Common Agency and the Revelation Principle

Michael Peters*
Department of Economics
University of Toronto
150 St. George St.
Toronto, Canada
M5S 3G7

April 27, 2000

Abstract

In the common agency problem multiple mechanism designers simultaneously attempt to control the behavior of a single privately informed agent. The paper shows that the allocations associated with equilibria relative to any ad hoc set of feasible mechanisms can be reproduced as equilibria relative to (some subset of) the set of menus. Furthermore, equilibria relative to the set of menus are weakly robust in the sense that it is possible to find continuation equilibria so that the equilibrium allocations persist even when the set of feasible mechanisms is enlarged.

In agency problems with a single principal, no loss of generality is involved when the principal is restricted to offer standard *direct mechanisms*. When there are multiple principals things are more complicated because of the fact that the agent knows about the mechanisms that have been offered by all the principles at the time that he communicates with any particular principal. A principal who recognizes this can construct a mechanism that makes the

*Thanks are due to Larry Epstein, Meg Meyer, Carolyn Pitchik, Lars Stole and Sergei Severinov for helpful comments. The financial support of the Social Sciences and Humanities Research Council of Canada is gratefully acknowledged.

allocation that he selects depend on the mechanisms that the other principals have offered.[1] To describe feasible mechanisms in full generality gives rise to an infinite regress since the principal's mechanism could depend on whether other principals' mechanisms depends on his mechanism, and so on. This problem is more that a simple theoretical curiosity. A variety of examples are cited in Martimort and Stole (1997), Peck (1994) and Epstein and Peters (1997) illustrating equilibrium allocations that cannot be supported unless these considerations are taken into account. The same sources also provide examples of allocations that are equilibrium allocations relative to the 'usual' set of direct mechanisms[2], but which do not survive when mechanism designers are allowed to use more complex message spaces that allow the principal to take advantage of agents' information about other mechanisms.

At a conceptual level, this problem has been resolved. Epstein and Peters (1997) provide a set of direct mechanisms and types that make it possible to use the revelation principle in multiple principal problems the same way that it is used in single principal problems. However, the set of mechanisms and types involved is complex and hard to apply because the agent's type must describe the mechanisms being used by the other principals, whether these other mechanisms depend on other principals mechanisms, and so on. The natural question to ask is whether there are restricted (but nonetheless interesting) economic environments in which the possibility that mechanisms depend on one another can be dealt with in a more tractable way.

Martimort and Stole (1997) make extensive use of the idea that principals could offer agents *menus* of alternatives. In their examples, the agent's preferences over alternatives in the principal's menu depend on the mechanisms that are offered by other principals. This makes it possible for the principal to design a mechanism that is responsive to the mechanism offered by the other principal. This paper shows that, in very general common agency problems in which there is only a *single* agent, principals do not need any more than menus of alternatives to respond to the mechanisms offered by the other principals.

In a single principal *single* agent setting, any contract is equivalent to a contract in which the principal offers the agent a menu of alternatives. The menu is simply all of the outcomes that the agent can attain by sending

[1] The basic problem with direct mechanisms in competing mechanism design problems was first observed by McAfee (1993) and Katz (1987).

[2] The 'usual' set of direct mechanisms simply ask agents to report information about their preferences, as is done, for example, in McAfee (1993)

different messages to the principal. This idea has been referred to as the *taxation principle* by Rochet (1986).[3] This paper shows that Rochet's taxation principle also works with multiple principals in the following sense: first, for any set of indirect mechanisms feasible for the mechanism designers, and for any equilibrium relative to that set, there is an equilibrium in (a subset of) *menus* which preserves the corresponding equilibrium allocation.[4] This theorem has two useful applications. First any 'interesting' equilibrium allocations can ultimately be discovered by considering equilibria relative to the set of menus - more complex mechanisms with high dimensional messages spaces like those described by Epstein and Peters (1997) are not required.[5] Secondly, alternative formulations of the set of feasible mechanisms for the principals can be evaluated by considering whether they generate equilibrium allocations that are also equilibrium allocations relative to the set of menus.[6]

This is not quite enough to guarantee that all equilibria relative to the set of menus are interesting. In principle, some equilibria might not survive the possibility that mechanism designers might dream up more complex and sophisticated mechanisms. This is an issue that does not arise in the single principal problem. The usual revelation principle translates each indirect mechanism and any form of agent behavior relative to that mechanism into a corresponding direct mechanism relative to which the agent reports truthfully and behaves as the principal recommends. Thus, once the principal has chosen the best direct mechanism, he can never gain by considering alternative indirect mechanisms. The reason is that any deviation to an indirect mechanism corresponds to a deviation that the principal could have taken within the set of direct mechanisms. When there are multiple principals, deviations into larger sets of mechanisms (for example mechanisms other than

[3]See also Guesnerie (1995) and Martimort and Stole (1997). It should be emphasized that the taxation principle applies only to situations in which there is a *single* agent. Offering menus to agents would be very restrictive when there are multiple agents (in an auction environment for example).

[4]The equilibrium may only be relative to a subset of the set of menus since the initial set of indirect mechanisms could be very restrictive. For example, forcing sellers to compete in price along when buyers have private information about their valuations.

[5]Again, it should be empasized that this result relies heavily on the assumption that there is a single agent. The argument in Epstein and Peters (1997) is designed for the case of multiple agents.

[6]For example, if equilibria are characterized where principals compete by offering, say, fixed prices, and these equilibrium allocations cannot be supported relative to the set of menus, then we can conclude that forcing principals to offer fixed prices is restrictive.

direct mechanisms or menus) may induce the agent to change the way he reports and behaves with *all* principals. In this sense the deviation may provide the principal a way of affecting the behavior of the other principals in a way that is unavailable to him in the original set of mechanisms. For this reason a second theorem is provided to show that equilibrium relative to the set of menus is *weakly robust* in the sense that there will always exist an equilibrium relative to the new set of feasible mechanisms such that equilibrium allocations relative to the set of menus are preserved.[7]

The approach presented in this paper is an alternative to the use of universal mechanisms and the revelation principle as developed in Epstein and Peters (1997). The approach here is simpler, but this does not come without cost. First, Epstein and Peters (1997) makes it possible to transform *every* indirect mechanism into a direct mechanism (with the appropriate types space) in a way that preserves payoffs. This approach provides a revelation principle exactly like the one for single agent problems. Menus can be used to preserve equilibrium allocations, but when arbitrary indirect mechanisms are transformed into menus, only the equilibrium path mechanisms can be transformed in a way that preserves payoff. The results here should not be interpreted as a revelation principle. Furthermore, when the set of direct mechanisms is enlarged in Epstein and Peters (1997), every equilibrium relative to the new enlarged set of mechanisms preserves the payoffs generated by equilibrium relative to the set of direct mechanisms (again, exactly as occurs with a single principal). No such result applies here, we only establish the existence of *some* equilibrium that preserves payoffs available with menus.

Finally, Martimort and Stole (1999b) and Martimort and Stole (1999a) present variants of the menus theorem in specialized environments involving restrictions to pure strategies, non-random contracts and restricted effort.

