other mechanisms m_2 and $m_1 \notin \{\bar{m}_1, \underline{m}_1\}$, one can extend α , π_1 , and β so that they constitute a continuation equilibrium. In particular, fixing (μ_1, π_1) , extend α as follows: if m_2 allows the agent to choose only \bar{x}_2 , then he sends the message \bar{a}_1 to \bar{m}_1 and \underline{a}_1 to \underline{m}_1 ; if instead m_2 allows the agent to choose only \underline{x}_2 , then he sends message \underline{a}_1 both to \bar{m}_1 and to \underline{m}_1 . It is easy to see that the strategy (μ_1, π_1) admits no profitable deviation given α and μ_2 —principal 1 is getting her maximal payoff of 4. Principal 2 is also maximizing her payoff given (μ_1, π_1) and α . Finally, given the updated belief $\beta(\bar{\theta}_1 | \bar{m}_1) = \beta(\underline{\theta}_1 | \underline{m}_1) = 1$, α also admits no profitable deviation.

There is, however, no equilibrium of G and no mechanism m_1 with the property that both types of principal 1 offer m_1 and the equilibrium social-choice function is f^* . If both types of principal 1 offer the same mechanism, the agent does not learn her type. Therefore, even if principal 2 allows the agent to choose between \bar{x}_2 and \underline{x}_2 , he cannot match the two principals' allocations as required by f^* : With positive probability, he will choose either \underline{x}_2 when $\theta_1 = \bar{\theta}_1$, or \bar{x}_2 when $\theta_1 = \underline{\theta}_1$.

We have therefore reached the following important conclusion. In contrast to non-delegation games with one informed principal, when studying the outcomes of non-delegation games with multiple informed principals one can no longer disregard that, in equilibrium, the principals may use their choices of mechanisms to signal their types.

4 Failure of the Extended Taxation Principle

In this section, I turn my attention to the literature on common agency with uninformed principals. An important result in this literature is the so called Extended Taxation Principle, or Menu Theorem (Peters (2001); Martimort and Stole (2002)). Loosely speaking, this Principle says the following: Given an arbitrary game with uninformed principals G^{un} , there exists a simpler game G^L in which principals are restricted to offering only menus of decisions, and such that any equilibrium social-choice function that arises in G^{un} also arises in G^L and vice versa. As I will momentarily show, when principals are informed, the Extended Taxation Principle does not apply; this happens for two reasons. First, an informed principal may use her indirect mechanisms not only to constrain the agent to choose from a specific menu of decisions, but also to signal her private information by framing such a menu in different ways. Second, in non-delegation games, an informed principal can also use her mechanisms to retain some ability to influence the final decisions through her messages, while revealing up front as little information as she wishes. Such a flexibility is, however, impossible if the principal can only offer menus of decisions.

To prove my claims, I first need to introduce the notion of menu and recall the statement of the Extended Taxation Principle. A menu is simply a ‘list’ of decisions that the agent can choose by naming one. Formally, we can define menus as follows.¹⁴

Definition 3 (Simple Menus) *For $i = 1, \dots, N$, a **menu** l_i is a mechanism such that (1) the message space of the agent, \mathcal{A}_{l_i} , is a nonempty subset of \mathcal{D}_i and (2) $l_i(d_i) = d_i$ for every $d_i \in \mathcal{A}_{l_i}$.*

Let the set of available menus for principal i be L_i , and let G^L be the ‘menu game’ in which each principal i is restricted to offering menus in L_i .

To be able to replicate all social-choice functions sustained by equilibria of a more general game G^{un} , in which uninformed principals compete in arbitrary mechanisms, the game G^L must be sufficiently comparable to G^{un} in terms of the principals’ ability to affect their decisions with the agent. Intuitively, if in G^{un} principal i can (not) offer a mechanism that allows the agent to induce a decision d_i , then in G^L principal i should (not) be able to offer a menu that contains d_i . This intuitive requirement is usually captured with the notion of enlargement (see also Pavan and Calzolari (2010)).

Definition 4 (Enlargement) *The game G^{un} is an enlargement of the game G^L if, for every $i = 1, \dots, N$, we have that for every menu $l_i \in L_i$ there exists a mechanism $m_i \in M_i$ such that $\text{Ra}(m_i) = \text{Ra}(l_i)$, and vice versa.*

The Extended Taxation Principle then says the following (Peters (2001); Martimort and Stole (2002)).

Extended Taxation Principle (Menu Theorem): Let G^{un} be an enlargement of G^L . There exists an equilibrium of G^{un} that sustains the social-choice function f if and only if there exists an equilibrium of G^L that sustains f .

¹⁴See Peters (2001), Martimort and Stole (2002), and Pavan and Calzolari (2010).

In the rest of this section, I provide two counterexamples to this result for games with informed principals—that is, if one replaces G^{un} with G or G^d . These examples also allow me to develop the intuition behind my main results in Sections 5.1 and 5.2.

The first example shows that the Extended Taxation Principle does not apply to delegation games G^d . The reason is as follows. In games with uninformed principals the Extended Taxation Principle holds because, by offering an arbitrary indirect mechanism, each principal is effectively restricting the agent to choosing from a specific menu of decisions. So, if in the menu game G^L each principal i can rely on an appropriately rich set of menus L_i , then the equilibrium outcomes of G^L should—and indeed are—the same as in the original game. In contrast, in a delegation game G^d , by choosing her mechanisms each informed principal not only defines a menu of available choices for the agent, but she may also signal some of her information by framing such a menu in a specific way. So in G^L , even if each principal can offer the same menus as she does in G^d through arbitrary mechanisms (as required by the notion of enlargement), she cannot frame them in different ways and therefore she loses part of the ability to signal her types. As a result, some of the equilibrium social-choice functions that arise in G^d may be unattainable in G^L .

Example 2 (Failure of the Extended Taxation Principle for G^d) *As in example 1, there are two principals, $\Theta_1 = \{\bar{\theta}_1, \underline{\theta}_1\}$ with $b(\bar{\theta}_1) = 1/2$, and $|\Theta_0| = |\Theta_2| = 1$. Their feasible allocations are $X_1 = \{\bar{x}_1, \underline{x}_1\}$ and $X_2 = \{\bar{x}_2, x'_2, \underline{x}_2\}$. For simplicity, both principals can only offer mechanisms of the form $m_i : \mathcal{A}_i \rightarrow X_i$ —i.e., m_i can induce only degenerate lotteries over X_i —where $\mathcal{A}_1 = \{\bar{a}_1, \underline{a}_1\}$ and $\mathcal{A}_2 = \{\bar{a}_2, a'_2, \underline{a}_2\}$. Table 4 reports the players' payoffs.*

$\bar{\theta}_1$	\bar{x}_2	x'_2	\underline{x}_2
\bar{x}_1	1, 1, 1	1, 3, 2	1, 1, 0
\underline{x}_1	0, 1, 0	0, 0, 3	0, 1, 0

$\underline{\theta}_1$	\bar{x}_2	x'_2	\underline{x}_2
\bar{x}_1	0, 1, 0	0, 0, 3	0, 1, 0
\underline{x}_1	1, 1, 0	1, 3, 2	1, 1, 1

Table 4: Payoffs of Example 2

I construct an equilibrium of G^d that sustains the social-choice function $f^*(\bar{\theta}_1) = (\bar{x}_1, \bar{x}_2)$ and $f^*(\underline{\theta}_1) = (\underline{x}_1, \underline{x}_2)$. In this equilibrium, principal 2 offers a mechanism m_2^* such that $\text{Ra}(m_2^*) = \{\bar{x}_2, \underline{x}_2\}$. Principal 1 offers the mechanism \bar{m}_1 if her type is $\bar{\theta}_1$, and \underline{m}_1 otherwise, where

$$\bar{m}_1(a) = \begin{cases} \bar{x}_1 & \text{if } a = \bar{a}_1 \\ \underline{x}_1 & \text{if } a = \underline{a}_1 \end{cases} \quad \text{and} \quad \underline{m}_1(a) = \begin{cases} \underline{x}_1 & \text{if } a = \bar{a}_1 \\ \bar{x}_1 & \text{if } a = \underline{a}_1 \end{cases}.$$

Then, on path the agent updates his beliefs to $\beta(\bar{\theta}_1 \mid \bar{m}_1, m_2^*) = \beta(\underline{\theta}_1 \mid \underline{m}_1, m_2^*) = 1$. Also, assume that these beliefs do not change after unilateral deviations of principal 2. The agent communicates using the following strategy α : given (\bar{m}_1, m_2^*) , he chooses \bar{x}_2 with principal 2 and sends the message \bar{a}_1 to \bar{m}_1 ; given (\underline{m}_1, m_2^*) , he chooses \underline{x}_2 with principal 2 and sends the message \bar{a}_1 to \underline{m}_1 . Moreover, whenever the mechanism of principal 2 lets the agent choose x'_2 , he does so and sends the message \underline{a}_1 both to \bar{m}_1 and to \underline{m}_1 . Again, one can appropriately define the agent's beliefs and behavior after deviations by principal 1 without affecting the message of the example. It is immediate to see that the strategy α admits no profitable deviation given β , μ_1 , and μ_2 . Also, μ_1 is a best reply to μ_2 and α . Finally, principal 2 has no profitable deviation: If she allows the agent to choose x'_2 , her payoff falls to zero because he then chooses the 'wrong' x_1 .

There is, however, no menu game G^L such that G^d is an enlargement of G^L , and f^* is an equilibrium social-choice function of G^L . On the one hand, if both types of principal 1 offer the same menu with \bar{x}_1 and \underline{x}_1 , then the agent will induce the same lottery over X_1 , independently of θ_1 . On the other hand, if either type of principal 1 offers the menu containing only her preferred x_1 , then principal 2 cannot be deterred from offering x'_2 .

