American Economic Journal: Microeconomics 2 (May 2010): 132–190
http://www.aeaweb.org/articles.php?doi=10.1257/mic.2.2.132

Truthful Revelation Mechanisms for Simultaneous Common Agency Games[†]

By Alessandro Pavan and Giacomo Calzolari*

We introduce new revelation mechanisms for simultaneous common agency games which, although they do not always permit a complete equilibrium characterization, do facilitate the characterization of the equilibrium outcomes that are typically of interest in applications. We then show how these mechanisms can be used in applications 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 all possible equilibrium outcomes, including those sustained by non-Markov strategies and/or mixed-strategy profiles. (*JEL* C72, D82, D86)

Many economic environments can be modelled as common agency games—that is, games where multiple principals contract simultaneously and noncooperatively with the same agent.[1] Despite their relevance for applications, the analysis of these games has been made difficult by the fact that one cannot safely assume that the agent selects a contract with each principal by simply reporting his "type" (i.e., his exogenous payoff-relevant information). In other words, the central tool of mechanism design theory—the Revelation Principle—is invalid in these games.[2] The reason is that the agent's preferences over the contracts offered by one principal depend not only on his type, but also on the contracts he has been offered by the other principals.[3]

* Pavan: Department of Economics, Northwestern University, 2001 Sheridan Road Arthur Andersen Hall 329 Evanston, IL 60208-2600 (e-mail: alepavan@northwestern.edu). Calzolari: Department of Economics University of Bologna, Piazza Scaravilli 2 40126 Bologna Italy and CEPR (e-mail: giacomo.calzolari@unibo.it). This is a substantial revision of an earlier paper that circulated under the same title. For useful discussions, we thank seminar participants at various conferences and institutions where this paper has been presented. A special thank is to Eddie Dekel, Mike Peters, Marciano Siniscalchi, Jean Tirole, and three anonymous referees for suggestions that helped us improve the paper. We are grateful to Itai Sher for excellent research assistance. For its hospitality, Pavan also thanks Collegio Carlo Alberto (Turin), where this project was completed.

† To comment on this article in the online discussion forum, or to view additional materials, visit the articles page at http://www.aeaweb.org/articles.php?doi=10.1257/mic.2.2.132.

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

[2] For the Revelation Principle, see, among others, Allan Gibbard (1973), Jerry Green and Jean-Jacques Laffont (1977), and Roger B. Myerson (1979). Problems with the Revelation Principle in games with competing principals have been documented in Michael L. Katz (1991), R. Preston McAfee (1993), James Peck (1997), Larry Epstein and Michael Peters (1999), Peters (2001), and David Martimort and Lars Stole (1997, 2002), among others. Recent work by Peters (2003, 2007); Andrea Attar, Gwenael Piaser, and Nicolàs Porteiro (2007a, 2007b); and Attar et al. (2008) has identified special cases in which these problems do not emerge.

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

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.[4] However, describing a mechanism requires an appropriate language. The main contribution of Epstein and Peters (1999) is in proving the existence of a universal language that is rich enough to describe all possible mechanisms. This language also permits one to identify a class of universal mechanisms with the property that any indirect mechanism can be embedded into this class. 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 this result is a remarkable contribution, the use of universal mechanisms in applications has been impeded 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 the fact that principal j's mechanism may also ask the agent to describe principal i's mechanism, as well as how this mechanism depends on principal j's mechanism, and so on, leading to the problem of infinite regress. The universal language is obtained as the limit of a sequence of enlargements of the message space, where at each enlargement the corresponding direct mechanism becomes more complex, and thus more difficult to describe and to use when searching for equilibrium outcomes.

The second solution, proposed by Peters (2001) and Martimort and Lars Stole (2002), is to restrict the principals to offering menus of contracts. These authors 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 offering menus of contracts that sustains the same outcomes. In this equilibrium, the principals simply offer the menus they would have offered through the equilibrium mechanisms of the original game, and then delegate to the agent the choice of the contracts. This result is referred to in the literature as the *Menu Theorem* and is the analog of the Taxation Principle for games with a single mechanism designer.[5]

The Menu Theorem has proved quite useful in certain applications. However, contrary to the Revelation Principle, it provides no indication as to which contracts the agent selects from the menus, nor does it permit one to restrict attention to a particular set of menus.[6]

The purpose of this paper is to show that, in most cases of interest for applications, one can still conveniently describe the agent's choice from a menu (equivalently, the outcome of his interaction with each principal) through a *revelation mechanism*. The structure of these mechanisms is, however, more general than the standard one for games with a single mechanism designer. Nevertheless, contrary to universal mechanisms, it does not lead to any infinite regress problem. In the revelation mechanisms we propose, the agent is asked to report his exogenous type

[4] A mechanism is simply a procedure for selecting a contract.

[5] The result is also referred to as the "Delegation Principle" (e.g., in Martimort and Stole 2002). For the Taxation Principle, see Jean Charles Rochet (1986) and Roger Guesnerie (1995).

[6] The only restriction discussed in the literature is that the menus should not contain dominated contracts (see Martimort and Stole 2002).

along with the endogenous payoff-relevant contracts chosen with the other principals. As is standard, a revelation mechanism is then said to be *incentive-compatible* if the agent finds it optimal to report such information truthfully.

Describing the agent's choice from a menu by means of an incentive-compatible revelation mechanism is convenient because it permits one to specify which contracts the agent selects from the menu in response to possible deviations by the other principals, without, however, having to describe such deviations (which would require the use of the universal language to describe the mechanisms offered by the deviating principals); what the agent is asked to report is only the contracts selected as a result of such deviations. This, in turn, can facilitate the characterization of the equilibrium outcomes.

The mechanisms described above are appealing because they capture the essence of common agency, i.e., the fact that the agent's preferences over the contracts offered by one principal depend not only on the agent's type, but also on the contracts selected with the other principals.[7] However, this property alone does not guarantee that one can always safely restrict the agent's behavior to depending only on such payoff-relevant information. In fact, when indifferent, the agent may also condition his choice on payoff-irrelevant information, such as the contracts included by the other principals in their menus, but which the agent decided not to select. Furthermore, when indifferent, the agent may randomize over the principals' contracts, inducing a correlation that cannot always be replicated by having the agent simply report to each principal his type along with the contracts selected with the other principals. As a consequence, not all equilibrium outcomes can be sustained through the revelation mechanisms described above. While we find these considerations intriguing from a theoretical viewpoint, we seriously doubt their relevance in applications.

Our concerns with mixed-strategy equilibria come from the fact that 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 nondegenerate examples in which the agent is made indifferent over some of the contracts offered by the principals; and 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 contracts that are most favorable to her (see the example discussed in Section IV).

We also have concerns about equilibrium outcomes sustained by a strategy for the agent that is not Markovian, i.e., that also depends on payoff-irrelevant information. These concerns 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. On the one hand, 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. On the other hand, it does not seem plausible that, *for a given contract with seller* j, the purchase from seller i would

[7] A special case is when preferences are separable, as in Andrea Attar et al. (2008), in which case they depend only on the agent's exogenous type.

depend on payoff-irrelevant information, such as which other contracts offered by seller *j* did the buyer decide not to choose.[8]

For most of the analysis here, we focus on outcomes sustained by pure-strategy profiles in which the agent's behavior in each relationship is Markovian.[9] We first show that any such outcome can be sustained by 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 in the game where principals can offer any menus. 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 mechanisms.

We then proceed by addressing the question of whether there exist environments in which making the assumption that the agent follows a Markov strategy is not only appealing, but actually unrestrictive. 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 never precludes the possibility of sustaining all equilibrium outcomes when information is complete, and the principals' preferences are sufficiently aligned. By sufficiently aligned we mean that, given the contracts signed with all principals other than *i*, the specific contract that the agent signs with principal *i* to punish a deviation by one of the other principals does not need to depend on the identity of the deviating principal. See the definition of "Uniform Punishment" in Section II. This property is always satisfied when there are only two principals. It is also satisfied when the principals are, for example, retailers competing "à la Cournot" in a downstream market. Each retailer's payoff then decreases with the quantity that 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 the following. It rules out the possibility that the equilibrium outcomes are sustained by the agent punishing a deviation, say, by principal *j*, by choosing the equilibrium contracts with all principals other than *i*, and then choosing with principal *i* a contract different from the equilibrium one. In games with incomplete information, allowing the agent to change his behavior with a nondeviating principal, despite the fact that he is selecting the equilibrium contracts with all the other principals, may be essential for punishing certain deviations. This in turn implies that Markov strategies need not support all equilibrium outcomes in such games. 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*

[8] That the agent's behavior is Markovian does not imply that the principals can be restricted to offering menus that contain only the contracts (e.g., the price-quantity pairs) that are selected in equilibrium. As is well known, the inclusion in the menu of "latent" contracts that are never selected in equilibrium may be essential to prevent deviations by other principals. See Gabriella Chiesa and Vincenzo Denicolo (2009) for an illustration.

[9] 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 2009), it shares with that definition the idea that the agent's behavior should depend only on payoff-relevant information.

to Equilibrium." This refinement simply requires that each type of agent selects the equilibrium contract with each principal when the latter offers the equilibrium menu, and when the contracts selected with the other principals are the equilibrium ones.[10] Again, in many real world situations, such behavior seems plausible.

While we find the restriction to pure-strategy-Markov equilibria reasonable and appealing for most applications, at the end of the paper, we also show how one can enrich our revelation mechanisms (albeit at the cost of an increase in complexity) to characterize equilibrium outcomes sustained by non-Markov strategies and/or mixed strategy profiles. For the former, it suffices to consider revelation mechanisms where, in addition to his type and the contracts he has selected with the other principals, the agent is asked to report the *identity of a deviating principal* (if any). For the latter, it suffices to consider *set-valued* revelation mechanisms that respond to each report about the agent's type and the contracts selected with the other principals with a set of contracts that are equally optimal for the agent among those available in the mechanism. Giving the same type of the agent the possibility of choosing different contracts in response to the same contracts selected with the other principals is essential to sustain certain mixed-strategy outcomes.

The remainder of the article is organized as follows. We conclude this section with a simple example that (gently) introduces the reader to the key ideas in the paper with as little formalism as possible. Section I describes the general contracting environment. Section II contains the main characterization results. Section III shows how our revelation mechanisms can be put to work in applications such as competition in nonlinear tariffs, menu auctions, and moral hazard settings. Section IV shows how the revelation mechanisms can be enriched to characterize equilibrium outcomes sustained by non-Markov strategies and/or mixed-strategy equilibria. Section V concludes. All proofs are in the Appendix.

Qualification.—While the approach here is similar in spirit to the one in Pavan and Calzolari (2009) for sequential common agency, there are important differences due to the simultaneity of contracting. First, the notion of Markov strategies considered here takes into account the fact that the agent, when choosing the messages to send to each principal, has not yet committed himself to any decision with any of the other principals. Second, in contrast to sequential games, the agent can condition his behavior on the entire profile of mechanisms offered by all principals. These differences explain why, despite certain similarities, the results here do not follow from the arguments in that paper.

A. *Simple Menu-Auction Example*

There are three players: a policymaker (the agent, A) and two lobbying domestic firms (principals P_1 and P_2). The policymaker must choose between a "protectionist" policy, $e = p$, and a "free-trade" policy, $e = f$ (e.g., opening the domestic market to foreign competition). To influence the policymaker's decision, the two firms can

[10] Note that this refinement is milder than the "conservative behavior" refinement of Attar et al. (2008).

make explicit commitments about their business strategy in the near future. We denote by $a_i \in \mathcal{A}_i = [0,1]$ the "aggressiveness" of firm i's business strategy, with $a_i = 1$ denoting the most aggressive strategy and $a_i = 0$ the least aggressive one. The aggressiveness of a firm's strategy should be interpreted as a proxy for a combination of its pricing policy, its investment strategy, the number of jobs the firm promises to secure, and similar factors.

The policymaker's payoff is a weighted average of domestic consumer surplus and domestic firms' profits. We assume that under a protectionist policy, welfare is maximal when the two domestic firms engage in fierce competition (i.e., when they both choose the most aggressive strategy). We also assume that the opposite is true under a free-trade policy. This could reflect the fact that, under a free-trade policy, large consumer surplus is already guaranteed by foreign supply, in which case the policymaker may value cooperation between the two firms.

We further assume that, absent any explicit contract with the government, the two firms cannot refrain from behaving aggressively. To make it simple, we assume that under a protectionist policy, P_1 has a dominant strategy in choosing $a_1 = 1$, in which case P_2 has an iteratively dominant strategy in also choosing $a_2 = 1$. Likewise, under a free-trade policy, P_2 has a dominant strategy in choosing $a_2 = 1$, in which case P_1 has an iteratively dominant strategy in also choosing $a_1 = 1$. By behaving aggressively, the two firms reduce their joint profits with respect to what they could obtain by "colluding," i.e., by setting $a_1 = a_2 = 0$.

Formally, the aforementioned properties can be captured by the following payoff structure:

$$u_1(e,a) = \begin{cases} a_1(1 - a_2/2) - a_2 & \text{if } e = p \\ a_1(a_2 - 1/2) - a_2 - 1 & \text{if } e = f \end{cases}$$

$$u_2(e,a) = \begin{cases} a_2(a_1 - 1/2) - a_1 & \text{if } e = p \\ a_2(1 - a_1/2) - a_1 - 1 & \text{if } e = f \end{cases}$$

$$v(e,a) = \begin{cases} 1 + a_2(2a_1 - 1) & \text{if } e = p \\ 10/3 + a_1(a_2 - 2) - a_2/2 & \text{if } e = f \end{cases},$$

where u_i denotes P_i's payoff, $i = 1,2$; v denotes the policymaker's payoff; and $a = (a_1, a_2)$.

What distinguishes this setting from most lobbying games considered in the literature is that payoffs are not restricted to being quasi-linear. As a consequence, the two lobbying firms respond to the choice of a policy e with an entire business plan as opposed to simply paying the policymaker a transfer t_i (e.g., a campaign contribution). Apart from this distinction, this is a canonical "menu-auction" setting à la Douglas B. Bernheim and Michael D. Whinston (1985, 1986a): the agent's action e is verifiable, preferences are common knowledge, and each principal can credibly commit to a contract $\delta_i : E \rightarrow \mathcal{A}_i$ that specifies a reaction (i.e., a business plan) for each possible policy $e \in E = \{p,f\}$.

In virtually all menu auction papers, it is customary to assume that the principals simply make take-it-or-leave-it offers to the agent. That is, they offer a single contract δ_i. Note that in games with complete information, a take-it-or-leave-it offer coincides with a standard direct revelation mechanism. It is easy to verify that, in the lobbying game in which the two firms are restricted to making take-it-or-leave-it offers, the only two pure-strategy equilibrium outcomes are: $e^* = p$ and $a_i^* = 1$, $i = 1, 2$, which yields each firm a payoff of $-1/2$ and the policymaker a payoff of 2; and $e^* = f$ and $a_i^* = 1$, $i = 1, 2$, which yields each firm a payoff of $-3/2$ and the policymaker a payoff of $11/6$. The proof is in the Appendix.

In an influential paper, Peters (2003) has shown that when a certain *no-externalities* condition holds, restricting the principals to making take-it-or-leave-it offers is inconsequential. Any outcome that can be sustained by allowing the principals to offer more complex mechanisms can also be sustained by restricting them to making take-it-or-leave-it offers. The no-externalities condition is often satisfied in quasi-linear environments (e.g., in Bernheim and Whinston's seminal 1986a menu-auction paper). However, it typically fails when a principal's action is the selection of an entire plan of action, such as a business strategy, as in the current example, or the selection of an incentive scheme, as in a moral hazard setting. In this case, restricting the principals to competing in take-it-or-leave-it offers (or equivalently, in standard direct revelation mechanisms) may preclude the possibility of characterizing interesting outcomes, as shown below.

A fully general approach would then require letting the principals compete by offering arbitrarily complex mechanisms. However, because ultimately a mechanism is just a procedure to select a contract, one can safely assume that each principal directly offers the agent a menu of contracts, and delegates to the agent the choice of the contract. In essence, this is what the Menu Theorem establishes. However, as anticipated above, this approach leaves open the question of which menus are offered in equilibrium, and how the different contracts in the menu are selected by the agent in response to the contracts selected with the other principals.

The solution offered by our approach consists in describing the agent's choice from a menu by means of a *revelation mechanism*. Contrary to the standard revelation mechanisms considered in the literature (where the agent simply reports his exogenous type), the revelation mechanisms we propose ask the agent also to report the (payoff-relevant) contracts selected with the other principals. Theorem 2 will show that any outcome of the menu game sustained by a pure-strategy equilibrium, in which the agent's strategy is *Markovian*, can also be sustained as a pure-strategy equilibrium outcome of the game in which the principals offer the revelation mechanisms described above.

In the lobbying game of this example, the policymaker's strategy is Markovian if, given any menu of contracts ϕ_i^M offered by firm i, and any contract $\delta_j : E \to \mathcal{A}_j$ by firm j, there exists a unique contract $\delta_i(\delta_j; \phi_i^M) : E \to \mathcal{A}_i$ such that the policymaker always selects the contract $\delta_i(\delta_j; \phi_i^M)$ from the menu ϕ_i^M when the contract he selects with firm j is δ_j, $j \neq i$. In other words, the choice from the menu ϕ_i^M depends only on the contract selected with the other firm, but not on payoff-irrelevant information, such as the other contracts included by firm j in her menu that the policymaker decided not to choose.

As anticipated in the introduction, while Markov strategies are appealing, they may fail to sustain certain outcomes. However, as Theorem 3 will show, this is never the case when the principals' preferences are sufficiently aligned (which is always the case when there are only two principals) and preferences are common knowledge, as in the example considered here. Moreover, as Proposition 4 will show, when effort is observable, as in menu-auctions, the revelation mechanisms can be further simplified by having the agent directly report to each principal the actions he is inducing the other principals to take in response to his choice of effort, as opposed to the contracts selected with the other principals. The idea is simple. For any given policy $e \in E$, the agent's preferences over the actions by principal i depend on the action by principal j. By implication, the agent's choice from any menu of contracts offered by P_i can be conveniently described through a mapping $\phi_i^r{:}E \times \mathcal{A}_j \rightarrow \mathcal{A}_i$ that specifies, for each *observable* policy $e \in E$, and for each *unobservable* action $a_j \in \mathcal{A}_j$ by principal j, an action $a_i \in \mathcal{A}_i$ that is as good for the agent as any other action a_i' that the agent can induce by reporting an action $a_j' \neq a_j$.[11] Furthermore, the agent's strategy can be restricted to being truthful in the sense that, in equilibrium, the agent correctly reports to each principal $i = 1,2$, the action a_j that will be taken by the other principal.

We conclude this example by showing how our revelation mechanisms can be used to sustain outcomes that can *not* be sustained with simple take-it-or-leave-it offers. To this aim, consider the following pair of revelation mechanisms.[12]

$$\phi_1^r(e,a_2) = \begin{cases} 1/2 & \text{if } e = p \ \forall a_2 \\ 1 & \text{if } e = f \ \forall a_2 \end{cases}, \quad \phi_2^r(e,a_1) = \begin{cases} 1 & \text{if } e = p \text{ and } a_1 > 1\backslash 2 \\ 0 & \text{if } e = p \text{ and } a_1 \leq 1\backslash 2 \\ 1 & \text{if } e = f \ \forall a_1. \end{cases}$$

Given these mechanisms, the policymaker optimally chooses a protectionist policy $e = p$. At the same time, the two firms sustain higher cooperation than under simple take-it-or-leave-it offers, thus obtaining higher total profits. Indeed, the equilibrium outcome is $e^* = p$, $a_1^* = 1/2$, $a_2^* = 0$ which yields P_1 a payoff of $1/2$, P_2 a payoff of $-1/2$, and the policymaker a payoff of 1. The key to sustaining this outcome is to have P_2 respond to the policy $e = p$ with a business strategy that depends on what P_1 does. Because P_2 cannot observe a_1 directly at the time she commits to her business plan, such a contingency must be achieved with the compliance of the policymaker.

Clearly, the same outcome can also be sustained in the menu game by having P_2 offer a menu that contains two contracts, one that responds to $e = p$ with $a_2 = 1$ and the other that responds to $e = p$ with $a_2 = 0$. The advantage of our mechanisms comes from the fact that they offer a convenient way of describing a principal's response to the other principals' actions that is compatible with the agent's

[11] When applied to games with no effort (i.e., to games where there is no action e that the agent has to take after communicating with the principals), these mechanisms reduce to mappings ϕ_i^r: $\mathcal{A}_j \rightarrow \mathcal{A}_i$ that specify a response by P_i (e.g., a price-quantity pair) to each possible action by P_j. Note that in these games, a contract for P_i simply coincides with an element of $\mathcal{A}_i$. In settings where the agent's preferences are not common knowledge, these mechanisms become mappings $\phi_i^r : \Theta \times \mathcal{A}_j \rightarrow \mathcal{A}_i$ according to which the agent is also asked to report his "type," i.e., his exogenous private information θ.

