Note

On the use of menus in sequential common agency [☆]

Giacomo Calzolari [a], Alessandro Pavan [b,*]

[a] *Department of Economics, University of Bologna, Piazza Scaravilli 2, 40126, Bologna, Italy*
[b] *Department of Economics, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208, USA*

Abstract

We illustrate, by means of two examples, why assuming the principals offer simple menus (i.e. collections of payoff-relevant alternatives) as opposed to more general mechanisms may preclude a complete characterization of the set of equilibrium outcomes in certain sequential contracting environments. We then discuss how refinements of the solution concept, or enrichments of the menus that allow for recommendations, may restore the possibility of using menus to obtain a complete equilibrium characterization.
© 2008 Elsevier Inc. All rights reserved.

1. Introduction

When multiple principals contract simultaneously with the same agent, Peters (2001) and Martimort and Stole (2002) have proved the following result: For any equilibrium relative to any game with arbitrary sets of mechanisms for the principals, there exists an equilibrium in the game in which the principals are restricted to offer the agent the menus of payoff-relevant alternatives they could have offered in the original game that sustains the same outcomes. This result is referred to in the literature as the Delegation Principle (or the Menu Theorem) and has proved useful in applications.

[☆] A previous version was entitled "On the Validity of Menu Theorems in Sequential Common Agency Games."

[*] Corresponding author.

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

Recent years have witnessed interest in environments in which contracting is sequential, in the sense that the agent contracts with his multiple principals at different points in time.[1] In light of this fast growing literature, it is important to understand whether restricting the principals to offer menus of payoff-relevant alternatives, as opposed to more general mechanisms, is without loss of generality also when contracting is sequential.

In this note we construct two examples that illustrate why simple menus may fail to sustain all equilibrium outcomes in certain sequential contracting environments.

The first example features a situation in which the principals observe the payoff-relevant decisions taken upstream, but not necessarily the mechanisms used to select them. In this setting, restricting the principals to offer menus may mean restricting the extent to which different principals can have diverging beliefs about the specific mechanism used upstream to select an off-equilibrium decision. When the agent's strategy is not Markov (i.e. it may depend on the entire upstream history), this means imposing restrictions on the principals' expectations about the agent's behavior downstream. Such restrictions may preclude the possibility of sustaining certain outcomes.

The second example features a situation in which all principals observe both the mechanisms and the payoff-relevant decisions selected upstream. The reason why simple menus fail in this example is that they do not permit the principals to use payoff-irrelevant information as a device to correlate their decisions. In the absence of alternative instruments such as sunspots or cheap talk messages, this means restricting the possible outcomes.

These examples warn against the use of simple menus in certain sequential contracting environments. However, there are situations of interest for applications in which the problems indicated by these examples never arise. Furthermore, there are ways of enriching the menus (for example allowing the principals to send each other, and/or the agent, recommendations about the decisions to take downstream) that may restore the possibility of using menus to sustain all equilibrium outcomes. We discuss some of these issues at the end of the note.

2. Simple menus

This section contains two examples that illustrate why simple menus may not sustain all equilibrium outcomes when contracting is sequential.

2.1. Out-of-equilibrium beliefs

Consider a game in which four principals contract sequentially with the same agent, A. The game has four stages. At each stage, a different principal, P_i, $i = 1, \ldots, 4$, contracts with A. Each principal must select a payoff-relevant alternative a_i (also referred to as a decision) from a set $\mathcal{A}_i$ of feasible alternatives. Depending on the application of interest, a_i can be a policy, a level of trade, or the decision to undertake a project.

The selection of a_i is obtained through a mechanism; the latter consists of a set of possible messages $\mathcal{M}_i$ along with a mapping $\phi_i : \mathcal{M}_i \to \mathcal{A}_i$ such that, when A sends the message $m_i \in \mathcal{M}_i$, P_i responds by selecting the alternative $a_i \in \mathcal{A}_i$.

In this example, the sets of feasible alternatives are $\mathcal{A}_i = \{b_i, c_i\}$, for $i = 1, \ldots, 3$, and $\mathcal{A}_4 = \{d, e, f\}$. The principals' and the agent's payoffs are described by the quintuples

---

[1] We refer the reader to Pavan and Calzolari (2007) for a discussion of the relevance of sequential contracting in applications.

Fig. 1. Payoffs.