Example 2 shows that in the menu game G^L principal 1 faces a trade-off, which is absent in G^d , between defining the menus of available decisions for the agent and signalling her type. To signal her type using only menus, principal 1 must exclude from one of them a ‘latent contract’ that, in the equilibrium of G^d , allows her to deter a deviation of principal 2 by relying on the self-interest of the agent.

The Extended Taxation Principle also fails in non-delegation games G . One reason of this failure is the same as in the case of delegation games: In the menu game G^L , principal i may not be able to restrict the available decisions for the agent to a specific menu, and at the same time signal as much information as she does in G . But in the case of non-delegation games, the issue may also be the opposite. Specifically, in G each principal can retain some power to influence the final decisions, through her messages, at the moment when the agent communicates with her and her opponents. Instead in G^L , to influence the agent’s decisions, each principal can only restrict his options to different menus; moreover, she must do so before the agent communicates with her opponents. So, in G^L a principal may face a trade-off between restricting the agent’s options to the decisions that her type prefers and revealing too much—rather than too little—information, which the agent could then use against her with her opponents. In contrast, in G each principal can reveal as little information as she wishes, while still influencing the final decisions. The next example illustrates these points.

Example 3 (Failure of the Extended Taxation Principle for G) *As in example 1, there are two principals, $\Theta_1 = \{\bar{\theta}_1, \underline{\theta}_1\}$ with $b(\bar{\theta}_1) = 1/2$, and $|\Theta_0| = |\Theta_2| = 1$. Their feasible allocations are $X_1 = \{\bar{x}_1, \underline{x}_1\}$ and $X_2 = \{x_2, x'_2, x''_2\}$; for simplicity, $\mathcal{D}_i = X_i$. Consider the game G where $M_2 = L_2$ and $M_1 = \{m_1 : \mathcal{A}_1 \times \mathcal{P}_1 \rightarrow X_1\}$ with $\mathcal{A}_1 = \{\bar{a}_1, a'_1, \underline{a}_1\}$ and $\mathcal{P}_1 = \{\bar{p}_1, \underline{p}_1\}$. Suppose that the set L_i allows principal i to ‘offer’ all nonempty subsets of X_i ; then G is an enlargement of G^L . Regarding players’ payoffs, consider the two possibilities in Table 2 and 3. In Table 2, both principal 2’s and the agent’s payoffs are independent of θ_1 (private values). Instead, in Table 3, the agent’s payoffs depend on θ_1 (common values).*

$\bar{\theta}_1$	x_2	x'_2	x''_2
\bar{x}_1	-1, 2, 3	7, 1, 2	-2, -2, 0
\underline{x}_1	1, -2, 0	4, 1, 2	3, 2, 3

$\underline{\theta}_1$	x_2	x'_2	x''_2
\bar{x}_1	2, 2, 3	3, 1, 2	1, -2, 0
\underline{x}_1	0, -2, 0	6, 1, 2	-1, 2, 3

Table 2: Private Values

$\bar{\theta}_1$	x_2	x'_2	x''_2
\bar{x}_1	-1, 2, 3	7, 1, 2	-2, -2, 0
\underline{x}_1	1, -2, 3	4, 1, 2	3, 2, 0

$\underline{\theta}_1$	x_2	x'_2	x''_2
\bar{x}_1	2, 2, 0	3, 1, 2	1, -2, 3
\underline{x}_1	0, -2, 0	6, 1, 2	-1, 2, 3

Table 3: Common Values

For both payoff specifications, I construct an equilibrium of G that sustains the social-choice function $f^*(\bar{\theta}_1) = (\bar{x}_1, x'_2)$ and $f^*(\underline{\theta}_1) = (\underline{x}_1, x'_2)$. In this equilibrium principal 2 offers the menu l_2^* with message space $\mathcal{A}_{l_2^*} = X_2$. Independently of her type, principal 1 offers the mechanism m_1^* where

$$m_1^*(a_1, \bar{p}_1) = \begin{cases} \underline{x}_1 & \text{if } a_1 = \bar{a}_1 \\ \bar{x}_1 & \text{if } a_1 = a'_1 \\ \underline{x}_1 & \text{if } a_1 = \underline{a}_1 \end{cases} \quad \text{and} \quad m_1^*(a_1, \underline{p}_1) = \begin{cases} \bar{x}_1 & \text{if } a_1 = \bar{a}_1 \\ \underline{x}_1 & \text{if } a_1 = a'_1 \\ \bar{x}_1 & \text{if } a_1 = \underline{a}_1 \end{cases} .$$

Moreover, $\pi_1(\bar{\theta}_1, m_1^*) = \bar{p}_1$ and $\pi_1(\underline{\theta}_1, m_1^*) = \underline{p}_1$. In words, by offering m_1^* , principal 1 allows the agent to tell her whether he chooses x'_2 or not, and then she chooses her best x_1 accordingly. Finally, the agent's strategy α is such that he sends the message a'_1 whenever he can choose x'_2 with principal 2; otherwise, he sends a different message to principal 1 and picks the best available option with principal 2. Formally,

$$\alpha(m_1^*, l_2) = \begin{cases} (a'_1, x'_2) & \text{if } x'_2 \in \mathcal{A}_{l_2} \\ (\bar{a}_1, x_2) & \text{if } \mathcal{A}_{l_2} = \{x_2, x'_2\} \\ (\bar{a}_1, x_2) & \text{if } \mathcal{A}_{l_2} = \{x_2\} \\ (\underline{a}_1, x'_2) & \text{if } \mathcal{A}_{l_2} = \{x'_2\} \end{cases} . \quad (2)$$

As in example 1, when $m_1 \neq m_1^*$ or $l_2 \neq l_2^*$, one can complete α , π_1 , and β to obtain a continuation equilibrium.

I claim that the strategy profile that I have just described—together with the belief $\beta(\theta_1 | m_1, l_2) = b(\theta_1)$ for all m_1, l_2 —is an equilibrium of G . First, both types of principal 1 get their maximal payoffs by playing π_1 against (α, μ_2) . Furthermore, α admits no profitable deviation because choosing x'_2 is the unique best reply of the agent to (μ_1, π_1) , and if x'_2 is not available, it is optimal for him to select x_2 (or x'_2) and then send any message to m_1^* . Therefore, (μ_1, π_1) is optimal for principal 1. To see that principal 2 has no profitable deviation, note that any menu l_2 that 'includes' x'_2 is equivalent to l_2^* ; on the other hand, the agent's strategy in (2) together with (μ_1, π_1) prevents principal 2 from gaining if she deviates to a menu l_2 that does not include x'_2 (i.e., such that $x'_2 \notin \mathcal{A}_{l_2}$).

However, no equilibrium of the menu game G^L can sustain f^* —this is true for both payoff specifications, but for different reasons as I explain shortly.¹⁵ Suppose in negation that such an equilibrium exists. Since f^* specifies a different x_1 depending on the type of principal 1, in the equilibrium of G^L $\bar{\theta}_1$ and $\underline{\theta}_1$ cannot offer the same menu; otherwise, the agent would (randomly) induce the same allocation with both types. It follows that, to sustain f^* , at least one of principal 1's types must offer only the allocation that she likes best; so, principal 1 must reveal her type up front. Finally, given his degenerate beliefs, the agent selects x'_2 with principal 2 if and only if she offers the menu l_2 that contains only x'_2 (i.e., with $\mathcal{A}_{l_2} = \{x'_2\}$). This offer strategy, however, is not optimal for principal 2: if she offers the menu l_2 that contains all options (i.e., with $\mathcal{A}_{l_2} = \{x_2, x'_2, x''_2\}$), then she is better off because the agent chooses either x_2 or x''_2 .¹⁶

¹⁵See also the discussion in Maskin and Tirole (1990-1992) about the differences between private- and common-value assumptions.

¹⁶In the common-value case, the agent would still choose x'_2 with principal 2, if her deviation to l_2 caused the agent not to update at all his beliefs about θ_1 even though principal 1 has not deviated from her strategy. Although this is technically possible, it seems implausible (see Fudenberg and Tirole (1991)).

Example 3 helps us better understand how G^L restricts, relative to G , the principals' ability to influence the final outcomes of the game. In G^L , to achieve the decision that best fits her type, principal 1 must offer a menu that includes only that decision, thereby revealing her type. But if she does so, she makes the agent take undesirable decisions with principal 2. The reason in the common-value case is simple: By revealing her type up front, principal 1 directly affects the agent's preference over her opponent's allocations. In the example, after discovering the true θ_1 , the agent knows whether x_2 or x_2'' is the best allocation that he can get with principal 2. The private-value case, instead, is more subtle. Even if principal 1 reveals her type up front, she has no direct effect on the agent's preference. Nonetheless, recall that the agent's preference over X_2 *endogenously* depends on x_1 . Therefore, if principal 1 reduces the agent's uncertainty about the final x_1 by committing to a menu, she can still induce an undesirable x_2 . In the example, when the agent is uncertain about the final x_1 that principal 1 will choose using her messages, he prefers the safe allocation x_2' with principal 2. But when the agent is not uncertain about x_1 because he himself chooses it from a menu, then if he can, he will choose either x_2 or x_2'' .

To summarize, the examples in this and the previous section have unveiled the main strategic components of the games with informed principals. Specifically, in delegation games, the principals use their mechanisms for *two* key purposes: (1) to specify menus of available decisions for the agent; (2) to signal their information. In games with multiple uninformed principals *only* the first purpose matters—as shown by the Extended Taxation Principle. In contrast, in games with multiple informed principals the second purpose becomes important, which is perhaps the key novelty of the present model. Moreover, in non-delegation games each principal also uses her mechanisms for a third purpose: to partake in choosing an allocation—acting on her private information—*after* the agent has communicated with her opponents.

