



Sequential contracting with multiple principals [☆]

Alessandro Pavan ^{a,*}, Giacomo Calzolari ^b

^a Department of Economics, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208-2600, United States

^b Department of Economics, University of Bologna, Piazza Scaravilli 2, 40125 Bologna, Italy

Received 7 April 2007; final version received 9 July 2008; accepted 16 July 2008

Available online 12 November 2008

Abstract

This paper considers dynamic games in which multiple principals contract sequentially and non-cooperatively with the same agent. We first show that when contracting is private, i.e. when downstream principals do not observe the mechanisms offered upstream and the decisions taken in these mechanisms, all PBE outcomes can be characterized through pure-strategy profiles in which the principals offer menus of contracts and delegate to the agent the choice of the contractual terms. We then show that, in most cases of interest for applications, the characterization of the equilibrium outcomes is further facilitated by the fact that the principals can be restricted to offer incentive-compatible extended direct mechanisms in which the agent reports the endogenous payoff-relevant decisions contracted upstream in addition to his exogenous private information. Finally we show how the aforementioned results must be adjusted to accommodate alternative assumptions about the observability of upstream histories and/or the timing of contracting examined in the literature.

© 2008 Elsevier Inc. All rights reserved.

JEL classification: D89; C72

Keywords: Sequential common agency; Mechanism design; Contracts; Endogenous types

[☆] This paper builds on a previous working paper that circulated under the title “A Markovian Revelation Principle for Common Agency Games.”

* Corresponding author. Fax: +1 847 491 7001.

E-mail addresses: alepavan@northwestern.edu (A. Pavan), giacomo.calzolari@unibo.it (G. Calzolari).

1. Introduction

There are many environments in which multiple principals contract sequentially and non-cooperatively with the same agent.¹ In politics, for example, a ruling administration (upstream principal) that signs a procurement contract with a defense supplier (agent) expects its counterpart to contract also with the next appointed administration (downstream principal). In organizations, an employer who hires a worker anticipates that the latter will leave after a while and be hired by other employers. In a regulatory environment, a multinational firm typically contracts first with domestic authorities (upstream principals) and then with foreign ones. In commerce, a seller (upstream principal) who sets up a menu of contract offers usually expects her buyers (agents) to purchase complementary products and services also from other vendors. In corporate finance, a venture capitalist (upstream principal) who offers a contract to an entrepreneur (agent) anticipates that the latter will borrow also from other lenders and then contract with suppliers, retailers and, perhaps, government agencies (downstream principals).²

Characterizing equilibrium outcomes in environments with competing principals is known to be difficult. Contrary to environments with a single principal, the agent cannot be assumed to select a contract by simply reporting his “type” i.e. his exogenous private information. First, the agent may have private information also about the mechanisms offered by other principals as well as, in the case of sequential contracting, the decisions taken in these mechanisms. Second, a principal may need to include in her mechanism also contracts that are selected by the agent only off-equilibrium to punish deviations by other principals.³

Despite a fast growing literature on sequential common agency, no general characterization results have been established for these games. This is what we aim to do in this paper by providing results that facilitate the construction of equilibrium outcomes in environments in which contracting is sequential.

We build our analysis on a benchmark model of *private contracting* in which principals do not observe other principals’ mechanisms, nor the decisions taken in these mechanisms. We also assume that the sequence of bilateral relationships is exogenous in that the agent cannot choose which principal to contract with at each date. Finally, we assume that some irreversible decisions are committed at each period.

Later in the paper, we discuss the intricacies that arise under alternative assumptions about the sequence of contracting and/or the observability of upstream histories and show how our results must be adapted to accommodate various alternative extensive forms examined in the literature.

Our first result shows that, when contracting is private, any social choice function⁴ that can be sustained with any ad-hoc strategy space for the principals can also be sustained by restricting the principals to offer menus of contracts and delegating to the agent the choice of the contractual terms. We also show that any equilibrium outcome in the game in which the principals can offer all possible menus is robust in the sense that it remains an equilibrium outcome in any game in

¹ In what follows, we refer to a principal (equivalently, a mechanism designer) as the party who offers the contract. We also adopt the convention of using masculine pronouns for the agent and feminine pronouns for the principals.

² Models of sequential contracting have been used in political economy by [5,13,29]; in sequential trade, by [1,2,8,20,31]; in regulation, by [3,12,16,17]; in labor relationships, by [15] and [23]; in financial contracting and insurance, by [7,25], and [32].

³ See [11,18,19,22,26–28] for a discussion of the problems with standard direct revelation mechanisms under simultaneous common agency and for possible solutions.

⁴ In the jargon of mechanism design and implementation theory, a social choice function (SCF) or, equivalently, an outcome function, is a mapping from states (here the agent’s exogenous private information) to decisions.

which the principals' strategy space is enlarged. These results show that the *menu theorems* of simultaneous common agency (cf. [28], and [19]) extend to sequential common agency, provided contracting is private.

Next, we prove the following two theorems, which are specific to sequential common agency. First, when lotteries are feasible, then all social choice functions can be sustained through *pure-strategy* profiles, i.e. by restricting the principals not to mix over their menus and the agent not to mix over the choice of a contract. Second, when information is complete (i.e. when the agent's type is common knowledge as typically assumed in menu auctions and moral hazard settings), then all deterministic social choice functions can be sustained through Markov-perfect equilibria, i.e. by restricting the agent's strategy to depend on any given history only through its payoff-relevant component. Because deterministic social choice functions naturally arise when one restricts attention to pure-strategy profiles and deterministic contracts, these results provide a possible justification for focusing on Markov-perfect equilibria in certain applications.⁵

Based on the aforementioned results, we then show how the characterization of all Markov-perfect equilibrium outcomes can be further simplified by restricting the principals to offer menus that can be conveniently described as *extended direct mechanisms*. In these mechanisms, the agent reports his *extended type*, i.e. his exogenous private information along with the endogenous payoff-relevant decisions contracted upstream. An extended direct mechanism is thus a collection of contracts, one for each extended type.

Describing a menu as a direct mechanism has proved useful in games with a single mechanism designer. The same can be done in sequential common agency by extending the notion of type to account for the fact that the agent's payoff-relevant information also includes endogenous decisions. However, there are two important differences with respect to games with a single principal. First, incentive-compatibility is endogenous: the agent's incentives to reveal his extended type depend on the mechanisms offered downstream. Incentive-compatibility must thus be established by backward induction. Second, an incentive-compatible mechanism must specify contracts also for extended types that have zero measure on the equilibrium path. This is because a mechanism must include also contracts that are used only off-equilibrium to punish deviations by upstream principals. In an extended direct mechanism, such out-of-equilibrium contracts are simply those designed for out-of-equilibrium extended types.

Another important difference with respect to the single-principal case is that, when the agent's strategy is not Markov, certain social choice functions cannot be sustained with extended direct mechanisms. With non-Markov strategies, a principal may need to give the same extended type of the agent a menu of contracts to choose from. This is because she may need the agent to punish deviations by upstream principals that altered the distribution of upstream payoff-relevant decisions but nevertheless led to equilibrium extended types. We find such a possibility intriguing from a theoretical viewpoint, but not particularly significant for the type of applications mentioned above.

Below, we conclude the introduction with a road map for our results and the related literature. The rest of the paper is then organized as follows. Section 2 describes the benchmark model of private contracting. Section 3 illustrates with an example why standard direct revelation mecha-

⁵ Note that a SCF can be deterministic even in the presence of stochastic outcomes. For example, in a moral hazard setting, a principal's decision coincides with the choice of a payment scheme that specifies a reward for the agent as a function of some (typically stochastic) performance measures. In these models, that a SCF is deterministic simply means that (a) the agent does not mix over his choice of effort and (b) the principals do not mix over their choice of a reward scheme. The final outcome, i.e. the payment to the agent, may however be stochastic.

nisms may fail to sustain certain outcomes and introduces the solution offered by our extended direct mechanisms. Section 4 contains the results for menus, while Section 5 contains the results for extended direct mechanisms and shows how to put these mechanisms to work in applications. Section 6 examines alternative extensive forms. Section 7 concludes. All proofs are either in Appendix A or in the Supplementary material.

1.1. A road map for the results and the related literature

Contrary to simultaneous common agency, when contracting is sequential there is no single model that fits all applications: sequential common agency can in fact be characterized by many different extensive forms depending on the assumptions one makes about the observability of upstream histories and the timing of the relevant decisions. Although it is difficult to have a single unified framework, it is however possible to have results that help characterize equilibria in various dynamic settings considered in applications. In what follows, we present a “road map” for our results and the related literature. This road map is organized around two dimensions that play a key role in most sequential contracting models: *history observability* and *timing*. Each of these two dimensions is associated with specific issues/difficulties with equilibrium characterization. Furthermore, different combinations of these two dimensions correspond to different environments considered in the literature.

Observability of upstream histories. First, consider an environment in which downstream principals observe upstream mechanisms, as for example in [1,5,17,21,31]. Contrary to private contracting, in these environments, restricting the principals to offer menus can be with loss of generality. The reason is that payoff-irrelevant details of the mechanisms can be used as correlation devices for the principals’ decisions. In the absence of alternative correlation devices, restricting the principals to offer menus then precludes the sustainability of certain outcomes. However, as discussed in Section 6, if one is interested only in pure-strategy profiles (the case considered in all applications that assume mechanism observability), then, not only can one restrict the principals to compete in menus, it is actually safe to restrict them to offer menus that can be described as extended direct mechanisms. This is so irrespective of whether the agent’s strategy is Markov. In fact, because the principals can observe the upstream mechanisms, they can punish upstream deviations by offering the agent a different mechanism. There is thus no need to offer the same extended type of the agent multiple contracts to choose from as a function of who deviated upstream. All pure-strategy profiles can thus be sustained with extended direct mechanisms.

Next, consider an environment in which downstream principals do not observe the mechanisms offered upstream, but observe the payoff-relevant decisions taken in these mechanisms. In this environment, restricting the principals’ strategy space may mean restricting the extent to which different principals can have different out-of-equilibrium beliefs about the mechanisms used upstream. When the agent’s strategy is non-Markov, this means restricting the principals’ beliefs about the agent’s behavior downstream. Such restrictions may preclude the possibility of sustaining certain outcomes, as shown in [9]. Assuming the principals compete in menus is thus not without loss of generality. However, as discussed in Section 6, the problems with out-of-equilibrium beliefs disappear if one restricts attention to Markov-perfect equilibria. This assumption is standard in the literature that assumes that upstream payoff-relevant decisions are observable (cf. [4,5,13]). Furthermore, because in these environments the principals directly observe the payoff-relevant decisions contracted upstream, there is no need for them to ask the agent to communicate such information. By implication, all Markov-perfect equilibrium outcomes can

be sustained with standard direct revelation mechanisms in which the agent simply reports his exogenous type.

Finally, consider an environment in which the downstream principals observe the entire upstream history, including the messages sent by the agent upstream. Because the only information that is private to the agent is his exogenous type, one may restrict the principals to offer standard direct revelation mechanisms. However, as shown in [6] for the single-principal case, what is not without loss of generality is restricting the agent's strategy to be truthful. The same is true with multiple principals.

While the observability of upstream histories is exogenous in the literature discussed above, a few recent papers consider environments in which the decision to disclose information to downstream principals is endogenous. Calzolari and Pavan [8], for example, derive general conditions for the optimality of privacy in sequential contracting, while [23] examines the interaction between information disclosure and career concerns in a common agency model of the labor market.

Timing of contracting. The timing of contracting can also vary significantly from one application to another. First, the sequence of bilateral relationships can be endogenously determined by the agent's participation decisions, instead of being exogenously fixed as in our benchmark model. This is the case, for example, when a buyer chooses in each period which seller to purchase from. The sequence of bilateral relationships is endogenous, for example, in [2,20,30]. In this setting, principals necessarily observe part of the upstream history. In fact, they observe at least the agent's upstream participation decisions. Furthermore, if in period t the agent decides to participate in principal j 's mechanism, then at any subsequent date, principal j necessarily knows the payoff-relevant decisions determined in period t . The problems with menus are then the same as in the case of observable upstream payoff-relevant decisions. However, because these problems emerge only when the players' strategies are non-Markov, it remains possible to use menus to characterize the entire set of social choice functions that can be sustained as Markov-perfect equilibria, as discussed in Section 6. As for direct mechanisms, restricting attention to truthful equilibria is not in general without loss of generality in these environments. The reason is that the agent contracts with the same principal multiple times. Unless the principals can commit to long-term mechanisms, having the agent report truthfully in each period precludes the possibility of sustaining certain outcomes. The problems with truth-telling are the same discussed above for fully observable histories. These problems however vanish if one assumes the agent contracts with each principal at most once, or if one assumes that preferences are common knowledge and restricts attention to pure-strategy equilibria, as in [5].