1 Basics

There is no single model of common agency in the literature, rather a variety of models all concerned with different aspects of the problem. The best known example of common agency is the model of Bernheim and Whinston (1987) or Bernheim and Whinston (1986) in which multiple principals attempt to influence the action taken by a single agent. The agent takes an action or

[7] Generally, not all equilibria relative to the new set of mechanisms will have this property.

effort that the principals all care about, while the principals can transfer money to the agent. The agent's effort is not contractible ex ante, but it is observable and principals are able to write enforceable contracts specifying the payments that they will make to the agent contingent on the effort that he or she does take. The agent chooses the action that maximizes the sum of his or her payments less the cost of effort. Knowing this, each principal offers a schedule that is a best reply to the schedules being offered by the other principals. The agent's effort is given a variety of interpretations, for example, Dixit, Grossman, and Helpman (1997) interpret the action as a government policy, like a set of environmental regulations, or an array of tariffs. They extend the arguments to consider the possibility that the agent's utility is non-linear in total payments. This extension is important because it admits the possibility that the principals take actions that are important to the agent. For example, when industry groups (the principals) lobby the government (the agent) for tariff protection, the outputs that they subsequently produce will depend both on the tariff the government sets, and potentially on the bribes that they are forced to pay (possibly in the form of employment guarantees and so forth) to get this tariff protection. If the government cares about the levels of output that the principals produce, its utility function will be non-linear in the payments that it receives.[8]

Parlour and Rajan (1997) provide another fruitful application of the common agency paradigm in which multiple lenders compete for a single borrower. The lenders can limit the amount they lend the borrower, but cannot control the borrowers transactions with other principals. Loan repayment is not enforceable, however, so the borrower (or agent) always has the option to default on the loans he has accepted.

Finally, the most basic competitive model can be interpreted as a common agency problem provided there is only a single buyer. For example, following Klemperer and Meyer (1989), sellers can commit themselves to the prices at which they trade. The buyer, acting as an agent, chooses the quantities that he or she wishes to buy from each of the sellers.

These models give the principals and the agent symmetric information about the agent's tastes since they are concerned primarily with the agent's effort. Biais, Martimort, and Rochet (1997) and Martimort and Stole (1999a) are concerned with problems where the agent has private information. In the first paper, multiple risk neutral traders in a financial market (the principals)

[8] See also Grossman and Helpman (1994).

buy and sell securities from a single risk averse trader (the agent) who has private information about the common value of the security being traded. The agent's non-contractible action is the quantity of the security that he or she buys from each of the risk neutral traders. The risk neutral traders offer non-linear pricing contracts that specify the per unit prices that they will charge or accept contingent on the quantity of the asset that the agent chooses to buy or sell from them. Martimort and Stole (1999a) consider a problem where multiple sellers (the principals) compete selling goods to a single buyer (the agent) who has private information about his or her taste. In their case, the quantity that the buyer purchases is contractible.

These models have the following general structure: there are n principals dealing with a single agent. Each principal $j \in \{1, \ldots n\}$ offers the agent a contract from some set $\mathcal{A}_j$, while the agent takes some effort from a set E. The set $\mathcal{A}_j$ consists of a some set of feasible mappings from E into the set of actions Y_j controlled by principal j. The agent's preferences are private information and are parameterized by elements in some set Ω. Principals commonly believe that the agent's preference parameter is distributed according to some distribution F on Ω.

Agents and principals have expected utility preferences. The payoff to principal $j \in \{1, \ldots, n\}$ is represented by $v_j : \prod_{k=1}^{n} \mathcal{A}_k \times E \times \Omega \longrightarrow [0,1]$. For the agent, payoffs are represented by the function $u : \prod_{k=1}^{n} \mathcal{A}_k \times E \times \Omega \longrightarrow [0,1]$.

The set $\mathcal{A}_j$ varies in different applications depending on the principal's ability to write contracts. In some cases the principal can make his action depend on every aspect of the agent's action, while in others he may be limited to contracts that depend on some component of the agents effort, or to actions that do not vary with the agent's action at all. In Bernheim and Whinston (1986) for example, the agent's effort is fully contractible ex post. The set of efforts E that the agent can take is interpreted as a tariff or some other government policy. The set Y is the set of feasible transfers of income from the principal to the agent. The principal is able to write contracts contingent of the full action of the agent, so for each principal j, $\mathcal{A}_j$ is the set of functions from E into Y describing the transfer that the principal will make to the agent for each possible effort level. The principal's payoff is $v_j(a_j(e), e)$ while the agent's payoff is $\sum_{j=1}^{n} a_j(e) - g(e)$ where $g(e)$ is the disutility associated with effort level e. In this formulation, observe that the principal's payoff depends only on his own action and the effort of the agent.

In Klemperer and Meyer (1989) on the other hand, the agent's (buyer's) effort is an n vector of quantities purchased from each seller, i.e., $e \in \mathbf{R}^n$, while the simple actions Y_j available to seller j are the per unit prices that they could charge. In this environment it is more natural to assume that seller j can only write contracts that specify price as a function of the quantity e_j that the buyer purchases from him. So $\mathcal{A}_j$ is just the set of maps from the j^{th} component of e into the set of per unit prices (or as Klemperer and Meyer put it, the set of *supply functions*).

In Biais, Martimort, and Rochet (1997), let $Q \subset \mathbf{R}$ be a set of feasible asset trades (allowing both positive and negative values so that either buying or selling is possible). Interpret e as a vector of n components $\{q_1, \ldots q_n\}$ of trades that the agent makes with each of the principals, so $E = Q^n$. Again, Y can represent the set of feasible transfers from the principal to the agent (and both positive and negative values for q and y are allowed - it would also be possible to interpret the principal as choosing unit prices). The principal cannot condition his transfer on transactions that the agent makes with other principals, so that $\mathcal{A}_j$ is the set of functions $a_j : E \to Y$ that depend only on the j^{th} component of e. The principal's contract describes the transfer from the principal to the agent for each possible quantity that the principal could buy from or sell to the agent. The asset being traded has a random value τ whose distribution depends on the agent's private information ω. Let U be the agent's utility for wealth function.[9] Then the agent's payoff function when his or her private information is ω is given by[10]

$$u\left(a_1, \ldots a_n, q_1, \ldots q_n, \omega\right) = \mathbf{E}_{\tau|\omega} U\left(\sum_{j=1}^{n} q_j \tau - \sum_{j=1}^{n} a_j\left(q_j\right)\right)$$

Principal j's payoff function is given by

$$v_j\left(a_1, \ldots a_n, q_1, \ldots q_n, \omega\right) = a_j\left(q_j\right) - q_j \mathbf{E}_{\tau|\omega} \tau$$

Finally, in Parlour and Rajan (1997) the principal offers some amount of credit $L \in \mathbf{R}$ and some interest rate $r \in \mathbf{R}$. So $Y = \mathbf{R}^2$. The agent decides whether or not to default on loans, so $E = \{0, 1\}$ with 0 interpreted

[9] Biais, Martimort, and Rochet (1997) assume that the agent has a utility function that has constant absolute risk aversion, but it is clearer here to use the more general form.

[10] Biais, Martimort, and Rochet (1997) assume that traders have a random endowment along with private information about the distribution of returns. To simplify, this is ignored in this expression.

as default. The principals action cannot depend at all on the agents effort so $\mathcal{A} = \mathbf{R}^2$ and the principal's payoff is $v_j(L,r,0) = -L$, $v_j(L,r,1) = rL$.

In each of these examples, each principal's payoff depends only indirectly on the contracts of the other principals. The indirect dependence comes about because the agent's effort depends on all these other actions. As Martimort and Stole (1997) have shown, difficult problems emerge when there is direct interdependence of principals' payoffs (this is illustrated by example shortly). This direct dependence is natural in many settings. For example, when industry groups lobby for tariffs or other government concessions, they often make commitments to future employment levels. These commitments should be modelled as part of the action that the government undertakes. However, these employment commitments will have output consequences that will directly affect the profitability of the firms' competitors. For this reason it seems sensible to make explicit allowance for this possibility in the payoff functions.

To avoid measure theoretic considerations, it will be assumed in the sequel that the sets $\mathcal{A}$, E and Ω are all finite.

2 Competition

The focus of this paper is on how to model competition among the principals. Ultimately the argument will be that it is 'safe' to model this competition as a competition where principals offer appropriately designed *menus* to the agent. To do this it is desirable to avoid imposing additional ad hoc restrictions on competition among principals - for example, forcing them to use non-random mechanisms, or focussing attention on pure strategy equilibria, even though these restrictions might be perfectly sensible in the context of specific applications. To deal with this, we revert to the formulation of Myerson (1979). The principals begin the process by designing *mechanisms* to guide subsequent communication. Without loss of generality, it can be assumed that the agent is bound by the mechanisms offered by the principals.[11] Once the principals have designed their mechanisms, agents communicate with the principals, possibly sending information about their types. The principals respond to the agents communications by taking actions, and possibly sending messages, according to the rules that they specified by their

[11] Endogenous participation decisions are simply incorporated into the agents uncontractible effort choice.

mechanism. The agent ends the process by taking whatever action he or she feels is appropriate - payoffs are then realized.