5 A Simple Characterization of Equilibrium Outcomes

In this section, I build on the insights from Sections 3 and 4 to construct new games that are simpler than the general games described in Section 2, but are capable of replicating all the equilibrium social-choice functions that arise in those general games. As noted, the new games should allow each principal to offer the agent appropriately defined choice sets as well as to signal her types, without introducing new trade-offs between these two goals.

As far as signalling is concerned, it takes a relatively simple form in equilibria in which principals play deterministic offer strategies. This important feature is shared by delegation and non-delegation games; it is therefore worth discussing it now. When principal i offers mechanisms as a deterministic function of her types, with each offer she simply signals that her type is in a specific subset of Θ_i . In other words, each deterministic offer strategy of principal i defines a partition of Θ_i , with each element of the partition containing all types that offer the same mechanism. In contrast, if principal i played a mixed offer strategy, she would give rise to a more complicated pattern of signals. Since each of her types could offer many mechanisms with different probabilities, we could no longer describe the resulting signal structure only in terms of Θ_i .

Therefore, to replicate all social-choice functions sustained by equilibria of G^d and G , the new games will have to allow each principal to achieve the form of signalling that I have just described, while offering the agent different choice sets. I begin by considering the delegation

games.

5.1 Delegation Games

For any delegation game G^d , I consider a new game that is equivalent to the standard menu game of Peters (2001), except that it also endows each principal with a set of ‘cheap-talk’ signals. I show that this new game can replicate all social-choice functions sustained by the equilibria of G^d . This result implies that we can study the equilibrium outcomes of any game G^d using a very simple extension of Peters’ menu game.

I refer to my new game as the ‘signal-menu’ game, which I denote by G^{SL} . In this game, each principal i does two things: (1) she offers menus from the set L_i (Definition 3), and (2) provided that her menu contains more than just the default decision \hat{d}_i —which means that she is actively participating in the game—she sends a signal from the set $S_i = \Theta_i$.¹⁷ After observing the menus and the signals of all principals, the agent chooses a decision from each menu. Therefore, the game G^{SL} works very similarly to Peters’ menu game except that, unless a principal is not actively participating in the game, she can also give the agent some information about her type using her signals. To capture this formally, for each principal i , let \hat{l}_i be the menu that contains only \hat{d}_i (i.e., $\mathcal{A}_{\hat{l}_i} = \{\hat{d}_i\}$). Then, in G^{SL} , the action space of principal i is $\{S_i \times L_i \setminus \hat{l}_i\} \cup \{\hat{l}_i\}$, and a signal-offer strategy μ_i specifies, for each θ_i , a lottery over i ’s action space. The agent’s strategy is a map that, for every profile of signals $\hat{\theta}$ and menus l , specifies a lottery over the profiles of messages \mathcal{A}_l . As for the original game G^d , the term equilibrium of G^{SL} refers to a weak PBE in which, after observing her type, every principal i chooses deterministically a menu and a signal—that is, if for every $i = 1, \dots, N$ and $\theta_i \in \Theta_i$, $\mu_i(\theta_i)$ is a degenerate lottery over i ’s action space.

To replicate all equilibrium social-choice functions that arise in G^d , the new game G^{SL} must satisfy two requirements: (1) in G^{SL} , the principals must be as able to affect the agent’s decisions by specifying his available choices as they are in G^d ; (2) in G^{SL} , the principals must also be as able to affect the agent’s decisions by signalling their information as they are in G^d . To satisfy requirement (1), which also applies to games with uninformed principals, I follow the literature using the notion of enlargement: I will say that G^d is an enlargement of G^{SL} if and only if they satisfy the conditions of Definition 4, where G^d replaces G^{run} and G^{SL} replaces G^L . Note that, to find the G^{SL} that satisfies this definition for a given G^d , we can treat G^d as a game with uninformed principals, derive the standard menu game that satisfies Definition 4, and then simply endow each principal with a set of signals equal to her set of types as suggested above. This last step should also ensure that requirement (2) is satisfied, given the simple form of signalling that principals can achieve in an equilibrium of G^d . Indeed, Theorem 1 below confirms this intuition.

Theorem 1 *Let G^d be an enlargement of G^{SL} . There exists an equilibrium of G^d that sustains the social-choice function f if and only if there exists an equilibrium of G^{SL} that sustains f .*

To understand the intuition behind Theorem 1, consider first an equilibrium of G^d . As noted, in such an equilibrium, by playing the mechanism m_i principal i signals that her type is in some subset $T_i(m_i)$ of Θ_i . So, in G^{SL} we can let all types in $T_i(m_i)$ offer the menu l_i containing the same

¹⁷Any other set S_i with $|S_i| \geq |\Theta_i|$ would work.

decisions that the agent can induce through m_i , and send him a common signal $\hat{\theta}_i \in T_i(m_i)$ —of course, unless $l_i = \hat{l}_i$.¹⁸ By replicating this procedure for each type of each principal, we construct a profile of offer strategies in G^{SL} . Note that according to this construction, after each on-path profile of menus and signals, the agent has the same information and available choices as he has after the corresponding on-path profile of mechanisms in the equilibrium of G^d . Therefore, it must be optimal for him to induce the same decisions as he does in G^d . Similarly, it must be optimal for each type of each principal to offer the menu and send the signal as specified. This intuitively explains why every social-choice function sustained by an equilibrium of G^d is also a social-choice function sustained by an equilibrium of G^{SL} .

Showing the converse is more delicate. By the definition of enlargement, each menu of decisions that principal i can offer in G^{SL} can also be offered in G^d through some indirect mechanism. However, this is not enough: It is also important that each principal can achieve, in G^d , the same degree of signaling that she achieves in the equilibrium of G^{SL} . On the one hand, in G^{SL} menus and signals are separated and principal i can combine them in any way she wants (provided that she is actively participating in the game). On the other hand, in G^d her signals are intertwined with the different ways in which she lets the agent’s messages map to a choice set. Nonetheless, the agent’s message spaces in G^d are rich enough to ensure that each principal can always design as many different indirect mechanisms as she needs to achieve the desired degree of signaling, while offering the same bundle of decisions. Because of this property, we can essentially reverse the construction that I sketched before, now going from an equilibrium of G^{SL} to an equilibrium of G^d that sustains the same social-choice function.

5.2 Non-Delegation Games

In this section, I consider the non-delegation game G . I will construct a new game that is simpler than G , but replicates all social-choice functions sustained by the equilibria of G .

My new game has two periods as does G and proceeds as follows. In the first period, each principal offers the agent a list of restricted direct mechanisms (RDMs), simultaneously with her opponents. Intuitively, an RDM of principal i is a standard direct mechanism that depends only on i ’s reports about her types, except that it restricts such reports to a subset of all her possible types. A list of RDMs of principal i is simply a collection of such RDMs, all restricting her reports to the same subset of types. In the second period of the game, after privately observing the offers of all principals, the agent chooses an RDM from each list. Finally, each principal submits her report before observing the RDM chosen by the agent.

I will now formally describe my new game, starting from the lists of RDMs. For every principal i and nonempty subset T_i of Θ_i , denote the set of RDMs restricted to T_i by

$$R^{T_i} = \{r^{T_i} : T_i \rightarrow \Delta(X_i)\}.$$

Intuitively, a list of RDMs restricted to T_i corresponds to a subset of R^{T_i} from which the agent chooses an option by simply naming it.

¹⁸Note that if in the equilibrium of G^d two (or more) types of principal i choose not to actively participate in the game—i.e., they both offer the unique mechanism m_i with $\text{Ra}(m_i) = \hat{d}_i$ —they have no possibility of differentiating themselves in the eyes of the agent. Therefore, they should also not have this possibility in G^{SL} after offering \hat{l}_i .

Definition 5 (List of RDMs) Fix principal i and a nonempty $T_i \subset \Theta_i$. A **list** of direct mechanisms restricted to T_i —denoted by λ^{T_i} —is a function $\lambda^{T_i} : \mathcal{A}_{\lambda^{T_i}} \rightarrow R^{T_i}$ such that (1) the message space of the agent, $\mathcal{A}_{\lambda^{T_i}}$, is a nonempty subset of R^{T_i} , and (2) λ^{T_i} is the identity function, i.e., $\lambda^{T_i}(r^{T_i}) = r^{T_i}$. Finally, the set of all possible lists of RDMs of principal i is Λ_i .

Based on this definition, I denote my new game by G^Λ . In this game, the strategy of principal i involves an offer strategy μ_i and a communication strategy π_i , where μ_i assigns to each θ_i a lottery over Λ_i and π_i assigns to each $(\theta_i, \lambda^{T_i})$ a lottery over the restricted set of reports T_i . The communication strategy α of the agent specifies, for each profile $\lambda^T = (\lambda^{T_1}, \dots, \lambda^{T_N})$, a lottery over \mathcal{A}_{λ^T} . As before, the term equilibrium of G^Λ refers to a weak PBE in which, for every principal i and type θ_i , $\mu_i(\theta_i)$ is a degenerate lottery over Λ_i .

To be able to replicate all social-choice functions that arise in the non-delegation game G , the new game G^Λ and the original one G must be sufficiently comparable. To denote when G^Λ and G are comparable, I will use an appropriate definition of enlargement as I did for delegation games in Section 5.1. To state this definition, I first need to show how we can combine any indirect mechanism m_i and communication strategy π_i in G to obtain a list of RDMs. For each message a_i of the agent, we can use $m_i(a_i, \cdot)$ and π_i to construct a direct mechanism r^{Θ_i} by defining the probability that r^{Θ_i} assigns to the allocation x_i given the report θ_i by

$$r^{\Theta_i}(x_i; \theta_i) = \sum_{\mathcal{P}_i} d_i(x_i) m_i(d_i; a_i, p_i) \pi_i(p_i; \theta_i, m_i). \quad (3)$$

If we now repeat the same construction for each message $a_i \in \mathcal{A}_i$, we identify a subset of R^{Θ_i} that we can use to define a list of RDMs λ^{Θ_i} following Definition 5. To denote the list of RDMs that results from combining m_i and π_i , I will use the notation $c(m_i, \pi_i)$. Finally, I will use the notation $\lambda^{T_i} \triangleleft \lambda^{\Theta_i}$ to denote that the list λ^{T_i} results from restricting every direct mechanism in the list λ^{Θ_i} to the subset of reports T_i .

