THE CARLO ALBERTO NOTEBOOKS

Collegio Carlo Alberto

Truthful Revelation Mechanisms for Simultaneous Common Agency Games

Alessandro Pavan
Giacomo Calzolari

Working Paper No. 85
December 2008
www.carloalberto.org

Truthful Revelation Mechanisms for Simultaneous Common Agency Games*

Alessandro Pavan† **Giacomo Calzolari**‡

November 2008

Abstract

This paper considers games in which multiple principals contract simultaneously with the same agent. We introduce a new class of revelation mechanisms that, although it does not always permit a complete equilibrium characterization, it facilitates the characterization of the equilibrium outcomes that are typically of interest in applications (those sustained by pure-strategy profiles in which the agent's behavior in each relationship is Markov, i.e., it depends only on payoff-relevant information such as the agent's type and the decisions he is inducing with the other principals). We then illustrate how these mechanisms can be put to work in environments such as menu auctions, competition in nonlinear tariffs, and moral hazard settings. Lastly, we show how one can enrich the revelation mechanisms, albeit at a cost of an increase in complexity, to characterize also equilibrium outcomes sustained by non-Markov strategies and/or mixed-strategy profiles.

JEL Classification Numbers: D89, C72.

Keywords: Mechanism design, contracts, revelation principle, menus, endogenous payoff-relevant information.

---

*This is a substantial revision of an earlier version that circulated with the same title. For discussions and useful comments, we thank seminar participants at various conferences and institutions where this paper has been presented and in particular Eddie Dekel, Mike Peters, Marciano Siniscalchi and Jean Tirole for suggestions that helped us improve the paper. We are grateful to Itai Sher for excellent research assistance. Pavan also thanks the ospitality of Collegio Carlo Alberto where part of this project was completed.

†Department of Economics, Northwestern University, Email: alepavan@northwestern.edu

‡Department of Economics, University of Bologna. E-mail: giacomo.calzolari@unibo.it

1 Introduction

It is by now well understood that in environments in which multiple principals[1] contract non-cooperatively with the same agent, the Revelation Principle[2] is invalid. The reason is that the agent's preferences over the decisions of one principal depend not only on his "type" (i.e. his exogenous private information) but also on the decisions induced with the other principals.[3]

Two solutions have been proposed in the literature. Epstein and Peters (1999) have suggested that the agent should communicate not only his type but also the mechanisms offered by the other principals. However, describing a mechanism requires an appropriate language. The main contribution of Epstein and Peters is in proving existence of a universal language that is rich enough to describe all possible mechanisms. This language also permits them to identify a class of universal mechanisms with the property that any indirect mechanism can be embedded into it. Since universal mechanisms have the agent truthfully report all his private information, they can be considered direct revelation mechanisms and therefore a *universal Revelation Principle* holds.

Although a remarkable contribution, the use of universal mechanisms in applications has been precluded by the complexity of the universal language. In fact, when asking the agent to describe principal j's mechanism, principal i has to take into account that principal j's mechanism may also ask the agent to describe principal i's mechanism, whether this mechanism depends on principal j's mechanism...and so on, leading to the so called "infinite regress" problem. The universal language is in fact obtained as the limit of a sequence of enlargements of the message space, where at each enlargement the corresponding direct mechanism becomes more complex to describe and hence more difficult to use when searching for equilibrium outcomes.

The second solution, proposed by Peters (2001) and Martimort and Stole (2002), is to restrict the principals to offer *menus* of contracts. They have shown that, for any equilibrium relative to any game with arbitrary sets of mechanisms for the principals, there exists an equilibrium in the game in which the principals are restricted to offer menus that sustains the same outcomes. In this equilibrium, the principals simply offer the menus in the range of the mechanisms they would have offered in the equilibrium of the indirect game and delegate to the agent the choice of the contractual terms. This result is referred to as the *Menu Theorem* and is the analog of the *Taxation Principle* for games with a single mechanism designer.[4]

[1] We refer to the players who offer the contracts either as the *principals* or as the *mechanism designers*. The two expressions are meant to be synonyms. Furthermore, we adopt the convention of using feminine pronouns for the principals and masculine pronouns for the agent.

[2] See, among others, Gibbard (1973), Green and Laffont (1977) and Myerson (1979).

[3] Problems with standard direct revelation mechanisms have been documented, among others, in Katz (1991), McAfee (1993), Peck (1997), Epstein and Peters (1999), Peters (2001) and Martimort and Stole (1997, 2002). Recent work by Peters (2003), Attar, Piaser and Porteiro (2007,a,b), and Attar, Majumadar, Piaser, and Porteiro (2007) has identified special cases in which these problems do not emerge.

[4] The result is also referred to as the "Delegation Principle" (e.g. Martimort and Stole, 2002). For the Taxation

The Menu Theorem has proved quite useful in applications. However, contrary to the Revelation Principle, it provides no indication on how the agent uses the different allocations in the menu as a function of his private information, nor does it permit one to restrict attention to any particular type of menus. This is what we aim at doing in this paper by showing that, in most cases of interest for applications, one can still conveniently describe the agent's interaction with each of his principals through *revelation mechanisms*. The structure of these mechanisms is however more general than the one for games with a single mechanism-designer. Nevertheless, contrary to universal mechanisms, it is not conductive to any "infinite regress." In the revelation mechanisms we propose, the agent is asked to report his exogenous type along with the endogenous (payoff-relevant) decisions he is inducing with the other principals. To fix ideas, let $V(\theta, \delta_i, \delta_{-i})$ denote the agent's payoff, with $\theta \in \Theta$ denoting the agent's type, $\delta_i \in \mathcal{D}_i$ denoting a decision for principal i and $\delta_{-i} \in \mathcal{D}_{-i}$ denoting a profile of decisions for all principals other than i.[5] An incentive-compatible revelation mechanism is a mapping $\phi_i^r : \Theta \times \mathcal{D}_{-i} \to \mathcal{D}_i$ with the property that for any (θ, δ_{-i}), $\phi_i^r(\theta, \delta_{-i}) \in \arg\max_{\delta_i \in \mathrm{Im}(\phi_i^r)} V(\theta, \delta_i, \delta_{-i})$, where $\mathrm{Im}(\phi_i^r)$ denotes the range of ϕ_i^r.

Describing the interaction of the agent with each of his principals through an incentive-compatible revelation mechanism is convenient because it permits one to specify which decisions the agent induces as a function of his type and the decisions he induces with the other principals. In particular, it permits one to specify which decisions the agent takes in response to deviations by any of the other principals. This in turn can give guidance on which outcomes can be sustained in equilibrium.

The mechanisms described above are appealing because they capture the essence of common agency, i.e., the fact that the agent's preferences vis a vis the decisions of each principal depend not only on his type θ but also on the decisions δ_{-i} he induces with the other principals.[6] However, this property does not guarantee that it is always without loss to restrict the agent's behavior to depend only on (θ, δ_{-i}). In fact, when indifferent, the agent may condition his behavior also on payoff-irrelevant information such as the decisions included in the menus offered by the other principals that he preferred not to select. Furthermore, when indifferent, the agent may randomize over the principals' decisions inducing a correlation that cannot always be replicated by having the agent simply report (θ, δ_{-i}) to each principal. As a consequence, not all outcomes can be sustained by restricting the principals to offer the simple revelation mechanisms described above.

While we find these considerations intriguing from a theoretical viewpoint, we seriously doubt their relevance in applications.

Principle, see Rochet (1986) and Guesnerie (1995).

[5] Depending on the application of interest, a decision can be a price-quantity pair, as in the case of competion in nonlinear tariffs, a reward scheme, as in menu auctions, or an incentive contract, as in moral hazard settings.

[6] A special case is when preferences are separable, as in Attar, Majumdar, Piasier, and Porteiro (2007), in which case they depend only on θ.

Our concerns with mixed-strategy equilibria are the usual ones: outcomes sustained by the agent mixing over the contracts offered by the principals or by the principals mixing over the menus they offer to the agent are typically not robust. Furthermore, when principals can offer *all possible* menus (including those containing lotteries over contracts), it is very hard to construct (non-degenerate) examples in which the agent is made indifferent over some of the contracts offered by the principals and, at the same time, no principal has an incentive to change the composition of her menu so as to break the agent's indifference and induce him to choose the decisions that are most favorable to her (see also the discussion in Section 5.2).

Our concerns with equilibrium outcomes sustained by a strategy for the agent that is not Markov, i.e., that it depends also on payoff-irrelevant information, are motivated by the observation that this type of behavior does not seem plausible in most real-world situations. Think of a buyer purchasing products or services from multiple sellers. While it is plausible that the quality/quantity purchased from seller i depends on the quality/quantity purchased from seller j (this is the intrinsic nature of the common agency problem which leads to the failure of the standard revelation principle), it does not seem plausible that, *for given choice with seller j*, the purchase from seller i depends on payoff-irrelevant information such as the other price-quantity offers in seller j's menu that the buyer decided not to choose.[7]

For most of the analysis, we thus focus on outcomes sustained by pure-strategy profiles in which the agent's behavior is Markov.[8] We first show that any such outcome can be sustained as a *truthful equilibrium* of the *revelation game*. We also show that, despite the fact that only certain menus can be offered in the revelation game, any truthful equilibrium is robust in the sense that its outcome can also be sustained by an equilibrium of the menu game. This guarantees that equilibrium outcomes in the revelation game are not artificially sustained by the fact that the principals are forced to choose from a restricted set of menus.

We then proceed by addressing the question of whether there exist environments in which restricting the agent's strategy to be Markov is not only appealing but actually without any loss of generality. Clearly, this is always the case when the agent's preferences are strict, for it is only when the agent is indifferent that his behavior may depend on payoff-irrelevant information. Furthermore, even when the agent can be made indifferent, restricting attention to Markov strategies is always without loss of generality when information is complete and when the principals' preferences are sufficiently aligned in the following sense: for any profile of decision δ_{-i}, and for any menu of decisions $D_i \subset \mathcal{D}_i$, there exists a decision $\delta_i \in D_i$ among those that are optimal for the agent given δ_{-i} such that the payoff of any principal P_j, $j \neq i$, under (δ_i, δ_{-i}) is (weakly) lower than under any

[7] Note that the fact that the agent's strategy is Markov does not imply that the principals can be restricted to offer menus that contain only the price-quantity pairs that are selected in equilibrium.

[8] Note that, while the definition of Markov strategy given here is different from the one considered in the literature on dynamic games (see e.g. Pavan and Calzolari, 2008), it shares with that definition the same spirit.

other profile (δ_i', δ_{-i}) such that δ_i' is optimal for the agent given δ_{-i}. This condition guarantees that, given δ_{-i}, the decision the agent induces with P_i to punish a deviation by one of P_i's opponents need not depend on the identity of the deviating principal. This property is trivially satisfied when there are only two principals. It is also satisfied, for example, when the principals are retailers competing "a la Cournot" in a downstream market (each retailer's payoff is then decreasing in the quantity the agent–here in the role of a common manufacturer–sells to any of the other principals).

As for the restriction to complete information, the only role that this restriction plays is to rule out the possibility that the equilibrium outcomes are sustained by the agent punishing a deviation, say by principal j, by choosing in state θ the equilibrium decisions $\delta_{-i}^*(\theta)$ with all principals other than i and then changing his behavior with principal i by inducing a decision $\delta_i \neq \delta_i^*(\theta)$. Allowing type θ to respond to the equilibrium decisions $\delta_{-i}^*(\theta)$ with a decision $\delta_i \neq \delta_i^*(\theta)$ may be necessary to discourage certain deviations by principal j. This in turn implies that Markov strategies need not be without loss of generality when information is incomplete. However, because this is the only complication that arises with incomplete information, we show that one can safely restrict attention to Markov strategies if one imposes a mild refinement on the solution concept which we call "*Conformity to Equilibrium.*" This refinement simply imposes that each type of the agent selects the equilibrium decision with each principal when the latter offers the equilibrium menu and the decisions the agent induces with the other principals are the equilibrium ones.[9] Again, in most real world situations, we find such a behavior plausible.

The rest of the paper is organized as follows. Section 2 describes the contracting environment. Section 3 contains the main characterization results. Section 4 shows how the simple revelation mechanisms described above can be put to work in applications such as competition in non-linear tariffs, menu auctions, and moral hazard settings. Section 5 shows how the revelation mechanisms can be enriched to characterize also equilibrium outcomes sustained by non-Markov strategies and/or mixed strategy equilibria. Section 6 concludes. All proofs are either in the Appendix or in the Supplementary Material.

Qualification. While the approach here is similar (in spirit) to the one in Pavan and Calzolari (2008) for sequential common agency, there are important differences due to the simultaneity of contracting. First, the notion of Markov strategies considered here is forward-looking instead of backward-looking and takes into account the fact that, when choosing which messages to send to principal i, the agent has not committed yet any decision with any of the other principals. Second, contrary to sequential games, the agent can condition his behavior not only on the mechanisms offered upstream but on the entire profile of mechanisms offered by all principals. These differences explain why, despite certain similarities, the results do not follow from the arguments in that paper.

[9] Note that this refinement is milder than the "conservative behavior" refinement considered in Attar, Majumdar, Piasier, and Porteiro (2007).

2 The environment

The following model encompasses essentially all variants of simultaneous common agency examined in the literature.

Players, actions and contracts. There are $n \in \mathbb{N}$ principals who contract simultaneously and non-cooperatively with the same agent, A. Each principal P_i, $i \in \mathcal{N} \equiv \{1, ..., n\}$, must select a contract y_i from a set of feasible contracts Y_i. A contract $y_i : E \to \Delta(\mathcal{A}_i)$ specifies (a lottery over) the actions $a_i \in \mathcal{A}_i$ that P_i will take in response to the agent's action/effort $e \in E$.[10] Both a_i and e may have different interpretations depending on the application of interest. When the agent is a buyer purchasing from multiple sellers, a_i stands for the price of seller i and e for a vector of quantities/qualities. When instead A is a politician lobbied by different interest groups, a_i represents a campaign contribution and E the set of relevant policies.

Depending on the environment, the set of feasible contracts Y_i may also be more or less restricted. For example, in a trade relationship, the price a_i of seller i may not depend on the quantities/qualities of other sellers.[11] In a moral hazard model, because e is not observable by the principals, the function $y_i : E \to \Delta(\mathcal{A}_i)$ is necessarily constant over E and an action $a_i \in \mathcal{A}_i$ represents a state-contingent payment that rewards the agent as a function of some exogenous performance measure correlated with the agent's effort. Finally, in certain environments, only deterministic contracts may be enforceable which can be captured by restricting each $y_i \in Y_i$ to respond to each $e \in E$ with a degenerate lottery $y_i(e)$ that assigns measure one to an element of $\mathcal{A}_i$.

Payoffs. Principal i's payoff is described by the function $u_i(e, a, \theta)$, whereas the agent's payoff by the function $v(e, a, \theta)$. The vector $a \equiv (a_1, ..., a_n) \in \mathcal{A} \equiv \prod_{i=1}^{n} \mathcal{A}_i$ denotes a profile of actions for the principals, while the variable θ denotes the agent's exogenous private information. The principals share a common prior over θ represented by the distribution F with support Θ. All players are expected-utility maximizers.

To avoid the usual measure-theoretic complications, we will often assume that $\mathcal{A}$, E and Θ are finite sets.

Mechanisms. Principals compete in mechanisms. A mechanism for P_i consists of a message (or communication) space $\mathcal{M}_i$ and a measurable mapping $\phi_i : \mathcal{M}_i \to \mathcal{D}_i$ where $\mathcal{D}_i \subseteq \Delta(Y_i)$ denotes a (compact) set of feasible lotteries over Y_i.[12] When A selects a message $m_i \in \mathcal{M}_i$, P_i thus randomizes

[10] Throughout, for any measurable set Ω, $\Delta(\Omega)$ will denote the set of probability measures over Ω. Furthermore, given any $\omega \in \Delta(\Omega)$, $Supp[\omega]$ will denote the support of ω.

[11] An exception is Martimort and Stole (2005).

[12] Again, depending on the application, the sets $\mathcal{D}_i$ may be more or less restricted. For example, in certain applications, it is customary to assume that, not only the contracts y_i must be deterministic, but also the lotteries over the contracts Y_i selected through the mechanism must be degenerate. More generally, the sets $\mathcal{D}_i$ incorporate

over Y_i according to the lottery $\delta_i = \phi_i(m_i) \in \mathcal{D}_i$.

In the following, we will refer to $a_i \in \mathcal{A}_i$ as the *action* and to δ_i as the *decision* for principal i. When A does not have any action to take after communicating with the principals (that is, when $|E| = 1$), δ_i reduces to a lottery over the set $\mathcal{A}_i$ of payoff-relevant actions.

To save on notation, in the sequel we will denote a mechanism simply by ϕ_i, thus dropping the specification of the message space $\mathcal{M}_i$ whenever this does not create confusion. Given a mechanism ϕ_i we then denote by $\text{Im}(\phi_i) \equiv \{\delta_i \in \Delta(Y_i) : \exists\, m_i \in \mathcal{M}_i \text{ s.t. } \phi_i(m_i) = \delta_i\}$ the set of lotteries in the range of ϕ_i.

For any common agency game Γ, we will then denote by Φ_i the set of feasible mechanisms for P_i, by $\phi \equiv (\phi_1, ..., \phi_n) \in \Phi \equiv \prod_{j=1}^{n} \Phi_j$ a profile of mechanisms for the n principals, and by $\phi_{-i} \equiv (\phi_1, ..., \phi_{i-1}, \phi_{i+1}, ..., \phi_n) \in \Phi_{-i} \equiv \prod_{j \neq i} \Phi_j$ a collection of mechanisms for all P_j with $j \neq i$.[13] As is standard, we assume that principals can fully commit to their mechanisms and that each principal cannot contract directly over the decisions by other principals.[14]

Timing. The sequence of events is the following.

- At $t = 0$, A learns θ.
- At $t = 1$, each P_i simultaneously and independently offers the agent a mechanism $\phi_i \in \Phi_i$.
- At $t = 2$, A privately sends a message $m_i \in \mathcal{M}_i$ to each P_i after observing the whole array of mechanisms ϕ. The messages $(m_1, ..., m_n)$ are sent simultaneously.[15]
- At $t = 3$, the contracts $(y_1, ..., y_n)$ are determined independently by the lotteries $(\phi_1(m_1), ..., \phi_n(m_n))$.
- At $t = 4$, A chooses $e \in E$ after observing the contracts $(y_1, ..., y_n)$.
- At $t = 5$, the principals' actions $(a_1, ..., a_n)$ are determined independently by the lotteries $(y_1(e), ..., y_n(e))$ and payoffs are realized.

---

all sorts of exogenous restrictions dictated by the environment under examination. What is important to us, is that the set of feasible lotteries $\mathcal{D}_i$ is a primitive of the environment, not a choice of P_i.

[13] We also define $\delta \equiv (\delta_1, ..., \delta_n) \in \mathcal{D} \equiv \prod_{j=1}^{n} \mathcal{D}_j$, $m \equiv (m_1, ..., m_n) \in \mathcal{M} \equiv \prod_{j=1}^{n} \mathcal{M}_j$, $y \equiv (y_1, ..., y_n) \in Y \equiv \prod_{j=1}^{n} Y_j$, $\delta_{-i} \in \mathcal{D}_{-i}$, $m_{-i} \in \mathcal{M}_{-i}$, and $y_{-i} \in Y_{-i}$ in the same way.

[14] As in Bernheim and Whinston (1986), this does not mean that P_i cannot reward the agent as a function of the actions he takes with the other principals. It simply means that P_i cannot make her contract $y_i : E \to \Delta(\mathcal{A}_i)$ contingent on the other principals' contracts y_{-i}, nor her mechanism ϕ_i contingent on the other principals' mechanisms ϕ_{-i}.

[15] As in Peters (2001) and Martimort and Stole (2002), we do not model the agent's participation decisions: these can be easily accommodated by adding to each mechanism a null contract that leads to the default decisions that are implemented in case of no participation such as, for example, no trade at a null price.

Strategies and equilibria. A strategy for P_i is a distribution $\sigma_i \in \Delta(\Phi_i)$ over the set of feasible mechanisms. As for the agent, a strategy $\sigma_A = (\mu, \xi)$ consists of a mapping $\mu : \Theta \times \Phi \to \Delta(\mathcal{M})$ that specifies a distribution over $\mathcal{M}$ for any (θ, ϕ), along with a mapping $\xi : \Theta \times \Phi \times \mathcal{M} \times Y \to \Delta(E)$ that specifies a distribution over effort for any (θ, ϕ, m, y).

Following Peters (2001), we will say that the strategy $\sigma_A = (\mu, \xi)$ constitutes a *continuation equilibrium* for Γ if for every (θ, ϕ, m, y), any $e \in Supp[\xi(\theta, \phi, m, y)]$ maximizes

$$\bar{V}(e; y, \theta) \equiv \int_{\mathcal{A}_1} \cdots \int_{\mathcal{A}_n} v\,(e, a, \theta)\, dy_1(e) \times \cdots \times dy_n(e)$$

and for every (θ, ϕ), any $m \in Supp[\mu(\theta, \phi)]$ maximizes

$$\int_{Y_1} \cdots \int_{Y_n} \max_{e \in E} \bar{V}(e; y, \theta) d\phi_1\,(m_1) \times \cdots \times d\phi_n(m_n).$$

For future reference, we denote by

$$V(\delta, \theta) \equiv \int_{Y_1} \cdots \int_{Y_n} \max_{e \in E} \bar{V}(e; y, \theta) d\delta_1 \times \cdots \times d\delta_n$$

the maximal payoff that A can obtain given the principals' decisions δ.

Denoting by $\rho_{\sigma_A}(\theta, \phi) \in \Delta(\mathcal{A} \times E)$ the distribution over outcomes induced by σ_A given θ and the profile of mechanisms ϕ, we then have that principal i's expected payoff when he chooses the strategy σ_i and the other principals and the agent follow (σ_{-i}, σ_A) is given by

$$U_i(\sigma_i; \sigma_{-i}, \sigma_A) \equiv \int_{\Phi_1} \cdots \int_{\Phi_n} \bar{U}_i(\phi; \sigma_A) d\sigma_1 \times \cdots \times d\sigma_n$$

where

$$\bar{U}_i(\phi; \sigma_A) \equiv \int_{\Theta} \int_{E} \int_{\mathcal{A}} u_i\,(e, a, \theta)\, d\rho_{\sigma_A}(\theta, \phi) dF(\theta).$$

A perfect Bayesian equilibrium for Γ is thus a strategy profile $\sigma \equiv (\{\sigma_i,\}_{i=1}^n, \sigma_A)$ such that σ_A is a continuation equilibrium and for every $i \in \mathcal{N}$,

$$\sigma_i \in \arg \max_{\tilde{\sigma}_i \in \Delta(\Phi_i)} U_i(\tilde{\sigma}_i; \sigma_{-i}, \sigma_A).$$

Throughout, we will denote the set of perfect Bayesian equilibria of Γ by $\mathcal{E}(\Gamma)$ and, for any $\sigma^* \in \mathcal{E}(\Gamma)$, the associated *social choice function* (SCF) by $\pi_{\sigma^*} : \Theta \to \Delta(\mathcal{A} \times E)$.

