

Galperti, Simone

Working Paper

Common agency with informed principals: Menus and signals

Discussion Paper, No. 1541

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Galperti, Simone (2011) : Common agency with informed principals: Menus and signals, Discussion Paper, No. 1541, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/59656>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



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
October 31, 2011

JEL Classification: Common agency, informed principals, Inscrutability Principle, Delegation Principle, menus, signals, direct revelation mechanisms.

Keywords: C72, D82, D83, D86



NORTHWESTERN
UNIVERSITY



Common Agency with Informed Principals: Menus *and* Signals

Simone Galperti*

Northwestern University

October 31, 2011

Abstract

I analyze common agency games in which the principals, and possibly the agent, have private information. I distinguish between games in which the principals delegate the final decisions to the agent, and games in which they retain some decision power after offering their mechanisms. I show that, in contrast with mechanism design models with *one* informed principal, Myerson's Inscrutability Principle fails when there are *many* informed principals. I also find that, in contrast with common agency models with *uninformed* principals, the Delegation Principle (Menu Theorem) fails when principals are *informed*. I then focus on Perfect Bayesian Equilibria in which principals offer their mechanisms without randomizing. I characterize the outcomes of arbitrary games with delegation as outcomes of a new game in which principals offer menus and send cheap-talk signals. Next, I characterize the outcomes of arbitrary games without delegation as outcomes of a new game in which principals offer menus of direct revelation mechanisms, to which they truthfully report their types.

KEYWORDS: Common agency, informed principals, Inscrutability Principle, Delegation Principle, menus, signals, direct revelation 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 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

The theory of common agency games has experienced a long development since Bernheim and Whinston's (1986b) seminal paper, and has achieved high levels of generality and complexity (see, e.g., Epstein and Peters (1999), Peters (2001), Martimort and Stole (2002), and Pavan and Calzolari (2010)). Until very recently—see Martimort and Moreira (2010)—this theory has always assumed that principals have no private information. There are, however, many contracting environments which correspond to common agency games with *informed* principals. For example, this class of games includes Martimort and Moreira's model of public good provision: in their model, contributors privately know how much they value the public good and offer the common provider contribution schemes to tie their payments to the actual production of such good. Alternatively, consider menu or package auctions as in Bernheim and Whinston's (1986a) model, in which bidders submit a bidding schedule to the auctioneer, who then chooses the allocation that maximizes his profits. It seems natural to extend Bernheim and Whinston's analysis to allow bidders to have private valuations as in standard auction settings.

In the present paper, I focus on characterizing the equilibrium outcomes—intended as maps from type profiles to allocations—of common agency games (hereafter, games) with informed principals. Following the path suggested by Peters (2001), Martimort and Stole (2002), and Pavan and Calzolari (2010) for games with uninformed principals, my goal is to identify a tractable class of games that allows to replicate all equilibrium outcomes of some underlying general game with informed principals. Similarly to those papers, my analysis identifies the key strategic roles that mechanisms play in such games.

The main novelty of my work is therefore to allow principals to have private information when they contract with their common agent. Because of this novel feature, my model departs from a standard model with uninformed principals also in another dimension: I consider both the case in which the principals can't communicate with their mechanisms, as well as the case in which they can do so. I refer to the first class of games as games with delegation, and to the second class as games without delegation.

Otherwise, my model is equivalent to a standard model with uninformed principals (Bernheim and Whinston (1986b)). Most importantly, I assume that contracting is private and that each principal can't do three things: 1) directly communicate with the other principals; 2) condition her mechanisms on the mechanisms of the other principals; and 3) condition her allocations on the allocations of the other principals. These restrictions are the essence of common agency games; they are natural assumptions in the contexts of menu auctions and public good provision. I also follow the literature in using (weak) Perfect Bayesian Equilibrium (PBE) as my solution concept. This concept imposes several refinements—sequential rationality and on-path Bayesian updating—which have important implications for my analysis.

By introducing private information on the principals' side, this paper also relates to the literature on mechanism design with *one* informed principal (Myerson (1983), Maskin and Tirole (1990-1992)). In his seminal work about the "theory of inscrutable mechanism selection," Myerson (1983) introduces a fundamental result known as the Inscrutability Principle. Roughly, this Principle says that if an informed principal can sustain a certain outcome by revealing some of her information through the mechanisms she offers, then she can also sustain the same outcome

with a sole, uninformative mechanism. The Inscrutability Principle is important because it implies that the signaling component of an informed principal's offers is irrelevant, and hence it is safe to assume that all types of the principal choose the same mechanism.

Given the evident usefulness of the Inscrutability Principle, in Section 3 I begin by investigating whether it holds also in contexts with multiple informed principals. Unfortunately, the answer is no. I argue this important point with an example that explains the fundamental difference between having many rather than one informed principal. Intuitively, if principal i reveals her type up front, then the agent may sustain an outcome that features some correlation between principal i 's type and the allocations with all the principals. To achieve such correlation, the agent clearly needs to learn principal i 's type before he communicates with all principals. This sequence of events is, however, impossible if principal i selects her mechanism in an inscrutable way.

The failure of the Inscrutability Principle in games with many informed principals has several conceptual and practical implications. It implies that, contrary to the case with one informed principal, signaling through mechanisms becomes an essential strategic component of the game. Therefore, if we want to characterize all equilibrium outcomes, we can't safely assume that each principal offers the same mechanism independently of her private information; we rather have to allow each principal to affect the agent's beliefs *before* he communicates with her opponents. Dealing with this signaling purpose of the principals' mechanisms is one of the main difficulties of the present analysis.

In Section 4, I turn my attention to another important principle, this time from the literature on common agency with uninformed principals: the Delegation Principle (Martimort and Stole (2002)), also known as the Menu Theorem (Peters (2001)). This principle says the following: any equilibrium outcome of a game in which principals compete in arbitrary indirect mechanisms is also an equilibrium outcome of a much simpler game in which principals offer menus of allocations and delegate the agent to pick one; the converse holds too. The Delegation Principle represents a key step to characterize the outcomes of standard games (see, e.g., Martimort and Stole (2002), and Pavan and Calzolari (2010)). Moreover, it provides important insights to understand the strategic nature of such games.

Similarly to the Inscrutability Principle, also the Delegation Principle would be extremely useful to characterize the outcomes of the games in this paper. Unfortunately, it fails when principals have private information. The first reason of this result is tightly linked to the failure of the Inscrutability Principle, and is easier to understand for games with delegation. In these games the informed principals use their mechanisms for two purposes: 1) define the set of allocations that the agent can choose, and 2) signal their private information. Relative to arbitrary indirect mechanisms, menus certainly allow the principals to achieve the first goal, but they may impose restrictions on the second goal. Intuitively, with arbitrary mechanisms principal i can offer the agent the same set of alternatives, but frame it in different ways to signal her private information.

In games without delegation, the Delegation Principle fails also for another, obvious reason. Without delegation, principal i can keep the agent uncertain about what her mechanism will select until after he communicates with all principals, and then let her type determine how this uncertainty resolves. In contrast, even if principal i can offer menus of lotteries over allocations,

she can't tie the outcome of these lotteries with her types.

In Section 5, I turn to the two main results of the paper. Building on the remarks about the Inscrutability and the Delegation Principles, I provide two characterizations of the outcomes of the games with informed principals—one for the games with delegation and one for those without delegation. However, in both cases I focus on equilibria in which principals do not randomize when offering their mechanisms—I call these equilibria "deterministic-offer equilibria" (D-O equilibria for short).¹ This kind of equilibria represents the focus of almost the entire literature on common agency due to their realism and simplicity. D-O equilibria involve a straightforward kind of signaling, which simplifies my characterizations: after observing principal i 's offer, the agent learns at most that i 's type is in a subset of all her possible types, but he assigns to the types in this subset the same relative likelihood as under his prior belief.

In the case of games with delegation, I show that the outcomes sustained by D-O equilibria of an underlying game coincide with the outcomes sustained by D-O equilibria of the following "signal-menu game." In such a game the principals offer menus of alternatives, as in the standard menu games, but they also privately send to the agent a "cheap talk" signal. Furthermore, each principal's set of signals equals her set of types (any set with the same cardinality would work).

Obtaining this equivalence result requires two conditions. First, similarly to the Delegation Principle, it is essential that in the new "signal-menu game" each principal have the same ability to contract as in the underlying game. To ensure this property, I rely on the standard notion of enlargement from the literature on games with uninformed principals (see, e.g., Pavan and Calzolari (2010)). Second, it is crucial that in the new game each principal also has the same ability to signal as in the underlying game. Using the principals' set of types as their set of signals does the job, if we focus on D-O equilibria. Note that, although my equivalence result covers only the outcomes sustained by D-O equilibria, it relies on an immediate and minor extension of the standard menu game of Peters (2001) and Martimort and Stole (2002).

In contrast, characterizing the outcomes of the games without delegation requires a more substantial departure from simple menus. My characterization, however, retains the spirit of menus and combines it with the classic notion of direct revelation mechanisms (hereafter, DRMs). In fact, as in games with delegation all indirect mechanisms are contractually equivalent to menus of allocations, in games without delegation they are contractually equivalent to menus of DRMs. By offering such a menu, principal i lets the agent pick a DRM, which she then plays by reporting some of her types. Menus of DRMs may, however, limit the principals' ability to signal relative to arbitrary indirect mechanisms. To solve this issue, I suggest to allow each principal to offer menus of *restricted* DRMs, in the sense that all DRMs in any such menu restrict principal i to report only a subset of her possible types. Using these menus, principal i can rely on the domain of the DRMs as a signaling device. Intuitively, suppose that in an equilibrium of the underlying game principal i reveals her two types by offering different indirect mechanisms. Then in my new game, principal i can send the same signal by offering two menus that commit her to report only her true type.

I derive my second main result using a new game in which principals offer menus of restricted DRMs. I show that the outcomes sustained by D-O equilibria of a game without delegation

¹A partial characterization of equilibria with mixed offer strategies is available upon request.

coincide with the outcomes sustained by D-O equilibria of my new game, that have the additional property that principals report truthfully their types to the DRMs that they offer. In fact, these "truthful" equilibria cover the entire set of outcomes sustained by D-O equilibria of the new game. Of course, the equivalence result relies on an appropriate notion of enlargement that makes the new game and the underlying game contractually equivalent.

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 Delegation Principle. Section 5 contains the two main characterization results of the paper. The last section concludes.

2 The Environment

In this Section I describe the class of common agency games that I study in this paper. As I mentioned in the Introduction, the main difference with respect to common agency games already examined in the literature consists in introducing private information on the principals' side. As a standard convention, I shall use feminine pronouns for principals and masculine ones for the agent.

Players: There are $N > 1$ principals, denoted by P_i for $i \in N$, and a common agent, denoted by A . Hereafter, N stands both for the number and for the set of principals, and $i = 0$ stands for the agent. As usual, each P_i has full bargaining power when she interacts with A . Each P_i offers A a mechanism, which is a procedure to select an outcome for their bilateral relationship; principals offer their mechanisms non-cooperatively.

Information: Before each P_i offers her mechanism, *each* player $i \in N \cup \{0\}$ is privately informed about some relevant aspect of their interaction. This information is represented by the type $\theta_i \in \Theta_i$, where Θ_i is a (nonempty) finite set; as usual, $\Theta = \times_{i=0}^N \Theta_i$ and $\Theta_{-i} = \times_{j \neq i} \Theta_j$ with corresponding elements θ and θ_{-i} . The players' types are drawn from Θ according to the joint probability distribution r , which is commonly known by all players. Given that player i 's type is θ_i , the function $r_i(\cdot | \theta_i)$ represents the *conditional* distribution that i assigns to the opponents' types in Θ_{-i} .

