Supplementary Material for
Sequential Contracting with Multiple Principals

Alessandro Pavan
Northwestern University

Giacomo Calzolari
University of Bologna

August 1, 2008

Abstract

This note contains additional material omitted in the paper. Section A1 contains an example that illustrates how extended direct mechanisms can be put to work to identify necessary and sufficient conditions for the sustainability of an outcome as a sequential common agency equilibrium. Section A2 contains the formal statements (with corresponding proofs) of the results discussed in Section 6 in the paper (Alternative extensive forms).

A1. Extended direct mechanisms: A buyer-sellers example

Consider a private contracting environment in which two sellers, P_1 and P_2, contract sequentially with a common buyer, A. The buyer is interested in purchasing two complementary products, one from each seller. An action $a_i = (s_i, t_i) \in \mathcal{A}_i = \{0,1\} \times \mathbb{R}$ for P_i thus consists of a decision to trade s_i along with a transfer t_i, with $s_i = 0$ in the case of no trade and $s_i = 1$ in the case of trade.[1] The buyer's preferences are described by the quasilinear function $v(a, \theta) = \theta(s_1 + s_2) + s_1 s_2 - t_1 - t_2$ where $\Theta = \{\underline{\theta}, \overline{\theta}\}$ with $\underline{\theta} > 1$ and $\Delta\theta \equiv (\overline{\theta} - \underline{\theta}) \in (0,1)$. The probability the buyer is a high type is $\Pr(\overline{\theta}) = p$. The sellers' payoffs are given by $u_i(a) = t_i - s_i$. It is common knowledge that the buyer contracts first with P_1 and then with P_2 (think of P_1 as a hardware supplier and of P_2 as a software provider). We assume that the buyer's participation in either relationship is voluntary and that the buyer can contract with P_2 after rejecting P_1's offer. In case the buyer rejects P_i's offer, the default contract $(0,0)$ with no trade and zero transfer is implemented.

In this setting, it seems reasonable to assume that the agent's behavior with P_2 depends only on the payoff-relevant decisions contracted upstream and not on things such as the mechanism used upstream or the message sent in this mechanism. We thus assume the agent's strategy is Markov.

Now suppose one is interested in understanding which SCFs $\pi : \Theta \to \Delta(\{0,1\}^2 \times \mathbb{R}^2)$ can be sustained as MPE when principals can offer any lottery over $\{0,1\} \times \mathbb{R}$. It then suffices to proceed as follows.[2]

First, consider downstream contracting. Because preferences are quasilinear, the transfer t_1 has no effect on the agent's preferences over $\mathcal{A}_2$. Without loss, we can thus simplify and let $\Theta_2^E = \Theta \times \{0,1\}$, with $\theta_2^{E,1} \equiv \underline{\theta}$, $\theta_2^{E,2} \equiv \overline{\theta}$, $\theta_2^{E,3} \equiv \underline{\theta} + 1$ and $\theta_2^{E,4} \equiv \overline{\theta} + 1$. Furthermore, because P_2 never finds it optimal to introduce randomizations over the decision to trade, we can restrict attention to deterministic extended direct mechanisms $\phi_2^D : \Theta_2^E \to \{0,1\} \times \mathbb{R}$, with $s_2(\theta_2^{E,i}) = s_2^i$

[1] In this representation, there is no effort, i.e. $|E_1| = |E_2| = 1$. Alternatively, one could assume that $E_i = \{0,1\}$ and that $\mathcal{A}_i = \mathbb{R}$ for each i. In this case, a contract $y_i : \{0,1\} \longrightarrow \mathbb{R}$ specifies a price for each decision $s_i \in \{0,1\}$. The two representations are clearly equivalent.

[2] In this example, we are restricing attention to MPE, but we not imposing any restriction on the set of feasible lotteries D_i. The approach illustrated here clearly applies also to environments where principals are restricted to offer deterministic contracts.

and $t_2(\theta_2^{E,i}) = t_2^i$ denoting respectively the decision to trade and the price asked to the extended type $\theta_2^{E,i}$, $i = 1, ..., 4$.

Now let β^i denote the (posterior) probability that P_2 assigns to $\theta_2^{E,i}$, with $\beta \equiv \left(\beta^1, \beta^2, \beta^3, \beta^4\right)$. Each β^i is derived from ϕ_1^D using Bayes' rule. With a slight abuse of notation, let $\delta_1(\theta) = \Pr(s_1 = 1 \mid \theta)$. We then have that $\beta^1 = (1-p)[1-\delta_1(\underline{\theta})]$, $\beta^2 = p[1-\delta_1(\overline{\theta})]$, $\beta^3 = (1-p)\delta_1(\underline{\theta})$ and $\beta^4 = p\delta_1(\overline{\theta})$. From standard results in contract theory (e.g. Maskin and Riley, 1986) we know that, in any optimal mechanism for P_2 the decision to trade is monotonic in θ_2^E so that $s_2^i \leq s_2^{i+1}$ $i = 1, 2, 3$, that all downward adjacent incentive compatibility constraints bind so that

$$\theta_2^{E,i} s_2^i - t_2^i = \theta_2^{E,i} s_2^{i-1} - t_2^{i-1}, \;\; i = 2, 3, 4, \tag{1}$$

and that the participation constraint for the lowest type binds so that $\underline{\theta} s_2^1 - t_2^1 = 0$.[3] Substituting the transfers

$$\begin{array}{l} t_2^1 = \underline{\theta} s_2^1, \quad t_2^2 = \overline{\theta} s_2^2 - \Delta\theta s_2^1, \quad t_2^3 = (\underline{\theta}+1) s_2^3 - (1-\Delta\theta) s_2^2 - \Delta\theta s_2^1 \\ t_2^4 = (\overline{\theta}+1) s_2^4 - \Delta\theta s_2^3 - (1-\Delta\theta) s_2^2 - \Delta\theta s_2^1 \end{array} \tag{2}$$

into P_2's payoff, we have that

$$U_2 = \sum_{i=1}^{4} \mathcal{W}^i(\delta_1) s_2^i \tag{3}$$

where $\mathcal{W}^i(\delta_1)$ denotes the virtual surplus of selling to type i, given the upstream decisions $\delta_1 \equiv (\delta_1(\overline{\theta}), \delta_1(\underline{\theta}))$:

$$\begin{aligned} \mathcal{W}^1 &\equiv \beta^1(\underline{\theta} - 1) - (1-\beta_1)\,\Delta\theta \\ \mathcal{W}^2 &\equiv \beta^2\left(\overline{\theta} - 1\right) - \left(\beta^3 + \beta^4\right)(1 - \Delta\theta) \\ \mathcal{W}^3 &\equiv \beta^3\underline{\theta} - \beta^4\Delta\theta \\ \mathcal{W}^4 &\equiv \beta^4\overline{\theta}, \end{aligned}$$

with $\beta^i = \beta^i(\delta_1)$. A mechanism ϕ_2^{D*} is thus an incentive-compatible best response to ϕ_1^{D*} if and only if (a) the allocation rule $s_2^i(\cdot)$ maximizes (3) subject to the monotonicity constraint $s_2^i \leq s_2^{i+1}$, $i = 1, 2, 3$ and (b) the transfers t_2^i are given by (2).[4]

Next, consider upstream contracting. When the allocation rule in ϕ_2^{D*} is monotonic and the transfers satisfy (2), the buyer's payoff at $t = 1$ satisfies the single-crossing property with respect to θ and δ_1. This in turn implies that the optimal mechanism $\phi_1^{D*} : \Theta \to \Delta(\{0,1\}) \times \mathbb{R}$ solves the following program

$$\max \;\; p[t_1(\overline{\theta}) - \delta_1(\overline{\theta})] + (1-p)[t_1(\underline{\theta}) - \delta_1(\underline{\theta})]$$

[3] Note that, because $\theta s_1(\theta) - t_1(\theta)$ is sunk, from the perspective of P_2, it is as if the buyer's reservation payoff is zero, for all θ_2^E.

[4] That ϕ_2^{D*} must necessarily solve the aforementioned program follows from the fact that P_2 can always make the agent strictly prefer to truthfully reveal his private information by breaking the agent's indifference by an $\varepsilon > 0$, for ε arbitrarily small.

subject to

$$[\overline{\theta}+(\overline{\theta}+1)s_2^4-t_2^4]\delta_1(\overline{\theta})+[\overline{\theta}s_2^2-t_2^2][1-\delta_1(\overline{\theta})]-t_1(\overline{\theta})= \tag{4}$$
$$[\overline{\theta}+(\overline{\theta}+1)s_2^4-t_2^4]\delta_1(\underline{\theta})+[\overline{\theta}s_2^2-t_2^2][1-\delta_1(\underline{\theta})]-t_1(\underline{\theta})$$

$$[\underline{\theta}+(\underline{\theta}+1)s_2^3-t_2^3]\delta_1(\underline{\theta})+[\underline{\theta}s_2^1-t_2^1][1-\delta_1(\underline{\theta})]-t_1(\underline{\theta})=\underline{\theta}s_2^1-t_2^1 \tag{5}$$

and

$$\delta_1(\bar{\theta})\geq\delta_1(\underline{\theta}). \tag{6}$$

Condition (4) guarantees that $\overline{\theta}$ is indifferent between $(\delta_1(\overline{\theta}),t_1(\overline{\theta}))$ and $(\delta_1(\underline{\theta}),t_1(\underline{\theta}))$, while condition (5) guarantees that $\underline{\theta}$ is indifferent between $(\delta_1(\underline{\theta}),t_1(\underline{\theta}))$ and the null contract $(0,0)$. The high type's participation is then guaranteed by (4) and (5) while incentive-compatibility for the low type is guaranteed by the monotonicity condition (6).

Equivalently, ϕ_1^{D*} maximizes

$$U_1=p\left[\delta_1(\overline{\theta})\bar{\mathcal{V}}(\delta_1)\right]+(1-p)\left[\delta_1(\underline{\theta})\underline{\mathcal{V}}(\delta_1)\right] \tag{7}$$

subject to $\delta_1(\bar{\theta})\geq\delta_1(\underline{\theta})$, where

$$\begin{aligned}&\bar{\mathcal{V}}(\delta_1)\equiv\overline{\theta}+\left[(\overline{\theta}+1)s_2^4-t_2^4\right]-\left[(\overline{\theta}s_2^2-t_2^2)\right]-1\\&\underline{\mathcal{V}}(\delta_1)\equiv\underline{\theta}+\left[(\underline{\theta}+1)s_2^3-t_2^3\right]-\left[\underline{\theta}s_2^1-t_2^1\right]-1\\&-\tfrac{p}{1-p}\{\Delta\theta+\left[(\overline{\theta}+1)s_2^4-t_2^4-(\overline{\theta}s_2^2-t_2^2)\right]-[(\underline{\theta}+1)s_2^3-t_2^3-(\underline{\theta}s_2^1-t_2^1)]\}\end{aligned}$$

Two observations are in order. First note that ϕ_2^{D*} must specify allocations also for extended types that may have zero measure on the equilibrium path (this is the case, for example, when $\delta_1(\underline{\theta})=0$ so that $\beta^3=0$). Second note that whether ϕ_1^{D*} is incentive-compatible or not depends on the mechanism ϕ_2^{D*} offered downstream. We thus have the following result.