Menus. A ***menu*** is a mechanism $\phi_i^M : \mathcal{M}_i^M \to \mathcal{D}_i$ whose message space $\mathcal{M}_i^M \subseteq \mathcal{D}_i$ is a subset of all possible decisions and whose mapping is the identity function, i.e. for any $\delta_i \in \mathcal{M}_i^M$, $\phi_i^M(\delta_i) = \delta_i$. In what follows, we denote by Φ_i^M the set of all possible menus for principal i and by Γ^M the "menu game" in which the set of feasible mechanisms for each P_i is Φ_i^M. The game Γ is an ***enlargement*** of Γ^M ($\Gamma \succcurlyeq \Gamma^M$) if for all $i \in \mathcal{N}$, (i) there exists an embedding $\alpha_i : \Phi_i^M \to \Phi_i$;[16]

[16] Formally, an embedding $\alpha_i : \Phi_i^M \to \Phi_i$ can here be thought of as an injective mapping such that, for any pair of mechanisms ϕ_i^M, ϕ_i with $\phi_i = \alpha_i(\phi_i^M)$, $\mathrm{Im}(\phi_i) = \mathrm{Im}(\phi_i^M)$.

and (ii) for any $\phi_i \in \Phi_i$, $\mathrm{Im}(\phi_i)$ is compact. A simple example of an enlargement of Γ^M is a game in which $\Phi_i \supseteq \Phi_i^M$ for all i. More generally, an enlargement is a game in which every Φ_i is larger than Φ_i^M in the sense that each menu ϕ_i^M is also present in Φ_i, although possibly with a different representation. The game in which the principals compete in menus is "focal" in the sense of the following theorem (cfr Peters, 2001, and Martimort and Stole, 2002).

Theorem 1 (Menus) *Let Γ be any enlargement of Γ^M. A SCF π can be sustained as an equilibrium of Γ if and only if it can be sustained as an equilibrium of Γ^M.*

When Γ is not an enlargement of Γ^M, for example because only certain menus can be offered in Γ, there may exist outcomes in Γ that cannot be sustained as equilibrium outcomes in Γ^M and vice-versa. In this case, one can still characterize all equilibrium outcomes of Γ using menus, but it is necessary to restrict the principals to offer only those menus that could have been offered in Γ : that is, the set of feasible menus for P_i must be restricted to $\tilde{\Phi}_i^M \equiv \{\phi_i^M : \mathrm{Im}(\phi_i^M) = \mathrm{Im}(\phi_i)$ for some $\phi_i \in \Phi_i\}$.

In the sequel we will restrict attention to environments in which all menus are feasible. The purpose of our results is to show that, in many applications of interest, one can restrict the principals to offer menus that can be conveniently described as incentive-compatible revelation mechanisms. This in turn may facilitate the characterization of the equilibrium outcomes.

3 Simple revelation mechanisms

Motivated by the arguments discussed in the introduction, in this section we focus on outcomes that can be sustained by pure-strategy profiles in which the agent's strategy is Markov.

Definition 1 *(i) Given the common agency game Γ, an equilibrium strategy profile $\sigma \in \mathcal{E}(\Gamma)$ is a* **pure-strategy equilibrium** *if (a) no principal randomizes over her mechanisms; (b) given any profile of mechanisms $\phi \in \Phi$ and any $\theta \in \Theta$, the agent does not randomize over the messages he sends to the principals.*

(ii) The agent's strategy σ_A is **Markov** *in Γ if and only if, for any $i \in \mathcal{N}$, $\phi_i \in \Phi_i$, $\theta \in \Theta$ and $\delta_{-i} \in \mathcal{D}_{-i}$, there exists a unique $\delta_i(\theta, \delta_{-i}; \phi_i) \in \mathrm{Im}(\phi_i)$ such that A always induces $\delta_i(\theta, \delta_{-i}; \phi_i)$ with P_i when the latter offers the mechanism ϕ_i, the agent's type is θ and the decisions A induces with the other principals are δ_{-i}.*

An equilibrium strategy profile is thus a pure-strategy equilibrium if no principal randomizes over her mechanisms and no type of the agent randomize over the messages to the principals. Note that the agent may however randomize over his choice of effort.

The agent's strategy σ_A in Γ is Markov if and only if the decisions the agent induces in each mechanism ϕ_i depend only on his type θ and the decisions δ_{-i} he is inducing with the other principals—and not on the particular profile of mechanisms (or menus) offered by the latter. As anticipated in the introduction, this definition is different from the one typically considered in dynamic games but it shares with the latter the idea that the agent's behavior should depend only on payoff-relevant information.

Definition 2 *(i) An* **incentive-compatible revelation mechanism** *is a mapping* $\phi_i^r : \mathcal{M}_i^r \to \mathcal{D}_i$*, with message space* $\mathcal{M}_i^r \equiv \Theta \times \mathcal{D}_{-i}$*, such that* $\mathrm{Im}(\phi_i^r)$ *is compact and, for any* $(\theta, \delta_{-i}) \in \Theta \times \mathcal{D}_{-i}$*,*

$$\phi_i^r(\theta, \delta_{-i}) \in \arg \max_{\delta_i \in \mathrm{Im}(\phi_i^r)} V(\delta_i, \delta_{-i}, \theta).$$

(ii) A **revelation game** Γ^r *is a game in which each principal's strategy space is* $\Delta(\Phi_i^r)$*, where* Φ_i^r *is the set of all revelation mechanisms for principal* i*.*

(iii) Given a profile of mechanisms $\phi^r \in \Phi^r$*, the agent's strategy is* **truthful** *in* ϕ_i^r *if, for any* $\theta \in \Theta$ *and any* $(m_i^r, m_{-i}^r) \in Supp[\mu(\theta, \phi_i^r, \phi_{-i}^r)]$*,*

$$m_i^r = (\theta, (\phi_j^r(m_j^r))_{j \neq i}).$$

(iv) An equilibrium strategy profile $\sigma^{r*} \in \mathcal{E}(\Gamma^r)$ *is a* **truthful equilibrium** *if, given any profile of mechanisms* ϕ^r *such that* $|\{j \in \mathcal{N} : \phi_j^r \notin Supp[\sigma_j^{r*}]\}| \leq 1$*, the agent's strategy is truthful in every* $\phi_i^r \in Supp[\sigma_i^{r*}]$*.*

In a revelation mechanism, the agent is thus asked to report his type θ along with the decisions δ_{-i} he is inducing with the other principals. Given a profile of mechanisms ϕ^r, the agent's strategy is truthful in ϕ_i^r if the message $m_i^r = (\theta, \delta_{-i})$ the agent reports to P_i coincides with his true type θ along with the true decisions $\delta_{-i} = \left(\phi_j(m_j)\right)_{j \neq i}$ that the agent induces with all principals other than i by sending the messages $m_{-i} \equiv (m_j)_{j \neq i}$. An equilibrium strategy profile is then said to be a truthful equilibrium if, whenever at most one principal deviated from equilibrium play, the agent reports truthfully to any of the non-deviating principals.

The following is our first characterization result.

Theorem 2 *Suppose the SCF* π *can be sustained as a pure-strategy equilibrium of* Γ^M *in which the agent's strategy is Markov. Then it can also be sustained as a truthful pure-strategy equilibrium of* Γ^r*. Furthermore, any SCF* π *that can be sustained as an equilibrium of* Γ^r *can also be sustained as an equilibrium of* Γ^M*.*

First, consider the "only if" part of the result. When the agent's choice from each menu depends only on his type θ and the decisions δ_{-i} he is inducing with the other principals, it is immediate that each principal can be restricted to offer the decisions $\delta_i(\theta, \delta_{-i}; \phi_i^{M*})$ that the agent would have selected from the equilibrium menu ϕ_i^{M*} for some (θ, δ_{-i}). Describing the menu of such decisions as a revelation mechanisms is then a convenient way of specifying which decisions the agent takes in response to each (θ, δ_{-i}). As illustrated in the next section, this often facilitates the characterization of the equilibrium allocations.

Next, consider the "if" part of the result. Despite the fact that Γ^r is not an enlargement of Γ^M, the result follows from arguments similar to those used to establish the Menu Theorem. The equilibrium σ^{M*} that sustains π in Γ^M features each principal offering the menus in the range of the equilibrium direct mechanism in Γ^r. When all principals offer the equilibrium menus, the agent implements the same decisions he would have implemented in Γ^r. When, instead, one principal, say P_i, deviates and offers a menu $\phi_i^M \notin Supp[\sigma_i^{M*}]$, the agent implements the same decisions he would have implemented in Γ^r had P_i offered the direct mechanism ϕ_i^r such that

$$\phi_i^r(\theta, \delta_{-i}) \in \arg \max_{\delta_i \in \mathrm{Im}(\phi_i^M)} V(\delta_i, \delta_{-i}, \theta) \qquad \forall (\theta, \delta_{-i}) \in \Theta \times \mathcal{D}_{-i}.$$

The behavior prescribed by the strategy σ_A^{M*} constructed this way is clearly rational for the agent in Γ^M. Furthermore, given σ_A^{M*}, no principal has an incentive to deviate.

Although in most applications restricting attention to Markov strategies seems perfectly reasonable, it is interesting to examine whether there exist special environments in which such a restriction is without any loss of generality. To address this question, we first need to introduce some notation. For any $k \in \mathcal{N}$, and any (δ, θ), let

$$\underline{U}_k(\delta, \theta) \equiv \int_Y \left[\int_{\mathcal{A}} u_k(a, \xi_k(\theta, y), \theta)\, dy_1(\xi_k(\theta, y)) \times \cdots \times dy_n(\xi_k(\theta, y)) \right] d\delta_1 \times \cdots \times d\delta_n \tag{1}$$

denote the minimal payoff for principal k that is compatible with the agent's rationality, where

$$\xi_k(\theta, y) \in \arg \min_{e \in E^*(\theta, y)} \left\{ \int_{\mathcal{A}} u_k(a, e, \theta)\, dy_1(e) \times \cdots \times dy_n(e) \right\} \tag{2}$$

and

$$E^*(\theta, y) \equiv \arg \max_{e \in E} \bar{V}(e; y, \theta).$$

Condition 1 (Uniform Punishment) *We say that the "Uniform Punishment" condition holds if for any $i \in \mathcal{N}$, $B \subseteq \mathcal{D}_i$, $\delta_{-i} \in \mathcal{D}_{-i}$, and $\theta \in \Theta$, there exists a $\delta_i' \in \arg\max_{\delta_i \in B} V(\delta_i, \delta_{-i}, \theta)$ such that for all $j \neq i$ and all $\hat{\delta}_i \in \arg\max_{\delta_i \in B} V(\delta_i, \delta_{-i}, \theta)$,*

$$\underline{U}_j(\delta_i', \delta_{-i}, \theta) \leq \underline{U}_j(\hat{\delta}_i, \delta_{-i}, \theta).$$

The condition says that the principals' preferences are sufficiently aligned in the sense that, given any menu of decisions $B \subseteq \mathcal{D}_i$ offered by P_i and any (θ, δ_{-i}), there exists a decision $\delta_i' \in B$ that is optimal for the agent given (θ, δ_{-i}) such that the payoff of any principal P_j, $j \neq i$, under δ_i' is lower than under any other decision $\delta_i \in B$ that is optimal for the agent.

We then have the following result.

Theorem 3 *Suppose one of the following holds:*

(a) for any $i \in \mathcal{N}$, $B \subseteq \mathcal{D}_i$, and $(\theta, \delta_{-i}) \in \Theta \times \mathcal{D}_{-i}$, $|\arg\max_{\delta_i \in B} V(\delta_i, \delta_{-i}, \theta)| = 1$;

(b) $|\Theta| = 1$ and the "Uniform Punishment" condition holds.

Then any SCF that can be sustained as a pure-strategy equilibrium of Γ^M can also be sustained as a pure-strategy equilibrium in which the agent's strategy is Markov.

Condition (a) says that the agent's preferences are "single-peaked" in the sense that, for any $(\theta, \delta_{-i}) \in \Theta \times \mathcal{D}_{-i}$ and any menu of decisions $B \subseteq \mathcal{D}_i$, there is a single decision in B that maximizes the agent's payoff. Clearly in this case the agent's strategy is necessarily Markov.

Condition (b) says that information is complete and that the principals' payoffs are sufficiently aligned in the sense of the Uniform Punishment condition. The role of this condition is to guarantee that, given δ_{-i}, the agent can punish any principal P_j, $j \neq i$, by taking the same decision with principal i. Note that this condition is satisfied, for example, when the agent is a manufacturer and the principals are retailers competing a' la Cournot in a downstream market; in fact, in this case

$$u_i = f(q_i + \sum_{k \neq i} q_k) q_i - t_i$$

where q_i denotes the quantity purchased by P_i, t_i the total payment made by P_i to the manufacturer, and $f : \mathbb{R}_+ \to \mathbb{R}$ the inverse demand function. In this environment, $|\Theta| = |E| = 1$. A (deterministic) contract δ_i is thus a degenerate lottery that assigns measure one to a single price-quantity pair $(t_i, q_i) \in \mathbb{R} \times \mathbb{R}_+$. It is then immediate that, given any menu $B \subseteq \mathbb{R} \times \mathbb{R}_+$ (i.e. any array of price-quantity pairs) offered by P_i, and any profile of contracts $(t_{-i}, q_{-i}) \in \mathbb{R}^{n-1} \times \mathbb{R}_+^{n-1}$ selected by the agent with the other principals, the contract $(t_i, q_i) \in B$ that minimizes P_j's payoff among those that are optimal for the agent given (t_{-i}, q_{-i}) is the one that entails the highest quantity q_i, and this is true for any P_j, $j \neq i$. The Uniform Punishment condition thus clearly holds in this environment.

The reason why one needs information to be complete in addition to enough alignment in the principals' preferences can be illustrated through the following example where $n = 2$ in which case the Uniform Punishment condition trivially holds. The sets of primitive decisions are $\mathcal{A}_1 = \{t, b\}$ and $\mathcal{A}_2 = \{l, r\}$. There is no effort so that a deterministic contract coincides with the choice of a decision $a_i \in \mathcal{A}_i$. There are two types of the agent, $\underline{\theta}$ and $\overline{\theta}$. The principals' common prior is that $\Pr(\theta = \overline{\theta}) = p < 1/5$. Payoffs, (u_1, u_2, v) are as in the following table:

$\theta = \underline{\theta}$

$a_1 \backslash a_2$	l			r		
t	2	1	1	2	0	0
b	1	0	1	1	2	2

$\theta = \overline{\theta}$

$a_1 \backslash a_2$	l			r		
t	2	2	2	-2	0	2
b	1	0	1	-2	1	1

Table 1

Consider the following (deterministic) SCF: if $\theta = \underline{\theta}$, then $a_1 = b$ and $a_2 = r$; if $\theta = \overline{\theta}$, then $a_1 = t$ and $a_2 = l$. This SCF can be sustained as a (pure-strategy) equilibrium of the menu game in which the agent's strategy is non-Markov. The equilibrium features P_1 offering the menu $\phi_1^{M*} = \{t, b\}$ and P_2 offering the menu $\phi_2^{M*} = \{l, r\}$. Clearly P_2 does not have profitable deviations because she is getting in each state her maximal feasible payoff. If P_1 deviates and offers $\{t\}$ then A selects (t, l) if $\theta = \underline{\theta}$ and (t, r) if $\theta = \overline{\theta}$ (given $(\underline{\theta}, t)$, A has strict preferences for l, whereas given $(\overline{\theta}, t)$, he is indifferent between l and r). A deviation to $\{t\}$ thus yields a payoff $U_1 = 2(1-p) - 2p = 2 - 4p$ to P_1 that is lower than her equilibrium payoff $U_1^* = 1 + p$ when $p > 1/5$. A deviation to $\{b\}$ is clearly never profitable for P_1, irrespective of the agent's behavior. Thus the SCF π^* described above can be sustained in equilibrium.

Now to see that this SCF cannot be sustained by restricting the agent's strategy to be Markov, first note that it is essential that ϕ_2^{M*} contains both l and r because in equilibrium A must choose different a_2 for different θ. Restricting the agent's strategy to be Markov then means that when P_2 offers the equilibrium menu, A necessarily chooses r if $\theta = \underline{\theta}$ and $a_1 = b$ and l if $\theta = \overline{\theta}$ and $a_1 = t$. Furthermore, because given $(\underline{\theta}, t)$, A strictly prefers l to r, A necessarily chooses l when $\theta = \underline{\theta}$ and $a_1 = t$. Given this behavior, if P_1 deviates and offers the menu $\phi_1^M = \{t\}$, she then gets a payoff $U_1 = 2(1-p) + 2p = 2 > U_1^*$.

The reason why, when information is incomplete, restricting the agent's strategy to be Markov may preclude the possibility of sustaining certain SCFs is that Markov strategies do not permit the same type of the agent, say θ', to punish a deviation by a principal P_j, $j \neq i$, by choosing with all principals other than i the equilibrium decisions $\delta_{-i}^*(\theta')$ and then choosing with P_i a decision $\delta_i \neq \delta_i^*(\theta')$ to punish P_j. Allowing type θ' to change his behavior in response to the equilibrium decisions $\delta_{-i}^*(\theta')$ may thus be essential to punish certain deviations. However, because this is the only reason why one needs $|\Theta| = 1$, the assumption of complete information can be dispensed with if one imposes the following refinement on the agent's continuation equilibrium.

Condition 2 (Conformity to Equilibrium) *Let Γ be any simultaneous common agency game. Given any pure-strategy equilibrium $\sigma^* \in \mathcal{E}(\Gamma)$, let ϕ^* denote the equilibrium mechanisms and $\delta^*(\theta)$ the equilibrium decisions implemented when the agent's type is θ. We say that the agent's strategy in σ^* satisfies the "Conformity to Equilibrium" condition if, for any i, θ, ϕ_{-i} and $m \in Supp[\mu(\theta, \phi_i^*, \phi_{-i})]$,*

$$(\phi_j(m_j))_{j \neq i} = \delta_{-i}^*(\theta) \Longrightarrow \phi_i^*(m_i) = \delta_i^*(\theta).$$

In words, the agent's strategy satisfies the Conformity to Equilibrium condition if each type of the agent induces the equilibrium decision $\delta_i^*(\theta)$ with each principal P_i when the latter offers the equilibrium mechanism ϕ_i^* and the agent induces the equilibrium decisions $\delta_{-i}^*(\theta)$ with the other principals. In many applications, this seems a mild requirement. We then have the following result.

Theorem 4 *Suppose the principals' payoffs are sufficiently aligned in the sense of the Uniform Punishment condition. Suppose in addition that the SCF π can be sustained as a pure-strategy equilibrium $\sigma^{M*} \in \mathcal{E}(\Gamma^M)$ in which the agent's strategy σ_A^{M*} satisfies the "Conformity to Equilibrium" condition. Then, π can also be sustained as a pure-strategy equilibrium $\tilde{\sigma}^{M*} \in \mathcal{E}(\Gamma^M)$ in which the agent's strategy $\tilde{\sigma}_A^{M*}$ is Markov.*

By implication, when the principals' payoffs are sufficiently aligned (e.g. when $n = 2$), even if $|\Theta| > 1$, any SCF π that can be sustained as a pure-strategy equilibrium of Γ^M in which the agent's strategy satisfies the "Conformity to Equilibrium" condition, can also be sustained as a truthful pure-strategy equilibrium of the revelation game Γ^r.

At this point, it is useful to contrast our results with those in Peters (2003, 2007) and Attar, Majumdar, Piasier, and Porteiro (2007). Peters (2003, 2007) considers environments in which a certain "no-externality condition" holds and shows that in these environments all *pure-strategy* equilibria can be characterized by restricting the principals to offer standard direct revelation mechanisms $\phi_i : \Theta \longrightarrow \mathcal{D}_i$.[17] The no-externality condition requires that (i) each principal's payoff be independent of the other principals' actions a_{-i} and (ii) that conditional on choosing effort in a certain equivalence class $\hat{E}$,[18] the agent's preferences over any set of decisions $B \subseteq \mathcal{A}_i$ by principal i be independent of the particular effort the agent chooses in $\hat{E}$, of his type θ, and of the actions a_{-i} he induces with the other principals. Attar, Majumdar, Piasier, and Porteiro (2007) show that in environments in which the only feasible contracts are deterministic, all action spaces are finite, and the agent's preferences are "separable" and "generic", condition (i) in Peters can be dispensed with: all equilibrium outcomes of the menu game (including mixed-strategy ones) can be sustained also in the game in which the principals' strategy space consists of all standard direct revelation mechanisms. Separability requires that the agent's preferences over the decisions of any of the principals be independent of his choice of effort and of the decisions of the other principals. Genericity requires that the agent never be indifferent between any pair of effort choices and/or any pair of decisions by any of the principals.[19] Combined these restrictions guarantee that the

[17] A standard direct revelation mechanism reduces to a take-it-or-leave-it-offers–i.e. to a degenerate menu consisting of a single contract $y_i : E \longrightarrow \Delta(Y_i)$–when the agent does not possess any exogenous private information, i.e. when $|\Theta| = 1$.

[18] In the language of Peters, an equivalence class $\hat{E} \subseteq E$ is a subset of E such that any feasible contract of P_i must respond to each $e, e' \in \hat{E}$ with the same action, i.e. $y_i(e) = y_i(e')$ for any $e, e' \in \hat{E}$.

[19] Formally, separability requires that any type θ of the agent who strictly prefers a_i to a_i' when the decisions by all principals other than i are a_{-i} and his choice of effort is e also strictly prefers a_i to a_i' when the decisions taken by all principals other than i are a_{-i}' and his choice of effort is e', for any $(a_{-i}, e), (a_{-i}', e') \in \mathcal{A}_{-i} \times E$. Genericity requires that , given any $(\theta, a_i) \in \Theta \times \mathcal{A}_i$, $v(\theta, a_i, a_{-i}, e) \neq v(\theta, a_i, a_{-i}', e')$ for any $(e, a_{-i}), (e, a_{-i}) \in E \times \mathcal{A}_{-i}$ with $(e, a_{-i}) \neq (e, a_{-i})$. Note that in general separability is neither weaker nor stronger than condition (ii) in Peters (2003, 2007). In fact, separability requires the agent's preferences over P_i's actions to be independent of e, whereas

messages each type of the agent sends to any of his principals do not depend on the messages he sends to the other principals; restricting attention to standard direct revelation mechanisms is then clearly without loss.

Compared to these results, our result in Theorem 2 does not require any restriction on the players' preferences. Provided one is willing to restrict attention to equilibria in which the agent's strategy is Markov, then all pure-strategy equilibrium outcomes can be characterized through a simple generalization of the class of standard direct revelation mechanisms in which the agent reports the decisions δ_{-i} in addition to his type θ. Because in most applications of interest, assuming the agent strategy is Markov is appealing, Theorem 2 thus provides a possible route to equilibrium characterization that does not require any restriction on the players' preferences. Theorem 3 in turn guarantees that restricting attention to Markov strategies is not only appealing but actually without any loss of generality when either the agent's preferences are single-peaked or information is complete and the principals' preferences are sufficiently aligned in the sense of the Uniform Punishment condition.

Our results are thus complementary to those in Peters (2003, 2007) and Attar, Majumdar, Piasier, and Porteiro (2007) in the sense that they are particularly useful precisely in environments in which one cannot restrict attention neither to simple take-it-or-leave-it offers nor standard direct revelation mechanisms. For example, consider a pure adverse selection setting with only deterministic contracts,[20] as in the baseline model of Attar, Majumdar, Piasier, and Porteiro (2007). Then condition (a) in Theorem 3 is equivalent to the "genericity" condition in their paper. If, in addition, preferences are separable (in the sense described above), then Theorem 1 in their paper implies that all equilibrium outcomes can be sustained by restricting the principals to offer standard direct revelation mechanisms. If, instead, they are not separable,[21] then all equilibrium outcomes can still be characterized restricting the principals to offer direct revelation mechanisms but the latter must be extended to allow the agent to report the decisions δ_{-i} he is inducing with the other principals in addition to his type θ.