The sequentiality in the principals-agent decisions may also vary across applications. In our benchmark model, a pair of irreversible decisions, one for the principal, the other for the agent, is committed in each period. There are environments in which instead the principals offer their mechanisms sequentially but where the agent takes decisions only after having observed all mechanisms. Clearly, if the principals do not observe the mechanisms offered by the other principals, these environments are strategically equivalent to simultaneous common agency. In this case, we know from [28] and [19] that all equilibrium outcomes can be sustained with menus. In [26], we show that if one restricts attention to equilibria in which the agent's strategy satisfies the analog of the Markov property described above,⁶ then all equilibrium outcomes can be

⁶ In simultaneous games, the agent's strategy is said to be Markov if the decisions the agent selects from each menu depend on his type and the decisions he selects with the other principals, but not on the menus offered by the latter.

characterized by (i) restricting the principals to offer menus that can be conveniently described as incentive-compatible extended direct mechanisms and (ii) by restricting the agent to follow a truthful strategy. The definition of extended direct mechanisms is adjusted to take into account that all decisions are determined simultaneously: the agent is thus asked to report his type along with the decisions he is inducing (through the messages he is sending) with any of the other principals. In the Supplementary material, we show that the same mechanisms sustain all pure-strategy Markov equilibrium outcomes in sequential games in which the principals observe the mechanisms offered upstream before offering their own mechanism and in which the agent takes decisions with each principal at the end, after having observed the mechanisms offered by all principals.⁷ This extensive form corresponds to the environments examined, for example, in [3, 7,17,29,32].

The extensive forms discussed above are clearly only a subset of the many extensive forms considered in the literature. For example, the repeated common agency games of [5] and [24] do not belong to any of these cases.⁸ Nor does the “agenda setting” game of [5].⁹ However, most applications combine elements from the various extensive forms discussed above. We thus expect our results to be of guidance in other settings as well.

2. The private contracting model

Players, actions and contracts. There are $n \in \mathbb{N}$ principals who contract sequentially and non-cooperatively with the same agent, A . Each principal P_i is indexed by the date $i \in \mathcal{N} \equiv \{1, \dots, n\}$ at which she contracts with the agent. Each P_i must select a contract $x_i: E_i \rightarrow \mathcal{A}_i$ from a set X_i of feasible contracts. A contract specifies the action $a_i \in \mathcal{A}_i$ that P_i will take in response to the agent’s choice of action/effort $e_i \in E_i$. Both E_i and \mathcal{A}_i may have different interpretations depending on the application under examination. In the relationship between a buyer and a seller, a_i may represent the price that the seller charges to the buyer when the latter chooses quantity/quality e_i . Similarly, when A performs a task on behalf of P_i , a_i may represent the payment that P_i promises as a function of the agent’s performance. Depending on the environment, the set of feasible contracts X_i may also be more or less restricted. For example, in standard moral hazard models, e_i is assumed to be the agent’s private information, in which case x_i is constant over E_i while the decision a_i should be interpreted as a payment scheme that rewards the agent as a function of some performance measure correlated with the agent’s effort. In contrast, in menu auctions, e_i is assumed to be verifiable in which case x_i specifies the transfer that P_i pays to the agent as a function of the “policy” that the agent selects.¹⁰ A profile of contracts will be denoted by $x \equiv (x_1, \dots, x_n)$. Similarly, profiles of principals’ actions and of agent’s efforts will be denoted by $a \equiv (a_1, \dots, a_n)$ and $e \equiv (e_1, \dots, e_n)$ respectively. We assume the sets E_i and \mathcal{A}_i do not depend on upstream decisions.

Payoffs. All players have expected utility preferences. A principal’s payoff is represented by the function $u_i(\theta, e, a)$. Similarly, the agent’s payoff is described by the function $v(\theta, e, a)$.

⁷ The reasons why menus or extended direct mechanisms may fail to sustain certain mixed-strategy equilibrium outcomes are the same as in environments with observable upstream mechanisms.

⁸ In these papers, the agent simultaneously contracts with multiple principals at each period.

⁹ In this game, upstream decisions determine the sets of feasible decisions downstream.

¹⁰ We assume that x_i does not depend on the agent’s effort at dates $t \neq i$. The results can however easily accommodate the case where $x_i: \prod_{i=1}^n E_i \rightarrow \mathcal{A}_i$ and where the agent chooses $e = (e_1, \dots, e_n)$ only at $t = n + 1$. See Section 6 for a discussion.

The variable θ denotes the agent’s exogenous private information. Principals share a common prior about θ represented by the distribution F with support Θ . To avoid measure-theoretic considerations, the sets Θ , $E \equiv \prod_i E_i$ and $\mathcal{A} \equiv \prod_i \mathcal{A}_i$ will be assumed to be finite.

Mechanisms. Principals compete in mechanisms. A mechanism for P_i consists of a message space \mathcal{M}_i , a set of signals \mathcal{S}_i , and a mapping $\phi_i : \mathcal{M}_i \rightarrow \Delta(X_i \times \mathcal{S}_i)$; when A sends a message $m_i \in \mathcal{M}_i$, P_i responds by selecting a contract from X_i and sending the agent a signal $s_i \in \mathcal{S}_i$.¹¹ The role of these signals is to control the agent’s posterior beliefs over X_i so as to fashion his effort decisions.¹² Indeed, the mechanism ϕ_i can also be seen as a mapping $\phi'_i : \mathcal{M}_i \rightarrow \Delta(X_i \times \Delta(X_i))$, where $\Delta(X_i)$ denotes the set of the agent’s posterior beliefs over X_i . In turn, such a mechanism is equivalent to one where P_i randomizes over $\Delta(X_i)$, then informs the agent of the result of such randomization—i.e. of the particular lottery $d_i \in \Delta(X_i)$ selected—and finally picks a contract x_i from X_i according to the lottery d_i after the agent chooses e_i . Letting Y_i denote the set of (feasible) *stochastic* contracts $y_i : E_i \rightarrow \Delta(\mathcal{A}_i)$, we can then suppress the signals s_i and describe a mechanism as a mapping $\phi_i : \mathcal{M}_i \rightarrow D_i$ such that, when A selects a message $m_i \in \mathcal{M}_i$, P_i randomizes over Y_i according to the lottery $\delta_i = \phi_i(m_i) \in \Delta(Y_i)$ and then informs A of the contract $y_i : E_i \rightarrow \Delta(\mathcal{A}_i)$ selected by the lottery δ_i before the agent chooses effort e_i . We denote by $D_i \subseteq \Delta(Y_i)$ the set of feasible lotteries over Y_i and by $\text{Im}(\phi_i) \equiv \{\delta_i \in \Delta(Y_i) : \exists m_i \in \mathcal{M}_i \text{ s.t. } \phi_i(m_i) = \delta_i\}$ the set of lotteries in the range of ϕ_i . Once again, depending on the application of interest, the set of feasible lotteries D_i may be more or less restricted. For example, in certain applications, it is customary to restrict the principals to offer deterministic mechanisms: this can be accommodated by restricting D_i to contain only degenerate lotteries over deterministic contracts $x_i : E_i \rightarrow \mathcal{A}_i$. More generally, the set of feasible lotteries D_i incorporates all sorts of exogenous restrictions dictated by the environment under examination. What is important to us, is that this set is a primitive of the environment, not a choice of P_i .

To save on notation, in the sequel we will often denote a mechanism by ϕ_i , thus dropping the specification of the message space \mathcal{M}_i , when this does not create confusion. We then let Φ_i denote the set of feasible mechanisms for P_i and $\phi \equiv (\phi_1, \dots, \phi_n)$ and $\phi_{-i} \equiv (\phi_1, \dots, \phi_{i-1}, \phi_{i+1}, \dots, \phi_n)$ denote respectively a profile of mechanisms for the n principals and a collection of mechanisms for all P_j with $j \neq i$.¹³ As is standard, we assume that principals can fully commit to their mechanisms and that each principal cannot contract directly over the mechanisms, or the contracts, of the other principals.

Timing. The sequence of events is the following:

- At date 0, A privately learns θ .
- At date i , with $i = 1, \dots, n$, P_i secretly offers the agent a mechanism $\phi_i \in \Phi_i$. A then chooses a message m_i from \mathcal{M}_i , the lottery $\phi_i(m_i) \in \Delta(Y_i)$ determines the contract y_i , and finally given y_i , A chooses effort e_i and P_i ’s action is determined by the lottery $y_i(e_i) \in \Delta(\mathcal{A}_i)$. None of the principals P_j with $j \neq i$ observes $(\phi_i, m_i, y_i, e_i, a_i)$.
- At date $n + 1$, the game ends.

¹¹ Throughout, for any measurable set Z , $\Delta(Z)$ will denote the set of probability measures over Z . Furthermore, given any $\delta \in \Delta(Z)$, we will denote by $\text{Supp}[\delta]$ the support of δ .

¹² See also [28] for a discussion of the role of signals in a mechanism.

¹³ Given any collection of sets $\{Z_i\}_{i=1}^n$, the following notation will be used throughout the paper: $z_{-i} \equiv (z_1, \dots, z_{i-1}, z_{i+1}, \dots, z_n) \in Z_{-i} \equiv \prod_{j \neq i} Z_j$; $z_i^- \equiv (z_k)_{k=1}^{i-1} \in Z_i^- \equiv \prod_{k=1}^{i-1} Z_k$; $z_i^+ \equiv (z_k)_{k=i+1}^n \in Z_i^+ \equiv \prod_{k=i+1}^n Z_k$; $Z_n^- \equiv Z_n^+ \equiv \emptyset$.

Although not explicitly modeled, the analysis can easily accommodate the agent’s decision (not) to participate in a mechanism. It suffices to add to each mechanism a “null” contract that specifies the default actions that are implemented in case of non-participation, such as no trade at a null price.

Strategies and equilibrium. Let Γ summarize the common agency game described above. A strategy for P_i in Γ is simply a probability measure $\sigma_i \in \Delta(\Phi_i)$ over the set of feasible mechanisms Φ_i . The agent’s (behavioral) strategy at date i given the history $h_i \equiv (h_i^-, \phi_i) \in \mathcal{H}_i$, where $h_i^- \equiv (\theta, e_i^-, a_i^-, \phi_i^-, m_i^-, y_i^-)$ denotes the upstream history, will be denoted by $\sigma_A(h_i)$; the strategy $\sigma_A(h_i)$ consists of a probability measure $\mu(h_i) \in \Delta(\mathcal{M}_i)$ over the messages \mathcal{M}_i along with a probability measure $\xi(h_i, m_i, y_i) \in \Delta(E_i)$ over effort conditional on h_i, m_i and the realized contract y_i . Finally, we denote by σ_A the agent’s complete strategy in Γ —with generic behavioral strategy $\sigma_A(h_i)$ at date i —and by $\sigma \equiv (\{\sigma_i\}_{i=1}^n, \sigma_A)$ an entire strategy profile for the agent and the n principals. Letting $\mathcal{E}(\Gamma)$ denote the set of *perfect Bayesian equilibria* (PBE) for the sequential common agency game Γ , for any $\sigma^* \in \mathcal{E}(\Gamma)$, we then denote by $\pi_{\sigma^*} : \Theta \rightarrow \Delta(\mathcal{A} \times E)$ the *social choice function* (SCF) induced by σ^* . Hereafter, unless otherwise specified, when we refer to equilibrium, we mean PBE. However, it is immediate that all the results for the benchmark model of private contracting remain valid even if one considers *sequential equilibrium* as the solution concept.¹⁴

3. Difficulties with standard direct mechanisms: an example

Consider the following simple environment with no asymmetric information and no effort. Each principal $i = 1, 2$ must select a decision $a_i \in \mathcal{A}_i$, with $\mathcal{A}_1 = \{t, b\}$ and $\mathcal{A}_2 = \{l, r\}$.¹⁵ Payoffs, respectively for P_1, P_2 , and A , are given by the triples in the following table:

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

Game 1

Now suppose that principals were restricted to offer standard direct revelation mechanisms. Because in this example the agent does not possess any exogenous private information, a direct revelation mechanism simply coincides with a decision $a_i \in \mathcal{A}_i$. The set of equilibrium outcomes would then coincide with the set of Nash equilibria in the simultaneous game between the principals only: (t, l) and (t, r) .

Next consider a game in which the principals’ strategy space is the set of all indirect mechanisms with message space $\mathcal{M}_i = \{0, 1\}$. In this game, (b, l) can also be sustained as an equilibrium outcome. The equilibrium features P_1 offering the mechanism that responds to both messages with the decision b and P_2 offering the mechanism that responds to $m_2 = 0$ with l and to $m_2 = 1$ with r . In equilibrium, the agent chooses $m_2 = 0$ with P_2 thus implementing the outcome (b, l) . The role of the decision r in P_2 ’s mechanism is to block a possible deviation to t by P_1 : if the latter were to deviate to t , the agent would choose $m_2 = 1$ with P_2 giving a payoff of 0 to P_1 .

¹⁴ Note that a PBE requires to specify also beliefs for each player. However, because with private contracting these beliefs are always pinned down by Bayes’ rule, hereafter we denote a PBE simply by its strategy profile σ^* .

¹⁵ In this example, a contract simply coincides with the choice of a payoff-relevant decision $a_i \in \mathcal{A}_i$.

Clearly, the same outcome (b, l) can be sustained in the game in which the principals offer menus, but not in the game in which they offer standard direct revelation mechanisms. The problem with these mechanisms is that they may not be responsive enough to possible deviations in upstream relationships. This problem can be addressed by considering more general direct mechanisms in which the agent reports, in addition to his type, the endogenous payoff-relevant decisions contracted upstream. In this example with no exogenous private information, an extended direct mechanism for P_1 simply coincides with an element of \mathcal{A}_1 , whereas for P_2 with a mapping $\phi_2^D: \mathcal{A}_1 \rightarrow \mathcal{A}_2$. It is then immediate that the following mechanisms sustain (b, l) : P_1 chooses b , whereas P_2 offers the mechanism $\phi_2^{D*}(t) = r$ and $\phi_2^{D*}(b) = l$. Note that ϕ_2^{D*} is incentive-compatible both on and off equilibrium: whatever decision a_1 is taken upstream, the agent has the incentives to report it truthfully to P_2 .