Allocations and Mechanisms: By interacting with A , each P_i induces an allocation from the (nonempty) finite set Y_i : for example, y_i can be a price-quantity pair for a transaction between a buyer and a seller, or a piece of regulation together with a monetary contribution by a lobbyist. P_i and A choose an allocation in Y_i by communicating through a certain language. I shall model such a language with a finite² set of messages for each P_i and for A , denoted by \mathcal{P}_i and \mathcal{A}_i . I assume that $\Theta_i \subset \mathcal{P}_i$, so that P_i can at least communicate what she privately knows, and that $\times_{j=0}^N \Theta_j \subset \mathcal{A}_i$ for every $i \in N$, so that the agent can at least communicate to each P_i what he privately knows as well as what he may learn about the environment by contracting with all principals.

A mechanism m_i specifies how the communication between P_i and A leads to a certain allocation. As usual, an indirect mechanism is a map $m_i : \mathcal{A}_i \times \mathcal{P}_i \rightarrow \mathcal{D}_i = \Delta(Y_i)$, that associates to each pair (a_i, p_i) a *decision* $d_i \in \mathcal{D}_i$ which determines (possibly randomly) one

²The assumption that the sets of messages are finite obviously entails a loss of generality, but it guarantees that continuation equilibria exists after any principals' offers.

allocation in Y_i . I shall denote by M_i the space of all indirect mechanisms of P_i ; M_i represents all the possible communication devices that P_i can offer A to select a certain decision with his collaboration.³ By changing the sets of messages and the sets of allocations, we obtain different spaces of mechanisms, and hence different contracting environments. For future reference, denote the range of any mechanism m_i by

$$\text{Im}(m_i) \equiv \{d_i \in \mathcal{D}_i : \exists (a_i, p_i), m_i(a_i, p_i) = d_i\},$$

and let $m_i(y_i; a_i, p_i) = d_i(y_i)$ be the probability that m_i assigns to y_i given the messages (a_i, p_i) . Of course, A can decide to accept any subset of the mechanisms that the principals offer him. Instead of explicitly modeling this decision, I assume that each message space \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 $m_i(\hat{a}_i, \cdot)$ may still depend on P_i 's messages, because P_i may still affect the final decision in \mathcal{D}_i even if the agent refuses to contract with her.⁴

As it is usually assumed in common agency models, principals can fully commit to their mechanisms. Furthermore, principals can't directly communicate among themselves, and for every P_i , neither m_i can directly depend on m_{-i} nor y_i on y_{-i} . This latter assumption is an important difference between my environment and those in Peters (2010), Peters and Szentes (2011), and Celik and Peters (2011).

Timing: The timing of the game is as follows:

- $t = 0$: each player i observes θ_i and forms the belief $r_i(\cdot | \theta_i)$;
- $t = 1$: all principals simultaneously and non-cooperatively offer A the mechanisms m_1, \dots, m_N : whereas A observes all mechanisms, each P_i does not observe m_{-i} (private contracting);
- $t = 2$: both A and each P_i simultaneously,⁵ non-cooperatively, and privately send their messages to the respective mechanisms, which in turn independently deliver an allocation y . The players enjoy their payoffs and the game is over.

Strategies:

- *Agent*: a (behavior) strategy α of A is a (measurable) function $\alpha : \Theta_0 \times M \rightarrow \Delta(\mathcal{A})$ which selects a lottery over message profiles a as a function of the type θ_0 and the offered mechanisms m . Hereafter, I shall refer to α as the "communication strategy" of A .

³Whenever required, Σ_{M_i} denotes a σ -algebra of measurable subsets of M_i . Also, 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 $d \equiv (d_1, \dots, d_N)$ respectively. I shall restrict the use the symbol $|$ to denote conditional probabilities.

⁴Here I follow Maskin and Tirole (1992) in allowing for the case of no pre-existing mechanism, in which case m_i can be a function at most of P_i 's messages. I won't consider the case of renegotiation of a mechanism, where m_i is a "full-fledged" mechanism.

⁵I assume that P_i can't observe m_{-i} at any stage in the game. However, one could consider the possibility that the principals choose their messages after observing the agent's messages. This may be more interesting or realistic, but would make the analysis more complicated.

- *Principal i*: a (behavior) strategy σ_i of P_i is a pair of (measurable) functions (μ_i, π_i) . The function $\mu_i : \Theta_i \rightarrow \Delta(M_i)$ represents a random choice of mechanisms as a function of the type θ_i ; the function $\pi_i : \Theta_i \times M_i \rightarrow \Delta(\mathcal{P}_i)$ selects a lottery over messages p_i as a function of the type θ_i and the offered mechanism m_i . Hereafter, I shall refer to μ_i as the *offer strategy* and to π_i as the *communication strategy* of P_i .

Updating: Upon observing a realization m of the offer strategy profile μ , A updates his belief about the principals' types. The combination of μ and the exogenous distribution r over Θ defines a joint probability measure over $M \times \Theta$. Using this joint distribution, we can define the agent's updated beliefs at the beginning of $t = 2$ as the (regular) conditional probability $\rho(\cdot | \theta_0, m)$ on Θ_{-0} , induced by μ and r .

Common Agency Game: I shall let G^M denote the common agency game in which P_1, \dots, P_N compete in offering A mechanisms, each from the respective set M_i . The game G^M is an incomplete information game with imperfectly observable actions.

Preferences and Payoffs: Each player $i \in N \cup \{0\}$ is an expected utility maximizer whose von Neumann-Morgenstern utility function is $u_i : Y \times \Theta \rightarrow \mathbb{R}$. Using u_i , let $U_0(\alpha, \pi; \theta_0, m, \rho)$ be A 's payoff 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 ρ . Similarly, let $U_i(\pi, \alpha; \theta_i, m_i, \mu_{-i})$ be P_i 's payoff from the communication strategy profile (α, π) , given that her type is θ_i , she offers m_i , and her opponents play μ_{-i} . Finally, let $\bar{U}_i(\mu; \theta_i, \pi, \alpha)$ be P_i 's initial payoff from the profile $\sigma = (\mu, \pi, \alpha)$, given that her type is θ_i .

Equilibrium Concept: In this paper I focus on (weak) PBE of the game G^M :

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

1. For every $\theta_0 \in \Theta_0$ and $m \in M$,

$$U_0(\alpha^*(\theta_0, m), \pi^*; \theta_0, m, \rho^*) \geq U_0(a, \pi^*; \theta_0, m, \rho^*)$$

for every $a \in \mathcal{A}$;

2. For every P_i , $\theta_i \in \Theta_i$, and $m_i \in M_i$,

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

for every $p_i \in \mathcal{P}_i$;

4. For every P_i and $\theta_i \in \Theta_i$,

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

for every $m_i \in M_i$;

5. On path, ρ^* is a (regular) conditional probability induced by μ^* and r .

I will denote by $\mathbb{E}(G^M)$ the set of equilibria of G^M . As in Peters (2001), Martimort and Stole (2002), Pavan and Calzolari (2009), Celik and Peters (2011), and Peters (2011), the existence of equilibria of G^M is assumed here. However, note that the finiteness of the sets of messages

ensure that after any profile of offered mechanisms m , there exists a profile of communication strategies (π, α) , possibly mixed, that satisfy conditions 1 and 2 in Definition 1. Hereafter, I shall refer to such a profile (π, α) as a "continuation equilibrium."

Social Choice Functions: Let $\mathcal{F} := \{f : \Theta \rightarrow \Delta(Y)\}$ be the set of all social choice functions (henceforth SCF). Furthermore, let f_{σ^*} be the SCF sustained by $(\sigma^*, \rho^*) \in \mathbb{E}(G^M)$, and let $\mathcal{F}(G^M)$ be the set of all SCFs that are sustained by elements of $\mathbb{E}(G^M)$.

3 Failure of the Inscrutability Principle

A fundamental result in the literature on mechanism design with *one* informed principal is the so called "Inscrutability Principle". Myerson discovered this principle in his 1983 seminal paper, and since then the literature has relied on it (see, e.g., Maskin and Tirole). 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) The reason is as follows: in models with only one principal and one agent, if the principal can sustain a sequentially rational allocation by revealing up front some of her private information through the mechanisms that she offers—and if her set of feasible mechanisms includes all direct mechanisms—then she can also achieve the same allocation by appropriately designing a sole, completely uninformative, direct mechanism. This new mechanism precludes changing agent's prior belief *before* communication occurs. Nonetheless, it replicates the original mapping of the principal's and the agent's types into allocations, and makes both players willing to truthfully reveal their private information.

The Inscrutability Principle is extremely useful because it allows us to completely disregard the potential role of the principal's mechanisms as signalling devices. Therefore, it is very important for the present paper to understand whether such a Principle applies also to environments with more than one informed principal. If the answer were affirmative, it would dramatically simplify our analysis of any game G^M , because we could safely assume that each principal follows the same offer strategy, independently of her type. Unfortunately, the Inscrutability Principle fails when the agent contracts simultaneously with many principals. The next example proves this claim.

Example 1 *There are only two principals, P_1 and P_2 . Let $\Theta_1 = \{\bar{\theta}_1, \underline{\theta}_1\}$, $|\Theta_1| = |\Theta_2| = 1$, and assume that $r(\bar{\theta}_1) = 1/2$. For each principal there are two possible allocations, i.e., $Y_1 = \{\bar{y}_1, \underline{y}_1\}$ and $Y_2 = \{\bar{y}_2, \underline{y}_2\}$. The sets of messages are $\mathcal{P}_1 = \{\underline{\theta}_1, \bar{\theta}_1\}$, $\mathcal{A}_1 = \{\underline{\theta}_1, \bar{\theta}_1\}$, $\mathcal{A}_2 = \{\underline{\theta}_1, \bar{\theta}_1\}$, and for simplicity $\mathcal{P}_2 = \{p\}$. Table 1 reports the players' payoffs: in each cell the first entry is P_1 's payoff, the second is P_2 's payoff, and the third is A 's payoff.*

$\bar{\theta}_1$			$\underline{\theta}_1$		
y_1/y_2	\bar{y}_2	\underline{y}_2	y_1/y_2	\bar{y}_2	\underline{y}_2
\bar{y}_1	2, 3, 3	1, -1, 0	\bar{y}_1	-1, 3, -1	1, -1, 0
\underline{y}_1	1, 0, 0	-1, 2, 1	\underline{y}_1	1, 0, 0	2, 2, 3

Table 1: Example 1 payoffs

Consider the following equilibrium in which P_1 reveals her type up front. P_1 's offer strategy is such that $\mu_1(\bar{\theta}_1) = \bar{m}_1$ and $\mu_1(\underline{\theta}_1) = \underline{m}_1$, where

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

P_1 's communication strategy is

$$\pi_1(\theta_1, \bar{m}_1) = \begin{cases} \bar{\theta}_1 & \text{if } \theta_1 = \bar{\theta}_1 \\ \underline{\theta}_1 & \text{if } \theta_1 = \underline{\theta}_1 \end{cases} \quad \text{and} \quad \pi_1(\theta_1, \underline{m}_1) = \begin{cases} \bar{\theta}_1 & \text{if } \theta_1 = \bar{\theta}_1 \\ \underline{\theta}_1 & \text{if } \theta_1 = \underline{\theta}_1 \end{cases} .$$

For $m_1 \notin \{\bar{m}_1, \underline{m}_1\}$, extend π_1 in any sequentially rational way. P_2 offers a mechanism m_2 such that $m_2(\bar{\theta}_1) = \bar{y}_2$ and $m_2(\underline{\theta}_1) = \underline{y}_2$. As far as A is concerned, when he is offered m_2 and either \bar{m}_1 or \underline{m}_1 , let

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

For other m_2 and if $m_1 \notin \{\bar{m}_1, \underline{m}_1\}$, extend α in any sequentially rational way: in particular, holding σ_1 fixed, if P_2 "offers" only \bar{y}_2 , A sends message $\bar{\theta}_1$ to \bar{m}_1 and $\underline{\theta}_1$ to \underline{m}_1 , and if P_2 "offers" only \underline{y}_2 , A sends message $\underline{\theta}_1$ to \bar{m}_1 and $\bar{\theta}_1$ to \underline{m}_1 . Note that according to (1), A selects different allocations through m_2 depending on whether P_1 offers \bar{m}_1 or \underline{m}_1 .