Example A1. *The outcome* $\pi^*=(\delta_1^*(\cdot),t_1^*(\cdot),s_2^*(\cdot),t_2^*(\cdot))$ *can be sustained as a MPE of* Γ^M *if and only if:*

(I) given $\delta_1^*(\cdot)$, $s_2^*(\cdot)$ *maximizes (3) subject to the monotonicity condition* $s_2^i\leq s_2^{i+1}$, $i=1,2,3$, *while* $t_2^*(\cdot)$ *solves (2);*

(II) given $s_2^*(\cdot)$ *and* $t_2^*(\cdot)$, $\delta_1^*(\cdot)$ *maximizes (7) subject to the monotonicity condition* $\delta_1(\bar{\theta})\geq\delta_1(\underline{\theta})$ *while* $t_1^*(\cdot)$ *solves (4) and (5).*

Extended direct mechanisms thus offer the possibility of using familiar techniques from games with a single mechanism designer to characterize necessary and sufficient conditions for equilibrium outcomes in sequential common agency. The preceding example illustrated such a possibility in a very simple way. In certain applications, the characterization of these conditions may require the use of the techniques from the multi-dimensional screening literature. This need not always be simple. However, when this is the case, assuming the principals offer menus instead of direct mechanisms does not simplify the analysis. In fact, the difficulties with multi-dimensional screening simply stem from the difficulty of controlling for the optimality of the agent's behavior. This is something one has to deal with, irrespective of how the menus are described.

A2. Alternative extensive forms: Formal results

A2-1. Observability of upstream payoff-relevant decisions

Consider an environment in which principals observe upstream payoff-relevant decisions before choosing their mechanisms. Let $\bar{\Gamma}^D$ denote the game in which the principals offer ***standard direct revelation mechanisms*** $\phi_i^D : \Theta \to D_i$ as opposed to extended direct mechanisms. We then have the following result.

Theorem 5 (Observable decisions). *(Part I: Menus) Let* $\Gamma \succcurlyeq \Gamma^M$*. Any SCF that can be sustained as a MPE of* Γ *can also be sustained as a MPE of* Γ^M*. Furthermore, any SCF that can be sustained as an equilibrium of* Γ^M *(not necessarily in Markov strategies) can also be sustained as an equilibrium of* Γ.

(Part II: Direct Mechanisms) Any SCF that can be sustained as a MPE of Γ^M *can be sustained as a pure-strategy truthful MPE of* $\bar{\Gamma}^D$*. Furthermore, any SCF that can be sustained as a MPE of* $\bar{\Gamma}^D$ *can also be sustained as a MPE of* Γ^M.

As with Theorem 4 in the main text, the result in Part (II) presumes that $D_i = \Delta(Y_i)$ for all i, which guarantees that outcomes in Γ^M sustained by mixed strategies can be sustained in $\bar{\Gamma}^D$ through mechanisms that respond to θ with lotteries over contracts. In environments in which not all possible lotteries are feasible, i.e. $D_i \subsetneqq \Delta(Y_i)$ for some i, the result in Part (II) must be replaced by the following: Any SCF 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 $\bar{\Gamma}^D$.

Proof of Theorem 5. Part I: Menus. First, consider the claim that any SCF π that can be sustained as a MPE of Γ can also be sustained as a MPE of Γ^M. The proof follows from the same steps used to establish Part 1 of Theorem 1 in the paper, with the following two (minor) adjustments. (i) The transformation of the principals' strategies indicated in that proof must now be done for any (e_t^-, a_t^-). (ii) The principals' strategies are now sustained by beliefs λ^M over upstream histories that satisfy Bayes' rule on the equilibrium path, whereas for any out-of-equilibrium (e_t^-, a_t^-), $t = 2, ..., n$, satisfy

$$\hat{\lambda}(e_t^-, a_t^-) = \hat{\lambda}^M(e_t^-, a_t^-) \quad (8)$$

where $\hat{\lambda}(e_t^-, a_t^-)$ and $\hat{\lambda}^M(e_t^-, a_t^-)$ denote the marginal distribution of λ and λ^M over Θ, respectively in the original game Γ and in the menu game Γ^M. Because the agent's strategy is Markov and (e_t^-, a_t^-) is public information, any profile of beliefs with these properties makes the principals' strategies sequentially optimal.

Next, consider the claim that any SCF that can be sustained as an equilibrium of Γ^M (not necessarily in Markov strategies) can also be sustained as an equilibrium of Γ.. The proof parallels that of Part 2 in Theorem 1. In the following, we simply construct a profile of beliefs that sustains the principals' strategies.

For any $i = 1, ..., n$, let $\mathcal{H}_i^-$ and $\mathcal{H}_i^{M-}$ denote the sets of all possible upstream histories, respectively in Γ and in Γ^M, and $\Sigma(\mathcal{H}_i^-)$ and $\Sigma(\mathcal{H}_i^{M-})$ the corresponding Borel sigma algebras. For any (e_i^-, a_i^-), let $\varkappa_i(e_i^-, a_i^-) \in \Delta(\mathcal{H}_i^-)$ and $\varkappa_i^M(e_i^-, a_i^-) \in \Delta(\mathcal{H}_i^{M-})$ denote P_i's beliefs about upstream histories, respectively in Γ and in Γ^M. If (e_i^-, a_i^-) is on the equilibrium path, then $\varkappa_i(e_i^-, a_i^-)$ is obtained from Bayes' rule using the equilibrium strategy profile σ. If instead (e_i^-, a_i^-) is an out-of-equilibrium observation, then $\varkappa_i(e_i^-, a_i^-)$ is constructed as follows. For any measurable set of histories $H_i^{M-} \in \Sigma(\mathcal{H}_i^{M-})$ in Γ^M, let $\Xi_i(H_i^{M-}) \in \Sigma(\mathcal{H}_i^-)$ denote the measurable set of histories in Γ that are obtained by substituting each history

$$h_i^{M-} = \left(\theta, e_i^-, a_i^-, \phi_i^{M-}, \delta_i^-, y_i^-\right)$$

in H_i^{M-} with the history

$$f_i(h_i^{M-}) \equiv \left(\theta, e_i^-, a_i^-, (\alpha_l(\phi_l^M))_{l=1}^{i-1}, (\tilde{\alpha}_l(\delta_l))_{l=1}^{i-1}, y_i^-\right).$$

The history $f_i(h_i^{M-})$ is simply the "translation" of the history h_i^{M-} using the embedding α_i. For any out-of-equilibrium (e_i^-, a_i^-), then let $\varkappa_i(e_i^-, a_i^-)$ be the unique beliefs that satisfy

$$\varkappa_i(\Xi_i(H_i^{M-})|e_i^-, a_i^-) = \varkappa_i^M(H_i^{M-}|e_i^-, a_i^-) \;\forall H_i^{M-} \in \Sigma(\mathcal{H}_i^{M-}).$$

Together with these beliefs, the strategy profile σ constructed from σ^M following the steps indicated in the proof of Theorem 1 constitutes an equilibrium for Γ which sustains the same outcomes as σ^M.

Part II: Direct Mechanisms. The proof parallels that of Theorem 4. The equilibrium strategy profiles σ^D and σ^M are sustained by any beliefs that are consistent with Bayes' rule on the equilibrium path, whereas for any out-of-equilibrium (e_t^-, a_t^-), satisfy

$$\hat{\lambda}^D(e_t^-, a_t^-) = \hat{\lambda}^M(e_t^-, a_t^-)$$

where $\hat{\lambda}^D(e_t^-, a_t^-)$ and $\hat{\lambda}^M(e_t^-, a_t^-)$ denote the marginal distributions of λ^D and λ^M over Θ, respectively in the revelation game $\bar{\Gamma}^D$ and in the menu game Γ^M. ■

A2-2. Observability of upstream mechanisms

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

Theorem 6 (Observable mechanisms). *(Part I: Menus) Let $\Gamma \succcurlyeq \Gamma^M$. For any $\sigma \in \mathcal{E}(\Gamma)$ in which all principals' strategies are pure, there exists a $\sigma^M \in \mathcal{E}(\Gamma^M)$ that sustains the same outcomes. Furthermore, any SCF that can be sustained as an equilibrium of Γ^M can be sustained as an equilibrium of Γ.*

(Part II: Direct Mechanisms) For any $\sigma^M \in \mathcal{E}(\Gamma^M)$ in which all principals' strategies are pure, there exists a pure-strategy truthful MPE $\sigma^D \in \mathcal{E}(\Gamma^D)$ that sustains the same outcomes.

Once again, the result in Part (II) presumes $D_i = \Delta(Y_i)$ for all i. When this is not the case, Part (II) must be replaced by the following: For any $\sigma^M \in \mathcal{E}(\Gamma^M)$ in which both the principals' and the agent's strategies are pure, there exists a pure-strategy truthful MPE $\sigma^D \in \mathcal{E}(\Gamma^D)$ that sustains the same outcomes.

Proof of Theorem 6. Part I: Menus. The proof is in two steps.

Step 1. We want to prove that, for any $\sigma \in \mathcal{E}(\Gamma)$ in which all principals' strategies are pure, there exists a $\sigma^M \in \mathcal{E}(\Gamma^M)$ that sustains the same outcomes.

Let $\sigma_i(\phi_i^-)$ denote the unique mechanism offered by P_i when the profile of upstream mechanisms is ϕ_i^-. Next, consider the game Γ_i in which P_i is restricted to offer menus, whereas all other principals have the same strategy space as in Γ. Now consider the following strategy profile $\mathring{\sigma}$ for Γ_i. For all principals P_j with $j < i$, simply let $\mathring{\sigma}_j = \sigma_j$. For P_i, let $\mathring{\sigma}_i$ be the strategy that maps each ϕ_i^- into the menu ϕ_i^M whose image is $\text{Im}(\phi_i^M) = \text{Im}(\sigma_i(\phi_i^-))$. Finally, for any P_j with $j > i$, $\mathring{\sigma}_j$ is as follows. If ϕ_j^- is such that at $t = i$, $\phi_i^M = \mathring{\sigma}_i(\phi_i^-)$, then

$$\mathring{\sigma}_j(\phi_i^-, \phi_i^M, \phi_{i+1}, ..., \phi_{j-1}) = \sigma_j(\phi_i^-, \sigma_i(\phi_i^-), \phi_{i+1}, ..., \phi_{j-1}).$$

If instead, $\phi_i^M \neq \mathring{\sigma}_i(\phi_i^-)$, then

$$\mathring{\sigma}_j(\phi_i^-, \phi_i^M, \phi_{i+1}, ..., \phi_{j-1}) = \sigma_j(\phi_i^-, \alpha_i(\phi_i^M), \phi_{i+1}, ..., \phi_{j-1}),$$

where $\alpha_i(\phi_i^M)$ is the embedding of ϕ_i^M into Φ_i.

Next, consider the agent. At any $t < i$, $\mathring{\sigma}_A(h_t) = \sigma_A(h_t)$ for any h_t. If at $t = i$, P_i offers the menu $\phi_i^M = \mathring{\sigma}_i(\phi_i^-)$, then at any downstream information set A induces the same outcomes that he would have induced in Γ had P_i offered the mechanism $\sigma_i(\phi_i^-)$, in the sense defined in the proof of Theorem 1. If, instead, P_i offers a mechanism $\phi_i^M \neq \mathring{\sigma}_i(\phi_i^-)$, then starting from $t = i$, at any subsequent information set, A behaves according to σ_A as if the game were Γ and the mechanism offered by P_i were $\alpha_i(\phi_i^M)$.