[12] Note that, because e is observable, these mechanisms only need to be incentive compatible with respect to a_j.

incentives. This simplification often facilitates the characterization of the equilibrium outcomes, as will be shown in the other examples in Section III.

I. 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 noncooperatively with the same agent, A. Each principal P_i, $i \in \mathcal{N} \equiv \{1, \dots, n\}$, must select a contract δ_i from a set of feasible contracts $\mathcal{D}_i$. A contract δ_i: $E \rightarrow \mathcal{A}_i$ specifies the action $a_i \in \mathcal{A}_i$ that P_i will take in response to the agent's action/effort $e \in E$. Both a_i and e may have different interpretations depending on the application of interest. When A is a policymaker lobbied by different interest groups, e typically represents a policy, and a_i may represent either a campaign contribution (as in Bernheim and Whinston 1986a) or a plan of action (as in the non-quasi-linear example of the previous section). When A is a buyer purchasing from multiple sellers, a_i may represent the price of seller i and e a vector of quantities/qualities purchased from the multiple sellers. Alternatively, as is typically assumed in models of competition in nonlinear tariffs, one can directly assume that $a_i = (t_i, q_i)$ is a price-quantity pair and then suppress e by letting E be a singleton (see, for example, the analysis in Section III).

Depending on the environment, the set of feasible contracts $\mathcal{D}_i$ may also be more or less restricted. For example, in certain trading environments, it can be appealing to assume that the price a_i of seller i cannot depend on the quantities/qualities of other sellers.[13] In a moral hazard setting, because e is not observable by the principals, each contract $\delta_i \in \mathcal{D}_i$ must respond with the same action $a_i \in \mathcal{A}_i$ to each e; in this case, a_i represents a state-contingent payment that rewards the agent as a function of some exogenous (and here unmodelled) performance measure that is correlated with the agent's effort. What is important to us is that the set of feasible contracts $\mathcal{D}_i$ is a primitive of the environment and not a choice of principal i.

Payoffs.—Principal i's payoff, $i = 1, \dots, n$, is described by the function $u_i(e, a, \theta)$, whereas the agent's payoff is described by the function $v(e, a, \theta)$. The vector $a \equiv (a_1, \dots, a_n) \in \mathcal{A} \equiv \times_{i=1}^{n} \mathcal{A}_i$ 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 that θ is drawn from the distribution F with support Θ. All players are expected-utility maximizers.

Mechanisms.—Principals compete in mechanisms. A mechanism for P_i consists of a (measurable) message space $\mathcal{M}_i$ along with a (measurable) mapping ϕ_i: $\mathcal{M}_i \rightarrow \mathcal{D}_i$. The interpretation is that when A sends the message $m_i \in \mathcal{M}_i$, P_i then responds by selecting the contract $\delta_i = \phi_i(m_i) \in \mathcal{D}_i$. Note that when there is no action that

[13] Such a departure is allowed in Calzolari and Devicolo (2009) and in Martimort and Stole (2005).

the agent must take after communicating with the principals (that is, when E is a singleton, as in the literature on competition in nonlinear schedules), δ_i reduces to a payoff-relevant action $a_i \in \mathcal{A}_i$, such as a price-quantity pair.

To save on notation, in the sequel, we will denote a mechanism simply by ϕ_i, thus dropping the specification of its message space $\mathcal{M}_i$ whenever this does not create any confusion. For any mechanism ϕ_i, we will then denote by $\text{Im}(\phi_i) \equiv \{\delta_i \in \mathcal{D}_i : \exists\, m_i \in \mathcal{M}_i \text{ s.t. } \phi_i(m_i) = \delta_i\}$ the range of ϕ_i, i.e., the set of contracts that the agent can select by sending different messages.

For any common agency game Γ, we will then denote by Φ_i the set of feasible mechanisms for P_i, by $\phi \equiv (\phi_1, \ldots, \phi_n) \in \Phi \equiv \times_{j=1}^n \Phi_j$ a profile of mechanisms for the n principals, and by $\phi_{-i} \equiv (\phi_1, \ldots, \phi_{i-1}, \phi_{i+1}, \ldots, \phi_n) \in \Phi_{-i} \equiv \times_{j \neq i} \Phi_j$ a profile of mechanisms for all $P_j, j \neq i$.[14] As is standard, we assume that principals can fully commit to their mechanisms and that each principal can neither communicate with the other principals,[15] nor make her contract contingent on the contracts by other principals.[16]

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 $m = (m_1, \ldots, m_n)$ are sent simultaneously.[17]
- At $t = 3$, A chooses an action $e \in E$.
- At $t = 4$, the principals' actions $a = \delta(e) \equiv (\delta_1(e), \ldots, \delta_n(\text{e}))$ are determined by the contracts $\delta = (\phi_1(m_1), \ldots, \phi_n(m_n))$, and payoffs are realized.

Strategies and Equilibria.—A (mixed) strategy for P_i is a distribution $\sigma_i \in \Delta(\Phi_i)$ over the set of feasible mechanisms. As for the agent, a (behavioral) 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 $\xi : \Theta \times \Phi \times \mathcal{M} \to \Delta(E)$ that specifies a distribution over effort for any (θ, ϕ, m).

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), any $e \in \text{Supp}[\xi(\theta, \phi, m)]$ maximizes $v(e, \delta(e), \theta)$, where $\delta = \phi(m)$; and, for every (θ, ϕ), any $m \in \text{Supp}[\mu(\theta, \phi)]$ maximizes $V(\phi(m), \theta) \equiv \max_{e \in E} v(e, \delta(e), \theta)$ with $\delta = \phi(m)$.

[14] We also define $\delta \equiv (\delta_1, \ldots, \delta_n) \in \mathcal{D} \equiv \times_{j=1}^n \mathcal{D}_j$, $m \equiv (m_1, \ldots, m_n) \in \mathcal{M} \equiv \times_{j=1}^n \mathcal{M}_j$, $\delta_{-i} \in \mathcal{D}_{-i}$, $m_{-i} \in \mathcal{M}_{-i}$ in the same way.

[15] A notable exception is Peters and Cristián Troncoso-Valverde (2009).

[16] As in Bernheim and Whinston (1986a), 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 $\delta_i : E \to \mathcal{A}_i$ contingent on the other principals' contracts δ_{-i}, nor her mechanism ϕ_i contingent on the other principals' mechanisms ϕ_{-i}. A recent paper that allows for these types of contingencies is Peters and Balazs Szentes (2008).

[17] 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 no trade at a null price.

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

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

where

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

A perfect Bayesian equilibrium for Γ is then 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 $\mathcal{E}(\Gamma)$. For any equilibrium $\sigma^* \in \mathcal{E}(\Gamma)$, we will then denote by $\pi_{\sigma^*} : \Theta \rightarrow \Delta(\mathcal{A} \times E)$ the associated *social choice function* (SCF).[18]

Menus.—A *menu* is a mechanism $\phi_i^M : \mathcal{M}_i^M \rightarrow \mathcal{D}_i$ whose message space $\mathcal{M}_i^M \subseteq \mathcal{D}_i$ is a subset of all possible contracts 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 of feasible contracts for P_i, and by Γ^M the "menu game" in which the set of feasible mechanisms for each P_i is Φ_i^M. We will then say that the game Γ is an *enlargement* of Γ^M ($\Gamma \succcurlyeq \Gamma^M$) if for all $i \in \mathcal{N}$, there exists an embedding $\alpha_i : \Phi_i^M \rightarrow \Phi_i$;[19] and for any $\phi_i \in \Phi_i$, $\text{Im}(\phi_i)$ is compact. A simple example of an enlargement of Γ^M is a game in which each Φ_i is a superset of Φ_i^M. More generally, an enlargement is a game in which each 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 (Peters 2001; Martimort and Stole 2002).

THEOREM 1 (Menus): *Let* Γ *be any enlargement of* Γ^M. *A social choice function* π *can be sustained by an equilibrium of* Γ *if and only if it can be sustained by an equilibrium of* Γ^M.

[18] In the jargon of the mechanism design/implementation literature, a social choice function $\pi : \Theta \rightarrow \Delta(\mathcal{A} \times E)$ is simply an outcome function, which specifies, for each state of nature θ, a joint distribution over payoff-relevant decisions (a, e).

[19] For our purposes, an embedding $\alpha_i : \Phi_i^M \rightarrow \Phi_i$ can be thought of as an injective mapping such that, for any pair of mechanisms ϕ_i^M, ϕ_i with $\phi_i = \alpha_i(\phi_i^M)$, $\text{Im}(\phi_i) = \text{Im}(\phi_i^M)$.

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 restricting 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 $\tilde{\Phi}_i^M \equiv \{\phi_i^M : \text{Im}(\phi_i^M) = \text{Im}(\phi_i)$ for some $\phi_i \in \Phi_i\}$.

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

Remark.—To ease the exposition, throughout the entire main text, we restrict our attention to settings where principals offer simple menus that contain only *deterministic* contracts, i.e., mapping $\delta_i : E \rightarrow \mathcal{A}_i$. All our results apply verbatim to more general settings where the principals can offer the agent menus of *lotteries over stochastic contracts*; it suffices to reinterpret each δ_i as a lottery over a set of stochastic contracts $Y_i = \{y_i : E \rightarrow \Delta(\mathcal{A}_i)\}$, where each y_i responds to each effort choice by the agent with a distribution over $\mathcal{A}_i$. Note that, in general, even if one restricts one's attention to pure-strategy profiles (i.e., to strategy profiles in which the principals do not mix over the menus they offer to the agent and where the agent does not mix over the messages he sends to the principals), allowing the principals to offer lotteries over stochastic contracts may be essential to sustain certain outcomes. The reason is that such lotteries create uncertainty about the principals' responses to the agent's effort, which permits one to sustain a wider range of equilibrium effort choices (see Peters 2001, for a few examples). All proofs in the Appendix consider these more general settings.

II. Simple Revelation Mechanisms

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

DEFINITION 1:

(i) Given *the common agency game* Γ, *an equilibrium strategy profile* $\sigma \in \mathcal{E}(\Gamma)$ *is a* pure-strategy equilibrium *if*

- no *principal randomizes over her mechanisms; and*
- 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* Markovian *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* $\delta_i(\theta, \delta_{-i} ; \phi_i) \in \text{Im}(\phi_i)$, *such that A always selects* $\delta_i(\theta, \delta_{-i} ; \phi_i)$ *with* P_i *when the latter offers the mechanism* ϕ_i, *the agent's type is* θ, *and the contracts A selects 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 randomizes over the messages he sends to the principals. Note, however, that the agent may randomize over his choice of effort.

The agent's strategy σ_A in Γ is Markovian if and only if the contracts the agent selects in each mechanism depend only on his type and the contracts which he selects with the other principals, but not on the particular profile of mechanisms (or menus) offered by those principals. 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* (*incentive-compatible*) *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* $(m_i^r, m_{-i}^r) \in \mathrm{Supp}[\mu(\theta, \phi_i^r, \phi_{-i}^r)]$,

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

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

In a revelation mechanism, the agent is thus asked to report his type θ along with the contracts δ_{-i} he is selecting with the other principals. Given a profile of mechanisms ϕ^r, the agent's strategy is then said to be truthful in ϕ_i^r if the message $m_i^r = (\theta, \delta_{-i})$, which the agent sends to P_i, coincides with his true type θ together with the true contracts $\delta_{-i} = (\phi_j(m_j))_{j \neq i}$ that the agent selects with all principals other than i by sending the messages $m_{-i} \equiv (m_j)_{j \neq i}$. Finally, an equilibrium strategy profile is said to be a truthful equilibrium if, whenever no more than a single principal deviates from equilibrium play, the agent reports truthfully to any of the nondeviating principals.

The following is our first characterization result.

THEOREM 2:

(i) Suppose *that the social choice function* π *can be sustained by a pure-strategy equilibrium of* Γ^M *in which the agent's strategy is Markovian. Then* π *can also be sustained by a truthful pure-strategy equilibrium of* Γ^r.

(ii) Furthermore, *any social choice function* π *that can be sustained by an equilibrium of* Γ^r *can also be sustained by an equilibrium of* Γ^M.

Consider first part (i). When the agent's choice from each menu depends only on his type θ and the contracts δ_{-i} that he is selecting with the other principals, one can easily see that, in equilibrium, each principal can be restricted to offering a menu ϕ_i^{M*} such that

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

It is then also easy to see that, starting from such an equilibrium, one can construct a truthful equilibrium for the revelation game that sustains the same outcomes.

Next, consider part (ii). Despite the fact that Γ^r is *not* an enlargement of Γ^M, the result follows essentially from the same arguments that establish the Menu Theorem. The equilibrium that sustains the SCF π in Γ^M is constructed from σ^{r*} by having each principal offering the menu ϕ_i^{M*} that corresponds to the range of the equilibrium mechanism ϕ_i^{r*} of Γ^r. When in Γ^M all principals offer the equilibrium menus, the agent then implements the same outcomes he would have implemented in Γ^r. When, instead, one principal (let us say P_i) deviates and offers a menu $\phi_i^M \notin \mathrm{Supp}[\sigma_i^{M*}]$, the agent implements the same outcomes he would have implemented in Γ^r had P_i offered a direct mechanism ϕ_i^r, such that

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

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 it seems reasonable to assume that the agent's strategy is Markovian, it is also important to understand whether there exist environments in which such an assumption is not a restriction. To address this question, we first need to introduce some notation. For any $k \in \mathcal{N}$ and any (δ, θ), let

$$E^*(\delta, \theta) \equiv \underset{e \in E}{\arg\max}\, v(e, \delta(e), \theta)$$

denote the set of effort choices that are optimal for type θ given the contracts δ. Then let

$$\underline{U}_k(\delta, \theta) \equiv \min_{e \in E^*(\delta, \theta)} u_k(e, \delta(e), \theta)$$

denote the lowest payoff that the agent can inflict to principal k by following a strategy that is consistent with the agent's own-payoff-maximizing behavior.

CONDITION 1 (Uniform Punishment): *We say that the "Uniform Punishment" condition holds if, for any* $i \in \mathcal{N}$, *compact set of contracts* $B \subseteq \mathcal{D}_i$, $\delta_{-i} \in \mathcal{D}_{-i}$, *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* $j \neq i$, *all* $\hat{\delta}_i \in \arg\max_{\delta_i \in B} V(\delta_i, \delta_{-i}, \theta)$,

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

This condition says that the principals' preferences are sufficiently aligned in the following sense. Given any menu of contracts B offered by P_i and any (θ, δ_{-i}), there exists a contract $\delta_i' \in B$ that is optimal for type θ given δ_{-i}, and which uniformly minimizes the payoff of any principal other than i. By this we mean the following. The payoff of any principal $P_j, j \neq i$, under δ_i' is (weakly) lower than under any other contract $\delta_i \in$ B that is optimal for the agent given (θ, δ_{-i}).

We then have the following result:

THEOREM 3: *Suppose that at least one of the following conditions holds:*

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

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

Then any social choice function π *that can be sustained by a pure-strategy equilibrium of* Γ^M *can also be sustained by a pure-strategy equilibrium in which the agent's strategy is Markovian.*

Condition (i) 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 contracts $B \subseteq \mathcal{D}_i$, there is a single contract in B that maximizes the agent's payoff. Clearly, in this case the agent's strategy is necessarily Markovian.

Condition (ii) 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 contract with principal i. Note that this condition would be satisfied, for example, when the agent is a manufacturer and the principals are retailers competing à la Cournot in a downstream market. In this case,

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

where q_i denotes the quantity sold to P_i, t_i denotes the total payment made by P_i to the manufacturer, and $f: \mathbb{R}_+ \to \mathbb{R}$ denotes the inverse demand function in the downstream market. In this environment, $|\Theta| = |E| = 1$. A contract δ_i is thus a simple price-quantity pair $(t_i, q_i) \in \mathbb{R} \times \mathbb{R}_+$. One can then immediately see that, given any menu $B \subseteq \mathbb{R} \times \mathbb{R}_+$ (i.e., any array of price-quantity pairs or, equivalently, any tariff) 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 (for any $j \neq i$) among those that are optimal for the agent given (t_{-i}, q_{-i}) is

Table 1

$\theta = \underline{\theta}$			$\theta = \overline{\theta}$		
$a_1 \backslash a_2$	l	r	$a_1 \backslash a_2$	l	r
t	2 1 1	2 0 0	t	2 2 2	–2 0 2
b	1 0 1	1 2 2	b	1 0 1	–2 1 1

the one that entails the highest quantity q_i. The Uniform Punishment condition thus clearly holds in this environment.

The reason why the result in Theorem 3 requires information to be complete in addition to having enough alignment in the principals' payoffs, can be illustrated through the following example where $n = 2$, in which case the Uniform Punishment condition trivially holds. The sets of actions are $\mathcal{A}_1 = \{t, b\}$ and $\mathcal{A}_2 = \{l, r\}$. There is no effort in this example, and hence a contract simply coincides with the choice of an element of $\mathcal{A}_i$. 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 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 by a (pure-strategy) equilibrium of the menu game in which the agent's strategy is non-Markovian. 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 in each state she is getting her maximal feasible payoff. If P_1 deviates and offers $\{t\}$, then A selects (t, l) if $\theta = \underline{\theta}$ and (t, r) if $\theta = \overline{\theta}$. Note that, 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 being Markovian, 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 being Markovian then means that when P_2 offers the equilibrium menu, A necessarily chooses r if $(\theta, a_1) = (\underline{\theta}, b)$, and l if $(\theta, a_1) = (\overline{\theta}, t)$. Furthermore, because given $(\underline{\theta}, t)$, A strictly prefers l to r, A necessarily chooses l when $(\theta, a_1) = (\underline{\theta}, t)$. Given this behavior, if P_1 deviates and offers the menu $\phi_1^M = \{t\}$, she then induces A to select $a_2 = l$ with P_2 irrespective of θ, which gives P_1 a payoff $U_1 = 2 > U_1^*$.

The reason why, when information is incomplete, restricting the agent's strategy to be Markovian may preclude the possibility of sustaining certain social choice functions is the following. Markov strategies do not permit the same type of the agent (let us say θ') to punish a deviation by a principal (let us say $P_j, j \neq i$) by choosing with all principals other than i the equilibrium contracts $\delta_{-i}^*(\theta')$, and then choosing with P_i a contract $\delta_i \neq \delta_i^*(\theta')$. As the example above illustrates, it may be essential in order to punish certain deviations to allow a type to change his behavior with a principal, even if the contracts he selects with all other principals coincide with the equilibrium ones. However, because this is the only reason that one needs information to be complete for the result in Theorem 3, it turns out that the assumption of complete information can be dispensed with if one imposes the following refinement on the agent's behavior:

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 contracts selected 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* $m \in Supp[\mu(\theta, \phi_i^*, \phi_{-i})]$,

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

That is, the agent's strategy satisfies the Conformity to Equilibrium condition if each type of the agent θ selects the equilibrium contract $\delta_i^*(\theta)$ with each principal P_i when the latter offers the equilibrium mechanism ϕ_i^*, and the agent selects the equilibrium contracts $\delta^*_{-i}(\theta)$ with the other principals. Consider the same example described above and assume that the principals compete in menus, i.e., $\Gamma = \Gamma^M$. Take the equilibrium in which P_1 offers the degenerate menu $\{t\}$, and P_2 offers the menu $\{l, r\}$. Given the equilibrium menus, both types select $a_2 = l$ with P_2. One can immediately see that this outcome can be sustained by a strategy for the agent that satisfies the "Conformity to Equilibrium" condition. It suffices that, whenever P_2 offers the equilibrium menu $\{l, r\}$, then each type θ selects the contract $a_2 = l$ with P_2, when selecting the equilibrium contract $a_1 = t$ with P_1. Note that this refinement does not require that the agent does not change his behavior with a nondeviating principal; in particular, should P_1 deviate and offer the menu $\{t, b\}$, then type $\underline{\theta}$ would, of course, select $a_1 = b$ with P_1, and then also change the contract with P_2 to $a_2 = r$. What this refinement requires is simply that each type of the agent continues to select the equilibrium contract with a nondeviating principal *conditional on choosing the equilibrium contracts with the remaining principals*. In many applications, this property seems to us 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 social choice function* π *can be sustained by 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, irrespective of whether information is complete or incomplete,* π *can also be sustained by a pure-strategy equilibrium* $\tilde{\sigma}^{M*} \in \mathcal{E}(\Gamma^M)$ *in which the agent's strategy* $\tilde{\sigma}_A^{M*}$ *is Markovian.*

At this point, it is useful to contrast our results with those in Peters (2003, 2007) and Attar et al. (2008). 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 offering standard direct revelation mechanisms $\phi_i : \Theta \rightarrow \mathcal{D}_i$.[20] The no-externality condition

[20] A standard direct revelation mechanism reduces to a take-it-or-leave-it-offer, i.e., to a degenerate menu consisting of a single contract $\delta_i : E \rightarrow \mathcal{A}_i$, when the agent does not possess any exogenous private information,

requires that (i) each principal's payoff be independent of the other principals' actions a_{-i}, and (ii) conditional on choosing effort in a certain equivalence class $\hat{E}$,[21] the agent's preferences over any set of actions $B \subseteq \mathcal{A}_i$ by principal i be independent of the particular effort the agent chooses in $\hat{E}$, of his type θ, and of the other principals' actions a_{-i}. Attar et al. (2008) show that in environments in which only deterministic contracts are feasible, all action spaces are finite, and the agent's preferences are "separable" and "generic," condition (i) in Peters (2003) can be dispensed with: any equilibrium outcome of the menu game (including those sustained by mixed strategies) can also be sustained as an equilibrium outcome 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 actions of any of the principals be independent of the effort choice and of the actions of the other principals. Genericity requires that the agent never be indifferent between any pair of effort choices and/or any pair of contracts by any of the principals.[22] Taken together, these restrictions guarantee that the messages that each type of the agent sends to any of the principals do not depend on the messages he sends to the other principals. It is then clear that, in these settings, restricting attention to standard direct revelation mechanisms never precludes the possibility of sustaining all outcomes.