More generally, when information is complete, as in the example above, any deterministic SCF that can be sustained with any arbitrary strategy space for the principals can also be sustained by restricting the principals to offer extended direct mechanisms. The same is true with incomplete information and/or stochastic SCFs, provided one restricts the agent's strategy to be Markov.

To see why, with non-Markov strategies, the direct mechanisms described above may fail to sustain certain outcomes, suppose that the principals' strategy space is now the set of all stochastic mechanisms that map $\mathcal{M}_i = \{0, 1, 2\}$ into lotteries over \mathcal{A}_i . Continue to assume that the agent observes a_1 before contracting with P_2 . The following is then an equilibrium. P_1 offers the mechanism ϕ_1^* that maps each message m_1 into the lottery that gives t and b with equal probabilities. On her part, P_2 offers a mechanism that responds to $m_2 = 0$ with a lottery δ that gives l with probability equal to $1/4$, to $m_2 = 1$ with a lottery λ that gives l with probability $1/2$, and to $m_2 = 2$ with a (degenerate) lottery μ that gives r with certainty. In equilibrium, A chooses $m_2 = 0$ conditional on t and $m_2 = 1$ conditional on b . These mechanisms are sustained by a non-Markov strategy according to which, whenever P_1 deviates from the equilibrium distribution over \mathcal{A}_1 , the agent chooses $m_2 = 2$ instead of $m_2 = 0$ after t is realized (note that the agent is indifferent about a_2 when $a_1 = t$).

Clearly, this outcome can also be sustained in the game in which the principals offer menus (of lotteries). However, it cannot be sustained in the game in which they offer extended direct mechanisms. In fact, to prevent a deviation from P_1 it is essential that the agent be given the possibility of choosing the (out-of-equilibrium) lottery μ with P_2 in response to a deviation by P_1 that altered the distribution upstream but nonetheless led to an equilibrium decision. This cannot be done with our extended direct mechanisms.¹⁶

When downstream principals observe part of the upstream history (such as the mechanisms offered upstream and/or the decisions taken in these mechanisms), there are additional reasons why restricting the principals' strategy space may preclude the possibility of sustaining certain outcomes. In these environments, even the restriction to menus is not always without loss of generality, as discussed in Section 6.

4. Menus

In this section we first show that the menu theorems of simultaneous common agency extend to sequential common agency when contracting is private. We then show that, contrary to simul-

¹⁶ Similar issues arise with incomplete information, as discussed in Section 4.

taneous games, all equilibrium outcomes can be characterized by restricting attention to strategy profiles in which the principals do not mix over their menus and in which the agent does not mix over the contracts he selects with each principal. Finally, we show that, when information is complete, all deterministic social choice functions can be sustained through equilibria in which the agent’s strategy is Markov. Proving these results is not only useful for applications, but also a key step for the results in the next section.

Definition 1. A *menu* is a mechanism $\phi_i^M : \mathcal{M}_i^M \rightarrow D_i$ such that (a) $\mathcal{M}_i^M \subseteq D_i$, and (b) for any $\delta_i \in \mathcal{M}_i^M$, $\phi_i^M(\delta_i) = \delta_i$.

A menu is thus a mechanism whose message space is equal to its image and whose mapping is the identity function. In what follows, we denote by Φ_i^M the set of all possible menus for principal i and by Γ^M the “menu game” in which the set of feasible mechanisms for each P_i is Φ_i^M .

Now consider any enlargement of the menu game, that is, a game in which principals have “more” mechanisms than in Γ^M .

Definition 2. The game Γ is an *enlargement* of Γ^M ($\Gamma \succ \Gamma^M$) if for all $i \in \mathcal{N}$,

- (i) there exists an embedding $\alpha_i : \Phi_i^M \rightarrow \Phi_i$;¹⁷
- (ii) 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 $\Phi_i \supseteq \Phi_i^M$ for all i . More generally, an enlargement is a game in which every Φ_i is larger than Φ_i^M in the sense that each menu ϕ_i^M is also present in Φ_i , although possibly with a different representation.¹⁸

Theorem 1. Let $\Gamma \succ \Gamma^M$. A SCF π can be sustained as an equilibrium of Γ if and only if it can be sustained as an equilibrium of Γ^M .

First consider the “only if” part of the result. The idea behind the proof is simple and can be sketched as follows. Suppose the SCF π can be sustained as an equilibrium of Γ and let σ be the supporting strategy profile. Now suppose in Γ^M each principal’s strategy is such that, for any set of menus $R_i \subseteq \Phi_i^M$,

$$\sigma_i^M(R_i) = \sigma_i \left(\bigcup_{\phi_i^M \in R_i} \Phi_i(\phi_i^M) \right) \tag{1}$$

where $\Phi_i(\phi_i^M) \equiv \{\phi_i \in \Phi_i : \text{Im}(\phi_i) = \text{Im}(\phi_i^M)\}$. This strategy consists in offering each menu ϕ_i^M with a probability equal to the total probability assigned by the original strategy σ_i in Γ to the set of all mechanisms in Γ whose image coincides with that of ϕ_i^M . When in the menu game Γ^M all

¹⁷ Formally, the embedding $\alpha_i : \Phi_i^M \rightarrow \Phi_i$ is an injective mapping such that, for any pair of mechanisms ϕ_i^M, ϕ_i with $\phi_i = \alpha_i(\phi_i^M)$, the following are true: (a) $\text{Im}(\phi_i) = \text{Im}(\phi_i^M)$; (b) there exists an injective function $\tilde{\alpha}_i : \mathcal{M}_i^M \rightarrow \mathcal{M}_i$ from the message space of ϕ_i^M to the message space of ϕ_i such that $\phi_i(\tilde{\alpha}_i(\delta_i)) = \delta_i$ for any $\delta_i \in \mathcal{M}_i^M$.

¹⁸ The requirement that each mechanism ϕ_i has a compact image guarantees that the agent’s best response is well defined.

downstream principals are expected to follow the strategy given in (1), in the continuation game that starts after P_i offers the menu ϕ_i^M it is clearly optimal for A to induce the same outcomes he would have induced in Γ had P_i offered one of the mechanisms in $\Phi_i(\phi_i^M)$. This also implies that after P_i offers the menu ϕ_i^M , it is optimal for the agent to use the conditional distribution

$$\sigma_i(\phi_i \mid \Phi_i(\phi_i^M))$$

to determine his behavior *at any subsequent information set*.¹⁹ Furthermore, when each P_i and A follow the strategies described above, the distribution over $E \times \mathcal{A}$ is the same as in Γ . Starting from σ one can thus construct an equilibrium σ^M for Γ^M that sustains the same outcomes.

Next, consider the “if” part of the result. Suppose the SCF π can be sustained as an equilibrium of Γ^M and let σ^M denote the sustaining strategy profile. Then there always exists an equilibrium σ in Γ that sustains the same outcomes. The agent’s strategy σ_A is constructed by “extending” the original strategy σ_A^M over Γ , as follows. For any history $h_t = ((\phi_i, m_i, y_i, e_i, a_i)_{i=1}^{t-1}, \phi_t)$, σ_A induces the same joint distribution over $E_i \times \mathcal{A}_i$ as the strategy σ_A^M in Γ^M after the history $h_t^M = ((\phi_i^M, \delta_i, y_i, e_i, a_i)_{i=1}^{t-1}, \phi_t^M)$ where the history h_t^M is constructed from the history h_t replacing each mechanism ϕ_i with the menu ϕ_i^M whose image is the same as that of ϕ_i (i.e. $\text{Im}(\phi_i^M) = \text{Im}(\phi_i)$) and each message m_i with the lottery $\delta_i = \phi_i(m_i)$. As for the principals, each strategy σ_i is simply the “translation” of the original strategy σ_i^M in Γ^M using the embedding α_i , i.e. $\sigma_i = \alpha_i(\sigma_i^M)$.²⁰ Given the principals’ strategies $(\sigma_i)_{i=1}^n$, the agent’s strategy σ_A is sequentially rational for the agent. Furthermore, given (σ_A, σ_{-i}) , no principal has an incentive to deviate from the strategy σ_i .

Any equilibrium σ^M of the menu game is thus *weakly robust* in the sense of [28]: for any enlargement Γ of Γ^M , there is an extension of the agent’s strategy σ_A^M over Γ such that it remains optimal for the principals to offer the same equilibrium menus as in the original game and for the agent to induce the same outcomes.²¹

When Γ is not an enlargement of Γ^M , because the environment imposes certain restrictions on the sets Φ_i , 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 becomes necessary to restrict the principals to offer only those menus that could have been offered in Γ : that is, the set of feasible menus for P_i must be restricted to $\tilde{\Phi}_i^M \equiv \{\phi_i^M : \text{Im}(\phi_i^M) = \text{Im}(\phi_i) \text{ for some } \phi_i \in \Phi_i\}$. For simplicity, in the sequel we will restrict attention to environments in which the set of feasible menus for each principal is the entire set of *all* possible menus.

The aforementioned results show that the menu theorems of simultaneous common agency extend to sequential common agency when contracting is private. The reason why the results do not follow directly from those theorems is twofold. First, the decisions the agent takes with his upstream principals are irreversible at the time he contracts with the downstream principals. This means that the agent’s behavior at period t can be conditioned not only on the mechanisms of-

¹⁹ The existence of such conditional measures as well as the specification of how the agent “translates” his behavior in Γ in his behavior in Γ^M is described in Appendix A.

²⁰ Formally, for any measurable set $R \subseteq \Phi_i$, $\sigma_i(R) = \sigma_i^M(\tilde{\Phi}_i^M)$, where $\tilde{\Phi}_i^M = \{\phi_i^M : \alpha_i(\phi_i^M) \in R\}$.

²¹ Equilibria in Γ^M are weakly robust, but not necessarily *strongly robust*: one may also be able to construct an enlargement Γ and an extension σ_A of σ_A^M over Γ such that $\sigma = (\sigma_A, \{\alpha_i(\sigma_i^M)\}_{i=1}^n)$ is not an equilibrium of Γ . However, what seems important to us is that any equilibrium outcome in Γ^M remains an equilibrium outcome also in Γ . That it can be sustained by *any* strategy profile $(\sigma_A, \{\alpha_i(\sigma_i^M)\}_{i=1}^n)$ in which σ_A is an extension of σ_A^M is not essential.

ferred upstream but also on the endogenous payoff-relevant decisions taken in these mechanisms. Second, at the time the agent commits a decision with principal t , he has not seen yet the mechanisms offered by the downstream principals. These differences do not pose serious problems. However, they require an adaptation of the arguments used in simultaneous games.

4.1. Pure strategies

The next result, which is specific to sequential contracting, goes a step further by showing that, in settings in which $D_i = \Delta(Y_i)$, i.e. in environments in which principals can offer any lottery over Y_i , the characterization of the equilibrium outcomes is further simplified by the fact that one can restrict attention to pure-strategy profiles.

Definition 3. A strategy profile $\sigma \in \mathcal{E}(\Gamma)$ is a *pure-strategy equilibrium* if and only if it is an equilibrium in which (a) no principal i mixes over Φ_i ; (b) after any history h_i , the agent does not mix over $\mathcal{M}_i, \forall i \in \mathcal{N}$.

While in a pure-strategy equilibrium, the agent does not mix over the messages he sends to the principals, he may however mix over effort.

Theorem 2. Suppose $D_i = \Delta(Y_i)$ for all i . A SCF π can be sustained as an equilibrium of Γ^M only if it can be sustained as a pure-strategy equilibrium.

Suppose the SCF π is sustained by a mixed-strategy equilibrium σ^M in which P_i randomizes over Φ_i^M according to σ_i^M and in which, given some ϕ_i^M , the agent randomizes over the different lotteries in ϕ_i^M according to $\mu(h_i^-, \phi_i^M)$.²² The same SCF can be sustained by an equilibrium $\hat{\sigma}^M$ in which P_i offers with probability one the menu $\hat{\phi}_i^M$ that contains the compound lotteries (indexed by h_i^-)

$$\int_{\phi_i^M \in \Phi_i^M} \int_{\delta_i \in \mathcal{M}_i^M} \delta_i d\mu(h_i^-, \phi_i^M) d\sigma_i^M \tag{2}$$

that can be obtained by mixing with distribution σ_i^M over the compound lotteries $\int \delta_i d\mu(h_i^-, \phi_i^M)$ that the agent would have induced, for each h_i^- , by following his original strategy $\mu(h_i^-, \phi_i^M)$. Holding constant σ_i^+ , when the upstream history is h_i^- and P_i offers the menu $\hat{\phi}_i^M$, it is then clearly optimal for A to choose with probability one the lottery $\delta_i(h_i^-, \phi_i^M)$ given in (2); at any downstream information set, the agent then adjusts his behavior to induce the same outcomes as in the original equilibrium σ^M .²³

When instead P_i offers any menu $\phi_i^M \neq \hat{\phi}_i^M$, let the agent choose with probability one the lottery $\delta_i(h_i^-, \phi_i^M)$ in $\text{Supp}[\sigma_A^M(h_i^-, \phi_i^M)]$ that minimizes P_i 's expected payoff taking into account that at any subsequent information set the agent's behavior will continue to be determined by the original strategy σ_A^M . When all other principals are expected to follow the same strategy as in σ^M , the strategy $\hat{\sigma}_A^M$ constructed this way is clearly sequentially optimal for the agent.