This completes the description of $\mathring{\sigma}_A$ at the information sets which are relevant for equilibrium. For all other information sets (i.e. those associated to upstream deviations by the agent), simply let $\mathring{\sigma}_A$ specify any behavior that is sequentially optimal for A given the payoff-relevant variables θ_t^E and the downstream principals' strategy profile $\mathring{\sigma}_t^+$. Given $(\mathring{\sigma}_i^+)_{i=1}^n$, the strategy $\mathring{\sigma}_A$ is clearly sequentially optimal for the agent at any information set. Thus consider the optimality of the principals' strategies. After any ϕ_j^-, $j = 1, ..., n$, beliefs about upstream histories are necessarily pinned down by Bayes' rule using the agent's strategy $\mathring{\sigma}_A$. This follows from the "no signal what you do not know" property of PBE: the observation of ϕ_j^- conveys no information about the agent's behavior in these mechanisms which hence must be assumed to have been consistent with what prescribed by the equilibrium strategy. Given these beliefs, the principals' strategies are sequentially rational. We conclude that the strategy profile $\mathring{\sigma}$ with the associated beliefs described above is an equilibrium for Γ_i and induces the same outcomes as σ in Γ.

Starting from $t = 1$ and proceeding forward, one can then apply the arguments described above to any $i = 1, ..., n$ to construct a pure-strategy equilibrium of Γ^M that implements the same outcomes as σ.

Step 2. We now prove that for any $\sigma^M \in \mathcal{E}(\Gamma^M)$ there exists a $\sigma \in \mathcal{E}(\Gamma)$ that sustains the same outcomes.

First consider the agent. The strategy σ_A is constructed by extending the strategy σ_A^M over Γ exactly as in the proof of Theorem 1. Next, consider the principals. For any t, let $\sigma_t(\phi_t^-) = \alpha_t(\sigma_t^M(\phi_t^{M-}))$, where $\alpha_t(\sigma_t^M(\cdot))$ denotes the mixed strategy over Φ_t obtained from the mixed strategy σ_t^M using the embedding α_t, while ϕ_t^{M-} denotes the profile of upstream menus that is obtained from ϕ_t^- by letting each ϕ_j^M be the menu whose image is $\text{Im}(\phi_j^M) = \text{Im}(\phi_j)$, $j = 1, ..., t-1$. The strategy profile σ constructed this way, along with the beliefs for the principals that are obtained from Bayes' rule using σ_A, is an equilibrium of Γ and sustains the same outcomes as σ^M.

Part II: Direct Mechanisms. We show that, for any $\sigma^M \in \mathcal{E}(\Gamma^M)$ in which all principals' strategies are pure, there exists a pure-strategy truthful MPE $\sigma^D \in \mathcal{E}(\Gamma^D)$ that sustains the same outcomes. Note that the agent's strategy in Γ^M need not be Markov—which explains why the proof does not follow directly from the same arguments used to establish Theorem 4.

Consider a game Γ_J in which $\Phi_j = \Phi_j^D$ for all $j \in J$ while $\Phi_j = \Phi_j^M$ for all $j \in \mathcal{N} \backslash J$, for some $J \subset \mathcal{N} \cup \{\varnothing\}$. We prove the result by showing that given any equilibrium $\sigma \in \mathcal{E}(\Gamma_J)$ in which all principals' strategies are pure, there exists an equilibrium $\mathring{\sigma} \in \mathcal{E}(\Gamma_{J'})$ that also has the property that all principals' strategies are pure and that sustains the same outcomes as σ, for any $J' = J \cup \{t\}$ with $t \in \mathcal{N} \backslash J$.

For any ϕ_t^-, let $\Theta_t^E(\phi_t^-) \subseteq \Theta_t^E$ denote the set of extended types that are consistent with σ_A (i.e. that can be generated by using σ_A recursively in Γ_J starting from $i = 1$ and proceeding forward). For any $\theta_t^E \in \Theta_t^E(\phi_t^-)$, then let $\eta(\theta_t^E, \phi_t^-) \in \Delta(\mathcal{H}_t^-)$ denote the conditional distribution over $\mathcal{H}_t^-$ that is obtained from Bayes' rule using the agent's strategy σ_A in Γ_J and conditioning on the event that the extended type in period t is θ_t^E and the mechanisms offered upstream are ϕ_t^-.

Now consider the following (pure) strategy for P_t in $\Gamma_{J'}$. For any profile of upstream mechanisms ϕ_t^-, let $\phi_t^M = \sigma_t(\phi_t^-)$ denote the equilibrium menu offered by P_t in Γ_J in response to ϕ_t^-. Then the extended direct mechanism $\phi_t^D = \mathring{\sigma}_t(\phi_t^-)$ that P_t offers in $\Gamma_{J'}$ in response to ϕ_t^- is such that

$$\phi_t^D(\theta_t^E) = \begin{cases} \int_{h_t^- \in \mathcal{H}_t^-} \int_{\delta_t \in \mathcal{M}_t^M} \delta_t d\mu(h_t^-, \sigma_t(\phi_t^-)) d\eta(\theta_t^E, \phi_t^-) & \text{if } \theta_t^E \in \Theta_t^E(\phi_t^-) \\ \delta_t \in \arg\max_{\delta_t' \in \mathrm{Im}(\sigma_t(\phi_t^-))} \mathring{V}(\theta_t^E, \phi_t^-, \sigma_t(\phi_t^-), \delta_t', \sigma_t^+) & \text{if } \theta_t^E \notin \Theta_t^E(\phi_t^-) \end{cases} \tag{9}$$

where $\mathring{V}(\theta_t^E, \phi_t^-, \sigma_t(\phi_t^-), \delta_t', \sigma_t^+)$ denotes the agent's continuation payoff in Γ given $(\theta_t^E, \phi_t^-, \sigma_t(\phi_t^-), \delta_t', \sigma_t^+)$. Note that the agent's continuation payoff now depends also on upstream mechanisms; this is because the latter now determine which mechanisms will be offered downstream. The mechanism ϕ_t^D described in (9) thus responds to each $\theta_t^E \in \Theta_t^E(\phi_t^-)$ with the same distribution over Y_t that A would have induced in the menu $\sigma_t(\phi_t^-)$ when his extended type is θ_t^E and the mechanisms offered upstream are ϕ_t^-. For any other $\theta_t^E \notin \Theta_t^E(\phi_t^-)$, the mechanism simply responds by giving the agent one of the lotteries in the menu $\phi_t^M = \sigma_t(\phi_t^-)$ that would have been optimal for θ_t^E given the mechanisms $(\phi_t^-, \sigma_t(\phi_t^-))$ and the profile of strategies σ_t^+ for the downstream principals in Γ_J.

Now consider the following strategy profile $\mathring{\sigma}$ for $\Gamma_{J'}$. For all principals P_j with $j < t$, simply let $\mathring{\sigma}_j = \sigma_j$. For P_t, let $\mathring{\sigma}_t$ be the strategy described above. Finally, for any P_j with $j > t$, $\mathring{\sigma}_j$ is as follows. If ϕ_j^- is such that in period t, P_t offered the mechanism $\phi_t^D = \mathring{\sigma}_t(\phi_t^-)$, then

$$\mathring{\sigma}_j(\phi_t^-, \phi_t^D, \phi_{t+1}, ..., \phi_{j-1}) = \sigma_j(\phi_t^-, \sigma_t(\phi_t^-), \phi_{t+1}, ..., \phi_{j-1}).$$

If instead, $\phi_t^D \neq \mathring{\sigma}_t(\phi_t^-)$, then

$$\mathring{\sigma}_j(\phi_t^-, \phi_t^D, \phi_{t+1}, ..., \phi_{j-1}) = \sigma_j(\phi_t^-, \phi_t^M, \phi_{t+1}, ..., \phi_{j-1}).$$

where ϕ_t^M is the menu whose image is $\mathrm{Im}(\phi_t^M) = \mathrm{Im}(\phi_t^D)$.

Next, consider the agent. At any $j < t$, $\mathring{\sigma}_A(h_j) = \sigma_A(h_j)$ for any h_j. If in period t, P_t offers the mechanism $\phi_t^D = \mathring{\sigma}_t(\phi_t^-)$, A truthfully reports his extended type and then at any subsequent information set, he induces the same outcomes that he would have induced in Γ_J had P_t offered the menu $\sigma_t(\phi_t^-)$. Formally, for any $y_t \in Supp[\phi_t^D(\theta_t^E)]$, let $\zeta(y_t; \theta_t^E, \phi_t^-, \sigma_t(\phi_t^-)) \in \Delta(\mathcal{H}_t^- \times \Delta(Y_t))$ denote the conditional distribution over the profiles $(h_t^-, \delta_t) \in \mathcal{H}_t^- \times \Delta(Y_t)$ in Γ_J that is obtained from Bayes' rule using the strategy σ_A, conditioning on the event that the contract selected in period t is y_t, that the agent's extended type is θ_t^E and that the mechanisms offered at $t = 1, ..., t$ are $(\phi_t^-, \sigma_t(\phi_t^-))$. In the continuation game that starts after the realization of the contract y_t, A then uses the conditional distribution $\zeta(y_t; \theta_t^E, \phi_t^-, \sigma_t(\phi_t^-))$ 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, and before choosing e_t, the history had been $(h_t^-, \sigma_t(\phi_t^-), \delta_t)$.

Finally, consider the continuation game that starts after P_t offers a mechanism $\phi_t^D \neq \mathring{\sigma}_t(\phi_t^-)$. Starting from period t, at any subsequent information set, A behaves according to σ_A as if the game were Γ_J and the menu offered by P_t were ϕ_t^M, where ϕ_t^M is the menu whose image is $\mathrm{Im}(\phi_t^M) = \mathrm{Im}(\phi_t^D)$.

This completes the description of $\mathring{\sigma}_A$ at the information sets which are relevant for equilibrium. For all other information sets (i.e. those associated to upstream deviations by the agent), simply let $\mathring{\sigma}_A$ specify any behavior that is sequentially optimal for A given the payoff-relevant variables θ_t^E and the downstream principals' strategy profile $\mathring{\sigma}_t^+$. Given the principals' strategies, the strategy $\mathring{\sigma}_A$ is sequentially optimal for the agent at any information set.

Next, consider the optimality of the principals' strategies. Given $(\mathring{\sigma}_A, \mathring{\sigma}_{-i})$, the optimality of $\mathring{\sigma}_i$ follows from the same arguments as in the proof of Part I–Step 1. We conclude that the strategy profile $\mathring{\sigma}$ with the associated beliefs $\mathring{\lambda}$ obtained from $\mathring{\sigma}$ using Bayes' rule, is an equilibrium for $\Gamma_{J'}$ and induces the same outcomes as σ in Γ_J.