$(u_1, u_2, u_3, u_4, u_A)$ in Fig. 1; note that, when all players' payoffs are independent of the decisions taken after period $t = 1, \ldots, 3$, the tree has been cut to highlight directly the final payoffs.

Before choosing her mechanism, each downstream principal observes all the payoff-relevant decisions taken upstream. Furthermore each downstream principal observes all upstream mechanisms, with the exception of the mechanism selected by P_1. None of the principals observes the messages sent by the agent to the other principals.

Now consider the game Γ in which the sets of feasible mechanisms are $\Phi_1 = \{\phi_1^{b_1}, \underline{\phi}_1^{c_1}, \bar{\phi}_1^{c_1},$ $\phi_1^{b_1,c_1}\}$ for P_1, $\Phi_i = \{\phi_i^{b_i}, \phi_i^{c_i}, \phi_i^{b_i,c_i}\}$ for P_2 and P_3, and $\Phi_4 = \{\phi_4^{d}, \phi_4^{e}, \phi_4^{f}, \phi_4^{d,e}, \phi_4^{d,f}, \phi_4^{e,f},$ $\phi_4^{d,e,f}\}$ for P_4. Superscripts denote the image (i.e. the range) of the mechanism; for example, $\text{Im}(\phi_i^{d,e}) = \{d, e\}$. Note that, for any $i \neq 1$, and any menu of payoff-relevant alternatives $\psi_i \in 2^{\mathcal{A}_i} \setminus \{\emptyset\}$ there exists one and only one mechanism $\phi_i^{\psi_i} \in \Phi_i$ such that $\text{Im}(\phi_i^{\psi_i}) = \psi_i$. In contrast, for $i = 1$, there are two mechanisms $\underline{\phi}_1^{c_1}$ and $\bar{\phi}_1^{c_1}$ whose range is $\{c_1\}$.

In this game, the outcome $(b_1, b_2, \bar{b}_3, d)$ can be sustained as a pure-strategy perfect Bayesian equilibrium. It suffices to take *any* strategy profile with the following properties: P_1 offers the mechanism $\phi_1^{b_1}$; regardless of a_1, P_2 offers $\phi_2^{b_2}$; regardless of (a_1, a_2, ϕ_2), P_3 offers $\phi_3^{b_3}$; regardless of (ϕ_2, ϕ_3), P_4 offers $\phi_4^{d,e,f}$ if $(a_1, a_2, a_3) = (c_1, b_2, b_3)$ and ϕ_4^{d} otherwise; at $t = 1$, A chooses c_1 if P_1 offers $\phi_1^{b_1,c_1}$; at $t = 2$, A chooses c_2 if P_2 offers $\phi_2^{b_2,c_2}$; at $t = 3$, A chooses c_3 if P_3 offers $\phi_3^{b_3,c_3}$; at $t = 4$, when offered the mechanism $\phi_4^{d,e,f}$, A chooses e if $\phi_1 = \underline{\phi}_1^{c_1}$, f if $\phi_1 = \bar{\phi}_1^{c_1}$ and d if either $\phi_1 = \phi_1^{b_1}$ or $\phi_1 = \phi_1^{b_1,c_1}$. Clearly, this is only a partial description of the strategy profile; however, such a description contains all information that is relevant for the result we want to establish. The aforementioned strategy profile is sustained by the following out-of-equilibrium beliefs: after observing c_1, P_2 believes that $\phi_1 = \bar{\phi}_1^{c_1}$, P_3 believes that $\phi_1 = \phi_1^{b_1,c_1}$, and P_4 believes that $\phi_1 = \underline{\phi}_1^{c_1}$.

We claim that the outcome (b_1, b_2, b_3, d) cannot be sustained in the "menu game" Γ^M in which the principals are restricted to offer the menus of payoff-relevant alternatives they could have offered in Γ and delegate to the agent the choice of the decisions. We prove the result by showing that this outcome requires that, after observing a deviation to c_1, the supports of P_2's, P_3's, and P_4's beliefs about the mechanism used by P_1 not overlap, which clearly cannot be the case in Γ^M.