Definition 6 (Enlargement of G^Λ) A game G is an enlargement of G^Λ if and only if for every principal i three conditions hold: (1) for every mechanism $m_i \in M_i$ and communication strategy π_i in G , there is a list $\lambda^{\Theta_i} \in \Lambda_i$ such that $c(m_i, \pi_i) = \lambda^{\Theta_i}$; (2) for every $\lambda^{\Theta_i} \in \Lambda_i$, there is an $m_i \in M_i$ and a π_i in G such that $\lambda^{\Theta_i} = c(m_i, \pi_i)$; (3) $\lambda^{T_i} \in \Lambda_i$ if and only if $\lambda^{T_i} \triangleleft \lambda^{\Theta_i}$ for some $\lambda^{\Theta_i} \in \Lambda_i$.

As noted, the combination of any indirect mechanism m_i and communication strategy π_i in G results in a list of RDMs λ^{Θ_i} . So, condition (1) requires that such a list be feasible for principal i in G^Λ . Conversely, condition (2) requires that a list λ^{Θ_i} be infeasible for principal i in G^Λ if in G she has no mechanism m_i and communication strategy π_i whose combination results in λ^{Θ_i} . Finally, condition (3) requires that the whole set Λ_i could be obtained by considering every feasible list of unrestricted direct mechanisms λ^{Θ_i} and by restricting all its mechanism to T_i , for every nonempty subset T_i of Θ_i .

Since in G^Λ the principals will play RDMs, it would be desirable to restrict attention to equilibria in which—at least on path—they report truthfully their type to their RDMs. The next definition and lemma show that one can do so without loss of generality.

Definition 7 (Truthful Equilibrium) An equilibrium of G^Λ is a (principal) **truthful equilibrium** if, for every principal i and $\theta_i \in \Theta_i$, we have $\pi_i(\theta_i, \mu_i(\theta_i)) = \theta_i$.

Lemma 1 *If there exists an equilibrium of G^Λ that sustains the social-choice function f , then there exists a truthful equilibrium that sustains f .*

Truthful equilibria represent a simple and intuitive class of equilibria of G^Λ . Moreover, the set of social-choice functions sustained by truthful equilibria of G^Λ coincides with the set of social-choice functions sustained by the equilibria of G .

Theorem 2 *Let G be an enlargement of G^Λ . There exists an equilibrium of G that sustains the social-choice function f if and only if there exists a truthful equilibrium of G^Λ that sustains f .*

Intuitively, the new game G^Λ can replicate all social-choice functions sustained by equilibria of G , because in G^Λ each principal maintains the same ability to influence the final decisions with the agent that she has in G . To see this, recall that in G each principal uses her indirect mechanisms for two purposes: to specify a set of decisions from which she and the agent will choose one using their messages, and to signal her information.

Consider the first purpose. On the one hand, if principal i retains some decision power through her messages, it means that she only lets the agent use his messages to select a subset of decisions from which she effectively chooses the final one. Furthermore, when principal i sends her messages, she only knows her type and the mechanism that she has offered to the agent. Therefore, to retain all her decision power, principal i only needs to assign to any message of the agent a subset that only contains one decision for each of her possible types. On the other hand, principal i may want to let the agent select different subsets of decisions depending on his exogenous type and endogenous information about her opponents (recall the discussion in Section 3 about the failure of the Inscrutability Principle). To do so, principal i could offer mechanisms that assign a subset of decisions to each report of the agent about his exogenous and endogenous information. Such a report, however, would involve a description of the mechanisms offered by i 's opponents and of how these mechanisms respond to reports about i 's mechanisms, causing an infinite regress. To avoid this problem, principal i could simply offer the agent a list of subsets of decisions and let him directly choose among these subsets. It is easy to see that any list of subsets, each containing one decision for each type of the principal, is equivalent to a list of direct mechanisms, each letting principal i choose a final decision by reporting her type.

Consider now the signalling purpose that indirect mechanisms have in G . As noted, in an equilibrium of G , by choosing a mechanism principal i at most signals that her true type is in a specific subset, say T_i , of her set of types Θ_i . So, if principal i could offer a list of direct mechanisms that commit her to reporting only the types in T_i , she could signal to the agent that her true type is in T_i . This is the basic idea behind letting principal i offer lists of RDMs, so that she can use the domain of these RDMs as her signalling device.

These remarks and Theorem 2 also help us understand why it is without loss of generality, with regard to the equilibrium social-choice functions, to assume that in G the principals communicate with their mechanisms simultaneously with the agent.¹⁹ To see this, consider a game G' that is identical to G , except that in G' each principal chooses her messages *after* observing the message that the agent has sent to her mechanism. From this observation principal i may learn something about the agent's type and her opponents' types and mechanisms. Conditional

¹⁹I limit myself to an informal argument because the cost of introducing new notation and definitions outweighs their benefit here.

on this information, however, i 's choices of messages depend only on her type. Thus, for the same reason that applies to G , principal i only needs lists of RDMs to be able to influence the final decisions with the agent as she does in G' . Furthermore, principal i can always offer lists of RDMs such that, conditional on what she learns from the agent's choice of an RDM, she wants to truthfully report her type to it. Therefore, we lose nothing if each principal has to report her types before observing the agent's choice of an RDM. Also, note that the agent only cares about the distribution over allocations that his choices of RDMs can induce. For these reasons, if an equilibrium of G' sustains the social-choice function f , then there is a truthful equilibrium of G^A that sustains f . Theorem 2 then implies that every social-choice function that arises in G' also arises in G .

There is a similarity between the logic behind the last remark and that behind the Inscrutability Principle. Roughly speaking, the Inscrutability Principle says that the single informed principal need not convey any information to the agent, because her information has no impact on the agent's exogenous information, which is what ultimately matters for the final decisions. Therefore, the Principle justifies restricting attention to simpler games in which the informed principal must offer the same mechanism independently of her type. Similarly, the previous remark says that the endogenously informed *agent* need not convey any information to the principals, because his information has no impact on the principals' exogenous information, which again ultimately determines their final decisions. Again, this justifies restricting attention to simpler games in which each informed principal communicates with her mechanism simultaneously with the agent.

6 Conclusions

In this paper, I have considered common-agency games in which principals have private information. I have shown that the set of equilibrium outcomes of these games depends on the principals' ability to signal their information through their choices of mechanisms. This property explains why both the Inscrutability and the Extended Taxation Principle fail in these games. Finally, I have shown how to characterize the equilibrium outcomes of general delegation and non-delegation games, using new games that are simpler than the original ones but ensure that each principal retains the same ability to influence the final decisions of the agent.

Common-agency games with informed principals provide a useful framework to analyze many situations of economic interest, including competition in menu auctions, in oligopolistic markets, or in lobbying contests. Applying my results to study how the outcomes in these situations depend on the principals' private information—e.g., in terms of efficiency or ability to aggregate dispersed information—represents an interesting avenue for future research.

7 Appendix

7.1 Proof of Theorem 1

Part 1: (\Rightarrow) Let (σ^*, β^*) be an equilibrium of G^d that sustains the social-choice function f_{σ^*} . I shall construct an equilibrium $(\sigma^{**}, \beta^{**})$ of G^{SL} that sustains $f_{\sigma^{**}} = f_{\sigma^*}$.

Given the offer strategy μ_i^* of principal i , let $\mu_i^*(\Theta_i)$ be the set of mechanisms that are offered by some θ_i in equilibrium. For every θ_i , it is possible to identify the subset $T_i(\mu_i^*(\theta_i)) = (\mu_i^*)^{-1}(\mu_i^*(\theta_i)) \subset \Theta_i$. Since μ_i^* is deterministic, the collection $\{T_i(\mu_i^*(\theta_i))\}_{\theta_i \in \Theta_i}$ is a partition of Θ_i with at most $|\Theta_i|$ elements. Define the function $\vartheta_i : \mu_i^*(\Theta_i) \rightarrow \Theta_i$ such that, for every $T_i(\mu_i^*(\theta_i))$, $\vartheta_i(\mu_i^*(\theta_i))$ is the signal $\hat{\theta}_i \in T_i(\mu_i^*(\theta_i))$ with the lowest index—this choice is just for convenience. Since G^d is an enlargement of G^{SL} , for every $m_i \in M_i^d$ let $l_i(m_i)$ be the menu l_i such that $\mathcal{A}_{l_i} = \text{Ra}(m_i)$. So, for $\theta_i \in \Theta_i$, let μ_i^{**} be defined as follows: if $l_i(\mu_i^*(\theta_i)) = \hat{l}_i$, then $\mu_i^{**}(\theta_i) = \hat{l}_i$, otherwise

$$\mu_i^{**}(\hat{\theta}_i, l_i; \theta_i) = \begin{cases} (\hat{\theta}_i, l_i) & \text{if } \hat{\theta}_i = \vartheta_i(\mu_i^*(\theta_i)) \text{ and } l_i = l_i(\mu_i^*(\theta_i)) \\ 0 & \text{else.} \end{cases}$$

The same construction applies to every principal i .