Iterating across all periods, starting from $t = 1$ and proceeding forward by letting $J' = J \cup \{t+1\}$, then gives a pure-strategy truthful equilibrium of Γ^D that implements the same outcomes as σ^M. ■

Note that, contrary to the benchmark model of private contracting and to the case of observable decisions considered above, the result in Part (II) in Theorem 6 does not have a converse. There may exist SCFs that can be sustained as equilibria of Γ^D and that cannot be sustained as equilibria of Γ^M. To see this, consider the following example where $n = 2$, $|\Theta| = |E_i| = 1$, $i = 1, 2$, $\mathcal{A}_1 = \{t, b\}$ and $\mathcal{A}_2 = \{l, r\}$. 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	r
t	1 3 0	3 3 4
b	2 0 5	2 2 3

Game A1

For simplicity, assume that only deterministic mechanisms are feasible so that $D_i = \mathcal{A}_i$, $i = 1, 2$.

Now consider the revelation game Γ^D. Here a direct mechanism for P_1 coincides with the choice of an element of $\mathcal{A}_1$ whereas a direct mechanism for P_2 is a mapping $\phi_2^D : \mathcal{A}_1 \to \mathcal{A}_2$. The following is then a pure-strategy equilibrium for Γ^D. P_2 reacts to the direct mechanism of P_1 that selects t with the mechanism that responds to both t and b with l and to the mechanism that selects b with the mechanism that responds to both t and b with r. Given this reaction, P_1 in equilibrium chooses the mechanism that selects b. The equilibrium outcome is thus (b, r).

Next consider the menu game Γ^M. Suppose P_1 offers the menu $\{t, b\}$. Because l is weakly dominated for P_2, there are only two possible outcomes in the continuation game between A and P_2 that starts after P_1 offers $\{t, b\}$. In the first one, A selects t and P_2 selects r. In the second, A selects t and P_2 randomizes over l and r, respectively with probability $1/6$ and $5/6$. In both cases, P_1 obtains a payoff of $16/6 > 2$. It follows that the SCF that selects (b, r) with certainty cannot be sustained as an equilibrium in the menu game because P_1 has a profitable deviation.

A2-3. Endogenous sequence of bilateral relationships

Consider the following game with endogenous sequence of contractual relationships. There are $T < \infty$ periods. In each period, *all* principals simultaneously offer the agent a mechanism $\phi_{i,t}$ from a set $\Phi_{i,t}$. The agent chooses at most one mechanism, say $\phi_{i,t}$, to participate in, then chooses a message $m_{i,t}$ from $\mathcal{M}_{i,t}$ and a contract $y_{i,t}$ is selected by the lottery $\phi_{i,t}(m_{i,t}) \in \Delta(Y_{i,t})$. Given $y_{i,t}$, the agent then chooses an action $e_{i,t}$ from $E_{i,t}$ and finally the contract $y_{i,t} : E_{i,t} \to \Delta(\mathcal{A}_{i,t})$ determines P_i's decision. The agent may, or may not, participate in a mechanism offered by the same principal multiple times. For those principals who are not selected in period t, simply let $e_{j,t} = e^{\varnothing}_{j,t}$ and $a_{j,t} = a^{\varnothing}_{j,t}$, where $(e^{\varnothing}_{j,t}, a^{\varnothing}_{j,t})$ are the exogenous default decisions that are implemented in the absence of contracting, such as no trade at a null price.

Payoffs, respectively for P_i, $i = 1, ..., n$, and for A continue to be denoted by $u_i(\theta, e, a)$ and $v(\theta, e, a)$, with $e_\tau \equiv (e_{1,\tau}, ..., e_{n,\tau})$ and $a_\tau \equiv (a_{1,\tau}, ..., a_{n,\tau})$ now denoting an entire profile of payoff-relevant decisions for period τ, one for each possible bilateral relationship, and $e \equiv (e_1, ..., e_T)$ and $a \equiv (a_1, ..., a_T)$.

For any $t = 1, ..., T$, any $i = 1, ..., n$ and any upstream history h^-_t, let $z_{i,t} = f_{i,t}(h^-_t)$ denote the elements of h^-_t that are observed by P_i in period t.[5] The function $f_{i,t} : \mathcal{H}^-_t \to Z_{i,t}$ maps each possible upstream history $h^-_t \in \mathcal{H}^-_t$ into an observation $z_{i,t} \in Z_{i,t}$, where $Z_{i,t} \equiv \{z_{i,t} : z_{i,t} = f_{i,t}(h^-_t), h^-_t \in \mathcal{H}^-_t\}$. As in the benchmark model, contracting is *private* in the sense that principals do not observe other principals' mechanisms, nor the messages, the contracts, or the decisions taken in these mechanisms. These restrictions are embedded in the mappings $f_{i,t}$.

For any $z_{i,t} \in Z_{i,t}$, let $\psi(z_{i,t})$ denote the payoff-relevant component of $z_{i,t}$, that is, the part of the agent's extended type $\theta^E_t = (\theta, e^-_t, a^-_t)$ that is observed by P_i at date t. Note that the agent's extended type now contains profiles of payoff-relevant decisions $e_\tau \equiv (e_{1,\tau}, ..., e_{n,\tau})$ and $a_\tau \equiv (a_{1,\tau}, ..., a_{n,\tau})$, one for each bilateral relationship, with $e^-_t \equiv (e_\tau)^{t-1}_{\tau=1}$ and $a^-_t \equiv (a_\tau)^{t-1}_{\tau=1}$.

Principal i's behavioral strategy in period t is now described by the distribution $\sigma_{i,t}(z_{i,t}) \in \Delta(\Phi_{i,t})$ over the mechanisms in $\Phi_{i,t}$. The agent's behavioral strategy $\sigma_A(h^-_t, \phi_t)$ given the upstream history $h^-_t \in \mathcal{H}^-_t$ and the profile of mechanisms $\phi_t \equiv (\phi_{1,t}, ..., \phi_{n,t})$ offered in period t, is decomposed as follows: $w^t(h^-_t, \phi_t) \in \Delta(\mathcal{N} \cup \varnothing)$ denotes the agent's participation strategy; $\mu_t(h^-_t, \phi_t, I_t) \in \Delta(\tilde{\mathcal{M}}_t)$ denotes the agent's message strategy after he chooses to participate in principal I_t's mechanism, where $I_t \in \mathcal{N} \cup \varnothing$ denotes the identity of the principal selected in period t and $\tilde{\mathcal{M}}_t \equiv \prod_i (\mathcal{M}_{i,t} \cup \varnothing)$; finally, $\xi(h^-_t, \phi_t, I_t, m_t, y_t) \in \Delta(\tilde{E}_t)$ denotes the agent's effort strategy, with $\tilde{E}_t \equiv \prod_i (E_{i,t} \cup \varnothing)$.[6]

Definition A1. *Principal i's strategy in period t is Markov if and only if, for any $z_{i,t}, z'_{i,t} \in Z_{i,t}$ such that $\psi(z_{i,t}) = \psi(z'_{i,t})$, $\sigma_{i,t}(z_{i,t}) = \sigma_{i,t}(z'_{i,t})$.*

[5] A history h^-_t now also includes the agent's upstream participation decisions.

[6] The vector $m_t \equiv (m_{1,t}, ..., m_{n,t})$ denotes the profile of messages sent by the agent in period t, with $m_{j,t} = \varnothing$ for any $j \neq I_t$. Similarly, $y_t \equiv (y_{1,t}, ..., y_{n,t})$ and $e_t \equiv (e_{1,t}, ..., e_{n,t})$ denote, respectively, the vector of contracts and the vector of effort choices, for period t, with $y_{j,t}, e_{j,t} = \varnothing$ for any $j \neq I_t$.

The agent's strategy in period t is Markov if and only if the following are true:

(a) for any (h_t^-, ϕ_t) *and* $(\tilde{h}_t^-, \phi_t)$ *such that* θ_t^E *is the same in* h_t^- *and* $\tilde{h}_t^-$, $w^t(h_t^-, \phi_t) = w^t(\tilde{h}_t^-, \phi_t)$;

(b) for any (h_t^-, ϕ_t, I_t) *and* $(\tilde{h}_t^-, \tilde{\phi}_t, I_t)$ *such that* θ_t^E *is the same in* h_t^- *and* $\tilde{h}_t^-$ *and* $\phi_{I_t,t}$ *is the same in* ϕ_t *and* $\tilde{\phi}_t$, $\mu_t(h_t^-, \phi_t, I_t) = \mu_t(\tilde{h}_t^-, \tilde{\phi}_t, I_t)$;[7]

(c) for any $(h_t^-, \phi_t, I_t, m_t, y_t)$ *and* $(\tilde{h}_t^-, \tilde{\phi}_t, I_t, \tilde{m}_t, y_t)$ *such that* θ_t^E *is the same in* h_t^- *and* $\tilde{h}_t^-$, $\xi(h_t^-, \phi_t, I_t, m_t, y_t) = \xi(\tilde{h}_t^-, \tilde{\phi}_t, I_t, \tilde{m}_t, y_t)$.

An equilibrium $\sigma \in \mathcal{E}(\Gamma)$ *is a MPE if and only if all players' strategies are Markov at any* $t = 1, ..., T$.

Theorem 7 (Endogenous sequence). *(Part I: Menus) Let* $\Gamma \succcurlyeq \Gamma^M$.[8] *Any SCF that can be sustained as a MPE of* Γ *can also be sustained as a MPE of* Γ^M. *Furthermore, any SCF that can be sustained as an equilibrium of* Γ^M *(not necessarily in Markov strategies) can also be sustained as an equilibrium of* Γ.

(Part II: Direct Mechanisms) Suppose the agent can contract with each principal at most once. Then any SCF that can be sustained as a MPE of Γ^M *can also be sustained as a truthful MPE of* Γ^D. *Furthermore, any SCF that can be sustained as a MPE of* Γ^D *can also be sustained as a MPE of* Γ^M.

Proof of Theorem 7. Part (I). The proof is in two steps and combines arguments from the proofs of Theorems 1 and 5.

Step 1. We want to show that given any MPE $\sigma \in \mathcal{E}(\Gamma)$, there exists a MPE $\sigma^M \in \mathcal{E}(\Gamma^M)$ that sustains the same outcomes as σ. The arguments here are similar to those in the proof of Theorem 1. The only differences come from the fact that (a) one has to adjust the replication arguments to take into account that the principals' strategies are now contingent on what they have observed upstream and (b) that one must specify supporting beliefs for the principals' strategies.

Consider the partition game $\Gamma^{\mathcal{Q}_{i,t}}$ in which, in period t, P_i chooses a cell $Q_{i,t}$ from the partition $\mathcal{Q}_{i,t}$ of $\Phi_{i,t}$ simultaneously with the other principals choosing their mechanisms $\phi_{j,t}$ from $\Phi_{j,t}$, $j \neq i$. Given $(Q_{i,t}, (\phi_{j,t})_{j \neq i})$, A first selects a mechanism $\phi_{i,t}$ from $Q_{i,t}$ and then, given the profile $(\phi_{i,t}, (\phi_{j,t})_{j \neq i})$, he chooses which mechanism to participate in. The choice of $\phi_{i,t}$ is observed by P_i, but not by the other principals. For any other principal and any other date, the choice set in $\Gamma^{\mathcal{Q}_{i,t}}$ is the same as in Γ; that is, for any $(j, \tau) \neq (i, t)$, the strategy space for P_j at date τ remains $\Phi_{j,\tau}$.