To see this, note that P_1 (weakly) prefers b_1 to c_1 if and only if c_1 is followed by b_2 and b_3. Hence, for (b_1, b_2, b_3, d) to be an equilibrium outcome in Γ^M, it must be that, after observing a

deviation to c_1, P_2 and P_3 offer menus that contain, respectively, b_2 and b_3, and that A selects b_2 and b_3 in each of these menus. Furthermore, because A strictly prefers c_2 to b_2 after c_1 and strictly prefers c_3 to b_3 after (c_1, b_2), it must be that P_2 offers the (degenerate) menu $\{b_2\}$ after observing c_1 and that P_3 offers the (degenerate) menu $\{b_3\}$ after observing the menu $\{b_2\}$ and the decisions (c_1, b_2). Because P_2 can always guarantee herself a payoff of 2 by choosing c_2 after c_1, for her to offer the menu $\{b_2\}$ it must be that she expects A to choose b_3 with P_3 and f with P_4. Similarly, because P_3 can always guarantee herself a payoff of 2 by choosing c_3 after (c_1, b_2), for her to offer the menu $\{b_3\}$ it must be that she expects A to choose d with P_4. Lastly, because P_4 can always guarantee herself a payoff of 2 by choosing e after (c_1, b_2, b_3), for her to offer any other menu, it must be that this menu contains e and that she expects A to choose e from the menu.

We conclude that any strategy profile that sustains (b_1, b_2, b_3, d) as an equilibrium in Γ^M must satisfy the following properties: P_2 offers the menu $\{b_2\}$ after observing c_1; P_3 offers the menu $\{b_3\}$ after observing the menu $\{b_2\}$ and the decisions (c_1, b_2); P_4 offers the menu $\{e, f, g\}$ after observing the menus $\{b_2\}$ and $\{b_3\}$ and the decisions (c_1, b_2, b_3).

Now, given the decisions (c_1, b_2, b_3) and the upstream menus $\{b_2\}$ and $\{b_3\}$, the agent's behavior at $t = 4$ may vary only on the basis of the particular menu containing c_1 offered at $t = 1$. Because in Γ^M there are only two such menus, there are only two possible behavioral strategies that A can follow at $t = 4$ given the decisions (c_1, b_2, b_3) and the menus $\{b_2\}$ and $\{b_3\}$. It is thus impossible that P_2, P_3 and P_4 expect A to choose respectively f, d, and e with probability one when offered the menu $\{e, f, g\}$. The outcome (b_1, b_2, b_3, d) can thus be sustained in Γ but not in the "menu game" Γ^M.

2.2. Correlation devices

We now illustrate a second reason why menus may not sustain all possible outcomes. Consider an environment in which three principals contract sequentially with the same agent. The sets of feasible alternatives are $\mathcal{A}_1 = \{t\}$ and $\mathcal{A}_2 = \mathcal{A}_3 = [0, 1]$. For simplicity, assume P_1's and P_3's payoffs are constant over $\mathcal{A} \equiv \mathcal{A}_1 \times \mathcal{A}_2 \times \mathcal{A}_3$, whereas P_2's and A's payoffs are respectively $u_2 = 2a_2a_3 + (1 - a_2)(1 - a_3)$ and $u_A = a_3(1 + a_2)$. In this environment, all principals observe all mechanisms selected upstream. Whether they also observe the payoff-relevant decisions taken in these mechanisms is not important in this example but, to fix ideas, assume they do.

Now consider a game Γ in which the sets of feasible mechanisms are $\Phi_1 = \{\bar{\phi}_1, \underline{\phi}_1\}$ and $\Phi_i = \{\phi_i^{a_i}\}_{a_i \in \mathcal{A}_i}$, $i = 2, 3$, where $\phi_i^{a_i}$ is a mechanism such that $\text{Im}(\phi_i^{a_i}) = a_i$.

The following is a perfect Bayesian equilibrium in Γ: P_1 randomizes over $\bar{\phi}_1$ and $\underline{\phi}_1$ with probability $q \in (0, 1)$ and $1 - q$, respectively; given $\bar{\phi}_1$, P_2 chooses $a_2 = 1$, whereas given $\underline{\phi}_1$ she chooses $a_2 = 0$; at $t = 3$, regardless of (ϕ_2, a_2), P_3 chooses $a_3 = 1$ if $\phi_1 = \bar{\phi}_1$ and $a_3 = 0$ if $\phi_1 = \underline{\phi}_1$. The equilibrium outcome is $(t, 1, 1)$ with probability q and $(t, 0, 0)$ with probability $1 - q$. This outcome cannot be sustained in the "menu game" Γ^M in which the principals offer the menus of payoff-relevant alternatives they could have offered in Γ and delegate to the agent the choice of the decisions. The reason is that Γ^M does not permit P_1 to correlate the other principals' decisions. The role of P_1 as a correlation device is key to sustain the outcome described above and cannot be replicated by P_2. In fact, P_2 finds it optimal to match the decision taken by P_3—and hence to respond to the mechanism selected by P_1 anticipating how P_3 responds to it—but is never willing to mix over $\mathcal{A}_2$, for she strictly prefers $(a_2, a_3) = (1, 1)$ to $(a_2, a_3) = (0, 0)$.