It is easy to see that the strategy σ_1 is sequentially rational given α and μ_2 , and P_1 is getting her maximal payoff of 2. P_2 is also maximizing her payoff given μ_1 and α . Finally, the updated belief $\rho(\bar{\theta}_1 | \bar{m}_1) = \rho(\underline{\theta}_1 | \underline{m}_1) = 1$, α satisfies sequential rationality. The equilibrium SCF is then

$$f_{\sigma_1, \sigma_2, \alpha}(\theta_1) = \begin{cases} (\bar{y}_1, \bar{y}_2) & \text{if } \theta_1 = \bar{\theta}_1 \\ (\underline{y}_1, \underline{y}_2) & \text{if } \theta_1 = \underline{\theta}_1 \end{cases}$$

However, there is no mechanism m_1 such that if $\mu_1(\bar{\theta}_1) = \mu_1(\underline{\theta}_1) = m_1$, the resulting equilibrium, if any, sustains the same SCF. Even if P_2 allows A to choose between \bar{y}_2 and \underline{y}_2 , since A doesn't perfectly infer P_1 's type, he can't correctly match the two principals' allocations: with positive probability, he will end up choosing either \underline{y}_2 when $\theta_1 = \bar{\theta}_1$, or \bar{y}_2 when $\theta_1 = \underline{\theta}_1$.

The example highlights why the Inscrutability Principle fails when there is more than one principal. If P_1 allows A to learn something about θ_1 through the mechanisms she offers, then A optimally chooses different contracts with P_2 . More generally, if in equilibrium the principals

reveal part of their information up front, then A may combine these pieces of information to sustain a SCF that exhibits some correlation between θ and the final allocations. However, A can't achieve the same degree of coordination when, say, P_1 offer strategy is uninformative: A may learn θ_1 after observing the consequence of his message to m_1 , but such a discovery happens too late, since A has to communicate with P_2 before seeing that consequence.

Moreover, it is possible to construct examples in which, holding fixed the strategies of P_{-i} , each P_i can sustain the *same* allocation, with regard to her bilateral relationship with A , with an informative as well as an uninformative offer strategies. However, A 's behavior with the other principals may depend on which strategy P_i uses. Therefore, an equilibrium may exist if P_i uses an informative offer strategy, but it may collapse if P_i uses an uninformative strategy. This observation highlights a second difference between the games in this paper and the principal-agent games originally studied by Myerson: the non-cooperative nature of the competition among the principals.

The bottom line is then as follows. In contrast to what happens in models with one informed principal and one agent, when there are multiple informed principals contracting with the same agent, we can no longer disregard the signalling component that the principals' offers may assume in equilibrium.

4 Failure of the Delegation Principle

I shall now move my attention towards the literature on common agency with uniformed principals. An important result there is the so called Delegation Principle, or Menu Theorem (see Peters (2001) and Martimort-Stole (2002)). The purpose of this Section is to show that the Delegation Principle does not apply when principals are informed. A difference with standard common agency games is that in the present model the principals communicate with their mechanism, and thus affect the final allocation at a second stage. However, this important difference is not the key reason why the Delegation Principle doesn't apply here. In fact, I will show that even if we don't allow the principals to send messages to their mechanisms, the Delegation Principle still fails.

To prove my claims, I first need to introduce the notion of menu, and recall the statement of the Delegation Principle, or Menu Theorem. Consistently with Peters (2001), Martimort and Stole (2002), and Pavan and Calzolari (2009), I define a menu as follows:

Definition 2 (Simple Menus) For $i \in N$, a *menu* is a mechanism $l_i : \mathcal{A}_i \rightarrow \mathcal{D}_i$ with the following properties:

1. $\mathcal{A}_i \in 2^{\mathcal{D}_i} \setminus \emptyset$;
2. l_i depends only on A 's messages;
3. l_i is the identity function, i.e., $l_i(d_i) = d_i$.

Properties (1) and (3) together correspond to the standard definition of menu— l stands for "list" of objects. Property (2) instead says that P_i can't participate in the selection of y_i by sending messages to the mechanism. Of course, this doesn't mean that P_i can't influence at all the outcome of her interaction with A ; it just precludes one way by which she can do so. As

usual, L_i is the collection of all menus that P_i has available, and G^L is the "menu game" in which each P_i offers menus from L_i .

The important contribution of Peters (2001) and Martimort and Stole (2002) is that they identify a game G that is much simpler than any general game with uninformed principals G' , and has the property that any equilibrium SCF of G' is also an equilibrium SCF of G and vice versa ($\mathcal{F}(G') = \mathcal{F}(G)$). This equivalence result requires, of course, that the underlying game G' and the new game G be comparable from the outset. The notion of enlargement is the usual way to capture this requirement: in the present environment, I shall say that G^M is an enlargement of G^L ($G^M \succcurlyeq G^L$) if for every $i \in N$ there exists an embedding⁶ $\omega_i : L_i \rightarrow M_i$ (cfr. Pavan-Calzolari (2009)). A formally equivalent, although false, statement of the Delegation Principle (Menu Theorem) for the present environment is (cfr. Peters (2001) and Martimort and Stole (2002)):

Delegation Principle (Menu Theorem): Let G^M be an enlargement of G^L , and f be a social choice function. There exists an equilibrium (σ^*, ρ^*) of G^M such that $f_{\sigma^*} = f$ if and only if there exists an equilibrium (σ^{**}, ρ^{**}) of G^L such that $f_{\sigma^{**}} = f$.

I shall provide two counterexamples to this statement for environments with informed principals. In the first example, in the second example instead only the agent can play the principals' mechanisms. These examples also help understand the fundamental novelties implied by having informed principals.

Example 2 *As in Example 1, there are only two principals, P_1 and P_2 , $\Theta_1 = \{\underline{\theta}_1, \bar{\theta}_1\}$, $|\Theta_0| = |\Theta_2| = 1$, and $r(\bar{\theta}_1) = 1/2$. Furthermore, $Y_1 = \{y_1, y'_1, y''_1\}$ and $Y_2 = \{y_2, y'_2, y''_2\}$. The set of menus for P_i is L_i and, for the sake of simplicity, let $\mathcal{D}_i = Y_i$. Consider an enlargement G^M where*

$$M_1 = \{m_1 : \mathcal{A}_1 \times \mathcal{P}_1 \rightarrow Y_1 \mid \mathcal{A}_1 = \{u, c, d\}, \mathcal{P}_1 = \Theta_1\}^7$$

and $M_2 = L_2$.

As far as the players' payoffs are concerned, consider the two possibilities reported in Table 2 and 3. In Table 2, both P_2 's and A 's payoffs are independent of θ_1 ("full" private values). Instead, in Table 3 A 's payoffs depend on θ_1 (common values); note that, conditional θ_1 , A 's preference over P_2 's contracts is independent of y_1 .

$\bar{\theta}_1$				$\underline{\theta}_1$			
y_1/y_2	y_2	y'_2	y''_2	y_1/y_2	y_2	y'_2	y''_2
y_1	-1, 2, 3	7, 1, 2	-2, -2, 0	y_1	2, 2, 3	3, 1, 2	1, -2, 0
y'_1	0, 1, 4	5, 1, 5	1, 1, 4	y'_1	-1, 1, 4	5, 1, 5	0, 1, 4
y''_1	1, -2, 0	4, 1, 2	3, 2, 3	y''_1	0, -2, 0	6, 1, 2	-1, 2, 3

Table 2: "full" private values

⁶For example, $\omega_i : L_i \rightarrow M_i$ can be an injective function such that if $m_i = \omega_i(l_i)$, then for every $p_i, p'_i \in \mathcal{P}_i$, $m_i(\cdot, p_i) = m_i(\cdot, p'_i)$ and $\text{Im}(m_i(p_i)) = \text{Im}(l_i)$.

$\bar{\theta}_1$

y_1/y_2	y_2	y'_2	y''_2
y_1	-1, 2, 3	7, 1, 2	-2, -2, 0
y'_1	0, 1, 4	5, 1, 3	1, 1, 1
y''_1	1, -2, 3	4, 1, 2	3, 2, 0

 $\underline{\theta}_1$

y_1/y_2	y_2	y'_2	y''_2
y_1	2, 2, 0	3, 1, 2	1, -2, 3
y'_1	-1, 1, 1	5, 1, 3	0, 1, 4
y''_1	0, -2, 0	6, 1, 2	-1, 2, 3

Table 3: common values

For both payoffs' specifications, consider the following profile (σ, ρ) . P_2 offers the menu l_2 with $\mathcal{A}_{l_2} = Y_2$. P_1 's strategy satisfies $\mu_1(\bar{\theta}_1) = \mu_1(\underline{\theta}_1) = m_1$, where

$$m_1(a_1, \bar{\theta}_1) = \begin{cases} y''_1 & \text{if } a_1 = u \\ y_1 & \text{if } a_1 = c \\ y'_1 & \text{if } a_1 = d \end{cases} \quad \text{and} \quad m_1(a_1, \underline{\theta}_1) = \begin{cases} y_1 & \text{if } a_1 = u \\ y''_1 & \text{if } a_1 = c \\ y_1 & \text{if } a_1 = d \end{cases} .$$

Moreover, $\pi_1(\theta_1, m_1) = \theta_1$ for every θ_1 . Finally, A 's strategy α satisfies:

$$\alpha(m_1, l_2) = \begin{cases} (c, y'_2) & \text{if } y'_2 \in \mathcal{A}_{l_2} \\ (u, y_2) & \text{if } \mathcal{A}_{l_2} = \{y_2, y''_2\} \\ (u, y_2) & \text{if } \mathcal{A}_{l_2} = \{y_2\} \\ (d, y''_2) & \text{if } \mathcal{A}_{l_2} = \{y''_2\} \end{cases} . \quad (2)$$

Complete α (and π_1) when $m'_1 \neq m_1$ or $l'_2 \neq l_2$ in any sequentially rational way.