[7] If $I_t = \varnothing$, then $\phi_{I,t} = \varnothing$.

[8] The game Γ is an enlargement of Γ^M if, for any $i = 1, ..., n$, and any $t = 1, ..., T$, the following are true:

(a) $\text{Im}(\phi_{i,t})$ is compact, for any $\phi_{i,t} \in \Phi_{i,t}$;

(b) there exists a injective mapping $\alpha_{i,t} : \Phi_{i,t}^M \to \Phi_{i,t}$ such that, for any pair of mechanisms $\phi_{i,t}^M, \phi_{i,t}$ with $\phi_{i,t} = \alpha_{i,t}(\phi_{i,t}^M)$,(i) $\text{Im}(\phi_{i,t}^M) = \text{Im}(\phi_{i,t})$, and (ii) there exists an injective function $\tilde{\alpha}_{i,t} : \mathcal{M}_{i,t}^M \to \mathcal{M}_{i,t}$ such that $\phi_{i,t}^M(\delta_{i,t}) = \delta_{i,t} = \phi_i(\tilde{\alpha}_i(\delta_{i,t}))$ for any $\delta_{i,t} \in \mathcal{M}_{i,t}^M$;

(c) there exists an injective mapping $\mathring{\alpha}_{i,t} : Z_{i,,t}^M \to Z_{i,t}$ from the set of possible signals $Z_{i,t}^M$ in Γ^M to the set of possible signals $Z_{i,t}$ in Γ.

Now let $\mathcal{Q}_{i,t}$be the partition of $\Phi_{i,t}$ given by the equivalence relation

$$\phi_{i,t} \sim_{i,t} \phi'_{i,t} \Longleftrightarrow \mathrm{Im}(\phi_{i,t}) = \mathrm{Im}(\phi'_{i,t}). \tag{10}$$

Following the same construction as in Step 1 in the proof of Theorem 1, it is easy to see that there exists an equilibrium $\hat{\sigma}$ for $\Gamma^{\mathcal{Q}_{i,t}}$ which sustains the same outcomes as σ in Γ. In this equilibrium, all P_j with $j \neq i$, follow the same strategy as in Γ, i.e. $\widehat{\sigma}_j = \sigma_j$. As for P_i, at any $\tau \neq t$ and for any $z_{i,\tau} \in Z_{i,\tau}$, $\widehat{\sigma}_i(z_{i,\tau}) = \sigma_i(z_{i,\tau})$.[9] In period t, for any $z_{i,t} \in Z_{i,t}$, P_i randomizes over any subset R of $\mathcal{Q}_{i,t}$ whose union is measurable with probability

$$\widehat{\sigma}_i(R\ ;\ z_{i,t}) = \sigma_i(\bigcup R\ ;\ z_{i,t}).$$

The agent's strategy is such that at any $\tau < t$, $\widehat{\sigma}_A(h_\tau) = \sigma_A(h_\tau)$ for any $h_\tau \in \mathcal{H}_\tau$. In period t, given any $(Q_{i,t}, (\phi_{j,t})_{j\neq i})$, A uses the conditional probability distribution $\sigma_{i,t}(\cdot|Q_{i,t}; z_{i,t})$ to select a mechanism $\phi_{i,t}$ from $Q_{i,t}$. At any subsequent informational set, A then behaves as if the game were Γ and the mechanism offered by P_i were $\phi_{i,t}$. As far as beliefs are concerned, at any information set, all principals have the same marginal beliefs over upstream payoff-relevant information as in Γ (note that, on the equilibrium path, this is consistent with principals' beliefs be obtained from Bayes rule). Because all players' strategies in $\hat{\sigma}$ are Markov, given these beliefs, all principals' strategies are sequentially rational.

Next, consider the game $\Gamma^M_{i,t}$ in which, in period t, P_i's choice set is $\Phi^M_{i,t}$, whereas for any $(j,\tau) \neq (i,t)$, P_j's choice set in period τ is the same as in Γ. Now, for any $\tau = 1, ..., T$, let $Z^M_{j,\tau}$ denote the set of possible signals that P_j can receive in $\Gamma^M_{i,t}$ in period τ, with $Z^M_{j,\tau} = Z_{j,\tau}$ for any (j,τ) such that either $j \neq i$, or $\tau \leq t$.

Because all players' strategies are Markov in $\hat{\sigma}$, starting from $\hat{\sigma}$ and following essentially the same construction as in Step 2 in the proof of Theorem 1, one can show that there exists a MPE $\mathring{\sigma} \in \mathcal{E}(\Gamma^M_{i,t})$ that sustains the same outcomes as σ. We refer the reader to that proof for the details of how to construct the strategies in $\mathring{\sigma}$ from the strategies in $\hat{\sigma}$. The only important observation is that, given the menu $\phi^M_{i,t}$ offered by P_i in period t, the agent uses the conditional distribution $\sigma_{i,t}(\cdot|Q_{i,t}(\phi^M_{i,t}); z_{i,t})$ to determine not only the messages to send to P_i in case he decides to participate in $\phi^M_{i,t}$ but also his participation decision. That is, given any profile of mechanisms $(\phi^M_{i,t}, (\phi_{j,t})_{j\neq i})$, A uses the conditional probability distribution $\sigma_{i,t}(\cdot|Q_{i,t}(\phi^M_{i,t}); z_{i,t})$ to select in his mind a mechanism $\phi_{i,t}$ from $Q_{i,t}(\phi^M_{i,t}) \equiv \{\phi_{i,t} : \mathrm{Im}(\phi_{i,t}) = \mathrm{Im}(\phi^M_{i,t})\}$ and then uses the original strategy $w^t(h^-_t, (\phi_{i,t}, (\phi_{j,t})_{j\neq i}))$ for Γ to determine his participation decision. At all subsequent information sets, the construction of $\mathring{\sigma}$ parallels that of $\hat{\sigma}$ in $\Gamma^{\mathcal{Q}_{i,t}}$.

The principals' strategies in $\mathring{\sigma}$ can be sustained by beliefs $\mathring{\lambda}_{j,\tau}(z^M_{j,\tau}) \in \Delta(\mathcal{H}^-_\tau)$ over upstream histories that satisfy the following properties.

[9] Formally, for any $\tau > t$, $z_{i,\tau}$ now includes the cell $Q_{i,t}$. However, because to any $\phi_{i,t}$ corresponds a unique cell $Q_{i,t}$, we can drop $Q_{i,t}$ from $z_{i,\tau}$.

Case (i). If $z_{j,\tau}^M$ is such that, given the mechanisms $(\phi_{j,l})_{l=1}^{\tau-1}$ offered by P_j upstream, the decisions in $z_{j,\tau}^M$ are consistent with $\mathring{\sigma}_A$ and $(\mathring{\sigma}_k)_{k\neq j}$, then $\mathring{\lambda}_{j,\tau}(z_{j,\tau}^M)$ are obtained from Bayes' rule using $\mathring{\sigma}_A$ and $(\mathring{\sigma}_k)_{k\neq j}$. For all P_j with $j \neq i$, these beliefs necessarily have the same marginal distribution over Θ_τ^E as in $\Gamma^{\mathcal{Q}_{i,t}}$ given $z_{j,\tau} = z_{j,\tau}^M$. Clearly, the same is true for P_i if $\tau \leq t$, but not necessarily if $\tau > t$. In fact, if $\tau > t$, then P_i's posterior beliefs about θ_τ^E in $\Gamma_{i,t}^M$ after P_i offered the menu $\phi_{i,t}^M$ in period t are a convex combination of the beliefs she would have had in $\Gamma^{\mathcal{Q}_{i,t}}$ had she offered $Q_{i,t}(\phi_{i,t}^m)$ in period t. More precisely, let $z_{i,\tau} = (\left(z_{i,\tau}^M \backslash \phi_{i,t}^m\right) \wedge \phi_{i,t}) \in Z_{i,\tau}$ denote the observation that is obtained from $z_{i,\tau}^M$ by substituting the mechanism $\phi_{i,t}^m$ with $\phi_{i,t}$. Similarly, let $z_{i,\tau} = (\left(z_{i,\tau}^M \backslash \phi_{i,t}^m, \delta_{i,t}\right) \wedge \phi_{i,t}, m_{i,t}) \in Z_{i,\tau}$ denote the observation that is obtained from $z_{i,\tau}^M$ by substituting the mechanism $\phi_{i,t}^m$ and the message $\delta_{i,t}$ with $\phi_{i,t}$ and $m_{i,t}$. Now let $\mathring{\Lambda}_{i,\tau}$ and $\Lambda_{i,\tau}$ denote P_i's marginal beliefs over Θ_τ^E, respectively in $\Gamma_{i,t}^M$ in $\Gamma^{\mathcal{Q}_{i,t}}$. First, suppose the agent did not participate in P_i's mechanism in period t, so that $I_t \neq i$. Then P_i's posterior beliefs over Θ_τ^E in period $\tau > t$ satisfy

$$\mathring{\Lambda}_{i,\tau}(z_{i,\tau}^M) = \int_{\Phi_{i,t}(\phi_{i,t}^m)} \Lambda_{i,t}(\left(z_{i,\tau}^M \backslash \phi_{i,t}^m\right) \wedge \phi_{i,t}) d\eta_{i,t}(\phi_{i,t}|z_{i,\tau}^M)$$

where $\eta_{i,t}(\phi_{i,t}|z_{i,\tau}^M)$ denote P_i's beliefs that the agent in period t behaved as if the game were $\Gamma^{\mathcal{Q}_{i,t}}$ and selected $\phi_{i,t}$ from $\Phi_{i,t}(\phi_{i,t}^m)$, given $z_{i,\tau}^M$. Next, suppose $I_t = i$ and let $\mathcal{M}_{i,t}(\delta_{i,t})$ denote the set of messages in $\Phi_{i,t}(\phi_{i,t}^m)$ that lead to the lottery $\delta_{i,t}$. Then P_i's posterior beliefs over Θ_τ^E in period $\tau > t$ satisfy

$$\mathring{\Lambda}_{i,\tau}(z_{i,\tau}^M) = \int_{\Phi_{i,t}(\phi_{i,t}^m)} \int_{\mathcal{M}_{i,t}(\delta_{i,t})} \Lambda_{i,t}(\left(z_{i,\tau}^M \backslash \phi_{i,t}^m, \delta_{i,t}\right) \wedge \phi_{i,t}, m_{i,t}) d\gamma_{i,t}(\phi_{i,t}, m_{i,t}|z_{i,\tau}^M)$$

where $\gamma_{i,t}(\phi_{i,t}, m_{i,t}|\phi_{i,t}^M, \delta_{i,t})$ denote P_i's beliefs that the agent in period t behaved as if the game were $\Gamma^{\mathcal{Q}_{i,t}}$, he selected $\phi_{i,t}$ from $\Phi_{i,t}(\phi_{i,t}^m)$, and then sent the message $m_{i,t}$. This difference in beliefs with respect to $\Gamma^{\mathcal{Q}_{i,t}}$ is due to the fact that the choice of the mechanism $\phi_{i,t}$ from $\Phi_{i,t}(\phi_{i,t}^m)$ and of the message $m_{i,t}$ from $\mathcal{M}_{i,t}(\delta_{i,t})$ is now only in the agent's mind and is thus not directly observed by P_i.