Now consider the agent's strategy α^{**} . If a profile of signals $\hat{\theta}$ and menus l in on path (i.e., it equals $\mu^{**}(\theta_{-0})$ for some θ_{-0}), then $\vartheta_i^{-1}(\hat{\theta}_i)$ is the mechanism $\mu_i^*(\hat{\theta}_i)$ in M^d with $l_i(\mu_i^*(\hat{\theta}_i)) = l_i$. So, consider the profile of mechanisms $m(\hat{\theta}, l)$ in which $m_i(\hat{\theta}_i, l_i) = \vartheta_i^{-1}(\hat{\theta}_i)$ and $m_i(\hat{l}_i) = \hat{m}_i$ (i.e., the mechanism with $\text{Ra}(\hat{m}_i) = \{\hat{d}_i\}$). Then, for every $\theta_0 \in \Theta_0$ and $d \in \mathcal{A}_l$, let

$$\alpha^{**}(d; \theta_0, \hat{\theta}_{-0}, l) = \sum_{\{a \in \mathcal{A} | m(\hat{\theta}, l)(a) = d\}} \alpha^*(a; \theta_0, m(\hat{\theta}, l)).$$

Suppose instead that $(\hat{\theta}, l)$ is off path. For every principal i , if $(\hat{\theta}_i, l_i) = \mu_i^{**}(\theta_i)$ for some θ_i let $\tilde{m}_i = \vartheta_i^{-1}(\hat{\theta}_i)$, and if $l_i = \hat{l}_i$ let $\tilde{m}_i = \hat{m}_i$. Otherwise, let \tilde{m}_i be any mechanism in M_i^d such that $\text{Ra}(\tilde{m}_i) = \mathcal{A}_{l_i}$. Given such a profile \tilde{m} , for every θ_0 and $d \in \mathcal{A}_l$, let

$$\alpha^{**}(d; \theta_0, \hat{\theta}, l) = \sum_{\{a \in \mathcal{A} | \tilde{m}(a) = d\}} \alpha^*(a; \theta_0, \tilde{m}).$$

I claim that $f_{\sigma^*}(\theta) = f_{\sigma^{**}}(\theta)$ for every θ . This is immediate because, for every θ_{-0} , the principals give the agent the same selection of profiles of decisions, and for every l on path, every θ_0 induces the same distribution over profiles of decisions as in the original equilibrium after the corresponding profile $\mu^*(\theta_{-0})$.

Now consider the agent's updated belief β^{**} . Suppose the profile $(\hat{\theta}, l)$ is on path. This means that $\mathbf{1}_{[\mu^{**}(\theta_{-0}) = (\hat{\theta}, l)]} > 0$ for some θ_{-0} . Thus, we have for every θ_{-0}

$$\begin{aligned} \beta^{**}(\theta_{-0} | \theta_0, \hat{\theta}, l) &= \frac{\mathbf{1}_{[\mu^{**}(\theta_{-0}) = (\hat{\theta}, l)]} b(\theta_{-0} | \theta_0)}{\sum_{\theta'_{-0}} \mathbf{1}_{[\mu^{**}(\theta'_{-0}) = (\hat{\theta}, l)]} b(\theta'_{-0} | \theta_0)} \\ &= \frac{\mathbf{1}_{[\mu^*(\theta_{-0}) = m(\hat{\theta}, l)]} b(\theta_{-0} | \theta_0)}{\sum_{\theta'_{-0}} \mathbf{1}_{[\mu^*(\theta'_{-0}) = m(\hat{\theta}, l)]} b(\theta'_{-0} | \theta_0)} = \beta^*(\theta_{-0} | \theta_0, m(\hat{\theta}, l)), \end{aligned}$$

where $m(\hat{\theta}, l)$ is the profile of mechanisms that I used before to define α^{**} on path. For every $(\hat{\theta}, l)$ off path, construct \tilde{m} as before and let $\beta^{**}(\cdot | \theta_0, \hat{\theta}, l) = \beta^*(\cdot | \theta_0, \tilde{m})$.

Given the agent's belief β^{**} , it follows that the strategy α^{**} is sequentially rational. Take any $(\hat{\theta}, l)$ on path—the same argument applies if $(\hat{\theta}, l)$ is off path. Given $(\hat{\theta}, l)$, the agent's set of available decisions and his belief over Θ_{-0} coincide with the set and belief that he had after observing $m(\hat{\theta}, l)$ in the original equilibrium σ^* ; moreover, α^{**} induces the same distribution over decisions as did α^* . Therefore, the agent cannot have a profitable deviation from α^{**} .

Given μ_{-i}^{**} and α^{**} , no principal i has a profitable deviation from μ_i^{**} . Suppose to the contrary that for some θ_i there exists an action $(\hat{\theta}_i, l_i)$ (or \hat{l}_i) such that

$$\bar{U}_i((\hat{\theta}_i, l_i), \mu_{-i}^{**}, \alpha^{**}; \theta_i) > \bar{U}_i(\mu_i^{**}(\theta_i), \mu_{-i}^{**}, \alpha^{**}; \theta_i).$$

Now, if $(\hat{\theta}_i, l_i)$ is on the path of play under μ_i^{**} , then it means that l_i corresponds to some $\tilde{m}_i = \mu_i^*(\tilde{\theta}_i)$ for $\tilde{\theta}_i \notin T_i(\mu_i^*(\theta_i))$. Otherwise, let \tilde{m}_i be the indirect mechanism I used in constructing α^{**} off path. Because the probability over outcomes that θ_i induces after deviating to $(\hat{\theta}_i, l_i)$ given μ_{-i}^{**} and α^{**} is identical to that induced after deviating to \tilde{m}_i given μ_{-i}^* and α^* , it must be that

$$\bar{U}_i(\tilde{m}_i, \mu_{-i}^*, \alpha^*; \theta_i) > \bar{U}_i(\mu_i^*(\theta_i), \mu_{-i}^*, \alpha^*; \theta_i).$$

A contradiction.

Part 2: (\Leftarrow) Let $(\sigma^{**}, \beta^{**})$ be an equilibrium of G^{SL} . I now derive an equilibrium (σ^*, β^*) of G^d , and show that $f_{\sigma^{**}} = f_{\sigma^*}$.

Consider the strategy μ_i^{**} for principal i . Let $\text{Ra}(\mu_i^{**})$ be the set of signals and menus that principal i can offer according to μ_i^{**} . For every $(\hat{\theta}_i, l_i) \in \text{Ra}(\mu_i^{**})$, let $S_i(l_i)$ be the set of signals that principal i sends while offering l_i . Now construct $|S_i(l_i)|$ indirect mechanisms $m_i(l_i)$ with $\text{Ra}(m_i(l_i)) = \mathcal{A}_i$. This can be done because G^d is an enlargement of G^{SL} and $|\Theta_i| \leq |\mathcal{A}_i|$ for every $i = 1, \dots, N$. Specifically, choose $|\Theta_i|$ messages in \mathcal{A}_i and label the elements in \mathcal{A}_i , other than the ‘default’ option \hat{x}_i , each time starting with a different element from the set of $|\Theta_i|$ selected messages. And when you get to \hat{x}_i , assign all remaining messages in \mathcal{A}_i to this option. This procedure delivers $|\Theta_i|$ different indirect mechanisms whose image equals the menu l_i . Then, assign one of these mechanisms to each $\hat{\theta}_i \in S_i(l_i)$ and denote it by $m_i(\hat{\theta}_i, l_i)$. Now, with a slight abuse of notation, let principal i strategy in G^d be defined as $\mu_i^*(\theta_i) = m_i(\mu_i^{**}(\theta_i))$, unless $\mu_i^{**}(\theta_i) = \hat{l}_i$ in which case let $\mu_i^*(\theta_i) = \hat{m}_i$. Finally, apply the same construction to obtain the offer strategy of every principal i .

Now consider the agent’s strategy α^* . If m is on path, then let $\mu^{**}((\mu^*)^{-1}(m)) = \times_{i=1}^N \mu_i^{**}((\mu_i^*)^{-1}(m_i))$. For every $d \in \text{Ra}(m)$ let $\mathcal{A}(d) = \{a \in \mathcal{A} : m(a) = d\}$ and choose one message profile $a(d) \in \mathcal{A}(d)$. Then for every θ_0 , let

$$\alpha^*(\hat{a}; \theta_0, m) = \begin{cases} \alpha^{**}(d; \theta_0, \mu^{**}((\mu^*)^{-1}(m))) & \text{if } \hat{a} = a(d) \\ 0 & \text{else} \end{cases}.$$

Now suppose \hat{m} is not on path. For every principal i , if $\hat{m}_i = \mu_i^*(\theta_i)$ for some θ_i let $(\tilde{\theta}_i, \tilde{l}_i) = \mu_i^{**}(\theta_i)$, and if $\hat{m}_i = \hat{m}_i$ let $\tilde{l}_i = \hat{l}_i$. Otherwise, let $(\tilde{\theta}_i, \tilde{l}_i)$ be any pair with $\mathcal{A}_i = \text{Ra}(\hat{m}_i)$. Use this construction to define the profile $(\tilde{\theta}, \tilde{l})$. Then, for every $\hat{d} \in \text{Ra}(\hat{m})$, let $\mathcal{A}(\hat{d}) = \{a \in \mathcal{A} : \hat{m}(a) = \hat{d}\}$ and choose one message profile $a(\hat{d}) \in \mathcal{A}(\hat{d})$. Then for every θ_0 , let

$$\alpha^*(a; \theta_0, \hat{m}) = \begin{cases} \alpha^{**}(\hat{d}; \theta_0, \tilde{\theta}, \tilde{l}) & \text{if } a = a(\hat{d}) \\ 0 & \text{else} \end{cases}.$$

I claim that $f_{\sigma^{**}} = f_{\sigma^*}$. This is because any profile of principals’ types θ_{-0} offers the agent the same choice set under μ^{**} and μ^* . Furthermore, given an array of mechanisms, the agent conditions the distribution over decisions on the specific designs of such indirect mechanisms as he was conditioning on the signals in the equilibrium of G^{SL} .