Compared to these results, the result in Theorem 2 does not require any restriction on the players' preferences. On the other hand, it requires restricting attention to equilibria in which the agent's strategy is Markovian. This restriction is, however, inconsequential either when the agent's preferences are single-peaked or when information is complete and the principals' preferences are sufficiently aligned in the sense of the Uniform Punishment condition. Our results complement those in Peters (2003, 2007) and Attar et al. (2008) in the sense that they are particularly useful precisely in environments in which one cannot restrict attention either to simple take-it-or-leave-it offers or to standard direct revelation mechanisms.

For example, consider a pure adverse selection setting, as in the baseline model of Attar et al. (2008).[23] Then condition (i) 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 Attar et al. (2008) guarantees that all equilibrium outcomes can be sustained by restricting the principals to offering standard direct revelation mechanisms. Assuming that preferences are separable, however, can be too restrictive. For example, it rules out the possibility that a buyer's preferences

i.e., when $|\Theta| = 1$.

[21] In the language of Peters (2003, 2007), 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., $\delta_i(e) = \delta_i(e')$ for any $e, e' \in \hat{E}$.

[22] 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 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 (2003, 2007) requires 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 $|\Theta| = 1$).

[23] A pure adverse selection setting is one with no effort, i.e., where $|E| = 1$.

for the quality/quantity of a seller's product might depend on the quality/quantity of the product purchased from another seller. In cases like these, all equilibrium outcomes can still be characterized by restricting the principals to offering direct revelation mechanisms; however, the latter must be enriched to allow the agent to report the contracts (i.e., the terms of trade) that he has selected with the other principals, in addition to his exogenous private information.

Also note that when action spaces are continuous, as is typically assumed in most applications, Attar et al. (2008) need to impose a restriction on the agent's behavior. This restriction, which they call "conservative behavior," consists in requiring that, after a deviation by P_k, each type θ of the agent continues to choose the equilibrium contracts $\delta^*_{-k}(\theta)$ with the nondeviating 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 prove useful in applications in which imposing the "conservative behavior" property seems too restrictive.

III. Using Revelation Mechanisms in Applications

Equipped with the results established in the preceding section, we now consider three canonical applications of the common agency model: competition in nonlinear tariffs with asymmetric information, menu auctions, and a moral hazard setting. The purpose of this section is to show how the revelation mechanisms introduced in this paper can facilitate the analysis of these games by helping one identify the necessary and sufficient conditions for the equilibrium outcomes.

A. *Competition in Nonlinear 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; a contract δ_i for principal i thus consists of a price-quantity pair $(t_i, q_i) \in \mathcal{A}_i \equiv \mathbb{R} \times \mathcal{Q}$, where $\mathcal{Q} = [0, \overline{Q}]$ denotes the set of feasible quantities.[24]

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, and where θ denotes the buyer's type. The two sellers share a common prior that θ is drawn from an absolutely continuous cumulative distributions function F with support $\Theta = [\underline{\theta}, \overline{\theta}]$, $\underline{\theta} > 0$, and log-concave density f strictly positive for any $\theta \in \Theta$. The sellers' payoffs are given by $u_i(a, \theta) = t_i - C(q_i)$, with $C(q) = q^2/2$, $i = 1, 2$.

We assume that the buyer's choice to participate in seller i's mechanism has no effect on his possibility to participate in seller j's mechanism. In other words, the buyer can choose to participate in both mechanisms, only in one of the two, or in none (In the literature, this situation is referred to as "nonintrinsic" common

[24] 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 $\delta_i : \mathcal{Q} \to \mathbb{R}$ that specifies a price for each possible quantity $q \in \mathcal{Q}$. Given a pair of tariffs $\delta = (\delta_1, \delta_2)$, 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.

agency.) In the case where A decides not to participate in seller i's mechanism, the default contract $(0,0)$ with no trade and zero transfer is implemented.

Following the pertinent literature, we assume that only deterministic mechanisms $\phi_i : \mathcal{M}_i \rightarrow \mathcal{A}_i$ are feasible. Because the agent's payoff is strictly decreasing in t_i, any such mechanism is strategically equivalent to a (possibly nonlinear) *tariff* T_i, such that, for any q_i, $T(q_i) = \min\{t_i : (t_i, q_i) \in \text{Im}(\phi_i)\}$ if $\{t_i : (t_i, q_i) \in \text{Im}(\phi_i)\} \neq \varnothing$, and $T(q_i) = \infty$ otherwise.[25]

The question of interest is which tariffs will be offered in equilibrium and, even more importantly, what are the corresponding quantity schedules $q_i^* : \Theta \rightarrow \mathcal{Q}$ that they support. Following the discussion in the previous sections, we focus on pure-strategy equilibria in which the buyer's behavior is Markovian.

The purpose of this section is to show how our results can help address these questions. To do this, we first show how our revelation mechanisms can help identify necessary and sufficient conditions for the sustainability of schedules $q_i^* : \Theta \rightarrow \mathcal{Q}$, $i = 1,2$, as equilibrium outcomes. Next, we show how these conditions can be used to prove that there is no equilibrium that sustains the schedules $q^c : \Theta \rightarrow \mathcal{Q}$ that maximize the sellers' joint payoffs. These schedules are referred to in the literature as "collusive schedules." Last, we identify sufficient conditions for the sustainability of differentiable schedules.

Necessary and Sufficient Conditions for Equilibrium Schedules.—By Theorem 2, the quantity schedules $q_i^*(\cdot)$, $i = 1,2$, can be sustained by a pure-strategy equilibrium of Γ^M in which the agent's strategy is Markovian *if and only if* they can be sustained by 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 quantity q_i when he purchases the equilibrium quantity $q_j^*(\theta)$ from seller j, $j \neq i$. In what follows, we restrict our attention to equilibrium schedules $(q_i^*(\cdot))_{i=1,2}$ for which the corresponding marginal valuation functions $\mathrm{m}_i(\cdot)$ are strictly increasing, $i = 1,2$.[26] These schedules can be characterized by restricting attention to revelation mechanisms with the property that $\phi_i^r(\theta, q_j, t_j) = \phi_i^r(\theta', q_j', t_j')$ whenever $\theta + \lambda q_j = \theta' + \lambda q_j'$.[27] With an abuse of notation, hereafter, we denote such mechanisms by $\phi_i^r = (\tilde{q}_i(\theta_i), \tilde{t}_i(\theta_i))_{\theta_i \in \Theta_i}$, where

$$\Theta_i \equiv \{\theta_i \in \mathbb{R} : \theta_i = \theta + \lambda q_j, \theta \in \Theta, q_j \in \mathcal{Q}\}$$

[25] Clearly, any such tariff is also equivalent to a menu of price-quantity pairs (see also Peters, 2001, 2003).

[26] Note that this is necessarily the case when $(q_i^*(\cdot))_{i=1,2}$ are the collusive schedule (described below). More generally, the restriction to schedules for which the corresponding marginal valuation functions $m_i(\cdot)$ are strictly increasing simplifies the analysis by guaranteeing that these functions are invertible.

[27] Clearly, restricting attention to such mechanisms would not be appropriate if either $m_i(\cdot)$ were not invertible; or the principals' payoffs also depended on θ and (q_j, t_j). In the former case, to sustain the equilibrium schedules, a mechanism may need to respond to the same m_i with a contract that also depends on θ. In the latter case, a mechanism may need to punish a deviation by the other principal with a contract that depends not only on m_i, but also on (θ, q_i, t_i).

denotes the set of marginal valuations that the agent may possibly have for P_i's quantity. Note that these mechanisms specify price-quantity pairs also for marginal valuations θ_i that may have zero measure on the equilibrium path. This is because sellers may need to include in their menus price-quantity pairs that are selected only off equilibrium to punish deviations by other sellers.[28] In the literature, these price-quantity pairs are typically obtained by extending the principals' tariffs outside the equilibrium range (see, e.g., Martimort 1992). However, identifying the appropriate extensions can 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, because the set of marginal valuations Θ_i is a compact interval, and the function $\tilde{v}(\theta_i, q) \equiv \theta_i q$ is equi-Lipschitz continuous and differentiable in θ_i, and satisfies the increasing-difference property, the mechanism $\phi_i^r = (\tilde{q}_i(\cdot), \tilde{t}_i(\cdot))$ is incentive-compatible if and only if the function $\tilde{q}_i(\cdot)$ is nondecreasing and the function $\tilde{t}_i(\cdot)$ satisfies

$$(1) \qquad \tilde{t}_i(\theta_i) = \theta_i \tilde{q}_i(\theta_i) - \int_{\min \Theta_i}^{\theta_i} \tilde{q}_i(s)\, ds - K_i \ \ \forall \theta_i \in \Theta_i,$$

where K_i is a constant.[29] 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 $\theta_i \in \Theta_i$, such that an agent with marginal valuation θ_i strictly prefers the null contract $(0,0)$ to the contract $(\tilde{q}_i(\theta_i), \tilde{t}_i(\theta_i))$, there exists another pair of mechanisms $(\phi_i^{r\prime})_{i=1,2}$, such that: for any $\theta_i \in \Theta_i$, the agent weakly prefers the contract $(\tilde{q}_i{}'(\theta_i), \tilde{t}_i{}'(\theta_i))$ to the null contract $(0,0)$, $i = 1,2$; and $(\phi_i^{r\prime})_{i=1,2}$ sustains the same outcomes as $(\phi_i^r)_{i=1,2}$.[30] From (1), we can therefore restrict K_i to be positive.

Now, given any pair of incentive-compatible mechanisms $(\phi_i^r)_{i=1,2}$, let $\overline{U}_i$ denote the maximal payoff that each P_i can obtain given the opponent's mechanism ϕ_j^r, $j \neq i$, while satisfying the agent's rationality. This can be computed by solving the following program:

$$\tilde{\mathcal{P}}: \begin{cases} \max\limits_{q_i(\cdot), t_i(\cdot)} \displaystyle\int_{\underline{\theta}}^{\bar{\theta}} \left[t_i(\theta) - \frac{q_i(\theta)^2}{2} \right] dF(\theta) \\ \text{s.t.} \\ \theta q_i(\theta) + v_i^*(\theta, q_i(\theta)) - t_i(\theta) \geq \theta q_i(\hat{\theta}) + v_i^*(\theta, q_i(\hat{\theta})) - t_i(\hat{\theta}) \, \forall (\theta, \hat{\theta}) \quad \text{(IC)} \\ \theta q_i(\theta) + v_i^*(\theta, q_i(\theta)) - t_i(\theta) \geq v_i^*(\theta, 0) \ \forall \theta \quad \text{(IR)}, \end{cases}$$

where, for any $(\theta, q) \in \Theta \times \mathcal{Q}$,

$$(2) \quad v_i^*(\theta, q) \equiv (\theta + \lambda q)\tilde{q}_j(\theta + \lambda q) - \tilde{t}_j(\theta + \lambda q) = \int_{\min \Theta_j}^{\theta + \lambda q} \tilde{q}_j(s)\, ds + K_j, j \neq i$$

[28] These allocations are sometimes referred to as "latent contracts;" see, e.g., Gwenael Piaser 2007.
[29] This result is standard in mechanism design; see, e.g., Paul R. Milgrom and Ilya R. Segal (2002).
[30] The result follows from replication arguments similar to those that establish Theorem 2.

denotes the maximal payoff that type θ obtains with principal P_j, $j \neq i$, when he purchases a quantity q from principal P_i. The payoff $\overline{U}_i$ is thus computed using the standard revelation principle, but taking into account the fact that, given the incentive-compatible mechanism ϕ_j^r offered by P_j, the *total* value that each type θ assigns to the quantity q purchased from P_i is $\theta q + v_i^*(\theta, q)$. Note that, in general, one should not presume that P_i can guarantee herself the payoff $\overline{U}_i$, even if $\overline{U}_i$ can be obtained without violating the agent's rationality. In fact, when the agent is indifferent, he could refuse to follow P_i's recommendations, thus giving P_i a payoff smaller than $\overline{U}_i$. The reason that, in this particular environment, P_i can guarantee herself the maximal payoff $\overline{U}_i$ is twofold: she is not personally interested in the contracts the agent signs with P_j; and the agent's payoff for any contract (q_i, t_i) is quasi-linear and has the increasing-difference property with respect to (θ, q_i). As we show in the Appendix, taken together these properties imply 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 $\overline{U}_i$.

Next, let

$$(3) \qquad V^*(\theta) \equiv \theta[q_1^*(\theta) + q_2^*(\theta)] + \lambda q_1^*(\theta) q_2^*(\theta) - \tilde{t}_1(m_1(\theta_1)) - \tilde{t}_2(m_2(\theta_2))$$

denote the equilibrium payoff that each type θ obtains by truthfully reporting to each principal the equilibrium marginal valuation $m_i(\theta) = \theta + \lambda q_j^*(\theta)$. The necessary and sufficient conditions for the sustainability of the pair of schedules $(q_i^*(\cdot))_{i=1}^2$ by an equilibrium can then be stated as follows:

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

(i) *for any marginal valuation $\theta_i \in [m_i(\underline{\theta}), m_i(\overline{\theta})]$, $\tilde{q}_i(\theta_i) = q_i^*(m_i^{-1}(\theta_i))$, $i = 1, 2$;*[31]

(ii) *for any $\theta \in \Theta$ and any pair $(\theta_1, \theta_2) \in \Theta_1 \times \Theta_2$,*

$$V^*(\theta) = \sup_{(\theta_1, \theta_2) \in \Theta_1 \times \Theta_2} \{\theta[\tilde{q}_1(\theta_1) + \tilde{q}_2(\theta_2)] + \lambda \tilde{q}_1(\theta_1)\tilde{q}_2(\theta_2) - \tilde{t}_1(\theta_1) - \tilde{t}_2(\theta_2)\},$$

where the functions $\tilde{t}_i(\cdot)$ are the ones defined in (1) *with $K_i = \tilde{K}_i$, $i = 1, 2$, and where the function $V^*(\cdot)$ is the one defined in* (3)*; and*

(iii) *each principal's equilibrium payoff satisfies*

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

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

where $\overline{U}_i$ is the value of the program $\tilde{\mathcal{P}}$ defined above.

Condition (i) guarantees that, on the equilibrium path, the mechanism ϕ_i^{r*} assigns to each θ the equilibrium quantity $q_i^*(\theta)$. Condition (ii) guarantees that each type θ finds it optimal to truthfully report to each principal his equilibrium marginal valuation $m_i(\theta)$. The fact that each type θ also finds it optimal to participate follows from the fact that $\tilde{K}_i \geq 0$. Finally, condition (iii) 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, condition (iii) directly guarantees that the equilibrium payoff for each principal coincides with the maximal payoff that the principal can obtain, given the opponent's mechanism, and without violating the agent's rationality. As explained above, because P_i can always guarantee herself the payoff $\overline{U}_i$, condition (iii) is not only sufficient, but also necessary.

When $\lambda > 0$, and the function $v_i^*(\theta, q)$ in (2), is differentiable in θ (which is the case, for example, when the schedule $\tilde{q}_j(\cdot)$ is continuous), the program $\tilde{\mathcal{P}}$ 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, 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 the constraints (IC) and (IR) in program $\tilde{\mathcal{P}}$ 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 value of program $\tilde{\mathcal{P}}$ then coincides with the value of the following program:

$$\tilde{\mathcal{P}}^{new} : \begin{cases} \max\limits_{q_i(\cdot), K_i} \displaystyle\int_{\underline{\theta}}^{\overline{\theta}} h_i(q_i(\theta); \theta)\, dF(\theta) - K_i \\ \text{s.t. } K_i \geq 0 \text{ and } q_i(\cdot) \text{ is nondecreasing'} \end{cases}$$

where

$$h_i(q; \theta) \equiv \theta q + [v_i^*(\theta, q) - v_i^*(\theta, 0)] - \frac{q^2}{2} - \frac{1 - F(\theta)}{f(\theta)}[q + \tilde{q}_j(\theta + \lambda q) - \tilde{q}_j(\theta)] \tag{6}$$

with

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

We now proceed by showing how the characterization of the necessary and sufficient conditions given above can be used to establish a few interesting results.

Nonimplementability of the Collusive Schedules.—It has long been noted that when the sellers' products are complements ($\lambda > 0$), it may be impossible to sustain the collusive schedules with a noncooperative equilibrium. However, this result has been established by restricting the principals to offering twice continuously differentiable tariffs $T : \Theta \to \mathbb{R}$, thus leaving open the possibility that it is merely a consequence of a technical assumption.[32] The approach suggested here permits one to verify that this result is true more generally.

PROPOSITION 2: *Let $q^c : \Theta \to \mathbb{R}$ be the function defined by*

$$q^c(\theta) \equiv \frac{1}{1 - \lambda}\left(\theta - \frac{1 - F(\theta)}{f(\theta)}\right) \forall\theta.$$

Assume that the sellers' products are complements ($\lambda > 0$), and $q^c(\theta) \in \text{int}(\mathcal{Q})$ for all $\theta \in \Theta$.[33] *The schedules $(q_i(\cdot))_{i=1}^2$ that maximize the sellers' joint profits are given by $q_i(\theta) = q^c(\theta)$ for all θ, $i = 1, 2$, and cannot be sustained by any equilibrium of in which the agent's strategy is Markovian.*

The proof in the Appendix uses the characterization of Proposition 1. By relying only on incentive compatibility, Proposition 2 guarantees that the aforementioned impossibility result is *by no means* a consequence of the assumptions one makes about the differentiability of the tariffs, or about the way one extends the tariffs outside the equilibrium range.

Sufficient Conditions for Differentiable Schedules.—We conclude this application by showing how the conditions in Proposition 1 can be used to construct equilibria supporting differentiable quantity schedules.