I claim that $\sigma = (\mu_1, \mu_2, \pi_1, \alpha)$, together with $\rho(\theta_1 | m_1, l_2) = r(\theta_1)$ for all m_1, l_2 , is an equilibrium of G^M . First of all, both types of P_1 are getting their maximal payoffs by playing π_1 against α . Furthermore, α is sequentially rational because choosing y'_2 is the unique best reply to (μ_1, π_1) , and if y'_2 is not available, it is optimal to select y_2 (or y''_2) and then send any message to m_1 . Therefore (μ_1, π_1) is optimal for P_1 . To see that P_2 has no profitable deviation, note that any menu l'_2 that "includes" y'_2 is equivalent to l_2 ; on the other hand, A 's strategy in (2) together with (μ_1, π_1) prevents P_2 from gaining if she deviates to l'_2 with $\mathcal{A}_{l'_2} \in 2^{\{y_2, y''_2\}} \setminus \emptyset$. The equilibrium SCF is then

$$f_\sigma(\theta_1) = \begin{cases} (y_1, y'_2) & \text{if } \theta_1 = \bar{\theta}_1 \\ (y''_1, y'_2) & \text{if } \theta_1 = \underline{\theta}_1 \end{cases} .$$

However, no equilibrium of G^L can sustain f_σ —this is true for both payoffs' specifications, but for different reasons that I shall explain shortly. Suppose on the contrary that such an equilibrium existed. Since f_σ deterministically specifies a different allocation between A and P_1 depending on P_1 's type, P_1 's offer strategy in G^L must satisfy $\mu_1(\bar{\theta}_1) \neq \mu_1(\underline{\theta}_1)$: otherwise, A would (randomly) induce the same allocation for $\bar{\theta}_1$ and $\underline{\theta}_1$. It follows that P_1 must reveal her type up front. In addition, since A strictly prefers the allocation y'_1 , P_1 should never allow A to choose such an allocation, i.e., $y'_1 \notin \mathcal{A}_{l_1}$. Therefore, for at least one of her types, P_1 must offer only the allocation that she likes best. Finally, given his degenerate beliefs, A selects y'_2 with P_2

if and only if $\mu_2 = l_2$ with $\mathcal{A}_{l_2} = \{y'_2\}$. But this offer strategy is not optimal for P_2 .⁸ if P_2 offers the menu l'_2 with $\mathcal{A}_{l'_2} = \{y_2, y'_2, y''_2\}$, A will then choose either y_2 or y''_2 , making P_2 strictly better off.

Example 2 helps understand the first limitation of restricting attention to G^L . When P_1 can offer only menus, she entirely transfers the decision power into the hands of the agent. Thus, P_1 has to rely only on his self-interest to achieve the decision that best matches her type. However, if P_1 constrains A 's choices to menus containing only the allocations that her type likes best, then P_1 may lead A to make undesirable decisions with P_{-1} for two reasons.⁹ In the "full" private values case, although A may update his beliefs about θ_1 upon observing l_1 , this updating has no effect on his preference and hence his decisions. Nonetheless, A 's preference over Y_2 *endogenously* depend on y_1 . Therefore, when P_1 offers A a subset of Y_1 , she reduces A 's uncertainty about their bilateral relationship, thereby inducing a y_2 that she dislikes. In the common value case, instead, A cares directly about the type of P_1 . Therefore, if P_1 reveals too much information about her type, then A refines his preference enough to discover that some other y'_2 is better than his current choice. Clearly, none of these issues arise when the principals can communicate with their mechanisms, but doing so requires more than simple menus of allocations.

In light of the previous remarks, it seems natural to wonder whether the principals' ability to communicate with their mechanisms is the only reason why the Delegation Principle fails. Since such a flexibility may be unappealing or unrealistic in some context (see, e.g., Martimort and Moreira (2010)), it is important to know what happens if, from the outset, all mechanisms can depend only on A 's messages. As in games with uninformed principals, in such an environment each P_i delegates the final decisions to the agent.

Definition 3 (Delegation) *The common agency game G^M is called a game with **delegation**, denoted G^D , if for every $i \in N$ and $m_i \in M_i$, $m_i : \mathcal{A}_i \rightarrow \mathcal{D}_i$.*

Note that in G^D the principals don't play any active role after offering their mechanisms, therefore there is no principals' communication strategy. The solution concept in Definition 1 naturally extends to G^D .

If we focus on games with delegation, we can better understand another important novelty caused by having informed principals: namely, the signaling content that each mechanism can assume in equilibrium. The next example illustrates this point and shows that, even in games with delegation, restricting attention to standard menus may not allow for full generality.

Example 3 *Again, there are only two principals, P_1 and P_2 , $\Theta_1 = \{\bar{\theta}_1, \underline{\theta}_1\}$, and $r(\bar{\theta}_1) = 1/2$, whereas P_2 and A are uninformed. P_1 's set of allocations is $Y_1 = \{\bar{y}_1, \underline{y}_1\}$ and P_2 's is $Y_2 = \{y_2, y'_2, y''_2\}$. Table 4 reports the players' payoffs for every θ_1 . For simplicity, P_1 and P_2 compete*

⁸Here I am using the assumption that if P_2 deviates and P_1 doesn't, A still uses Bayes rule to update his beliefs about θ_1 . See Fudenberg and Tirole (1991).

⁹These reasons resemble the effects of private and common values assumptions that Maskin and Tirole (1990-1992) discuss in their papers.

in offering A mechanisms of the form $m_i : \mathcal{A}_i \rightarrow Y_i$, where $\mathcal{A}_1 = \{\bar{\theta}_1, \underline{\theta}_1\}$ and $\mathcal{A}_1 = \{\bar{\theta}_1, \underline{\theta}_1, c\}$.

$\bar{\theta}_1$				$\underline{\theta}_1$			
y_1/y_2	y_2	y'_2	y''_2	y_1/y_2	y_2	y'_2	y''_2
\bar{y}_1	1, 1, 1	1, 3, 2	1, 1, 0	\bar{y}_1	0, 1, 0	0, 0, 3	0, 1, 0
\underline{y}_1	0, 1, 0	0, 0, 3	0, 1, 0	\underline{y}_1	1, 1, 0	1, 3, 2	1, 1, 1

Table 4: Example 3 payoffs

Consider the following equilibrium (σ, ρ) . P_2 offers m_2 such that $\text{Im}(m_2) = \{y_2, y'_2\}$. P_1 's strategy satisfies $\mu_1(\bar{\theta}_1) = \bar{m}_1$ and $\mu_1(\underline{\theta}_1) = \underline{m}_1$ with

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

Then, on path A 's beliefs ρ will assign probability one to $\bar{\theta}_1$ after \bar{m}_1 and probability one to $\underline{\theta}_1$ after \underline{m}_1 (as before, these beliefs remain unchanged after any unilateral deviation of P_2). Finally, A 's communication strategy α is as follows: if P_1 and P_2 offer (\bar{m}_1, m_2) , A "chooses" y_2 and sends the message $\bar{\theta}_1$; whereas if P_1 and P_2 offer (\underline{m}_1, m_2) , A "chooses" y'_2 and sends the message $\bar{\theta}_1$. Moreover, if P_2 deviates to a mechanism that allow A to choose y'_2 , then A will do so, and will send the message $\underline{\theta}_1$ both to \bar{m}_1 and to \underline{m}_1 .

It is immediate to see that the (partially)¹⁰ defined strategy α is sequentially rational for A , given ρ and μ_1 and μ_2 . Also μ_1 is a best reply to μ_2 and α . Similarly, P_2 has no profitable deviation: even if P_2 allowed A to choose y'_2 , A would choose the "wrong" contract with P_1 , causing P_2 's payoff to drop to zero. The equilibrium SCF is

$$f_\sigma(\theta_1) = \begin{cases} (\bar{y}_1, y_2) & \text{if } \theta_1 = \bar{\theta}_1 \\ (\underline{y}_1, y'_2) & \text{if } \theta_1 = \underline{\theta}_1 \end{cases}.$$

I claim that no equilibrium of a menu game G^L , with $G^D \succneq G^L$, can sustain the same SCF. On the one hand, if both types of P_1 offer the same menu containing both \bar{y}_1 and \underline{y}_1 , A will induce the same lottery over $\{\bar{y}_1, \underline{y}_1\}$, independently of θ_1 . On the other hand, if either type of P_1 offers the menu containing only her preferred allocation, then P_2 can't be deterred from offering y'_2 .

Example 3 suggests another limitation of G^L that makes this class of games inadequate to replicate all equilibrium SCFs of G^D (and a fortiori of G^M). From Section 3, we know that signaling is necessary to sustain certain equilibrium SCFs of G^D . However, the principals' ability to signal their type can be poorer in G^L than in the underlying game G^D , therefore some equilibrium SCFs of G^D may disappear in G^L . In G^L , to define what A can chose and possibly signal her type, each P_i can use only one instrument: the set of allocations that she

¹⁰ A 's beliefs and behavior after P_1 's deviations can be appropriately defined without affecting the message of the example.

offers A . Such a restriction can involve some trade-offs. In Example 3, if P_1 and P_2 compete in menus, to signal her type P_1 must exclude from one of her offers a "latent contract" that, in the equilibrium of G^D , deterred P_2 's deviations by leveraging on A 's self-interest. In contrast, in G^D such a trade-off doesn't arise because each P_i can frame her mechanisms in different ways to signal her type, while offering A the same choice set.

To summarize, the examples in this section and Section 3 have brought to light the main strategic components of common agency games with informed principals. Compared to games with uninformed principals, when principals are informed and there is delegation, the mechanisms play *two* key roles *in equilibrium*. The first one is to specify a set of allocations among which A is allowed to choose—this is the *only* role when principals are uninformed, as the Delegation Principle shows. The second one is to signal the principals' private information, a role that can clearly appear only when principals are informed, and is perhaps the key novel ingredient of the present model. Furthermore, in games without delegation, mechanisms play a third key role: namely, they allow each principal to further act on her private information *after* the agent has communicated with her opponents.

These observations suggest how to extend the intuition behind the Delegation Principle to identify a class of games that is simpler than G^D or G^M , but allows to replicate all equilibrium SCFs of G^D and G^M . Intuitively, this class should allow principals to offer appropriately defined choice sets as well as to send "cheap-talk" signals, so that each P_i has enough flexibility to reveal different pieces of information, while offering A the same choice set.

5 Deterministic-Offer Equilibria: A Simple Characterization

In this section I focus on equilibria of the underlying games G^D and G^M in which principals' offer strategies are deterministic functions of their types. This kind of equilibria represents the focus of almost the entire literature on common agency—to the best of my knowledge, there is no paper that actually constructs an equilibrium in which principals mix over mechanisms. From a pragmatic perspective, deterministic-offer equilibria are clearly more realistic and appealing. Furthermore, as we shall see, they are amenable to a simple and intuitive characterization.

I begin by formally defining deterministic-offer equilibria as follows.

Definition 4 (D-O equilibrium) *A profile $(\sigma, \rho) \in \mathbb{E}(G^M)$ (respectively $\mathbb{E}(G^D)$) is a **deterministic - offer equilibrium** if for every $i \in N$ and $\theta_i \in \Theta_i$, $\mu_i(\theta_i) = m_i$ for some $m_i \in M_i$.¹¹*

Note that Definition 4 doesn't require that the principals' or the agent's communication strategy be deterministic.

In D-O equilibria the kind of signaling that the principals achieve through their mechanisms is relatively simple, and it involves a straightforward inference problem for the agent. When P_i offer strategy is deterministic, with each mechanism she simply signals that her type is to

¹¹The notation $\mu_i(\theta_i) = m_i$ stands for $\mu_i(\theta_i) = \delta_{m_i}$ (δ_{m_i} is the Dirac measure on m_i). For simplicity, I shall use the same convention to describe pure strategies.

a certain subset of Θ_i —this subset can be a singleton or contain many types. Therefore, upon observing a profile m , A needs only to infer which types θ_{-0} are consistent with m and the equilibrium behavior; the relative likelihood of these θ_{-0} is then simply pin down by A 's prior belief. Thus, in contrast to an equilibrium with mixed offer strategies, A never has to worry about the fact that different profiles θ_{-0} may play the same m but with different probabilities.

Hence, to replicate at least the SCFs that result from D-O equilibria of any G^D and G^M , the "cheap-talk" signals of our new class of games should allow each P_i to control the flow of information towards A as described above. I shall first consider the games with delegation, G^D , and then the games without delegation, G^M .