Also note that, when action spaces are continuous, as typically assumed in many applied papers, Attar, Majumdar, Piasier, and Porteiro (2007) need to impose a restriction on the agent's behavior. This restriction, which they call "conservative behavior" consists in requiring that, after

condition (ii) in Peters only requires them to be independent of the particular effort the agent chooses in a given equivalence class. On the other hand, condition (ii) in Peters imposes that the agent's preferences over P_i's actions be independent of the agent's type, whereas such a dependence is allowed by separability. The two conditions are however equivalent in standard moral hazard settings (i.e. when effort is completely unobservable so that $\hat{E} = E$ and information is complete so that $|\Theta| = 1$).

[20] A pure adverse selction setting with only deterministic contracts is an environment in which $|E| = 1$ and where, for all i, $\mathcal{D}_i$ contains only degenerate lotteries that assign measure one to one of the elements of $\mathcal{A}_i$.

[21] As, for example, in the case of a buyer whose preferences for the quality/quantity of the product/service of seller i depend on the quality/quantity purchased from seller j.

a deviation by P_k, each type θ of the agent continues to choose the equilibrium decisions $\delta^*_{-k}(\theta)$ with the non-deviating principals whenever this is compatible with the agent's rationality. This restriction is stronger than the "Conformity to Equilibrium" condition introduced above. Hence, even with separable preferences, the more general revelation mechanisms introduced here may turn useful in applications in which imposing the "conservative behavior" property seems too restrictive.

4 Using revelation mechanisms in applications

Equipped with the results established in the preceding session, we now show how revelation mechanisms can be put to work in applications to identify necessary and sufficient conditions for the sustainability of outcomes as common agency equilibria. We consider three cases of interest: competition in non-linear tariffs, menu auctions, and a (simplified version of a standard) moral hazard setting.

4.1 Competition in non-linear tariffs

Consider an environment in which P_1 and P_2 are two sellers providing two differentiated products to a common buyer, A. In this environment, there is no effort and an action for principal i consists of a price-quantity pair $(t_i, q_i) \in \mathcal{A}_i \equiv \mathbb{R} \times \mathcal{Q}$, where $\mathcal{Q} = [0, \bar{Q}]$ denotes the set of feasible quantities.[22]

The buyer's payoff is given by $v(a, \theta) = \theta(q_1 + q_2) + \lambda q_1 q_2 - t_1 - t_2$, where λ parametrizes the degree of complementarity/substitutability between the two products, while θ denotes the buyer's type which is assumed to be distributed over the interval $\Theta = [\underline{\theta}, \bar{\theta}]$, $\underline{\theta} > 0$, with log-concave density f strictly positive for any θ. The sellers' payoffs are given by $u_i(a, \theta) = t_i - C(q_i)$, with $C(q) = q^2/2$.

The buyer can participate in one mechanism without participating in the other (in the literature this is referred to as non-intrinsic common agency). In the case A decides not to participate in P_i's mechanism, the default contract $(0, 0)$ with no trade and zero transfer is implemented.

Following the pertinent literature, we assume here that only deterministic mechanisms $\phi_i : \mathcal{M}_i \longrightarrow \mathcal{A}_i$ are feasible. Note that any such mechanism is strategically-equivalent to a *non-linear tariff* T_i such that, for any q_i, $T(q_i) = \min\{t_i : (t_i, q_i) \in \mathrm{Im}(\phi_i)\}$ if $\{t_i : (t_i, q_i) \in \mathrm{Im}(\phi_i)\} \neq \varnothing$ and $T(q_i) = \infty$ otherwise. It is also immediate that any such tariff is equivalent to a menu of price-quantity pairs (see also Peters, 2001, 2004). The question of interest is which tariffs will be offered in equilibrium and how they can be conveniently characterized using incentive-compatibility.

[22] An alternative way of modelling this environment is the following. The set of primitive actions for each principal i consists of the set $\mathbb{R}$ of all possible prices. A contract for P_i then consists of a tariff $y_i : \mathcal{Q} \longrightarrow \mathbb{R}$ that specifies a price for each possible quantity $q \in \mathcal{Q}$. Given a pair of tariffs $y = (y_1, y_2)$, the agent's effort then consists of the choice of a pair of quantities $e = (q_1, q_2) \in E = \mathcal{Q}^2$. While the two approaches ultimately lead to the same results, we find the one proposed in the text more parsimonious.

Following the discussion in the previous sections, we focus on pure-strategy equilibria in which the agent's strategy is Markov.

First, we show how revelation mechanisms help identify necessary and sufficient conditions for the implementability of schedules $q_i^* : \Theta \longrightarrow \mathcal{Q}$. Next, we show how these conditions can be used to prove that there is no equilibrium that sustains the schedules $q^c : \Theta \longrightarrow \mathcal{Q}$ that maximize the sellers' joint payoffs (these schedules are referred to in the literature as the "collusive schedules"). Last, we identify sufficient conditions for the sustainability of differentiable schedules.

Necessary and sufficient conditions for the implementability of schedules.

By Theorem 2, the schedules $q_i^*(\cdot)$, $i = 1, 2$, can be sustained as a pure-strategy equilibrium of Γ^M in which the agent's strategy is Markov *if and only if* they can be sustained as a pure-strategy truthful equilibrium of Γ^r. Now let

$$m_i(\theta) \equiv \theta + \lambda q_j^*(\theta)$$

denote type θ's *marginal valuation* for q_i when he purchases the equilibrium quantity $q_j^*(\theta)$ from P_j, $j \neq i$. In what follows we restrict attention to schedules $q^*(\cdot) = (q_i^*(\cdot))_{i=1,2}$ for which the functions $m_i(\cdot)$ are strictly increasing, $i = 1, 2$. Because for any (compact) collection of price-quantity pairs $B \subseteq \mathcal{A}_i$ and any pair (θ, q_j, t_j) and (θ', q_j', t_j') such that $\theta + \lambda q_j = \theta' + \lambda q_j'$

$$\arg \max_{(q_i,t_i) \in B} v(\theta, q_j, t_j, q_i, t_i) = \arg \max_{(q_i,t_i) \in B} v(\theta', q_j', t_j', q_i, t_i),$$

and because there are no direct externalities between the two principals, it is immediate that it suffices to consider revelation mechanisms with the property that $\phi_i^r(\theta, q_j, t_j) = \phi_i^r(\theta', q_j', t_j')$ whenever $\theta + \lambda q_j = \theta' + \lambda q_j'$. In the sequel we thus restrict attention to such mechanism which, with a slight abuse of notation, we denote by $\phi_i^r = (\tilde{q}_i(\tilde{\theta}_i), \tilde{t}_i(\tilde{\theta}_i))_{\tilde{\theta}_i \in \tilde{\Theta}_i}$, where

$$\tilde{\Theta}_i \equiv \{\tilde{\theta}_i \in \mathbb{R} : \tilde{\theta}_i = \theta + \lambda q_i, \ \theta \in \Theta, \ q_i \in \mathcal{Q}\}$$

denotes the set of the agent's possible marginal valuations for P_i's quantity. Note that these mechanisms specify price-quantity pairs also for $\tilde{\theta}_i$ that have zero measure on the equilibrium path. As discussed in the literature, sellers may need to include in their menus also allocations that are selected only off equilibrium to punish deviations by other sellers.[23] These allocations are typically obtained by extending the principals' tariffs outside the equilibrium range. Identifying the appropriate extensions can however be quite complicated. One of the advantages of the approach suggested here is that it permits one to use incentive-compatibility to describe such extensions.

Now note that, because the set $\tilde{\Theta}_i$ is an interval and because the function $\tilde{v}(\tilde{\theta}, q) \equiv \tilde{\theta} q$ is equi-Lipschitz continuous and differentiable in $\tilde{\theta}$ and satisfies the increasing-difference property, from standard results in mechanism design (see e.g. Milgrom and Segal, 2002), the mechanism

[23] Such allocations are also referred to as "latent contracts" (see, e.g. Piasier, 2007).

$\phi_i^r = (\tilde{q}_i(\cdot), \tilde{t}_i(\cdot))$ is incentive-compatible if and only if the function $\tilde{q}_i(\cdot)$ is non-decreasing and the function $\tilde{t}_i(\cdot)$ satisfies

$$\tilde{t}_i(\tilde{\theta}_i) = \tilde{\theta}_i \tilde{q}_i(\tilde{\theta}_i) - \int_{\min \tilde{\Theta}_i}^{\tilde{\theta}_i} \tilde{q}_i(s) ds - K_i \quad \forall \tilde{\theta}_i \in \tilde{\Theta}_i,$$

where K_i is a constant. Next note that for any pair of mechanisms $(\phi_i^r)_{i=1,2}$ for which there exists an $i \in \mathcal{N}$ and a $\tilde{\theta}_i \in \tilde{\Theta}_i$ such that an agent with marginal valuation $\tilde{\theta}_i$ strictly prefers the null contract $(0,0)$ to the contract $(\tilde{q}_i(\tilde{\theta}_i), \tilde{t}_i(\tilde{\theta}_i))$, there exists another pair of mechanisms $(\phi_i^{r\prime})_{i=1,2}$ such that (i) any $\tilde{\theta}_i \in \tilde{\Theta}_i$ weakly prefers the contract $(\tilde{q}_i(\tilde{\theta}_i), \tilde{t}_i(\tilde{\theta}_i))$ to the contract $(0,0)$, $i = 1,2$, and (ii) $(\phi_i^{r\prime})_{i=1,2}$ sustains the same outcomes as $(\phi_i^r)_{i=1,2}$.[24] It is thus without loss of generality to restrict $K_i \geq 0$. We then have the following result.

Proposition 1 *The schedules $q_i^*(\cdot)$, $i = 1,2$, can be sustained as a pure-strategy equilibrium of Γ^M in which the agent's strategy is Markov if and only if there exist non-decreasing functions $\tilde{q}_i : \tilde{\Theta}_i \to \mathcal{Q}$ and scalars $\tilde{K}_i \geq 0$, $i = 1,2$, such that the following conditions hold:*

(a) for any $\tilde{\theta}_i \in [m_i(\underline{\theta}), m_i(\bar{\theta})]$, $\tilde{q}_i(\tilde{\theta}_i) = q_i^*(m_i^{-1}(\tilde{\theta}_i))$, $i = 1,2$;[25]

(b) for any $\theta \in \Theta$ and any pair $(\tilde{\theta}_1, \tilde{\theta}_2) \in \tilde{\Theta}_1 \times \tilde{\Theta}_2$

$$V^*(\theta) = \sup_{\tilde{\theta}_1, \tilde{\theta}_2} \left\{ \theta \left[\tilde{q}_1(\tilde{\theta}_1) + \tilde{q}_2(\tilde{\theta}_2) \right] + \lambda \tilde{q}_1(\tilde{\theta}_1) \tilde{q}_2(\tilde{\theta}_2) - \tilde{t}_1(\tilde{\theta}_1) - \tilde{t}_2(\tilde{\theta}_2) \right\}$$

where $V^*(\theta) \equiv \theta \left[q_1^*(\theta) + q_2^*(\theta) \right] + \lambda q_1^*(\theta) q_2^*(\theta) - \tilde{t}_1(m_1(\theta_1)) - \tilde{t}_2(m_2(\theta_2))$ and, for $i = 1,2$,

$$\tilde{t}_i(\tilde{\theta}_i) \equiv \tilde{\theta}_i \tilde{q}(\tilde{\theta}_i) - \int_{\min \tilde{\Theta}_i}^{\tilde{\theta}_i} \tilde{q}_i(s) ds - \tilde{K}_i; \tag{3}$$

(c) each principal i's equilibrium payoff satisfies

$$U_i^* \equiv \int_{\underline{\theta}}^{\bar{\theta}} \left[\tilde{t}_i(m_i(\theta)) - \tfrac{q_i^*(\theta)^2}{2} \right] dF(\theta) = \bar{U}_i \tag{4}$$

where $\bar{U}_i$ is the value of the following program

$$\max_{q_i(\cdot), t_i(\cdot)} \quad \textstyle\int_{\underline{\theta}}^{\bar{\theta}} [t_i(\theta) - \tfrac{q_i(\theta)^2}{2}] dF(\theta)$$

s.t.

$$\theta q_i(\theta) + v_i^*(\theta, q_i(\theta)) - t_i(\theta) \geq \theta q_i(\hat{\theta}) + v_i^*(\theta, q_i(\hat{\theta})) - t_i(\hat{\theta}) \ \forall (\theta, \hat{\theta}) \qquad (IC)$$

$$\theta q_i(\theta) + v_i^*(\theta, q_i(\theta)) - t_i(\theta) \geq v_i^*(\theta, 0) \ \forall \theta \qquad (IR)$$

where, for any $(\theta, q) \in \Theta \times \mathcal{Q}$, $v_i^*(\theta, q) \equiv (\theta + \lambda q)\, \tilde{q}_j(\theta + \lambda q) - \tilde{t}_j(\theta + \lambda q) = \int_{\min \tilde{\Theta}_j}^{\theta + \lambda q} \tilde{q}_j(s) ds + \tilde{K}_j$, $j \neq i$.

[24] The result follows from replication arguments similar to those that lead to Theorem 2.

[25] This condition also implies that $q_i^*(\cdot)$ are nondecreasing, $i = 1,2$.

Condition (a) guarantees that, on the equilibrium path, the mechanism ϕ_i^{r*} assigns to each θ the equilibrium quantity $q_i^*(\theta)$. Condition (b) guarantees that each type θ finds it optimal to truthfully report to each principal the marginal valuation $m_i(\theta) = \theta + \lambda q_j^*(\theta)$. That each θ also finds it optimal to participate follows from the fact that $\tilde{K}_i \geq 0$. Finally, condition (c) guarantees that no principal has a profitable deviation. Instead of specifying a reaction by the agent to any possible pair of mechanisms and then checking that, given this reaction and the mechanism offered by the other principal, no P_i has a profitable deviation, the program in condition (c) gives directly the maximal payoff that each P_i can obtain given the opponent's mechanism without violating the agent's rationality. To compute the payoff $\bar{U}_i$, the program in (c) uses the standard revelation principle, taking into account that the value that each type θ assigns to q_i is $\theta q_i + v_i^*(\theta, q_i)$, where $v_i^*(\theta, q)$ denotes the maximal payoff that the agent can obtain with P_j when his type is θ and the quantity purchased from P_i is q_i. Note that, in general, this approach is not correct: in fact, it presumes that, when indifferent, the agent follows the recommendations by P_i. Permitting the agent to deviate from P_i's recommendations can however be essential to sustain certain SCFs, for it makes fewer deviations profitable. The reason why, in this particular environment, P_i can guarantee herself the payoff $\bar{U}_i$ is twofold: (i) she is not personally interested in the decisions the agent takes with P_j and (ii) the agent's payoff for (q_i, t_i) is quasilinear and has the increasing-difference property with respect to (θ, q_i). As we show in the appendix, this implies that, given the mechanism $\phi_j^{r*} = (\tilde{q}_j(\cdot), \tilde{t}_j(\cdot))$ offered by P_j, there always exists an incentive-compatible mechanism $\phi_i^r = (\tilde{q}_i(\cdot), \tilde{t}_i(\cdot))$ such that, given (ϕ_j^{r*}, ϕ_i^r), *any* sequentially rational strategy σ_A^r for the agent yields P_i a payoff arbitrarily close to $\bar{U}_i$. This explains why condition (c) is not only sufficient but also necessary.

When $\lambda > 0$ and the function $v_i^*(\theta, q)$ is differentiable in θ (which is the case for example when the schedule $\tilde{q}_j(\cdot)$ is continuous), the program in condition (c) has a simple solution. The fact that the mechanism $\phi_j^{*r} = (\tilde{q}_j(\cdot), \tilde{t}_j(\cdot))$ is incentive-compatible implies that the function $g_i(\theta, q) \equiv \theta q + v_i^*(\theta, q) - v_i^*(\theta, 0)$ is equi-Lipschitz continuous and differentiable in θ, it satisfies the increasing-difference property, and is increasing in θ. It follows that a pair of functions $q_i : \Theta \to \mathcal{Q}$, $t_i : \Theta \to \mathbb{R}$ satisfies (IC) and (IR) if and only if $q_i(\cdot)$ is nondecreasing and, for any $\theta \in \Theta$,

$$t_i(\theta) = \theta q_i(\theta) + [v_i^*(\theta, q_i(\theta)) - v_i^*(\theta, 0)] - \int_{\underline{\theta}}^{\theta} [q_i(s) + \tilde{q}_j(s + \lambda q_i(s)) - \tilde{q}_j(s)]ds - K_i, \tag{5}$$

with $K_i \geq 0$. The program in condition (c) then reduces to

$$\max_{q_i(\cdot), K_i} \int_{\underline{\theta}}^{\bar{\theta}} h_i(q_i(\theta); \theta) dF(\theta) - K_i \tag{6}$$

$$\text{s.t. } K_i \geq 0 \text{ and } q_i(\cdot) \text{ is nondecreasing}$$

where

$$h_i(q; \theta) \equiv \theta q + [v_i^*(\theta, q) - v_i^*(\theta, 0)] - \tfrac{q^2}{2} - \tfrac{1-F(\theta)}{f(\theta)} [q + \tilde{q}_j(\theta + \lambda q) - \tilde{q}_j(\theta)] \tag{7}$$

with

$$v_i^*(\theta, q) - v_i^*(\theta, 0) = \int_{\theta}^{\theta+\lambda q} \tilde{q}_j(s)ds.$$

Equipped with these tools, one can then establish for example the following couple of results.

Non-implementability of the collusive schedules.

It has long been noted that when two products are complements ($\lambda > 0$), it may be impossible to sustain the collusive schedules $q^c(\theta)$ as a non-cooperative equilibrium.[26] However, this result has been established restricting the principals to offer twice continuously differentiable tariffs, thus leaving open the possibility that it is merely a consequence of a technical assumption.[27]

The approach suggested here permits one to verify that this result is true more generally.

Proposition 2 *Suppose $\lambda > 0$. There exists no equilibrium in which the agent's strategy is Markov that sustains the collusive schedules*

$$q_i(\theta) = q^c(\theta)\ \forall\theta,\ \ i = 1, 2.$$

The proof uses the characterization of Proposition 1. By relying only on incentive-compatibility, it guarantees that the aforementioned impossibility result is *by no means* a consequence of the assumptions one makes on the differentiability of the tariffs, or on the way one extends the tariffs outside the equilibrium range.

Sufficient conditions for the implementability of differentiable schedules.

We conclude by showing how the conditions in Proposition 1 specialize in the case of differentiable schedules and can be used to construct equilibria.

Proposition 3 *Let $q^* : \Theta \longrightarrow \mathcal{Q}$ be a non-decreasing function satisfying the following differential equation*

$$\lambda \left[q^*(\theta)(1-\lambda) - \theta + 2\left(\tfrac{1-F(\theta)}{f(\theta)}\right)\right] \frac{dq^*(\theta)}{d\theta} = \theta - \tfrac{1-F(\theta)}{f(\theta)} - q^*(\theta)(1-\lambda) \tag{8}$$

with boundary condition $q^*(\bar{\theta}) = \bar{\theta}/(1-\lambda)$. Then let $\tilde{q} : \tilde{\Theta} \longrightarrow \mathcal{Q}$ be the function defined by

$$\tilde{q}(\tilde{\theta}) \equiv \begin{cases} 0 & \text{if} \quad \tilde{\theta} < m(\underline{\theta}) \\ q^*(m^{-1}(\tilde{\theta})) & \text{if} \quad \tilde{\theta} \in [m(\underline{\theta}), m(\bar{\theta})] \\ q^*(\bar{\theta}) & \text{if} \quad \tilde{\theta} > m(\bar{\theta}), \end{cases} \tag{9}$$

with $m(\theta) \equiv \theta + \lambda q^*(\theta)$. If, for any $\theta \in int(\Theta)$, the function $h(\cdot; \theta) : \mathcal{Q} \to \mathbb{R}$ defined by

$$h(q; \theta) \equiv \theta q + \int_{\theta}^{\theta+\lambda q} \tilde{q}(\tilde{\theta})d\tilde{\theta} - q^2/2 - \tfrac{1-F(\theta)}{f(\theta)}[q + \tilde{q}(\theta + \lambda q) - \tilde{q}(\theta)] \tag{10}$$

[26] The collusive schedules solve the following pointwise maximization problem: $\max_{q_1,q_2}\{\theta\,[q_1+q_2] + \lambda q_1 q_2 - \frac{1}{2}(q_1^2 + q_2^2) - \frac{1-F(\theta)}{f(\theta)}\,[q_1+q_2]\}$.

[27] In the approach followed in the literature, twice differentiability is assumed to guarantee that a seller's best response can be obtained as a solution to a well-behaved optimization problem (e.g. Martimort 1992).

is quasiconcave in q, then the schedules $q_i(\cdot) = q^*(\cdot)$, $i = 1, 2$, can be sustained as a symmetric pure-strategy equilibrium of Γ^M in which the agent's strategy is Markov.

The result in Proposition 3 thus offers a convenient two-step procedure to construct equilibrium schedules. The first step consists in solving the differential equation given in (8). The second step consists in checking whether the function $h(\cdot)$ constructed using the solution $q^*(\cdot)$ to (8) is quasiconcave. If this is the case, the pair of schedules $q_i(\cdot) = q^*(\cdot)$, $i = 1, 2$, is implementable.

4.2 Menu auctions

Consider now a menu auction environment a' la Bernheim and Whinston (1985, 1986a): in these models, the agent's effort is verifiable, preferences are common knowledge (i.e. $|\Theta| = 1$) and each principal is restricted to offer only deterministic mechanisms.[28]

Definition 3 *The environment is deterministic if and only if, for any $i \in \mathcal{N}$, the set $\mathcal{D}_i$ contains only degenerate lotteries that assign measure one to deterministic contracts $y_i : E \longrightarrow \mathcal{A}_i$.*

In virtually all menu auction papers, it is customary to assume that principals make take-it-or-leave-it offers to the agent, that is, they offer a single contract $y_i : E \longrightarrow \mathcal{A}_i$. Peters (2003) shows that any equilibrium in take-it-or-leave-it offers is robust; furthermore when the *no-externalities* condition holds, any outcome that can be sustained with richer menus can also be sustained with take-it-or-leave-it offers. The no-externalities condition typically holds in environments in which the principals' decisions are monetary transfers to the agent (as in Bernheim and Whinston) and where payoffs are quasi-linear in money. When instead a principal's action is the selection of a policy, or of a reward package that includes a non-monetary compensation such as the transfer of an asset, assuming the no-externalities condition holds is restrictive. In this case, principals must be allowed to offer menus of contracts. The question is then how to identify the menus that sustain the equilibrium outcomes.