For any measurable set X, Let $\triangle(X)$ denote the set of probability measures on X. When a topology on $\triangle(X)$ is needed, the topology of weak convergence is assumed. For any $x \in \triangle(X)$, let $\operatorname{supp} x$ be the support of the distribution x. Let C be the (measurable) set of messages that the agent can send to a principal. Let R be the (measurable) set of responses that principals can send back to the agent. These message spaces are intended to be quite general in the degree and nature of the communication about the other principals' mechanisms that they permit. For example, the messages could allow the agent to communicate the mechanisms being used by other principals.

The role of the principal's messages is to provide the agent with information about the action that the principal has chosen. An indirect mechanism γ_j for principal j specifies a measurable map from C into $\triangle(\mathcal{A}\times\mathcal{R})$. For each message the agent sends, the principal responds by committing himself to a joint distribution over actions and messages. We illustrate one of the roles these messages might play in an example below. If the set of messages is finite, then each message that the agent receives from the principal induces a posterior belief about the action that the principal has taken. A posterior belief is a probability distribution over $\mathcal{A}$, so in this sense each mechanism induces a mapping from C into $\triangle(\mathcal{A}\times\triangle(\mathcal{A}))$.

More generally, if $\mathcal{R}$ is a some measurable space of messages and $P \in \triangle(\mathcal{A}\times\mathcal{R})$, then for any integrable function $x : \mathcal{A}\times\mathcal{R} \to \mathbf{R}$, there is a $\mathcal{R}$ measurable function $\mathbf{E}[x|r,P]$ satisfying $\int_D \mathbf{E}[x|r]\,dP = \int_D x dP$ for all $D \subset \mathcal{R}$, interpreted as the conditional expectation of x given r.[12] For any measurable subset A of $\mathcal{A}$, let K_A be its characteristic function (i.e., $K_A(a,r) = 1\; if\; a \in A;\ 0$ otherwise). Then $\mathbf{E}[K_A|r,P]$ defines a conditional probability measure in $\triangle(\mathcal{A})$, interpreted as the posterior probability for A conditional on the message r given the joint distribution P. Once again, observe that the principal's mechanism induces a joint probability distribution over actions and posterior beliefs on the part of the agent. The agent's effort choice depends on his posterior belief about the action the principal has chosen, so from the principal's point of view, a mechanism delivers payoffs according to the joint distribution of actions and beliefs that it generates.

This idea is used in the main theorem below. If the principal uses some

[12] See for example, Burrill (1972, p 393).

randomization that results in the message r being sent to the agent, the message induces a posterior belief in $b \in \triangle(\mathcal{A})$. This mechanism could be replaced by one that sends the message b to the agent, then randomizes over actions according to the distribution b.[13] Whatever effort the agent took in response to the message r will also be a best reply in response to the message b.

Let Γ be a topological space of feasible indirect mechanisms. Let $\gamma = \{\gamma_1 \dots \gamma_n\} \in \Gamma^n$ refer to the entire array of mechanisms offered by the principals. In the sequel I will continue to abuse notation as needed by letting γ_j refer to both the mechanism that the seller offers, and the mapping from messages into actions that is associated with this mechanism.

Agent behavior in each mechanism depends on the agent's valuation and on the mechanisms that he or she observes being offered by the other principals. A *communications strategy* is a measurable mapping $\tilde{c} : \Omega \times \Gamma^n \rightarrow \triangle(C^n)$ that describes the (probability distribution over) messages that the agent will send to the principals as a function of the agent's type and the array of mechanisms that he is offered by the principals. A decision strategy $\tilde{\pi} : \Omega \times \Gamma^n \times C^n \times R^n \rightarrow \triangle(E)$ is a measurable mapping that describes the probability distribution the agent will use to choose his action as a function of his type, the array of mechanisms that he has been offered, the messages he has sent, and the array of messages received from the principals. The pair $(\tilde{c}, \tilde{\pi})$ together constitute a continuation strategy for the agent. Say that the continuation strategy $(\widetilde{c}, \tilde{\pi})$is a *continuation equilibrium* if for every array of mechanisms $\gamma \in \Gamma^n$ offered by the principals and for almost every (with respect to F) valuation $\omega \in \Omega$, (i) the randomization $\tilde{c}(\omega, \gamma)$ maximizes

$$\int \cdots \int \left\{ \int u(a_1, \dots a_n, e, \omega)\, d\tilde{\pi}(\omega, \gamma_1, \dots \gamma_n, c_1, \dots c_n, r_1, \dots r_n) \right\} d\gamma_1(c_1) \dots d\gamma_n(c_2),$$

and; (ii) for every array of messages $(c_1, \dots, c_n)$ the decision strategy $\tilde{\pi}(\omega, \gamma, c',)$ maximizes

$$\mathbf{E}\left[u(a_1, \dots a_n, e, \omega) | r_1, \dots r_n\right]$$

for $\gamma_1(c_1) \dots \gamma_n(c_n)$-almost every array of messages $(r_1, \dots r_n)$.

The first condition requires that the communications strategy that the agent uses maximizes his expected utility for all possible arrays of contract

[13]Recall that the principal commits him or herself to this randomization ex ante.

offers, given the decision strategy $\tilde{\pi}$ that he is using to respond to principals' messages. The second condition requires that the decision strategy that the agent uses maximize his expected utility conditional on the information conveyed by the principals' signals for every array of signals the agent expects to see with positive probability. One important observation is that there is no need to worry about signals that occur with zero probability because of the fact that the principal is able to commit him or herself to a joint randomization over actions and signals. In this sense there are no 'off-equilibrium' moves to worry about at the point where the agent selects his or her action.

When we want to emphasize the underlying set of indirect mechanisms Γ we refer to$(\widetilde{c}, \tilde{\pi})$ as a continuation equilibrium *relative to* Γ. When we wish to emphasize a particular array of mechanisms, we refer to $(\tilde{c}(\cdot, \gamma), \tilde{\pi}(\cdot, \cdot, \gamma))$as a continuation equilibrium relative to γ.

The key to the standard (one principal) revelation principle, is that composing a mechanism with agents' strategies yields a mapping from valuations into actions, or in other words, a 'direct mechanism'. In the context of this paper, there is a corresponding composition. Let

$$m_{\tilde{c},\tilde{\pi}} : \Omega \times \Gamma^n \longrightarrow \triangle(\mathcal{A}^n \times E), \tag{1}$$

be the joint distribution of actions and efforts induced by a particular continuation equilibrium $(\tilde{c}, \tilde{\pi})$.

These definitions make it possible to describe the agent's and principal's payoffs. Suppose that the competing firms choose the randomizations $\delta = \{\delta_1 \ldots \delta_n\}$ with $\delta_k \in \triangle(\Gamma)$ for all $j = 1 \ldots n$, and that agent behavior is described by the continuation equilibrium strategy $(\widetilde{c}, \tilde{\pi})$. The agent's payoff depends on the mechanisms that principals offer. Let $\gamma = \{\gamma_1 \ldots \gamma_n\}$ be any array of mechanisms. Then the payoff that the agent gets is

$$U_\Gamma(\gamma, \tilde{c}, \tilde{\pi}, \omega) =$$

$$\int \cdots \int u(a_1 \ldots a_n, e, \omega)\, dm_{\tilde{c},\tilde{\pi}}(\omega, \gamma) \tag{2}$$

The principal who chooses the randomization δ_j expects the payoff

$$V_\Gamma(\delta_j; \delta_{-j}, \widetilde{c}, \tilde{\pi}) \equiv \tag{3}$$

$$\int \cdots \int \bar{v}(\gamma_j, \gamma_{-j})\;\; d\delta_j(\gamma_j)\, d\delta_{-j}(\gamma_{-j})$$

where

$$\bar{v}(\gamma_j, \gamma_{-j}) \equiv \int\int \cdots \int v(a_1 \ldots a_n, e, \omega)\, dm_{\tilde{c},\tilde{\pi}}(\omega, (\gamma_j, \gamma_{-j}))\, dF(\omega)$$

In an obvious notation, $d\delta_{-j}(\gamma_{-j}) = d\delta_1 \cdots d\delta_{j-1} \cdot d\delta_{j+1} \cdots d\delta_n$.