²² Recall that $\mu(h_i, \phi_i^M) \in \Delta(\mathcal{M}_i^M)$ denotes the agent's communication strategy with P_i .

²³ The details of how the agent adjusts his behavior downstream to induce the same distribution over $E \times \mathcal{A}$ as in the original equilibrium σ^M are in Appendix A.

Furthermore, given this strategy, no principal has an incentive to deviate. Iterating across all i , it is then possible to construct a pure-strategy equilibrium that implements the same outcomes as σ^M .

The point here is that any randomization induced by a principal mixing over her mechanisms or by the agent mixing over his messages can be replicated by having the principal offer a single menu of lotteries over contracts and having the agent selecting deterministically a lottery as a function of his upstream contractual history. For the transformation described above to work, it is however essential that any principal can offer any lottery in $\Delta(Y_i)$, which explains the qualification in the theorem. This is an assumption that we will maintain in most of the subsequent analysis. We will however be careful to clarify how our results extend to environments in which $D_i \subsetneq \Delta(Y_i)$ for some i , a restriction that is common, for example, in applications that assume that only (degenerate lotteries over) deterministic contracts are feasible.

The possibility of restricting attention to pure-strategy profiles is appealing for two reasons. First, it is reminiscent of a similar result for games with a single mechanism designer. Second, it is common practice in applications to restrict attention to pure-strategy equilibria—Theorem 2 provides a possible justification for such a practice.

It is also important to note that such a result does not have a counterpart in simultaneous common agency. In fact it is essential that the agent takes decisions with each principal after having committed irreversible decisions upstream and before having seen the mechanisms offered downstream. To see this, consider a simple environment with only two principals, with no adverse selection and with no effort, so that $|\Theta| = |E_i| = 1, i = 1, 2$. Suppose in addition that the sets of primitive decisions are $\mathcal{A}_1 = \{t, b\}$ and $\mathcal{A}_2 = \{l, r\}$. The payoffs, respectively, for P_1, P_2 and A , are as in the following table:

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

Game 2

The social choice function that selects (t, l) and (b, r) respectively with probability $q \in (0, 1)$ and $1 - q$ can be sustained as an equilibrium both in our sequential game (with private contracting) and in the canonical simultaneous version of the game considered in the literature. In both games, with probability $q \in (0, 1)$, P_1 offers the degenerate menu that contains the lottery that selects t with certainty, whereas with probability $1 - q$ she offers the degenerate menu that contains the lottery that selects b with certainty. On her part, P_2 offers the menu that contains both the lottery that gives l with certainty and the lottery that gives r with certainty. The agent selects l with P_2 if $a_1 = t$ and r if $a_1 = b$. It is immediate to see that, in the sequential game, the same outcome can be sustained by a pure-strategy equilibrium in which P_1 offers a degenerate menu containing the lottery that gives t and b respectively with probability q and $1 - q$. This is not true in the simultaneous version of the game. Indeed, to sustain the same outcomes with a pure-strategy profile, it is necessary that P_1 offers a menu that contains both the degenerate lottery that gives t with certainty and the degenerate lottery that gives b with certainty. But then A strictly prefers to induce (b, r) rather than randomizing over (b, r) and (t, l) .

4.2. Markov-perfect equilibria

In applications, it is also customary to restrict attention to equilibria in which a player's strategy depends on the upstream history only through its payoff relevant component.²⁴ With private contracting, this is necessarily the case for principals. In what follows, we examine the implications of assuming such a property holds also for the agent's strategy.

Formally, for any upstream history $h_i^- = (\theta, \phi_i^-, m_i^-, y_i^-, e_i^-, a_i^-)$, let

$$\theta_i^E \equiv (\theta, e_i^-, a_i^-) \in \Theta_i^E \equiv \Theta \times E_i^- \times \mathcal{A}_i^-$$

denote the payoff-relevant component of h_i^- . For a reason that will become clear in the next section, hereafter we refer to θ_i^E as the agent's *extended type*.

Definition 4. The agent's strategy σ_A is *Markov* at $t = i$ if and only if, for any ϕ_i and any pair of upstream histories h_i^-, \hat{h}_i^- with $h_i^- = (\theta_i^E, \phi_i^-, m_i^-, y_i^-)$ and $\hat{h}_i^- = (\theta_i^E, \hat{\phi}_i^-, \hat{m}_i^-, \hat{y}_i^-)$,

$$\sigma_A(h_i^-, \phi_i) = \sigma_A(\hat{h}_i^-, \phi_i).$$

A *Markov-perfect equilibrium* (hereafter, MPE) is an equilibrium in which σ_A is Markov at any t .

Assuming the agent's strategy is Markov seems appealing. However, it is important to understand what social choice functions cannot be sustained with these strategies. The next result provides an answer and can be seen as a possible justification for restricting attention to Markov perfect equilibria in certain environments.

Theorem 3. Suppose information is complete (i.e. $|\Theta| = 1$). Then any deterministic SCF π that can be sustained as an equilibrium of Γ^M can also be sustained as a pure-strategy MPE.

The idea behind the proof is the following. When preferences are common knowledge and the SCF π is deterministic, there is a unique sequence of equilibrium decisions and hence a unique sequence of equilibrium extended types. Now, assume the agent's strategy in the equilibrium σ^{M*} that supports π is not Markov. Then consider the alternative strategy profile $\tilde{\sigma}^M$ in which all principals offer the same menus as in σ^{M*} and in which the agent behaves as follows. At any period t , if the extended type θ_t^E is the equilibrium one, the agent implements the equilibrium decisions for that extended type, independently of which particular upstream history h_t^- led to θ_t^E . If instead, θ_t^E is not the equilibrium extended type, then let $j \leq t - 1$ be the first date at which a departure from the sequence of equilibrium decisions occurred. Starting from period t , at any downstream information set, the agent then chooses among the decisions that are sequentially optimal for him, those that minimize the payoff of principal j . Given this Markov strategy for the agent, no principal has a profitable deviation. The strategy profile $\tilde{\sigma}^M$ is thus an equilibrium for Γ^M and sustains the same outcomes as σ^{M*} .

In the argument sketched above, it is essential that the agent be able to identify (and punish) an upstream principal who deviated from equilibrium play simply by looking at the extended type. This is always possible when information is complete and the SCF π is deterministic. When instead $|\Theta| > 1$ and/or the SCF π is stochastic, it may be necessary to have the agent condition his behavior not only on θ_t^E , but also on payoff-irrelevant information such as the mechanisms

²⁴ See, among others, [13] and [5].

offered upstream. This permits the agent to punish deviations that altered the *distribution* of upstream payoff-relevant decisions but nevertheless led to equilibrium extended types.

It is also worth noticing that even if one is interested in characterizing only deterministic SCFs, it may be necessary to allow for menus that contain (non-degenerate) lotteries. To see this, assume that $\mathcal{A}_1 = \{t, b\}$ and $\mathcal{A}_2 = \{l, m, r\}$. The payoffs, respectively for P_1 , P_2 and A are given by the triples (u_1, u_2, v) in the following table:

$a_1 \backslash a_2$	l	m	r
t	1 2 1	0 1 0	0 0 0
b	2 3 1	-2 4 0	3 0 2

Game 3

The outcome (t, l) cannot be sustained by restricting attention to menus that contain only degenerate lotteries. Indeed, to punish a possible deviation by P_1 to b , the equilibrium menu offered by P_2 should also contain m . However, conditional on b , A necessarily prefers l over m , which implies that it is impossible to punish P_1 's deviation while satisfying the agent's rationality. On the other hand, the outcome (t, l) can be sustained by having P_2 offer any menu that together with the equilibrium contract l contains a lottery that gives l with probability p , m with probability q , and r with probability $1 - p - q$, where (p, q) is any pair of positive real numbers that satisfies $q \leq (1 - p)/2$ and $q \geq (2 - p)/5$. The first bound on q guarantees that, choosing this lottery is incentive-compatible for the agent given b , while the second ensures that P_1 does not find it profitable to deviate.

5. Extended direct mechanisms

Building on the results in the previous section, we now show that, in most cases of interest for applications, the characterization of the equilibrium outcomes can be further simplified by restricting the principals to offer menus that can be conveniently described as extended direct mechanisms.

Definition 5. An *extended direct mechanism* is a mapping $\phi_i^D : \Theta_i^E \rightarrow D_i$. A *revelation game* Γ^D is a game in which the principals' strategy space is $\Delta(\Phi_i^D)$, where Φ_i^D denotes the set of all possible extended direct mechanisms for $P_i, i = 1, \dots, n$.

Extended direct mechanisms can thus be thought of as menus whose allocations are indexed by the payoff-relevant component of the agent's upstream history. Extended direct mechanisms are thus the analogue of standard direct revelation mechanisms with the only difference being that they may specify contracts also for extended types that have zero measure on the equilibrium path.

Next, let $\tilde{V}(h_i^-, \delta_i, \sigma_i^+)$ denote the maximal payoff that A can obtain in the continuation game that starts at $t = i$ when the upstream history is $h_i^- = (\theta_i^E, \phi_i^-, m_i^-, y_i^-)$, he chooses a lottery δ_i with P_i , and the downstream principals' strategy profile is σ_i^+ . Clearly, $\tilde{V}(h_i^-, \delta_i, \sigma_i^+)$ depends on h_i^- only through θ_i^E ; in the following, we thus denote the agent's continuation payoff by $V(\theta_i^E, \delta_i, \sigma_i^+)$.

Definition 6.

- (i) Fix a strategy profile σ_i^+ for the downstream principals. A mechanism ϕ_i^D is *incentive-compatible* if and only if, for any $\theta_i^E \in \Theta_i^E$ and any $\delta'_i \in \text{Im}(\phi_i^D)$,

$$V(\theta_i^E, \phi_i^D(\theta_i^E), \sigma_i^+) \geq V(\theta_i^E, \delta'_i, \sigma_i^+).$$

- (ii) The agent’s strategy is *truthful* in ϕ_i^D if and only if, for any upstream history h_i^- , the agent truthfully reports θ_i^E .²⁵ Given a strategy profile $\sigma^D \equiv (\{\sigma_i^D\}_{i=1}^n, \sigma_A^D)$ for Γ^D , σ_A^D is said to be truthful if and only if it is truthful in every $\phi_i^D \in \text{Supp}[\sigma_i^D]$, $\forall i \in \mathcal{N}$.
- (iii) A *truthful equilibrium* for Γ^D is an equilibrium in which each mechanism $\phi_i^D \in \text{Supp}[\sigma_i^D]$ is incentive-compatible and σ_A^D is truthful.

Whether the mechanism ϕ_i^D is incentive-compatible depends on the downstream principals’ strategy profile σ_i^+ . Contrary to games with a single mechanism designer, incentive-compatibility must thus be established by backward induction. Also note, given a mechanism ϕ_i^D , the agent’s strategy is truthful in ϕ_i^D if and only if the agent truthfully reports *any* extended type. This implies that when a deviation from equilibrium occurred at some date $t < i$, the agent still reports his extended type truthfully to P_i .

The following result relates the set of equilibrium outcomes that can be sustained with extended direct mechanisms to the set of outcomes that can be sustained with menus. In virtue of the results in the previous section, it should be clear that this is the relevant comparison.

Theorem 4. Any SCF π that can be sustained as a MPE of Γ^M can also be sustained as a pure-strategy truthful MPE of Γ^D . Furthermore, any SCF that can be sustained as an equilibrium of Γ^D can also be sustained as an equilibrium of Γ^M .

Consider the first part. Suppose there exists a MPE σ^{M*} in the menu game that sustains the SCF π . By Theorem 2, without loss of generality one can assume that σ^{M*} is a pure-strategy profile. That the agent’s strategy in σ^{M*} is Markov in turn implies that, for any θ_i^E and for any menu ϕ_i^M , there exists a unique lottery $\delta_t(\theta_i^E, \phi_i^M)$ such that the agent selects the lottery $\delta_t(\theta_i^E, \phi_i^M)$ from the menu ϕ_i^M when his extended type is θ_i^E , irrespective of the upstream history h_i^- that has conducted to θ_i^E . The equilibrium that sustains π in Γ^D is then constructed by having each principal offer the direct mechanism ϕ_i^{D*} that responds to each θ_i^E with the lottery $\delta_t(\theta_i^E, \phi_i^{M*})$, where ϕ_i^{M*} is the equilibrium menu in Γ^M . When offered the equilibrium mechanism ϕ_i^{D*} , the agent then responds by reporting θ_i^E truthfully and then selecting effort according to σ_A^M as if the game were Γ^M , the menu offered by P_t were ϕ_i^{M*} and the message sent to P_t were $\delta_t(\theta_i^E, \phi_i^{M*})$. When instead the agent is offered a mechanism $\phi_i^D \neq \phi_i^{D*}$, the agent behaves according to σ_A^M as if the game were Γ^M and the menu offered by P_t were ϕ_i^M , where $\text{Im}(\phi_i^M) = \text{Im}(\phi_i^D)$. Given the aforementioned strategies for the principals, the agent’s (Markov) strategy described above is clearly sequentially optimal. Furthermore, given $(\sigma_A^{D*}, \sigma_{-i}^{D*})$, no principal has an incentive to deviate and offer a mechanism $\phi_i^D \neq \phi_i^{D*}$. We conclude that σ^{D*} is a (pure-strategy) MPE of Γ^D and sustains the same outcomes as σ^{M*} in Γ^M .