Consider now the agent's beliefs. Suppose first that the profile m is on path according to μ^* . Then, for every θ_0 and θ_{-0}

$$\begin{aligned}\beta^*(\theta_{-0}|\theta_0, m) &= \frac{\mathbf{1}_{[\mu^*(\theta_{-0})=m]}b(\theta_{-0}|\theta_0)}{\sum_{\theta'_{-0}} \mathbf{1}_{[\mu^*(\theta'_{-0})=m]}b(\theta'_{-0}|\theta_0)} \\ &= \frac{\mathbf{1}_{[\mu^{**}(\theta_{-0})=\mu^{**}((\mu^*)^{-1}(m))]}b(\theta_{-0}|\theta_0)}{\sum_{\theta'_{-0}} \mathbf{1}_{[\mu^{**}(\theta'_{-0})=\mu^{**}((\mu^*)^{-1}(m))]}b(\theta'_{-0}|\theta_0)} = \beta^{**}(\theta_{-0}|\theta_0, \mu^{**}((\mu^*)^{-1}(m))).\end{aligned}$$

If instead \hat{m} is off path. Let $(\tilde{\theta}, \tilde{l})$ be the profile of signals and menus that I constructed before to define α^* off path, and let

$$\beta^*(\theta_{-0}|\theta_0, \hat{m}) = \beta^{**}(\theta_{-0}|\theta_0, \tilde{\theta}, \tilde{l}).$$

I claim that the profile (σ^*, β^*) is an equilibrium of G^d . With regard to the agent's strategy, after any profile m that is on path, the agent has the same beliefs and induces the same distributions over profiles of decisions as after the profile of menus and signals $\mu^{**}((\mu^*)^{-1}(m))$. Hence, he cannot have any profitable deviation. If \hat{m} is off path, the same argument applies. With regard to the principals' strategies, suppose some θ_i of some principal i has a profitable deviation to a mechanism \hat{m}_i , that is

$$\bar{U}_i(\hat{m}_i, \mu_{-i}^*, \alpha^*; \theta_i) > \bar{U}_i(\mu_i^*(\theta_i), \mu_{-i}^*, \alpha^*; \theta_i).$$

Then it means that the way that the agent interprets θ_i 's deviation in forming his beliefs and the resulting distribution over decisions make θ_i strictly better off. On the other hand, by construction the deviation to \hat{m}_i is equivalent to offering \tilde{l}_i with $\mathcal{A}_{\tilde{l}_i} = \text{Ra}(\tilde{m}_i)$ and send the signal $\tilde{\theta}_i$ that I used in the construction of α^* and β^* above. Since the agent is responding to μ_{-i}^* as he was responding to μ_{-i}^{**} , it follows that

$$\bar{U}_i((\tilde{\theta}_i, \tilde{l}_i), \mu_{-i}^{**}, \alpha^{**}; \theta_i) > \bar{U}_i(\mu_i^{**}(\theta_i), \mu_{-i}^{**}, \alpha^{**}; \theta_i).$$

A contradiction.

7.2 Proof of Lemma 1

Let (σ^*, β^*) be an equilibrium of G^Λ . I will construct a truthful equilibrium of G^Λ that sustains $f_{\sigma^{**}} = f_{\sigma^*}$.

Consider first principal i . For every θ_i , let $\lambda_*^{T'_i}(\theta_i) = \mu_i^*(\theta_i)$ and let $T_i(\theta_i) = (\mu_i^*)^{-1}(\lambda_*^{T'_i}(\theta_i)) \subset \Theta_i$ —clearly, $\lambda_*^{T'_i}(\theta'_i) = \lambda_*^{T'_i}(\theta_i)$ if and only if $\theta'_i \in T_i(\theta_i)$. Now for each θ_i , construct a new list $\lambda_*^{T_i}(\theta_i)$ as follows: for each $r^{T'_i} \in \mathcal{A}_{\lambda_*^{T'_i}(\theta_i)}$, let $\rho^{\Theta_i}(r^{T'_i})(\cdot) = r^{T'_i}(\cdot) \circ \pi_i^*(\cdot, \lambda_*^{T'_i}(\theta_i))$ ²⁰ and let $\rho^{T_i}(r^{T'_i}) = \rho^{\Theta_i}(r^{T'_i})|_{T_i(\theta_i)}$ which is an RDM in $R^{T_i(\theta_i)}$. Repeating this operation and collecting the resulting RDMS, we obtain the new list $\lambda_*^{T_i}(\theta_i)$. Clearly, $\lambda_*^{T_i}(\theta_i) = \lambda_*^{T_i}(\theta'_i)$ if and only if $\theta'_i \in$

²⁰Given a RDM $r^{T_i}(\cdot)$ and a communication strategy $\pi_i(\cdot; \lambda^{T_i})$, the operation \circ simply reduces the compound lottery defined by $r^{T_i}(\cdot)$ and $\pi_i(\cdot; \lambda^{T_i})$ to obtain a direct mechanism as in (3).

$T_i(\theta_i)$. So, for every θ_i let $\mu_i^{**}(\theta_i) = \lambda_{**}^{T_i}(\theta_i)$, and $\pi_i^{**}(\theta_i, \mu_i^{**}(\theta_i)) = \theta_i$, whereas for $\lambda^{T_i} \neq \mu_i^{**}(\theta_i)$, let $\pi_i^{**}(\theta_i; \lambda^{T_i}) = \pi_i^*(\theta_i; \lambda^{T_i})$. Finally, repeat the same construction for every principal i .

Now consider the agent. Suppose first that λ^T is on path under μ^{**} . Then, for every θ_0 and every $r^T \in \mathcal{A}_{\lambda^T}$, let

$$\alpha^{**}(r^T; \theta_0, \lambda^T) = \sum_{\{r^{T'} \in \mathcal{A}_{\lambda_*^{T'}} : \rho^T(r^{T'}) = r^T\}} \alpha^*(r^{T'}; \theta_0, \lambda_*^{T'}),$$

where $\lambda_*^{T'} = \times_{i=1}^N \mu_i^*((\mu_i^{**})^{-1}(\lambda_*^{T'}))$. Also, using Bayes rule it follows that on path, for every θ ,

$$\begin{aligned} \beta^{**}(\theta_{-0} | \theta_0, \lambda^T) &= \frac{\mathbf{1}_{[\mu^{**}(\theta_{-0}) = \lambda^T]} b(\theta_{-0} | \theta_0)}{\sum_{\theta_{-0}} \mathbf{1}_{[\mu^{**}(\theta'_{-0}) = \lambda^T]} b(\theta'_{-0} | \theta_0)} \\ &= \frac{\mathbf{1}_{[\mu^*(\theta_{-0}) = \lambda_*^{T'}]} b(\theta_{-0} | \theta_0)}{\sum_{\theta_{-0}} \mathbf{1}_{[\mu^*(\theta'_{-0}) = \lambda_*^{T'}]} b(\theta'_{-0} | \theta_0)} = \beta^*(\theta_{-0} | \theta_0, \lambda_*^{T'}), \end{aligned}$$

where again $\lambda_*^{T'} = \times_{i=1}^N \mu_i^*((\mu_i^{**})^{-1}(\lambda_*^{T'}))$.

I claim that $f_{\sigma^{**}} = f_{\sigma^*}$. To see this, note that for every θ , both through (μ^*, π^*) and through (μ^{**}, π^{**}) , the principals' essentially give the agent the possibility to choose among the same profiles of lotteries over allocations; furthermore, the agent actually induces the same distribution over such profiles under α^* as well as under α^{**} .

To complete the construction of the equilibrium $(\sigma^{**}, \beta^{**})$, for every λ^T off path let $\alpha^{**}(\cdot; \cdot, \lambda^T) = \alpha^*(\cdot; \cdot, \lambda^T)$ and $\beta^{**}(\cdot | \cdot, \lambda^T) = \beta^*(\cdot | \cdot, \lambda^T)$. Using the assumption that (σ^*, β^*) is an equilibrium of G^Λ , it follows that the profile $(\sigma^{**}, \beta^{**})$ is also an equilibrium of G^Λ . Furthermore, $(\sigma^{**}, \beta^{**})$ is a truthful equilibrium.

7.3 Proof of Theorem 2

Part 1: (\Rightarrow) Let (σ^*, β^*) be an equilibrium of G sustaining f . I will construct a truthful equilibrium $(\sigma^{**}, \beta^{**})$ of G^Λ that sustains f .

Given the offer strategy μ_i^* of principal i , let $\mu_i^*(\Theta_i)$ be the set of mechanisms that are offered by some θ_i in the equilibrium. For every θ_i identify the subset $T_i(\mu_i^*(\theta_i)) = (\mu_i^*)^{-1}(\mu_i^*(\theta_i)) \subset \Theta_i$, that is the set of types of principal i who offer the same mechanism $\mu_i^*(\theta_i)$. Using the communication strategy π_i^* and the function $c(\cdot, \pi_i^*)$,²¹ for every θ_i recover the list of DMs $\lambda^{\Theta_i}(\theta_i) = c(\mu_i^*(\theta_i), \pi_i^*)$ that corresponds to the mechanism offered by θ_i under μ_i^* . Finally, for every θ_i identify the list of RDM $\lambda^{T_i}(\theta_i)$ in $R^{T_i(\mu_i^*(\theta_i))}$ by restricting every r^{Θ_i} in $\lambda^{\Theta_i}(\theta_i)$ to $T_i(\mu_i^*(\theta_i))$. By construction, for every θ_i and $\theta'_i \in T_i(\mu_i^*(\theta_i))$, $\lambda^{T_i}(\theta_i) = \lambda^{T_i}(\theta'_i)$. Therefore, let principal i 's offer strategy in G^Λ be $\mu_i^{**}(\theta_i) = \lambda^{T_i}(\theta_i)$. Finally, repeat the same construction for every principal i to obtain the profile of offer strategies μ^{**} .