5.1 Games with Delegation

In the case of games with delegation, I augment the standard menu game of Peters (2001) by endowing each principal with the set of signals $S_i = \Theta_i$.¹² I then show that the set of SCFs sustained by D-O equilibria of G^D coincides with the set of SCFs sustained by D-O equilibria of a game in which principals compete in offering menus and signal their types through S_i . Therefore, as long as we are interested in D-O equilibria of an underlying game G^D , a very simple extension of Peters' (2001) menu game is enough to replicate the entire set of possible SCFs.

I shall refer to my new menu game in which each P_i sends signals in S_i as the "signal-menu game," denoted G^{SL} . In G^{SL} , P_i 's action space is $S_i \times L_i$, where L_i is the usual set of menus of P_i . A signal-offer strategy of P_i is then a map μ_i from Θ_i to $\Delta(S_i \times L_i)$. On the other hand, A 's strategy is now a map that, for every profile of signals θ and menus l , specifies a lottery in $\Delta(\mathcal{A}_l)$.

As usual, to replicate all the equilibrium SCFs of G^D , G^{SL} must be comparable to G^D from the outset in terms of the contracting possibilities that the two games allow. To capture this requirement, I follow the literature again and use the notion of enlargement.

Definition 5 ($G^D \succcurlyeq G^{SL}$) *G^D is an enlargement of G^{SL} if and only if for every $i \in N$ we have that for every menu $l_i \in L_i$ there exists a mechanism $m_i \in M_i$ such that $\text{Im}(m_i) = \text{Im}(l_i)$ and vice versa.*

It is important to note that Definition 5 is essentially the same as the standard notion of enlargement (see Section 4) as it requires the same fundamental conditions for G^D and G^{SL} to be comparable. Hence, if we start from any G^D and we want to identify the game G^{SL} such that $G^D \succcurlyeq G^{SL}$, we have to perform the same exercise as if we were dealing with a standard game with uninformed principal and had to derive the contractually equivalent, standard menu game.

Following Definition 4, an equilibrium of G^{SL} is a D-O equilibrium if every P_i chooses deterministically what menu she offers and what signal she sends after observing her type θ_i : if for every $i \in N$ and every $\theta_i \in \Theta_i$, $\mu_i(\theta_i) = (\hat{\theta}_i, l_i)$ for some $(\hat{\theta}_i, l_i) \in S_i \times L_i$.

I am now ready to state my main characterization result for games with informed principals and delegation.

¹²It will become clear in a moment that the essential property of S_i is that $|S_i| = |\Theta_i|$.

Theorem 1 *Let G^D be an enlargement of G^{SL} and f be a social choice function. There exists a deterministic-offer equilibrium (σ^*, ρ^*) of G^D such that $f_{\sigma^*} = f$ if and only if there exists a deterministic-offer equilibrium (σ^{**}, ρ^{**}) of G^{SL} such that $f_{\sigma^{**}} = f$.*

The intuition behind Theorem 1 is as follows. In a D-O equilibrium of G^D , by playing m_i , P_i signals that her type is in some $T_i(m_i) \subset \Theta_i$. Now in G^{SL} , all types in $T_i(m_i)$ can equivalently contract with and signal to A by offering him the same bundle of decisions that A can choose through m_i , and by sending him a common signal, e.g., *one* type in $T_i(m_i)$. In G^D , P_i 's offer strategy can be at most fully revealing—i.e., $|T_i(m_i)| = 1$ —but at the same time all types of P_i could offer A the same bundle of decisions (recall Example 3). Again, in G^{SL} each P_i can equivalently interact with A by offering him such a bundle of decisions, and by simply telling him her type.

Consider now a D-O equilibrium of G^{SL} . There always exists an outcome-equivalent equilibrium in G^D for two reasons: first, by the definition of enlargement, each bundle of decisions that P_i can offer in G^{SL} can also be offered in G^D through some indirect mechanism. Furthermore, the communication language in G^D is sufficiently rich to design enough different indirect mechanisms that, while leading to the same bundle of decisions, allow P_i to achieve the same degree of signaling as in the equilibrium of G^{SL} .

5.2 Games without Delegation

In this section, I drop the restriction that principals can't communicate with their mechanism, but I still focus on D-O equilibria. My goal is again to identify a class of games, possibly simpler than G^M , that allows us to replicate all the SCFs sustained by D-O equilibria of G^M . As I argued in Section 4, each P_i 's ability to communicate with her mechanisms requires moving beyond the standard notion of menus. However, I only need to replace standard menus of decisions with menus of direct revelation mechanisms to achieve my goal. I proceed in two steps: first, I show that menus of direct revelation mechanisms preserve the same ability to contract that principals have in games without delegation; I then provide my characterization result.

To identify a tractable class of games that serves our goal, we first need to understand how indirect mechanisms work in G^M . I suggest thinking of a mechanism $m_i : \mathcal{A}_i \times \mathcal{P}_i \rightarrow \mathcal{D}_i$ as an array of objects as in Figure 1.




Figure 1: Graphical representation of m_i .

Each entry of this array is simply an element of \mathcal{D}_i , which P_i and A select by independently submitting a pair of messages. If we now consider the rows of the array, we see that m_i identifies a collection of maps from P_i 's messages to decisions in \mathcal{D}_i : for each m_i this collection is

$$H_i(m_i) = \{h_i(m_i) : \mathcal{P}_i \rightarrow \mathcal{D}_i \mid h_i(\cdot; m_i) = m_i(a_i, \cdot), a_i \in \mathcal{A}_i\}.$$

Also, we can think of A as choosing the row of the array by submitting his message. Thus, m_i also defines a map

$$v_i(\cdot; m_i) : \mathcal{A}_i \rightarrow H_i(m_i)$$

for A , where $v_i(a_i; m_i) = m_i(a_i, \cdot)$. In short, we can interpret any indirect mechanism m_i as working in two steps: first, it matches each message of A with a "partial" mechanism $h_i(m_i)$ for P_i ; second, it uses this partial mechanism to translate P_i 's messages into decisions in \mathcal{D}_i . Note that since A and P_i send their messages simultaneously, P_i must play all partial mechanisms in $H_i(m_i)$ in the same way. Finally, $H_i = \bigcup_{m_i \in M_i} H_i(m_i)$ represents the collection of all partial mechanisms that P_i can combine to create any indirect mechanism in M_i —graphically, H_i is the collection of all rows that P_i can stack to form an array as in Figure 1.

Due to the Revelation Principle, direct revelation mechanisms represent a class of particular interest in the mechanism-design literature. In fact, they play an important role also in the present environment. For each P_i , define the collection of direct revelation mechanisms (hereafter, DRMs) as

$$H_i^R = \{h_i^R : \Theta_i \rightarrow \Delta(Y_i)\}.$$

Note that although $\Theta_i \subset \mathcal{P}_i$, formally H_i^R is not a subset of H_i because each h_i is a function of the whole set \mathcal{P}_i . Nonetheless, given any communication strategy π_i , we can match any partial mechanism $h_i \in H_i$ to a DRM $h_i^R \in H_i^R$ in the usual way: for every $y_i \in Y_i$, let

$$h_i^R(y_i; \theta_i) = \sum_{\mathcal{P}_i} d_i(y_i) h_i(d_i; p_i, m_i) \pi_i(p_i; \theta_i, m_i).$$

Hereafter, to concisely represent this transformation of $h_i(m_i)$ into the corresponding h_i^R through π_i , I shall use the notation $h_i(m_i) \triangleright \pi_i(m_i)$. Given π_i , if we iteratively apply this transformation to all the elements of $H_i(m_i)$, we obtain the list $\lambda^R(m_i, \pi_i)$ of DRMs that A can induce by communicating to m_i : formally,

$$\lambda^R(m_i, \pi_i) = \{h_i^R \in H_i^R : \exists a_i \in \mathcal{A}_i, v_i^R(a_i; m_i, \pi_i) = h_i^R\},$$

where $v_i^R(a_i; m_i, \pi_i) = v_i(a_i; m_i) \triangleright \pi_i(m_i)$. It is easy to see that, although two indirect mechanisms m_i and m'_i may look completely different, there may exist a communication strategy π_i such that $\lambda^R(m_i, \pi_i) = \lambda^R(m'_i, \pi_i)$.

I can now introduce the notion of "list" of DRMs:

Definition 6 (List of DRMs) *A list of direct revelation mechanisms for principal P_i , denoted l_i^R , is a function $l_i^R : \mathcal{A}_{l_i^R} \rightarrow H_i^R$ with the following properties:*

1. $\mathcal{A}_{l_i^R} \in 2^{H_i^R} \setminus \emptyset$;
2. l_i^R is the identity function, i.e., $l_i^R(h_i^R) = h_i^R$ for all $h_i^R \in \mathcal{A}_{l_i^R}$.

Let L_i^R be the collection of lists of DRMs of P_i . Figure 2 represents graphically how any l_i^R works: intuitively, if we compare Figure 1 and Figure 2, we see that any l_i^R is still an array of decisions in \mathcal{D}_i , but now the rows of Figure 1 have become direct revelation mechanisms, which A chooses by simply naming one.




Figure 2: Graphical representation of l_i^R .

Since any pair (m_i, π_i) corresponds to a list of DRMs, the set L_i^R contains all contracting opportunities that P_i can rely on in G^M . As I highlighted in Sections 3 and 4, however, we can't disregard that P_i 's indirect mechanisms are contracting devices as well as signaling devices. Therefore, although focussing on L_i^R does not restrict P_i 's ability to contract, it may still limit her ability to signal relative to M_i . Intuitively, in an equilibrium of G^M P_i may reveal her two types by playing two different indirect mechanisms that π_i turns into the same list l_i^R . Clearly, if constrained to L_i^R , P_i would face a trade-off between offering the same l_i^R and signaling what she knows.

Again, to replicate all SCFs sustained by D-O equilibria of G^M , I enrich each P_i 's space L_i^R to rule out any trade-off between contracting and signalling. In fact, there is a simple and intuitive way to do so without introducing any dedicated set of signals. The idea is to let each P_i offer lists of *restricted* DRMs (hereafter, RDRMs) that commit her to report only types in a subset of Θ_i . More formally, for every $T_i \in 2^{\Theta_i} \setminus \emptyset$ let

$$H_i^{RR}(T_i) = \{h_i^{RR} : T_i \rightarrow \Delta(Y_i)\}.$$

If we now rewrite Definition 6 using $H_i^{RR}(T_i)$ in place of H_i^R , we obtain all lists of RDRMs that commit P_i to report types in T_i , denoted $L_i^{RR}(T_i)$. And if we repeat the exercise for every $T_i \in 2^{\Theta_i} \setminus \emptyset$, we get the collection L_i^{RR} of all lists of RDRMs.

Figure 2 helps understand the intuition behind offering lists of RDRMs. Graphically, an element of $L_i^{RR}(T_i)$ is simply an array of DRMs as in Figure 2 from which we delete the columns corresponding to all types $\theta_i \notin T_i$. Now recall that in a D-O equilibrium of G^M , when P_i offers the mechanism m_i , at most she signals that her type is in some subset T_i of Θ_i . Then P_i can signal the same amount of information by offering a list of RDRMs that commit her to report only the types in T_i . In other words, suppose again that in G^M θ_i and θ'_i offer different mechanisms m_i and m'_i , but π_i implies that m_i and m'_i boil down to the same l_i^R . Then, within L_i^{RR} , θ_i and θ'_i can still differentiate themselves by offering $l_i^{RR}(T_i)$ and $l_i^{RR}(T'_i)$ such that $\theta_i \notin T'_i$, and vice versa.