²⁵ Formally, for any $h_i^- = (\theta_i^E, \phi_i^-, m_i^-, y_i^-)$, $\hat{\theta}_i^E \in \text{Supp}[\mu(h_i^-, \phi_i^D)] \implies \hat{\theta}_i^E = \theta_i^E$.

Note that this result presumes $D_i = \Delta(Y_i)$ for all i . As discussed in the previous section, this is necessary to guarantee that any SCF π that can be sustained as a MPE of Γ^M can also be sustained as a pure-strategy MPE. However, what matters for the possibility of sustaining π as a MPE of Γ^D is only the purity of the agent’s strategy in Γ^M ; that the principals’ strategies are also pure is not important. To understand this, note that when certain lotteries are not feasible, i.e. when $D_i \subsetneq \Delta(Y_i)$ for some i , then it may be impossible to replicate the outcomes induced in Γ^M by the agent mixing over the different contracts in a menu with a direct mechanism that simply asks the agent to report his extended type.²⁶ In environments in which there are restrictions on the sets of feasible lotteries, the result in the first part of the theorem must thus be replaced with the following: Any SCF π that can be sustained as a MPE of Γ^M in which the agent’s strategy is pure can also be sustained as a truthful MPE of Γ^D .

Next, consider the second part of the theorem. Take any $\sigma^D \in \mathcal{E}(\Gamma^D)$. Irrespective of whether the agent’s strategy in σ^D is Markov and of whether σ^D is a pure- or mixed-strategy equilibrium, there always exists a $\sigma^M \in \mathcal{E}(\Gamma^M)$ that sustains the same outcomes. Note that this result does not follow from Theorem 1: in fact, in general, Γ^D is not an enlargement of Γ^M (nor is Γ^M an enlargement of Γ^D). The equilibrium σ^M that sustains π in Γ^M is constructed by having each principal offer each menu ϕ_i^M with the same probability she would have offered all direct mechanisms with the same image as ϕ_i^M . That is, σ_i^M is constructed from σ_i^D using the same transformation as in (1). As for the agent, if P_i offered a menu ϕ_i^M whose image coincides with the image of one of the direct mechanisms in Φ_i^D , then A induces the same outcomes he would have induced in Γ^D had P_i offered²⁷

$$\Phi_i^D(\phi_i^M) \equiv \{\phi_i^D \in \Phi_i^D : \text{Im}(\phi_i^D) = \text{Im}(\phi_i^M)\}.$$

If instead, P_i offered a menu ϕ_i^M that is not in the image of any of the mechanisms in Φ_i^D , then A behaves according to σ_A^D as if the game were Γ^D and the mechanism offered by P_i were ϕ_i^D , where ϕ_i^D is such that

$$\phi_i^D(\theta_i^E) \in \arg \max_{\delta_i \in \text{Im}(\phi_i^M)} V(\theta_i^E, \delta_i, \sigma_i^+) \quad \forall \theta_i^E \in \Theta_i^E.$$

These strategies for the principals and the agent clearly constitute an equilibrium for Γ^M that induces the same outcomes of σ^D .

That any equilibrium outcome of the revelation game is also an equilibrium outcome of the menu game is important because it guarantees that the outcomes that one obtains by restricting the principals to offer extended direct mechanisms are not artificially sustained by the impossibility for the principals to offer certain menus that are not available in the revelation game.

Describing menus as direct revelation mechanisms has proved very convenient in games with a single mechanism designer. The same approach can be used to characterize MPE outcomes in sequential common agency. We illustrate such a possibility in a canonical buyer–sellers example with incomplete information in the Supplementary material.

²⁶ The impossibility of sustaining all outcomes with direct mechanisms when $D \subset \Delta(Y)$ is clearly not specific to common agency; it also applies to settings with a single principal.

²⁷ That is, A uses $\sigma_i^D(\phi_i^D | \Phi_i^D(\phi_i^M))$ to determine his behavior at any subsequent information set.

6. Alternative extensive forms

The model considered so far assumes that contracting is *private*: the mechanism offered by principal i , the contract selected by the agent and the decisions taken in response to this contract, are all information that is private to A and P_i , in the sense that it cannot be observed by any of the other principals. It was also assumed that a pair of payoff-relevant decisions (one for the agent, the other for the principal) is committed at each date instead of some decisions being taken only after all principals have offered their mechanisms. Finally, the sequence of bilateral relationships was exogenous in the sense that the agent could not choose which principal to contract with at each date.

Although these assumptions seem reasonable for many applications, it is important to understand how the preceding results must be adapted to accommodate alternative extensive forms examined in the literature. This is what we do in this section. Hereafter, we summarize the key insights. The formal statements of the results and their corresponding proofs can be found in the Supplementary material.

Observability of upstream payoff-relevant decisions. Consider an environment in which principals observe upstream payoff-relevant decisions before choosing their mechanisms.²⁸

As long as one restricts attention to Markov-perfect equilibria, this extension poses no problems to our characterization results: all Markov-perfect equilibrium outcomes can be characterized by restricting the principals to offer either menus or extended direct revelation mechanisms. Furthermore, because the choice of a mechanism is now contingent on the observable upstream payoff-relevant decisions (e_i^-, a_i^-) , when using direct mechanisms, there is no need to ask the agent to report such information. One can thus drop (e_i^-, a_i^-) from the agent's message and restrict attention to standard direct revelation mechanisms in which the agent simply reports his exogenous type θ .

If, instead, one is also interested in equilibrium outcomes sustained by non-Markov strategies, then restricting the principals to offer menus (or direct mechanisms) may preclude a complete characterization. The reason is that restricting the principals' strategy space means restricting the extent to which principals' out-of-equilibrium beliefs can differ one from the other.²⁹ The fact that, in an indirect game, the same menu can be offered with different mechanisms may permit downstream principals to have different beliefs about the particular mechanism used upstream to select an out-of-equilibrium decision. When the agent's strategy is non-Markov, this means allowing the principals to have different expectations about the agent's behavior in downstream relationships, a property that may be essential to sustain certain outcomes (see [9] for an example that illustrates such a possibility). For the same reason, the set of equilibrium outcomes of the menu game is no longer a superset of the set of equilibrium outcomes of the revelation game—the same menu can in fact be offered through multiple direct mechanisms; what remains true is that the set of *Markov-perfect* equilibrium outcomes is the same in the two games, as shown in Theorem 5 in the Supplementary material.

Observability of upstream mechanisms. Next, consider an environment in which every P_i , $i = 2, \dots, n$, observes the mechanisms ϕ_i^- offered upstream before choosing her own mechanism. As in the benchmark model, P_i does not observe $(m_i^-, y_i^-, e_i^-, a_i^-)$.

²⁸ Upstream decisions are observable in [1,5,17,21,31]. In these papers, the observability of upstream decisions is exogenous. In contrast, [8] and [23] examine models in which the observability of upstream decisions is controlled by upstream principals.

²⁹ Note that there are no out-of-equilibrium beliefs in the benchmark model of private contracting.

In this setting, restricting the principals' strategy space may mean restricting the extent to which payoff-irrelevant distinctions among mechanisms can be used as correlation devices for the principals' decisions. This may preclude a complete characterization of the equilibrium outcomes (see [9] for an illustration). However, one can safely restrict the principals to offer menus if one is interested only in equilibria in which the principals do not mix—to the best of our knowledge, the case considered in all papers that assume observability of upstream mechanisms.

One way of restoring the possibility of using menus to sustain also mixed-strategy equilibria is to allow for public randomizing devices, as in the case of correlated equilibria.³⁰ Alternatively, one may allow the principals to send to each other cheap talk messages whose role is to replicate the role of payoff-irrelevant distinctions among mechanisms that are used as correlation devices.

Furthermore, all equilibrium outcomes that can be sustained in the menu game by restricting the principals' strategies to be pure can also be sustained in the revelation game. Contrary to the case of private contracting, this is true irrespective of whether the agent's strategy is Markov. This is because, in this setting, downstream principals can observe directly a deviation by an upstream principal. They can thereby respond to such a deviation by changing the contract offered to each extended type. There is thus no need to offer the same extended type of the agent a menu of contracts to choose from as a function of what happened upstream—see Theorem 6 in the Supplementary material.

However, note that, in this environment, equilibria in direct mechanisms may fail to be robust, in the sense of [28]: there may exist SCFs that can be sustained as equilibria of the revelation game that cannot be sustained as equilibria of the menu game. The reason is that direct mechanisms restrict a principal's ability to offer the agent out-of-equilibrium allocations. These allocations may be essential to induce a certain behavior by the agent and the downstream principals in the continuation game that starts after a mechanism has been announced. By implication, a principal may have a profitable deviation in the menu game even if she did not have it in the revelation game (see the Supplementary material for an example). Such out-of-equilibrium allocations play no role under private contracting, for they do not affect the behavior of the downstream principals. As discussed after Theorem 4, in this case, any deviation to a menu that could not be offered in the revelation game can be punished by having the agent behave as if the principal offered a direct mechanism that gives to each extended type a contract that is optimal for that extended type among those that are in the menu, *holding fixed the mechanisms offered downstream*. The possibility of holding fixed the mechanisms offered downstream however disappears in environments in which upstream mechanisms are observable. By implication, certain deviations to menus that could not be offered as direct mechanisms may be impossible to punish, thus making equilibrium outcomes in the revelation game not robust. However, this problem with direct mechanisms disappears when direct mechanisms “span” all possible menus³¹—a case that often arises in settings with a continuum of types. Because in this case $\Gamma^D \succneq \Gamma^M$, then all equilibrium outcomes of Γ^D are also equilibrium outcomes of Γ^M .

Endogenous sequence of bilateral relationships. Consider the following variant of our benchmark contracting game. There are $T > 2$ periods. In each period, *all* principals simultaneously offer the agent a mechanism and the agent chooses at most one mechanism to participate in. The agent may participate in a mechanism offered by the same principal multiple times. As in

³⁰ A similar idea has been explored in [14] in the context of a simultaneous multi-agent–multi-principal setting.

³¹ That is, when for any $\phi_i^M \in \Phi_i^M$, there exists a $\phi_i^D \in \Phi_i^D$ such that $\text{Im}(\phi_i^D) = \text{Im}(\phi_i^M)$.

the benchmark model, contracting is private so that the principals do not observe other principals' mechanisms, nor the messages, the contracts, or the decisions taken in these mechanisms.

This setting may correspond, for example, to an environment in which at each period, a consumer decides which seller to use for the provision of a certain good or service.³²

The difficulties with this extensive form come from the fact that the agent's continuation payoff at date t may now depend not only on the upstream payoff-relevant decisions but also on the mechanisms, the messages and the contracts selected upstream. In fact, such payoff-irrelevant information may determine which mechanisms will be offered downstream. Furthermore, because the principals observe some of the decisions taken upstream (e.g. the agent's participation decisions), the problems with out-of-equilibrium beliefs discussed above for the case of observable actions carry over to this environment.

These problems, however, disappear if one restricts attention to equilibria in which not only the agent's strategy but also the principals' strategies are Markov. All MPE outcomes can in fact be characterized with menus. They can also be characterized as truthful equilibria, but only if the agent is restricted to contract with each principal at most once. If instead the agent can contract with the same principal multiple times, then having the agent report truthfully is in general not without loss of generality. The reasons are the same as in the literature that assumes a single principal who lacks commitment. As shown in [6], one can safely restrict the principal to offer direct mechanisms but should not restrict the agent to report truthfully in each period. The same is true with multiple principals. One can characterize all MPE outcomes by having the principals compete in extended direct mechanisms, but should not assume the agent reports truthfully to each principal. Restricting attention to truthful equilibria is however fine if the agent does not possess any private information vis a vis the principal; this is the case when preferences are common knowledge (i.e. $|\Theta| = 1$) and when all players' strategies are pure as, for example, in [5].

Sequential offering as opposed to sequential contracting. Sequential offering refers to an environment in which principals offer their mechanisms sequentially, but the agent takes payoff-relevant decisions only after observing the mechanisms offered by all principals. This is in contrast to sequential contracting where some payoff-relevant decisions are committed at each date.³³

First, consider an environment similar to the one examined in the benchmark model but where the agent's effort is chosen only at the very end, say at $t = n + 1$. At $t = i$, for $i = 1, \dots, n$, P_i offers a mechanism $\phi_i : \mathcal{M}_i \rightarrow D_i$, where D_i is the set of feasible lotteries over the contracts $y_i : E \rightarrow \Delta(\mathcal{A}_i)$, with $e \in E$ now denoting some common effort.³⁴ After sending the message m_i , the agent may or may not observe the realization y_i of the lottery $\delta_i = \phi_i(m_i)$. As in the benchmark model, principals do not observe other principals' mechanisms, nor the messages the agent sends to other principals, nor the contracts selected in response to these messages.

It is immediate that this extensive form poses no problems to our characterization results. It suffices to adjust the definition of extended type to take into account that, because the decisions (e, a) are now chosen only at $t = n + 1$, at $t = i$, the component of the upstream history that is

³² This is the environment examined, for example, in [2,20,30].

³³ Sequential offering has been examined, among others, in [2,7,17,20,29,30,32].