Now consider the communication strategies of the principals and the agent on path. For every θ_i of principal i , let $\pi_i^{**}(\theta_i, \mu_i^{**}(\theta_i)) = \theta_i$. With regard to the agent, given λ^T on path, let $m(\lambda^T) = \mu^*((\mu^{**})^{-1}(\lambda^T))$, and for every $r^T \in \mathcal{A}_{\lambda^T}$, let $\mathcal{A}(r^T) = \{a : (m_i(\lambda^{T_i})(a_i, \cdot) \circ \pi_i(\cdot, m(\lambda^{T_i})))|_{T_i}\}_{i=1}^N = r^T\}$. In words, $\mathcal{A}(r^T)$ is the set of messages that allowed the agent to induce a profile of DMs

²¹Recall the discussion before Definition 6.

(given π^*) such that, if we restrict each of them to the corresponding T_i , we obtain r^T . Then for each θ_0 , let

$$\alpha^{**}(r^T; \theta_0, \lambda^T) = \sum_{a \in \mathcal{A}(r^T)} \alpha^*(a; \theta_0, m(\lambda^T)).$$

To derive the on-path beliefs of the agent, we can use Bayes rule as follows: for every θ_{-0}, θ_0 and λ^T on path

$$\begin{aligned} \beta^{**}(\theta_{-0} | \theta_0, \lambda^T) &= \frac{\mathbf{1}_{[\mu^{**}(\theta_{-0}) = \lambda^T]} b(\theta_{-0} | \theta_0)}{\sum_{\theta'_{-0}} \mathbf{1}_{[\mu^{**}(\theta'_{-0}) = \lambda^T]} b(\theta'_{-0} | \theta_0)} \\ &= \frac{\mathbf{1}_{[\mu^*(\theta_{-0}) = m(\lambda^T)]} b(\theta_{-0} | \theta_0)}{\sum_{\theta'_{-0}} \mathbf{1}_{[\mu^*(\theta'_{-0}) = m(\lambda^T)]} b(\theta'_{-0} | \theta_0)} = \beta^*(\theta_{-0} | \theta_0, m(\lambda^T)). \end{aligned}$$

I claim that $f_{\sigma^*} = f_{\sigma^{**}}$. Fix any $\theta \in \Theta$ and let m^* and λ_{**}^T be the profiles of indirect mechanisms and lists of RDMs offered by θ_{-0} under σ^* and σ^{**} , we have

$$\begin{aligned} f_{\sigma^*}(x; \theta) &= \sum_{\mathcal{A}} \prod_{i=1}^N [\sum_{\mathcal{P}_i} d_i(x_i) m_i^*(d_i; a_i, p_i) \pi_i^*(p_i; \theta_i, m_i^*)] \alpha^*(a; \theta_0, m^*) \\ &= \sum_{r^T \in \mathcal{A}_{\lambda_{**}^T}} \prod_{i=1}^N r^{T_i}(x_i; \theta_i) \sum_{a \in \mathcal{A}(r^T)} \alpha^*(a; \theta_0, m^*) = f_{\sigma^{**}}(x; \theta). \end{aligned}$$

Using this result, we can see that (π^{**}, α^{**}) is a continuation equilibrium given any λ^T on path and β^{**} . This is because, given any (θ_0, λ^T) on path, the agent has the same belief over θ_{-0} and, given π^{**} , he expects the same distribution over outcomes by playing α^{**} as when he was playing α^* after $(\theta_0, m(\lambda^T))$. Furthermore, given π^{**} the agent can't induce any distribution over outcomes that he couldn't induce after $m(\lambda^T)$ given π^* . Similarly, each θ_i of principal i can't have a profitable deviation from her truthful strategy after offering λ^{T_i} , because any such deviation was available also in the original game after offering $m_i(\lambda^{T_i})$.

If λ^T is off path and involves deviations by more than one principal, choose any belief for the agent $\beta^{**}(\cdot | \cdot, \lambda^T)$ and equilibrium of the resulting continuation game—finiteness of \mathcal{A}_{λ^T} and Θ_{-0} ensures the existence of at least one continuation equilibrium. Now consider λ^T off path induced by the deviation of one single principal i . Suppose first that θ_i of principal i offers $\lambda^{T_i} = \mu_i^{**}(\theta'_i)$ for $\theta'_i \neq \theta_i$. Then the agent's strategy and beliefs are defined as on path. Let $\pi_i^{**}(\theta_i, \lambda^{T_i}) \in \Delta(T_i)$ be any best reply to α^{**} and σ_{-i}^{**} . We have that, for any principal i , no type θ_i can profit from such a deviation because it was already available to θ_i in the equilibrium of the original game G . The remaining case is when λ^T is off path (in the sense that it can't occur with positive probability according to σ_i^{**}) and only principal i has deviated for some θ_i . I claim that there must exist a continuation equilibrium that makes such a deviation unprofitable. Suppose to the contrary that θ_i of principal i deviates to offering $\hat{\lambda}^{T_i}$ off path, whereas all other principals follow σ_{-i}^{**} , and that for any choice of $\beta^{**}(\cdot | \cdot, (\hat{\lambda}^{T_i}, \lambda^{T_{-i}}))$ —where $\lambda^{T_{-i}}$ results from μ_{-i}^{**} —any continuation equilibrium can't deter θ_i from such a deviation. I will argue that then (σ^*, β^*) can't be an equilibrium of G , which is a contradiction. To see this recall that, since G is an enlargement of G^Λ , for such an $\hat{\lambda}^{T_i}$ there exist an indirect mechanism \tilde{m}_i and a communication strategy $\tilde{\pi}_i(\tilde{m}_i)$ for such a mechanism, such that restricting the DMs in the list $c(\tilde{m}_i, \tilde{\pi}_i)$ to T_i delivers $\hat{\lambda}^{T_i}$. Hence, consider the mechanism \hat{m}_i such that, for every $\theta_i \in T_i$, there

is a message $p_i(\theta_i)$ such that $\hat{m}_i(\cdot, p_i(\theta_i)) = \tilde{m}_i(\cdot, \tilde{\pi}_i(\theta_i, \tilde{m}_i))$ and, for every $p_i \notin \{p_i(\theta_i)\}_{\theta_i \in T_i}$, $\hat{m}_i(\cdot, p_i) = \hat{m}_i(\cdot, p_i(\theta'_i))$ where θ'_i is the element of T_i with the lowest index—this choice is just for convenience. Then, for type θ_i of principal i , offering \hat{m}_i in G followed by the message $p_i(\theta_i)$ is essentially equivalent to offering $\hat{\lambda}^{T_i}$ in G^Λ followed by the report θ_i . Now, given any realization m_{-i}^* according to μ_{-i}^* , the set of distribution over outcomes that the agent can induce after (\hat{m}_i, m_{-i}^*) given $(\hat{\pi}_i, \pi_{-i}^*)$ is the same as after the corresponding profile $(\hat{\lambda}^{T_i}, \lambda^{T_{-i}})$ given π^{**} , if we define $\hat{\pi}_i$ to play $p_i(\theta_i)$ after \hat{m}_i with the same probability with which π_i^{**} plays θ_i after $\hat{\lambda}^{T_i}$. By assumption, it is not possible to find a belief β^{**} and a strategy α^{**} so that α^{**} is a best reply to π^{**} and vice versa, and θ_i of principal i doesn't profit by deviating to $\hat{\lambda}^{T_i}$. But then, given any specification of the agent's belief after (\hat{m}_i, m_{-i}^*) , there can't be a strategy $\hat{\alpha}$ that delivers, together with $\hat{\pi}_i$, a continuation equilibrium such that θ_i of principal i is deterred from offering \hat{m}_i followed by $\hat{\pi}_i$. This contradicts the assumption that (σ^*, β^*) is an equilibrium of G . We conclude that it is possible to extend the on-path strategies and beliefs that I have constructed above to obtain an equilibrium $(\sigma^{**}, \beta^{**})$ of G^Λ such that $f_{\sigma^{**}} = f_{\sigma^*}$.

Part 2: (\Leftarrow) Suppose $(\sigma^{**}, \beta^{**})$ is a truthful equilibrium of G^Λ that sustains f . I will construct an equilibrium (σ^*, β^*) of G that sustains f .

Consider principal i and her offer strategy μ_i^{**} . For every $\lambda^{T_i} \in \mu_i^{**}(\Theta_i)$, consider the mechanism $m(\lambda^{T_i}) \in M_i$ constructed as follows. Since G is an enlargement of G^Λ , there exist \tilde{m}_i and $\tilde{\pi}_i$ such that restricting the DMs in the list $c(\tilde{m}_i, \tilde{\pi}_i)$ to T_i delivers λ^{T_i} . Since $|\Theta_i| \leq |\mathcal{P}_i|$, assign to each $\theta_i \in \Theta_i$ a message $p_i(\theta_i)$. Now let $m(\lambda^{T_i})(\cdot, p_i(\theta_i)) = \tilde{m}_i(\cdot, \tilde{\pi}_i(\theta_i, \tilde{m}_i))$ for every $\theta_i \in T_i$, and for every $p_i \notin \{p_i(\theta_i)\}_{\theta_i \in T_i}$, let $m(\lambda^{T_i})(\cdot, p_i) = m(\lambda^{T_i})(\cdot, p_i(\theta'_i))$ where θ'_i is the element of T_i with the lowest index—again, this choice is just for convenience. There are two cases to consider. The first case corresponds to $|T_i| > 1$. In this case, each $m(\lambda^{T_i})$ signals the same amount of information about principal i 's type as λ^{T_i} because λ^{T_i} and $m(\lambda^{T_i})$ restrict principal i to choose a message in T_i and in $\{p_i(\theta_i)\}_{\theta_i \in T_i}$, respectively. The second case corresponds to $|T_i| = 1$, which implies that λ^{T_i} is essentially a simple menu. In this case, we can match different θ_i s to different mechanisms m_i s, even when these mechanisms correspond to the same menu, as in the proof of Theorem 1. So, denote by $\lambda^{T_i}(\theta_i)$ the list of RDMS that θ_i offers under μ_i^{**} and let $m(\lambda^{T_i}(\theta_i))$ be the indirect mechanism that we have just constructed. Then, for every principal i and θ_i , define $\mu_i^*(\theta_i) = m(\lambda^{T_i}(\theta_i))$.