Given the aforementioned beliefs, the (behavioral) strategies $\mathring{\sigma}_{j,\tau}(z_{j,\tau}^M) = \hat{\sigma}_{j,\tau}(z_{j,\tau}^M)$ for all (j, τ) such that either $j \neq i$ or $\tau < t$ are clearly sequentially optimal.[10] Thus consider $j = i$ and $\tau > t$. Because the strategy $\hat{\sigma}_{i,\tau}$ was Markov in $\Gamma^{\mathcal{Q}_{i,t}}$, then $\hat{\sigma}_{i,\tau}(z_{i,\tau}) = \hat{\sigma}_{i,\tau}(z'_{i,\tau})$ for any $z_{i,\tau}$ and $z'_{i,\tau}$ that contain the same payoff-relevant information, i.e. such that $\psi(z_{i,\tau}) = \psi(z'_{i,\tau})$. Now, suppose $z_{i,\tau}^M$ is such that $I_t \neq i$ and let $Z_{i,\tau}(z_{i,\tau}^M)$ denote the set of observations $z_{i,\tau} \in Z_{i,\tau}$ such that $z_{i,\tau} = (\left(z_{i,\tau}^M \backslash \phi_{i,t}^m\right) \wedge \phi_{i,t}) \in Z_{i,\tau}$, with $\phi_{i,t} \in \Phi_{i,t}(\phi_{i,t}^m)$. Clearly $\psi(z_{i,\tau}) = \psi(z'_{i,\tau})$ for any pair $z_{i,\tau}, z'_{i,\tau} \in Z_{i,\tau}(z_{i,\tau}^M)$. That in $\Gamma^{\mathcal{Q}_{i,t}}$ the strategy $\hat{\sigma}_{i,\tau}$ was Markov implies that $\hat{\sigma}_{i,\tau}(z_{i,\tau})$ was optimal for any $z_{i,\tau} \in Z_{i,\tau}(z_{i,\tau}^M)$ and hence for any beliefs $\Lambda_{i,t}(\left(z_{i,\tau}^M \backslash \phi_{i,t}^m\right) \wedge \phi_{i,t})$, with $\phi_{i,t} \in \Phi_{i,t}(\phi_{i,t}^m)$. Because $\mathring{\Lambda}_{i,\tau}(z_{i,\tau}^M)$ is a convex combination of $\Lambda_{i,t}(\left(z_{i,\tau}^M \backslash \phi_{i,t}^m\right) \wedge \phi_{i,t})$, with $\phi_{i,t} \in \Phi_{i,t}(\phi_{i,t}^m)$, this necessarily implies that $\mathring{\sigma}_{i,\tau}(z_{i,\tau}^M) = \hat{\sigma}_{i,\tau}(z_{i,\tau})$, with $z_{i,\tau} \in Z_{i,\tau}(z_{i,\tau}^M)$, is sequentially optimal.

[10] Recall that for these (j, τ), $Z_{j,\tau}^M = Z_{j,\tau}$.

Next, suppose that $z^M_{i,\tau}$ is such that $I_t = i$ and let $Z_{i,\tau}(z^M_{i,\tau})$ denote the set of observations $z_{i,\tau} \in Z_{i,\tau}$ such that $z_{i,\tau} = (\left(z^M_{i,\tau} \backslash \phi^m_{i,t}, \delta_{i,t}\right) \wedge \phi_{i,t}, m_{i,t})$, with $\phi_{i,t} \in \Phi_{i,t}(\phi^m_{i,t})$ and $\phi_{i,t}(m_{i,t}) = \delta_{i,t}$. The same arguments as for $I_t \neq i$ imply that the strategy $\mathring{\sigma}_{i,\tau}(z^M_{i,\tau}) = \hat{\sigma}_{i,\tau}(z_{i,\tau})$, with $z_{i,\tau} \in Z_{i,\tau}(z^M_{i,\tau})$, is sequentially optimal.

Case (ii). Next, suppose the observation $z^M_{j,\tau}$ indicates that a departure from equilibrium play occurred by either A or some P_j, $j \neq i$. Then let $\mathring{\lambda}_{j,\tau}(z^M_{j,\tau})$ be any beliefs that are consistent with[11] $z^M_{j,\tau}$ and satisfy $\mathring{\Lambda}_{j,\tau}(z^M_{j,\tau}) = \Lambda_{j,\tau}(z_{j,\tau})$, where $\mathring{\Lambda}_{j,\tau}$ and $\Lambda_{j,\tau}$ denote P_j's marginal beliefs over Θ^E_τ, respectively in $\Gamma^M_{i,t}$ conditional on $z^M_{j,\tau}$ and in $\Gamma^{\mathcal{Q}_{i,t}}$ conditional on $z_{j,\tau}$, where $z_{j,\tau}$ is any signal that contains the same payoff-relevant information as $z^M_{j,\tau}$.

Because $\Gamma^M_{i,t} \succcurlyeq \Gamma^M$ and because $\mathring{\sigma}$ is a MPE of $\Gamma^M_{i,t}$, one can keep iterating the same construction described above across all i and all t, starting from $t = 1$ and proceeding forward. This gives a MPE $\sigma^M \in \mathcal{E}(\Gamma^M)$ that sustains the same outcomes as σ.

Step 2. Next, we prove that, given any $\sigma^M \in \mathcal{E}(\Gamma^M)$ (not necessarily in Markov strategies) there exists a $\sigma \in \mathcal{E}(\Gamma)$ that sustains the same outcomes as σ^M. The construction parallels that in the proof of Theorems 1 and 5.

First, consider the agent. The strategy σ_A is constructed from σ^M_A as in the proof of Theorem 1. After any history $h_t = (\theta, (\phi_\tau, I_\tau, m_\tau, y_\tau, e_\tau, a_\tau)^{t-1}_{\tau=1}, \phi_t)$, the agent behaves according to σ^M_A (in the same sense as in the proof of Theorem 1 in the main text) as if the game were Γ^M and the history were $h^M_t = (\theta, (\phi^M_\tau, I_\tau, \delta_\tau, y_\tau, e_\tau, a_\tau)^{t-1}_{\tau=1}, \phi^M_t)$ where the history h^M_t is obtained from h_t replacing $((\phi_\tau)^{t-1}_{\tau=1}, \phi_t)$ with $((\phi^M_\tau)^{t-1}_{\tau=1}, \phi_t)$ and $(m_\tau)^{t-1}_{\tau=1}$ with $(\delta_\tau)^{t-1}_{\tau=1}$, where each $\phi^M_{j,\tau}$ in h^M_t is the menu whose image is $\text{Im}(\phi^M_{j,\tau}) = \text{Im}(\phi_{j,\tau})$ and where $\delta_{j,\tau} = \phi_{j,\tau}(m_{j,\tau})$.[12]

Next, consider the principals. For any t, any i and any $z_{i,t} \in Z_{i,t}$, let $\sigma_i(z_{i,t}) = \alpha_i(\sigma^M_i(\beta(z_{i,t})))$, where $\alpha_i(\sigma^M_i)$ is the distribution over Φ_i obtained from σ^M_i using the embedding α_i and where $z^M_{i,t} = \beta(z_{i,t})$ is the observation obtained from $z_{i,t}$, using the same transformation of $\phi_{i,\tau}$ and $m_{i,\tau}$ indicated above for the agent.

The principals' strategies are supported by the following beliefs. For any t, let $\mathcal{H}^-_t$ and $\mathcal{H}^{M-}_t$ denote the sets of all possible upstream histories, respectively in Γ and in Γ^M, and $\Sigma(\mathcal{H}^-_t)$ and $\Sigma(\mathcal{H}^{M-}_t)$ denote the corresponding Borel sigma algebras. For any $z_{i,t}$ and $z^M_{i,t}$, let $\varkappa_{i,t}(z_{i,t}) \in \Delta(\mathcal{H}^-_t)$ and $\varkappa^M_{i,t}(z^M_{i,t}) \in \Delta(\mathcal{H}^{M-}_t)$ denote P_i's period-t beliefs about upstream histories, respectively in Γ and in Γ^M. If $z_{i,t}$ is such that, given the mechanisms $(\phi_{i,\tau})^{t-1}_{\tau=1}$ offered by P_i upstream, the decisions in $z_{i,t}$ are consistent with σ_A and $(\sigma_k)_{k \neq i}$, then $\varkappa_{i,t}(z_{i,t})$ is obtained from Bayes' rule using σ_A and $(\sigma_k)_{k\neq i}$. Otherwise, $\varkappa_{i,t}(z_{i,t})$ are constructed as follows. For any measurable set of upstream histories $H^{M-}_t \in \Sigma(\mathcal{H}^{M-}_t)$ in Γ^M, let $\Xi_t(H^{M-}_t) \in \Sigma(\mathcal{H}^-_t)$ denote the measurable set of histories in

[11] The beliefs $\mathring{\lambda}_{j,\tau}(z^M_{j,\tau}) \in \Delta(\mathcal{H}^-_\tau)$ are consistent with $z^M_{i,t}$ if they assign positive measure only to upstream histories h^-_t such that $f_{i,t}(h^-_t) = z_{i,t}$.

[12] For any principal i not selected in period τ, $\delta_{i,\tau}, y_{i,\tau} = \varnothing$.

Γ that are obtained by substituting each history

$$h_t^M = (\theta, (\phi_\tau^M, I_\tau, \delta_\tau, y_\tau, e_\tau, a_\tau)_{\tau=1}^{t-1})$$

in H_t^{M-} with the family of histories $f_t(h_t^{M-}) \in \Sigma(\mathcal{H}_t^-)$ such that, each history

$$h_t^- = (\theta, (\phi_\tau, I_\tau, m_\tau, y_\tau, e_\tau, a_\tau)_{\tau=1}^{t-1})$$

in $f_t(h_t^{M-})$ has the following properties: (a) $(\theta, (I_\tau, y_\tau, e_\tau, a_\tau)_{\tau=1}^{t-1})$ is the same as in h_t^M; (b) each $\phi_{i,\tau}$ is such that $\mathrm{Im}(\phi_{i,\tau}) = \mathrm{Im}(\phi_{i,\tau}^M)$; each $m_{i,\tau}$ is such that $m_{i,\tau} = \varnothing$ if $\delta_{i,\tau} = \varnothing$ and $\phi_{i,\tau}(m_{i,\tau}) = \delta_{i,\tau}$ if $\delta_{i,\tau} \neq \varnothing$. For any out-of-equilibrium $z_{i,t}$, then let $\varkappa_i(z_{i,t})$ be the unique beliefs that are consistent with $z_{i,t}$ and satisfy

$$\varkappa_{i,t}(\Xi_t(H_t^{M-}) \mid z_{i,t}) = \varkappa_{i,t}^M(H_t^{M-} \mid \beta(z_{i,t}))\ \forall H_t^{M-} \in \Sigma(\mathcal{H}_t^{M-})$$

where $z_{i,t}^M = \beta(z_{i,t})$ is obtained from $z_{i,t}$, using the transformation of $\phi_{i,\tau}$ and $m_{i,\tau}$ indicated above for the agent. With these beliefs, the strategy σ_i given by $\sigma_i(z_{i,t}) = \alpha_i(\sigma_i^M(\beta(z_{i,t})))$ for any $z_{i,t}$ is sequentially rational for P_i, given σ_A and $(\sigma_k)_{k\neq i}$.