³⁴ One may think of e as a vector of decisions $e = (e_1, \dots, e_n)$. Depending on the application of interest, the set of feasible contracts Y_i may be more or less restricted. For example, when e_j stands for a decision specific to the relationship between A and P_j , it may be reasonable to assume that y_i must not depend on e_j . This could be the case, for example, if e_j represents the quantity/quality of the product of seller j and if the latter could not be observed by P_i .

relevant for the agent's continuation payoff becomes the profile of contracts selected upstream along with the agent's exogenous type. That is, all the results in the benchmark model extend to this environment by letting $\theta_i^E \equiv (\theta, y_i^-)$, or $\theta_i^E \equiv (\theta, \delta_i^-)$, depending on whether at every $t < i$, the agent observes the contract y_t or only the lottery $\delta_t = \phi_t(m_t)$.

Next, consider an environment in which principals offer their mechanisms sequentially, but where the agent sends the messages (m_1, \dots, m_n) simultaneously at $t = n + 1$. Given the contracts $(y_i)_{i=1}^n$ selected by the lotteries $(\delta_i)_{i=1}^n = (\phi_i(m_i))_{i=1}^n$, the agent then chooses an action $e \in E$ and finally the principals' decisions $(a_i)_{i=1}^n$ are determined by the contracts $(y_i)_{i=1}^n$. If the principals do not observe other principals' mechanisms, this setting is strategically equivalent to simultaneous common agency; the menu theorems of [28] and [19] then guarantee that the entire set of equilibrium outcomes can be characterized by restricting the principals to offer menus. In many applications, it is however appealing to restrict attention to equilibria in which the decisions the agent induces with each principal P_i depend on the menu offered by P_i , the agent's type θ and the decisions δ_{-i} the agent is inducing with the other principals (but not the menus, or more generally the mechanisms, offered by the latter). Imposing this property is analog to restricting the agent's strategy to be Markov in a sequential contracting game. In [26], we show that any SCF that can be sustained by an equilibrium in which the agent's strategy satisfies the aforementioned Markov property can also be sustained as a truthful equilibrium of a game in which the principals offer incentive-compatible extended direct mechanisms. The definition of these mechanisms in simultaneous games is adjusted to take into account that decisions are determined simultaneously: instead of reporting the payoff-relevant decisions (e_i^-, a_i^-) contracted upstream, the agent is asked to report (in addition to θ) the lotteries δ_{-i} he is inducing with the other principals. The agent's strategy is then truthful if the decisions δ_{-i} he reports to P_i coincide with the true decisions $(\phi_j(m_j))_{j \neq i}$ induced with all principals other than i by sending the messages $m_{-i} = (m_j)_{j \neq i}$.³⁵

In the Supplementary material, we show that the same mechanisms also permit one to sustain all *pure-strategy* MPE outcomes in sequential games in which downstream principals observe the mechanisms offered upstream before choosing their own mechanism.³⁶ The reasons why extended direct mechanisms, or menus, fail to sustain certain mixed-strategy equilibrium outcomes are the same discussed for environments in which upstream mechanisms are publicly observable.

7. Concluding remarks

We discussed the intricacies that emerge in environments in which multiple principals contract sequentially and non-cooperatively with the same agent and provided characterization results useful for applications. Our benchmark model was one of private contracting in which downstream principals do not observe the mechanisms offered upstream, nor the decisions taken in these mechanisms. We also assumed that the sequence of bilateral relationships was exogenous and that a pair of payoff-relevant decisions, one for the principal, the other for the agent, was committed at each period.

³⁵ When the agent's strategy is not Markov, we show that (i) pure-strategy equilibrium outcomes can be characterized by having the agent report the identity of a deviating principal (if any) in addition to his type and the decisions he is inducing with the other principals, and (ii) that all mixed-strategy equilibrium outcomes can be characterized through incentive-compatible set-valued direct mechanisms in which the agent is offered multiple (payoff-equivalent) decisions to choose from as a function of his type and the decisions he is inducing with the other principals.

³⁶ Clearly, the same outcomes can also be characterized assuming the principals compete in menus.

For this environment, we first showed that all equilibrium outcomes can be characterized by assuming the principals compete in menus, thus proving that the menu theorems of simultaneous common agency extend to this environment. We then proceed by showing that, when lotteries are feasible, then all equilibrium outcomes can be sustained through pure-strategy profiles, as in games with a single mechanism designer (but not in simultaneous common agency). We also showed that when information is complete, any deterministic social choice function (that is, any outcome sustained by the agent not mixing over effort and the principals not mixing over their contracts) can always be sustained as a Markov-perfect equilibrium (that is, by restricting each player's strategy to depend only on payoff-relevant information).

Starting from these results, we then introduced a class of direct mechanisms in which the agent is asked to report his exogenous type along with the payoff-relevant decisions contracted upstream. We showed that all MPE outcomes of the menu game are also MPE outcomes in the game in which principals offer these direct mechanisms. The advantage of these mechanisms is that they permit one to use techniques from standard mechanism design (i.e. incentive-compatibility) to identify necessary and sufficient conditions for equilibrium outcomes. There are however two differences with respect to standard mechanism design. First, incentive compatibility must be established by backward induction: whether a mechanism is incentive-compatible or not depends on the mechanisms offered downstream. Second, a mechanism must specify incentive-compatible allocations also for extended types that have zero measure on the equilibrium path: this is because out-of-equilibrium allocations may be necessary to punish upstream deviations.

Finally, we discussed the problems with menus and extended direct mechanisms that emerge in environments in which downstream principals observe the mechanisms and/or the payoff-relevant decisions selected upstream, or when the sequence of bilateral relationships is endogenously determined by the agent's participation decisions. Building on the results for the benchmark model, we proposed solutions that, although do not always permit a complete equilibrium characterization, allow one to characterize the outcomes that are typically of interest for applications. While the various extensive forms considered here do not exhaust all the cases examined in the literature, we expect our results to be useful for many applications of sequential common agency.

Appendix A

Proof of Theorem 1. The proof is in two parts. Part 1 proves that for any equilibrium σ of Γ , there exists an equilibrium σ^M of Γ^M that implements the same outcomes. Part 2 proves the converse.

Part 1. Let \mathcal{Q}_i be a generic partition of Φ_i and denote by $Q_i \in \mathcal{Q}_i$ an element of \mathcal{Q}_i . Consider now a partition game $\Gamma^{\mathcal{Q}_i}$ in which at $t = i$, P_i chooses an element Q_i of \mathcal{Q}_i , then A selects a mechanism ϕ_i from \mathcal{Q}_i , sends a message $m_i \in \mathcal{M}_i$ to P_i and finally, given the realization y_i of the lottery $\delta_i = \phi_i(m_i)$, chooses effort e_i . At any other date $j \neq i$, both P_j and A have exactly the same choice sets as in Γ . Now, let Γ_i^M be a game with the same structure as Γ , except that at $t = i$, the strategy space for P_i is $\Delta(\Phi_i^M)$.

The proof of Part 1 is in three steps. Step 1 identifies a partition of Φ_i that makes the agent indifferent between any two mechanisms in the same cell $Q_i \in \mathcal{Q}_i$ and then constructs an equilibrium $\hat{\sigma}$ of $\Gamma^{\mathcal{Q}_i}$ that implements the same outcomes as σ . Step 2 uses the construction in Step 1 to derive an equilibrium $\hat{\sigma}$ in Γ_i^M which also implements the same outcomes as σ . Finally, Step 3 shows how the previous two steps can be applied recursively to construct an equilibrium σ^M for Γ^M that implements the same outcomes as σ .

Step 1. Start by considering a generic partition \mathcal{Q}_i of Φ_i that consists of measurable sets³⁷ and consider the following strategy profile for $\Gamma^{\mathcal{Q}_i}$. For P_i , let $\hat{\sigma}_i \in \Delta(\mathcal{Q}_i)$ be the probability measure over \mathcal{Q}_i generated by the original strategy σ_i in Γ . That is, for any subset R of \mathcal{Q}_i whose union is measurable,

$$\hat{\sigma}_i(R) = \sigma_i\left(\bigcup R\right).$$

For any P_j , with $j \neq i$, simply let $\hat{\sigma}_j = \sigma_j$. Next consider the agent. For any $t < i$ and for any h_t , let A 's strategy be the same as in σ , that is, $\hat{\sigma}_A(h_t) = \sigma_A(h_t)$. At $t = i$, given any history $h_i = (\phi_i^-, m_i^-, y_i^-, e_i^-, a_i^-, Q_i)$, A selects each mechanism from \mathcal{Q}_i using the regular conditional probability distribution $\sigma_i(\cdot | Q_i)$.³⁸ At any subsequent information set, A then simply behaves according to the original strategy σ_A , as if at the beginning of $t = i$, the history had been $h_i = (\phi_i^-, m_i^-, y_i^-, e_i^-, a_i^-, \phi_i)$.

Now, fix the agent's strategy $\hat{\sigma}_A$. Whatever the partition \mathcal{Q}_i , the principals' strategies described above constitute an equilibrium for the game $\Gamma^{\mathcal{Q}_i}(\hat{\sigma}_A)$ among the principals.

In the following, we identify a partition \mathcal{Q}_i that makes the strategy $\hat{\sigma}_A$ described above sequentially optimal for the agent. To this purpose, let \sim_i be the partition defined by the following equivalence relation. For any two mechanisms $\phi_i, \phi'_i \in \Phi_i$,

$$\phi_i \sim_i \phi'_i \iff \text{Im}(\phi_i) = \text{Im}(\phi'_i). \tag{3}$$

Clearly, the partition generated by (3) consists of measurable sets. It is also immediate that, in the partition game $\Gamma^{\mathcal{Q}_i}$, $\hat{\sigma}_A$ is a sequentially rational best response to the principals' strategy profile $(\hat{\sigma}_1, \dots, \hat{\sigma}_n)$. We conclude that, for any equilibrium σ of Γ , the partition game $\Gamma^{\mathcal{Q}_i}$ —where \mathcal{Q}_i is the partition generated by the equivalence relation \sim_i as given in (3)—admits an equilibrium $\hat{\sigma}$ that implements the same outcomes as σ .

Step 2. We now prove that starting from $\hat{\sigma}$, one can construct an equilibrium $\hat{\sigma}$ for the game Γ_i^M that implements the same outcomes. Start with P_i . Now let $\text{Im}(Q_i)$ denote the image of any of the mechanisms in Q_i , and for any $\phi_i^M \in \Phi_i^M$, let $Q_i(\phi_i^M) \in \mathcal{Q}_i$ denote the cell defined by

$$Q_i(\phi_i^M) \equiv \{\phi_i \in \Phi_i : \text{Im}(\phi_i) = \text{Im}(\phi_i^M)\}.$$

Then, for any measurable set $K \subseteq \Phi_i^M$, let

$$\hat{\sigma}_i(K) = \hat{\sigma}_i(\tilde{Q}_i),$$

where $\tilde{Q}_i \equiv \{Q_i \in \mathcal{Q}_i : \text{Im}(Q_i) = \text{Im}(\phi_i^M) \text{ for some } \phi_i^M \in K\}$. For all principals P_j with $j \neq i$, simply let $\hat{\sigma}_j = \hat{\sigma}_j$. Next, consider the agent. At any $t < i$, $\hat{\sigma}_A(h_t) = \hat{\sigma}_A(h_t)$ for any h_t . Starting from $t = i$, for any $\phi_i^M \in \Phi_i^M$, at any subsequent information set, A then induces the same outcomes he would have induced in $\Gamma^{\mathcal{Q}_i}$ had P_i offered the cell $Q_i(\phi_i^M)$. Formally, let

$$\xi(h_i^-, Q_i) \equiv \int_{\Phi_i} \int_{\mathcal{M}_i} \phi_i(m_i) d\mu(h_i^-, \phi_i) d\sigma_i(\phi_i | Q_i)$$

³⁷ In the sequel, we assume that Φ_i is a Polish space and whenever we talk about measurability, we mean with respect to the Borel σ -algebra Σ on Φ_i .

³⁸ Assuming that Φ_i is a Polish space endowed with the Borel σ -algebra Σ , the existence and (almost) uniqueness of such a conditional probability distribution follows from standard results (e.g. in [10, Theorem 10.2.2, p. 345]).

denote the distribution over Y_i that A would have induced in $\Gamma^{\mathcal{Q}_i}$ by following the strategy $\hat{\sigma}_A$, given (h_i^-, Q_i) . Next, let $\mathcal{M}_i^{\Phi_i}$ be the union of all the messages that A can send to P_i in Γ and for any $y_i \in \text{Supp}[\zeta(h_i^-, Q_i)]$, let $\beta(y_i; h_i^-, Q_i) \in \Delta(\Phi_i \times \mathcal{M}_i^{\Phi_i})$ denote the conditional distribution over $\Phi_i \times \mathcal{M}_i^{\Phi_i}$ that is obtained from Bayes' rule using the strategy $\hat{\sigma}_A$, conditioning on the event that the contract selected at $t = i$ is y_i and the upstream history at the beginning of period i is (h_i^-, Q_i) . Then in Γ_i^M , for any (h_i^-, ϕ_i^M) , the agent's mixed strategy $\hat{\mu}(h_i^-, \phi_i^M) \in \mathcal{M}_i^M$ over the messages in ϕ_i^M induces the same distribution $\zeta(h_i^-, Q_i(\phi_i^M))$ over the set of contracts Y_i as the strategy $\hat{\sigma}_A(h_i^-, Q_i(\phi_i^M))$ in $\Gamma^{\mathcal{Q}_i}$. In the continuation game that starts after the realization of the contract y_i , A then uses the conditional distribution $\beta(y_i; h_i^-, Q_i(\phi_i^M))$ to determine his downstream behavior. That is, at any downstream information set, A behaves according to the strategy $\hat{\sigma}_A$ as if in $\Gamma^{\mathcal{Q}_i}$, A selected the mechanism ϕ_i from $Q_i(\phi_i^M)$ and the message m_i from $\mathcal{M}_i^{\Phi_i}$.