Say that $(\widetilde{c}, \tilde{\pi}, \delta_1, \ldots \delta_n)$ is an *equilibrium* relative to $(\mathcal{M}, \Gamma)$, or simply that $(\Gamma, \widetilde{c}, \tilde{\pi}, \delta_1 \ldots \delta_n)$ is an equilibrium for $\mathcal{M}$, if: $(\widetilde{c}, \tilde{\pi})$ is a continuation equilibrium and for each $j = 1, \ldots n$ and

$$\delta_j \in \arg\max_{\delta' \in \Delta(\Gamma)} V_\Gamma(\delta'; \delta_{-j}, \widetilde{c}, \tilde{\pi}).$$

It should be apparent from this that the actions that any particular principal takes, and the effort level that is induced depend not only on the agent's valuation but also on the mechanisms that have been offered by the other principals. It might seem reasonable to model common agency in the familiar way by having principals choose allocations in response to the agent's reports about preferences. As actions and efforts will generally depend on mechanisms offered by other firms, it should be clear that the allocations and effort induced by some models of indirect competition will not be reproducible with this kind of direct mechanism. This was illustrated initially by the example in Peck (1995). A complete series of examples illustrating the pitfalls of this approach are given in Martimort and Stole (1997) for the common agency case.

3 Example

The interactions that occur between principals and the agent are complex, so this section provides an example consistent with the formalism of Bernheim and Whinston to illustrate some of the principles involved. The example also illustrates the sense in which restricting the principals to simple 'direct' mechanisms can be restrictive. There is an allocation that can be supported as an equilibrium when principals are allowed to use menus, but which cannot be supported when principals are restricted to use direct mechanisms of the kind used by Bernheim and Whinston. The example is an adaptation of one provided by Martimort and Stole (1997).

Example 1 *There are two principals and a single agent. The feasible actions for the principals are given by* $Y = \{a_1, a_2, a_3\}$. *The agent can take one of two* *possible effort levels given by $E = \{e_1, e_2\}$. The payoff functions are given in the following table:*

	e_1				e_2		
	a_1	a_2	a_3		a_1	a_2	a_3
a_1	$0,0,0$	$2,0,2$	$-1,5,3$	a_1	$0,0,0$	$0,0,0$	$0,0,0$
a_2	$0,2,2$	$1,1,1$	$0,0,0$	a_2	$0,0,0$	$3/2,3/2,3/2$	$0,0,0$
a_3	$5,-1,3$	$0,0,0$	$0,0,1$	a_3	$0,0,0$	$0,0,0$	$0,0,0$

where the first entry in each cell is the payoff to the principal 1 (who chooses the row), the second entry is the payoff to principal 2 (who chooses the column), and the third entry is the agent's payoff (the agent chooses the matrix).

The first matrix is a slightly altered version of the payoffs used in the example developed by Martimort and Stole (1997) to illustrate problems with the revelation principal. Following Bernheim and Whinston (1986) it is assumed that the principal can observe the effort that the agent takes ex post, and write enforceable contracts contingent on this effort level so $\mathcal{A}$ is the set of maps from E into Y. A natural way to model this problem, following Bernheim and Whinston (1986), is simply to assume that each principal chooses a contract from $\mathcal{A}$, and to focus on pure strategy equilibria of the ensuing normal form game. There is a pure strategy equilibrium for this game in which each principal offers a contract which chooses action a_3 no matter what effort the agent takes. In this equilibrium each principal gets 0 and the agent gets 1. It is straightforward to check that this is the only outcome that can be supported when all players use pure strategies. This is restrictive in a crude way since it rules out contracts where the principals respond to some effort levels with random actions. For the moment, suppose that such random contracts are infeasible.

As Martimort and Stole (1997) point out, alternative outcomes may be supportable if principals are allowed to use more complex mechanisms. For example, if principals can offer menus, then the outcome where each player gets the payoff 3/2 is also feasible. The principals achieve this outcome by communicating with the agents before selecting a contract. This works as follows. Each principal offers the agent a choice between two contracts. The first contract responds to every effort level that the agent takes with the action a_2. The second contract responds to every effort that the agent undertakes with the action a_3. Call these contracts γ_2 and γ_3 respectively. The agent then communicates a choice of one of these contracts to each principal. Since random mechanisms are assumed away at this point, there is no need for either of the principals to communicate further with the agents, so the game ends with the agent choosing a level of effort. The principals carry out the actions they have committed themselves to given the choices of the agent and the effort level that the agent undertakes. On the continuation equilibrium path associated with this equilibrium, the agent selects contract γ_2 from each principal and chooses the effort level e_2. The principals respond with action a_2, and all payoffs are $3/2$.

In the Bernheim Whinston framework, this outcome could not be an equilibrium. Either principal could break it up by offering a contract that rewards the agent with action a_1 if he selects effort e_1. If the agent were presented with this kind of contract, he could do better by changing his effort to e_1. However in the menus framework, the agent would also select the contract γ_3 (which responds to every effort level with action a_3) from the non-deviating principal. The net result is that the deviating principal's payoff would then fall from $3/2$ to -1. Thus the menu of contracts makes the deviation unprofitable.

In the formal model described above, principals choose a contract from $\triangle(\mathcal{A} \times \mathcal{R})$ which would allow the principal to respond to some effort levels with a random action. This example nicely illustrates the role of random mechanisms. Focus again on the putative equilibrium in which principals offer the menu $\{\gamma_2, \gamma_3\}$ as above. As mentioned, this supports the payoff $3/2$ for all players when random mechanisms are ruled out. Now, suppose that random mechanisms are feasible and consider the following deviation: the principal offers the agent a lottery that implements contract γ_2 (which responds to every effort with action a_2) with high probability and with complementary probability implements a contract γ' that responds to both effort levels with the action a_1. When the deviator makes this offer, the agent thinks it is very likely that the deviating principal will implement contract γ_2. If this happens the agent wants the other principal to use the action a_2 and so the agent chooses γ_2 from the non-deviating principal just as he did before the deviation. At this point the principal randomizes according to his commitment. Whichever contract the randomizing device selects, the deviating principal will want to inform the agent about the outcome, so that the agent takes the correct action. This is where the subsequent communication from the principal to agent comes into play. The principal simply sends a message to the agent informing him about the contract he has implemented.[14] So with very small probability the principal reverts to action a_1, the agent responds with effort e_1 and the deviating principal's payoff rises.[15]

4 Menus

In the single principal single agent problem, any direct mechanism assigns an outcome to each report the agent makes about his type. No matter how complex the type space is for the agent, the range of this mapping obviously constitutes a menu of alternatives from which the agent can choose. This works in reverse as well, since every menu of alternatives can be written as an incentive compatible direct mechanism that yields the same outcome. It is natural to enquire whether this approach might work in the common agency problem when there is a single agent.

Indirect mechanisms with large message spaces create a fundamentally new opportunity when there are multiple principals. By presenting the same menu with different names the principal can change the effort the agent takes, and change the selection that the agent makes from other principal's menus. In this sense, indirect mechanisms give the principal some control over what other principals do. Forcing the principal to present the agent with a single menu severely restricts the principal's control in all but the simplest problems. So the simple analogy between indirect mechanisms and menus breaks down with multiple principals.