Furthermore, given the principals' strategies $(\sigma_i)_{i=1}^n$ constructed above, the agent's strategy σ_A is clearly sequentially rational. We conclude that $\sigma \in \mathcal{E}(\Gamma)$. That σ implements the same SCF as σ^M is then immediate.

Proof of Part (II). The proof is in two steps.

Step 1. Consider an environment in which the agent contracts with each principal at most once. We want to show that given any MPE $\sigma^M \in \mathcal{E}(\Gamma^M)$, there exists a MPE $\sigma^D \in \mathcal{E}(\Gamma^D)$ that sustains the same SCF as σ^M. To ease the exposition, hereafter we allow the principals to offer mechanisms also in periods subsequent to the one they contracted with the agent. This is clearly inconsequential for the arguments below.

Let Γ_J denote a game in which $\Phi_{j,\tau} = \Phi_{j,\tau}^D$ for all $(j,\tau) \in J$, while $\Phi_{j,\tau} = \Phi_{j,\tau}^M$ for all $(j,\tau) \in \mathcal{R} \backslash J$, for some $J \subset \mathcal{R} \cup \{\varnothing\}$, where $\mathcal{T} \equiv \{1, ..., T\}$ and $\mathcal{R} \equiv (\mathcal{N} \times \mathcal{T})$. We prove the result by showing that, given any MPE $\sigma \in \mathcal{E}(\Gamma_J)$, there exists an MPE $\tilde{\sigma} \in \mathcal{E}(\Gamma_{J'})$, with $J' = J \cup \{i, t\}$ for some $\{i, t\} \in \mathcal{R} \backslash J$, that sustains the same outcomes.

That the agent's strategy in σ is Markov implies that, for any $\phi_{i,t}^M \in \Phi_{i,t}^M$, there is a single probability distribution $\delta_{i,t}(\theta_t^E, \phi_{i,t}^M) \in \Delta(Y_{i,t})$ over $Y_{i,t}$ such that, conditional on having decided to participate in $\phi_{i,t}^M$, whatever the particular upstream history h_t^- that conducted to θ_t^E, A always induces the distribution $\delta_{i,t}(\theta_t^E, \phi_{i,t}^M)$ when his extended type is θ_t^E.

The MPE $\tilde{\sigma}$ that sustains π in $\Gamma_{J'}$ is obtained from σ as follows. For any $\tau \neq t$, all players' (Markov) strategies are the same as in σ. For $\tau = t$, if $j \neq i$, then $\tilde{\sigma}_{j,t} = \sigma_{j,t}$. If instead $j = i$, then $\tilde{\sigma}_{i,t}$ is obtained from $\sigma_{i,t}$ as follows. For any menu $\phi_{i,t}^M$, let $\phi_{i,t}^D = g_{i,t}(\phi_{i,t}^M)$ be the direct mechanism given by

$$\phi_{i,t}^D(\theta_t^E) = \delta_{i,t}(\theta_t^E, \phi_{i,t}^M) \quad \forall \theta_t^E \in \Theta_t^E.$$ [13]

Now, let $\Phi^D_{i,t}(g_{i,t}) \equiv \{\phi^D_{i,t} : \phi^D_{i,t} = g_{i,t}(\phi^M_{i,t}),\ \phi^M_{i,t} \in \Phi^M_{i,t}\}$. After any $z_{i,t} \in Z^{J'}_{i,t}$, P_i uses his original behavioral strategy $\sigma_i(z_{i,t})$ to randomize over $\Phi^D_{i,t}$; formally, for any measurable subset $K \subseteq \Phi^D_{i,t}$

$$\tilde{\sigma}_i(K; z_{i,t}) = \sigma_i(B_K; z_{i,t})$$

where $B_K \equiv \{\phi^M_{i,t} \in \Phi^M_i : g_{i,t}(\phi^M_{i,t}) \in K\}$. Clearly, any menu in B_K is payoff-equivalent for the agent. Given any profile of mechanisms $(\phi^D_{i,t}, (\phi_{j,t})_{j\neq i})$ with $\phi^D_{i,t} \in \Phi^D_{i,t}(g_{i,t})$, A then uses the conditional distribution $\sigma_i(\cdot \mid B_{\phi^D_{i,t}})$ to determine his participation decision. That is, with probability $\sigma_i(\phi^M_{i,t} \mid B_{\phi^D_{i,t}})$, A behaves according to the participation strategy $w^t(\theta^E_t, \phi^M_{i,t}, (\phi_{j,t})_{j\neq i})) \in \Delta(\mathcal{N} \cup \varnothing)$ as if the game were Γ_J and the mechanisms offered by the principals were $\phi_t = (\phi^M_{i,t}, (\phi_{j,t})_{j\neq i})$. If the lottery $w^t(\theta^E_t, \phi^M_{i,t}, (\phi_{j,t})_{j\neq i}))$ selects P_i, A reports his extended type truthfully to P_i. If instead, $w^t(\theta^E_t, \phi^M_{i,t}, (\phi_{j,t})_{j\neq i}))$ selects a P_j with $j \neq i$, then A uses the same Markov strategy as in Γ_J to select which messages to send to P_j. In either case, the agent's choice of effort is governed by the same Markov strategy as in Γ_J.

Next, consider a $(\phi^D_{i,t}, (\phi_{j,t})_{j\neq i})$ such that $\phi^D_{i,t} \notin \Phi^D_{i,t}(g_{i,t})$.Then, at any downstream information set A behaves as if the game were Γ_J and the menu offered by P_i were $\phi^M_{i,t}$ where $\phi^M_{i,t}$ is the menu whose image is $\mathrm{Im}(\phi^M_{i,t}) = \mathrm{Im}(\phi^D_{i,t})$.

The principals' strategies in $\tilde{\sigma}$ can be sustained by beliefs over upstream histories that satisfy the (analog of the) properties described in the proof of Part 1—Step 1.[14] Along with these beliefs, the strategy profile $\tilde{\sigma}$ is a MPE for $\Gamma_{J'}$ and sustains the same outcomes as σ in Γ_J.

Iterating across all i, t gives the result.

Step 2. We now prove that for any MPE $\sigma^D \in \mathcal{E}(\Gamma^D)$, there exists a MPE $\sigma^M \in \mathcal{E}(\Gamma^M)$ that sustains the same outcomes. The proof parallels that of Theorems 4 and 5.

Let Γ_J denote a game in which $\Phi_{j,\tau} = \Phi^M_{j,\tau}$ for all $(j,\tau) \in J$, while $\Phi_{j,\tau} = \Phi^D_{j,\tau}$ for all $(j,\tau) \in \mathcal{R}\backslash J$, for some $J \subset \mathcal{R} \cup \{\varnothing\}$ with $\mathcal{R} \equiv \mathcal{N} \times \mathcal{T}$. We prove the result by showing that, given any MPE $\sigma \in \mathcal{E}(\Gamma_J)$, there exists an MPE $\tilde{\sigma} \in \mathcal{E}(\Gamma_{J'})$, with $J' = J \cup \{i,t\}$ for some $\{i,t\} \in \mathcal{R}\backslash J$, that sustains the same outcomes.

The (Markov) strategy profile $\tilde{\sigma}$ is constructed from σ as follows. For any $(j,\tau) \neq (i,t)$, $\tilde{\sigma}_{j,\tau} = \sigma_{j,\tau}$. For $(j,\tau) = (i,t)$, the strategy $\tilde{\sigma}_{i,t}$ is such that, for any measurable set $R \subseteq \Phi^M_{i,t}$ and any $z_{i,t} \in Z_{i,t}$

$$\tilde{\sigma}_{i,t}(R \mid z_{i,t}) = \sigma_{i,t}\left(\bigcup_{\phi^M_{i,t} \in R} \{\phi^D_{i,t} : \mathrm{Im}(\phi^D_{i,t}) = \mathrm{Im}(\phi^M_{i,t})\} \mid z_{i,t}\right).$$

Next, consider the agent. Let

$$\bar{\Phi}^M_{i,t} \equiv \{\phi^M_{i,t} : \mathrm{Im}(\phi^M_{i,t}) = \mathrm{Im}(\phi^D_{i,t}) \text{ for some } \phi^D_{i,t} \in \Phi^D_{i,t}\}$$

[14] Take a $z_{i,\tau}$ such that, given $(\phi_{i,l})^{\tau-1}_{l=1}$, $z_{i,\tau}$ is consistent with σ_A and σ_{-i}. If $\tau > t$ and $I_t = i$, then it is no longer true that P_i's marginal beliefs over Θ^E_τ are a convex combination of her beliefs in Γ^J. However, because in this case A will never contract again with P_i, this is irrelevant for the result.

and for any $\phi_{i,t}^M \in \bar{\Phi}_{i,t}^M$, let $\Phi_{i,i}^D(\phi_{i,t}^M) \equiv \{\phi_{i,t}^D : \text{Im}(\phi_{i,t}^D) = \text{Im}(\phi_{i,t}^M)\}$. At any $\tau \neq t$, $\tilde{\sigma}_A$ induces the same behavior as σ_A in Γ_J (recall that σ_A is Markov). At $\tau = t$, for any $(\phi_{i,t}^M, (\phi_{j,t})_{j\neq i})$ such that $\phi_{i,t}^M \in \bar{\Phi}_{i,t}^M$, A uses the conditional distribution $\sigma_i(\cdot \mid \Phi_{i,i}^D(\phi_{i,t}^M))$ to determine his participation decision. That is, with probability $\sigma_i(\phi_{i,t}^D \mid \Phi_{i,i}^D(\phi_{i,t}^M))$, A behaves according to the participation strategy $w^t(\theta_t^E, \phi_{i,t}^D, (\phi_{j,t})_{j\neq i})) \in \Delta(\mathcal{N} \cup \varnothing)$ as if the game were Γ_J and the mechanisms offered by the principals were $(\phi_{i,t}^D, (\phi_{j,t})_{j\neq i})$. In case the lottery $w^t(\theta_t^E, \phi_{i,t}^M, (\phi_{j,t})_{j\neq i}))$ selects P_i, A then also induces the same distribution over $Y_{i,t}$ as in Γ_J given $(\theta_t^E, \phi_{i,t}^D)$, where $\phi_{i,t}^D$ is the same mechanism selected by the distribution $\sigma_i(\cdot \mid \Phi_{i,i}^D(\phi_{i,t}^M))$. If instead, $w^t(\theta_t^E, \phi_{i,t}^M, (\phi_{j,t})_{j\neq i}))$ selects a P_j with $j \neq i$, then A uses the same Markov strategy as in Γ_J to select which messages to send to P_j. In either case, the agent's choice of effort is governed by the same Markov strategy as in Γ_J.