One approach is offered by Theorem 2. A profile of decisions (e^*, a^*) can then be sustained as a pure-strategy equilibrium in which the agent's strategy is Markov if and only if there exists a profile of incentive-compatible revelation mechanisms[29] $\phi_i^{r*} : Y_{-i} \longrightarrow Y_i$, $i = 1, ..., n$, and a profile of contracts $y^* \equiv (y_1^*, ..., y_n^*)$ such that (i) $\phi_i^{r*}(y_{-i}^*) = y_i^*$; (ii) given y^*, $e^* \in \arg\max_e \bar{V}(e; y^*)$ and $y_i^*(e^*) = a_i^*$, $i = 1, ..., n$; (iii) given any contract $y_i \neq y_i^*$, there exists a profile of contracts $y_{-i} = (y_1, ..., y_{i-1}, y_{i+1}, ..., y_n)$ such that, for any $j \in \mathcal{N} \backslash \{i\}$, $y_j = \phi_j^{r*}(y_{-j-i}, y_i)$, where $y_{-j-i} \equiv (y_l)_{l \neq i,j}$, and a level of effort $e \in \arg\max_e \bar{V}(e; y_{-i}, y_i)$ such that $u_i(e, a) \leq u_i(e^*, a^*)$, with $a = y(e)$ and

[28] See also Dixit, Grossman and Helpman (1997), Biais, Martimort and Rochet (1997), Parlour and Rajan (2001), and Segal and Whinston, (2003).

[29] Because the environment is deterministic, we find it convenient to denote the domain of ϕ_i^r with Y_{-i} and its codomain with Y_i instead of $\mathcal{D}_{-i}$ and $\mathcal{D}_i$ respectively.

$V(e,a) \geq V(e',a')$ for any (e',a') such that there exists a $y'_{-i} \in \mathrm{Im}(\phi^{r*}_{-i})$ for which $a' = y'(e')$, where $y' = (y'_{-i}, y_i)$.

This approach uses incentive-compatibility over *contracts*, i.e. it specifies a contract for the agent as a function of the contracts offered by the other principals. An alternative and perhaps more convenient approach is to think of the principals offering revelation mechanisms that respond directly to the primitive actions a_{-i} taken by the other principals.

Definition 4 *Let $\mathring{\Phi}^r_i$ denote the set of mechanisms $\mathring{\phi}^r_i : E \times \mathcal{A}_{-i} \to \mathcal{A}_i$ such that, for any $e \in E$ and any $a_{-i}, a'_{-i} \in \mathcal{A}_{-i}$*

$$v(\phi^r_i(e, a_{-i}), e, a_{-i}) \geq v(\phi^r_i(e, a'_{-i}), e, a_{-i}).$$

The idea is simple. In settings in which the no-externalities condition fails, for any given $e \in E$, the agent's preferences over the contracts offered by principal i depend on the decisions a_{-i} by the other principals. By implication, any menu of contracts by P_i can be conveniently described through a mapping $\mathring{\phi}^r_i$ that specifies, for each *observable* e and for each *unobservable* a_{-i}, an action a_i that is optimal for the agent among those that the agent can induce by reporting different a'_{-i}.[30] We then have the following result.

Proposition 4 *Let $\mathring{\Gamma}^r$ be the game in which P_i's strategy space is $\Delta(\mathring{\Phi}^r_i)$, $i = 1, ..., n$. Suppose the environment is deterministic. A SCF π can be sustained as a pure-strategy equilibrium of Γ^M in which the agent's strategy is Markov if and only if it can be sustained as a pure-strategy truthful equilibrium of $\mathring{\Gamma}^r$.*

Using the direct mechanisms of Definition 4, the necessary and sufficient conditions for the implementability of the decisions (e^*, a^*) can then be stated as follows. There exists a profile of mechanisms $\mathring{\phi}^{r*}$ such that (i) $a^*_i = \mathring{\phi}^{r*}_i(e^*, a^*_{-i})$ with $v(e^*, a^*) \geq v(e', a')$ for any $e' \in E$ and any $a' \in \mathcal{A}$ such that $a'_j = \mathring{\phi}^{r*}_j(e', \hat{a}_{-j})$ for some $\hat{a}_{-j} \in \mathcal{A}_{-j}$; (ii) for any i and any contract $y_i \in Y_i$, there exists a profile of decisions (e, a) with $a_i = y_i(e)$ and $a_j = \mathring{\phi}^{r*}_j(e, a_{-j})$ for all $j \neq i$ such that $u_i(e, a) \leq u_i(e^*, a^*)$ and $v(e, a) \geq v(e', a')$ for any $e' \in E$ and any $a' \in \mathcal{A}$ such that $a'_i = y_i(e')$ and $a'_j = \mathring{\phi}^{r*}_i(e', \hat{a}_{-j})$ for some $(\hat{a}_{-j}) \in \mathcal{A}_{-j}$.

To illustrate how these conditions can facilitate the construction of equilibria in environments in which the no-externalities condition fails, consider a setting with two principals where $E = \{e_1, e_2\}$ and $\mathcal{A}_i = [0, 1]$, $i = 1, 2$. Payoffs are given by the following functions, where $i, j = 1, 2$, $j \neq i$:

$$u_i(e, a) = \begin{cases} a_i(1 - a_j/2) - a_j & \text{if } e = e_i \\ a_i(a_j - 1/2) - a_j & \text{if } e = e_j \end{cases} \qquad v(e, a) = \begin{cases} 1 + a_2(2a_1 - 1) & \text{if } e = e_1 \\ 2 + a_1(a_2 - 2) - a_2/2 & \text{if } e = e_2 \end{cases}$$

[30] When the agent's preferences are not common knowledge, these mechanisms must be replaced by $\mathring{\phi}^r_i : E \times \mathcal{A}_{-i} \times \Theta \to \mathcal{A}_i$, with $\mathring{\phi}^r_i(e, a_{-i}, \theta) \in \arg\max_{a_i \in \mathcal{A}_i(e; \mathring{\phi}^r_i)} v(e, a_i, a_{-i}, \theta)$ for any $(e, a_{-i}, \theta) \in E \times \mathcal{A}_{-i} \times \Theta$, where $\mathcal{A}_i(e; \mathring{\phi}^r_i) \equiv \{a_i : a_i = \mathring{\phi}^r_i(e, a_{-i}, \theta),\ a_{-i} \in \mathcal{A}_{-i},\ \theta \in \Theta\}$.

One can interpret this environment as one in which a politician, or a regulator, must choose between two policies, e_1 and e_2, and where two firms (P_1 and P_2) must choose the "aggressiveness" of their business strategies, with $a_i = 1$ denoting the most aggressive strategy and $a_i = 0$ the least aggressive one. When $e = e_i$ firm i has a dominant strategy in choosing $a_i = 1$ in which case the other firm has an (iteratively) dominant strategy in choosing $a_j = 1$. However, by behaving aggressively, firms reduce their payoffs with respect to what they could obtain by "colluding", i.e. by playing $a_1 = a_2 = 0$. The politician's payoff is domestic welfare (some weighted average of consumer surplus and domestic profits). This in turn depends on the aggressiveness of the two firms' strategies and on the policy e. Think of policy e_2 as opening the domestic market to foreign competition, while policy e_1 as protectionism. While under protectionism welfare is maximal when the two domestic firms behave aggressively, the opposite is true under foreign competition.[31]

Notice that in this environment, condition (b) of Theorem 3 holds so that restricting attention to Markov strategies is without loss: by implication, any pure-strategy equilibrium outcome can be characterized as a truthful equilibrium of the revelation game.

In the lobbying game in which the two firms are restricted to make take-it-or-leave-it offers to the politician (these offers specify a business strategy for each policy e), the only two pure-strategy equilibrium outcomes are: (i) $e^* = e_1$ and $a_i^* = 1$, $i = 1, 2$ which yields each firm a payoff of $-1/2$ and the policy maker a payoff of 2; and (ii) $e^* = e_2$ and $a_i^* = 1/2$, $i = 1, 2$ which yields P_1 a payoff of $-1/2$, P_2 a payoff of $-1/8$ and the policy maker a payoff of 1 (the proof is in the Appendix).

When, instead, firms can offer menus of contracts, the following outcome can also be sustained as an equilibrium: $e^* = e_1$, $a_1^* = 1/2$, $a_2^* = 0$. This can seen using Theorem 3 and Proposition 4. Suppose, for example, that the principals offer following mechanisms which clearly satisfy conditions (i) and (ii) above:

$$\mathring{\phi}_1^r(e, a_2) = \begin{cases} 1/2 & \text{if } e = e_1 \ \forall a_2 \\ 1 & \text{if } e = e_2 \ \forall a_2 \end{cases}, \quad \mathring{\phi}_2^r(e, a_1) = \begin{cases} 1 & \text{if } e = e_1 \text{ and } a_1 \geq 1/2 \\ 0 & \text{if } e = e_1 \text{ and } a_1 < 1/2 \\ 1 & \text{if } e = e_2 \ \forall a_1 \end{cases}$$

Given these mechanisms, firms induce the policy maker to choose the protectionist policy e_1 while at the same time achieving higher cooperation than under simple take-it-or-leave-it offers, thus obtaining higher total profits. The key to sustaining this outcome is to have P_2 respond to the policy e_1 with a business strategy that depends on what P_1 does. Because P_2 cannot observe a_1 directly, such a contingency must be achieved with the compliance of the agent. A revelation mechanism is then a convenient way of describing P_2's response to P_1's strategy that is compatible with the agent's incentives.

[31] There may be several explanations for this. For example, the politician may value cooperation between the two domestic firms when high consumer surplus is guaranteed by foreign supply, while it may prefer low cooperation when the entire supply comes from the two domestic firms.

4.3 Moral hazard

We now turn to environments in which the agent's effort is not observable. In these environments, a principal's action consists of an incentive scheme that specifies a reward to the agent as a function of some (verifiable) performance measure that is correlated with the agent's effort. Depending on the application of interest, the reward can be a monetary payment, the transfer of an asset, the choice of a policy, or a combination of any of the above.

At a first glance, using revelation mechanisms may appear prohibitively complex in this setting due to the fact that the agent must report an entire array of incentive schemes to each principal. However, as long as for any array of incentive schemes, the choice of optimal effort for the agent is unique, things simplify significantly. Indeed, it suffices to attach to each incentive scheme a label, say an integer a_j, and then simply have the agent reports an array of integers a_{-i} along with his payoff type θ. In fact, because for each array of incentive schemes, the choice of effort is unique, all players' preferences can be expressed in reduced form directly over $\mathcal{A}$. The analysis of incentive-compatibility then proceeds in the familiar way.

To illustrate, considered the following (simplified version of a) standard moral-hazard setting. There are two principals and two effort levels, $\underline{e}$ and $\bar{e}$. As in Bernheim and Whinston (1986,b), the agent's preferences are common knowledge so that $|\Theta| = 1$. Each principal i must choose an incentive scheme a_i from the set $\mathcal{A}_i = \{a^l, a^m, a^h\}$, $i = 1, 2$, where a^l stands for a low-power, a^m for a medium-power and a^h for a high-power incentive scheme.[32]

Instead of specifying for each player an utility function over (w, e), where $w \equiv (w_i)_{i=1}^n$ is an array of rewards (e.g. monetary transfers from the principals to the agent), in the following table we describe directly the players' expected payoffs (u_1, u_2, v) as a function of the agent's effort and the principals' incentive schemes.

$e = \underline{e}$

$a_1 \backslash a_2$	a^h	a^m	a^l
a^h	1 2 2	1 3 1	1 6 0
a^m	2 2 2	2 3 4	2 6 1
a^l	3 2 0	3 3 1	3 6 4

$e = \bar{e}$

$a_1 \backslash a_2$	a^h	a^m	a^l
a^h	4 5 4	4 5 5	4 4 3
a^m	5 5 5	5 5 1	5 4 0
a^l	6 5 2	6 5 0	6 4 0

Table 2

Note that there are no direct externalities between the principals: given e, $u_i(e, a_i, a_j)$ is independent of a_j, $j \neq i$, meaning that P_i is interested in the incentive scheme offered by P_j only because the latter influences the agent's choice over effort. Nevertheless, the *no-externalities* condition of Peters

[32] That the set of feasible incentive schemes is finite in this example is clearly only to shorten the exposition. The same logic applies to settings in which each $\mathcal{A}_i$ has the cardinality of the continuum; in this case, an incentive scheme can be indexed, for example, by an integer $a_i \in [0, 1]$.

(2003) fails here because the agent's preferences over the schemes offered by P_i depend on the incentive scheme offered by P_j; by implication, restricting the principals to offer a single incentive scheme may preclude the possibility of sustaining certain outcomes, as we verify below.[33] Also note that payoffs are such that the agent prefers a high effort to a low effort if and only if at least one of the two principals offered a high-power incentive scheme. The players' payoffs (U_1, U_2, V) can thus be written in reduced form as a function of the (a_1, a_2) only.

$a_1 \backslash a_2$	a^h	a^m	a^l
a^h	4 5 4	4 5 5	4 4 3
a^m	5 5 5	2 3 4	2 6 1
a^l	6 5 2	3 3 1	3 6 4

Table 3

Now suppose the principals were restricted to offer a single incentive scheme to the agent. The unique pure-strategy equilibrium outcome would be $(a^h, a^m, \bar{e})$ with associated expected payoffs $(4, 5, 5)$.

When, instead, principals are allowed to offer menus of incentive schemes, the outcome $(a^m, a^h, \bar{e})$ can also be sustained as a pure-strategy equilibrium outcome.[34] The advantage of menus stems from the fact that they give the agent the possibility of punishing deviations by principal j by selecting a different incentive scheme with principal i. Because the agent's preferences over P_i's incentive schemes in turn depend on the incentive scheme selected by P_j, these menus can be conveniently described as mappings $\phi_i^r : \mathcal{A}_j \longrightarrow \mathcal{A}_i$ with the property that, for any a_j, $\phi_i^r(a_j) \in \arg\max_{a_i \in \text{Im}(\phi_i^r)} V(a_i, a_j)$. The following mechanisms then support $(a^m, a^h, \bar{e})$ as a truthful equilibrium:

$$\phi_1^{r*}(a_2) = \begin{cases} a^h \text{ if } a_2 = a^l, a^m \\ a^m \text{ if } a_2 = a^h \end{cases} \qquad \phi_2^{r*}(a_1) = \begin{cases} a^h \text{ if } a_2 = a^h, a^m \\ a^l \text{ if } a_2 = a^l \end{cases}$$

Given these mechanisms, it is strictly optimal for the agent to choose (a^m, a^h) and then to select $e = \bar{e}$. Furthermore, given ϕ_{-i}^{r*}, it is immediate that no principal i has a profitable deviation, which verifies that $(a^m, a^h, \bar{e})$ can be supported as an equilibrium.

[33] See Attar, Piaser and Porteiro, (2007a) and Peters (2007) for the appropriate version of the no-externalities condition in models with non-contractable effort and Attar, Piaser, and Porteiro (2007b) for an alternative set of conditions.

[34] Note that the possibility of sustaining $(a^m, a^h, \bar{e})$ is appealing because $(a^m, a^h, \bar{e})$ yields a Pareto improvement with respect to $(a^h, a^m, \bar{e})$.

5 Enriched mechanisms

Suppose now one is interested in SCFs that cannot be sustained by restricting the agent's strategy to be Markov or in SCFs that cannot be sustained by restricting the players' strategies to be pure. The question we address in this section is whether there exist natural ways of enriching the simple revelation mechanisms introduced above so as to characterize such SCFs, while at the same time avoiding the "infinite regress" problem of universal revelation mechanisms.

First, we consider pure-strategy equilibrium outcomes sustained by non-Markov strategies. Next, we turn to mixed-strategy equilibria.

Although the revelation mechanisms presented here are more complex than the ones considered in the previous sections, they still permit one to conceptualize the role that the agent plays vis a vis each of his principals thus potentially facilitating the characterization of equilibrium outcomes in applications.

5.1 Non-Markov strategies

We first introduce a new class of revelation mechanisms that permits one to accommodate non-Markov strategies and adjust the notion of truthful equilibria accordingly. We then prove that any pure-strategy equilibrium outcome that can be sustained in the menu game can also be sustained as a truthful equilibrium in the new revelation game.

Definition 5 *(i) Let* $\hat{\Gamma}^r$ *denote the revelation game in which each principal's strategy space is* $\Delta(\hat{\Phi}_i^r)$*, where* $\hat{\Phi}_i^r$ *is the set of revelation mechanisms* $\hat{\phi}_i^r : \hat{\mathcal{M}}_i^r \to \mathcal{D}_i$ *with message space* $\hat{\mathcal{M}}_i^r \equiv \Theta \times \mathcal{D}_{-i} \times \mathcal{N}_{-i}$ *with* $\mathcal{N}_{-i} \equiv \mathcal{N} \backslash \{i\} \cup \{0\}$*, such that* $\mathrm{Im}(\hat{\phi}_i^r)$ *is compact and, for any* $(\theta, \delta_{-i}, k) \in \Theta \times \mathcal{D}_{-i} \times \mathcal{N}_{-i}$*,*

$$\hat{\phi}_i^r(\theta, \delta_{-i}, k) \in \arg \max_{\delta_i \in \mathrm{Im}(\hat{\phi}_i^r)} V(\delta_i, \delta_{-i}, \theta).$$

(ii) Given a profile of mechanisms $\hat{\phi}^r \in \hat{\Phi}^r$*, the agent's strategy is truthful in* $\hat{\phi}_i^r$ *if and only if, for any* $\theta \in \Theta$ *and any* $(\hat{m}_i^r, \hat{m}_{-i}^r) \in Supp[\mu(\theta, \hat{\phi}^r)]$*,*

$$\hat{m}_i^r = (\theta, (\hat{\phi}_j^r(\hat{m}_j^r))_{j \neq i}, k), \text{ for some } k \in \mathcal{N}_{-i}.$$

An equilibrium strategy profile $\sigma^{r*} \in \mathcal{E}(\hat{\Gamma}^r)$ *is a truthful equilibrium if and only if, given any profile of mechanisms* $\hat{\phi}^r$ *such that* $|\{j \in \mathcal{N} : \hat{\phi}_j^r \notin Supp[\sigma_j^{r*}]\}| \leq 1$*, the agent's strategy is truthful in every* $\hat{\phi}_i^r \in Supp[\sigma_i^{r*}]$ *with* $k = 0$ *if* $\hat{\phi}_j^r \in Supp[\sigma_j^{r*}]$ *for all* $j \in \mathcal{N}$ *and* $k = l$ *if* $\hat{\phi}_j^r \in Supp[\sigma_j^{r*}]$ *for all* $j \neq l$ *and* $\hat{\phi}_l^r \notin Supp[\sigma_l^{r*}]$ *for some* $l \in \mathcal{N}$*.*

The interpretation is that the agent is now asked to report to each P_i the identity $k \in \mathcal{N}_{-i}$ of a deviating principal, in addition to (θ, δ_{-i}), with $k = 0$ in the absence of any deviation. Because the identity of a deviating principal is not payoff-relevant, a revelation mechanism $\hat{\phi}_i^r$ is incentive-compatible only if, for any $(\theta, \delta_{-i}) \in \Theta \times \mathcal{D}_{-i}$ and any $k, k' \in \mathcal{N}_{-i}$, $V(\phi_i^r(\theta, \delta_{-i}, k), \theta, \delta_{-i}) = V(\phi_i^r(\theta, \delta_{-i}, k'), \theta, \delta_{-i})$. As we show below, allowing a principal's response to (θ, δ_{-i}) to depend on the identity of a deviating principal may be essential to sustain certain outcomes when the agent's strategy is not Markov.

An equilibrium strategy profile is then said to be a truthful equilibrium of the new revelation game $\hat{\Gamma}^r$ if, whenever no more than one principal deviated from equilibrium play, the agent truthfully reports to any of the non-deviating principals his true type θ, the decisions he is inducing with the other principals, and the identity k of the deviating principal. We then have the following result.

Theorem 5 *Any SCF π that can be sustained as a pure-strategy equilibrium of Γ^M can also be sustained as a pure-strategy truthful equilibrium of $\hat{\Gamma}^r$. Furthermore, any SCF π that can be sustained as an equilibrium of $\hat{\Gamma}^r$ can also be sustained as an equilibrium of Γ^M.*

Consider the"only if" part of the result (the "if" part follows from essentially the same arguments as in the proof of Theorem 2).[35] The key step in the proof consists in showing that if the SCF π can be sustained as a pure-strategy equilibrium of Γ^M, it can also be sustained by a continuation equilibrium σ_A^{M*} with the following property. For any $k \in \mathcal{N}$, $\theta \in \Theta$ and $\delta_k \in \mathcal{D}_k$, there exists a unique profile of decisions $\delta_{-k}(\theta, \delta_k) \in \mathcal{D}_{-k}$ such that A always selects $\delta_{-k}(\theta, \delta_k)$ with all principals other than k when his type is θ, the decision A induces with P_k is δ_k, and k is the only deviating principal. In other words, the decisions that the agent induces with the non-deviating principals depend on the decision δ_k of the deviating principal but not on the menus offered by the latter. The decisions $\delta_{-k}(\theta, \delta_k)$ belong to those that minimize the payoff of the deviating principal P_k among those in the equilibrium menus of the non-deviating principals that are optimal for type θ given δ_k. The rest of the proof then follows quite naturally. When the agent reports to P_i that no deviation occurred—i.e. when he reports that his type is θ, that the decisions he is inducing with the other principals are $\delta_{-i}^*(\theta)$ and that $k = 0$—the revelation mechanism $\hat{\phi}_i^{r*}$ responds with the equilibrium decision $\delta_i^*(\theta)$. When instead, the agent reports that principal k deviated and that, as a result, the agent is inducing the decision δ_k with P_k and the decisions $(\delta_j(\theta, \delta_k))_{j \neq i,k}$ with the other principals, the mechanism ϕ_i^{r*} responds with the decision $\delta_i(\theta, \delta_k)$ that, together with $(\delta_j(\theta, \delta_k))_{j \neq i,k}$, minimizes the payoff of the deviating principal P_k.[36] Given the equilibrium mechanisms $\hat{\phi}_{-k}^{r*}$, following a truthful strategy is clearly optimal for the agent. Furthermore, given $\hat{\sigma}_A^{r*}$, a principal P_k who expects all other principals to offer the equilibrium mechanisms $\hat{\phi}_{-k}^{r*}$ cannot do better than offering the equilibrium mechanism $\hat{\phi}_i^{r*}$ herself. We conclude that if the SCF π can be

[35] Note that in general $\hat{\Gamma}^r$ is not an enlargement of Γ^M (certain menus in Γ^M may not be available in Γ^r); nor is Γ^M an enlargement of $\hat{\Gamma}^r$ (the same menu can be offered through multiple revelation mechanisms).

[36] This is only a partial description of the equilibrium mechanisms $\hat{\phi}^{r*}$ and of the continuation equilibrium σ_A^{r*}. The complete description is in the appendix.

sustained as a pure-strategy equilibrium of Γ^M it can also be sustained as a pure-strategy *truthful* equilibrium of $\hat{\Gamma}^r$.

To see why, with non-Markov strategies, it may be essential to condition a principal's response to (θ, δ_{-i}) on the identity of a deviating principal, consider the following example where $n = 3$, $|\Theta| = |E| = 1$, $\mathcal{A}_1 = \{t, m, b\}$, $\mathcal{A}_2 = \{l, r\}$, $\mathcal{A}_3 = \{s, d\}$ and payoffs (u_1, u_2, u_3, v) as in the following table.