Now, consider the on-path communication strategies of the principals and the agent. First, for every principal i and θ_i , let $\pi_i^*(\theta_i, m(\lambda^{T_i}(\theta_i))) = p_i(\theta_i)$ where $p_i(\theta_i)$ was defined before. Now consider the agent. Given any profile m^* on path, recover the corresponding $\lambda^T(m^*)$ given the above construction of m^* . Also, for every $r^T \in \mathcal{A}_{\lambda^T(m^*)}$, let $\mathcal{A}(r^T) = \{a : ((m_i^*(a_i, \cdot) \circ \pi_i^*(\cdot, m_i^*))|_{T_i})_{i=1}^N = r^T\}$ and choose *one* message $a(r^T) \in \mathcal{A}(r^T)$. Then for each θ_0 , define

$$\alpha^*(\tilde{a}; \theta_0, m^*) = \begin{cases} \alpha^{**}(r^T; \theta_0, \lambda^T(m^*)) & \text{if } \tilde{a} = a(r^T) \\ 0 & \text{else} \end{cases}.$$

The agent's on-path beliefs follow from Bayes rule: for every θ_{-0}, θ_0 and m^* on path

$$\begin{aligned} \beta^*(\theta_{-0} | \theta_0, m^*) &= \frac{\mathbf{1}_{[\mu^*(\theta_{-0})=m^*]} b(\theta_{-0} | \theta_0)}{\sum_{\theta'_{-0}} \mathbf{1}_{[\mu^*(\theta'_{-0})=m^*]} b(\theta'_{-0} | \theta_0)} \\ &= \frac{\mathbf{1}_{[\mu^{**}(\theta_{-0})=\lambda^T(m^*)]} b(\theta_{-0} | \theta_0)}{\sum_{\theta'_{-0}} \mathbf{1}_{[\mu^{**}(\theta'_{-0})=\lambda^T(m^*)]} b(\theta'_{-0} | \theta_0)} = \beta^{**}(\theta_{-0} | \theta_0, \lambda^T(m^*)). \end{aligned}$$

Given the construction of (μ^*, π^*) and α^* on path, it follows that $f_{\sigma^*} = f_{\sigma^{**}}$. This is because each type θ_i of each principal i is offering the same options to the agent and is constraining herself to the same communication possibilities as in the profile (μ^{**}, π^{**}) , and—given that—the agent is inducing the same distribution over maps from the principals’ messages to outcomes as under α^{**} . This last observation implies that (π^*, α^*) is a continuation equilibrium given the agent’s belief β^* .

It remains to show that π^* , α^* , and β^* can be extended off path to obtain an equilibrium (σ^*, β^*) of G . If m is off path because more than one principal i deviated, let $\beta^*(\cdot | \cdot, m)$ be any belief of the agent and let (π^*, α^*) be any equilibrium of the resulting continuation game. Now suppose \hat{m} is off path because only principal i deviated for some θ_i . If $\hat{m}_i = \mu_i^*(\tilde{\theta}_i)$ for some $\tilde{\theta}_i \neq \theta_i$, then the agent’s strategy and beliefs are as if \hat{m} is on path. Hence, for θ_i , offering \hat{m}_i followed by some communication strategy—given α^* and (μ_{-i}^*, π_{-i}^*) in G —is equivalent to offering $\lambda^{T_i}(\tilde{\theta}_i)$, again followed by some communication strategy, given α^{**} and $(\mu_{-i}^{**}, \pi_{-i}^{**})$ in G^Λ . Therefore, θ_i can’t gain by offering \hat{m}_i . Now consider the case with $\hat{m}_i \neq \mu_i^*(\tilde{\theta}_i)$ for all $\tilde{\theta}_i$. Suppose that for any specification of $\beta^*(\cdot | \cdot, \hat{m})$, any continuation equilibrium can’t deter θ_i from deviating to \hat{m}_i . Consider the communication strategy $\hat{\pi}_i(\cdot, \hat{m}_i)$ in any such continuation equilibrium, and let $\hat{\lambda}^{\Theta_i} = c(\hat{m}_i, \hat{\pi}_i)$. Hence, for θ_i of principal i , offering \hat{m}_i followed by $\hat{\pi}_i$ in G is equivalent to offering $\hat{\lambda}^{\Theta_i}$ in G^Λ followed by a truthful report of θ_i . Given the profile σ_{-i}^{**} and any λ^{T-i} resulting from μ_{-i}^{**} , the set of distributions over outcomes that the agent can induce after $(\hat{\lambda}^{\Theta_i}, \lambda^{T-i})$ given π^{**} is the same as the set that the agent can induce after the corresponding profile (\hat{m}_i, m_{-i}^*) given $(\hat{\pi}_i, \pi_{-i}^*)$. By assumption, for any specification of $\beta^*(\cdot | \cdot, (\hat{m}_i, m_{-i}^*))$ there is no $\alpha^*(\cdot, (\hat{m}_i, m_{-i}^*))$ that is a best reply to $\hat{\pi}_i$ and vice versa—given σ_{-i}^* —and can deter the deviation to \hat{m}_i by θ_i of principal i . But then, there is no specification of $\beta^{**}(\cdot | \cdot, (\hat{\lambda}^{\Theta_i}, \lambda^{T-i}))$ for which a continuation equilibrium can be constructed that deters θ_i of principal i from deviating to λ^{Θ_i} in G^Λ . This implies that $(\sigma^{**}, \beta^{**})$ can’t be an equilibrium of G^Λ . A contradiction that implies that it is possible to complete π^*, α^* and β^* off path to make any deviation from the on path behavior constructed above unprofitable, and thus obtaining an equilibrium (σ^*, β^*) of G .

References

- [1] Bernheim, D. and M. Whinston (1986a), *Menu Auctions, Resource Allocation, and Economic Influence*. Quarterly Journal of Economics, Vol. 101, No. 1, 1-32
- [2] Bernheim, D. and Whinston, M. (1986b), *Common Agency*. Econometrica, Vol. 54, No.4, pp. 923-942.
- [3] Celik, G. and M. Peters (2011), *Reciprocal Relationships and Mechanism Design*, Mimeo, University of British Columbia.
- [4] Epstein, L. and Peters, M. (1999), *A Revelation Principle for Competing Mechanisms*. Journal of Economic Theory, 88, 119-160.
- [5] Gibbard, A., (1973), *Manipulation of Voting Schemes: A General Result*. Econometrica, Vol. 41, No. 4, pp. 587-601.

- [6] Green, J., Laffont, J.-J. (1977), *Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods*. *Econometrica*, Vol. 45, No. 2, pp. 427-38.
- [7] Guesneries, R. (1995), *A Contribution to the Pure Theory of Taxation*. Cambridge, UK. Cambridge University Press.
- [8] Katz, M. L. (1991), *Game-Playing Agents: Unobservable Contracts as Precommitments*. *RAND Journal of Economics*, Vol. 22, No. 3, pp. 307-28.
- [9] Laffont, J.-J., Tirole, J., (1988), *The Dynamics of Incentive Contracts*. *Econometrica*, Vol. 56, No. 5, pp. 1153-1175.
- [10] Martimort, D., Moreira, H., (2010), *Common Agency and Public Good Provision under Asymmetric Information*. *Theoretical Economics*, Vol. 5, No. 2, 159-213.
- [11] Martimort, D. and L. Stole, (1997), *Communication Spaces, Equilibria Sets and the Revelation Principle under Common Agency*. Unpublished.
- [12] Martimort, D. and L. Stole, (2002), *The Revelation and Delegation Principles in Common Agency Games*. *Econometrica* 70, 1659-1674.
- [13] Maskin, E. and J. Tirole, (1990), *The Principal-Agent Relationship with an Informed Principal, I: The Case of Private Values*. *Econometrica*, 58: 379-410.
- [14] Maskin, E. and J. Tirole, (1992), *The Principal-Agent Relationship with an Informed Principal, II: The Case of Common Values*. *Econometrica*, 60: 1-42.
- [15] McAfee, R. P. (1993), *Mechanism Design by Competing Sellers*, *Econometrica*. Vol. 61, No. 6, pp. 1281-312.
- [16] Myerson, R., (1983), *Mechanism Design with an Informed Principal*. *Econometrica*, 51: 1767-1797.
- [17] Myerson, R., (1997), *Game Theory: analysis of conflict*. Harvard University Press.
- [18] Mylovanov, T. and Tröger, T. (2011), *Informed-Principal Problems in Environments with Generalized Private Values*, *Theoretical Economics*, forthcoming.
- [19] Fudenberg, D. and Tirole, J., (1991), *Game Theory*. MIT Press.
- [20] Pavan, A. and G. Calzolari (2010), *Truthful Revelation Mechanisms for Simultaneous Common Agency Games*. *AEJ: Microeconomics*, 2(2), 132-190.
- [21] Peck, J. (1997), *A Note on Competing Mechanisms and the Revelation Principle*. Unpublished.
- [22] Peters, M., (2001), *Common Agency and the Revelation Principle*. *Econometrica*, 69: 1349-1372.
- [23] Peters, M., (2011), *A Revelation Principle for Competing Mechanism Games*. Mimeo, University of British Columbia.

- [24] Peters, M. and B. Szentes (2011), *Definable and Contractible Contracts*, *Econometrica* (forthcoming).
- [25] Rochet, J.-C. (1986), *Le Controle Des Equations Aux Derivees Partielles Issues de la Theorie Des Incitations*. PhD Thesis, Universite Paris IX.