PROPOSITION 3: *Fix $\lambda \in (0,1)$, and let $q^* : \Theta \to \mathbb{R}$ be the solution to the differential equation*

$$\lambda\left[q(\theta)(1 - \lambda) - \theta + 2\left(\frac{1 - F(\theta)}{f(\theta)}\right)\right]\frac{dq(\theta)}{d\theta} = \theta - \frac{1 - F(\theta)}{f(\theta)} - q(\theta)(1 - \lambda), \tag{7}$$

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

[33] Note that this also requires $\lambda < 1$.

with boundary condition $q(\overline{\theta}) = \overline{\theta}/(1 - \lambda)$. *Suppose that* $q^* : \Theta \to \mathbb{R}$ *is nondecreasing and such that* $q^*(\theta) \in \mathcal{Q}$ *for all* $\theta \in \Theta$, *with* $q^*(\underline{\theta}) \geq [\overline{\theta} - \underline{\theta}]/\lambda$. *Then, let* $\tilde{q} : [0, \overline{\theta} + \lambda\overline{Q}] \to \mathcal{Q}$ *be the function defined by*

$$\tilde{q}(s) \equiv \begin{cases} 0 & \text{if} \quad s < m(\underline{\theta}) \\ q^*(m^{-1}(s)) & \text{if} \quad s \in [m(\underline{\theta}), m(\overline{\theta})] \\ q^*(\overline{\theta}) & \text{if} \quad s > m(\overline{\theta}), \end{cases} \tag{8}$$

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

$$h(q;\theta) \equiv \theta q + \int_{\theta}^{\theta+\lambda q} \tilde{q}(s)\,ds - \frac{q^2}{2} - \frac{1 - F(\theta)}{f(\theta)}[q + \tilde{q}(\theta + \lambda q) - \tilde{q}(\theta)] \tag{9}$$

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

The result in Proposition 3 offers a two-step procedure to construct an equilibrium with differentiable quantity schedules. The first step consists in solving the differential equation given in (7). The second step consists of checking whether the solution is nondecreasing, satisfies the boundary condition $q^*(\underline{\theta}) \geq [\overline{\theta} - \underline{\theta}]/\lambda$, and is such that the function $h(\cdot\,;\theta)$ defined in (9) is quasi-concave. If these properties are satisfied, the pair of schedules $q_i(\cdot) = q^*(\cdot)$, $i = 1, 2$, can be sustained by an equilibrium in which the agent's strategy is Markovian. The equilibrium features each principal i offering the menu of price quantity pairs ϕ_i^{M*} whose image is given by $\mathrm{Im}(\phi_i^{M*}) = \{(q_i, t_i) : (q_i, t_i) = (q_i(\theta), t_i(\theta)), \theta \in \Theta\}$ with $q_i(\cdot) = q^*(\cdot)$ and $t_i(\cdot) = t^*(\cdot)$, where, for any $\theta \in \Theta$,

$$t^*(\theta) = \theta q^*(\theta) - \int_{\underline{\theta}}^{\theta} q^*(s)\left[1 - \lambda\frac{\partial q^*(s)}{\partial s}\right]ds. \tag{10}$$

B. *Menu Auctions*

Consider now a menu auction environment à la Bernheim and Whinston (1985, 1986a): the agent's effort is verifiable and preferences are common knowledge (i.e., $|\Theta| = 1$).[34] As illustrated in the example in the introduction, assuming that the principals offer a single contract to the agent may preclude the possibility of sustaining interesting outcomes when preferences are not quasi-linear (more generally, when Peters (2003) *no-externalities* condition is violated). The question of interest is then how to identify the menus that sustain the equilibrium outcomes.

[34] See also Avinash Dixit, Gene M. Grossman, and Elhanan Helpman (1997); Bruno Biais, Martimort, and Jean-Charles Rochet (1997); Christine A. Parlour and Uday Rajan (2001); and Segal and Whinston (2003).

One approach is offered by Theorem 2. A profile of decisions (e^*, a^*) can be sustained by a pure-strategy equilibrium in which the agent's strategy is Markovian *if and only if* there exists a profile of incentive-compatible revelation mechanisms ϕ^{r*} and a profile of contracts δ^* that together satisfy the following conditions:

- Each mechanism ϕ_i^{r*} responds to the equilibrium contracts δ_{-i}^* by the other principals with the equilibrium contract δ_i^*, i.e., $\phi_i^{r*}(\delta_{-i}^*) = \delta_i^*$.
- Each contract δ_i^* responds to the equilibrium choice of effort e^* with the equilibrium action a_i^*, i.e., $\delta_i^*(e^*) = a_i^*$.
- Given the contracts δ^*, the agent's effort is optimal, i.e., $e^* \in \arg\max_{e \in E} v(e, \delta^*(e))$.
- For any contract $\delta_i \neq \delta_i^*$ by principal i, there exists a profile of contracts δ_{-i} by the other principals and a choice of effort e for the agent such that:

(a) each contract δ_j, $j \neq i$, can be obtained by truthfully reporting $(\delta_i, \delta_{-i-j})$ to P_j, i.e., $\delta_j = \phi_j^{r*}(\delta_{-j-i}, \delta_i)$;[35]

(b) given (δ_i, δ_{-i}), the agent's effort e is optimal, i.e., $e \in \arg\max_{\hat{e} \in E} v(\hat{e}, (\delta_i(\hat{e}), \delta_{-i}(\hat{e})))$, and there exists no other profile of contracts $\delta'_{-i} \in \times_{j \neq i} \mathrm{Im}(\phi_j^{r*})$ and effort choice e' that together give the agent a payoff higher than $(e, \delta_i, \delta_{-i})$, i.e., $v(e, (\delta_i(e), \delta_{-i}(e))) \geq v(e', (\delta_i(e'), \delta'_{-i}(e')))$ for any $e' \in E$ and any $\delta'_{-i} \in \times_{j \neq i} \mathrm{Im}(\phi_j^{r*})$;

(c) the payoff that principal i obtains by inducing the agent to select the contract δ_i is smaller that her equilibrium payoff, i.e., $u_i(e, (\delta_i(e), \delta_{-i}(e))) \leq u_i(e^*, a^*)$.

The approach described above uses incentive compatibility over *contracts*, i.e., it is based on revelation mechanisms that ask the agent to report the contracts selected with other principals. As anticipated in the example in the introduction, a more parsimonious approach consists in having the principals offer revelation mechanisms that simply ask the agent to report the actions a_{-i} that will be taken by the other principals.

DEFINITION 3: *Let $\mathring{\Phi}_i^r$ denote the set of mechanisms $\mathring{\phi}_i^r : E \times \mathcal{A}_{-i} \rightarrow \mathcal{A}_i$, such that, for any $e \in E$, any $a_{-i}, \hat{a}_{-i} \in \mathcal{A}_{-i}$*

$$v(e, \mathring{\phi}_i^r(e, a_{-i}), a_{-i}) \geq v(e, \mathring{\phi}_i^r(e, \hat{a}_{-i}), a_{-i}).$$

The idea is simple. In settings in which Peters (2003) no-externalities condition fails, for given choice of effort $e \in E$, the agent's preferences over the actions a_i by principal P_i depend on the actions a_{-i} by the other principals. A revelation mechanism $\mathring{\phi}_i^r$ is then a convenient tool for describing principal i's response to each *observable* effort choice e by the agent and to each *unobservable* profile of actions a_{-i} by the other principals, which is compatible with the agent's incentives. This last property is guaranteed by requiring that, for any (e, a_{-i}), the action $a_i = \mathring{\phi}_i^r(e, a_{-i})$ specified by the mechanism $\mathring{\phi}_i^r$ is as good for the agent as any other action a_i' that the agent can induce by reporting a profile of actions $\hat{a}_{-i} \neq a_{-i}$.

[35] Here, $\delta_{-j-i} \equiv (\delta_l)_{l \neq i,j}$.

Note, however, that while it is appealing to assume that the action a_i that the agent induces P_i to take depends only on (e, a_{-i}), restricting the agent's behavior to satisfying such a property may preclude the possibility of sustaining certain social choice functions. The reason is similar to the one indicated above when discussing the limits of Markov strategies. Such a restriction is, nonetheless, inconsequential when the principals' preferences are sufficiently aligned in the sense of the following definition.

DEFINITION 4 (Punishment with the same action): *We say that the "Punishment with the same action" condition holds if, for any* $i \in \mathcal{N}$, *compact set of decisions* $B \subseteq \mathcal{A}_i$, $a_{-i} \in \mathcal{A}_{-i}$, *and* $e \in E$, *there exists an action* $a_i' \in \arg\max_{a_i \in B} v(e, a_i, a_{-i})$, *such that for all* $j \neq i$, *all* $\hat{a}_i \in \arg\max_{a_i \in B} v(e, a_i, a_{-i})$,

$$v_j(e, a_i', a_{-i}) \leq v_j(e, \hat{a}_i, a_{-i}).$$

This condition is similar to the "Uniform Punishment" condition introduced above. The only difference is that it is stated in terms of *actions* as opposed to *contracts*. This difference permits one to restrict the agent's choice from each menu to depending only on his choice of effort and the actions taken by the other principals. The two definitions coincide when there is no action the agent must undertake after communicating with the principals, i.e., when $|E| = 1$, for in that case, a contract by P_i coincides with the choice of an action a_i. Lastly, note that the "Punishment with the same action" condition always holds in settings with only two principals, such as in the lobbying example considered in the introduction. We then have the following result.

PROPOSITION 4: *Assume that the principals' preferences are sufficiently aligned in the sense of the "Punishment with the same action" condition. Let* $\mathring{\Gamma}^r$ *be the game in which* P_i*'s strategy space is* $\Delta(\mathring{\Phi}_i^r)$, $i = 1, \ldots, n$. *A social choice function* π *can be sustained by a pure-strategy equilibrium of* Γ^M *if and only if it can be sustained by a pure-strategy truthful equilibrium of* $\mathring{\Gamma}^r$.

The simplified structure of the mechanisms $\mathring{\phi}^r$ proposed above permits one to restate the necessary and sufficient conditions for the equilibrium outcomes as follows. The action profile (e^*, a^*) can be sustained by a pure-strategy equilibrium of Γ^M if and only if there exists a profile of mechanisms $\mathring{\phi}^{r*}$ that satisfy the following properties:

- $a_i^* = \mathring{\phi}_i^{r*}(e^*, a_{-i}^*)$ all $i = 1, \ldots, n$;
- $v(e^*, a^*) \geq v(e', a')$ for any $(e', a') \in E \times \mathcal{A}$, such that $a_j' = \mathring{\phi}_j^{r*}(e', \hat{a}_{-j})$, $\hat{a}_{-j} \in \mathcal{A}_{-j}$, all $j = 1, \ldots, n$;
- for any i and any contract $\delta_i : E \to \mathcal{A}_i$, there exists a profile of actions (e, a), such that
  - (a) $a_i = \delta_i(e)$;
  - (b) $a_j = \mathring{\phi}_j^{r*}(e, a_{-j})$ all $j \neq i$;
  - (c) $v(e, a) \geq v(e', a')$ for any $(e', a') \in E \times \mathcal{A}$, such that $a_i' = \delta_i(e')$ and $a_j' = \mathring{\phi}_i^{r*}(e', \hat{a}_{-j})$ for some $\hat{a}_{-j} \in \mathcal{A}_{-j}$; and
  - (d) $u_i(e, a) \leq u_i(e^*, a^*)$.

As illustrated in the example in the introduction, this more parsimonious approach often simplifies the characterization of the equilibrium outcomes.

C. *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 these.

At first glance, using revelation mechanisms may appear prohibitively complicated in this setting due to the fact that the agent must report an entire array of incentive schemes to each principal. However, things simplify significantly—as long as for any array of incentive schemes, the choice of optimal effort for the agent is unique. It suffices to attach a label, say, an integer, to each incentive scheme a_i, and then have the agent report to each principal an array of integers, one for each other principal, along with the 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 the set of incentive schemes $\mathcal{A}$. The analysis of incentive compatibility then proceeds in the familiar way.

To illustrate, consider the following simplified version of a standard moral-hazard setting. There are two principals and two effort levels, $\underline{e}$ and $\overline{e}$. As in Bernheim and Whinston (1986b), the agent's preferences are common knowledge, so that $|\Theta| = 1$. Each principal i must choose an incentive scheme a_i from the set of feasible schemes $\mathcal{A}_i = \{a^l, a^m, a^h\}$, $i = 1, 2$. Here, a^l stands for a low-power incentive scheme, a^m for a medium-power one, and a^h for a high-power one.[36]

The typical moral hazard model specifies a Bernoulli utility function for each player defined over (w, e), where $w \equiv (w_i)_{i=1}^n$ stands for an array of rewards (e.g., monetary transfers) from the principals to the agent, together with the description of how the agent's effort determines a probability distribution over a set of verifiable outcomes used to determine the agent's reward. Instead of following this approach, in Table 2, we describe the players' expected payoffs (u_1, u_2, v) as a function of the agent's effort and the principals' incentive schemes.

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 insofar as the latter influences the agent's choice of effort. Nevertheless, Peters (2003) no-externalities condition fails here because the agent's preferences over the incentive schemes offered by P_i depend on the incentive scheme offered by P_j. By implication, restricting the principals to offering a single incentive scheme may preclude the possibility of sustaining certain outcomes, as we

[36] 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 a real number.

Table 2—Example: Common Agency with Normal Hazzard

	$e = \underline{e}$									$e = \overline{e}$									
$a_1 \backslash a_2$	a^h			a^m			a^l			$a_1 \backslash a_2$	a^h			a^m			a^l		
a^h	1	2	2	1	3	1	1	6	0	a^h	4	5	4	4	5	5	4	4	3
a^m	2	2	2	2	3	4	2	6	1	a^m	5	5	5	5	5	1	5	4	0
a^l	3	2	0	3	3	1	3	6	4	a^l	6	5	2	6	5	0	6	4	0

verify below.[37] 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 has 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 (a_1, a_2) as follows:

Table 3—Reduced Form of Table 2

$a_1 \backslash a_2$	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

Now suppose the principals were restricted to offering a single incentive scheme to the agent (i.e., to competing in take-it-or-leave-it offers). The unique pure-strategy equilibrium outcome would be $(a^h, a^m, \overline{e})$ with associated expected payoffs $(4, 5, 5)$.

When the principals are instead allowed to offer menus of incentive schemes, the outcome $(a^m, a^h, \overline{e})$ can also be sustained by a pure-strategy equilibrium.[38] The advantage of offering menus stems from the fact that they give the agent the possibility of punishing a deviation by the other principal by selecting a different incentive scheme with the nondeviating principal. Because the agent's preferences over a principal's incentive schemes, in turn, depend on the incentive scheme selected by the other principal, these menus can be conveniently described as revelation mechanisms $\phi_i^r : \mathcal{A}_j \to \mathcal{A}_i$ with the property that, for any a_j, $\phi_i^r(a_j) \in \arg\max_{a_i \in \mathrm{Im}(\phi_i^r)} V(a_i, a_j)$. Now, consider the mechanisms

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

Given these mechanisms, it is strictly optimal for the agent to choose (a^m, a^h) and then to select $e = \overline{e}$. Furthermore, given ϕ_{-i}^{r*}, it is easy to see that principal i has no profitable deviation, $i = 1, 2$, which establishes that $(a^m, a^h, \overline{e})$ can be sustained in equilibrium.

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

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

IV. Enriched Mechanisms

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

First, we consider pure-strategy equilibrium outcomes sustained by a strategy for the agent that is not Markovian. Next, we turn to mixed-strategy equilibrium outcomes.

Although the revelation mechanisms presented below are more complex than the ones considered in the previous sections, they still permit one to conceptualize the role that the agent plays in each bilateral relationship, thus possibly facilitating the characterization of the equilibrium outcomes.

A. *Non-Markov Strategies*

Here, we introduce a new class of revelation mechanisms that permit us to accommodate non-Markov strategies. We then adjust the notion of truthful equilibria accordingly, and finally prove that *any* outcome that can be sustained by a pure-strategy equilibrium of the menu game can also be sustained by a truthful equilibrium of 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* $\hat{\phi}_i^r : \hat{\mathcal{M}}_i^r \rightarrow \mathcal{D}_i$ *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* $\theta \in \Theta$, *any* $(\hat{m}_i^r, \hat{m}_{-i}^r) \in Supp\,[\mu(\theta, \hat{\phi}^r)]$,

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

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

The interpretation is that, in addition to (θ, δ_{-i}), the agent is now asked to report to each P_i the identity $k \in \mathcal{N}_{-i}$ of a deviating principal, 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 shown below, allowing a principal to response to (θ, δ_{-i}) with a contract that depends on the identity of a deviating principal may be essential to sustain certain outcomes when the agent's strategy is not Markovian.

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 deviates from equilibrium play, the agent truthfully reports to any of the nondeviating principals his true type θ, the contracts he is selecting with the other principals, and the identity k of the deviating principal. We then have the following result:

THEOREM 5:

(i) *Any social choice function* π *that can be sustained by a pure-strategy equilibrium of* Γ^M *can also be sustained by a pure-strategy truthful equilibrium of* $\hat{\Gamma}^r$.

(ii) *Furthermore, any* π *that can be sustained by an equilibrium of* $\hat{\Gamma}^r$ *can also be sustained by an equilibrium of* Γ^M.

Part (ii) follows from essentially the same arguments that establish part (ii) in Theorem 2).[39] Thus, consider part (i). The key step in the proof consists in showing that if the SCF π can be sustained by a pure-strategy equilibrium of Γ^M, it can also be sustained by an equilibrium in which the agent's strategy σ_A^{M*} has the following property. For any principal P_k, $k \in \mathcal{N}$, any contract $\delta_k \in \mathcal{D}_k$, and any type $\theta \in \Theta$, there exists a unique profile of contracts $\delta_{-k}(\theta, \delta_k) \in \mathcal{D}_{-k}$, such that A always selects $\delta_{-k}(\theta, \delta_k)$ with all principals other than k when (a) his type is θ, (b) the contract A selects with P_k is δ_k, and (c) P_k is the only deviating principal. In other words, the contracts that the agent selects with the nondeviating principals depend on the contract δ_k of the deviating principal, but not on the menus offered by the latter. The contracts $\delta_{-k}(\theta, \delta_k)$ minimize the payoff of the deviating principal P_k from among those contracts in the equilibrium menus of the nondeviating principals that are optimal for type θ given δ_k.

The rest of the proof 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 contracts he is selecting with the other principals are the equilibrium ones $\delta_{-i}^*(\theta)$, and that $k = 0$—then the revelation mechanism $\hat{\phi}_i^{r*}$ responds with the equilibrium contract $\delta_i^*(\theta)$. In contrast, when the agent reports that principal k deviated and that, as a result of such deviation, the agent selected the contract δ_k with P_k and the contracts $(\delta_j(\theta, \delta_k))_{j \neq i,k}$ with the other nondeviating principals, then the mechanism ϕ_i^{r*} responds with the contract $\delta_i(\theta, \delta_k)$ that, together with the contracts $(\delta_j(\theta, \delta_k))_{j \neq i,k}$, minimizes the payoff of the deviating principal P_k.[40] Given the equilibrium mechanisms $\hat{\phi}_{-k}^{r*}$, following a truthful strategy in these mechanisms 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

[39] Note that, in general, $\hat{\Gamma}^r$ is not an enlargement of Γ^M since certain menus in Γ^M may not be available in Γ^r, nor is Γ^M an enlargement of $\hat{\Gamma}^r$ since $\hat{\Gamma}^r$ may contains multiple mechanisms that offer the same menu.

[40] This is only a partial description of the equilibrium mechanisms $\hat{\phi}^{r*}$ and of the agent's strategy σ_A^{r*}. The complete description is in the Appendix.

TABLE 4

	$a_3 = s$			$a_3 = d$	
$a_1 \backslash a_2$	l	r	$a_1 \backslash a_2$	l	r
t	1 4 4 5	1 5 0 4	t	1 0 5 4	1 1 1 3
m	1 1 1 0	1 5 1 0	m	1 1 1 0	1 0 5 5
b	1 1 1 0	1 0 1 0	b	1 1 5 0	1 5 0 5

mechanisms $\hat{\phi}^{r*}_{-k}$ cannot do better than offering the equilibrium mechanism $\hat{\phi}^{r*}_k$ herself. We conclude that if the SCF π can be sustained by a pure-strategy equilibrium of Γ^M, it can also be sustained by a pure-strategy *truthful* equilibrium of $\hat{\Gamma}^r$.

To see why it may be essential with non-Markov strategies 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 where payoffs (u_1, u_2, u_3, v) are as in Table 4.

Because there is no effort in this example, a contract δ_i here simply coincides with the choice of an element of $\mathcal{A}_i$. It is then easy to see that the outcome (t, l, s) can be sustained by a pure-strategy equilibrium of the menu game Γ^M. The equilibrium features each P_i offering the menu that contains all contracts in $\mathcal{A}_i$. 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. 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-Markovian: given the contracts $(a_2, a_3) = (r, d)$ with P_2 and P_3, the contract that the agent chooses with P_1 depends on the particular menus offered by P_2 and P_3. This type of behavior is essential to sustain the equilibrium outcome. By implication, (t, l, s) cannot be sustained by an equilibrium of the revelation game in which the principals offer the simple mechanisms $\phi^r_i : \mathcal{A}_{-i} \rightarrow \mathcal{A}_i$ considered in the previous sections.[41] The outcome (t, l, s) can, however, be sustained by 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 contracts a_{-i}.[42]

B. *Mixed Strategies*

We now turn to equilibria in which the principals randomize over the menus they offer to the agent and/or the agent randomizes over the contracts he selects from the menus.[43]

[41] In fact, any incentive-compatible mechanism ϕ^r_1 that permits the agent to select the equilibrium contract t with P_1 must satisfy $\phi^r_i(a_2, a_3) = t$ 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 fails to provide the agent with either the contract m that is necessary to punish a deviation by P_2, or the contract b that is necessary to punish a deviation by P_3.

[42] Consistently with the result in Theorem 3, note that the problems with simple revelation mechanisms $\phi^r_i : \mathcal{A}_{-i} \rightarrow \mathcal{A}_i$ 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 choose different contracts 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 not sufficiently aligned so that the contract the agent must select with P_1 to punish a deviation by P_2 cannot be the same as the one he selects to punish a deviation by P_3.

[43] Recall that the notion of pure-strategy equilibrium of Definition 1 allows the agent to mix over effort.