The strategy profile $\hat{\sigma}$ constructed this way is clearly an equilibrium for Γ_i^M and induces the same outcomes as σ in Γ .

Step 3. Since in the construction of the equilibrium $\hat{\sigma}$ for Γ_i^M , σ_i^- is kept fixed, Steps 1 and 2 can be applied recursively starting from $t = 1$ and proceeding forward to construct an equilibrium σ^M for Γ^M that implements the same outcomes as σ .

Part 2. We now prove that, given any equilibrium σ^M of Γ^M , there exists an equilibrium σ of Γ that implements the same outcomes and such that σ_A is an extension of σ_A^M over Γ .

First consider the principals. For any P_i , simply let $\sigma_i = \alpha_i(\sigma_i^M)$, where $\alpha_i(\sigma_i^M)$ is the distribution over Φ_i obtained from σ_i^M using the embedding α_i .

Next consider the agent. After any history $h_t = ((\phi_j, m_j, y_j, e_j, a_j)_{j=1}^{t-1}, \phi_t)$, the agent behaves according to σ_A^M as if the game were Γ^M and the history were $h_t^M = ((\phi_j^M, \delta_j, y_j, e_j, a_j)_{j=1}^{t-1}, \phi_t^M)$ where the history h_t^M is obtained from h_t replacing each pair (ϕ_j, m_j) with (ϕ_j^M, δ_j) where ϕ_j^M is the menu whose image is $\text{Im}(\phi_j^M) = \text{Im}(\phi_j)$ and $\delta_j = \phi_j(m_j)$. That is, for any measurable set of messages $M_t \subseteq \mathcal{M}_t$ in ϕ_t , the strategy $\sigma_A(h_t)$ is such that

$$\mu(M_t | h_t) = \mu(\beta(M_t) | h_t^M)$$

where $\mu(\cdot | h_t)$ and $\mu(\cdot | h_t^M)$ denote the distributions over \mathcal{M}_t and \mathcal{M}_t^M , respectively in Γ and in Γ^M , and

$$\beta(M_t) \equiv \{\delta_t: \delta_t = \phi_t(m_t), m_t \in M_t\}.$$

Given the realization y_t of the lottery $\phi_t(m_t)$, A then induces the same distribution over E_t that he would have induced in Γ^M given $(h_t^M, \phi_t(m_t), y_t)$.

The strategy σ_A constructed this way is an extension of σ_A^M over Γ . Furthermore, when A follows the strategy σ_A , no principal has incentive to deviate from $\sigma_i = \alpha_i(\sigma_i^M)$. We conclude that σ is an equilibrium of Γ and implements the same outcomes as σ^M . \square

Proof of Theorem 2. We want to show that, when $D_i = \Delta(Y_i)$ for all i , then for any $\sigma^M \in \mathcal{E}(\Gamma^M)$ there exists a pure-strategy equilibrium $\hat{\sigma}^M \in \mathcal{E}(\Gamma^M)$ that implements the same outcomes as σ^M .

Suppose $\sigma^M \in \mathcal{E}(\Gamma^M)$ is such that P_i mixes over different menus and/or, given some menu ϕ_i^M , the agent randomizes over the lotteries in ϕ_i^M . We prove that there exists another $\hat{\sigma}^M \in \mathcal{E}(\Gamma^M)$ in which all P_j with $j \neq i$ follow the strategy $\hat{\sigma}_j^M = \sigma_j^M$, P_i offers only one menu, $\hat{\phi}_i^M$, and, after any (h_i^-, ϕ_i^M) , there is a single lottery $\delta_i(h_i^-, \phi_i^M)$ that the agent selects from ϕ_i^M .

The equilibrium menu $\hat{\phi}_i^M$ is such that

$$\text{Im}(\hat{\phi}_i^M) = \text{cl} \left\{ \delta_i : \delta_i = \int_{\phi_i^M \in \Phi_i^M} \int_{\delta_i \in \mathcal{M}_i^M} \delta_i d\mu(h_i^-, \phi_i^M) d\sigma_i^M \text{ for some } h_i^- \in \mathcal{H}_i^- \right\}, \quad (4)$$

where $\text{cl}(X)$ denotes the closure of the set X . The menu $\hat{\phi}_i^M$ is thus the (closure of the) set of lotteries that can be constructed by mixing with distribution σ_i^M over the different lotteries that A would have induced in each menu ϕ_i^M , for some upstream history $h_i^- \in \mathcal{H}_i^-$.

Next, consider the following strategy for the agent. For any $t < i$ and any h_t , $\hat{\sigma}_A^M(h_t) = \sigma_A^M(h_t)$. At $t = i$, given h_i^- , if $\phi_i^M = \hat{\phi}_i^M$, then A selects with probability one the lottery

$$\delta_i(h_i^-, \hat{\phi}_i^M) \equiv \int_{\phi_i^M \in \Phi_i^M} \int_{\delta_i \in \mathcal{M}_i^M} \delta_i d\mu(h_i^-, \phi_i^M) d\sigma_i^M.$$

Note that, for any h_i^- , the distribution over Y_i induced by $\delta_i(h_i^-, \hat{\phi}_i^M)$ is the same as that induced by σ_i^M and σ_A^M in the original equilibrium σ^M . Now for any $y_i \in \text{Supp}[\delta_i(h_i^-, \hat{\phi}_i^M)]$, let $\lambda(y_i; h_i^-) \in \Delta(\Phi_i^M \times \Delta(Y_i))$ denote the joint distribution over $\Phi_i^M \times \Delta(Y_i)$ that is obtained from Bayes' rule conditioning on the event that the contract selected by P_i is y_i and using the original strategies σ_i^M and σ_A^M . Then, in the continuation game that starts after the realization of the contract y_i , A uses the distribution $\lambda(y_i; h_i^-)$ to determine his downstream behavior. That is, A behaves in any downstream information set according to the original strategy σ_A^M , as if P_i offered the menu ϕ_i^M and A selected the message $\delta_i \in \Delta(Y_i)$. Conditional on P_i offering $\hat{\phi}_i^M$, the distribution over the payoff-relevant decisions $(e_j, a_j)_{j=1}^n$ is then the same as in the original equilibrium σ^M .

If instead $\phi_i^M \neq \hat{\phi}_i^M$, then the particular lottery $\delta_i(h_i^-, \phi_i^M)$ that A selects from ϕ_i^M is any lottery $\delta_i \in \text{Supp}[\sigma_A^M(h_i^-, \phi_i^M)]$ that minimizes P_i 's expected payoff taking into account that at any subsequent information set the agent's behavior will be determined by the original strategy σ_A^M . Given $(\hat{\sigma}_i^M)_{i=1}^n$, the strategy $\hat{\sigma}_A^M$ is clearly sequentially optimal for the agent. Furthermore, given $(\hat{\sigma}_A^M, \hat{\sigma}_{-i}^M)$, no principal has an incentive to deviate from $\hat{\sigma}_i^M$. We conclude that $\hat{\sigma}^M$ is an equilibrium for Γ^M and induces the same outcomes as σ^M .

Iterating for all $i = 1, \dots, n$, starting from $i = 1$ and proceeding forward then gives the result. \square

Proof of Theorem 3. Assume $|\Theta| = 1$ and let $(e^*, a^*) \equiv (e_i^*, a_i^*)_{i=1}^n \in E \times \mathcal{A}$ denote the equilibrium decisions. We prove that if there exists a $\sigma^{M*} \in \mathcal{E}(\Gamma^M)$ that implements (e^*, a^*) , there also exists a pure-strategy equilibrium $\tilde{\sigma}^M \in \mathcal{E}(\Gamma^M)$ which implements the same outcomes and such that the agent's strategy is Markov at any $t \in \mathcal{N}$.

First, note that, because (e^*, a^*) is deterministic, for any mixed-strategy equilibrium that sustains (e^*, a^*) there exists a pure-strategy equilibrium that sustains the same outcomes. This is true even if $D_i \subsetneq \Delta(Y_i)$ for some i . Hence, without loss of generality, assume σ^{M*} is in pure strategies. Given σ^{M*} , let $\theta_i^{E*} \equiv (e_i^{E*}, a_i^{E*})$ denote the equilibrium extended type for $t = i$ and denote by $\mathcal{H}_i^-(\theta_i^{E*}; \sigma^{M*})$ the set of all possible equilibrium upstream histories that lead to θ_i^{E*} . Note that even if the SCF is deterministic and σ^{M*} is a pure-strategy profile, $\mathcal{H}_i^-(\theta_i^{E*}; \sigma^{M*})$ need not be a singleton.

Next, consider the strategy profile $\tilde{\sigma}^M$ in which all principals offer the same equilibrium menus as in σ^{M*} and in which the agent's (Markov) strategy $\tilde{\sigma}_A^M$ is constructed from σ_A^{M*} as follows.

Start from $t = n$. First suppose P_n offers the equilibrium menu ϕ_n^{M*} . If h_n^- is such that $\theta_n^E = \theta_n^{E*}$, that is, if the decisions taken upstream are the equilibrium decisions, then irrespective of which particular upstream history led to (e_n^-, a_n^-) , A always selects the same lottery $\delta_n(\theta_n^{E*}, \phi_n^{M*})$, where $\delta_n(\theta_n^{E*}, \phi_n^{M*})$ is any lottery that A would have selected in ϕ_n^{M*} after some history $h_n^- \in \mathcal{H}_n^-(\theta_n^{E*}; \sigma^{M*})$. After any of the contracts $y_n \in \text{Supp}[\delta_n(\theta_n^{E*}, \phi_n^{M*})]$ is realized, A then chooses the equilibrium effort e_n^* leading to the equilibrium decision a_n^* . Clearly (e_n^*, a_n^*) is the same no matter which particular equilibrium history $h_n^- \in \mathcal{H}_n^-(\theta_n^{E*}; \sigma^{M*})$ one considers and which particular contract y_n is realized.

If, instead, $\theta_n^E \notin \theta_n^{E*}$, that is, if the decisions taken upstream are different from the equilibrium decisions, then let $j \leq n - 1$ be the first date at which a departure from equilibrium occurred—that is, the unique date $j \leq n - 1$ such that $\theta_j^E = \theta_j^{E*}$ and $\theta_{j+1}^E \neq \theta_{j+1}^{E*}$. In this case, A selects a lottery $\delta_n(\theta_n^E, \phi_n^{M*})$, followed by the effort strategy $\xi(j) : Y_n \rightarrow \Delta(E_n)$, that minimizes P_j 's payoff among those that are sequentially optimal for A given θ_n^E . Neither $\delta_n(\theta_n^E, \phi_n^{M*})$ nor $\xi(j)$ depend on the particular upstream history that led to θ_n^E .

Next, suppose P_n offers a menu $\phi_n^M \neq \phi_n^{M*}$. Then let $\tilde{\sigma}_A^M(\theta_n^E, \phi_n^M)$ be any Markov strategy that is sequentially optimal for A and that minimizes the payoff of P_n .

Next move to $t = n - 1$ and construct the agent's Markov strategy $\tilde{\sigma}_A^M(\theta_{n-1}^E, \phi_{n-1}^M)$ following the same steps as for $t = n$ assuming that at $t = n$, P_n offers ϕ_n^{M*} and that A follows $\tilde{\sigma}_A^M(\theta_n^E, \phi_n^M)$.

Iterating backwards up to $t = 1$ gives a strategy $\tilde{\sigma}^M$ in which the agent's strategy is Markov at any date. It is immediate that, given $(\tilde{\sigma}_i^M)_{i=1}^n$, the strategy $\tilde{\sigma}_A^M$ is sequentially optimal for the agent. Furthermore, given $(\tilde{\sigma}_A^M, \tilde{\sigma}_{-i}^M)$ no principal has an incentive to deviate from $\tilde{\sigma}_i^M$. We conclude that $\tilde{\sigma}^M$ is an equilibrium for Γ^M and induces the same outcomes as σ^{M*} . \square

Proof of Theorem 4. Part I. We want to show that any SCF π that can be sustained as a MPE of Γ^M can also be sustained as a pure-strategy truthful MPE of Γ^D .

Let σ^M denote the equilibrium strategy profile that sustains π in Γ^M ; in virtue of Theorem 2, without loss of generality, we can assume that σ^M is a pure-strategy profile and denote by ϕ_t^{M*} the equilibrium menu for P_t .³⁹ That σ_A^M is Markov implies that, for any (θ_t^E, ϕ_t^M) , there is a single lottery $\delta_t(\theta_t^E, \phi_t^M)$ such that, whatever the particular upstream history h_t^- that has conducted to θ_t^E , A always chooses $\delta_t(\theta_t^E, \phi_t^M)$ from ϕ_t^M when his extended type is θ_t^E . We henceforth denote the agent's behavioral strategy in period t by $\sigma_A^M(\theta_t^E, \phi_t^M)$. To prove that there exists a pure-strategy truthful MPE of Γ^D that sustains the same outcomes as σ^M , we then proceed in two steps.