$a_3 = s$

$a_1 \backslash a_2$	l				r			
t	1	4	4	5	1	5	0	4
m	1	1	1	0	1	5	1	0
b	1	1	1	0	1	0	1	0

$a_3 = d$

$a_1 \backslash a_2$	l				r			
t	1	0	5	4	1	1	1	3
m	1	1	1	0	1	0	5	5
b	1	1	5	0	1	5	0	5

Table 3

For simplicity, assume that lotteries over contracts are not feasible and that only deterministic contracts can be offered: because there is no effort, a menu of deterministic contracts for each P_i then consists of a subset of $\mathcal{A}_i$. The outcome (t, l, s) can then be sustained as a pure-strategy equilibrium of the menu game Γ^M. The equilibrium features each P_i offering the entire menu $\mathcal{A}_i$. Given the equilibrium menus, the agent chooses (t, l, s). Any deviation by P_2 to the (degenerate) menu $\{r\}$ is punished by the agent choosing m with P_1 and d with P_3, whereas any deviation by P_3 to the degenerate menu $\{d\}$ is punished by the agent choosing b with P_1 and r with P_2. This strategy for the agent is clearly non-Markov: given the same actions $(a_2, a_3) = (r, d)$ with P_2 and P_3, the agent chooses different actions with P_1 as a function of the particular menus offered by P_2 and P_3. This behavior is essential to sustain the equilibrium outcome. By implication, (t, l, s) cannot be sustained as an equilibrium of the revelation game Γ^r in which the principals offer the simple mechanisms $\phi_i^r : \mathcal{A}_{-i} \to \mathcal{A}_i$ considered in the previous sections.[37]The outcome (t, l, s) can however be sustained as a truthful equilibrium of the more general revelation game $\hat{\Gamma}^r$ in which the agent reports the identity of the deviating principal in addition to the payoff-relevant decisions a_{-i}.[38]

[37] In fact, any incentive-compatible mechanism ϕ_1^r that permits the agent to induce the equilibrium decision t must satisfy $\phi_i^r(a_2, a_3) = t$ for any $(a_2, a_3) \neq (r, d)$; this is because the agent strictly prefers t to both m and b for any $(a_2, a_3) \neq (r, d)$. It follows that any such mechanism either fails to provide the agent with the decision m that is necessary to punish a deviation by P_2 or the decision b that is necessary to punish a deviation by P_3.

[38] Consistently with the result in Theorem 3, note that the problems with simple revelation mechansims $\phi_i^r : \mathcal{A}_{-i} \to \mathcal{A}_i$ emerge in this example only because (i) the agent is indifferent about P_1's response to $(a_2, a_3) = (r, d)$ so that he can be induced to choose different decisions with P_1 as a function of whether it is P_2 or P_3 who deviated from equilibrium play; (ii) the principals' payoffs are sufficiently asymmetric so that the decision the agent induces with P_1 to punish a deviation by P_2 cannot be the same as the one he induces to punish a deviation by P_3.

5.2 Mixed strategies

We now turn to equilibria in which the principals randomize over their mechanisms and/or the agent randomizes over the reports he sends to the principals.[39]

The reason why the simple mechanisms considered in Section 3 may fail to sustain certain mixed-strategy outcomes is that they do not permit the agent to induce different decisions with the same principal in response to the same decisions δ_{-i} he is inducing with the other principals. To illustrate, consider the following example in which $|\Theta| = |E| = 1$, $n = 2$, $\mathcal{A}_1 = \{t, b\}$ and $\mathcal{A}_2 = \{l, r\}$, and where payoffs, (u_1, u_2, v) are as in the following table:

$a_1 \backslash a_2$	l			r		
t	2	1	1	1	0	1
b	1	0	1	1	2	0

Table 4

Assume that $\mathcal{D}_i$ contains only degenerate lotteries over $\mathcal{A}_i$. The following is then an equilibrium in the menu game. Each principal offers the menu ϕ_i^{M*} whose image is the entire set $\mathcal{A}_i$. Given the equilibrium menus, the agent selects with equal probabilities the decisions (t, l), (b, l) and (t, r). Note that it is essential that $\mathcal{D}_i$ contains only degenerate lotteries. If P_1 could offer non-degenerate lotteries, she could do better by deviating and offer the lottery that gives t and b with equal probabilities. In this case, A would strictly prefer to induce l with P_2 thus giving P_1 a higher payoff. As anticipated in the introduction, we see this as a serious limitation on what can be implemented with mixed strategy equilibria: when neither the agent's nor the principals' preferences are flat (i.e. constant over $E \times \mathcal{A}$) and when all lotteries are feasible, it is very difficult to construct examples where the agent is indifferent over the lotteries offered by the principals (so that he can be induced to randomize) and, at the same time, no principal can benefit by breaking the agent's indifference by offering a different menu so as to induce the agent to choose only those lotteries that are more favorable to her.

Having said this, it is important to note that, while certain SCFs may not be sustained with the simple revelation mechanisms $\phi_i^r : \mathcal{D}_{-i} \longrightarrow \mathcal{D}_i$ of the previous sections, *any* SCF that can be sustained as a mixed strategy equilibrium in the menu game can also be sustained as a truthful equilibrium of an enriched revelation game in which the principals offer *set-valued* revelation

[39] Recall that the notion of pure-strategy equilibria given in Definition 1 allows the agent to mix over effort.

mechanisms $\tilde{\phi}_i^r : \Theta \times \mathcal{D}_{-i} \longrightarrow 2^{\mathcal{D}_i}$ such that, for any $(\theta, \delta_{-i}) \in \Theta \times \mathcal{D}_{-i}$,[40]

$$\tilde{\phi}_i^r(\theta, \delta_{-i}) = \arg \max_{\delta_i \in \mathrm{Im}(\tilde{\phi}_i^r)} V(\delta_i, \delta_{-i}, \theta)$$

The interpretation is that the agent first reports his type along with the decisions δ_{-i} that he is inducing with the other principals (possibly by mixing, or in response to a mixed strategy by one of the other principals); the mechanism then responds by giving the agent the decisions in $\tilde{\phi}_i^r$ that are optimal for type θ given δ_{-i}; finally, the agent selects a decision from the set $\tilde{\phi}_i^r(\theta, \delta_{-i})$ and this decision is implemented. In the example above, the equilibrium SCF can be sustained by having P_1 offer the mechanism $\tilde{\phi}_1^{r*}(l) = \{t, b\}$, $\tilde{\phi}_1^{r*}(r) = \{t\}$, and P_2 the mechanism $\tilde{\phi}_2^{r*}(t) = \{l, r\}$, $\tilde{\phi}_2^{r*}(b) = \{l\}$. Given the equilibrium mechanisms, with probability 1/3, the agent induces the decisions (t, l) by reporting l truthfully to P_1 and then choosing t from $\tilde{\phi}_1^{r*}(l)$ and by reporting t truthfully to P_2 and then choosing l from $\tilde{\phi}_2^{r*}(t)$, and so on. The equilibrium is truthful in the sense that the agent may well randomize over the decisions he is inducing with the principals, but once he has decided which decisions he wants to induce (i.e. for any given realization of his mixed strategy), he always reports these decisions truthfully to each principal.

Now note that, although a revelation mechanism is conveniently described by the correspondence $\tilde{\phi}_i^r : \Theta \times \mathcal{D}_{-i} \longrightarrow 2^{\mathcal{D}_i}$, formally such a mechanism is a standard single-valued mapping $\bar{\phi}_i^r : \mathcal{M}_i^r \to \mathcal{D}_i$ with message space $\tilde{\mathcal{M}}_i^r \equiv \Theta \times \mathcal{D}_{-i} \times \mathcal{D}_i$ such that[41]

$$\bar{\phi}_i^r(\theta, \delta_{-i}, \delta_i) = \begin{cases} \delta_i \text{ if } \delta_i \in \tilde{\phi}_i^r(\theta, \delta_{-i}), \\ \delta_i' \in \tilde{\phi}_i^r(\theta, \delta_{-i}) \text{ otherwise.} \end{cases}$$

These mechanisms are clearly incentive-compatible in the sense that, given (θ, δ_{-i}), the agent (weakly) prefers *any* decision in $\tilde{\phi}_i^r(\theta, \delta_{-i})$ to any decision that can be obtained by reporting (θ', δ'_{-i}). Furthermore, given any profile of mechanisms $\tilde{\phi}^r$, the decisions that are optimal for each type θ always belong to those that can be obtained by reporting truthfully to each principal.

Definition 6 *Let $\tilde{\Gamma}^r$ denote the revelation game in which each principal's strategy space is $\Delta(\tilde{\Phi}_i^r)$, where $\tilde{\Phi}_i^r$ is the class of set-valued incentive-compatible revelation mechanisms defined above. Given a mechanism $\tilde{\phi}_i^r \in \tilde{\Phi}_i^r$, the agent's strategy is truthful in $\tilde{\phi}_i^r$ if and only if, for any $\tilde{\phi}_{-i}^r \in \tilde{\Phi}_{-i}^r$, $\theta \in \Theta$ and $\tilde{m}^r \in Supp[\mu(\theta, \tilde{\phi}_i^r, \tilde{\phi}_{-i}^r)]$,*

$$\tilde{m}_i^r = (\bar{\phi}_1^r(\tilde{m}_1^r), ..., \bar{\phi}_i^r(\tilde{m}_i^r), ..., \bar{\phi}_n^r(\tilde{m}_n^r), \theta),$$

[40] With an abuse of notation, in the sequel, we denote by $2^{\mathcal{D}_i}$ the power set of $\mathcal{D}_i$, with the exclusion of the empty set.

For any set-valued mapping $f : \mathcal{M}_i \to 2^{\mathcal{D}_i}$, we then let $\mathrm{Im}(f) \equiv \{\delta_i \in \mathcal{D}_i : \exists\, m_i \in \mathcal{M}_i \text{ s.t. } \delta_i \in f(m_i)\}$ denote the range of f.

[41] The particular decision δ_i' associated to the message $m_i^r = (\theta, \delta_{-i}, \delta_i)$, with $\delta_i \notin \tilde{\phi}_i^r(\delta_{-i}, \theta)$ is not important: the agent never finds it optimal to choose any such message.

An equilibrium strategy profile $\tilde{\sigma}^r \in \mathcal{E}(\tilde{\Gamma}^r)$ is a truthful equilibrium if $\tilde{\sigma}^r_A$ is truthful in every $\tilde{\phi}^r_i \in \tilde{\Phi}^r_i$ for any $i \in \mathcal{N}$.

The agent's strategy is truthful in $\tilde{\phi}^r_i$ if the message $\tilde{m}^r_i = (\theta, \delta_{-i}, \delta_i)$ the agent sends to principal i coincides with his true type θ along with the true decisions $\delta_{-i} = \left(\bar{\phi}^r_j(\tilde{m}^r_j)\right)_{j \neq i}$ induced with the other principals by sending the messages $\tilde{m}^r_{-i}$ and the true decision $\delta_i = \bar{\phi}^r_i(\tilde{m}^r_i)$ that A induces with P_i by sending the message $\tilde{m}^r_i$. We then have the following result.

Theorem 6 *A SCF $\pi : \Theta \longrightarrow \Delta(E \times \mathcal{A})$ can be sustained as an equilibrium of Γ^M if and only if it can be sustained as a truthful equilibrium of $\tilde{\Gamma}^r$.*

The proof is similar to the one that establishes the Menu Theorems and is thus confined to the Supplementary Material. The reason why the result does not follow directly from the Menu Theorems is that $\tilde{\Gamma}^r$ is not an enlargement of Γ^M. In fact, the menus in the range of the revelation mechanisms in $\tilde{\Gamma}^r$ are only those that have the following property: for each δ_i in the menu there exists a (θ, δ_{-i}) such that, given (θ, δ_{-i}), δ_i is as good for the agent as any other decision in the menu.[42]

That the principals can be restricted to offer menus that have this property is not surprising; the proof however requires some work to show how the agent's and the principals' mixed strategies can be adjusted to preserve the same distribution over outcomes as in the original unrestricted menu game Γ^M. The value of Theorem 6 is however not in refining the existing Menu Theorems but in providing a convenient way of describing which decisions the agent finds it optimal to induce as a function of the decisions he induces with the other principals; this can facilitate the construction of the equilibrium outcomes in applications in which mixing plays a role.

6 Conclusions

We have shown how the equilibrium outcomes that are typically of interest in common agency games (those sustained by pure-strategy profiles in which the agent's behavior is Markov) can be conveniently characterized by having the principals offering revelation mechanisms in which the agent truthfully reports his type along with the decisions he is inducing with the other principals.

As compared to universal mechanisms, the approach proposed here has the advantage that it does not lead to the infinite regress problem, for it does not require the agent to describe the mechanisms offered by other principals.

[42] These menus are also different from the menus of *undominated* decisions considered in Martimort and Stole (2002). A menu for principal i is said to contain a dominated decision, say δ_i, if there exists another decision δ'_i in the menu such that, *whatever* the decisions δ_{-i} of the other principals, the agent's payoff under δ'_i is strictly higher than under δ_i.

As compared to the Menu Theorems, our results offer a convenient way of describing how the agent chooses from a menu as a function of “who he is” (his exogenous type) and “what he is doing” with the other principals (the decisions he induces in the other relationships). The advantage of describing the agent’s choices from a menu through revelation mechanisms comes from the fact that this often facilitates the characterization of the necessary and sufficient conditions for the sustainability of outcomes as common agency equilibria. We have illustrated such a possibility in a few cases of interest: menu auctions, moral hazard settings, and competition in non-linear tariffs with adverse selection.

We have also shown how the simple revelation mechanisms described above can be enriched (albeit at the cost of an increase in complexity) to characterize also outcomes sustained by non-Markov strategies and/or mixed strategy equilibria.

Throughout the analysis, we have maintained the assumption of a single agent. The idea of having each agent truthfully reporting a profile of payoff-relevant decisions in addition to his private information seems however promising also in environments with *multiple agents* (see, for example, Yamashita, 2007, for an exploration in this direction).

Appendix

Proof of Theorem 2. Part 1. We prove that if there exists a pure-strategy equilibrium σ^{M*} of Γ^M in which the agent’s strategy is Markov and that implements π, then there also exists a truthful pure-strategy equilibrium σ^{r*} of Γ^r that implements the same SCF.

Let ϕ^{M*} and σ_A^{M*} denote respectively the equilibrium menus and the continuation equilibrium that support π in Γ^M. Because σ_A^{M*} is Markov, then for any i and any $(\theta, \delta_{-i}, \phi_i^M)$ there exists a unique decision $\delta_i(\theta, \delta_{-i}; \phi_i^M) \in \mathrm{Im}(\phi_i^M)$ such that A always induces $\delta_i(\theta, \delta_{-i}; \phi_i^M)$ with P_i when the latter offers the menu ϕ_i^M, the agent’s type is θ, and the decisions A induces with the other principals are δ_{-i}. Finally let $\delta^*(\theta) = (\delta_i^*(\theta))_{i=1}^n$ denote the equilibrium decisions that type θ induces in Γ^M when all principals offer the equilibrium menus, i.e., when $\phi^M = (\phi_i^{M*})_{i=1}^n$.

Now consider the following strategy profile σ^{r*} for the revelation game Γ^r. Each principal i, $i \in \mathcal{N}$, offers the mechanism ϕ_i^{r*} such that

$$\phi_i^{r*}(\theta, \delta_{-i}) = \delta_i(\theta, \delta_{-i}; \phi_i^{M*}) \ \forall \ (\theta, \delta_{-i}) \in \Theta \times \mathcal{D}_{-i}.$$

The agent’s strategy σ_A^{r*} is such that, when $\phi^r = (\phi_i^{r*})_{i=1}^n$, then each type θ reports to each principal i the message $m_i^r = (\theta, \delta_{-i}^*(\theta))$ thus inducing the equilibrium decision $\delta_i^*(\theta)$ with each principal. Given the contracts y, then each type θ induces the same distribution over effort he would have induced in Γ^M had the contracts profile been y, the menus profile been ϕ^{M*}, and the lotteries profile been $\delta^*(\theta)$.

If, instead, ϕ^r is such that $\phi_j^r = \phi_j^{r*}$ for all $j \neq i$ whereas $\phi_i^r \neq \phi_i^{r*}$, then each type θ induces the same outcomes he would have induced in Γ^M had the menu profile been $\phi^M = ((\phi_j^{M*})_{j\neq i}, \phi_i^M)$ where ϕ_i^M is the menu whose image is $\mathrm{Im}(\phi_i^M) = \mathrm{Im}(\phi_i^r)$. That is, let $\delta(\theta; \phi^M)$ denote the decisions that type θ would have induced in Γ^M given ϕ^M. Then given ϕ^r, A induces the decision $\delta_i(\theta; \phi^M)$ with the deviating principal P_i and then reports to each non-deviating principal P_j the message $m_j^r = (\theta, \delta_{-j}(\theta; \phi^M))$ thus inducing the same decisions $\delta(\theta; \phi^M)$ as in Γ^M. In the continuation game that starts after the contracts y are realized, A then induces the same distribution over effort he would have induced in Γ^M given the contracts y, the menus ϕ^M and the decisions $\delta(\theta; \phi^M)$.

Finally, given any profile of mechanisms ϕ^r such that $|\{j \in \mathcal{N} : \phi_j^r \neq \phi_j^{r*}\}| > 1$, the strategy σ_A^{r*} prescribes that A induces the same outcomes he would have induced in Γ^M given ϕ^M, where ϕ^M is the profile of menus such that $\mathrm{Im}(\phi_i^M) = \mathrm{Im}(\phi_i^r)$ for all i.

The strategy σ_A^{r*} described above is clearly a truthful strategy. The optimality of such a strategy in Γ^r then follows directly from the optimality of the agent's strategy σ_A^{M*} in Γ^M together with the fact that $\mathrm{Im}(\phi_i^{r*}) \subseteq \mathrm{Im}(\phi_i^{M*})$ for all i.

Given the continuation equilibrium σ_A^{r*} it is then immediate that any principal P_i who expects all other principals P_j, $j \neq i$, to offer the mechanisms ϕ_{-i}^{r*} cannot do better than offering the equilibrium mechanism ϕ_i^{r*}. We conclude that the pure-strategy profile σ^{r*} constructed above is an equilibrium of Γ^r and sustains the same SCF π as the equilibrium σ^{M*} of Γ^M.

Part 2. We now prove the converse: if there exists an equilibrium σ^{r*} of Γ^r that sustains the SCF π, then there also exists an equilibrium σ^{M*} of Γ^M that sustains the same SCF.

First, consider the principals. For any $i \in \mathcal{N}$ and any $\phi_i^M \in \Phi_i^M$, let $\Phi_i^r(\phi_i^M) \equiv \{\phi_i^r \in \Phi_i^r : \mathrm{Im}(\phi_i^r) = \mathrm{Im}(\phi_i^M)\}$ denote the set of revelation mechanisms with the same range as ϕ_i^M (note that $\Phi_i^r(\phi_i^M)$ may well be empty). The strategy $\sigma_i^{M*} \in \Delta(\Phi_i^M)$ for P_i in Γ^M is then such that, for any set of menus $B \subseteq \Phi_i^M$

$$\sigma_i^{M*}(B) = \sigma_i^{r*}(\bigcup_{\phi_i^M \in B} \Phi_i^r(\phi_i^M)).$$

Next, consider the agent.

Case 1. Given any profile of menus $\phi^M \in \Phi^M$ such that, for any $i \in \mathcal{N}$, $\Phi_i^r(\phi_i^M) \neq \varnothing$, the strategy σ_A^{M*} induces the same distribution over $\mathcal{A} \times E$ as the strategy σ_A^{r*} in Γ^r given the event that $\phi^r \in \Phi^r(\phi^M) \equiv \prod_i \Phi_i^r(\phi_i^M)$. Precisely, let $\rho_{\sigma_A^{r*}} : \Theta \times \Phi^r \to \Delta(\mathcal{A} \times E)$ denote the distribution over outcomes induced by the strategy σ_A^{r*} in Γ^r. Then, for any $\theta \in \Theta$, $\sigma_A^{M*}(\theta, \phi^M)$ is such that

$$\rho_{\sigma_A^{M*}}(\theta, \phi^M) = \int_{\Phi^r} \rho_{\sigma_A^{r*}}(\theta, \phi^r) d\sigma_1^{r*}(\phi_1^r | \Phi_1^r(\phi_1^M)) \times \cdots \times d\sigma_n^{r*}(\phi_n^r | \Phi_n^r(\phi_n^M))$$

where, for any i, $\sigma_i^{r*}(\cdot | \Phi_i^r(\phi_i^M))$ denotes the regular conditional probability distribution over Φ_i^r generated by the original strategy σ_i^{r*} conditioning on ϕ_i^r belonging to $\Phi_i^r(\phi_i^M)$.

Case 2. If, instead, ϕ^M is such that there exists a $j \in \mathcal{N}$ such that $\Phi_i^r(\phi_i^M) \neq \varnothing$ for all $i \neq j$

while $\Phi_j^r(\phi_j^M) = \varnothing$, then let ϕ_j^r be any arbitrary revelation mechanism such that

$$\phi_j^r(\theta, \delta_{-j}) \in \arg\max_{\delta_j \in \mathrm{Im}(\phi_j^M)} V(\delta_j, \delta_{-j}, \theta) \qquad \forall (\theta, \delta_{-j}) \in \Theta \times \mathcal{D}_{-j}.$$

The strategy σ_A^{M*} then induces the same outcomes as the strategy σ_A^{r*} given ϕ_j^r and given $\phi_{-j}^r \in \Phi_{-j}^r(\phi_{-j}^M) \equiv \prod_{i \neq j} \Phi_i^r(\phi_i^M)$; that is, for any $\theta \in \Theta$,

$$\rho_{\sigma_A^{M*}}(\theta, \phi^M) = \int_{\Phi_{-j}^r} \rho_{\sigma_A^{r*}}(\theta, \phi_j^r, \phi_{-j}^r) d\sigma_1^{r*}(\phi_1^r | \Phi_1^r(\phi_1^M)) \times \cdots \times d\sigma_n^{r*}(\phi_n^r | \Phi_n^r(\phi_n^M)) \qquad (11)$$

Case 3. Finally, for any ϕ^M such that $|\{j \in \mathcal{N} : \Phi_j^r(\phi_j^M) = \varnothing| > 1$, simply let $\sigma_A^{M*}(\theta, \phi^M)$ be any strategy that is sequentially optimal for A given (θ, ϕ^M).

The fact that σ_A^{r*} was a continuation equilibrium for Γ^r guarantees that the strategy σ_A^{M*} constructed above is a continuation equilibrium for Γ^M. Furthermore, given σ_A^{M*}, any principal P_i who expects any other principal P_j, $j \neq i$, to follow the strategy σ_i^{M*} cannot do better than following the strategy σ_i^{M*}. We conclude that the strategy profile σ^{M*} constructed above is an equilibrium of Γ^M and sustains the same outcomes as σ^{r*} in Γ^r. ■

Proof of Theorem 3. When (a) holds, the result is immediate. In what follows we prove that when (b) holds, then if the SCF π can be sustained as a pure-strategy equilibrium σ^{M*} of Γ^M, it can also be sustained as a pure-strategy equilibrium $\hat{\sigma}^M$ in which the agent's strategy $\hat{\sigma}_A^M$ is Markov.

Let ϕ^{M*} denote the equilibrium menus under the strategy profile σ^{M*} and δ^* denote the equilibrium decisions that are implemented when all principals offer the equilibrium menus ϕ^{M*}.