Table 5

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

The reason why the simple mechanisms considered in Section II may fail to sustain certain mixed-strategy outcomes is that they do not permit the agent to select different contracts with the same principal in response to the same contracts δ_{-i} he is selecting with the other principals. To illustrate, consider the following example in which $|\Theta| = |E| = 1, n = 2, \mathcal{A}_1 = \{t, b\}$, $\mathcal{A}_2 = \{l, r\}$, and where payoffs (u_1, u_2, v) are as in Table 5.

Again, because there is no effort in this example, a contract for each P_i simply coincides with an element of $\mathcal{A}_i$. The following is then an equilibrium in the menu game. Each principal offers the menu ϕ_i^{M*} that contains all contracts in $\mathcal{A}_i$. Given the equilibrium menus, the agent selects with equal probabilities the contracts (t, l), (b, l), and (t, r). Note that, to sustain this outcome, it is essential that principals cannot offer lotteries over contracts. Indeed, if P_1 could offer a lottery over $\mathcal{A}_1$, she could do better by deviating from the strategy described above and offering the lottery that gives t and b with equal probabilities. In this case, A would strictly prefer to choose 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 constant over $E \times \mathcal{A}$, and when principals can offer lotteries over contracts, it is very difficult to construct examples where the agent is indifferent over some of the lotteries offered by the principals so that he can randomize, and no principal can benefit by breaking the agent's indifference so as to induce him to choose only those lotteries that are most favorable to her.

Nevertheless, it is important to note that, while certain *stochastic* SCFs may not be sustainable with the simple revelation mechanisms $\phi_i^r : \mathcal{D}_{-i} \to \mathcal{D}_i$ of the previous sections, *any* SCF that can be sustained by an equilibrium of the menu game can also be sustained by a truthful equilibrium of the following revelation game. The principals offer *set-valued* mechanisms $\tilde{\phi}_i^r : \Theta \times \mathcal{D}_{-i} \to 2^{\mathcal{D}_i}$ with the property that, for any $(\theta, \delta_{-i}) \in \Theta \times \mathcal{D}_{-i}$,[44]

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

The interpretation is that the agent first reports his type θ along with the contracts δ_{-i} that he is selecting with the other principals (possibly by mixing, or in response to a mixed strategy by the other principals). The mechanism then responds by offering

[44] With an abuse of notation, we will hereafter 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.

the agent the entire set $\tilde{\phi}_i^r(\theta, \delta_{-i})$ of contracts that are optimal for type θ, given δ_{-i}, out of those contracts that are available in $\tilde{\phi}_i^r$. Finally, the agent selects a contract from the set $\tilde{\phi}_i^r(\theta, \delta_{-i})$ and this contract is implemented.

In the example above, the equilibrium SCF can be sustained by having P_1 offer the mechanism $\tilde{\phi}_1^{r*}(l) = \{t, b\}$ and $\tilde{\phi}_1^{r*}(r) = \{t\}$; and by having P_2 offer the mechanism $\tilde{\phi}_2^{r*}(t) = \{l, r\}$ and $\tilde{\phi}_2^{r*}(b) = \{l\}$. Given the equilibrium mechanisms, with probability $1/3$ the agent then selects the contracts (t, l), with probability $1/3$ he selects the contracts (t, r), and with probability $1/3$ he selects the contracts (b, l). Note that a property of the mechanisms introduced above is that they permit the agent to select the equilibrium contracts by truthfully reporting to each principal the contracts selected with the other principals. For example, the contracts (t, l) can be selected by truthfully reporting l to P_1 and then choosing t from $\tilde{\phi}_1^{r*}(l)$, and by truthfully reporting t to P_2 and then choosing l from $\tilde{\phi}_2^{r*}(t)$. The equilibrium is thus *truthful* in the sense that the agent may well randomize over the contracts he selects with the principals, but once he has chosen which contracts he wants (i.e., for any realization of his mixed strategy), he always reports these contracts truthfully to each principal.

Next, note that while the revelation mechanisms introduced above are conveniently described by the correspondence $\tilde{\phi}_i^r : \Theta \times \mathcal{D}_{-i} \rightarrow 2^{\mathcal{D}_i}$, formally any such mechanism is a standard single-valued mapping $\overline{\phi}_i^r : \mathcal{M}_i^r \rightarrow \mathcal{D}_i$ with message space $\tilde{\mathcal{M}}_i^r \equiv \Theta \times \mathcal{D}_{-i} \times \mathcal{D}_i$ such that[45]

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

These mechanisms are clearly incentive-compatible in the sense that, given (θ, δ_{-i}), the agent strictly prefers *any* contract in $\tilde{\phi}_i^r(\theta, \delta_{-i})$ to any contract that can be obtained by reporting (θ', δ'_{-i}). Furthermore, as anticipated above, given any profile of mechanisms $\tilde{\phi}^r$, the contracts 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$[\mu(\theta, \tilde{\phi}_i^r, \tilde{\phi}_{-i}^r)]$,

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

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

[45] The particular contract 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.

The agent's strategy is thus said to be truthful in $\tilde{\phi}_i^r$ if the message $\hat{m}_i^r = (\theta, \delta_{-i}, \delta_i)$, which the agent sends to principal i, coincides with his true type θ along with the true contracts $\delta_{-i} = (\overline{\phi}_j^r(\tilde{m}_j^r))_{j \neq i}$ that the agent selects with the other principals by sending the messages $\tilde{m}_{-i}^r$, and the contract $\delta_i = \overline{\phi}_i^r(\tilde{m}_i^r)$ that A selects with P_i by sending the message $\tilde{m}_i^r$. We then have the following result:

THEOREM 6: *A social choice function* $\pi : \Theta \to \Delta(E \times \mathcal{A})$ *can be sustained by an equilibrium of* Γ^M *if and only if it can be sustained by a truthful equilibrium of* $\tilde{\Gamma}^r$.

The proof is similar to the one that establishes the Menu Theorems (e.g., Peters 2001). The reason that 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 that can be offered through the revelation mechanisms of $\tilde{\Gamma}^r$ are only those that satisfy the following property. For each contract 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 contract in the menu.[46] That the principals can be restricted to offering menus that satisfy this property should not be surprising. The proof, however, requires some work to show how the agent's and the principals' mixed strategies must be adjusted to preserve the same distribution over outcomes as in the 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 contracts the agent finds it optimal to choose as a function of the contracts he selects with the other principals. This, in turn, can facilitate the characterization of the equilibrium outcomes in applications in which mixed strategies are appealing.

V. Conclusions

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

When compared to universal mechanisms, the mechanisms proposed here have the advantage that they do not lead to the problem of infinite regress, for they do not require the agent to describe the mechanisms offered by the other principals.

When 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" (i.e., his exogenous type) and "what he is doing with the other principals" (i.e., the contracts he is selecting in the other relationships). The advantage of describing the agent's choice from a menu by means of a revelation mechanism is that this often facilitates the characterization of the necessary and sufficient conditions for the equilibrium

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

outcomes. We have illustrated such a possibility in a few cases of interest: competition in nonlinear tariffs with adverse selection; menu auctions; and moral hazard settings.

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 outcomes sustained by non-Markov strategies and/or mixed-strategy equilibria.

Throughout the analysis, we maintained the assumption that the multiple principals contract with a single common agent. Clearly, the results are also useful in games with multiple agents, provided that the contracts that each principal offers to each of her agents do not depend on the contracts offered to the other agents (see also Seungjin Han, 2006, for a similar restriction.) More generally, it has recently been noted that in games in which multiple principals contract simultaneously with three or more agents (or those in which principals also communicate among themselves), a "folk theorem" holds: all outcomes yielding each player a payoff above the Max-Min value can be sustained in equilibrium (Takuro Yamashita, 2007; and Peters and Christian Troncoso Valverde, 2009). While these results are intriguing, they also indicate that, to retain predictive power, it is now time for the theory of competing mechanisms to accommodate restrictions on the set of feasible mechanisms and/or on the agents' behavior. These restrictions should, of course, be motivated by the application under examination. For many applications, we find appealing the restriction imposed by requiring that the agents' behavior be Markovian. Investigating the implications of such a restriction for games with multiple agents is an interesting line for future research.

Appendix 1: Take-It-or-Leave-It-Offer Equilibria in the Menu-Auction Example in the Introduction

Assume that the principals are restricted to making take-it-or-leave-it offers to the agent, that is, to offering a single contract $\delta_i : E \to [0,1]$. Denote by e^* the equilibrium policy, and by $(\delta_i^*)_{i=1,2}$ the equilibrium contracts.

- We start by considering (pure-strategy) equilibria sustaining $e^* = p$. First, note that if an equilibrium exists in which $\delta_2^*(p) > 0$, then necessarily $\delta_1^*(p) = 1$. Indeed, if $\delta_1^*(p) < 1$, then P_1 could deviate and offer a contract δ_1, such that $\delta_1(p) = 1$ and $\delta_1(f) = \delta_1^*(f)$. Such a deviation would ensure that A strictly prefers $e = p$ and would give P_1 a strictly higher payoff. Thus, if $\delta_2^*(p) > 0$, then necessarily $\delta_1^*(p) = 1$. This result, in turn, implies that, if an equilibrium exists in which $\delta_2^*(p) > 0$, then necessarily $\delta_2^*(p) = 1$. Or else, P_2 could offer herself a contract δ_2, such that $\delta_2(p) = 1$ and $\delta_2(f) = \delta_2^*(f)$, ensuring that A strictly prefers $e = p$ and obtaining a strictly higher payoff. Finally, observe that there exists no equilibrium sustaining $e^* = p$ in which $\delta_2^*(p) = 0$. This follows directly from the fact that $v(p, \delta_1^*(p), 0) < v(f, a_1, a_2)$, for any $\delta_1^*(p)$ and any (a_1, a_2). We conclude that any equilibrium sustaining $e^* = p$ must be such that $\delta_i^*(p) = 1$, $i = 1,2$. That such an equilibrium exists follows from the fact that it can be sustained, for example, by the following contracts: $\delta_i^*(e) = 1$ all e, $i = 1,2$. Given δ_i^* and δ_2^*, A

strictly prefers $e = p$. Furthermore, when $a_{-i} = 1$, each P_i strictly prefers $e = p$, which guarantees that no principal has a profitable deviation.

- Next, consider equilibria sustaining $e^* = f$. In any such equilibrium, necessarily $\delta_1^*(f) > 1/2$. Indeed, suppose that there existed an equilibrium in which $\delta_1^*(f) \leq 1/2$. Then, necessarily $\delta_2^*(f) = 1$. This follows from (i) the fact that, for any a_2, $v(f, \delta_1^*(f), a_2) > 2$ whenever $\delta_1^*(f) \leq 1/2$; and (ii) the fact that, for any a_1, $v(p, a_1, 0) = 1$. Taken together these properties imply that, if $\delta_1^*(f) \leq 1/2$ and $\delta_2^*(f) < 1$, then P_2 could deviate and offer a contract such that $\delta_2(f) = 1$ and $\delta_2(p) = 0$. Such a contract would guarantee that A strictly prefers $e = f$ and, at the same time, would give P_2 a strictly higher payoff than the proposed equilibrium contract, which is clearly a contradiction. Hence, if an equilibrium existed in which $\delta_1^*(f) \leq 1/2$, then necessarily $\delta_2^*(f) = 1$, but P_1 would have a profitable deviation that consists in offering the agent a contract such that $\delta_1(f) = 1$ and $\delta_1(p) = 0$. Such a contract would induce A to select $e = f$ and would give P_1 a payoff strictly higher than the proposed equilibrium payoff, once again a contradiction. We thus conclude that, if an equilibrium sustaining $e^* = f$ exists, it must be such that $\delta_1^*(f) > 1/2$. But then, in any such equilibrium, necessarily $\delta_2^*(f) = 1$. This follows from the fact that, when $e = f$ and $a_1 > 1/2$, both A's and P_2's payoffs are strictly increasing in a_2. But if $\delta_2^*(f) = 1$, then necessarily $\delta_1^*(f) = 1$. Else, P_1 could deviate and offer a contract such that $\delta_1(f) = 1$ and $\delta_1(p) = 0$. Such a contract would guarantee that A strictly prefers $e = f$ and would give P_1 a payoff strictly higher than the one she obtains under any contract that sustains $e = f$ with $\delta_1(f) < 1$. We conclude that in any equilibrium in which $e^* = f$, necessarily $\delta_1^*(f) = \delta_2^*(f) = 1$. The following pair of contracts then supports the outcome $(f, 1, 1)$: $\delta_i^*(f) = 1$, and $\delta_i^*(p) = 0$, $i = 1, 2$. Note that, given δ_{-i}^*, there is no way P_i can induce A to switch to $e = p$. Furthermore, when $e = f$ and $a_{-i} = 1$, each P_i's payoff is maximized at $a_i = 1$. Thus, no principal has a profitable deviation.

Appendix 2: Omitted Proofs

As explained in Section I, to ease the exposition, throughout the main text, we restricted attention to settings where the principals offer the agent *deterministic* contracts. However, all our results apply to more general settings where the principals can offer the agent mechanisms that map messages into *lotteries* over *stochastic* contracts. All proofs here in the Appendix thus refer to these more general settings.

Below, we show how the model set up of Section I must be adjusted to accommodate these more general mechanisms and then turn to the proofs of the results in the main text.

Let Y_i denote the set of feasible *stochastic contracts* for P_i. A stochastic contract $y_i : E \to \Delta(\mathcal{A}_i)$ specifies a distribution over P_i's actions $\mathcal{A}_i$, one for each possible effort $e \in E$. Next, let $\mathcal{D}_i \subseteq \Delta(Y_i)$ denote a (compact) set of feasible *lotteries* over Y_i and denote, by $\delta_i \in \mathcal{D}_i$, a generic element of $\mathcal{D}_i$. Clearly, depending on the application of interest, the set $\mathcal{D}_i$ of feasible lotteries may be more or less restricted. For example, the deterministic environment considered in the main text corresponds to a setting where each set $\mathcal{D}_i$ contains only degenerate lotteries (i.e., Dirac measures) that assign probability one to contracts that responds to each effort $e \in E$ with a degenerate distribution over $\mathcal{A}_i$.

Given this new interpretation for $\mathcal{D}_i$, we then continue to refer to a mechanism as a mapping $\phi_i : \mathcal{M}_i \to \mathcal{D}_i$. However, note that, given a message $m_i \in \mathcal{M}_i$, a mechanism now responds by selecting a (stochastic) contract y_i from Y_i using the lottery $\delta_i = \phi_i(m_i) \in \Delta(Y_i)$. The timing of events must then be adjusted as follows:

- 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 $\phi = (\phi_1, \ldots, \phi_n)$. The messages $m = (m_1, \ldots, m_n)$ are sent simultaneously.
- At $t = 3$, the contracts $y = (y_1, \ldots, y_n)$ are drawn from the (independent) lotteries $\delta = (\phi_1(m_1), \ldots, \phi_n(m_n))$.
- At $t = 4$, A chooses $e \in E$ after observing the contracts $y = (y_1, \ldots, y_n)$.
- At $t = 5$, the principals' actions $a = (a_1, \ldots, a_n)$ are determined by the (independent) lotteries $(y_1(e), \ldots, y_n(e))$, and payoffs are realized.

Both the principals' and the agent's strategies continue to be defined as in the main text. However, note that the agent's effort strategy $\xi : \Theta \times \Phi \times \mathcal{M} \times \mathrm{Y} \to \Delta(E)$ is now contingent also on the realizations y of the lotteries $\delta = \phi(m)$. The strategy $\sigma_A = (\mu, \xi)$ is then said to be a continuation equilibrium if for every (θ, ϕ, m, y), any $e \in Supp[\xi(\theta, \phi, m, y)]$ maximizes

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

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

$$\int_{Y1} \cdots \int_{Y_n} \max_{e \in E} \overline{V}(e; y, \theta)\, d\phi_1(m_1) \times \cdots \times d\phi_n(m_n).$$

We then denote, by

$$V(\delta, \theta) \equiv \int_{Y1} \cdots \int_{Y_n} \max_{e \in E} \overline{V}(e; y, \theta)\, d\delta_1 \times \cdots \times d\delta_n,$$

the maximal payoff that type θ can obtain given the principals' lotteries δ. All results in the main text apply *verbatim* to this more general setting provided that (i) one reinterprets $\delta_i \in \Delta(\mathrm{Y}_i)$ as a *lottery* over the set of (feasible) stochastic contracts Y_i, as opposed to a deterministic contract $\delta_i : E \to \mathcal{A}_i$; and (ii) one reinterprets $V(\delta, \theta)$ as the agent's *expected* payoff given the lotteries δ, as opposed to his deterministic payoff.

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 Markovian and which implements π, then there also exists a truthful pure-strategy equilibrium σ^{r*} of Γ^r which 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 Markovian, then for any i and any $(\theta, \delta_{-i}, \phi_i^M)$, there exists a unique $\delta_i(\theta, \delta_{-i}; \phi_i^M) \in \text{Im}(\phi_i^M)$, such that A always selects $\delta_i(\theta, \delta_{-i}; \phi_i^M)$ with P_i when the latter offers the menu ϕ_i^M, the agent's type is θ, and the lotteries A selects with the other principals are δ_{-i}. Finally, let $\delta^*(\theta) = (\delta_i^*(\theta))_{i=1}^n$ denote the equilibrium lotteries that type θ selects 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 P_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 P_i the message $m_i^r = (\theta, \delta_{-i}^*(\theta))$, thus selecting $\delta_i^*(\theta)$ with each P_i. Given the contracts y selected by the lotteries $\delta^*(\theta)$, then each type θ chooses the same distribution over effort he would have selected 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 $\text{Im}(\phi_i^M) = \text{Im}(\phi_i^r)$. That is, let $\delta(\theta; \phi^M)$ denote the lotteries that type θ would have selected in Γ^M given ϕ^M. Then given ϕ^r, A selects the lottery $\delta_i(\theta; \phi^M)$ 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 lotteries $\delta(\theta; \phi^M)$ as in Γ^M. In the continuation game that starts after the contracts y are drawn, A then chooses the same distribution over effort he would have chosen in Γ^M given the contracts y, the menus ϕ^M, and the lotteries $\delta(\theta; \phi^M)$.

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 $\text{Im}(\phi_i^M) = \text{Im}(\phi_i^r)$ for all i.

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

Given the continuation equilibrium σ_A^{r*}, any principal P_i, who expects the other principals 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 a truthful 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 image 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 $B \subseteq \Phi_i^M$,

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

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 $\phi^r \in \Phi^r(\phi^M) \equiv \Pi_i \Phi_i^r(\phi_i^M)$. Precisely, let $\rho_{\sigma_A^{r*}} : \Theta \times \Phi^r \rightarrow \Delta(\mathcal{A} \times E)$ 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 the event that ϕ_i^r belongs 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 $\phi_{-j}^r \in \Phi_{-j}^r(\phi_{-j}^M) \equiv \Pi_{i \neq j} \; \Phi_i^r(\phi_i^M)$. That is, for any $\theta \in \Theta$,

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

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*} is 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_j^{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 condition (i) holds, the result is immediate. In what follows, we prove that when condition (ii) holds, then if the SCF π can be sustained by a pure-strategy equilibrium σ^{M*} of Γ^M, it can also be sustained by a pure-strategy equilibrium $\hat{\sigma}^M$ in which the agent's strategy $\hat{\sigma}_A^M$ is Markovian.

Let ϕ^{M*} denote the equilibrium menus under the strategy profile σ^{M*}, and let δ^* denote the equilibrium lotteries that are selected by the agent when all principals offer the equilibrium menus ϕ^{M*}.

Suppose that σ_A^{M*} is not Markovian. 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 $\underline{\phi}_{-i}^M, \overline{\phi}_{-i}^M \in \Phi_{-i}^M$, such that A selects $(\underline{\delta}_i, \delta'_{-i})$ when $\phi^M = (\tilde{\phi}_i^M, \underline{\phi}_{-i}^M)$ and $(\overline{\delta}_i, \delta'_{-i})$ when $\phi^M = (\tilde{\phi}_i^M, \overline{\phi}_{-i}^M)$, with $\underline{\delta}_i \neq \overline{\delta}_i$. Below, we show that when this is the case, then, starting from σ_A^{M*}, one can construct a Markovian continuation equilibrium $\hat{\sigma}_A^M$ which induces all principals to continue to offer the equilibrium menus ϕ^{M*} and sustains the same outcomes as σ_A^{M*}.