Step 1. Consider the game Γ_t in which $\Phi_i = \Phi_i^M$ for all $i \neq t$, whereas for $i = t$, $\Phi_i = \Phi_i^D$. The following is then an equilibrium for Γ_t . For any $i \neq t$, $\sigma_i = \sigma_i^M$, whereas for $i = t$, σ_i is the (pure) strategy that consists in offering the mechanism ϕ_t^{D*} given by

$$\phi_t^{D*}(\theta_t^E) = \delta_t(\theta_t^E, \phi_t^{M*}) \quad \forall \theta_t^E.$$

³⁹ Note that the Markov property of the agent's strategy is preserved by the replication arguments in the proof of Theorem 2. That is, for any $\sigma^M \in \mathcal{E}(\Gamma^M)$ in which the agent's strategy is Markov, there exists a pure-strategy profile $\hat{\sigma}^M \in \mathcal{E}(\Gamma^M)$ that sustains the same outcomes as σ^M and such that $\hat{\sigma}_A^M$ is Markov.

At any date $i \neq t$, the agent's (Markov) strategy is the same as in Γ^M , i.e. $\sigma_A(\theta_i^E, \phi_i^M) = \sigma_A^M(\theta_i^E, \phi_i^M)$. In period t , if $\phi_t^D = \phi_t^{D*}$, then A reports θ_t^E truthfully and then selects effort according to σ_A^M as if the game were Γ^M , the menu offered by P_t were ϕ_t^{M*} and the decision selected with P_t were $\delta_t(\theta_t^E, \phi_t^{M*})$. If instead $\phi_t^D \neq \phi_t^{D*}$, then the agent behaves according to σ_A^M as if the game were Γ^M and the menu offered by P_t were ϕ_t^M , where $\text{Im}(\phi_t^M) = \text{Im}(\phi_t^D)$. It is immediate that σ is a (pure-strategy) MPE for Γ_t , it sustains the same outcomes as σ^M and is such that (a) ϕ_t^{D*} is incentive-compatible and (b) the agent's strategy is truthful at period t .

Step 2. Following the same arguments as in Step 1, one can easily see that, starting from any game Γ_J in which $\Phi_i = \Phi_i^D$ for any $i \in J \subset \mathcal{N}$ while $\Phi_i = \Phi_i^M$ for $i \in \mathcal{N} \setminus J$, and from any (pure-strategy) MPE $\sigma \in \mathcal{E}(\Gamma_J)$ that sustains the same outcomes as σ^M and satisfies (a) and (b) for any $i \in J$, there exists a game $\Gamma_{J'}$ in which $\Phi_i = \Phi_i^D$ for any $i \in J'$ and $\Phi_i = \Phi_i^M$ for all $i \in \mathcal{N} \setminus J'$ with $J' = J \cup \{t\}$ for some $t \in \mathcal{N} \setminus J$ and a (pure-strategy) MPE $\sigma' \in \mathcal{E}(\Gamma_{J'})$ that sustains the same outcomes as σ^M and satisfies (a) and (b) for all $i \in J'$. Combining Step 1 and Step 2 gives the result.

Part II. We now show that for any $\sigma^D \in \mathcal{E}(\Gamma^D)$ there exists a $\sigma^M \in \mathcal{E}(\Gamma^M)$ that sustains the same outcomes.

Let Γ_J denote the game in which $\Phi_j = \Phi_j^M$ for all $j \in J$ while $\Phi_j = \Phi_j^D$ for all $j \in \mathcal{N} \setminus J$, for some $J \subset \mathcal{N} \cup \{\emptyset\}$. We want to show that given any equilibrium $\sigma \in \mathcal{E}(\Gamma_J)$, there exists an equilibrium $\tilde{\sigma} \in \mathcal{E}(\Gamma_{J'})$, with $J' = J \cup \{t\}$ for some $t \in \mathcal{N} \setminus J$, that sustains the same outcomes. The strategy profile $\tilde{\sigma}$ is constructed from σ as follows. For any $i \neq t$, $\tilde{\sigma}_i = \sigma_i$. For $i = t$, $\tilde{\sigma}_t$ is such that, for any measurable set $R \subseteq \Phi_t^M$,

$$\tilde{\sigma}_t(R) = \sigma_t\left(\bigcup_{\phi_t^M \in R} \Phi_t^D(\phi_t^M)\right)$$

where $\Phi_t^D(\phi_t^M) \equiv \{\phi_t^D : \text{Im}(\phi_t^D) = \text{Im}(\phi_t^M)\}$. Next, consider the agent. Let

$$\bar{\Phi}_t^M \equiv \{\phi_t^M : \text{Im}(\phi_t^M) = \text{Im}(\phi_t^D) \text{ for some } \phi_t^D \in \Phi_t^D\}.$$

At any $i < t$, $\tilde{\sigma}_A(h_i) = \sigma_A(h_i)$ for any h_i . Starting from $i = t$, for any $\phi_t^M \in \bar{\Phi}_t^M$, at any subsequent information set, A induces the same outcomes he would have induced in Γ_J had P_t offered one of the mechanisms in $\Phi_t^D(\phi_t^M)$. Formally, let

$$\zeta(h_t^-, \Phi_t^D(\phi_t^M)) \equiv \int_{\Phi_t^D} \int_{\Theta_t^E} \phi_t^D(\theta_t^E) d\mu(h_t^-, \phi_t^D) d\sigma_t(\phi_t^D | \Phi_t^D(\phi_t^M))$$

denote the distribution over Y_t generated in Γ_J by σ_A and σ_t , conditional on $(h_t^-, \Phi_t^D(\phi_t^M))$. Next, for any $y_t \in \text{Supp}[\zeta(h_t^-, \Phi_t^D(\phi_t^M))]$, let $\beta(y_t; h_t^-, \Phi_t^D(\phi_t^M)) \in \Delta(\Phi_t^D \times \Theta_t^E)$ denote the conditional distribution over $\Phi_t^D \times \Theta_t^E$ that is obtained from Bayes' rule using the strategies σ_t and σ_A , conditioning on the event that the contract selected in period t is y_t , that the upstream history at the beginning of period t is h_t^- , and that the mechanism selected by P_t belongs to $\Phi_t^D(\phi_t^M)$. Then in $\Gamma_{J'}$, for any (h_t^-, ϕ_t^M) with $\phi_t^M \in \bar{\Phi}_t^M$, the agent's mixed strategy $\tilde{\mu}(h_t^-, \phi_t^M) \in \mathcal{M}_t^M$ over the messages in ϕ_t^M induces the same distribution $\zeta(h_t^-, \Phi_t^D(\phi_t^M))$ over the set of contracts Y_t as the strategy σ_A in Γ_J given $(h_t^-, \Phi_t^D(\phi_t^M))$. In the continuation game that starts after the realization of the contract y_t , A then uses the conditional distribution $\beta(y_t; h_t^-, \Phi_t^D(\phi_t^M))$ to determine his downstream behavior. That is, at any downstream information set, A behaves according to the strategy σ_A as if the game were Γ_J , P_t selected the mechanism ϕ_t^D from $\Phi_t^D(\phi_t^M)$ and A sent the message θ_t^E to P_t .

If instead $\phi_t^M \notin \bar{\Phi}_t^M$, then starting from $i = t$, at any downstream information set A behaves according to σ_A (in the same sense as defined above) as if the game were Γ_J and the mechanism offered by P_t were ϕ_t^D , where ϕ_t^D is obtained from ϕ_t^M as follows:

$$\phi_t^D(\theta_t^E) \in \arg \max_{\delta_t \in \text{Im}(\phi_t^M)} V(\theta_t^E, \delta_t, \sigma_t^+) \quad \forall \theta_t^E \in \Theta_t^E.$$

Given $(\tilde{\sigma}_i)_{i=1}^n$, the strategy $\tilde{\sigma}_A$ is sequentially optimal for the agent. Furthermore, given $(\tilde{\sigma}_A^M, \tilde{\sigma}_{-i}^M)$ no principal has an incentive to deviate from $\tilde{\sigma}_i$. It follows that $\tilde{\sigma} \in \mathcal{E}(\Gamma_{J'})$. Because $\tilde{\sigma}$ sustains the same outcomes as σ , the result then follows by iterating across periods, starting from $t = 1$ and proceeding forward by letting $J' = J \cup \{t + 1\}$. \square

Acknowledgments

For discussions and useful comments, we thank seminar participants at various institutions and conferences where this paper has been presented. A special thank is for Eddie Dekel, Mike Peters, Marciano Siniscalchi, Jean Tirole, an associate editor and four anonymous referees for suggestions that helped us improve the paper. We are also thankful to Itai Sher for excellent research assistance.

Supplementary material

The online version of this article contains the formal statements and the proofs of the results in Section 6.

Please visit DOI: 10.1016/j.jet2008.07.003.

References

- [1] A. Acquisti, H.R. Varian, Conditioning prices on purchase history, *Marketing Sci.* 24 (3) (2005) 1–15.
- [2] P. Aghion, P. Bolton, Contracts as a barrier to entry, *Amer. Econ. Rev.* 77 (1987) 388–401.
- [3] D. Baron, Non-cooperative regulation of a non-localized externality, *RAND J. Econ.* 16 (1985) 553–568.
- [4] G. Bellettini, G. Ottaviano, Special interests and technological change, *Rev. Econ. Stud.* 72 (2005) 43–56.
- [5] D. Bergemann, J. Välimäki, Dynamic common agency, *J. Econ. Theory* 111 (1) (2003) 23–48.
- [6] H. Bester, R. Strausz, Imperfect commitment and the revelation principle: The single agent case, *Econometrica* 69 (4) (2001) 1077–1098.
- [7] D. Bizer, P. DeMarzo, Sequential banking, *J. Polit. Economy* 100 (1) (1992) 41–61.
- [8] G. Calzolari, A. Pavan, On the optimality of privacy in sequential contracting, *J. Econ. Theory* 130 (1) (2006) 168–204.
- [9] G. Calzolari, A. Pavan, On the use of menus in sequential common agency, *Games Econ. Behav.* 64 (1) (2008) 329–334.
- [10] R.M. Dudley, *Real Analysis and Probability*, Cambridge Stud. Adv. Math., vol. 74, Cambridge University Press, 2002.
- [11] L. Epstein, M. Peters, A revelation principle for competing mechanisms, *J. Econ. Theory* 88 (1999) 119–160.
- [12] A. Faure-Grimaud, D. Martimort, Regulatory inertia, *RAND J. Econ.* 34 (2003) 437–413.
- [13] G. Grossman, E. Helpman, Intergenerational redistribution with short-lived governments, *Econ. J.* 108 (1998) 1299–1329.
- [14] S. Han, Menu theorems for bilateral contracting, *J. Econ. Theory* 127 (1) (2006) 157–178.
- [15] A.K. Koch, E. Peyrache, Aligning ambition and incentives: Optimal contracts with career concerns, Mimeo, University of Bonn, 2003.
- [16] T. Lewis, D. Sappington, Oversight of long-term investment by short-lived regulators, *Int. Econ. Rev.* 32 (1991) 579–600.
- [17] D. Martimort, The multiprincipal nature of the government, *Europ. Econ. Rev.* 40 (1996) 673–685.

- [18] D. Martimort, L. Stole, Communication spaces, equilibria sets and the revelation principle under common agency, Mimeo, University of Toulouse, 1997.
- [19] D. Martimort, L. Stole, The revelation and delegation principles in common agency games, *Econometrica* 70 (2002) 1659–1674.
- [20] L.M. Marx, G. Shaffer, Predatory accommodation: Below-cost pricing without exclusion in intermediate goods markets, *RAND J. Econ.* 30 (1999) 22–43.
- [21] L.M. Marx, G. Shaffer, Bargaining power in sequential contracting, Mimeo, University of Rochester, 2004.
- [22] P. McAfee, Mechanism design by competing sellers, *Econometrica* 61 (6) (1993) 1281–1312.
- [23] A. Mukherjee, Career concerns, matching, and optimal disclosure policy, *Int. Econ. Rev.* 49 (2008) 1211–1250.
- [24] T. Olsen, G. Torsvik, Intertemporal common agency and organizational design: How much centralization, *Europ. Econ. Rev.* 39 (1995) 1405–1428.
- [25] M.V. Pauly, Overinsurance and public provision of insurance: The roles of moral hazard and adverse selection, *Quart. J. Econ.* 88 (1) (1974) 44–66.
- [26] A. Pavan, G. Calzolari, Truthful revelation mechanisms for simultaneous common agency games, Mimeo, Northwestern University and University of Bologna, 2008.
- [27] J. Peck, A note on competing mechanisms and the revelation principle, Mimeo, Ohio State University, 1997.
- [28] M. Peters, Common agency and the revelation principle, *Econometrica* 69 (2001) 1349–1372.
- [29] A. Prat, A. Rustichini, Sequential common agency, Discussion Paper 9895, Center for Economic Research, Tilburg University, 1998.
- [30] K. Spier, M.D. Whinston, On the efficiency of privately stipulated damages for breach of contract: Entry barriers, reliance, and renegotiation, *RAND J. Econ.* 26 (1995) 180–202.
- [31] C.R. Taylor, Consumer privacy and the market for customer information, *RAND J. Econ.* 35 (4) (2004) 631–650.
- [32] M. Tommasi, F. Weinschelbaum, The threat of insurance. On the robustness of principal–agent models, *J. Inst. Theoretical Econ.* 163 (3) (2007) 379–393.