In its simplest form, a menu is simply a mechanism with a message space equal to the set of actions $\mathcal{A}$ along with a mapping $\gamma^* : \mathcal{A} \to \mathcal{A}$ satisfying $\gamma^*(a) \in \{a, \underline{a}\}$ for all ω. In the example given above, the agent is given a choice between two contracts. To model this formally as an indirect mechanism in the manner of Section 2 simply think of the agent naming the contract that he wants. If it is one of the two that the principal wants to allow, the agent gets it. If not, the principal simply assigns one of the two contracts to the message arbitrarily. At first pass, there is no need for the principal to

[14]In a formal sense, the principal commits himself to a randomization in which the action he takes and the message he sends are perfectly correlated.

[15]There is a subtle difference between mixed strategies in the original Bernheim Whinston formulation and random mechanisms. If the principal chooses a contract from $\mathcal{A}$ randomly the agent knows exactly what action will correspond to each effort level he takes. If the principal chooses a random contract from $\triangle(\mathcal{A})$ the agent will only know what probability distribution over actions is associated with each effort level.

communicate further with the agent, so the principal's message space R can be set equal to the empty set.

As illustrated above, mechanisms formed by offering menus of this kind will generally be restrictive since principals might want to offer random mechanisms in which the agent is offered a distribution over the set of outcomes.[16] If the principal does randomize, he may also want to send a signal to the agent that is correlated with the action that he chooses in order to maintain some connection between his action and the effort that the principal makes. In the example given above, the principal simply wants to inform the agent about the action he is taking. The principal does this by committing himself to some joint randomization in $\triangle(\mathcal{A} \times \mathcal{R})$ over actions and messages for each message that he receives from the agent. If the set of messages is finite, then each message r that the agents receives from the principal induces a posterior belief, say $b \in \triangle(\mathcal{A})$ about the action that the principal has taken. Instead of sending the message r, the principal could send the message b, interpreted as a recommendation to the agent about the posterior belief that he should hold. If the agent believes the message, every action that is a best reply to r will also be a best reply to b. Since the principal can commit himself to a joint randomization, he can send recommendations that the agent will believe. He might do this by committing himself to a device that randomly selects a message from $\triangle(\mathcal{A})$ to send to the agent, then chooses output according to the distribution selected.

More generally, if $\mathcal{R}$ is some measurable space of messages and $P \in \triangle(\mathcal{A} \times \mathcal{R})$, then for any integrable function $x : \mathcal{A} \times \mathcal{R} \to \mathbf{R}$, there is a measurable function $\mathbf{E}[x|r, P]$ satisfying $\int_D \mathbf{E}[x|r]\, dP = \int_D x dP$ for all $D \subset \mathcal{R}$, interpreted as the conditional expectation of x given r.[17] For any measurable subset A of $\mathcal{A}$, let K_A be its characteristic function (i.e., $K_A(a, r) = 1$ if $a \in A$; 0 otherwise). Then $\mathbf{E}[K_A|r, P]$ defines a conditional probability measure in $\triangle(\mathcal{A})$, interpreted as the posterior probability for A conditional on the message r given the joint distribution P. Thus every mechanism that assigns a particular randomization in $\triangle(\mathcal{A} \times \mathcal{R})$ induces a randomization

[16] In the example given above the principal and agent have symmetric information, and the principal randomizes only to hide his action from the agent until after the agent has communicated with the other principal. With asymmetric information randomization is often used to make it costly for agents to lie about their types. A typical example is that a buyer who claims to have a low valuation will typically be offered a low price along with a relatively high probability of being rationed.

[17] See for example, Burrill (1972, p393)

in $\triangle(\mathcal{A}\times\triangle(\mathcal{A}))$ that would support the same distribution of effort on the part of the agent. This idea is used in the following theorem to provide a general message space for the principal.

Define $\mathcal{A}^*=\triangle(\mathcal{A}\times\triangle(\mathcal{A}))$. In the formulation that follows the principal will offer the agent the choice from a menu of actions in $\mathcal{A}^*$. For this reason, the set of messages C available to the agent will be equal to $\mathcal{A}^*$. Now define Γ^* to be the set of all mappings $\gamma^*:\mathcal{A}^*\to\mathcal{A}^*$ satisfying

$$\gamma^*(a)=\begin{cases} a & a\in P \\ \underline{a}' & a\notin P \end{cases}$$

for some *closed* subset $P\subset\mathcal{A}^*$, where $\underline{a}'$ is an arbitrary element of P. Each mechanism like this could perfectly well be described by referring to the menu P that defines the set. So Γ^* will henceforth be referred to as the set of menus.

The following theorem shows that arbitrary indirect mechanisms can be replaced with mechanisms in which principals offer the agent a menu of alternatives (from $\mathcal{A}^*$). This new menu is simply the set of alternatives that the agent could have induced in the original direct mechanism by sending appropriate messages to the principal.

Theorem 2 *Suppose that* Γ *is some space of indirect mechanisms; let* $(\widetilde{c},\tilde{\pi},\delta_1,\ldots\delta_n)$ *be an equilibrium relative to* Γ. *Then there is a map* $\psi:\Gamma\to\Gamma^*$, *an array of randomizations* $(\delta_1^*\ldots\delta_n^*)$ *over* Γ^* *and a continuation equilibrium* (c^*,π^*) *relative to* Γ^* *such that*

1.

$$U_{\Gamma^*}(\psi(\gamma),c^*,\pi^*,\omega)=U_\Gamma(\gamma,\tilde{c},\tilde{\pi},\omega)$$

for all ω and for all $\gamma\in\Gamma^n$; and

2. for each $\gamma_j\in\operatorname{supp}\delta_j$

$$V_{\Gamma^*}(\psi(\gamma_j);\delta^*_{-j},c^*,\pi^*)=V_\Gamma(\gamma_j;\delta_{-j},\widetilde{c},\tilde{\pi})$$

otherwise if $\gamma_j\notin\operatorname{supp}\delta_j$

$$V_{\Gamma^*}(\psi(\gamma_j);\delta^*_{-j},c^*,\pi^*)\le V_\Gamma(\delta_j;\delta_{-j},\widetilde{c},\tilde{\pi})$$

The mapping ψ transforms arbitrary indirect mechanisms into menus. The menu $\psi(\gamma_j)$ is simply the list of all the outcomes in $\mathcal{A}^*$ that the agent can induce by sending some message in C. The first condition asserts that there is a continuation equilibrium relative to the set of menus such that for *any* array of mechanisms the principals might offer in the original game, the agent's payoff is unchanged once these are all transformed into menus (using ψ). Condition 2 says that for principal j, if γ_j is on the *equilibrium path* of the original game, then seller j will get the same payoff offering the menu $\psi(\gamma_j)$ in the continuation equilibrium relative to menus, as he does offering γ_j in the original continuation equilibrium. While if γ_j is not on the equilibrium path of the original game, then principal j will get a payoff from the menu $\psi(\gamma_j)$ that is no higher than his equilibrium payoff in the original game. Thus, provided the set Γ is large enough that ψ spans the set of menus, the equilibrium payoffs relative to Γ can be preserved as equilibrium payoffs relative to the set of menus.[18]

It is convenient to contrast the result in theorem 2 with the standard revelation principle. In single mechanism designer problems, every indirect mechanism maps into a payoff equivalent direct mechanism where the agent reports truthfully. When the principal uses any direct mechanism in the image of this map, he attains the same payoff as he gets from some feasible indirect mechanism. It follows immediately that whenever an indirect mechanism lies on the equilibrium path of the original game, then its image in the set of menus will be more profitable than any other direct mechanism. Unfortunately there is no payoff equivalent mapping from indirect mechanisms into menus. Instead Theorem 2 creates a mapping from indirect menus into menus such that the principal's payoff by using the menu is never higher than it was in the original equilibrium. It follows immediately that the payoffs in the original equilibrium can be supported as equilibrium payoffs relative to the image of the original set of indirect mechanisms. If this original set of mechanisms is 'big' enough that its image is equal to the entire set of menus, then its equilibrium payoffs can be supported as equilibrium payoffs relative to the set of menus.