Suppose that σ_A^{M*} is not Markov. This means that there exists an $i \in \mathcal{N}$, a $\tilde{\phi}_i^M \in \Phi_i^M$, a $\delta'_{-i} \times \mathcal{D}_{-i}$ and a pair $\underline{\phi}_{-i}^M, \bar{\phi}_{-i}^M \in \Phi_{-i}^M$ such that A selects $(\underline{\delta}_i, \delta'_{-i})$ when $\phi^M = (\tilde{\phi}_i^M, \underline{\phi}_{-i}^M)$ and $(\bar{\delta}_i, \delta'_{-i})$ when $\phi^M = (\tilde{\phi}_i^M, \bar{\phi}_{-i}^M)$, with $\underline{\delta}_i \neq \bar{\delta}_i$. We then show that, starting from σ_A^{M*}, one can construct a Markov continuation equilibrium $\hat{\sigma}_A^M$ that induces all principals to continue to offer the equilibrium menus ϕ^{M*} and that sustains the same equilibrium decisions δ^* as σ_A^{M*}.

Case 1. First consider the case that $\tilde{\phi}_i^M = \phi_i^{M*}$ and $\delta'_{-i} = \delta^*_{-i}$. Let then $\hat{\sigma}_A^M$ be the strategy that coincides with σ_A^{M*} for all $\phi^M \neq (\tilde{\phi}_i^M, \underline{\phi}_{-i}^M), (\tilde{\phi}_i^M, \bar{\phi}_{-i}^M)$ and that prescribes that A selects δ^* both when $\phi^M = (\tilde{\phi}_i^M, \underline{\phi}_{-i}^M)$ and when $\phi^M = (\tilde{\phi}_i^M, \bar{\phi}_{-i}^M)$. In the continuation game that starts after the lotteries δ^* select the contracts y, $\hat{\sigma}_A^M$ prescribes that A induces the same distribution over effort he would have induced according to the original strategy σ_A^{M*} as if the menus offered had been ϕ^{M*}. It is immediate that the strategy $\hat{\sigma}_A^M$ is sequentially rational for the agent. It is also immediate that, given $\hat{\sigma}_A^M$, any principal P_j who expects any other principal P_l, $l \neq j$, to offer the equilibrium menu ϕ_l^{M*} cannot do better than offering the equilibrium menu ϕ_j^{M*}.

Case 2. Next consider the case that $\tilde{\phi}_i^M = \phi_i^{M*}$ and $\delta'_{-i} \neq \delta^*_{-i}$ (which implies that both $\underline{\phi}_{-i}^M$ and $\bar{\phi}_{-i}^M$ are necessarily different from ϕ_{-i}^{M*}). Let $\hat{\sigma}_A^M$ be the strategy that coincides with σ_A^{M*} for all

$\phi^M \neq (\tilde{\phi}_i^M, \underline{\phi}_{-i}^M)$,$(\tilde{\phi}_i^M, \bar{\phi}_{-i}^M)$ and that prescribes that A selects $(\delta_i', \delta_{-i}')$ both when $\phi^M = (\tilde{\phi}_i^M, \underline{\phi}_{-i}^M)$ and when $\phi^M = (\tilde{\phi}_i^M, \bar{\phi}_{-i}^M)$, where $\delta_i' \in \arg\max_{\delta_i \in \mathrm{Im}(\tilde{\phi}_i^M)} V(\delta_i, \delta_{-i}')$ is any decision such that, for all $j \neq i$,

$$\underline{U}_j(\delta_i', \delta_{-i}') \leq \underline{U}_j(\hat{\delta}_i, \delta_{-i}') \text{ for all } \hat{\delta}_i \in \arg \max_{\delta_i \in \mathrm{Im}(\tilde{\phi}_i^M)} V(\delta_i, \delta_{-i}').$$

By the Uniform Punishment condition, such a decision always exists. In the continuation game that starts after the lotteries $\delta = (\delta_i', \delta_{-i}')$ select the contracts y, A then selects effort $\xi_k(y)$, where

$$k \in \{j \in \mathcal{N} \backslash \{i\} : \phi_j^M \neq \phi_j^{M*}\}$$

is identity of one of the principals who deviated from equilibrium play, whereas $\xi_k(y)$ is the level of effort defined in (2). Clearly, when $\left|\{j \in \mathcal{N} \backslash \{i\} : \phi_j^M \neq \phi_j^{M*}\}\right| > 1$, the identity k of the deviating principal can be chosen arbitrarily. Once again, it is immediate that the strategy $\hat{\sigma}_A^M$ is sequentially rational for the agent and that, given $\hat{\sigma}_A^M$, any principal P_j who expects any other principal P_l, $l \neq j$, to offer the equilibrium menu ϕ_l^{M*} cannot do better than offering the equilibrium menu ϕ_l^{M*}.

Case 3. Lastly, consider the case that $\tilde{\phi}_i^M \neq \phi_i^{M*}$. Irrespective of whether $\delta_{-i}' = \delta_{-i}^*$ or $\delta_{-i}' \neq \delta_{-i}^*$, let $\hat{\sigma}_A^M$ be the strategy that coincides with σ_A^{M*} for all $\phi^M \neq (\tilde{\phi}_i^M, \underline{\phi}_{-i}^M)$,$(\tilde{\phi}_i^M, \bar{\phi}_{-i}^M)$ and that prescribes that A selects $(\delta_i', \delta_{-i}')$ both when $\phi^M = (\tilde{\phi}_i^M, \underline{\phi}_{-i}^M)$ and when $\phi^M = (\tilde{\phi}_i^M, \bar{\phi}_{-i}^M)$, where $\delta_i' \in \arg\max_{\delta_i \in \mathrm{Im}(\tilde{\phi}_i^M)} V(\delta_i, \delta_{-i}')$ is any decision such that

$$\underline{U}_i\left(\delta_i', \delta_{-i}'\right) \leq \underline{U}_i\left(\hat{\delta}_i, \delta_{-i}'\right) \text{ for all } \hat{\delta}_i \in \arg \max_{\delta_i \in \mathrm{Im}(\tilde{\phi}_i^M)} V(\delta_i, \delta_{-i}').$$

Again, $\hat{\sigma}_A^M$ is trivially sequentially rational for the agent and, given $\hat{\sigma}_A^M$, no principal has an incentive to deviate.

Note that the strategy $\hat{\sigma}_A^M$ constructed from σ_A^{M*} using the procedure described above has the property that, given any $\phi^M \in \Phi^M$ such that $\phi_i^M \neq \tilde{\phi}_i^M$ the behavior specified by $\hat{\sigma}_A^M$ is the same as that specified by the original strategy σ_A^{M*}. Furthermore, for any $\phi^M \in \Phi^M$, the decision the agent takes with any principal P_j, $j \neq i$, is the same as under the original strategy σ_A^{M*}. This implies that the procedure described above can be iterated for all $i \in \mathcal{N}$ and all $\tilde{\phi}_i^M \in \Phi_i^M$; this gives a strategy for the agent that is Markov and that induces all principals to continue to offer the equilibrium mechanisms. ■

Proof of Theorem 4. The proof follows from applying the same steps indicated in the proof of Theorem 3 to all $\theta \in \Theta$ and by noting that, when σ_A^{M*} satisfies the "Conformity to Equilibrium" condition, the following is true. For any $i \in \mathcal{N}$ there exists no pair $\underline{\phi}_{-i}^M, \bar{\phi}_{-i}^M \in \Phi_{-i}^M$ such that some type $\theta \in \Theta$ selects $(\underline{\delta}_i, \delta_{-i}^*(\theta))$ when $\phi^M = (\phi_i^{M*}, \underline{\phi}_{-i}^M)$ and $(\bar{\delta}_i, \delta_{-i}^*(\theta))$ when $\phi^M = (\phi_i^{M*}, \bar{\phi}_{-i}^M)$, with $\underline{\delta}_i \neq \bar{\delta}_i$. In other words, Case 1 in the proof of Theorem 3 never arises when the strategy σ_A^{M*} satisfies the "Conformity to Equilibrium" condition. This in turn guarantees that when one replaces the original strategy σ_A^{M*} with the strategy $\hat{\sigma}_A^M$ that is obtained from σ_A^{M*} iterating the steps in the proof of Theorem 3 for all $\theta \in \Theta$, all $i \in \mathcal{N}$, and all $\tilde{\phi}_i^M \in \Phi_i^M$, it remains optimal for each P_i to offer the equilibrium menu ϕ_i^{M*}. ■

Proof of Proposition 1. It is immediate that conditions (a)-(c) guarantee existence of a truthful equilibrium in the revelation game Γ^r sustaining the schedules $q_i^*(\cdot)$, $i = 1, 2$. Theorem 2 then implies that the same schedules can also be sustained in the menu game Γ^M.

Thus consider the necessity of these conditions. That conditions (a) and (b) are necessary follows directly from Theorem 2: If the schedules $q_i^*(\cdot)$, $i = 1, 2$, can be sustained as a pure-strategy equilibrium of Γ^M in which the agent's strategy is Markov, then they can also be sustained as a pure-strategy truthful equilibrium of Γ^r. As discussed in the main text, the same schedules can then also be sustained by a truthful (pure-strategy) equilibrium in which the mechanism offered by each principal i is such that $\phi_i^r(\theta, q_j, t_j) = \phi_i^r(\theta', q_j', t_j')$ whenever $\theta + \lambda q_j = \theta' + \lambda q_j'$. The definition of such an equilibrium then implies that there must exist a pair of mechanisms $\phi_i^{r*} = (\tilde{q}_i(\cdot), \tilde{t}_i(\cdot))$, $i = 1, 2$, such that $\tilde{q}_i(\cdot)$ is nondecreasing, $\tilde{t}_i(\cdot)$ satisfies (3), and conditions (a) and (b) in the proposition hold.

It remains to show that condition (c) is also necessary. To see this, first note that if there exists a pair of mechanisms $(\tilde{q}_i(\cdot), \tilde{t}_i(\cdot))_{i=1,2}$ and a truthful continuation equilibrium σ_A^r that sustain the schedules $q_i^*(\cdot)$, $i = 1, 2$, then this means that the schedules $q_i^*(\cdot)$ and $t_i^*(\cdot) \equiv \tilde{t}_i(m_i(\cdot))$, $i = 1, 2$, must satisfy the equivalent of the (IC) and (IR) constraints in the program of condition (c); this in turn means that necessarily $U_i^* \leq \bar{U}_i$, $i = 1, 2$. To prove the result it then suffices to show that if $U_i^* < \bar{U}_i$, then P_i has a profitable deviation. This can be shown by contradiction. Suppose there exists a truthful equilibrium $\sigma^r \in \mathcal{E}(\Gamma^r)$ sustaining the schedules $(q_i^*(\cdot))_{i=1,2}$ and such that $U_i^* < \bar{U}_i$, for some $i \in \mathcal{N}$. Then there also exists a (pure-strategy) equilibrium $\sigma^{M*} \in \mathcal{E}(\Gamma^M)$ sustaining the same schedules and such that (i) each P_i offers the menu ϕ_i^{M*} defined by $\text{Im}(\phi_i^{M*}) = \text{Im}(\phi_i^{r*})$ and (ii) each type θ selects the pair $(q_i^*(\theta), t_i^*(\theta))$ from each menu ϕ_i^{M*}, thus yielding P_i a payoff U_i^* (see the proof of part 2 of Theorem 2). We then show that, irrespective of which continuation equilibrium σ_A^{r*} one considers, P_i has a profitable deviation.

Case 1. Suppose that the schedules $q_i(\cdot)$ and $t_i(\cdot)$ that solve the program of condition (c) are such that the (IC) and (IR) constraints hold as strict inequalities for almost all θ. This immediately implies that if in Γ^M P_i deviates and offers the menu ϕ_i^M defined by $\text{Im}(\phi_i^M) = \{(q_i(\theta), t_i(\theta)) : \theta \in \Theta\}$ then in any continuation equilibrium, almost every type θ will necessarily select the pair $(q_i(\theta), t_i(\theta))$ from ϕ_i^M, thus giving P_i a payoff $\bar{U}_i > U_i^*$.[43]

[43] Note that, while almost every $\theta \in \Theta$ strictly prefers $(q_i(\theta), t_i(\theta))$ to any other pair $(q_i, p_i) \in \text{Im}(\phi_i^M)$, there may exist a positive-measure set of types θ' who, given $(q_i(\theta'), t_i(\theta'))$, is indifferent between inducing the decision $(\tilde{q}_j(\theta' + \lambda q_i(\theta')), \tilde{t}_j(\theta' + \lambda q_i(\theta'))$ with P_j or inducing another decision $(q_j, t_j) \in \text{Im}(\phi_j^{M*})$. The fact that P_i is not personally interested in (q_j, t_j) however implies that P_i's deviation to ϕ_i^M is profitable irrespective of how one specifies

Case 2. Next suppose that the schedules $q_i(\cdot)$ and $t_i(\cdot)$ that solve the program of condition (c) are such that the (IC) constraints hold as strict inequalities for almost all θ, but there exists a positive-measure set of types $\Theta' \subset \Theta$ such that, for any $\theta' \in \Theta'$ the (IR) constraint holds as an equality. In this case, a deviation to the menu ϕ_i^M of Case 1 need not be profitable for P_i, for each type $\theta' \in \Theta'$ could react by choosing not to participate. However, if this is the case, then P_i could offer another menu $\phi_i^{M\prime}$ such that $\mathrm{Im}(\phi_i^{M\prime}) = \{(q_i'(\theta), t_i'(\theta)) : \theta \in \Theta\}$, with $q_i'(\theta) = q_i(\theta)$ and $t_i'(\theta) = t_i(\theta) - \varepsilon$ for all $\theta \in \Theta$, with $\varepsilon > 0$. Clearly such a menu guarantees participation by all types. By choosing ε arbitrarily close to zero, P_i can then guarantee herself a payoff arbitrarily close to $\bar{U}_i$ and thus strictly higher than U_i^*, once again a contradiction.

Case 3. Finally, let $V_i(\theta, \theta') \equiv \theta q_i(\theta') + v_i^*(\theta, q_i(\theta')) - t_i(\theta')$ denote the payoff that type θ obtains by selecting the pair $(q_i(\theta'), t_i(\theta'))$ designed by P_i for type θ' and then selecting the pair $(\tilde{q}_j(\theta + \lambda q_i(\theta')), \tilde{t}_j(\theta + \lambda q_i(\theta'))$ with P_j, where $q_i(\cdot)$ and $t_i(\cdot)$ are again the schedules that solve the program of condition (c). Now suppose there exists a positive-measure set of types $\Theta_0 \subset \Theta$ such that for any $\theta \in \Theta_0$, there exists a $\theta' \in \Theta$ such that

$$V_i(\theta, \theta) = V_i(\theta, \theta')$$

with $q_i(\theta') \neq q_i(\theta)$,[44] whereas for any $\theta \in \Theta \backslash \Theta_0$,

$$V_i(\theta, \theta) > V_i(\theta, \hat{\theta}) \text{ for any } \hat{\theta} \in \Theta \text{ such that } q_i(\hat{\theta}) \neq q_i(\theta).$$

The set Θ_0 thus corresponds to the set of types θ for whom the pair $(q_i(\theta), t_i(\theta))$ is not strictly optimal, in the sense that there exists another pair $(q_i(\theta'), t_i(\theta'))$ with $(q_i(\theta'), t_i(\theta')) \neq (q_i(\theta), t_i(\theta))$ that is as good for type θ as the pair $(q_i(\theta), t_i(\theta))$.

Without loss, assume that $q_i(\cdot)$ and $t_i(\cdot)$ are such that each type $\theta \in \Theta$ strictly prefers the pair $(q_i(\theta), t_i(\theta))$ to the null contract $(0, 0)$ (as shown in Case 2 above, P_i can always adjust the original transfer schedule $t_i(\cdot)$ so as to guarantee that this property holds, while preserving incentive compatibility for all types and still obtaining a payoff $U_i > U_i^*$).

Now let $z : \Theta \rightrightarrows 2^\Theta$ be the correspondence defined by

$$z(\theta) = \{\theta' \in \Theta,\ \theta' \neq \theta : V_i(\theta, \theta) = V_i(\theta, \theta') \text{ and } q_i(\theta') \neq q_i(\theta)\}\ \forall \theta \in \Theta$$

and then let $z(\Theta) \equiv \mathrm{Im}(z)$ denote the range of $z(\cdot)$. In words, this correspondence maps each type $\theta \in \Theta$ into the set of types $\theta' \neq \theta$ that receive a contract $(q_i(\theta'), t_i(\theta'))$ different from the one $(q_i(\theta), t_i(\theta))$ designed for type θ but that nonetheless give type θ the same payoff as the contract $(q_i(\theta), t_i(\theta))$.

the agent's choice with P_j.

[44] Cearly if $q_i(\theta) = q_i(\theta')$, which also implies that $t_i(\theta) = t_i(\theta')$, then whether type θ selects the contract designed for him or that designed for type θ' is inconsequential for P_i's payoff.

Finally, let $g : \Theta \rightrightarrows 2^{\Theta}$ denote the correspondence defined by

$$g(\theta) = \{\theta' \in \Theta,\ \theta' \neq \theta :\ (q_i(\theta'), t_i(\theta')) = (q_i(\theta), t_i(\theta))\}\ \forall \theta \in \Theta.$$

This correspondence maps each type θ into the set of types $\theta' \neq \theta$ that, given the schedules $(q_i(\cdot), t_i(\cdot))$, receive the same price-quantity pair as type θ. Then, given any set $\Theta' \subset \Theta$, let

$$g(\Theta') \equiv \{\bigcup g(\theta) : \theta \in \Theta'\}$$

Starting from the schedules $q_i(\cdot)$ and $t_i(\cdot)$, then let $q_i'(\cdot)$ and $t_i'(\cdot)$ be a new pair of schedules such that $q_i'(\theta) = q_i(\theta)$ for all $\theta \in \Theta$, $t_i'(\theta) = t_i(\theta)$ for all $\theta \notin \Theta_0 \cup g(\Theta_0)$ while for any $\theta \in \Theta_0 \cup g(\Theta_0)$, $t_i'(\theta) = t_i(\theta) - \varepsilon$ with $\varepsilon > 0$. Clearly, if ε is chosen sufficiently small, then the new schedules $q_i'(\cdot)$ and $t_i'(\cdot)$ necessarily satisfy the (IC) and (IR) constraints in the program of condition (c) for all θ.

Now suppose that the original schedules $q_i(\cdot)$ and $t_i(\cdot)$ were such that $\{\Theta_0 \cup g(\Theta_0)\} \cap z(\Theta) = \varnothing$. Then the new schedules $q_i'(\cdot)$ and $t_i'(\cdot)$ guarantee that each type $\theta \in \Theta$ now strictly prefers the contract $(q_i'(\theta), t_i'(\theta))$ designed for him to any other contract $(q_i'(\theta'), t_i'(\theta')) \neq (q_i'(\theta), t_i'(\theta))$. This in turn implies that by choosing ε sufficiently small and offering the menu $\phi_i^{M\prime}$ such that $\mathrm{Im}(\phi_i^{M\prime}) = \{(q_i'(\theta), t_i'(\theta)) : \theta \in \Theta\}$, irrespective of the agent's continuation equilibrium σ_A^M, P_i can guarantee herself a payoff arbitrarily close to $\bar{U}_i$ and hence has profitable deviation.

Next suppose that $\{\Theta_0 \cup g(\Theta_0)\} \cap z(\Theta) \neq \varnothing$. This means that there exists a pair θ, θ' with $\theta \in \Theta_0$ and $\theta' \in z(\theta)$ such that either $\theta' \in \Theta_0$ or there exists another type $\theta'' \in \Theta_0$ such that $(q_i(\theta''), t_i(\theta'')) = (q_i(\theta'), t_i(\theta'))$ which in turn implies that $\theta'' \in z(\theta)$. Without loss, thus assume the former case. The schedules $q_i'(\cdot)$ and $t_i'(\cdot)$ constructed above then leave type θ indifferent between the contract $(q_i'(\theta), t_i'(\theta))$ designed for him and the contract $(q_i'(\theta'), t_i'(\theta'))$ designed for type θ'. The fact that the agent's payoff $\theta q_i + v_i^*(\theta, q_i) - v_i^*(\theta, 0)$ has the strict increasing-difference property with respect to (θ, q_i) however guarantees that $\theta \notin z(\theta')$: that is, if type θ is willing to take type θ''s contract, then it cannot be that type θ' is also willing to swap with type θ. The same property also implies that if $\theta'' \in z(\theta')$, with $\theta'' \neq \theta$, then necessarily $\theta \notin z(\theta'')$. That is, if type θ is indifferent between the contract designed for him and the contract designed for type θ' and if, at the same time, type θ' is indifferent between the contract designed for him and that designed for type θ'', then it cannot be that type θ'' is also indifferent between the contract designed for him and that designed for type θ. These properties in turn guarantee that the procedure that permits one to transform the schedules $q_i(\cdot)$ and $t_i(\cdot)$ into the schedules $q_i'(\cdot)$ and $t_i'(\cdot)$ can be iterated (without cycling) till no type is any longer indifferent.