3. Discussion

The equilibrium constructed in the first example is a (weak) perfect Bayesian equilibrium but neither a sequential nor a Markov-perfect equilibrium. Imposing such refinements may restore the possibility of using menus to sustain all outcomes. In this respect, the example highlights an important difference between simultaneous and sequential common agency. In the former, the validity of the Delegation Principle is independent of whether one is interested in all perfect Bayesian equilibrium outcomes or only in outcomes that can be sustained by imposing refinements such as sequential or Markov-perfect equilibrium. This is not the case in the latter.

Next, consider the second example. If public sunspots are available, then restricting the principals to offer simple menus may not pose any problem. The role of the example is to warn against the use of simple menus in environments in which alternative correlation devices are not available—another important difference with respect to the simultaneous case.

Also note that the notion of menus considered here is the one used in applications: a menu is a collection of payoff-relevant alternatives.[2] The problems highlighted by our examples vanish if one considers more general menus that allow the principals to send each other, and/or the agent, recommendations about the decisions to take downstream. For instance, in the first example, the two mechanisms $\underline{\psi}_1 = \{c_1\}$ and $\bar{\psi}_1 = \{c_1\}$ can be replaced by two menus that contain the same payoff-relevant decision but two different recommendations to the agent about the strategy to follow downstream. Provided that these recommendations are private (in the sense that they are not observed by the downstream principals), then the outcome (b_1, b_2, b_3, d) can be sustained also in the menu game with recommendations.[3] Similarly, the outcome in the second example can be sustained by letting the first principal send public (perfectly correlated) recommendations to the downstream principals.

These enriched menus are more similar to Myerson (1982) generalized direct revelation mechanisms than to simple menus such as price-quantity schedules, as typically used in applications. As shown in Peters (2001), allowing for such enriched menus may be necessary when the agent exerts some effort after communicating with the principals. In Peters' environment, contracting is simultaneous and recommendations are used to fashion the agent's beliefs about the principal's response to the agent's effort. Because this is the only role that recommendations play in simultaneous games, such recommendations can be dispensed with if one allows the principals to offer menus of lotteries over contracts (as opposed to menus of deterministic decisions).[4] The role of recommendations in sequential contracting is different: in the first example, recommendations are used to permit the downstream principals to have diverging beliefs about the agent's behavior in downstream relationships. In the second example, recommendations are used to correlate the principals' decisions. Furthermore, it may not suffice to introduce lotteries to dispense with such recommendations, as it can be seen by considering the second example where there is a single such lottery.

While our examples warn against the use of simple menus in certain environments, there are situations in which simple menus do sustain all equilibrium outcomes. As shown in Pavan and Calzolari (2007), this is always the case when contracting is private, i.e. when downstream principals observe neither the mechanisms nor the payoff-relevant decisions selected upstream.

---

[2] This is also the notion used by Martimort and Stole (2002) to establish the Delegation Principle.

[3] We thank a referee for suggesting this possibility.

[4] In a moral hazard setting, a decision should be interpreted as a contract that specifies the principal's action as a function of some verifiable performance measure correlated with the agent's effort.

Furthermore, even when contracting is not private, all equilibrium outcomes sustained by Markov strategies can be sustained with simple menus. Because Markov strategies are often considered a focal class, simple menus may retain a significant appeal in many applications, even if they do not sustain all equilibrium outcomes.

Acknowledgments

References

Martimort, D., Stole, L., 2002. The revelation and delegation principles in common agency games. Econometrica 70, 1659–1674.

Myerson, R., 1982. Optimal coordination mechanisms in generalized principal–agent problems. J. Math. Econ. 10, 67–81.

Pavan A., Calzolari, G., 2007. Sequential contracting with multiple principals. Mimeo, Northwestern University and the University of Bologna.

Peters, M., 2001. Common agency and the revelation principle. Econometrica 69, 1349–1372.