In light of the previous considerations, I suggest to consider a game, denoted G^{LRR} , in which principals compete in lists of RDRMs. In such a game each P_i offers an element of L_i^{RR} simultaneously with P_{-i} . After A has observed all the offers, he selects a RDRMs from each list, and at the same time P_i submits a report $\hat{\theta}_i$ from the set to which she committed. More formally, in G^{LRR} each P_i 's strategy is a pair of maps (μ_i, π_i) such that $\mu_i : \Theta_i \rightarrow \Delta(L_i^{RR})$ and π_i assigns to each (θ_i, l_i^{RR}) a lottery over the resulting $T_i \subset \Theta_i$; and, A 's strategy α assigns to each profile $l^{RR} \in L^{RR}$ a lottery over $\mathcal{A}_{l^{RR}}$.

As I did for the games with delegation G^D , I define when G^M and G^{LRR} are comparable using an appropriate notion of enlargement.

Definition 7 ($G^M \succcurlyeq G^{LRR}$) *A game G^M is an enlargement of G^{LRR} if and only if for every P_i two conditions hold:*

1. *For every $l_i^{RR} \in L_i^{RR}(\Theta_i)$, there is a $m_i \in M_i$ and a communication strategy π_i such that $l_i^{RR} = \lambda^R(m_i, \pi_i)$;*
2. *For every $m_i \in M_i$ and π_i there is an $l_i^{RR} \in L_i^{RR}(\Theta_i)$ such that $\lambda^R(m_i, \pi_i) = l_i^{RR}$.*

Note that Definition 7 imposes conditions only on lists of DRMs, rather than on all lists of RDRMs.

Since in G^{LRR} each P_i ends up playing a RDRM, it seems natural to consider equilibria in which, at least on path, P_i reports truthfully her θ_i to the mechanisms she offers.

Definition 8 (Truthful Equilibrium) *An equilibrium (σ, ρ) of G^{LRR} is a (**principal**) **truthful equilibrium** if, for every principal P_i , every $\theta_i \in \Theta_i$, and every $l_i^{RR} \in \text{Supp}[\mu_i(\theta_i)]$, we have $\pi_i(\theta_i, l_i^{RR}) = \theta_i$.*

The next lemma justifies focussing on truthful equilibria of G^{LRR} .

Lemma 1 *If there exists a deterministic-offer equilibrium (σ^*, ρ^*) of G^{LRR} that sustains the social choice function f_{σ^*} , then there exists a deterministic-offer, truthful, equilibrium (σ^{**}, ρ^{**}) of G^{LRR} such that $f_{\sigma^{**}} = f_{\sigma^*}$.*

Truthful equilibria in which principals play deterministic-offer strategies represent a simple and intuitive class of equilibria of G^{LRR} . Moreover, as the next theorem shows, the D-O truthful equilibria of G^{LRR} exactly replicate all SCFs sustained by D-O equilibria of G^M .

Theorem 2 *Let G^M be an enlargement of G^{LRR} and f be a social choice function. There exists a deterministic-offer equilibrium (σ^*, ρ^*) of G^M such that $f_{\sigma^*} = f$ if and only if there exists a deterministic-offer, truthful, equilibrium (σ^{**}, ρ^{**}) of G^{LRR} such that $f_{\sigma^{**}} = f$.*

Hence, we can interpret a game without delegation G^M as being strategically equivalent to a game in which principals first commit to menus of restricted DRMs, the agent then observes all menus and picks one DRM for each principal, which she plays by truthfully reporting her type.

6 Conclusions

I have studied the strategic nature of common agency games with informed principals, both in the case in which the principals delegate the final decisions to their agent, and in the case in which they retain some decision power after they offer their mechanisms.

I have shown that, contrary to what happens in mechanism design models with one informed principal, Myerson's (1983) Inscrutability Principle fails in games with many informed principals. This result is important because it shows that in such games signaling through mechanisms is an essential strategic component. Preserving the principals' ability to signal is therefore necessary to characterize all the equilibrium outcomes.

I have also shown that, contrary to what happens in common agency games with uninformed principals, Martimort and Stole's (2002) Delegation Principle—equivalently, Peters' (2001) Menu Theorem—fails when principals have private information. This conclusion implies that simple menus of alternatives don't capture the entire strategic nature of games with informed principals. Since signaling through mechanisms is essential, how a principal frames a set of alternatives matters too. Furthermore, simple menus deprive the principals of their ex-post decision power if the underlying game is without delegation.

Based on these considerations, I have provided two results that characterize the equilibrium outcomes of games with informed principals. I have focussed on the equilibria that are typically of interest in common agency games, namely those in which principals offer their mechanisms following pure strategies. For games with delegation, we can replicate such equilibrium outcomes using a standard menu game with the only novelty that each principal is endowed with a set of cheap-talk signals equal to her set of types. For games without delegation, we can replicate the equilibrium outcomes using a new game in which principals offer menus of restricted direct-revelation mechanisms, which they then play truthfully in equilibrium.

Although my results clarify the strategic essence of games with informed principals and dramatically reduce their inherent complexity, they represent only a first step in providing a complete and practical piece of machinery to characterize all equilibrium outcomes. A natural completion of the present analysis would be to extend the results to equilibria in which principals randomly select their mechanisms. Most importantly, however, future research should try to address the lack of structure that handicaps menus—Pavan and Calzolari (2010) constitutes a starting point of this line of research. Solving this issue is likely to improve the usefulness of common agency models with informed principals, especially for applied research.

7 Appendix

7.1 Proof of Theorem 1

Part 1: (\Rightarrow) Let (σ^*, ρ^*) be a D-O equilibrium of G^D . I shall construct a D-O equilibrium (σ^{**}, ρ^{**}) of G^{SL} , and then show that $f_{\sigma^*} = f_{\sigma^{**}}$.

Given the offer strategy μ_i^* of P_i , let $\mu_i^*(\Theta_i)$ be the set of mechanisms that are offered by some θ_i in the equilibrium σ^* . 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. For every $T_i(\mu_i^*(\theta_i))$, let $\vartheta_i(\mu_i^*(\theta_i))$ be the type $\theta_i \in T_i(\mu_i^*(\theta_i))$ with the lowest index. Since $G^D \succcurlyeq G^{SL}$, for every $m_i \in M_i$ let $l_i(m_i)$ be the menu l_i such that $\mathcal{A}_{l_i} = \text{Im}(m_i)$. So, for $\theta_i \in \Theta_i$ let μ_i^{**} be defined as follows

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

The same construction applies to every P_i .

Now consider the agent's strategy α^{**} . Suppose the profile (θ_{-0}, l) is on path. Let m^{*l} be the profile of mechanisms such that $m_i^* = \mu_i^*(\theta_i)$ for every $i \in N$. And for every $\theta_0 \in \Theta_0$ and $d \in \mathcal{A}_l$, let

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

Suppose instead that (θ_{-0}, \hat{l}) is not on path. Then, if $\hat{l}_i = \hat{l}_i(\mu_i^*(\theta_i))$ let $\hat{m}_i = \mu_i^*(\theta_i)$, otherwise let \hat{m}_i be any mechanism in M_i such that $\text{Im}(\hat{m}_i) = \mathcal{A}_{\hat{l}_i}$. Given the profile \hat{m} , for every θ_0 and $d \in \mathcal{A}_{\hat{l}}$, let

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

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 m^{*l} .

Now consider the agent's updated belief ρ^{**} . Suppose the profile $(\hat{\theta}_{-0}, l)$ is on path. Because $\mu^{**}(\hat{\theta}_{-0}, l; \theta_{-0}) = \prod_{i \in N} \mu_i^{**}(\hat{\theta}_i, l_i; \theta_i)$ and μ^{**} have finite support, we have for every θ_{-0}

$$\begin{aligned} \rho^{**}(\theta_{-0} \mid \theta_0, \hat{\theta}_{-0}, l) &= \frac{\Pr_{\mu^{**}}(\hat{\theta}_{-0}, l \mid \theta_{-0}) r_0(\theta_{-0} \mid \theta_0)}{\sum_{\theta'_{-0}} \Pr_{\mu^{**}}(\hat{\theta}_{-0}, l \mid \theta'_{-0}) r_0(\theta'_{-0} \mid \theta_0)} \\ &= \frac{\prod_{i \in N} \delta_{(\vartheta_i(\mu_i^*(\theta_i)), l_i)} r_0(\theta_{-0} \mid \theta_0)}{\sum_{\theta'_{-0}} \prod_{i \in N} \delta_{(\vartheta_i(\mu_i^*(\theta'_i)), l_i)} r_0(\theta'_{-0} \mid \theta_0)} \\ &= \frac{r_0(\theta_{-0} \mid \theta_0)}{\sum_{\theta'_{-0}} r_0(\theta'_{-0} \mid \theta_0)} = \rho^*(\theta_{-0} \mid \theta_0, \mu^*(\theta_{-0})). \end{aligned}$$

For every $(\hat{\theta}_{-0}, \hat{l})$ off path, construct \hat{m} as before and let $\rho^{**}(\cdot \mid \theta_0, \hat{\theta}_{-0}, \hat{l}) = \rho^*(\cdot \mid \theta_0, \hat{m})$.

Given the agent's belief ρ^{**} , it follows that the strategy α^{**} is sequentially rational. Take any $(\hat{\theta}_{-0}, l)$ either on path—the same argument applies if $(\hat{\theta}_{-0}, l)$ is off path. Given $(\hat{\theta}_{-0}, l)$, the agent's belief over Θ_{-0} coincides with the beliefs he had after observing $\mu^*(\hat{\theta}_{-0})$ in the original equilibrium σ^* , and α^{**} induces the same distribution over decisions as did α^* having the same set of alternatives available. Therefore, if the agent has a profitable deviation from α^{**} , then α^* could not be sequentially rational from the outset. A contradiction.

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

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

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

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

A contradiction.

Part 2: (\Leftarrow) Let (σ^{**}, ρ^{**}) be a D-O equilibrium of G^{SL} . I now derive a D-O equilibrium (σ^*, ρ^*) of G^D , and show that $f_{\sigma^{**}} = f_{\sigma^*}$.

Consider the strategy μ_i^{**} for P_i . Let $\Theta_i^{**} \times L_i^{**} = \text{Im}(\mu_i^{**})$ be the finite set of signals and menus that P_i can offer according to μ_i^{**} . For every $(\hat{\theta}_i, \hat{l}_i) \in \Theta_i^{**} \times L_i^{**}$, let $\Theta_i(\hat{l}_i)$ be the set of signals that P_i sends while offering \hat{l}_i . Now construct $|\Theta_i(\hat{l}_i)|$ indirect mechanisms $m_i(\hat{l}_i)$ with $\text{Im}(m_i(\hat{l}_i)) = \mathcal{A}_{\hat{l}_i}$. This can be done because $G^D \succcurlyeq G^{SL}$ and $|\Theta_i| \leq |\mathcal{A}_i|$ for every $i \in N$: choose $|\Theta_i|$ messages in \mathcal{A}_i and label the elements in $\mathcal{A}_{\hat{l}_i}$, other than the "default" option \hat{y}_i corresponding to rejecting \hat{l}_i , each time starting with a different element from the set of $|\Theta_i|$ selected messages. And when you get to \hat{y}_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 \hat{l}_i . Then, to every $\hat{\theta}_i \in \Theta_i(\hat{l}_i)$ associate the indirect mechanism $m_i(\hat{\theta}_i, \hat{l}_i)$. Now, with a slight abuse of notation, let P_i strategy in G^D be defined as $\mu_i^*(\theta_i) = m_i(\mu_i^{**}(\theta_i))$. And apply the same construction to every P_i .