To understand this point, it may help to consider more generally the difference between indirect mechanisms and menus. An indirect mechanism

[18]This is discussed further below. Ad hoc models of competition might be so restrictive that seller cannot offer arbitrary menus of alternatives. For example, if sellers have to offer a contract without communicating with buyers, as in Bernheim and Whinston (1986), then sellers can only offer degenerate menus consisting of a single element.

might admit many more messages than there are actions available to the principal. This makes it possible for the principal to offer the same menu of alternative actions in different ways. For example, suppose c and c' are two distinct messages in C. Mechanism γ chooses contract a if the agent says c and a' otherwise. Mechanism γ' chooses contract a if the agent says c' and choose a' otherwise. These two mechanisms provide the agent with exactly the same menu of alternatives. The reason that the principal might not be indifferent between the two indirect mechanisms is that the agent could respond differently to each of these offers by taking a different level of effort, choosing differently from other principals' menus, or even choosing differently between a and a'. Suppose that γ and γ' both lie off the equilibrium path relative to Γ. Furthermore, suppose that the principal strictly prefers the way the agent responds to γ to the way he responds to γ'. Both these mechanisms are translated into the same menu by the mapping ψ in Theorem 2. In the continuation equilibrium relative to Γ^* suppose that we have the agent respond by acting as if the offer were γ. Then the payoff that the principal gets in the transformed equilibrium when he offers $\psi(\gamma')$ relative to Γ^* will actually be greater than the payoff he gets by offering γ' relative to Γ. Of course, since both γ and γ' lie off the equilibrium path, this payoff will be smaller than the payoff he gets by offering the menu associated with an equilibrium contract.

To further illustrate these ideas, consider the following example.

Example 3 *There are again two principals and the agent has one of two possible* types*, either θ_1 or θ_2. The agent takes no effort in this example, so that the set of feasible contracts for the principal is simply the set of simple actions a_1 and a_2. The principals believe that either type is equally likely. The payoffs are given in the following table:*

	θ_1			θ_2	
	a_1	a_2		a_1	a_2
a_1	$0,0,0$	$-1,5,3$	a_1	$3,3,3$	$0,0,0$
a_2	$5,-1,3$	$0,0,0$	a_2	$0,0,0$	$0,0,0$

The agent takes no effort here, so contracts are just simple actions. Suppose initially that random contracts are infeasible and that the set of feasible indirect mechanisms Γ consists of all the menus *and* all direct mechanisms, where direct mechanisms respond to the agent's declaration about whether type is θ_1 or θ_2 with specific actions. That is, if $\mathcal{D}$ is the set of direct mechanisms, then $\Gamma = \Gamma^* \cup \mathcal{D}$. The equilibrium of interest has each principal offering the agent a *menu* consisting of a choice between actions a_1 and a_2. Call this mechanism γ^* to indicate that it is an *equilibrium* mechanism relative to the initial set of indirect mechanisms Γ. The continuation equilibrium has the property that when both principals offer γ^*, the agent chooses a_1 from both principals when his or her type is θ_2. When his or her type is θ_1, the agent randomizes, choosing a_1 from principal 1 and a_2 from principal 2, or a_1 from principal 2 and a_2 from principal 1 with equal probability. The principals' expected payoffs in this putative equilibrium in which each of them offers γ^* are $2\frac{1}{2}$ $(\Pr\left(\tilde{\theta} = \theta_1\right)\left(\left(\frac{1}{2}\right)5 + \left(\frac{1}{2}\right)(-1)\right) + \Pr\left(\tilde{\theta} = \theta_2\right)3)$.

There are two kinds of deviation in Γ that should be considered in order to verify that this is an equilibrium. One possible deviation is within the set of menus. Note that in the event that the agent has type θ_1, each principal would like to deviate by eliminating the option a_1 from their menu. Eliminating the option reduces the principal's payoff when the agent has type θ_2 to zero, yielding an expected payoff $2\frac{1}{2}$ which is no better than what the principal gets with the original mechanism.[19]

A more interesting deviation occurs when the principal offers a direct mechanism which promises the action a_1 if the agent reports that his type is θ_2 and to promises the action a_2 if the agent reports that his type is θ_1. Call this deviation γ'. Since γ^* is supposed to be an equilibrium, the continuation equilibrium has to be constructed in such a way that the deviation to γ' is unprofitable. Suppose that the agent is expected to respond by reporting type θ_2 to the deviating principal no matter what his true type really is, then choosing the action a_2 from the non-deviating principal when his type is θ_1 and the action a_1 from the non-deviator when his type is θ_2. In this case, the deviating principal's payoff falls to -1 when the agent has type θ_1 and this makes the deviation unprofitable.

How is the deviation γ' translated into the set of menus? Observe that the menu of actions available to the agent under γ' is the set of actions he or she can induce by making different reports about θ. This is just the set $\{a_1, a_2\}$, exactly as it is when the deviating principal offers γ^*. Thus both γ^* and γ' must 'translate' into the same menu in Γ^*, say γ^{**}. The agent's behavior in the continuation equilibrium relative to the original set of indirect

[19]The fact that the principal gets the same payoff before and after a deviation is unimportant - adding ε to the payoffs in the event type is θ_2 eliminates this indifference.

mechanisms is, however, very different under γ^* and γ', and it is only possible to assign a single continuation equilibrium to each menu in Γ^*. By using the continuation equilibrium associated with the original equilibrium offer γ^*, it is ensured that the payoff that the principal gets by deviating from γ^{**} to $\psi(\gamma')$ is no larger than the payoff that the principal gets when he deviates from γ^* to γ'. Since γ^* lies on the equilibrium path of the original game, the deviation relative to the set of menus Γ^* must be unprofitable.

The reader might wonder whether a fully payoff equivalent mapping could be achieved by allowing the principal to send the agent a recommendation about the continuation equilibrium that he wants him to play along with the menu of alternatives he wants to make available. This is the approach that Myerson (1981) uses for the single principal problem. This approach is not pursued here, primarily because it is unnecessary for the purposes of this paper. However, these sorts of messages from principal to agent raise difficult problems in the multi-principal context. For example, the recommendations that the principal makes to the agent will generally have to depend on the mechanisms that the other principals have offered, whether these mechanisms depend on the original principal's mechanism, and so on. So the infinite regress re-emerges. Furthermore, menus depend on the physical characteristics of the environment. The principal's recommendations to the agent about the continuation strategy he should play need to depend on the other principals' mechanisms. To provide a 'revelation principle' that works for any model of indirect competition, the principal would need to make recommendations in a language that works for any model of indirect competition. These are the concerns addressed in Epstein and Peters (1997) - these are also the concerns that this paper is trying to avoid.

The second point to keep in mind is that the theorem does not say that equilibrium payoffs for principals relative to Γ can be supported as equilibrium payoffs relative to Γ^*. The reason is that arbitrary models of indirect competition may restrict principals to mechanisms Γ that incorporate various implicit restrictions on the contracts that principals are allowed. Example 1 illustrates this. Restricting principals to offer contracts that specify what action they will take for each effort level they observe essentially involves a restriction that prevents the agent from communicating with the principal. The impact of this is to force the principal to offer only degenerate (one element) menus. This is expressed formally by observing that $\psi(\Gamma) \subset \Gamma^*$. This observation explains the various failures of the revelation principle that have been uncovered in the literature. For example it is not hard to imagine that if the image of Γ is 'small' enough in Γ^* then equilibrium outcomes relative to Γ^* will generate payoffs that cannot be attained by any combination of mechanisms in Γ. This is the basis of the example given in Peck (1995). Secondly, even if the payoffs in some equilibrium relative to Γ^* are attainable by combinations of mechanisms in Γ, these payoffs may not be supportable as equilibria relative to Γ. This is the basis of the example given in Martimort and Stole (1997).