We conclude that if there exists a pair of schedules $q_i(\cdot)$ and $t_i(\cdot)$ that solve the program in condition (c) in the proposition and yield P_i a payoff $\bar{U}_i > U_i^*$, then irrespective of how one specifies the agent's continuation equilibrium, P_i necessarily has a profitable deviation. This in turn proves that (c) is necessary. ■

Proof of Proposition 2. The collusive schedules solve the following pointwise maximization problem:

$$\max_{q_1,q_2}\{\theta\left[q_1+q_2\right]+\lambda q_1 q_2-\tfrac{1}{2}(q_1^2+q_2^2)-\tfrac{1-F(\theta)}{f(\theta)}\left[q_1+q_2\right]\}.$$

The solution to this program is given by[45]

$$q_i(\theta)=q^c(\theta)\equiv\tfrac{1}{1-\lambda}\left(\theta-\tfrac{1-F(\theta)}{f(\theta)}\right)\ \forall\theta,\ i=1,2.$$

To prove the result, we proceed by contradiction. Suppose there exists a pair of tariffs that sustains the collusive schedules as an equilibrium in which the agent's strategy is Markov. Using the result of Proposition 1, there then exists a pair of incentive-compatible mechanisms $\phi_i^r=(\tilde{q}_i(\cdot),\tilde{t}_i(\cdot))$ that satisfies conditions (a) and (b) in Proposition 1 with $q_i^*(\cdot)=q^c(\cdot)$, $i=1,2$. The fact that it is optimal for each θ to select the quantity $q^c(\theta)$ and pay $\tilde{t}_i(m(\theta))$ to each principal implies that, for $i=1,2$,

$$\begin{aligned}V^*(\theta) &= \sup_{\tilde{\theta}_1,\tilde{\theta}_2}\left\{\theta\left[\tilde{q}_1(\tilde{\theta}_1)+\tilde{q}_2(\tilde{\theta}_2)\right]+\lambda\tilde{q}_1(\tilde{\theta}_1)\tilde{q}_2(\tilde{\theta}_2)-\tilde{t}_1(\tilde{\theta}_1)-\tilde{t}_2(\tilde{\theta}_2)\right\}\\ &= \sup_{\tilde{\theta}_i\in\tilde{\Theta}_i}\left\{\theta\tilde{q}_i(\tilde{\theta}_i)+v_i^*(\theta,\tilde{q}_i(\tilde{\theta}_i))-\tilde{t}_i(\tilde{\theta}_i)\right\}\\ &= \sup_{\tilde{\theta}_i\in[m_i(\underline{\theta}),m_i(\bar{\theta})]}\left\{\theta\tilde{q}_i(\tilde{\theta}_i)+v_i^*(\theta,\tilde{q}_i(\tilde{\theta}_i))-\tilde{t}_i(\tilde{\theta}_i)\right\}\end{aligned}$$

where all equalities follow directly from the fact that the mechanisms $\phi_i^r=(\tilde{q}_i(\cdot),\tilde{t}_i(\cdot))$ are incentive-compatible and satisfy conditions (a) and (b) in Proposition 1. Because for any $\theta\in\Theta$ and any message $\tilde{\theta}_i\in[m_i(\underline{\theta}),m_i(\bar{\theta})]$, the marginal valuation $\theta+\lambda\tilde{q}_i(\tilde{\theta}_i)\in[m_j(\underline{\theta}),m_j(\bar{\theta})]$ and because $\tilde{q}_j(\cdot)$ is continuous over $[m_j(\underline{\theta}),m_j(\bar{\theta})]$, there exists a constant $M_i>0$ such that, for any $q\in[\tilde{q}_j(m_j(\underline{\theta})),\tilde{q}_j(m_j(\bar{\theta}))]=[q^c(\underline{\theta}),q^c(\bar{\theta})]$, the function $w_i(\cdot,q):\Theta\to\mathbb{R}$ defined by $w_i(\theta,q)\equiv\theta q+v_i^*(\theta,q)$ is M_i-Lipschitz continuous and differentiable and its derivative satisfies

$$\frac{\partial w_i(\theta,q)}{\partial\theta}=q+\tilde{q}_j(\theta+\lambda q)\leq 2\bar{Q}_i.$$

Using the envelope theorem, we then have that, if the mechanisms ϕ_1^r and ϕ_2^r satisfy conditions (a) and (b), then the functions $\tilde{t}_i(\cdot)$ must satisfy

$$\begin{aligned}\tilde{t}_i(m(\theta)) &= \theta q^c(\theta)+v_i^*(\theta,q^c(\theta))-\int_{\underline{\theta}}^{\theta}[q^c(s)+\tilde{q}_j(s+\lambda q^c(s))]ds-\hat{K}_i\\ &= \theta q^c(\theta)+[v_i^*(\theta,q^c(\theta))-v_i^*(\theta,0)]-\int_{\underline{\theta}}^{\theta}[q^c(s)+\tilde{q}_j(s+\lambda q^c(s))-\tilde{q}_j(s)]ds-K_i\end{aligned}$$

for some $\hat{K}_i,K_i\geq 0$, where the second equality comes from the fact that $v_i^*(\theta,0)=\int_{\min\tilde{\Theta}_i}^{\theta}\tilde{q}_j(s)ds+\tilde{K}_j=\int_{\underline{\theta}}^{\theta}\tilde{q}_j(s)ds+\tilde{K}_j$. This in turn implies that the equilibrium payoff U_i^* for each P_i can be written

[45] For simplicity, we assume that the solution is interior: $q^c(\theta)\in int(\mathcal{Q})$ for any θ.

as

$$U_i^* = \int_{\underline{\theta}}^{\bar{\theta}} h_i(q^c(\theta);\theta)dF(\theta) - K_i.$$

Now take an interval $\left[\theta',\theta''\right] \subset [\underline{\theta},\bar{\theta}]$ and, for any $\theta \in \left[\theta',\theta''\right]$, let $Q(\theta) \equiv \left[q^c(\theta)-\varepsilon, q^c(\theta)+\varepsilon\right]$, where $\varepsilon > 0$ is chosen so that, for any $\theta \in \left[\theta',\theta''\right]$ and any $q \in Q(\theta)$, $(\theta+\lambda q) \in [m(\underline{\theta}), m(\bar{\theta})]$. Next, note that, for any $\theta \in \left[\theta',\theta''\right]$, the function $h_i(\cdot;\theta)$ defined in (7) is continuously differentiable over $Q(\theta)$ with

$$\begin{aligned}
\frac{\partial h_i(q^c(\theta);\theta)}{\partial q} &= \theta + \lambda\tilde{q}_j(\theta+\lambda q^c(\theta)) - q^c(\theta) - \tfrac{1-F(\theta)}{f(\theta)}\left[1+\lambda\tfrac{\partial \tilde{q}_j(\theta+\lambda q^c(\theta))}{\partial\tilde{\theta}}\right] \\
&= \theta + (\lambda-1)\,q^c(\theta) - \tfrac{1-F(\theta)}{f(\theta)} - \tfrac{1-F(\theta)}{f(\theta)}\lambda\tfrac{\partial \tilde{q}_j(\theta+\lambda q^c(\theta))}{\partial\tilde{\theta}} < 0
\end{aligned}$$

where the inequality follows from the definition of $q^c(\theta)$ and from the fact that $\tilde{q}_j(\cdot)$ is strictly increasing over $[m(\underline{\theta}), m(\bar{\theta})]$. This means that there exists a non-decreasing schedule $q_i : \Theta \to \mathcal{Q}$ such that

$$\int_{\underline{\theta}}^{\bar{\theta}} h_i(q_i(\theta);\theta)dF(\theta) > \int_{\underline{\theta}}^{\bar{\theta}} h_i(q^c(\theta);\theta)dF(\theta). \tag{12}$$

Condition (12) then implies that P_i has a profitable deviation, which contradicts the assumption that $\tilde{\phi}_1^r$ and $\tilde{\phi}_2^r$ satisfy condition (c) in the proposition. We conclude that, when the agent's behavior is Markov, there exists no pair of tariffs that support the collusive schedules as an equilibrium. ■

Proof of Proposition 3. Let $q^*(\cdot)$ be the solution to the differential equation in (9) and $\tilde{q}(\cdot)$ the schedule given in (9). Using the result in Proposition 1, it suffices to show that there exists a scalar $\tilde{K} \geq 0$ such that the pair of schedules $\tilde{q}_i(\cdot) = \tilde{q}(\cdot)$, $i = 1,2$, along with the pair of schedules $\tilde{t}_i(\cdot) = \tilde{t}(\cdot)$, $i = 1,2$, with $\tilde{t}(\cdot)$ defined by

$$\tilde{t}(\tilde{\theta}) = \tilde{\theta}\tilde{q}(\tilde{\theta}) - \int_{\min\tilde{\Theta}}^{\tilde{\theta}} \tilde{q}(s)ds - \tilde{K} \quad \forall\tilde{\theta} \in \tilde{\Theta}$$

satisfy conditions (a)-(c) in Proposition 1. That these schedules satisfy condition (a) is immediate. Thus consider (b). Fix $\phi_j^{r*} = (\tilde{q}(\cdot), \tilde{t}(\cdot))$. Note that, for any $q \in \mathcal{Q}$, the function

$$g_i(\theta,q) \equiv \theta q + v_i^*(\theta,q) - v_i^*(\theta,0) = \theta q + \int_{\theta}^{\theta+\lambda q} \tilde{q}(s)ds$$

is equi-Lipschitz continuous in θ, has the strict increasing difference property, and satisfies the "convex-kink" condition of Assumption 1 in Ely (2001). Theorem 2 in Milgrom and Segal (2002) together with Theorem 2 in Ely (2001) then imply that, given ϕ_j^{r*}, the schedules $(q_i(\cdot), t_i(\cdot))$ satisfy the (IC) and (IR) constraints of condition (c) in Proposition 1 if and only if $q_i(\cdot)$ is nondecreasing and $t_i : \Theta \to \mathbb{R}$ is such that, for any $\theta \in \Theta$,

$$t_i(\theta) = \theta q_i(\theta) + [v_i^*(\theta, q_i(\theta)) - v_i^*(\theta,0)] - \int_{\underline{\theta}}^{\theta} [q_i(s) + \tilde{q}(s+\lambda q_i(s)) - \tilde{q}(s)]ds - K_i' \tag{13}$$

for some $K_i' \geq 0$. Now let $t^*(\cdot)$ be the schedule that is obtained from (13) letting $q_i(\cdot) = q^*(\cdot)$ and setting $K_i' = 0$. By construction, it then follows that each type θ prefers the allocation

$$(q^*(\theta), t^*(\theta), \tilde{q}(m(\theta)), \tilde{t}(m(\theta))) = (q^*(\theta), t^*(\theta), q^*(\theta), \tilde{t}(m(\theta)))$$

to any allocation (q_i, t_i, q_j, t_j) such that $(q_i, t_i) \in \{\left(q^*(\theta'), t^*(\theta')\right) : \theta' \in \Theta\} \cup (0,0)$, and $(q_j, t_j) \in \{(\tilde{q}(\tilde{\theta}'), \tilde{t}(\tilde{\theta}')) : \tilde{\theta}' \in \tilde{\Theta}\} \cup (0,0)$. This also means that the pair of schedules $q' : [m(\underline{\theta}), m(\bar{\theta})] \longrightarrow \mathcal{Q}$ and $t' : [m(\underline{\theta}), m(\bar{\theta})] \longrightarrow \mathbb{R}$ given by

$$q'(\tilde{\theta}) = q^*(m^{-1}(\tilde{\theta})) \text{ and } t'(\tilde{\theta}) = t^*(m^{-1}(\tilde{\theta}))$$

are incentive-compatible over $[m(\underline{\theta}), m(\bar{\theta})]$. But this means that the schedule $t'(\cdot)$ can also be written as

$$t'(\tilde{\theta}) \equiv \tilde{\theta} q'(\tilde{\theta}) - \int_{m(\underline{\theta})}^{\tilde{\theta}} q'(s) ds.$$

Clearly, if P_j offers the mechanism $\phi_j^{r*} = (\tilde{q}(\cdot), \tilde{t}(\cdot))$ and P_i offers the schedules $(q'(\cdot), t'(\cdot))$, it is optimal for each θ to participate and report $m(\theta)$ to each principal. Because for each $\tilde{\theta} \in [m(\underline{\theta}), m(\bar{\theta})]$, $q'(\tilde{\theta}) = \tilde{q}(\tilde{\theta})$ and because $\tilde{q}(\tilde{\theta}) = 0$ for any $\tilde{\theta} < m(\underline{\theta})$, we then have that, for any $\tilde{\theta} \in [m(\underline{\theta}), m(\bar{\theta})]$,

$$t'(\tilde{\theta}) = \tilde{t}(\tilde{\theta}) + \tilde{K}.$$

Furthermore, because for any $\tilde{\theta} > m(\bar{\theta})$, $\left(\tilde{q}(\tilde{\theta}), \tilde{t}(\tilde{\theta})\right) = \left(\tilde{q}(m(\bar{\theta})), \tilde{t}(m(\bar{\theta}))\right) = (q'(\bar{\theta}), t'(\bar{\theta}))$, it is immediate from the aforementioned results that when both principals offer the mechanism $\phi_i^{r*} = (\tilde{q}(\cdot), \tilde{t}(\cdot))$ $i = 1, 2$, with $\tilde{K} = 0$, each type θ finds it optimal to participate in both mechanisms and report $m(\theta)$ to each principal thus obtaining the equilibrium quantity $q^*(\theta)$. In other words, the pair of mechanisms $\phi_i^{r*} = (\tilde{q}(\cdot), \tilde{t}(\cdot))$, $i = 1, 2$, with $\tilde{K} = 0$, satisfies conditions (a) and (b) in the proposition.

It remains to show that, condition (c) also holds. Recall that, given $\phi_j^{r*} = (\tilde{q}(\cdot), \tilde{t}(\cdot))$, a pair of schedules $(q_i(\cdot), t_i(\cdot))$ satisfies the (IC) and (IR) constraints of Proposition 1 if and only if $q_i(\cdot)$ is nondecreasing and $t_i(\cdot)$ is as in (13). This means that the program of condition (c) is equivalent to that in (6). Because, for any $\theta \in int(\Theta)$, the function $h(\cdot; \theta)$ is maximized at $q = q^*(\theta)$, the solution to this program is the function $q^*(\cdot)$ along with $K_i = 0$. To see this, note that the fact that $q^*(\cdot)$ solves the differential equation in (8) implies that the function $h(\cdot; \theta)$ is differentiable at $q = q^*(\theta)$ with derivative

$$\frac{\partial h(q^*(\theta); \theta)}{\partial q} = \theta + \lambda \tilde{q}(\theta + \lambda q^*(\theta)) - q^*(\theta) - \tfrac{1 - F(\theta)}{f(\theta)} \left[1 + \lambda \tfrac{\partial \tilde{q}(\theta + \lambda q^*(\theta))}{\partial \tilde{\theta}}\right] = 0 \tag{14}$$

Together with the fact that $h(\cdot; \theta)$ is quasiconcave then gives the result. ■

Equilibria in the menu auction game of Section 4.2. Consider the menu auction environment of Section 4.2 and assume principals are restricted to make take-it-or-leave-it offers $a_i : E \to [0, 1]$.

Consider first pure-strategy equilibria in which e_1 is selected. In any such equilibrium, necessarily $a_1^*(\cdot)$ is such that $a_1^*(e_1) = 1$; otherwise P_1 could deviate and offer a contract $a_1(\cdot)$ such that $a_1(e_1) = a_1(e_2) = 1$, ensuring that A does not find it profitable to switch to $e = e_2$, and obtaining a higher payoff. But then necessarily $a_2^*(\cdot)$ must be such that $a_2^*(e_1) = 1$; otherwise P_2 could increase $a_2(e_1)$ ensuring A does not find it profitable to change action, and obtaining a higher payoff. We conclude that in any equilibrium in which e_1 is selected, $a_i^*(e_1) = 1$, $i = 1, 2$. That such an equilibrium exists follows from the fact that it can be sustained, for example, by the following contracts $a_i^*(e) = 1$, $i = 1, 2$, $e = e_1, e_2$.

Next, consider equilibria in which e_2 is selected. In these equilibria, necessarily $a_1^*(e_2) = 1/2$. To see this, first suppose that $a_1^*(e_2) < 1/2$. The agent's equilibrium payoff is then strictly higher than 1. But then, necessarily $a_2^*(e_2) = 1$; otherwise P_2 could deviate and offer a contract such that $a_2(e_2) = 1$ and $a_2(e_1) = 0$ which ensures that A does not benefit from switching to $e = e_1$ and gives P_2 a strictly higher payoff. This means that P_1's equilibrium payoff is strictly less than $-1/2$. But then P_1 has a profitable deviation that consists in setting $a_1(e) = 1$ for any e, which induces A to switch to e_1, raising P_1's payoff to at least $-1/2$.

Next, suppose that $a_1^*(e_2) > 1/2$. Then necessarily $a_2^*(e_2) = 1$; this follows from the fact that, given e_2, both the agent's and P_2's payoffs are increasing in a_2. But then necessarily $a_1^*(e_2) = 1$, for otherwise P_1 would have a profitable deviation that consists in setting $a_1(e_1) = a_1(e_2) = 1$ thus inducing A to change action. But this means that P_1's equilibrium payoff is exactly equal to $-1/2$. This in turn also implies that necessarily $a_2^*(e_1) = 1$, for otherwise P_1 would again have a profitable deviation by setting $a_1(e_1) = a_1(e_2) = 1$ which would induce A to switch to $e = e_1$. Furthermore, because A's equilibrium payoff is $1/2$, this also means that necessarily $a_1^*(e_1) \leq 1/4$; else A would benefit from deviating to $e = e_1$. But then P_2 has a profitable deviation that consists in offering a contract such that $a_2(e_1) = 0$ and $a_2(e_2) = 1$ which induces the agent to switch to $e = e_1$.

Thus necessarily $a_1^*(e_2) = 1/2$. Now, because P_1 can guarantee herself at least $-1/2$ by offering a contract such that $a_1(e_1) = a_1(e_2) = 1$ and inducing A to select $e = e_1$, it must be that $a_2^*(e_2) \leq 1/2$. Furthermore, for any $a_2(e_2) < 1/2$, given e_2, both the agent's and P_1's payoff are strictly decreasing in a_1; this implies that there cannot exist equilibria in which $a_2(e_2) < 1/2$. Hence, in any equilibrium in which e_2 is selected, necessarily $a_1^*(e_2) = a_2^*(e_2) = 1/2$. To see that such an equilibrium exists, it then suffices to note that it can be sustained, for example, by the following contracts: $a_1^*(e_1) = a_1^*(e_2) = 1/2$, $a_2^*(e_1) = 1$ and $a_2^*(e_2) = 1/2$. Given the contract offered by P_2, P_1 clearly does not have profitable deviations—this is true whatever the agent's strategy. Furthermore, given the contract offered by P_1, the agent is indifferent about e—this is true whatever the contract offered by P_2. It follows that any deviation by P_2 can be punished by having the agent switching to $e = e_1$. We conclude that an equilibrium that sustains $(e_2^*, a_1^* = a_2^* = 1/2)$ exists. ■

Proof of Theorem 5.

The proof is in two parts. Part 1 proves that if there exists a pure-strategy equilibrium σ^{M^*} of Γ^M that implements the SCF π, there also exists a truthful pure-strategy equilibrium σ^{r*} of $\hat{\Gamma}^r$ that implements the same outcomes. Part 2 proves that any SCF π that can be sustained as an equilibrium of $\hat{\Gamma}^r$ can also be sustained as an equilibrium of Γ^M.

Part 1. Let ϕ^{M*} and σ_A^{M*} denote respectively the equilibrium menus and the continuation equilibrium that support π in Γ^M. Then, for any i, let $\delta_i^*(\theta)$ denote the decision that A takes in equilibrium with P_i when his type is θ.

As a preliminary step, we prove the following result.

Lemma 1 *Suppose the SCF π can be sustained as a pure-strategy equilibrium of Γ^M. Then it can also be sustained as a pure-strategy equilibrium in which the agent's strategy satisfies the following property. For any $k \in \mathcal{N}$, $\theta \in \Theta$ and $\delta_k \in \mathcal{D}_k$, there exists a unique $\delta_{-k}(\theta, \delta_k) \in \mathcal{D}_{-k}$ such that A always selects $\delta_{-k}(\theta, \delta_k)$ with all principals other than k when P_k deviates from the equilibrium menu, the agent's type is θ, the decision A selects with P_k is δ_k, and any principal P_i, $i \neq k$, offered the equilibrium menu.*

Proof of Lemma 1. Let $\tilde{\phi}^M$ and $\tilde{\sigma}_A^M$ denote respectively the equilibrium menus and the continuation equilibrium that support π in Γ^M. Take any $k \in \mathcal{N}$ and for any $(\theta, \delta_k) \in \Theta \times \mathcal{D}_k$ let $\delta_{-k}(\theta, \delta_k)$ be any profile of decisions such that

$$\delta_{-k}(\theta, \delta_k) \in \arg \min_{\delta_{-k} \in D_{-k}(\theta, \delta_k; \tilde{\phi}_{-k}^M)} \underline{U}_k\left(\delta_k, \delta_{-k}, \theta\right) \tag{15}$$

where

$$D_{-k}(\theta, \delta_k; \tilde{\phi}_{-k}^M) \equiv \arg \max_{\delta_{-k} \in \mathrm{Im}(\tilde{\phi}_{-k}^M)} V(\delta_{-k}, \delta_k, \theta)$$

with $\mathrm{Im}(\tilde{\phi}_{-k}^M) \equiv \prod_{j \neq k} \mathrm{Im}(\tilde{\phi}_j^M)$. Now consider the following pure-strategy profile $\mathring{\sigma}^M$. For any $i \in \mathcal{N}$, $\mathring{\sigma}_i^M$ is the pure strategy that prescribes that P_i offers the same menu $\tilde{\phi}_i^M$ as under $\tilde{\sigma}^M$. The continuation equilibrium $\mathring{\sigma}_A^M$ is such that, when either $\phi_i^M = \tilde{\phi}_i^M$ for all i, or $|\{i \in \mathcal{N} : \phi_i^M \neq \tilde{\phi}_i^M\}| > 1$, then $\mathring{\sigma}_A^M(\theta, \phi^M) = \tilde{\sigma}_A^M(\theta, \phi^M)$, for any θ. When instead ϕ^M is such that $\phi_i^M = \tilde{\phi}_i^M$ for all $i \neq k$, while $\phi_k^M \neq \tilde{\phi}_i^M$ for some $k \in \mathcal{N}$, then each type θ selects a profile of decisions (δ_k, δ_{-k}) such that δ_k is the same decision that, given the menus $(\tilde{\phi}_{-k}^M, \phi_k^M)$, type θ would have selected with P_k according to the original strategy $\tilde{\sigma}_A^M$ whereas $\delta_{-k} = \delta_{-k}(\theta, \delta_k)$, as defined in (15). Given any profile of contracts y selected by the lotteries (δ_k, δ_{-k}), the effort the agent selects is then $\xi_k(\theta, y)$ as defined in (2).

It is immediate that the behavior prescribed by the strategy $\mathring{\sigma}_A^M$ is sequentially rational for the agent. Furthermore, given $\mathring{\sigma}_A^M$, a principal P_i who expects all other principals to offer the equilibrium menus $\tilde{\phi}_{-i}^M$ cannot do better than offering the equilibrium menu $\tilde{\phi}_i^M$. We conclude that $\mathring{\sigma}^M$ is a pure-strategy equilibrium of Γ^M that sustains the same SCF as $\tilde{\sigma}^M$. ■

Hence, without loss, assume σ^{M*} satisfies the property of Lemma 1. For any $i, k \in \mathcal{N}$ with $k \neq i$, and for any $(\theta, \delta_k) \in \Theta \times \mathcal{D}_k$, let $\delta_i(\theta, \delta_k)$ denote the unique decision that A selects with P_i when his type is θ, the decision taken with P_k is δ_k, and the menus offered are $\phi_j^M = \phi_j^{M*}$ for all $j \neq k$, and $\phi_k^M \neq \phi_k^{M*}$.

Next, consider the following strategy profile $\hat{\sigma}^{r*}$ for $\hat{\Gamma}^r$. Each principal offers a direct mechanism $\hat{\phi}_i^{r*}$ such that, for any $(\theta, \delta_{-i}, k) \in \Theta \times \mathcal{D}_{-i} \times \mathcal{N}_{-i}$,

$$\hat{\phi}_i^{r*}(\theta, \delta_{-i}, k) = \begin{cases} \delta_i^*(\theta) \text{ if } k = 0 \text{ and } \delta_{-i} = \delta_{-i}^*(\theta) \\ \delta_i(\theta, \delta_k) \text{ if } k \neq 0 \text{ and } \delta_{-i} \text{ is such that } \delta_j = \delta_j(\theta, \delta_k) \text{ for all } j \neq i, k \\ \delta_i \in \arg\max_{\delta_i' \in \mathrm{Im}(\phi_i^{M*})} V(\delta_{-i}, \delta_i', \theta) \text{ in all other cases.} \end{cases}$$

By construction, $\hat{\phi}_i^{r*}$ is incentive compatible. Now consider the following strategy $\hat{\sigma}_A^{r*}$ for the agent in $\hat{\Gamma}^r$.

(i) Given the equilibrium mechanisms $\hat{\phi}^{r*}$, each type θ reports a message $\hat{m}_i^r = (\theta, \delta_{-i}^*(\theta), 0)$ to each P_i. Given any profile of contracts y selected by the lotteries $\delta^*(\theta)$, the agent then mixes over E with the same distribution he would have used in Γ^M given $(\theta, \phi^{M*}, m^*(\theta), y)$ where $m^*(\theta) \equiv \delta^*(\theta)$ are the equilibrium messages that type θ would have sent in Γ^M given the equilibrium menus ϕ^{M*}.