Case 1: First consider the case in which $\tilde{\phi}_i^M = \phi_i^{M*}$ and $\delta'_{-i} = \delta^*_{-i}$. Then, let $\hat{\sigma}_A^M$ be the strategy that coincides with σ_A^{M*} for all $\phi^M \neq (\tilde{\phi}_i^M, \underline{\phi}_{-i}^M), (\tilde{\phi}_i^M, \overline{\phi}_{-i}^M)$ and that prescribes that A selects δ^* both when $\phi^M = (\tilde{\phi}_i^M, \underline{\phi}_{-i}^M)$ and when $\phi^M = (\tilde{\phi}_i^M, \overline{\phi}_{-i}^M)$. In the continuation game that starts after the lotteries δ^*, select the contracts y, $\hat{\sigma}_A^M$ then prescribes that A induces the same distribution over effort he would have induced according to the original strategy σ_A^{M*} had the menus offered been ϕ^{M*}. Clearly, if the strategy σ_A^{M*} was sequentially rational, so is $\hat{\sigma}_A^M$. Furthermore, it is easy to see 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 continuing to offer the equilibrium menu ϕ_j^{M*}.

Case 2: Next, consider the case in which $\tilde{\phi}_i^M = \phi_i^{M*}$, but where $\delta'_{-i} \neq \delta^*_{-i}$ (which implies that both $\underline{\phi}_{-i}^M$ and $\overline{\phi}_{-i}^M$ are necessarily different from ϕ_{-i}^{M*}). For any $j \in \mathcal{N}$, any $\delta \in \mathcal{D}$, let $\underline{U}_j(\delta)$ denote the lowest payoff that the agent can inflict to principal P_j, without violating his rationality. This payoff is given by

$$(12)\quad \underline{U}_j(\delta) \equiv \int_Y \left[\int_{\mathcal{A}} u_j(a, \xi_j(y))dy_1(\xi_j(y)) \times \cdots \times dy_n(\xi_j(y))\right] d\delta_1 \times \cdots \times d\delta_n,$$

where for any $y \in Y$,

$$(13)\qquad \xi_j(y) \in \arg\min_{e \in E^*(y)} \left\{\int_{\mathcal{A}} u_j(a, e)dy_1(e) \times \cdots \times dy_n(e)\right\}$$

with

$$E^*(y) \equiv \arg\max_{e \in E} \left\{\int_{\mathcal{A}} v(a, e)dy_1(e) \times \cdots \times dy_n(e)\right\}.$$

Now, let $\hat{\sigma}_A^M$ be the strategy that coincides with σ_A^{M*} for all $\phi^M \neq (\tilde{\phi}_i^M, \underline{\phi}_{-i}^M), (\tilde{\phi}_i^M, \overline{\phi}_{-i}^M)$ and that prescribes that A selects $(\delta'_i, \delta'_{-i})$ both when $\phi^M = (\tilde{\phi}_i^M, \underline{\phi}_{-i}^M)$ and when $\phi^M = (\tilde{\phi}_i^M, \overline{\phi}_{-i}^M)$, where $\delta_i' \in \arg\max_{\delta_i \in \mathrm{Im}(\tilde{\phi}_i^M)} V(\delta_i, \delta'_{-i})$ is any contract 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 contract always exists. In the continuation game that starts after the lotteries, $\delta = (\delta'_i, \delta'_{-i})$ selects the contracts y, A then selects effort $\xi_k(y)$, where

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

is the identity of one of the deviating principals, and where $\xi_k(y)$ is the level of effort defined in (13). Clearly, when $|\{j \in \mathcal{N} \setminus \{i\} : \phi_j^M \neq \phi_j^{M*}| > 1$, the identity k of the deviating principal can be chosen arbitrarily. Once again, it is easy to see 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 continuing to offer the equilibrium menu ϕ_l^{M*}.

Case 3: Lastly, consider the case where $\tilde{\phi}_i^M \neq \phi_i^{M*}$. 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 $\phi^M \neq (\tilde{\phi}_i^M, \underline{\phi}_{-i}^M)$, $(\tilde{\phi}_i^M, \overline{\phi}_{-i}^M)$ and that prescribes that A selects $(\delta'_i, \delta'_{-i})$ both when $\phi^M = (\tilde{\phi}_i^M, \underline{\phi}_{-i}^M)$ and when $\phi^M = (\tilde{\phi}_i^M, \overline{\phi}_{-i}^M)$, where $\delta'_i \in \arg\max_{\delta_i \in \mathrm{Im}(\tilde{\phi}_i^M)} V(\delta_i, \delta'_{-i})$ is any contract such that

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

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

This completes the description of the strategy $\hat{\sigma}_A^M$. Now, 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 lottery over contracts that the agent selects with any principal P_j, $j \neq i$, is the same as under the original strategy σ_A^{M*}. When combined together, these properties imply that the procedure described above can be iterated for all $i \in \mathcal{N}$, all $\tilde{\phi}_i^M \in \Phi_i^M$. This gives a new strategy for the agent that is Markovian, that induces all principals to continue to offer the equilibrium menus ϕ^{M*}, and that implements the same outcomes as σ_A^{M*}.

PROOF OF THEOREM 4:

The result follows from the same construction as in the proof of Theorem 3, now applied to each $\theta \in \Theta$, 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 $\underline{\phi}^M_{-i}, \overline{\phi}^M_{-i} \in \Phi^M_{-i}$, such that some type $\theta \in \Theta$ selects $(\underline{\delta}_i, \delta^*_{-i}(\theta))$ when $\phi^M = (\phi^{M*}_i, \underline{\phi}^M_{-i})$ and $(\overline{\delta}_i, \delta^*_{-i}(\theta))$ when $\phi^M = (\phi^{M*}_i, \overline{\phi}^M_{-i})$, with $\underline{\delta}_i \neq \overline{\delta}_i$. In other words, Case 1 in the proof of Theorem 3 is never possible when the strategy σ^{M*}_A satisfies the "Conformity to Equilibrium" condition. This, in turn, guarantees that when one replaces the original strategy σ^{M*}_A with the strategy $\hat{\sigma}^M_A$ obtained from σ^{M*}_A, iterating the steps in the proof of Theorem 3 for all $\theta \in \Theta$, all $i \in \mathcal{N}$, and all $\tilde{\phi}^M_i \in \Phi^M_i$, it remains optimal for each P_i to offer the equilibrium menu ϕ^{M*}_i.

PROOF OF PROPOSITION 1:

One can immediately see that conditions (i)–(iii) 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 by an equilibrium of the menu game Γ^M.

The proof below establishes the necessity of these conditions. That conditions (i) and (ii) are necessary follows directly from Theorem 2. If the schedules $q^*_i(\cdot)$, $i = 1,2$, can be sustained by a pure-strategy equilibrium of Γ^M, in which the agent's strategy is Markovian, then they can also be sustained by 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 is such that $\phi^r_i(\theta, q_j, t_j) = \phi^r_i(\theta', q'_j, t'_j)$, 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 $\phi^{r*}_i = (\tilde{q}_i(\cdot), \tilde{t}_i(\cdot))$, $i = 1,2$, such that $\tilde{q}_i(\cdot)$ is nondecreasing, $\tilde{t}_i(\cdot)$ satisfies (1), and conditions (i) and (ii) in the proposition hold.

It remains to show that condition (iii) 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^r_A that sustain the schedules $q^*_i(\cdot)$, $i = 1,2$, in Γ^r, then it must be that the schedules $q^*_i(\cdot)$ and $t^*_i(\cdot) \equiv \tilde{t}_i(m_i(\cdot))$, $i = 1,2$, satisfy the equivalent of the (IC) and (IR) constraints of program $\tilde{\mathcal{P}}$ in the main text. In turn, this means that necessarily $U^*_i \leq \overline{U}_i$, $i = 1,2$. To prove the result it then suffices to show that if $U^*_i < \overline{U}_i$, then P_i has a profitable deviation.

This property can be established by contradiction. Suppose that there exists a truthful equilibrium $\sigma^r \in \mathcal{E}(\Gamma^r)$ which sustains the schedules $(q^*_i(\cdot))_{i=1,2}$ and such that $U^*_i < \overline{U}_i$, for some $i \in \mathcal{N}$. Then there also exists a (pure-strategy) equilibrium σ^{M*} of Γ^M which sustains the same schedules and such that each P_i offers the menu ϕ^{M*}_i defined by $\text{Im}(\phi^{M*}_i) = \text{Im}(\phi^{r*}_i)$, and each type θ selects the contract $(q^*_i(\theta), t^*_i(\theta))$ from each menu ϕ^{M*}_i, thus giving P_i a payoff U^*_i. (See the proof of part 2 of Theorem 2.) Below, we, however, show that this cannot be the case, Irrespective of which continuation equilibrium σ^{M*}_A one considers, P_i has a profitable deviation, which establishes the contradiction.

Case 1: Suppose that the schedules $q_i(\cdot)$ and $t_i(\cdot)$ that solve the program $\tilde{\mathcal{P}}$ defined in the main text are such that the set of types $\theta \in \Theta$, who strictly prefer the contract $(q_i(\theta), t_i(\theta))$ to any other contract $(q_i, p_i) \in \{(q_i(\theta'), t_i(\theta')) : \theta' \in \Theta, \theta' \neq \theta\} \cup \{(0,0)\}$, in the sense defined by the IC and IR constraints, has (probability) measure one. When this is the case, principal P_i has a profitable deviation in Γ^M that consists of offering the menu ϕ_i^M defined by $\text{Im}(\phi_i^M) = \{(q_i(\theta), t_i(\theta)) : \theta \in \Theta\}$. Irrespective of which particular continuation equilibrium σ_A^{M*} one considers, given $(\phi_i^M, \phi_{-i}^{M*})$, almost every type θ must necessarily choose the contract $(q_i(\theta), t_i(\theta))$ from ϕ_i^M, thus giving P_i a payoff $\overline{U}_i > U_i^*$.[47]

Case 2: Next, suppose that the schedules $q_i(\cdot)$ and $t_i(\cdot)$ that solve the program $\tilde{\mathcal{P}}$ are such that almost every $\theta \in \Theta$ strictly prefers the contract $(q_i(\theta), t_i(\theta))$ to any other contract $(q_i, p_i) \in \{(q_i(\theta'), t_i(\theta')) : \theta' \in \Theta, \theta' \neq \theta\}$, again in the sense defined by the IC constraints. However, now suppose that 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 by P_i to the menu whose image is $\text{Im}(\phi_i^M) = \{(q_i(\theta), t_i(\theta)) : \theta \in \Theta\}$ need not be profitable for P_i. In fact, any type $\theta' \in \Theta'$ could punish such a deviation by choosing not to participate (equivalently, by choosing the null contract $(0,0)$). However, if this is the case, then P_i could offer the menu $\phi_i^{M'}$, such that $\text{Im}(\phi_i^{M'}) = \{(q_i'(\theta), t_i'(\theta)) : \theta \in \Theta\}$, where, for any $\theta \in \Theta$, $q_i'(\theta) \equiv q_i(\theta)$ and $t_i'(\theta) \equiv t_i(\theta) - \varepsilon$, $\varepsilon > 0$. Clearly, any such menu guarantees participation by all types. Furthermore, by choosing $\varepsilon > 0$ small enough, P_i can guarantee herself a payoff arbitrarily close to $\overline{U}_i > U_i^*$, 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 contract $(q_i(\theta'), t_i(\theta'))$ specified by the schedules $q_i(\cdot)$ and $t_i(\cdot)$ for type θ', and then selecting the contract $(\tilde{q}_j(\theta + \lambda q_i(\theta')), \tilde{t}_j(\theta + \lambda q_i(\theta'))$ with principal P_j, where $q_i(\cdot)$ and $t_i(\cdot)$ are, again, the schedules that solve program $\tilde{\mathcal{P}}$ in the main text. Now, suppose that the schedules $q_i(\cdot)$ and $t_i(\cdot)$ are such that there exists a positive-measure set of types $\Theta_0 \subset \Theta$, such that for any $\theta \in \Theta_0$, there exists a $\theta' \in \Theta$, and such that

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

with $q_i(\theta') \neq q_i(\theta)$;[48] and 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 contract $(q_i(\theta), t_i(\theta))$ is not strictly optimal, in the sense that there exists another contract $(q_i(\theta'), t_i(\theta'))$ with $(q_i(\theta'), t_i(\theta')) \neq (q_i(\theta), t_i(\theta))$ that is as good for type θ as the contract $(q_i(\theta), t_i(\theta))$.

[47] Note that while almost every $\theta \in \Theta$ strictly prefers $(q_i(\theta), t_i(\theta))$ to any other pair $(q_i, p_i) \in \text{Im}(\phi_i^M) \cup \{(0,0)\}$, there may exist a positive-measure set of types θ', who, given $(q_i(\theta'), t_i(\theta'))$, are indifferent between choosing the contract $(\tilde{q}_j(\theta' + \lambda q_i(\theta')), \tilde{t}_j(\theta' + \lambda q_i(\theta'))$ with P_j or choosing another contract $(q_j, t_j) \in \text{Im}(\phi_j^{M*})$. 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 the agent's choice with P_j.

[48] Clearly, 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.

Without loss of generality, assume that the schedules $q_i(\cdot)$ and $t_i(\cdot)$ are such that each type $\theta \in \Theta$ strictly prefers the contract $(q_i(\theta), t_i(\theta))$ to the null contract $(0,0)$. As shown in Case 2, when this property is not satisfied, there always exists another pair of schedules $q'_i(\cdot)$ and $t'_i(\cdot)$ that guarantee participation by all types, preserve incentive compatibility for all θ, and yield P_i a payoff $U_i > U_i^*$.

Now, given $q_i(\cdot)$ and $t_i(\cdot)$, let $z : \Theta \rightrightarrows \Theta \cup \{\varnothing\}$ be the correspondence defined by

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

and denote by $z(\Theta) \equiv \operatorname{Im}(z)$ the range of $z(\cdot)$. 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 $(q_i(\theta), t_i(\theta))$ specified by $q_i(\cdot), t_i(\cdot)$ for type θ, but which, nonetheless, gives type θ the same payoff as the contract $(q_i(\theta), t_i(\theta))$.

Next, let $g : \Theta \rightrightarrows \Theta \cup \{\varnothing\}$ denote the correspondence defined by

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

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 contract as type θ. Finally, given any set $\Theta' \subset \Theta$, let

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

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 $\theta \notin \Theta_0 \cup g(\Theta_0)$, and, for any $\theta \in \Theta_0 \cup g(\Theta_0)$, $t'_i(\theta) = t_i(\theta) - \varepsilon$ with $\varepsilon > 0$.[49] Clearly, if $\varepsilon > 0$ is chosen sufficiently small, then the new schedules $q'_i(\cdot)$ and $t'_i(\cdot)$ continue to satisfy the (IC) and (IR) constraints of program $\tilde{\mathcal{P}}$ 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 $t'_i(\cdot)$ constructed above guarantee that each type $\theta \in \Theta$ now strictly prefers the contract $(q'_i(\theta), t'_i(\theta))$ to any other contract $(q'_i(\theta'), t'_i(\theta')) \neq (q'_i(\theta), t'_i(\theta))$. This, in turn, implies that, irrespective of the agent's continuation equilibrium σ_A^M, P_i can guarantee herself a payoff arbitrarily close to $\overline{U}_i$ by choosing $\varepsilon > 0$ sufficiently small and offering the menu $\phi_i^{M'}$, such that $\operatorname{Im}(\phi_i^{M'}) = \{(q'_i(\theta), t'_i(\theta)) : \theta \in \Theta\}$. Thus, starting from ϕ_i^{M*}, P_i has again a profitable deviation.

Next, suppose that $\{\Theta_0 \cup g(\Theta_0)\} \cap z(\Theta) \neq \varnothing$. Note that this also implies that $\Theta_0 \cap z(\Theta) \neq \varnothing$. To see this, note that for any $\hat{\theta} \in g(\Theta_0) \cap z(\Theta)$, with $\hat{\theta} \notin \Theta_0$, there exists a $\theta' \in \Theta_0$ such that $(q_i(\theta'), t_i(\theta')) = (q_i(\hat{\theta}), t_i(\hat{\theta}))$. But then, by definition of z, $\theta' \in z(\Theta)$. That $\Theta_0 \cap z(\Theta) \neq \varnothing$, in turn, implies that, given the new schedules $q'_i(\cdot)$ and $t'_i(\cdot)$, there must still exist at least one type $\theta \in \Theta_0$ together with a type $\tilde{\theta} \in z(\theta)$

[49] Note that $\Theta_0 \cup g(\Theta_0)$ represents the set of types who are either willing to change contract, or receive the same contract as another type who is willing to change.

such that type θ is indifferent between the contract $(q_i'(\theta), t_i'(\theta))$ designed for him and the contract $(q_i'(\tilde{\theta}), t_i'(\tilde{\theta})) \neq (q_i'(\theta), t_i'(\theta))$ designed for type $\tilde{\theta}$. However, 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) guarantees that $\theta \notin z(\tilde{\theta})$. That is, if type θ is indifferent between the contract designed for him and the contract designed for type $\tilde{\theta}$, then it cannot be the case that type $\tilde{\theta}$ is also indifferent between the contract designed for him and that designed for type θ. Clearly, the same property also implies that for any $\theta'' \in z(\tilde{\theta})$, with $\theta'' \neq \theta$, then necessarily $\theta \notin z(\theta'')$. That is, if type θ is willing to swap contract with type $\tilde{\theta}$ and if, at the same time, type $\tilde{\theta}$ is willing to swap contract with type θ'', then it cannot be the case that type θ'' is also willing to swap contract with type θ. These properties, in turn, guarantee that the procedure described above 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) until 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 $\tilde{\mathcal{P}}$ in the main text and yield P_i a payoff $\overline{U}_i > U_i^*$, then irrespective of how one specifies the agent's continuation equilibrium σ_A^{M*}, P_i necessarily has a profitable deviation. This, in turn, proves that condition (iii) is necessary.

PROOF OF PROPOSITION 2:

Suppose that the principals collude so as to maximize their joint profits. In any mechanism that is individually rational and incentive compatible for the agent, the principals' joint profits are given by[50]

$$(14) \quad \int_{\underline{\theta}}^{\overline{\theta}} \Big\{ \theta[q_1(\theta) + q_2(\theta)] + \lambda q_1(\theta) q_2(\theta) - \frac{1}{2}[q_1(\theta)^2 + q_2(\theta)^2] - \frac{1-F(\theta)}{f(\theta)}[q_1(\theta) + q_2(\theta)] \Big\} dF(\theta) - \underline{U},$$

where $\underline{U} = \underline{\theta}[q_1(\underline{\theta}) + q_2(\underline{\theta})] + \lambda q_1(\underline{\theta}) q_2(\underline{\theta}) - t(\underline{\theta}) \geq 0$ denotes the equilibrium payoff of the lowest type. It is easy to see that, under the assumptions in the proposition, the schedules $(q_i(\cdot))_{i=1}^2$ that maximize (14) are those that maximize pointwise the integrand function and are given by $q_i(\theta) = q^c(\theta)$, all θ, $i = 1, 2$. The fact that these schedules can be sustained in a mechanism that is individually rational and incentive compatible for the agent, and that gives zero surplus to the lowest type, follows from the following properties:

- the agent's payoff $\theta(q_1 + q_2) + \lambda q_1 q_2$ is increasing in θ and satisfies the strict increasing-difference property in (θ, q_i), $i = 1, 2$; and
- the schedules $q_i(\cdot)$, $i = 1, 2$, are nondecreasing (see, e.g., Diego Garcia 2005).

Next, consider the result that the collusive schedules cannot be sustained by a noncooperative equilibrium in which the agent's strategy is Markovian. This result

[50] The result is standard and follows from the fact that the agent's payoff $\theta(q_1 + q_2) + \lambda q_1 q_2$ is equi-Lipschitz continuous and differentiable in θ (see, e.g., Milgrom and Segal 2002).

is established by contradiction. Suppose, on the contrary, that there exists a pair of tariffs $T_i : \mathcal{Q} \to \mathbb{R}$, $i = 1,2$, that sustain the collusive schedules as an equilibrium in which the agent's strategy is Markovian. Using the result in Proposition 1, this means that there exists a pair of nondecreasing functions $\tilde{q}_i : \Theta_i \to \mathcal{Q}$, $i = 1,2$, and a pair of scalars $\tilde{K}_i \geq 0$, $i = 1,2$, that satisfy conditions (i)–(iii) in Proposition 1, with $q_i^*(\cdot) = q^c(\cdot)$, $i = 1,2$. In particular, for any $\theta \in \Theta$, any $i = 1,2$, it must be that

$$(15)\quad V^*(\theta) = \sup_{\{(\theta_1,\theta_2)\in\Theta_1\times\Theta_2} \{\theta[\tilde{q}_1(\theta_1) + \tilde{q}_2(\theta_2)] + \lambda\tilde{q}_1(\theta_1)\tilde{q}_2(\theta_2) - \tilde{t}_1(\theta_1) - \tilde{t}_2(\theta_2)\}$$

$$= \sup_{\theta_i\in\Theta_i} \{\theta\tilde{q}_i(\theta_i) + v_i^*(\theta,\tilde{q}_i(\theta_i)) - \tilde{t}_i(\theta_i)\}$$

$$= \sup_{\theta_i\in[m_i(\underline{\theta}),m_i(\overline{\theta})]} \{\theta\tilde{q}_i(\theta_i) + v_i^*(\theta,\tilde{q}_i(\theta_i)) - \tilde{t}_i(\theta_i)\},$$

where the functions $\tilde{t}_i(\cdot)$ are the ones defined in (1) with $K_i = \tilde{K}_i$, $i = 1,2$, and where the function $V^*(\cdot)$ is the one defined in (3). Note that all equalities in (15) follow directly from the fact that the mechanisms $\phi_i^r = (\tilde{q}_i(\cdot),\tilde{t}_i(\cdot))$, $i = 1,2$, are incentive-compatible and satisfy conditions (i) and (ii) in Proposition 1.