The third point to note is that the theorem given here is considerable weaker than the one in Epstein and Peters (1997). In that paper conditions are given under which every mechanism in Γ (on or off the equilibrium path) is translated into a payoff equivalent mechanism in the universal set of mechanisms (again the translation here is not payoff equivalent off the equilibrium path).

5 Weak Robustness

The wide use made of direct mechanisms (or menus) in single mechanism designer problems stems from two facts - first, direct mechanisms are simple. Perhaps more important, once principals find the best direct mechanism they can be sure that they cannot improve their profits any more by experimenting with other more complex mechanisms. Neither menus, nor direct mechanisms have this property in competing mechanism problems. Martimort and Stole (1997) provide an example that illustrates this for direct mechanisms. To see why this property fails when principals offer menus, consider Example 3 given above, and suppose we model the competition in that environment by allowing each of the principals to offer menus of actions. The equilibrium allocation described is still supported with this form of competition - recall that each principal offers the agent the option of choosing either action a_1 or a_2. The agent responds by choosing action a_1 when his type is θ_2; otherwise, when his type is θ_1 the agent randomizes choosing a_1 from principal 1 and a_2 from principal 2, or a_2 from principal 2 and a_1 from principal 1 with equal probability. The agent's payoff is always 3, while the principals' payoffs are $2\frac{1}{2}$.

Now suppose we expand the set of mechanisms to allow principals to offer direct mechanisms. As before, suppose that a deviating principal tries to exploit this by offering a direct mechanism that promises a_2 when the agent reports his type as θ_1 and promises action a_1 when the agent reports his type as θ_2. Contrary to the previous example, however, suppose that the agent is expected to report his type truthfully to the principal. This is incentive compatible provided that when the agent has type θ_1 he chooses action a_1 from the non-deviating principal's menu. The deviating principal's payoff rises to 5 when the agent has type θ_1 which makes the deviation profitable.

This example also suggests a way around this problem, since the ability of the principal to increase his profits by offering more complicated mechanisms seems to depend on the way that the continuation equilibrium is assigned. Some assignments increase profits, but as indicated in the previous discussion of this example, there are alternative way to assign the continuation equilibrium so the deviating principal's profits do not increase. It turns out that when there is a single agent, the property is true more generally. The point of this section is to illustrate this idea - if principals are constrained to use menus of the kind described above, then there will always be a way to assign the continuation equilibrium so that this equilibrium persists when the set of mechanisms is enlarged.

To illustrate this idea consider two feasible sets of indirect mechanisms Γ and Γ^1. Say that $\Gamma^1 \succcurlyeq \Gamma$ (Γ^1 is *bigger* than Γ) if there exists an embedding $\alpha' : \Gamma \rightarrow \Gamma^1$. This idea is natural, there are more mechanisms in Γ^1 than in Γ. Now take two models of competition with feasible sets of mechanisms Γ, and Γ^1 and such that $\Gamma^1 \succcurlyeq \Gamma$. Let the associated continuation equilibria be $(\tilde{c}, \tilde{\pi})$ and $(\tilde{c}^1, \tilde{\pi}^1)$ respectively. As before, let $m_{\tilde{c},\tilde{\pi}}(\cdot, \gamma)$ and $m_{\tilde{c}^1,\tilde{\pi}^1}(\cdot, \gamma^1)$ be the joint distributions on actions and effort induced by these continuation equilibria when the array of mechanisms on offer are γ and γ^1 respectively. The continuation equilibrium $(\Gamma^1, \tilde{c}^1, \tilde{\pi}^1)$ is said to *extend* $(\Gamma, \tilde{c}, \tilde{\pi})$ if there is an embedding $\alpha : \Gamma \rightarrow \Gamma^1$ such that for all γ in Γ^n

$$m_{\tilde{c},\tilde{\pi}}(\cdot, \gamma) = m_{\tilde{c}^1,\tilde{\pi}^1}(\cdot, \alpha(\gamma))$$

The generalization of the idea behind direct mechanisms in the single mechanism designer problem is to imagine principals exploring more complex mechanisms (in Γ^1) that are not envisioned by the 'ad hoc' model of competition specified in the economic model of competition (i.e., Γ). A sensible model of competition is one that has the property that the equilibrium allocations that it describes do not disappear when new, more complicated mechanisms are added to the principals' feasible sets.

Definition 4 (Strong Robustness) *An equilibrium* $(\Gamma, \tilde{c}, \tilde{\pi}, \delta)$ *is said to be* strongly robust *if for* any *extension* $(\Gamma^1, \tilde{c}^1, \pi^1)$ *for which* Γ^1 *is compact* *metric,* $(\Gamma^1, \tilde{c}^1, \tilde{\pi}^1, \alpha[\delta])$ *is an equilibrium, where* $\alpha[\delta]$ *is the randomization on* Γ^1 *induced by* δ *and* α*.*

Strongly robust equilibria are analyzed in Epstein and Peters (1997). As shown in the example above, equilibria in menus are not strongly robust. If menus have desirable properties at all they must be weaker than full robustness.

Definition 5 (Weak Robustness) *An equilibrium* $(\Gamma, \tilde{c}, \tilde{\pi}, \delta)$ *is said to be* weakly robust *if for every compact metric* $\Gamma^1 \succcurlyeq \Gamma$ *there exists an extension* $(\Gamma^1, \tilde{c}^1, \tilde{\pi}^1)$ *such that* $(\Gamma^1, \tilde{c}^1, \tilde{\pi}^1, \alpha[\delta])$ *is an equilibrium, where* $\alpha[\delta]$ *is the randomization on* Γ^1 *induced by* δ *and* α*.*

Theorem 6 *In the common agency problem with a single agent, let* $(\Gamma^*, c^*, \pi^*, \delta^*)$ *be an equilibrium. Then* $(\Gamma^*, c^*, \pi^*, \delta^*)$ *is weakly robust.*

To see the idea behind the proof, start with an equilibrium relative to Γ^* and for simplicity consider a pure strategy equilibrium in which each principal offers a single menu of alternatives. Now expand the set of mechanism to Γ and consider the possibility that one principal deviates while all others continue to offer menus. For any arbitrary mechanism γ that this deviating principal offers from the new set of feasible mechanisms Γ, the image $\gamma(C)$ constitutes a new 'menu' of alternative actions that the agent can choose from. The agents choice from this and the other principals menus, as well as his effort choice, determines his or her payoff. Typically there will be one (or more) mechanisms γ in Γ that offer the agent the same menu of alternatives as $\gamma^* \in \Gamma^*$. Ordinarily these will yield the principal a different payoff than the payoff he gets from γ^* because the agent will choose different actions from each menu and different effort. This is precisely why the principal might like the larger set of mechanisms. Yet since the overall menu of actions and efforts available to the agent is the same when he is offered γ as it is when he is offered γ^* (since the other principals continue to offer the same menus as before), the maximum payoff that he can attain from both mechanisms must be the same. This means that a new continuation equilibrium can be constructed by having the agent choose from the 'menu' associated with γ in exactly the same way as he does when he chooses from the menu γ^*. Then, in this new continuation equilibrium, the principal must get the same payoff when he offers γ as he does when he offers γ^*. With this construction, every payoff that the principal can attain in Γ coincides with a payoff that he could have attained in Γ^*. It follows immediately that whatever was optimal for the principal relative to Γ^* will then be optimal relative to Γ.