Now consider the agent's strategy α^* . If m is on path, then let $\mu^{**}((\mu^*)^{-1}(m)) = (\mu_i^{**}((\mu_i^*)^{-1}(m_i)))_{i \in N}$. For every $d \in \text{Im}(m)$ 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. Let $\hat{l}(\hat{m})$ be the array of menus with $\mathcal{A}_{\hat{l}_i} = \text{Im}(\hat{m}_i)$. Again, for every $\hat{d} \in \text{Im}(\hat{m})$, 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

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

where the profile of signals $\hat{\theta}_{-0}$ is chosen as follows. For every P_i , if $\hat{m}_i = \mu_i^*(\theta_i)$ for some θ_i , then let $\hat{\theta}_i = \mu_i^{**}((\mu_i^*)^{-1}(\hat{m}_i))|_{\Theta_i}$. Otherwise, let $\hat{\theta}_i$ be any element of Θ_i .

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}

$$\begin{aligned}\rho^*(\theta_{-0}|\theta_0, m) &= \frac{\Pr_{\mu^*}(m|\theta_{-0})r_0(\theta_{-0}|\theta_0)}{\sum_{\theta'_{-0}} \Pr_{\mu^*}(m|\theta'_{-0})r_0(\theta'_{-0}|\theta_0)} \\ &= \frac{\Pr_{\mu^{**}}(\mu^{**}((\mu^*)^{-1}(m))|\theta_{-0})r_0(\theta_{-0}|\theta_0)}{\sum_{\theta'_{-0}} \Pr_{\mu^{**}}(\mu^{**}((\mu^*)^{-1}(m))|\theta'_{-0})r_0(\theta'_{-0}|\theta_0)} \\ &= \rho^{**}(\theta_{-0}|\theta_0, \mu^{**}((\mu^*)^{-1}(m))).\end{aligned}$$

If instead \hat{m} is off path. Let $\hat{l}(\hat{m})$ be the array of menus with $\mathcal{A}_{\hat{l}_i} = \text{Im}(\hat{m}_i)$, and let

$$\rho^*(\theta_{-0}|\theta_0, \hat{m}) = \rho^{**}(\theta_{-0}|\theta_0, \hat{\theta}_{-0}, \hat{l}(\hat{m})),$$

where the signal profile $\hat{\theta}_{-0}$ is chosen as in the construction of α^* off path.

I claim that the profile (σ^*, ρ^*) 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 P_i has a profitable deviation to a mechanism \tilde{m}_i , that is

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

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 \tilde{m}_i is equivalent to offering \tilde{l}_i with $\mathcal{A}_{\tilde{l}_i} = \text{Im}(\tilde{m}_i)$ and send the signal $\tilde{\theta}_i$ that I used in the construction of α^* and ρ^* above. Since the agent is responding to μ_{-i}^* as he was responding to μ_{-i}^{**} , it follows that

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

A contradiction.

7.2 Proof of Lemma 1

Let (σ^*, ρ^*) be a D-O equilibrium of G^{LRR} . I shall construct a D-O, truthful, equilibrium of G^{LRR} that sustains $f_{\sigma^{**}} = f_{\sigma^*}$.

Consider first P_i . For every θ_i , let $l_i^{RR*}(\theta_i) = \mu_i^*(\theta_i)$ and let $T_i(l_i^{RR*}(\theta_i)) = (\mu_i^*)^{-1}(l_i^{RR*}(\theta_i)) \subset \Theta_i$ —clearly, $l_i^{RR*}(\theta'_i) = l_i^{RR*}(\theta_i)$ if and only if $\theta'_i \in T_i(l_i^{RR*}(\theta_i))$. Now for each θ_i , construct a new list $l_i^{RR**}(\theta_i)$ as follows: for each $h_i^{RR*} \in \mathcal{A}_{l_i^{RR*}(\theta_i)}$, let $h_i^R(h_i^{RR*}) = h_i^{RR*} \triangleright \pi_i^*(l_i^{RR*}(\theta_i))$ and let $h_i^{RR**}(h_i^{RR*}) = h_i^R(h_i^{RR*})|_{T_i(l_i^{RR*}(\theta_i))}$; the new $l_i^{RR**}(\theta_i)$ is the list consisting of these h_i^{RR**} . Clearly, $l_i^{RR**}(\theta_i) = l_i^{RR*}(\theta_i)$ if and only if $\theta_i \in T_i(l_i^{RR*}(\theta_i))$. So, for every θ_i let $\mu_i^{**}(\theta_i) = l_i^{RR**}(\theta_i)$, and $\pi_i^{**}(\theta_i, \mu_i^{**}(\theta_i)) = \theta_i$, whereas for $l_i^{RR} \neq \mu_i^{**}(\theta_i)$, let $\pi_i^{**}(\theta_i; l_i^{RR}) = \pi_i^*(\theta_i; l_i^{RR})$. Finally, repeat the same construction for every P_i .

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

$$\alpha^{**}(h^{RR}; \theta_0, l^{RR}) = \sum_{\{h^{RR*} \in \mathcal{A}_{l^{RR*}} : h^{RR**}(h^{RR*}) = h^{RR}\}} \alpha^*(h^{RR*}; \theta_0, l^{RR*}),$$

where $l_i^{RR*} = \mu_i^*((\mu_i^{**})^{-1}(l^{RR}))$. Also, using Bayes rule it follows that on path, for every θ ,

$$\begin{aligned} \rho^{**}(\theta_{-0} | \theta_0, l^{RR}) &= \frac{\Pr_{\mu^{**}}(l^{RR} | \theta_{-0}) r_0(\theta_{-0} | \theta_0)}{\sum_{\theta_{-0}} \Pr_{\mu^{**}}(l^{RR} | \theta'_{-0}) r_0(\theta'_{-0} | \theta_0)} \\ &= \frac{\Pr_{\mu^*}(l^{RR*} | \theta_{-0}) r_0(\theta_{-0} | \theta_0)}{\sum_{\theta_{-0}} \Pr_{\mu^*}(l^{RR*} | \theta'_{-0}) r_0(\theta'_{-0} | \theta_0)} = \rho^*(\theta_{-0} | \theta_0, l^{RR*}), \end{aligned}$$

where again $l_i^{RR*} = \mu_i^*((\mu_i^{**})^{-1}(l^{RR}))$.

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 each Y_i ; 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 (σ^{**}, ρ^{**}) , for every l^{RR} off path let $\alpha^{**}(\cdot; \cdot, l^{RR}) = \alpha^*(\cdot; \cdot, l^{RR})$ and $\rho^{**}(\cdot | \cdot, l^{RR}) = \rho^*(\cdot | \cdot, l^{RR})$. Using the assumption that (σ^*, ρ^*) is an equilibrium of G^{LRR} , it follows that such a profile (σ^{**}, ρ^{**}) is also an equilibrium of G^{LRR} , which is D-O and truthful.

7.3 Proof of Theorem 2

Part 1: Let (σ^*, ρ^*) be the D-O equilibrium of G^M sustaining the SCF f . The proof proceeds by constructing the desired equilibrium (σ^{**}, ρ^{**}) of G^{LRR} , and concludes by showing that $f_{\sigma^{**}} = f$.

Given the offer strategy μ_i^* of P_i , let $\mu_i^*(\Theta_i)$ be the set of mechanisms that are offered by some θ_i in the equilibrium σ^* . For every θ_i it is possible to identify the subset $T_i(\mu_i^*(\theta_i)) = (\mu_i^*)^{-1}(\mu_i^*(\theta_i)) \subset \Theta_i$, that is the set of types of P_i who offer the same mechanism $\mu_i^*(\theta_i)$. Using the communication strategy π_i^* and the mapping λ^R , recover the list of DRMs $l_i^R(\theta_i) = \lambda^R(\mu_i^*(\theta_i), \pi_i^*)$ that corresponds to the mechanism offered by θ_i under μ_i^* . Finally, for every θ_i identify the RDRM $l_i^{RR}(\theta_i)$ in $H_i^R(T_i(\mu_i^*(\theta_i)))$ by restricting every h_i^R in $l_i^R(\theta_i)$ to $h_i^{RR} = h_i^R|_{T_i(\mu_i^*(\theta_i))}$. By construction, it follows that for every θ_i and $\theta'_i \in T_i(\mu_i^*(\theta_i))$, $l_i^{RR}(\theta_i) = l_i^{RR}(\theta'_i)$. Therefore, let P_i 's offer strategy in G^{LRR} be $\mu_i^{**}(\theta_i) = l_i^{RR}(\theta_i)$. And repeat the same construction for every P_i .

Now consider the communication strategies of the principals and the agent on path. For every θ_i of P_i , let $\pi_i^{**}(\theta_i, \mu_i^{**}(\theta_i)) = \theta_i$. With regard to the agent, given l^{RR} on path, let $m(l^{RR}) = \mu^*((\mu^{**})^{-1}(l^{RR}))$ and $T(l^{RR}) = T(\mu^*((\mu^{**})^{-1}(l^{RR})))$, and for every $h^{RR} \in \mathcal{A}_{l^{RR}}$, let $\mathcal{A}(h^{RR}) = \{a | v^R(a; m(l^{RR}), \pi^*)|_{T(l^{RR})} = h^{RR}\}$. In words, $\mathcal{A}(h^{RR})$ is the set of messages that allowed the agent to induce a profile of DRMs (given π^*) such that, if we restrict each of them to $T_i(l_i^{RR})$, we obtain h^{RR} . Then for each θ_0 , let

$$\alpha^{**}(h^{RR}; \theta_0, l^{RR}) = \sum_{a \in \mathcal{A}(h^{RR})} \alpha^*(a; \theta_0, m(l^{RR})).$$

Since the strategy μ^{**} is deterministic, we can use Bayes rule to derive the on-path beliefs of the agent as follows: for every θ_{-0}, θ_0 and l^{RR} on path

$$\begin{aligned} \rho^{**}(\theta_{-0} | \theta_0, l^{RR}) &= \frac{\Pr_{\mu^{**}}(l^{RR} | \theta_{-0}) r_0(\theta_{-0} | \theta_0)}{\sum_{\theta'_{-0}} \Pr_{\mu^{**}}(l^{RR} | \theta'_{-0}) r_0(\theta'_{-0} | \theta_0)} \\ &= \frac{\Pr_{\mu^*}(m(l^{RR}) | \theta_{-0}) r_0(\theta_{-0} | \theta_0)}{\sum_{\theta'_{-0}} \Pr_{\mu^*}(m(l^{RR}) | \theta'_{-0}) r_0(\theta'_{-0} | \theta_0)} = \rho^*(\theta_{-0} | \theta_0, m(l^{RR})). \end{aligned}$$

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

$$\begin{aligned} f_{\sigma^*}(y; \theta) &= \sum_{\mathcal{A}} \sum_{\mathcal{P}} m^*(y; a, p) \pi^*(p; \theta_{-0}, m^*) \alpha^*(a; \theta_0, m^*) \\ &= \sum_{h^{RR} \in \mathcal{A}_{l^{RR**}}} \prod_{i \in N} h_i^{RR}(y_i; \theta_i) \sum_{a \in \mathcal{A}(h^{RR})} \alpha^*(a; \theta_0, m^*) = f_{\sigma^{**}}(y; \theta). \end{aligned}$$

Using this result, we can see that (π^{**}, α^{**}) is a continuation equilibrium given any l^{RR} on path and ρ^{**} . This is because, given any (θ_0, l^{RR}) 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 $(\theta_0, m(l^{RR}))$. Furthermore, given π^{**} the agent can't induce any distribution over outcomes that he couldn't induce after $m(l^{RR})$ given π^* . Similarly, each θ_i of P_i can't have a profitable deviation from her truthful strategy after offering l_i^{RR} , because any such deviation was available also in the original game after offering $m_i(l_i^{RR})$.