Next, note that the property for any message $\theta_i \in [m_i(\underline{\theta}),m_i(\overline{\theta})]$, and any $\theta \in \Theta$, the marginal valuation $\theta + \lambda\tilde{q}_i(\theta_i) \in [m_j(\underline{\theta}),m_j(\overline{\theta})]$, combined with the property that the schedule $\tilde{q}_j(\cdot), j \neq i$, is continuous over $[m_j(\underline{\theta}),m_j(\overline{\theta})]$, implies that, given any $\theta_i \in [m_i(\underline{\theta}),m_i(\overline{\theta})]$, the agent's payoff

$$w_i(\theta;\theta_i) \equiv \theta\tilde{q}_i(\theta_i) + v_i^*(\theta,\tilde{q}_i(\theta_i)) - \tilde{t}_i(\theta_i)$$

$$= \theta\tilde{q}_i(\theta_i) + \int_{\min\Theta_j}^{\theta+\lambda\tilde{q}_i(\theta_i)} \tilde{q}_j(s)ds + \tilde{K}_j - \tilde{t}_i(\theta_i)$$

is M_i-Lipschitz continuous and differentiable in θ with derivative

$$\frac{\partial w_i(\theta;\theta_i)}{\partial\theta} = \tilde{q}_i(\theta_i) + \tilde{q}_j(\theta + \lambda\tilde{q}_i(\theta_i)) \leq 2\overline{Q} \equiv M_i.$$

Standard envelope theorem results (see, e.g., Milgrom and Segal 2002) then imply that the value function

$$W_i(\theta) \equiv \sup_{\theta_i\in[m_i(\underline{\theta}),m_i(\overline{\theta})]} \{\theta\tilde{q}_i(\theta_i) + v_i^*(\theta,\tilde{q}_i(\theta_i)) - \tilde{t}_i(\theta_i)\}$$

is Lipschitz continuous with derivative almost everywhere given by

$$(16)\quad \frac{\partial W_i(\theta)}{\partial\theta} = \tilde{q}_i(\theta_i^*) + \tilde{q}_j(\theta + \lambda\tilde{q}_i(\theta_i^*)) = q^c(m^{-1}(\theta_i^*)) + \tilde{q}_j(\theta + \lambda\tilde{q}_i(\theta_i^*)),$$

where $\theta_i^* \in \arg\max_{\theta_i \in [m_i(\underline{\theta}), m_i(\overline{\theta})]} \{\theta \tilde{q}_i(\theta_i) + v_i^*(\theta, \tilde{q}_i(\theta_i)) - \tilde{t}_i(\theta_i)\}$ is an arbitrary maximizer for type θ. The fact that the mechanisms $(\tilde{q}_i(\cdot), \tilde{t}_i(\cdot))$, $i = 1, 2$, satisfy conditions (i) and (ii) in Proposition 1, however, implies that

$$m(\theta) \in \arg\max_{\theta_i \in [m_i(\underline{\theta}), m_i(\overline{\theta})]} \{\theta \tilde{q}_i(\theta_i) + v_i^*(\theta, \tilde{q}_i(\theta_i)) - \tilde{t}_i(\theta_i)\}.$$

Using (16) and property (i) in Proposition 9, the agent's value function can then be rewritten as

$$(17) \qquad W_i(\theta) = \theta q^c(\theta) + v_i^*(\theta, q^c(\theta)) - \tilde{t}_i(m(\theta))$$
$$= \int_{\underline{\theta}}^{\theta} [q^c(s) + \tilde{q}_j(s + \lambda q^c(s))] ds + W_i(\underline{\theta}).$$

We thus conclude that the functions $\tilde{t}_i(\cdot)$ must satisfy

$$(18) \quad \tilde{t}_i(m(\theta)) = \theta q^c(\theta) + v_i^*(\theta, q^c(\theta)) - \int_{\underline{\theta}}^{\theta} [q^c(s) + \tilde{q}_j(s + \lambda q^c(s))] ds - W_i(\underline{\theta})$$
$$= \theta q^c(\theta) + [v_i^*(\theta, q^c(\theta)) - v_i^*(\theta, 0)]$$
$$- \int_{\underline{\theta}}^{\theta} [q^c(s) + \tilde{q}_j(s + \lambda q^c(s)) - \tilde{q}_j(s)] ds - W_i(\underline{\theta}) + \tilde{K}_{j.}$$

Note that the second equality follows from the fact that $v_i^*(\theta, 0) = \int_{\min \Theta_i}^{\theta} \tilde{q}_j(s) ds + \tilde{K}_j = \int_{\underline{\theta}}^{\theta} \tilde{q}_j(s) ds + \tilde{K}_j$. Also note that necessarily $B_i \equiv W_i(\underline{\theta}) - \tilde{K}_j \geq 0, i = 1, 2$; otherwise, given ϕ_1^r and ϕ_2^r, type $\underline{\theta}$ would be strictly better off participating only in principal P_j's mechanism, $j \neq i$. Using (18), principal i's equilibrium U_i^* can then be expressed as

$$U_i^* = \int_{\underline{\theta}}^{\overline{\theta}} h_i(q^c(\theta); \theta) dF(\theta) - B_i,$$

where $h_i(q; \theta)$ is the function defined in (6).

We are finally ready to establish the contradiction. Below, we show that, given $\phi_j = (\tilde{q}_j(\cdot), \tilde{t}_j(\cdot)), j \neq i$, the value $\overline{U}_i$ of program $\mathcal{P}$, as defined in the main text, is strictly higher than U_i^*. This contradicts the assumption made above that the pair of mechanisms $\phi_i = (\tilde{q}_i(\cdot), \tilde{t}_i(\cdot)), i = 1, 2$, satisfies condition (iii) of Proposition 1.

Take an arbitrary interval $[\theta', \theta''] \subset (\underline{\theta}, \overline{\theta})$ and, for any $\theta \in [\theta', \theta'']$, let $Q(\theta) \equiv [q^c(\theta) - \varepsilon, q^c(\theta) + \varepsilon]$, where $\varepsilon > 0$ is chosen, so that, for any $\theta \in [\theta', \theta'']$ and any $q \in Q(\theta), (\theta + \lambda \text{q}) \in [m(\underline{\theta}), m(\overline{\theta})]$. Note that, for any $\theta \in [\theta', \theta'']$, the function $h_i(\cdot; \theta)$ defined in (6) is continuously differentiable over $Q(\theta)$ with

$$\frac{\partial h_i(q^c(\theta);\theta)}{\partial q} = \theta + \lambda \tilde{q}_j(\theta + \lambda q^c(\theta)) - q^c(\theta) - \frac{1 - F(\theta)}{f(\theta)}[1 + \lambda \frac{\partial \tilde{q}_j(\theta + \lambda q^c(\theta))}{\partial \tilde{\theta}_j}]$$

$$= \theta - (1 - \lambda) q^c(\theta) - \frac{1 - F(\theta)}{f(\theta)} - \frac{1 - F(\theta)}{f(\theta)} \lambda \frac{\partial \tilde{q}_j(\theta + \lambda q^c(\theta))}{\partial \tilde{\theta}_j} < 0,$$

where the inequality follows from the definition of $q^c(\theta)$ and from the fact that $\tilde{q}_j(\cdot)$ is strictly increasing over $[m(\underline{\theta}), m(\overline{\theta})]$. The last result implies that there exists a nondecreasing schedule $q_i : \Theta \to \mathcal{Q}$ such that

$$\int_{\underline{\theta}}^{\overline{\theta}} h_i(q_i(\theta);\theta) dF(\theta) > \int_{\underline{\theta}}^{\overline{\theta}} h_i(q^c(\theta);\theta) dF(\theta), \tag{19}$$

and

$\theta + \lambda q_i(\hat{\theta}) \in [m(\underline{\theta}), m(\overline{\theta})]$ for all $(\theta, \hat{\theta}) \in \Theta^2$. Now, let $t_i : \Theta \to \mathbb{R}$ be the function that is obtained from $q_i(\cdot)$ using (5) and setting $K_i = 0$. That is, for any $\theta \in \Theta$,

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

It is easy to see that the pair of functions $q_i(\cdot), t_i(\cdot)$ constructed above satisfies all the IR constraints of program $\tilde{\mathcal{P}}$. To see that they also satisfy all the IC constraints, note that the agent's payoff under truthtelling is

$$X(\theta) \equiv \theta q_i(\theta) + [v_i^*(\theta, q_i(\theta)) - v_i^*(\theta, 0)] - t_i(\theta)$$

$$= \int_{\underline{\theta}}^{\theta} [q_i(s) + \tilde{q}_j(s + \lambda q_i(s)) - \tilde{q}_j(s)] ds,$$

whereas the payoff that type θ obtains by mimicking type $\hat{\theta}$ is

$$R(\theta;\hat{\theta}) \equiv \theta q_i(\hat{\theta}) + [v_i^*(\theta, q_i(\hat{\theta})) - v_i^*(\theta, 0)] - t_i(\hat{\theta})$$

$$= \theta q_i(\hat{\theta}) + \int_{\theta}^{\theta + \lambda q_i(\hat{\theta})} \tilde{q}_j(s) ds - t_i(\hat{\theta}).$$

Now, for any $(\theta, \hat{\theta}) \in \Theta^2$, let $\Phi(\theta;\hat{\theta}) \equiv X(\theta) - R(\theta;\hat{\theta})$. Note that, for any $\hat{\theta}$, $\Phi(\cdot;\hat{\theta})$ is Lipschitz continuous and its derivative, wherever it exists, satisfies

$$\frac{\partial \Phi(\theta;\hat{\theta})}{\partial \theta} = q_i(\theta) + \tilde{q}_j(\theta + \lambda q_i(\theta)) - [q_i(\hat{\theta}) + \tilde{q}_j(\theta + \lambda q_i(\hat{\theta}))].$$

Because $q_i(\cdot)$ and $\tilde{q}_j(\cdot)$ are both nondecreasing, we then have that, for all $\hat{\theta}$, a.e. θ, $(\partial \Phi(\theta;\hat{\theta})/\partial \theta)(\theta - \hat{\theta}) \geq 0$. Because, for any θ, $\Phi(\theta;\theta) = 0$, this implies that, for all $(\theta, \hat{\theta}) \in \Theta^2$, $\Phi(\theta; \hat{\theta}) = \int_{\hat{\theta}}^{\theta} (\partial\Phi(s; \hat{\theta})/\partial\theta) \geq 0$, which establishes that $q_i(\cdot), t_i(\cdot)$ is indeed incentive compatible.

Now, it is easy to see that principal i's payoff under $q_i(\cdot), t_i(\cdot)$ is

$$U_i = \int_{\underline{\theta}}^{\overline{\theta}} \left[t_i(\theta) - \frac{q_i(\theta)^2}{2}\right] dF(\theta) = \int_{\underline{\theta}}^{\overline{\theta}} h_i(q_i(\theta); \theta) dF(\theta),$$

which, by construction, is strictly higher than U_i^*. This, in turn, implies that, given the mechanism $\phi_j^r = (\tilde{q}_j(\cdot), \tilde{t}_j(\cdot))$, the value $\overline{U}_i$ of program $\tilde{\mathcal{P}}$ is necessarily higher than U_i^*. Hence, any pair of mechanisms $\phi_i = (\tilde{q}_i(\cdot), \tilde{t}_i(\cdot))$, $i = 1, 2$, that satisfy conditions (i) and (ii) in Proposition 1, necessarily fail to satisfy condition (iii). Because conditions (i)–(iii) are necessary, we conclude that there exists no equilibrium in which the agent's strategy is Markovian that sustains the collusive schedules.

PROOF OF PROPOSITION 3:

The result is established using Proposition 1. Below, we show that the pair of quantity schedules $\tilde{q}_i(\cdot) = \tilde{q}(\cdot)$, $i = 1, 2$, together with the pair of transfer schedules $\tilde{t}_i(\cdot) = \tilde{t}(\cdot)$, $i = 1, 2$, satisfies conditions (i) and (ii) in Proposition 1, where $\tilde{q} : [0, \overline{\theta} + \lambda\overline{Q}] \to \mathcal{Q}$ is the function defined in (8), and where $\tilde{t} : [0, \overline{\theta} + \lambda\overline{Q}] \to \mathbb{R}$ is the function defined by

$$\tilde{t}(s) = s\tilde{q}(s) - \int_0^s \tilde{q}(s)ds \qquad \forall s \in [0, \overline{\theta} + \lambda\overline{Q}].$$

That these schedules satisfy condition (i) is immediate. Thus, consider condition (ii). Fix $\phi_j^{r*} = (\tilde{q}_j(\cdot), \tilde{t}_j(\cdot))$. Note that, given any $q \in \mathcal{Q}$, the function $g_i(\cdot, q)$: $\Theta \to \mathbb{R}$ defined by

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

is Lipschitz continuous with derivative bounded uniformly over q, and satisfies the "convex-kink" condition of Assumption 1 in Jeffrey C. Ely (2001). This last property follows from the assumption that $\underline{\theta} + \lambda q^*(\underline{\theta}) \geq \overline{\theta}$. Combining Theorem 2 of Milgrom and Segal (2002) with Theorem 2 of Ely (2001), it is easy to verify that the schedules $q_i : \Theta \to \mathcal{Q}$ and $t_i : \Theta \to \mathbb{R}$ satisfy all the (IC) and (IR) constraints of program $\tilde{\mathcal{P}}$ if and only if $q_i(\cdot)$ is nondecreasing and $t_i(\cdot)$ satisfies

$$t_i(\theta) = \theta q_i(\theta) + [v_i^*(\theta, q_i(\theta)) - v_i^*(\theta, 0)] - \int_{\underline{\theta}}^{\theta} [q_i(s) + \tilde{q}(s + \lambda q_i(s)) - \tilde{q}(s)]ds - K_i' \tag{20}$$

for all $\theta \in \Theta$, with $K_i' \geq 0$.

Next, let $t^* : \Theta \to \mathbb{R}$ be the function that is obtained from (20), letting $q_i(\cdot) = q^*(\cdot)$ and setting $K_i' = 0$. Note that this function reduces to the one in (10) after a simple change in variable. The fact that $q_i(\cdot)$ and $t_i(\cdot)$ satisfy all the IC and IR constraints of program $\tilde{\mathcal{P}}$, together with the fact that the mechanism $\phi_j^{r*} = (\tilde{q}_j(\cdot), \tilde{t}_j(\cdot))$ is incentive compatible and individually rational for each $\theta_j \in \Theta_i$, in turn, implies 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 $(q_i, t_i) \in \{(q^*(\theta'), t^*(\theta')) : \theta' \in \Theta\} \cup (0,0)$, and $(q_j, t_j) \in \{(\tilde{q}(\theta_j), \tilde{t}(\theta_j)) : \theta_j \in \Theta_j\} \cup (0,0)$. But this also means that the schedules $q' : [m(\underline{\theta}), m(\overline{\theta})] \to \mathcal{Q}$ and $t' : [m(\underline{\theta}), m(\overline{\theta})] \to \mathbb{R}$ given by

$$q'(s) \;\equiv\; q^*(m^{-1}(s)) \text{ and } t'(s) \;\equiv\; t^*(m^{-1}(s))$$

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

$$t'(s) \;\equiv\; sq'(s) \;-\; \int_{m(\underline{\theta})}^{s} q'(x)dx.$$

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

$$t'(s) \;=\; \tilde{t}(s).$$

Furthermore, because for any $s > m(\overline{\theta}), (\tilde{q}(s), \tilde{t}(s)) = (\tilde{q}(m(\overline{\theta})), \tilde{t}(m(\overline{\theta}))) = (q'(\overline{\theta}), t'(\overline{\theta}))$, it immediately follows from the aforementioned results that, when both principals offer the mechanism $\phi_i^{r*} = (\tilde{q}_i(\cdot), \tilde{t}_i(\cdot)), i = 1,2$, each type θ finds it optimal to participate in both mechanisms and report $s = m(\theta)$ to each principal. Note that, in so doing, each type θ obtains the equilibrium quantity $q^*(\theta)$ and pays the equilibrium price $\tilde{t}(m(\theta)) = t^*(\theta)$ to each principal.

We have thus established that the pair of mechanisms $\phi_i^{r*} = (\tilde{q}_i(\cdot), \tilde{t}_i(\cdot)), i = 1,2$, satisfies conditions (i) and (ii) in Proposition 1. To complete the proof, it remains to show that they also satisfy condition (iii). For this purpose, recall that, given $\phi_j^{r*} = (\tilde{q}_j(\cdot), \tilde{t}_j(\cdot))$, a pair of schedules $q_i : \Theta \to \mathcal{Q}$ and $t_i : \Theta \to \mathbb{R}$ satisfies the (IC) and (IR) constraints of program $\tilde{\mathcal{P}}$ if and only if the function $q_i(\cdot)$ is nondecreasing and the function $t_i(\cdot)$ is as in (20). This, in turn, means that the value of program $\tilde{\mathcal{P}}$ coincides with the value of program $\tilde{\mathcal{P}}^{new}$, as defined in the main text. Now, note that for any $\theta \in \text{int}(\Theta)$, the function $h(\cdot;\theta)$: $\mathcal{Q} \to \mathbb{R}$ is maximized at $q = q^*(\theta)$. To see this, note that the fact that $q^*(\cdot)$ solves the differential equation in (7) implies that the function $h(\cdot;\theta)$ is differentiable at $q = q^*(\theta)$ with derivative

(21)
$$\frac{\partial h(q^*(\theta);\theta)}{\partial q} = \theta + \lambda\tilde{q}(\theta + \lambda q^*(\theta)) - q^*(\theta) - \frac{1 - F(\theta)}{f(\theta)}\left[1 + \lambda\frac{\partial\tilde{q}(\theta + \lambda q^*(\theta))}{\partial\theta_i}\right] = 0.$$

Together with the fact that $h(\cdot;\theta)$ is quasiconcave, this property implies that $h(q;\theta)$ is maximized at $q = q^*(\theta)$. This implies that the solution to the program $\tilde{\mathcal{P}}^{new}$ is the function $q^*(\cdot)$ along with $K_i = 0$. However, by construction, the payoff U_i^* that principal P_i obtains in equilibrium by offering the mechanism ϕ_i^{r*} is

$$U_i^* = \int_{\underline{\theta}}^{\overline{\theta}}\left[\tilde{t}(m(\theta)) - \frac{\tilde{q}(m(\theta))^2}{2}\right]dF(\theta) = \int_{\underline{\theta}}^{\overline{\theta}}\left[t^*(\theta) - \frac{q^*(\theta)^2}{2}\right]dF(\theta)$$
$$= \int_{\underline{\theta}}^{\overline{\theta}} h(q^*(\theta);\theta)dF(\theta) = \overline{U}_i,$$

where $\overline{U}_i$ is the value of program $\tilde{\mathcal{P}}^{new}$ (and hence of program $\tilde{\mathcal{P}}$ as well). We thus conclude that the pair of mechanisms $\phi_i^{r*} = (\tilde{q}_i(\cdot), \tilde{t}_i(\cdot))$, $i = 1, 2$, satisfies condition (iii), which completes the proof.