(ii) Given any profile of mechanisms $\hat{\phi}^r$ such that $\hat{\phi}_i^r = \hat{\phi}_i^{r*}$ for all $i \neq k$, while $\hat{\phi}_k^r \neq \hat{\phi}_k^{r*}$ for some $k \in \mathcal{N}$, let δ_k denote the decision that type θ would have induced with P_k in Γ^M had the menus offered been $\phi^M = (\phi_{-k}^{M*}, \phi_k^M)$ where ϕ_k^M is the menu whose image $\mathrm{Im}(\phi_k^M) = \mathrm{Im}(\hat{\phi}_k^r)$. The strategy $\hat{\sigma}_A^{r*}$ then prescribes that type θ reports to P_k any message m_k^r such that $\phi_k^r(m_k^r) = \delta_k$ and then reports to any other principal P_i, $i \neq k$, the message $\hat{m}_i^r = (\theta, \delta_{-i}, k)$, with

$$\delta_{-i} = (\delta_k, (\delta_j(\theta, \delta_k))_{j \neq i,k}).$$

Given any contracts y selected by the lotteries $\delta = (\delta_k, \delta_j(\theta, \delta_k)_{j \neq k})$, A then selects effort $\xi_k(\theta, y)$, as defined in (2).

(iii) Finally, for any profile of mechanisms $\hat{\phi}^r$ such that the $|\{i \in \mathcal{N} : \hat{\phi}_i^r \neq \hat{\phi}_i^{r*}\}| > 1$, simply let $\hat{\sigma}_A^r(\theta, \phi^r)$ be any strategy that is sequentially optimal for A given $(\theta, \hat{\phi}^r)$.

The behavior prescribed by the strategy $\hat{\sigma}_A^{r*}$ is clearly a continuation equilibrium. Furthermore, given $\hat{\sigma}_A^{r*}$, any principal P_i who expects all other principals to offer the equilibrium mechanisms $\hat{\phi}_{-i}^{r*}$ cannot do better than offering the equilibrium mechanism $\hat{\phi}_i^{r*}$, for any $i \in \mathcal{N}$. We conclude that the strategy profile $\hat{\sigma}^{r*}$ in which each P_i offers the mechanism $\hat{\phi}_i^{r*}$ and A follows the strategy $\hat{\sigma}_A^*$ is a truthful pure-strategy equilibrium of $\hat{\Gamma}^r$ and sustains the same SCF π as σ^{M*} in Γ^M.

Part 2. We now prove that if there exists an equilibrium $\hat{\sigma}^r$ of $\hat{\Gamma}^r$ that sustains the SCF π, then there also exists an equilibrium σ^{M*} of Γ^M that sustains the same SCF. For any $i \in \mathcal{N}$ and any $\phi_i^M \in \Phi_i^M$, let $\hat{\Phi}_i^r(\phi_i^M) \equiv \{\hat{\phi}_i^r \in \Phi_i^r : \mathrm{Im}(\hat{\phi}_i^r) = \mathrm{Im}(\phi_i^M)\}$ denote the set of revelation mechanisms with the same image as ϕ_i^M. The proof then follows from the same steps as in the proof of Part 2 in Theorem 2 replacing the mappings $\Phi_i^r(\cdot)$ with the mappings $\hat{\Phi}_i^r(\cdot)$ and with the following adjustment for *Case 2*. For any ϕ^M such that there exists a $j \in \mathcal{N}$ such that $\hat{\Phi}_i^r(\phi_i^M) \neq \varnothing$ for all $i \neq j$ while $\hat{\Phi}_j^r(\phi_j^M) = \varnothing$, let $\hat{\phi}_j^r$ be any arbitrary revelation mechanism such that

$$\hat{\phi}_j^r(\theta, \delta_{-j}, k) \in \arg \max_{\delta_j \in \mathrm{Im}(\phi_j^M)} V(\delta_j, \delta_{-j}, \theta) \qquad \forall\, (\theta, \delta_{-j}, k) \in \Theta \times \mathcal{D}_{-j} \times \mathcal{N}_{-j}.$$

For any $\theta \in \Theta$, the strategy $\sigma_A^{M*}(\theta, \phi^M)$ then induces the same distribution over outcomes as the strategy $\hat{\sigma}_A^{r*}$ given $\hat{\phi}_j^r$ and given $\hat{\phi}_{-j}^r \in \hat{\Phi}_{-j}^r(\phi_{-j}^M) \equiv \prod_{i \neq j} \hat{\Phi}_i^r(\phi_i^M)$ in the sense of (11). ■

References

[1] Attar, A., Piaser G. and N. Porteiro, 2007a, "Negotiation and take-it or leave-it in common agency with non-contractible actions," *Journal of Economic Theory*, 135(1), 590-593.

[2] Attar, A. Piaser G., and N. Porteiro, 2007b, "A Note on Common Agency Models of Moral Hazard," *Economics Letters*, forthcoming.

[3] Attar, A., Majumadar D., Piaser G., and N. Porteiro, 2007, "Common Agency Games with Separable Preferences," mimeo, University of Venice.

[4] Bernheim, B. D. and M. Whinston, 1985, "Common Marketing Agency as a Device for Facilitating Collusion," *RAND Journal of Economics* (16), 269-81.

[5] Bernheim, D. and M. Whinston, 1986a, "Menu Auctions, Resource Allocations and Economic Influence," *Quarterly Journal of Economics*, 101, 1-31.

[6] Bernheim, D. and M. Whinston, 1986b, "Common Agency," *Econometrica*, 54(4), 923-942.

[7] Biais, B., D. Martimort, and J.-C. Rochet, 2000, "Competing Mechanisms in a Common Value Environment," *Econometrica*, 68, 799-837.

[8] Calzolari, G., 2004, "Incentive Regulation of Multinational Enterprises," *International Economic Review*, 45 (1), 257-282.

[9] Dixit, A., G. Grossman, and E. Helpman, 1997 "Common agency and coordination: General theory and application to government policymaking," *Journal of Political Economy*, 752-769.

[10] Dudley, R. M., 2002, *Real Analysis and Probability*, Cambridge Studies in Advanced Mathematics No. 74, Cambridge University Press.

[11] Ely, Jeffrey, 2001, "Revenue Equivalence without Differentiability Assumptions," mimeo Northwestern University.

[12] Epstein, L. and M. Peters, 1999, “A Revelation Principle for Competing Mechanisms,” *Journal of Economic Theory*, 88, 119-160.

[13] Gibbard, A., 1977, “Manipulation of Voting Schemes,” *Econometrica*, 41, 587-601.

[14] Green, J. and J-J. Laffont, 1977, “Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods,” *Econometrica*, 45, 427-438.

[15] Grossman, G. M., and H., Elhanan, 1994, "Protection for Sale," *American Economic Review*, 84: 833-50.

[16] Guesnerie, R. “A Contribution to the Pure Theory of Taxation,” Cambridge University Press.

[17] Katz, M., 1991, “Game Playing Agents: Unobservable Contracts as Precommitments,” *Rand Journal of Economics*, 22(1), 307-328.

[18] Klemperer, P. D., and M. A. Meyer , 1989 “Supply Function Equilibria in Oligopoly under Uncertainty,” *Econometrica*, 57(6), 1243-77.

[19] Martimort, D., 1992, “ Multi-Principaux avec Sélection Adverse,” *Annales d’Economie et de Statistique*, 28, 1-38.

[20] Martimort, D., 1996, “Exclusive Dealing, common agency, and multi-principals incentive theory,” *RAND Journal of Economics,* 27(1), 1-31.

[21] Martimort, D. and L. Stole, 1997, “Communication Spaces, Equilibria Sets and the Revelation Principle under Common Agency,” mimeo, University of Toulouse.

[22] Martimort, D. and L. Stole, 2002, “The Revelation and Delegation Principles in Common Agency Games,” *Econometrica* 70, 1659-1674.

[23] Martimort and Stole, 2003, “Contractual Externalities and Common Agency Equilibria”, *Advances in Theoretical Economics*: Vol. 3: No. 1, Article 4.

[24] Martimort, D. and L. Stole, 2005, “Common Agency Games with Common Screening Devices,” mimeo, University of Chicago GSB and Toulouse University.

[25] McAfee, P., 1993, “Mechanism Design by Competing Sellers,” *Econometrica*, 61(6), 1281-1312.

[26] Mezzetti, C., 1997, “Common Agency with Horizontally Differentiated Principals,” *Rand Journal of Economics*, 28: 323-345.

[27] Milgrom, P. and I Segal, 2002, “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica*, 70(2), 583–601.

[28] Myerson, R., 1979, "Incentive Compatibility and the Bargaining Problem," *Econometrica*, 47, 61-73.

[29] Olsen, T. and P. Osmundsen, 2001 "Strategic tax competition: implications of national ownership," *Journal of Public Economics*, Vol. 81, no. 2, 253-277.

[30] Parlour C.A. and U. Rajan, 2001 "Competition in Loan Contracts," *American Economic Review*, 91(5): 1311-28.

[31] Pavan, A. and G. Calzolari, 2008, "Sequential Contracting with Multiple Principals," *Journal of Economic Theory*, forthcoming.

[32] Peck, J., 1997, "A Note on Competing Mechanisms and the Revelation Principle," Ohio State University mimeo.

[33] Peters, M., 2001, "Common Agency and the Revelation Principle," *Econometrica*, 69, 1349-1372.

[34] Peters, M., 2003, "Negotiations versus take-it-or-leave-it in common agency," *Journal of Economic Theory*, 111, 88-109.

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

[36] Piaser, G., 2007, "Direct Mechanisms, Menus and Latent Contracts," mimeo, University of Venice.

[37] Rochet, J-C., 1986, "Le Controle Des Equations Aux Derivees Partielles Issues de la Theorie Des Incitations," PhD Thesis Universite' Paris IX.

[38] Segal, I. and M. Whinston, 2003, "Robust Predictions for Bilateral Contracting with Externalities," *Econometrica*, 71, 757-792.

[39] Yamashita, T., 2007, "A Revelation Principle and A Folk Theorem without Repetition in Games with Multiple Principals and Agents," mimeo Hitotsubashi University.

Supplementary Material for
Truthful Revelation Mechanisms for Simultaneous Common Agency Games

Alessandro Pavan Northwestern University

Giacomo Calzolari University of Bologna

November, 2008

Proof of Theorem 6

The proof is in 2 parts.[1] Part 1 proves that for any equilibrium σ^M of Γ^M, there exists an equilibrium $\tilde{\sigma}^r$ of $\tilde{\Gamma}^r$ that implements the same outcomes. Part 2 proves the converse.

Part 1. Let $\mathcal{Q}_i$ be a generic partition of Φ_i^M and denote by $Q_i \in \mathcal{Q}_i$ a generic element of $\mathcal{Q}_i$. Consider now a partition-game $\Gamma^{\mathcal{Q}}$ in which each P_i chooses an element of $\mathcal{Q}_i$, then A selects a profile of menus $\phi^M = (\phi_1^M, ..., \phi_n^M)$ one from each Q_i, chooses the lotteries δ and given the contracts y determined by the lotteries δ, he finally chooses effort $e \in E$.

The proof of part 1 is in two steps. Step 1 identifies a collection of partitions $\mathcal{Q} = (\mathcal{Q}_i)_{i\in\mathcal{N}}$ such that the agent's payoff is the same for any pair of menus $\phi_i^M, \phi_i^{M\prime} \in Q_i$, $i = 1, ..., n$. It then shows that, for any $\sigma^M \in \mathcal{E}(\Gamma^M)$ there exists a $\mathring{\sigma} \in \mathcal{E}(\Gamma^{\mathcal{Q}})$ which implements the same outcomes. Step 2 uses the equilibrium $\mathring{\sigma}$ of $\Gamma^{\mathcal{Q}}$ constructed in Step 1 to prove existence of a truthful equilibrium $\tilde{\sigma}^r$ of $\tilde{\Gamma}^r$ which also supports the same outcomes as σ^M.

Step 1. Take a generic collection of partitions $\mathcal{Q} = (\mathcal{Q}_i)_{i\in\mathcal{N}}$, one for each Φ_i^M, $i = 1, ..., n$ with $\mathcal{Q}_i$ consisting of measurable sets.[2] Consider the following strategy profile $\mathring{\sigma}$ for the partition game $\Gamma^{\mathcal{Q}}$. For any P_i, let $\mathring{\sigma}_i \in \Delta(\mathcal{Q}_i)$ be the distribution over $\mathcal{Q}_i$ induced by the equilibrium strategy σ_i^M of Γ^M. That is, for any subset R_i of $\mathcal{Q}_i$ whose union is measurable,

$$\mathring{\sigma}_i(R_i) = \sigma_i^M(\bigcup R_i).$$

Next consider the agent. For any $Q = (Q_1, ..., Q_n) \in \prod_{i\in\mathcal{N}} \mathcal{Q}_i$, A selects the menus ϕ^M from $\prod_{i\in\mathcal{N}} Q_i$ using the distribution $\mathring{\sigma}_A(\cdot|Q) \equiv \sigma_1^M(\cdot|Q_1) \times \cdots \times \sigma_n^M(\cdot|Q_n)$, where for each Q_i, $\sigma_i^M(\cdot|Q_i)$ is the

[1] The notation follows that in the paper.

[2] In the sequel, we assume that any set of mechanisms Φ_i^M is a Polish space and whenever we talk about measurability, we mean with respect to the Borel σ-algebra Σ on Φ_i^M.

regular conditional distribution over Φ_i^M that is obtained from the equilibrium strategy σ_i^M of P_i conditioning on $\phi_i^M \in Q_i$.[3] After selecting the menus ϕ^M, A follows the same behavior prescribed by the strategy σ_A^M for Γ^M.

Now, fix the agent's strategy $\tilde{\sigma}_A$ as described above. It is immediate that, whatever the partitions $\mathcal{Q}$, the strategies $(\mathring{\sigma}_i)_{i\in\mathcal{N}}$ constitute an equilibrium for the game $\Gamma^{\mathcal{Q}}(\mathring{\sigma}_A)$ among the principals.

In what follows, we identify a collection of partitions $\mathcal{Q}$ that make $\mathring{\sigma}_A$ sequentially optimal for the agent. Consider the equivalence relation $\sim_i$ defined as follows: given any two menus ϕ_i^M and $\phi_i^{M\prime}$,

$$\phi_i^M \sim_i \phi_i^{M\prime} \Longleftrightarrow h_\theta(\delta_{-i};\phi_i^M) = h_\theta(\delta_{-i};\phi_i^{M\prime})\ \forall(\theta,\delta_{-i}),$$

where, for any mechanism ϕ_i, $h_\theta(\delta_{-i};\phi_i) \equiv \arg\max_{\delta_i\in \mathrm{Im}(\phi_i)} V(\delta_i,\delta_{-i},\theta)$.

Now, let $\mathcal{Q} = (\mathcal{Q}_i)_{i\in\mathcal{N}}$ be the collection of partitions generated by the equivalence relations $\sim_i$, $i = 1, ..., n$. It is immediate that, in the partition game $\Gamma^{\mathcal{Q}}$, $\mathring{\sigma}_A$ is sequentially optimal for A. We conclude that for any $\sigma^M \in \mathcal{E}(\Gamma^M)$ there exists a $\mathring{\sigma} \in \mathcal{E}(\Gamma^{\mathcal{Q}})$ which implements the same outcomes as σ^M.

Step 2. We next prove that starting from $\mathring{\sigma}$, one can construct a truthful equilibrium $\tilde{\sigma}^r$ for $\tilde{\Gamma}^r$ that also implements the same outcomes as σ^M in Γ^M. For any $i \in \mathcal{N}$ and $Q_i \in \mathcal{Q}_i$, let $h_\theta(\delta_{-i};Q_i) \equiv h_\theta(\delta_{-i};\phi_i^M)$ for some $\phi_i^M \in Q_i$. Since for any two menus $\phi_i^M, \phi_i^{M\prime} \in Q_i$, $h_\theta(\delta_{-i};\phi_i^M) = h_\theta(\delta_{-i};\phi_i^{M\prime})$ for all (θ,δ_{-i}), then $h_\theta(\delta_{-i};Q_i)$ is uniquely determined by Q_i. Now, for any $Q_i \in \mathcal{Q}_i$, let $\tilde{\phi}_i^r\Big|_{Q_i} \in \tilde{\Phi}_i^r$ denote the revelation mechanism given by

$$\tilde{\phi}_i^r(\theta,\delta_{-i}) = h_\theta(\delta_{-i};Q_i) \quad \forall\,(\theta,\delta_{-i}) \in \Theta\times\mathcal{D}_{-i}. \tag{1}$$

For any set of mechanisms $B \subseteq \tilde{\Phi}_i^r$, then let $\mathcal{Q}_i(B) \equiv \{Q_i \in \mathcal{Q}_i : \tilde{\phi}_i^r\Big|_{Q_i} \in B\}$ denote the set of corresponding cells in $\mathcal{Q}_i$. The strategy $\tilde{\sigma}_i^r \in \Delta(\tilde{\Phi}_i^r)$ for P_i is given by

$$\tilde{\sigma}_i^r(B) = \mathring{\sigma}_i(\mathcal{Q}_i(B))\ \forall B \subseteq \tilde{\Phi}_i^r.$$

Next, consider the agent. Given any profile of mechanisms $\tilde{\phi}^r \in \tilde{\Phi}^r$, let $Q(\tilde{\phi}^r) = (Q_i(\tilde{\phi}_i^r))_{i\in\mathcal{N}} \in \prod_{i\in\mathcal{N}} \mathcal{Q}_i$ denote the profile of cells in $\Gamma^{\mathcal{Q}}$ such that, for any $i \in \mathcal{N}$, the cell $Q_i(\tilde{\phi}_i^r)$ is such that $h_\theta(\delta_{-i};Q_i) = \tilde{\phi}_i^r(\delta_{-i},\theta)$ for any $(\theta,\delta_{-i}) \in \Theta\times\mathcal{D}_{-i}$. Now, let $\tilde{\sigma}_A^r$ be any truthful strategy that implements the same distribution over $\mathcal{A}\times E$ as $\mathring{\sigma}_A$ given $Q(\phi^r)$. Precisely, let $\rho_{\sigma_A} : \Theta\times\Phi \to \Delta(\mathcal{A}\times E)$ denote the distribution over outcomes induced by σ_A in Γ. Then $\tilde{\sigma}_A^r$ is any truthful strategy such that, for any $(\theta,\tilde{\phi}^r) \in \Theta\times\tilde{\Phi}^r$,

$$\rho_{\tilde{\sigma}_A^r}(\theta,\tilde{\phi}^r) = \rho_{\mathring{\sigma}_A}(\theta,Q(\tilde{\phi}^r)) \equiv \int_{\Phi_1^M}\cdots\int_{\Phi_n^M} \rho_{\sigma_A^M}(\theta,\phi^M)d\sigma_1^M(\phi_1^M|Q_1(\tilde{\phi}_1^r))\times\cdots\times d\sigma_n^M(\phi_n^M|Q_n(\tilde{\phi}_n^r)).$$

[3] Assuming that each Φ_i^M is a Polish space endowed with the Borel σ-algebra Σ_i, the existence of such a conditional probability measure follows from Theorem 10.2.2 in Dudley (2002, p. 345).

The strategy $\tilde{\sigma}_A^r$ is clearly optimal for A. Furthermore, given $\tilde{\sigma}_A^r$, the strategy profile $(\tilde{\sigma}_i^r)_{i\in\mathcal{N}}$ is an equilibrium for the game among the principals. We conclude that $\tilde{\sigma}^r = (\tilde{\sigma}_A^r, (\tilde{\sigma}_i^r)_{i\in\mathcal{N}})$ is an equilibrium for $\tilde{\Gamma}^r$ and sustains the same outcomes as σ^M in Γ^M.

Part 2. We now prove the converse: Given an equilibrium $\tilde{\sigma}^r$ of $\tilde{\Gamma}^r$ that sustains the SCF π, there exists an equilibrium σ^M of Γ^M that sustains the same SCF.

For any $i \in \mathcal{N}$, let $\alpha_i : \tilde{\Phi}_i^r \to \Phi_i^M$ denote the injective mapping defined by the relation

$$\mathrm{Im}(\alpha_i(\tilde{\phi}_i^r)) = \mathrm{Im}(\tilde{\phi}_i^r) \ \forall \tilde{\phi}_i^r \in \tilde{\Phi}_i^r$$

and $\alpha_i(\tilde{\Phi}_i^r) \subset \Phi_i^M$ denote the range of $\alpha_i(\cdot)$. For any $\phi_i^M \in \alpha_i(\tilde{\Phi}_i^r)$, then let $\alpha_i^{-1}(\phi_i^M)$ denote the unique revelation mechanism such that $\mathrm{Im}(\tilde{\phi}_i^r) = \mathrm{Im}(\phi_i^M)$.

Now consider the following strategy for the agent in Γ^M. For any ϕ^M such that, for all $i \in \mathcal{N}$, $\phi_i^M \in \alpha_i(\tilde{\Phi}_i^r)$, let σ_A^M be such that $\rho_{\sigma_A^M}(\theta, \phi^M) = \rho_{\tilde{\sigma}_A^r}(\theta, \alpha^{-1}(\phi^M))$, where $\alpha^{-1}(\phi^M) \equiv (\alpha_i^{-1}(\phi_i^M))_{i=1}^n$. If instead ϕ^M is such that $\phi_j^M \in \alpha_j(\tilde{\Phi}_j^r)$ for all $j \neq i$, while for i, $\phi_i^M \notin \alpha_i(\tilde{\Phi}_i^r)$, then let σ_A^M be such that $\rho_{\sigma_A^M}(\theta, \phi^M) = \rho_{\tilde{\sigma}_A^r}(\theta, \tilde{\phi}_i^r, (\alpha_j^{-1}(\phi_j^M)_{j\neq i})$ where $\tilde{\phi}_i^r$ is any revelation mechanism that satisfies

$$\tilde{\phi}_i^r\,(\theta, \delta_{-i}) = h_\theta(\delta_{-i}; \phi_i^M) \qquad \forall\,(\theta, \delta_{-i}) \in \Theta \times \mathcal{D}_{-i}.$$

Finally, for any ϕ^M such that $|\{j \in \mathcal{N} : \phi_j^M \notin \alpha_j(\tilde{\Phi}_j^r)\}| > 1$, simply let σ_A^M be any rational response for the agent given (θ, ϕ^M). It is immediate that the strategy σ_A^M constitutes a continuation equilibrium for Γ^M.

Now consider the following strategy profile for the principals For any $i \in \mathcal{N}$, let $\sigma_i^M = \alpha_i(\tilde{\sigma}_i^r)$, where $\alpha_i(\tilde{\sigma}_i^r)$ denotes the randomization over Φ_i^M obtained from the strategy $\tilde{\sigma}_i^r$ using the mapping α_i. Formally, for any measurable set $B \subseteq \Phi_i^M$, $\sigma_i^M(B) = \tilde{\sigma}_i^r(\{\tilde{\phi}_i^r : \alpha_i(\tilde{\phi}_i^r) \in B\})$. It is straight forward to see that any principal P_i who expects the agent to follow the strategy σ_A^M and any other principal P_j to follow the strategy $\sigma_j^M = \alpha_j(\tilde{\sigma}_j^r)$ cannot do better than following the strategy $\sigma_i^M = \alpha_i(\tilde{\sigma}_i^r)$. We conclude that σ^M is an equilibrium of Γ^M and sustains the same SCF π as $\tilde{\sigma}^r$ in $\tilde{\Gamma}^r$.