If l^{RR} is off path and involves deviations by more than one principal, choose any belief for the agent $\rho^{**}(\cdot | \cdot, l^{RR})$ and equilibrium of the resulting continuation game—finiteness of $\mathcal{A}_{l^{RR}}$ and Θ_{-0} ensures the existence of at least one continuation equilibrium. Now consider l^{RR} off path induced by the deviation of one single principal P_i . Suppose first that θ_i of P_i offers $l_i^{RR} = \mu_i^{**}(\theta'_i)$ for $\theta'_i \neq \theta_i$. Then the agent's strategy and beliefs are defined as on path, and let $\pi_i^{**}(\theta_i, l_i^{RR}) \in \Delta(T_i(\mu_i^*(\theta'_i)))$ be any best reply to α^{**} and σ_{-i}^{**} . Then for any P_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^M . The remaining case is when l^{RR} is off path (in the sense that it can't occur with positive probability according to σ_i^{**}) and only P_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 P_i deviates to offering \hat{l}_i^{RR} off path, whereas all other principals follow σ_{-i}^{**} , and that for any choice of $\rho^{**}(\cdot | \cdot, \hat{l}^{RR})$ any continuation equilibrium can't deter θ_i from such a deviation. I claim that then (σ^*, ρ^*) can't be an equilibrium of G^M . To see this, recall that since $G^M \succneq G^{LRR}$, for such an \hat{l}_i^{RR} there exist an indirect mechanism \tilde{m}_i and a communication strategy $\tilde{\pi}_i(\tilde{m}_i)$ for such a mechanism, such that $\lambda^R(\tilde{m}_i, \tilde{\pi}_i)|_{T_i(\hat{l}_i^{RR})} = \hat{l}_i^{RR}$. Hence, consider the mechanism \hat{m}_i such that, for every $\theta_i \in T_i(\hat{l}_i^{RR})$, $\hat{m}_i(\theta_i, \cdot) = \tilde{m}_i(\tilde{\pi}_i(\theta_i, \tilde{m}_i), \cdot)$ and, for every other $p_i \in \mathcal{P}_i \setminus T_i(\hat{l}_i^{RR})$, $\hat{m}_i(p_i, \cdot) = \hat{m}_i(\theta'_i, \cdot)$ where θ'_i is the element of $T_i(\hat{l}_i^{RR})$ with the lowest index. Then for type θ_i of P_i offering \hat{m}_i in G^M is essentially equivalent to offering \hat{l}_i^{RR} in G^{LRR} and then follow the same communication strategy in both games. Now, given the strategy σ_{-i}^* and 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{l}^{RR} given π^{**} . By assumption, it is not possible to find a belief ρ^{**} and a strategy α^{**} so that α^{**} is a

best reply to π^{**} and vice versa, and θ_i of P_i doesn't profit by deviating to \hat{l}_i^{RR} . 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 P_i is deterred from offering \hat{m}_i followed by $\hat{\pi}_i$. This contradicts the assumption that (σ^*, ρ^*) is an equilibrium of G^M , and the first claim I started with follows. We conclude that it is possible to extend the on-path strategies and beliefs that I have constructed above to obtain an equilibrium (σ^{**}, ρ^{**}) of G^{LRR} such that $f_{\sigma^{**}} = f_{\sigma^*}$.

Part 2: Suppose (σ^{**}, ρ^{**}) is a D-O truthful equilibrium of G^{LRR} that sustains the SCF f . I shall show that there exists a D-O equilibrium (σ^*, ρ^*) of G^M that sustains the same SCF.

Consider P_i and her offer strategy μ_i^{**} . For every $l_i^{RR} \in \mu_i^{**}(\Theta_i)$, consider the mechanism $m_i^{l_i^{RR}} \in M_i$ constructed as follows. Since $G^M \succcurlyeq G^{LRR}$, there exist \tilde{m}_i and $\tilde{\pi}_i$ such that $\lambda^R(\tilde{m}_i, \tilde{\pi}_i)|_{T_i(l_i^{RR})} = l_i^{RR}$, where $T_i(l_i^{RR})$ is the set of $\hat{\theta}_i$ that l_i^{RR} allows P_i to send. Since $\Theta_i \subset \mathcal{P}_i$, for each $\theta_i \in T_i(l_i^{RR})$, let $m_i^{l_i^{RR}}(\theta_i, \cdot) = \tilde{m}_i(\tilde{\pi}_i(\theta_i, \tilde{m}_i), \cdot)$, and for every other $p_i \in \mathcal{P}_i \setminus T_i(l_i^{RR})$, let $m_i^{l_i^{RR}}(p_i, \cdot) = m_i^{l_i^{RR}}(\theta'_i, \cdot)$ where θ'_i is the element of $T_i(l_i^{RR})$ with the lowest index. There are two cases to consider. The first case corresponds to $|T_i(l_i^{RR})| > 1$. In this case, each $m_i^{l_i^{RR}}$ signals the same amount of information about P_i 's type as l_i^{RR} because both $m_i^{l_i^{RR}}$ and l_i^{RR} essentially restrict P_i to choose a message in $T_i(l_i^{RR})$. The second case corresponds to $|T_i(l_i^{RR})| = 1$, which implies that l_i^{RR} 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 $l_i^{RR}(\theta_i)$ the list of RDRMs that θ_i offers under μ_i^* and let $m_i^{l_i^{RR}(\theta_i)}$ be the indirect mechanism we have just constructed. Then, for every P_i and θ_i , define $\mu_i^*(\theta_i) = m_i^{l_i^{RR}(\theta_i)}$.

Now consider the on-path communication strategies of the principals' and the agent's. First, for every P_i and θ_i , let $\pi_i^*(\theta_i, m_i^{l_i^{RR}(\theta_i)}) = \theta_i$. Now consider the agent. Given any profile m^* on path, recover the corresponding $l^{RR}(m^*)$ given the above construction of m^* . Also, for every $h^{RR} \in \mathcal{A}_{l^{RR}(m^*)}$, let $\mathcal{A}(h^{RR}) = \{a | v^R(a; m^*, \pi^*)|_{T(l^{RR}(m^*))} = h^{RR}\}$ and choose *one* message $a(h^{RR}) \in \mathcal{A}(h^{RR})$. Then for each θ_0 , define

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

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

$$\begin{aligned} \rho^*(\theta_{-0} | \theta_0, m^*) &= \frac{\Pr_{\mu^*}(m^* | \theta_{-0}) r_0(\theta_{-0} | \theta_0)}{\sum_{\theta'_{-0}} \Pr_{\mu^*}(m^* | \theta'_{-0}) r_0(\theta'_{-0} | \theta_0)} \\ &= \frac{\Pr_{\mu^{**}}(l^{RR}(m^*) | \theta_{-0}) r_0(\theta_{-0} | \theta_0)}{\sum_{\theta'_{-0}} \Pr_{\mu^{**}}(l^{RR}(m^*) | \theta'_{-0}) r_0(\theta'_{-0} | \theta_0)} = \rho^{**}(\theta_{-0} | \theta_0, l^{RR}(m^*)). \end{aligned}$$

Given the construction of (μ^*, π^*) and α^* on path, it follows that $f_{\sigma^*} = f_{\sigma^{**}}$ because each type θ_i of each P_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 ρ^* .

It remains to show that π^* , α^* , and ρ^* can be extended off path to obtain an equilibrium (σ^*, ρ^*) of G^M . If m is off path because more than one P_i deviated, let $\rho^*(\cdot | \cdot, m)$ 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 P_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^M , is equivalent to offering $l_i^{RR}(\tilde{\theta}_i)$, again followed by some communication strategy, given α^{**} and $(\mu_{-i}^{**}, \pi_{-i}^{**})$ in G^{LRR} . 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 any $\tilde{\theta}_i$. Suppose that for any specification of $\rho^*(\cdot | \cdot, \hat{m})$, 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 $\hat{l}_i^R = \lambda^R(\hat{m}_i, \hat{\pi}_i)$. Hence, for θ_i of P_i offering \hat{m}_i followed by $\hat{\pi}_i$ in G^M is equivalent to offering \hat{l}_i^R in G^{LRR} followed by a truthful report of θ_i . Given the profile σ_{-i}^{**} and any l_{-i}^{RR**} resulting from μ_{-i}^{**} , the set of distributions over outcomes that the agent can induce after $(\hat{l}_i^R, l_{-i}^{RR**})$ 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 $\rho^*(\cdot | \cdot, (\hat{m}_i, m_{-i}^*))$ there is no $\alpha^*(\cdot, (\hat{m}_i, m_{-i}^*))$ that is a best reply to $\hat{\pi}_i$ and vice versa, and can deter the deviation by θ_i of P_i to \hat{m}_i . But then, there is no specification of $\rho^{**}(\cdot | \cdot, (\hat{l}_i^R, l_{-i}^{RR**}))$ for which a continuation equilibrium can be constructed that deters θ_i of P_i from deviating to l_i^R in G^{LRR} . This implies that (σ^{**}, ρ^{**}) can't be an equilibrium of G^{LRR} . A contradiction that implies that it is possible to complete π^* and α^* off path to make any deviation from the on path behavior constructed above unprofitable, and thus obtaining an equilibrium (σ^*, ρ^*) of G^M .

References

- [1] Aliprantis, C.D. and Border, K.C. (2006), *Infinite dimensional analysis: a Hitchhiker's guide*. Springer Verlag, 3rd Edition.
- [2] Bernheim, D. and M. Whinston (1986a), *Menu Auctions*. Quarterly Journal of Economics, Vol. 101, No. 1, 1-32
- [3] Bernheim, D. and Whinston, M. (1986b), *Common Agency*. Econometrica, Vol. 54, No.4, pp. 923-942.
- [4] Celik, G. and M. Peters (2011), *Reciprocal Relationships and Mechanism Design*, Mimeo, University of British Columbia.
- [5] Epstein, L. and Peters, M. (1999), *A Revelation Principle for Competing Mechanisms*. Journal of Economic Theory, 88, 119-160.
- [6] Laffont, J.-J., Tirole, J., (Sep., 1988), *The Dynamics of Incentive Contracts*. Econometrica, Vol. 56, No. 5, pp. 1153-1175.

- [7] Martimort, D., Moreira, H., (2010), *Common Agency and Public Good Provision under Asymmetric Information*. Theoretical Economics, Vol. 5, No. 2, 159-213.
- [8] Martimort, D. and L. Stole, (2002), *The Revelation and Delegation Principles in Common Agency Games*. Econometrica 70, 1659-1674.
- [9] Maskin, E. and J. Tirole, (1990), *The Principal-Agent Relationship with an Informed Principal, I: The Case of Private Values*. Econometrica, 58: 379-410.
- [10] Maskin, E. and J. Tirole, (1992), *The Principal-Agent Relationship with an Informed Principal, II: The Case of Common Values*. Econometrica, 60: 1-42.
- [11] Myerson, R., (1983), *Mechanism Design with an Informed Principal*. Econometrica, 51: 1767-1797.
- [12] Myerson, R., (1997), *Game Theory: analysis of conflict*. Harvard University Press.
- [13] Fudenberg, D. and Tirole, J., (1991), *Game Theory*. MIT Press.
- [14] Pavan, A. and G. Calzolari (2010), *Truthful Revelation Mechanisms for Simultaneous Common Agency Games*. AEJ: Microeconomics, 2(2), 132-190
- [15] Peters, M., (2001), *Common Agency and the Revelation Principle*. Econometrica, 69: 1349-1372.
- [16] Peters, M., (2011), *A Revelation Principle for Competing Mechanism Games*. Mimeo, University of British Columbia.
- [17] Peters, M. and B. Szentes (2011), *Definable and Contractible Contracts*, Econometrica (forthcoming).