Next, consider a $(\phi_{i,t}^M, (\phi_{j,t})_{j\neq i})$ such that $\phi_{i,t}^M \notin \bar{\Phi}_{i,t}^M$. At any downstream information set A behaves as if the game were Γ_J and the direct mechanism offered by P_i were $\phi_{i,t}^D$ where $\phi_{i,t}^D$ is obtained from $\phi_{i,t}^M$ as follows:

$$\phi_{i,t}^D(\theta_t^E) \in \arg \max_{\delta_{i,t} \in \text{Im}(\phi_{i,t}^M)} V(\theta_t^E, \delta_{i,t}, \sigma_t^+) \quad \forall \theta_t^E \in \Theta_t^E$$

where $V(\theta_t^E, \delta_{i,t}, \sigma_t^{D+})$ denotes the agent's continuation payoff in $\Gamma_{J'}$ when his extended type is θ_t^E, he chooses to participate in P_i's mechanism and the principals' downstream strategies are σ_t^+.[15]

Because all players' strategies are Markov, the principals' strategies in $\tilde{\sigma}$ can be sustained by beliefs over upstream histories that satisfy the analog of the properties in Part 1—Step 1. Together with these beliefs, the strategy profile $\tilde{\sigma}$ is a MPE for $\Gamma_{J'}$ and sustains the same outcomes as σ in Γ_J. ■

A2-4. Sequential offering as opposed to sequential contracting

Finally, consider an environment in which principals offer their mechanisms sequentially, but where the agent sends the messages $(m_1, ..., m_n)$ simultaneously at $t = n + 1$. Assume that any P_t, $t = 2, ..., n$, observes the mechanisms ϕ_t^- selected upstream before choosing her own mechanism. A (pure) strategy for P_i thus consists of a function $\sigma_i : \Phi_i^- \to \Phi_i$ such that $\sigma_i(\phi_i^-)$ is the mechanism offered by P_i when the profile of upstream mechanisms is ϕ_i^-.

Since the agent's decisions are now taken only at the end of the game, the definition of extended type must be modified as follows. For any $i = 1, ..., n$, let $\theta_i^E \equiv (\theta, \delta_{-i})$ with $\delta_{-i} \equiv (\delta_j)_{j\neq i}$. From the perspective of P_i, the agent's extended type thus consists of his exogenous type θ along with the lotteries δ_{-i} he is inducing at $t = n + 1$ with the other principals. An extended direct mechanism $\phi_i^D : \Theta_i^E \to D_i$ is then defined as in the benchmark model. The definition of incentive-compatibility and truthful equilibrium must however be adjusted as follows. Let $V(\theta, \delta)$ denote the maximal payoff that type θ can obtain by choosing the lotteries δ.

[15] Because all principals' strategies are Markov, V depends on any upstream history only through θ_t^E.

Definition A2. *(i) A mechanism ϕ_i^D is incentive-compatible if and only if, for any $\theta_i^E \in \Theta_i^E$,*

$$\phi_i^D(\theta_i^E) \in \arg\max_{\delta_i \in \mathrm{Im}(\phi_i^D)} V(\theta_i^E, \delta_i)$$

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

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

(iii) A strategy profile $\sigma^D \in \mathcal{E}(\Gamma^D)$ is a pure-strategy truthful equilibrium of Γ^D if and only if it is a pure-strategy equilibrium in which, given any profile of mechanisms ϕ^D such that $|\{j \in \mathcal{N} : \phi_j^D \neq \sigma_j^D(\phi_j^{D-})\}| \leq 1$, the agent's strategy is truthful in every mechanism ϕ_i^D for which $\phi_i^D = \sigma_j^D(\phi_j^{D-})$.

A mechanism ϕ_i^D is thus incentive-compatible if and only if, conditional on being a type θ and choosing the lotteries δ_{-i} with all principals other than i, the lottery $\delta_i = \phi_i^D(\theta_i^E)$ that the agent obtains by reporting $\theta_i^E \equiv (\theta, \delta_{-i})$ truthfully to P_i leads to an expected payoff for the agent that is at least as high as the one that he obtains by reporting any other $\hat{\theta}_i^E \in \Theta_i^E$. Given a profile ϕ^D of extended direct mechanisms, the agent's strategy is then truthful in ϕ_i^D if the message each type θ sends to P_i coincides with his true type along with the true decisions $\delta_{-i} = \phi_j^D(m_j^D))_{j \neq i}$ that he induces (by sending the messages m_{-i}^D) to the other principals. A strategy profile $\sigma^D \in \mathcal{E}(\Gamma^D)$ is a pure-strategy truthful equilibrium of Γ^D if and only if, whenever at most one principal deviated from her equilibrium strategy (i.e. offered a mechanism $\phi_j^D \neq \sigma_j^D(\phi_j^{D-})$), the agent's strategy at $t = n + 1$ is truthful in the mechanisms of any of the principals who conformed to the equilibrium strategy.

The following is then a natural adaptation of the notion of Markov strategies to this setting.

Definition A3. *Let Γ be a game with arbitrary choice sets for the principals. Given any pure-strategy profile $\sigma \in \mathcal{E}(\Gamma)$, we say that the agent's strategy σ_A is Markov with P_i if and only if, for any $\theta \in \Theta$, $\delta_{-i} \in D_{-i}$ and $\phi_i \in \Phi_i$, there exists a unique lottery $\delta_i(\theta, \delta_{-i}; \phi_i) \in \mathrm{Im}(\phi_i)$ such that A always selects $\delta_i(\theta, \delta_{-i}; \phi_i)$ with P_i when the latter offers the mechanism ϕ_i, the agent's type is θ and the decisions A induces with the other principals are δ_{-i}. We then say that the agent's strategy is Markov if and only if it is Markov with all P_i, $i \in \mathcal{N}$.*

We then have the following result.

Theorem 8 (Sequential offering). *(Part I: Menus) Let $\Gamma \succcurlyeq \Gamma^M$. For any $\sigma \in \mathcal{E}(\Gamma)$ in which all principals' strategies are pure, there exists a $\sigma^M \in \mathcal{E}(\Gamma^M)$ that sustains the same outcomes. Furthermore, any SCF π that can be sustained as an equilibrium of Γ^M can be sustained as an equilibrium of Γ.*

(Part II: Direct Mechanisms) For any pure-strategy equilibrium $\sigma^M \in \mathcal{E}(\Gamma^M)$ in which the agent's strategy is Markov, there exists a pure-strategy truthful equilibrium $\sigma^D \in \mathcal{E}(\Gamma^D)$ that sustains the same outcomes.

Proof of Theorem 8. Part I: Menus. The proof parallels that of Part I in Theorem 6 and is thus omitted (one can easily verify that the proof is actually simpler when the agent takes decisions only at $t = n+1$).

Part II: Direct Mechanisms. We show that, for any pure-strategy $\sigma^M \in \mathcal{E}(\Gamma^M)$ in which the agent's strategy is Markov, there exists a pure-strategy truthful equilibrium $\sigma^D \in \mathcal{E}(\Gamma^D)$ that sustains the same outcomes.

Consider a game Γ_J in which $\Phi_j = \Phi_j^D$ for all $j \in J$ while $\Phi_j = \Phi_j^M$ for all $j \in \mathcal{N} \backslash J$, for some $J \subset \mathcal{N} \cup \{\varnothing\}$. We prove the result by showing that given any pure-strategy equilibrium $\sigma \in \mathcal{E}(\Gamma_J)$ in which the agent's strategy is Markov there exists a pure-strategy equilibrium $\mathring{\sigma} \in \mathcal{E}(\Gamma_{J'})$ in which the agent's strategy is also Markov that sustains the same outcomes as σ, for any $J' = J \cup \{t\}$ with $t \in \mathcal{N} \backslash J$. The construction of $\mathring{\sigma}$ will also reveal that the strategy profile σ^D obtained from σ^M by iterating across all t,starting from $t = 1$ and moving forward, is such that σ_A^D is truthful.

Consider the following (pure) strategy for P_t in $\Gamma_{J'}$. For any profile of upstream mechanisms ϕ_t^-, let $\phi_t^M = \sigma_t(\phi_t^-)$ denote the equilibrium menu that P_t would have offered in Γ_J in response to ϕ_t^-. The extended direct mechanism $\phi_t^D = \mathring{\sigma}_t(\phi_t^-)$ that P_t offers in $\Gamma_{J'}$ in response to ϕ_t^- is such that, for any $\theta_t^E \in \Theta_t^E$,

$$\phi_t^D(\theta_t^E) = \delta_t(\theta, \delta_{-t}; \sigma_t(\phi_t^-))$$

Clearly, $\phi_t^D = \mathring{\sigma}_t(\phi_t^-)$ is incentive-compatible. Now consider the following strategy profile $\mathring{\sigma}$ for $\Gamma_{J'}$. For all principals P_j with $j < t$, simply let $\mathring{\sigma}_j = \sigma_j$. For P_t, let $\mathring{\sigma}_t$ be the strategy described above. Finally, for any P_j with $j > t$, $\mathring{\sigma}_j$ is constructed from σ_j as follows. If ϕ_j^- is such that in period t, P_t offered the mechanism $\phi_t^D = \mathring{\sigma}_t(\phi_t^-)$, then

$$\mathring{\sigma}_j(\phi_t^-, \phi_t^D, \phi_{t+1}, ..., \phi_{j-1}) = \sigma_j(\phi_t^-, \sigma_t(\phi_t^-), \phi_{t+1}, ..., \phi_{j-1}).$$

If instead, $\phi_t^D \neq \mathring{\sigma}_t(\phi_t^-)$, then

$$\mathring{\sigma}_j(\phi_t^-, \phi_t^D, \phi_{t+1}, ..., \phi_{j-1}) = \sigma_j(\phi_t^-, \phi_t^M, \phi_{t+1}, ..., \phi_{j-1}).$$

where ϕ_t^M is the menu whose image is $\mathrm{Im}(\phi_t^M) = \mathrm{Im}(\phi_t^D)$.

Next, consider the agent. Given any profile of mechanisms $(\phi_t^-, \phi_t^D, \phi_{t+1}, ..., \phi_n)$ such that $\phi_t^D = \mathring{\sigma}_t(\phi_t^-)$, at $t = n+1$ each type θ of the agent induces the same outcomes he would have induced in Γ_J had the mechanisms offered been $(\phi_t^-, \sigma_t(\phi_t^-), \phi_{t+1}, ..., \phi_n)$. Note that this can be achieved by reporting $(\theta, (\phi_j(m_j))_{j \neq t})$ truthfully to P_t. If, instead, $\phi_t^D \neq \mathring{\sigma}_t(\phi_t^-)$, then A induces the same outcomes he would have induced in Γ_J had the mechanisms offered been $(\phi_t^-, \phi_t^M, \phi_{t+1}, ..., \phi_n)$, where ϕ_t^M is the menu whose image is $\mathrm{Im}(\phi_t^M) = \mathrm{Im}(\phi_t^D)$. Clearly, this strategy is sequentially optimal for the agent. Furthermore, given $(\mathring{\sigma}_A, \mathring{\sigma}_{-i})$, no principal has a profitable deviation. We conclude that the strategy profile $\mathring{\sigma}$ constructed this way is an equilibrium for $\Gamma_{J'}$ and induces the same outcomes as σ in Γ_J.

Iterating across all periods, starting from $t = 1$ and letting $J = \{\varnothing\}$ and proceeding forward by letting $J' = J \cup \{t+1\}$, gives a pure-strategy truthful equilibrium of Γ^D that sustains the same outcomes as σ^M. ■

References

Pavan, A. and G. Calzolari, (2008), "Truthful Revelation Mechanisms for Simultaneous Common Agency Games," w.p. Northwestern University and the University of Bologna.