PROOF OF PROPOSITION 4:

Consider the "only if" part of the result. Starting from any pure-strategy equilibrium σ^M of Γ^M, one can construct another pure-strategy equilibrium $\hat{\sigma}^M$ that sustains the same SCF π, but in which the agent's strategy $\hat{\sigma}_A^M$ satisfies the following property. Given any $i \in \mathcal{N}$, any menu ϕ_i^M, and any action profile (e, a_{-i}), there exists a unique action $a_i(e, a_{-i}; \phi_i^M) \in \mathcal{A}_i$, such that the agent always chooses a contract δ_i from ϕ_i^M which responds to effort e with the action $a_i(e, a_{-i}; \phi_i^M)$, when the contracts the agent selects with the other principals respond to the same effort choice with the actions a_{-i}. The proof for this step follows from arguments similar to those that establish Theorem 3. Given $\hat{\sigma}^M$, it is then easy to construct a pure-strategy truthful equilibrium $\mathring{\sigma}^*$ of $\mathring{\Gamma}^r$ that sustains the same SCF. The proof for this step follows from arguments similar to those that establish Theorem 2. The only delicate part is in specifying how the agent reacts off-equilibrium to a revelation mechanism $\mathring{\phi}_i^r \neq \mathring{\phi}_i^{r*}$. In the proof of Theorem 2, it was assumed that the agent responds to an off-equilibrium mechanism $\phi_i^r \neq \phi_i^{r*}$ as if the game were Γ^M and P_i offered the menu whose image is $\text{Im}(\phi_i^M) = \text{Im}(\phi_i^r)$. However, in the new revelation game $\mathring{\Gamma}^r$, the image $\text{Im}(\mathring{\phi}_i^r)$ of a direct revelation mechanism $\mathring{\phi}_i^r$ is a subset of $\mathcal{A}_i$ as opposed to a menu of contracts. This, nonetheless, does not pose any problem. It suffices to proceed as follows. Given any direct mechanism $\mathring{\phi}_i^r$, and any effort choice e, let $\mathcal{A}_i(e;\phi_i^r) \equiv \{a_i : a_i = \phi_i^r(e, a_{-i}),\ a_{-i} \in \mathcal{A}_{-i}\}$ denote the set of responses to effort choice e that the agent can induce in $\mathring{\phi}_i^r$ by reporting different messages $a_{-i} \in \mathcal{A}_{-i}$. Given any mechanism $\mathring{\phi}_i^r$, then let $\phi_i^M = \chi(\mathring{\phi}_i^r)$ denote the menu of contracts whose image is $\text{Im}(\phi_i^M) = \{\delta_i \in \mathcal{D}_i : \delta_i(e) \in \mathcal{A}_i(e;\mathring{\phi}_i^r) \text{ all } e \in E\}$. Clearly, for any (e, a_{-i}), the maximum payoff that the agent can guarantee himself in Γ^M, given the menu ϕ_i^M, is the same as in $\mathring{\Gamma}^r$ given $\mathring{\phi}_i^r$. The rest of the proofs then parallels that of Theorem 2, by having the agent react to any mechanism $\mathring{\phi}_i^r \neq \mathring{\phi}_i^{r*}$ as if the game were Γ^M and P_i offered the menu $\phi_i^M = \chi(\mathring{\phi}_i^r)$.

Next, consider the "if" part of the result. The proof parallels that of part (ii) of Theorem 2 using the mapping $\chi : \mathring{\Phi}_i^r \to \Phi_i^M$ defined above to construct the equilibrium menus, and the mapping $\varphi : \Phi_i^M \to \mathring{\Phi}_i^r$ defined below to construct the agent's reaction to any off-equilibrium menu $\phi_i^M \neq \phi_i^{M*}$. Let $\varphi : \Phi_i^M \to \mathring{\Phi}_i^r$ be any arbitrary function that maps each menu ϕ_i^M into a direct mechanism $\mathring{\phi}_i^r = \varphi(\phi_i^M)$ with the following property

$$\mathring{\phi}_i^r(e, a_{-i}) \in \underset{a_i \in \{\hat{a}_i : \hat{a}_i = \delta_i(e), \delta_i \in \mathrm{Im}(\phi_i^M)\}}{\arg\max} v(e, a_i, a_{-i}) \quad \forall (e, a_{-i}) \in E \times \mathcal{A}_{-i}.$$

The agent's reaction to any menu $\phi_i^M \neq \phi_i^{M*}$ is then the same as if the game were $\mathring{\Gamma}^r$ and P_i offered the direct mechanism $\mathring{\phi}_i^r = \varphi(\phi_i^M)$. The rest of the proof is based on the same arguments as in the proof of part (ii) of Theorem 2 and is omitted for brevity.

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 by an equilibrium of $\hat{\Gamma}^r$ can also be sustained by 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 contract that A takes in equilibrium with P_i when his type is θ.

As a preliminary step, we establish the following result.

LEMMA 1: *Suppose the SCF π can be sustained by a pure-strategy equilibrium of Γ^M. Then it can also be sustained by 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 (i) P_k deviates from the equilibrium menu, (ii) the agent's type is θ, (iii) the lottery over contracts A selects with P_k is δ_k, and (iv) any principal P_i, $i \neq k$, offers 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 (δ, θ), let $\underline{U}_k(\delta, \theta)$ denote the lowest payoff that the agent can inflict to principal P_k, without violating his rationality. This payoff is given by

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

where, for any $y \in Y$,

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

with

$$E^*(y,\theta) \equiv \arg\max_{e\in E} \left\{ \int_{\mathcal{A}} v(a,e,\theta) dy_1(e) \times \cdots \times dy_n(e) \right\}.$$

Next, for any $(\theta,\delta_k) \in \Theta \times \mathcal{D}_k$, let

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

denote the set of lotteries in the menus $\tilde{\phi}^M_{-k}$ that are optimal for the agent, given (θ,δ_k), where $\mathrm{Im}(\tilde{\phi}^M_{-k}) \equiv \times_{j\neq k} \mathrm{Im}(\tilde{\phi}^M_j)$. Then, for any $(\theta,\delta_k) \in \Theta \times \mathcal{D}_k$, let $\delta_{-k}(\theta,\delta_k) \in \mathcal{D}_{-k}$ be any profile of lotteries such that

$$\delta_{-k}(\theta,\delta_k) \in \arg\min_{\delta'_{-k}\in D_{-k}(\theta,\delta_k;\tilde{\phi}^M_{-k})} \underline{U}_k(\delta_k,\delta'_{-k},\theta). \tag{23}$$

Now, consider the following pure-strategy profile $\mathring{\sigma}^M$. For any $i \in \mathcal{N}$, $\mathring{\sigma}^M_i$ is the pure strategy that prescribes that P_i offers the same menu $\tilde{\phi}^M_i$ as under $\tilde{\sigma}^M$. The continuation equilibrium $\mathring{\sigma}^M_A$ is such that, when either $\phi^M_i = \tilde{\phi}^M_i$ for all i, or $|\{i \in \mathcal{N} : \phi^M_i \neq \tilde{\phi}^M_i\}| > 1$, then $\mathring{\sigma}^M_A(\theta,\phi^M) = \tilde{\sigma}^M_A(\theta,\phi^M)$, for any θ. When, instead, ϕ^M is such that $\phi^M_i = \tilde{\phi}^M_i$ for all $i \neq k$, while $\phi^M_k \neq \tilde{\phi}^M_k$ for some $k \in \mathcal{N}$, then each type θ selects the profile of lotteries (δ_k,δ_{-k}) defined as follows. δ_k is the same lottery that type θ would have selected with P_k according to the original strategy $\tilde{\sigma}^M_A$, given the menus $(\tilde{\phi}^M_{-k},\phi^M_k)$; $\delta_{-k} = \delta_{-k}(\theta,\delta_k)$ is the profile of lotteries defined in (23). 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 (22). It is immediate that the behavior prescribed by the strategy $\mathring{\sigma}^M_A$ is sequentially rational for the agent. Furthermore, given $\mathring{\sigma}^M_A$, a principal P_i who expects all other principals to offer the equilibrium menus $\tilde{\phi}^M_{-i}$ cannot do better than offering the equilibrium menu $\tilde{\phi}^M_i$. We conclude that $\mathring{\sigma}^M$ is a pure-strategy equilibrium of Γ^M and 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 lottery that A selects with P_i when his type is θ, the contract selected with P_k is δ_k, and the menus offered are $\phi^M_j = \phi^{M*}_j$ for all $j \neq k$, and $\phi^M_k \neq \phi^{M*}_k$.

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

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

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

• 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*}.

• 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 lottery that type θ would have selected 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 with 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 (22).

• Finally, for any profile of mechanisms $\hat{\phi}^r$ such that $|\{i \in \mathcal{N} : \hat{\phi}_i^r \neq \hat{\phi}_i^{r*}\}| > 1$, simply let $\hat{\sigma}_A^r(\theta, \phi^r)$ be any strategy that is sequentially rational 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 follows from the same arguments as in the proof of Part 2 in Theorem 2. It suffices to replace the mappings $\Phi_i^r(\cdot)$ with the mappings $\hat{\Phi}_i^r(\cdot)$ and then make the following adjustment to Case 2. For any profile of menus ϕ^M for which there exists a $j \in \mathcal{N}$ such that $\hat{\Phi}_i^r(\phi_i^M) \neq \varnothing$ for all $i \neq j$, and $\hat{\Phi}_j^r(\phi_j^M) = \varnothing$, let $\hat{\phi}_j^r$ be any arbitrary revelation mechanism, such that

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

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

PROOF OF THEOREM 6:

The proof is in two parts. 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 $\mathcal{Q}_i$. Now, consider a partition-game $\Gamma^{\mathcal{Q}}$ in which first each principal P_i chooses an element of $\mathcal{Q}_i$. After observing the collection of cells $Q = (Q_i)_{i=1}^n$, the agent then selects a profile of menus $\phi^M = (\phi_1^M, \dots, \phi_n^M)$, one from each cell Q_i, then chooses the lotteries over contracts δ, and finally, given the contracts y selected by the lotteries δ, chooses effort $e \in E$.

The proof of Part 1 is in two steps. Step 1 identifies a collection of partitions $\mathcal{Q}^Z = (\mathcal{Q}_i^Z)_{i\in\mathcal{N}}$ such that the agent's payoff is the same for any pair of menus $\phi_i^M, \phi_i^{M'} \in Q_i^Z$, $i = 1, \dots, n$. It then shows that, for any $\sigma^M \in \mathcal{E}(\Gamma^M)$ there exists a $\mathring{\sigma} \in \mathcal{E}(\Gamma^{\mathcal{Q}^Z})$ that implements the same outcomes. Step 2 uses the equilibrium $\mathring{\sigma}$ of $\Gamma^{\mathcal{Q}^Z}$ 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, $i = 1, \dots, n$ with $\mathcal{Q}_i$ consisting of measurable sets.[51] 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 $\mathcal{Q}_i$ the union of whose elements is measurable,

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

Next, consider the agent. For any $Q = (Q_1, \dots, Q_n) \in \times_{i\in\mathcal{N}}\mathcal{Q}_i$, A selects the menus ϕ^M from $\times_{i\in\mathcal{N}} Q_i$ 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 regular conditional distribution over Φ_i^M that is obtained from the equilibrium strategy σ_i^M of P_i conditioning on $\phi_i^M \in Q_i$.[52] 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, irrespective of 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 $\mathcal{Q}^Z$ that make $\mathring{\sigma}_A$ sequentially rational for the agent. Consider the equivalence relation $\sim_i$ defined as follows: given any two menus ϕ_i^M and $\phi_i^{M'}$,

$$\phi_i^M \sim_i \phi_i^{M'} \Leftrightarrow Z_\theta(\delta_{-i}; \phi_i^M) = Z_\theta(\delta_{-i}; \phi_i^{M'}) \quad \forall(\theta, \delta_{-i}),$$

[51] 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.

[52] 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, 345).

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

Now, let $\mathcal{Q}^Z = (\mathcal{Q}_i^Z)_{i\in\mathcal{N}}$ be the collection of partitions generated by the equivalence relations $\sim_i$, $i = 1,\ldots,n$. It follows immediately that, in the partition game $\Gamma^{\mathcal{Q}^Z}$, $\mathring{\sigma}_A$ is sequentially rational for A. We conclude that for any $\sigma^M \in \mathcal{E}(\Gamma^M)$ there exists a $\mathring{\sigma} \in \mathcal{E}(\Gamma^{\mathcal{Q}^Z})$ 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 sustains the same outcomes as σ^M in Γ^M. To simplify the notation, hereafter, we drop the superscrips Z from the partitions $\mathcal{Q}$, with the understanding that $\mathcal{Q}$ refers to the collection of partitions generated by the equivalence relations $\sim_i$ defined above. For any $i \in \mathcal{N}$, any $Q_i \in \mathcal{Q}_i$, and any $(\theta, \delta_{-i}) \in \Theta \times \mathcal{D}_{-i}$, then let $Z_\theta(\delta_{-i};Q_i) \equiv Z_\theta(\delta_{-i};\phi_i^M)$ for some $\phi_i^M \in Q_i$. Since for any two menus ϕ_i^M, $\phi_i^{M'} \in Q_i$, $Z_\theta(\delta_{-i};\phi_i^M) = Z_\theta(\delta_{-i};\phi_i^{M'})$ for all (θ,δ_{-i}), then $Z_\theta(\delta_{-i};Q_i)$ is uniquely determined by Q_i. Now, for any $Q_i \in \mathcal{Q}_i$, let $\tilde{\phi}_i^r|_{Q_i} \in \tilde{\Phi}_i^r$ denote the revelation mechanism given by

$$\tilde{\phi}_i^r(\theta,\delta_{-i}) = Z_\theta(\delta_{-i};Q_i) \quad \forall(\theta,\delta_{-i}) \in \Theta \times \mathcal{D}_{-i}. \tag{24}$$

For any set of mechanisms $B \subseteq \tilde{\Phi}_i^r$, then let $\mathcal{Q}_i(B) \equiv \{Q_i \in \mathcal{Q}_i : \tilde{\phi}_i^r|_{Q_i^Z} \in B\}$ 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

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

Next, consider the agent. Given any profile of mechanisms $\tilde{\phi}^r \in \tilde{\Phi}^r$, let $Q(\tilde{\phi}^r) = (Q_i(\tilde{\phi}_i^r))_{i\in\mathcal{N}} \in \times_{i\in\mathcal{N}}\mathcal{Q}_i$ 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 $Z_\theta(\delta_{-i};Q_i(\tilde{\phi}_i^r)) = \tilde{\phi}_i^r(\delta_{-i},\theta)$ 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 $Q(\phi^r)$. That is, for any $(\theta,\tilde{\phi}^r) \in \Theta \times \tilde{\Phi}^r$,

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

$$\times \cdots \times d\sigma_n^M(\phi_n^M | Q_n(\tilde{\phi}_n^r)).$$

The strategy $\tilde{\sigma}_A^r$ is clearly sequentially rational 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

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

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 $\text{Im}(\tilde{\phi}_i^r) = \text{Im}(\phi_i^M)$.

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

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

Finally, for any ϕ^M such that $|\{j \in \mathcal{N} : \phi_j^M \notin \alpha_j(\tilde{\Phi}_j^r)\}| > 1$, simply let $\sigma_A^M(\theta, \phi^M)$ be any sequentially rational response for the agent given (θ, ϕ^M). It immediately follows 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 easy 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 $\tilde{\Gamma}^r$.

REFERENCES

Attar, Andrea, Gwenael Piaser, and Nicolàs Porteiro. 2007a. "A Note on Common Agency Models of Moral Hazard." *Economics Letters*, 95(2): 278–84.

Attar, Andrea, Gwenael Piaser, and Nicolàs Porteiro. 2007b. "Negotiation and Take-It or Leave-It in Common Agency with Non-Contractible Actions." *Journal of Economic Theory*, 135(1): 590–93.

Attar, Andrea, Dipjyoti Majumdar, Gwenael Piaser, and Nicolàs Porteiro. 2008. "Common Agency Games: Indifference and Separable Preferences." *Mathematical Social Sciences*, 56(1): 75–95.

Bernheim, B. Douglas, and Michael D. Whinston. 1985. "Common Marketing Agency as a Device for Facilitating Collusion." *RAND Journal of Economics*, 16(2): 269–81.

Bernheim, B. Douglas, and Michael D. Whinston. 1986a. "Common Agency." *Econometrica*, 54(4): 923–42.

Bernheim, B. Douglas, and Michael D. Whinston. 1986b. "Menu Auctions, Resource Allocation, and Economic Influence." *Quarterly Journal of Economics*, 101(1): 1–31.

Biais, Bruno, David Martimort, and Jean-Charles Rochet. 2000. "Competing Mechanisms in a Common Value Environment." *Econometrica*, 68(4): 799–837.

Calzolari, Giacomo. 2004. "Incentive Regulation of Multinational Enterprises." *International Economic Review*, 45(1): 257–82.

Calzolari, Giacomo and Vincenzo Devicoló. 2009. "Competition with Exclusive Contracts and Market Share Discounts." Center for Economic and Policy Research (CEPR) Discussion Paper 7613.

Chiesa, Gabriella, and Vincenzo Denicolo. 2009. "Trading with a Common Agent under Complete Information: A Characterization of Nash Equilibria." *Journal of Economic Theory*, 144(1): 296–311.

Dixit, Avinash, Gene M. Grossman, and Elhanan Helpman. 1997. "Common Agency and Coordination: General Theory and Application to Government Policy Making." *Journal of Political Economy*, 105(4): 752–69.

Dudley, R. M. 2002. *Real Analysis and Probability*. Cambridge, UK: Cambridge University Press.

Ely, Jeffrey C. 2001. "Revenue Equivalence without Differentiability Assumptions." Unpublished.

Epstein, Larry G., and Michael Peters. 1999. "A Revelation Principle for Competing Mechanisms." *Journal of Economic Theory*, 88(1): 119–60.

Garcia, Diego. 2005. "Monotonicity in Direct Revelation Mechanisms." *Economics Letters*, 88(1): 21–26.

Gibbard, Allan. 1973. "Manipulation of Voting Schemes: A General Result." *Econometrica*, 41(4): 587–601.

Green, Jerry, and Jean-Jacques Laffont. 1977. "Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods." *Econometrica*, 45(2): 427–38.

Guesnerie, Roger. 1995. *A Contribution to the Pure Theory of Taxation*. Cambridge, UK: Cambridge University Press.

Han, Seungjin. 2006. "Menu Theorems for Bilateral Contracting." *Journal of Economic Theory*, 131(1): 157–78.

Katz, Michael L. 1991. "Game-Playing Agents: Unobservable Contracts as Precommitments." *RAND Journal of Economics*, 22(3): 307–28.

Martimort, David. 1992. "Multi-Principaux Avec Anti-Selection." [Multiprincipals under adverse selection.] *Annales d'Economie et de Statistique*, 1992(28): 1–37.

Martimort, David, and Lars Stole. 1997. "Communication Spaces, Equilibria Sets and the Revelation Principle under Common Agency." Unpublished.

Martimort, David, and Lars Stole. 2002. "The Revelation and Delegation Principles in Common Agency Games." *Econometrica*, 70(4): 1659–73.

Martimort, David, and Lars Stole. 2003. "Contractual Externalities and Common Agency Equilibria." *Advances in Theoretical Economics*, 3(1): Berkeley Electronic Press.

Martimort, David, and Lars Stole. 2005. "Public Contracting in Delegated Agency Games" Unpublished.

McAfee, R. Preston. 1993. "Mechanism Design by Competing Sellers." *Econometrica*, 61(6): 1281–312.

Milgrom, Paul, and Ilya Segal. 2002. "Envelope Theorems for Arbitrary Choice Sets." *Econometrica*, 70(2): 583–601.

Myerson, Roger B. 1979. "Incentive Compatibility and the Bargaining Problem." *Econometrica*, 47(1): 61–73.

Parlour, Christine A., and Uday Rajan. 2001. "Competition in Loan Contracts." *American Economic Review*, 91(5): 1311–28.

Pavan, Alessandro, and Giacomo Calzolari. 2009. "Sequential Contracting with Multiple Principals." *Journal of Economic Theory*, 144(2): 503–31.

Peck, James. 1997. "A Note on Competing Mechanisms and the Revelation Principle." Unpublished.

Peters, Michael. 2001. "Common Agency and the Revelation Principle." *Econometrica*, 69(5): 1349–72.

Peters, Michael. 2003. "Negotiation and Take It or Leave It in Common Agency." *Journal of Economic Theory*, 111(1): 88–109.

Peters, Michael. 2007. "Erratum to 'Negotiation and Take It or Leave It in Common Agency.'" *Journal of Economic Theory*, 135(1): 594–95.

Peters, Michael, and Balázs Szentes. 2008. "Definable and Contractible Contracts." Unpublished.

Peters, Michael, and Christian Troncoso Valverde. 2009. "A Folk Theorem for Competing Mechanisms." Unpublished.

Piaser, Gwenael. 2007. "Direct Mechanisms, Menus and Latent Contracts." Unpublished.

Rochet, Jean-Charles. 1986, "Le Controle Des Equations Aux Derivees Partielles Issues de la Theorie Des Incitations." PhD thesis. Université Paris IX.

Segal, Ilya, and Michael D. Whinston. 2003. "Robust Predictions for Bilateral Contracting with Externalities." *Econometrica*, 71(3): 757–91.

Yamashita, Takuro. 2007. "A Revelation Principle and A Folk Theorem without Repetition in Games with Multiple Principals and Agents." Unpublished.