6 Applications

Despite the fact that menus in $\mathcal{A}^*$ are relatively tractable, they are considerably more complicated than the menus that have been used in the literature. They involve menus of randomizations and messages sent from the principal to the agent. The reason for the apparent complexity of menus here is the fact that the revelation principle presented has been designed to encompass a broad variety of different applications. In more specific environments the set of menus could be simplified considerably. There are two possible approaches. One approach is simply to argue that certain contracts are infeasible. For example, Bernheim and Whinston (1987) and Dixit, Grossman, and Helpman (1997) deal with environments where agents have no private information. They rule out the possibility that the principal could respond to the agent's effort with a random choice of actions. Since the agent already knows what action the principal will take in response to any effort, there is no need to allow the principal to send the agent additional messages. A menu is then simply a set of mappings from the agent's effort space into the set of actions feasible for the principal. A reasonable justification for this approach is that the agent can observe the *outcome* of a randomizing device, but may have a much more difficult time observing and verifying the random device itself. Specific properties of the environment could then be exploited to try to determine whether attention could be restricted to an even more reduced subset of the set of menus - for example simple non-linear pricing schemes.

For example, Martimort and Stole (1999a) consider an environment similar to the one in Klemperer and Meyer (1989) in which a single agent produces output for two different principals. One way to interpret their problem is to model the simple actions of the principals as the per unit prices that they offer for output, while the agent's effort is the quantity of output that is produced for each principal. They rule out random contracts and mixed strategy equilibrium, so that the set $\mathcal{A}_j$ of feasible contracts for principal j in this interpretation is the set of mappings from the output that the agent produces for principle j into the set of possible per unit prices, i.e., the set of non-linear pricing schemes. Their 'taxation principle' shows that the payoffs in every (pure strategy) equilibrium can be reproduced by having each principal offer a single non-linear pricing scheme. In this sense, they provide an environment in which equilibrium allocations with more general mechanisms can be supported without menus (of non-linear pricing schemes).

It is tempting to argue that because certain kinds of contracts are not observed in practise, they can be safely ignored. One advantage of the approach taken here is that it makes it possible to ask whether certain properties of the environment might explain *why* complex contracts are not observed. To illustrate consider the simplest model of price competition. To keep the argument simple, suppose there are two sellers selling to a single buyer. Each seller can contract on his own price, and the quantity that he sells to the buyer (though the buyer cannot be compelled to make the trade). So the set of simple actions available to each seller is simply the set of all feasible price quantity combinations. The buyer has no private information, but always has the option of refusing to deal with either or both of the sellers. To model this, let $E = \{0, 1, 2\}$ where $e = 0$ indicates that the buyer chooses not to trade with either seller, $e = 1$ means he trades with seller 1, and so on. It seems reasonable to assume that the buyer's utility function is such that he doesn't care about the actions of a particular seller unless he trades with him. The set of feasible contracts is then simply the set of price quantity combinations the seller can offer a buyer who chooses to trade with him. The buyer's payoff is some function $u(p_1, q_1, p_2, q_2, e)$ of the prices and quantities offered by the two sellers, while sellers' payoff is assumed to be revenue from the transaction less some non-negative cost $c(q)$.

By theorem 2 and 6 we should model the competition among the sellers by having them simultaneously offer the buyer menus from $\Gamma^* = \triangle(\mathcal{A} \times \triangle(\mathcal{A}))$. To help think about elements in Γ^* imagine a random device as a map from some probability space (Ω, F) into $\mathcal{A} \times \triangle(\mathcal{A})$. For each possible value of the random variable $\omega \in \Omega$ the seller chooses a contract in $\mathcal{A}$ and sends a message to the buyer consisting of a probability distribution over $\mathcal{A}$. By Theorem 2 attention can be restricted to contracts where the seller sends the true probability distribution over actions conditional on the message (in other words, we can safely assume that the buyer believes the seller's messages). Since the buyer and seller only care about what happens when the buyer trades with the seller in the simple example, each value ω for the random variable generates a price quantity choice and a distribution over the set of feasible price quantity choices. The buyer learns the distribution over prices and quantities that the seller has chosen then decides whether to trade with the seller. In effect, the formalism takes the simple model in which each seller offers the buyer a price quantity pair and generalizes to allow each seller to offer a price quantity lottery. Full generality is achieved by allowing each seller to offer the buyer a *menu* of price quantity lotteries.

It is easy to see that if the buyer is risk averse (utility is quasi-concave) then for any lottery that the seller offers, there will be some price quantity contract that the gives the buyer the same expected utility no matter what his trading choice. With appropriate restrictions on the utility functions, lotteries won't be offered in equilibrium - so there is no harm in ruling them out and assuming that the buyer knows the seller's price quantity offer exactly at the time he makes his participation choice.

Further restrictions on utility can be used to specialize the menus. For example suppose that buyer utility is given by

$$u\left(p_1, q_1, p_2, q_2, e\right) =$$

$$e_1 \min\left\{\left(p_1 - v\right) q_1, \left(p_1 - v\right) \overline{q}\right\} + \left(1 - e_1\right) \min\left\{\left(p_2 - v\right) q_1, \left(p_2 - v\right) \overline{q}\right\}$$

where $e_1 \in \{0, 1\}$ is 1 when the buyer chooses to buy from seller 1 and $\overline{q}$ is the buyers inelastic demand. This is the Bertrand model of price competition. If one of the sellers offers a menu where two different quantities are offered at the same price, the buyer will always pick the larger quantity (provided it is less than or equal to $\overline{q}$). It follows that in any equilibrium in which sellers offer menus to the buyers, the same allocation could be achieved by replacing the sellers' menus with simple non-linear pricing schemes.

These considerations are specific to the application being considered, and go well beyond the scope of this paper.

7 Conclusion

In the case where many principals try to control the incentives of a single agent, the competition among them can be effectively modelled by having them offer alternative menus to agents. This model will be 'effective' in two ways. First, equilibria from arbitrary competitive models can be represented as equilibria relative to (a subset of) the set of menus. Secondly, equilibria relative to the set of all menus are weakly robust. So, equilibria in menus need not fall apart when the set of mechanisms is extended.

Despite the fact that common agency is now a common way of modelling incentive problems, an obvious shortcoming of the analysis in this paper stems from the fact that there is only a single agent. Unfortunately, the methods presented here do not suggest any extension to the multiple agent case. Even if there is only a single principal, menus offer little advantage with multiple agents, because the choice from the menu is a 'joint' choice (of all the agents) not an individual choice. A key part of the argument in Theorem 2 relies on the fact that when a single agent faces the same choice set offered in two different ways, his optimal choice from it will always give him the same payoff - even when this choice varies with the way the choice set if offered. When multiple agents are offered the same normal form game in different ways, it is quite possible that the equilibrium that they choose (and consequently every agents payoff) will vary. Once the principal is required to send instructions to the agents about how to play a game, an infinite regress emerges - the principal's recommendation about how to play should depend on whether the other principals' recommendations depend on whether the principal's recommendation depends... and so on. It is possible to deal with this (see Epstein and Peters (1997)) but the appropriate set of mechanisms is complex.

8 Appendix - Proofs

We begin with some basic definitions and relationships. For any subset $P \subset \mathcal{A}^*$ define $\beta(P) : \mathcal{A}^* \to \mathcal{A}^*$ such that

$$\beta(P)(a) = \begin{cases} a & \text{if } a \in P \\ \underline{a}' & \text{otherwise} \end{cases} \tag{4}$$

In words, $\beta(P)$ is a function that responds to any joint distribution (action) in P with the same joint distribution (over actions and messages), otherwise it responds with $\underline{a}'$ where $\underline{a}'$ is some arbitrary action in P. With this understanding, we will refer to $\beta(P)$ as a menu of alternatives from $\mathcal{A}^*$.

Each mechanism γ_j and report c produces a joint distribution over simple actions and messages for the principal. Each message that the agent receives induces some posterior belief in $\triangle(\mathcal{A})$ about the simple action that the principal has undertaken. Thus each report that the agent makes to the principal induces a pure *action* in $\mathcal{A}^* \equiv \triangle(\mathcal{A} \times \triangle(\mathcal{A}))$. With a slight abuse of notation, define the set of actions (in $\mathcal{A}^*$) that the agent can induce by sending some message to the principal as $\gamma_j(C) \subset \mathcal{A}^*$. Recall that this set is closed in $\mathcal{A}^*$ by assumption.

The agent chooses his effort after observing the message associated with the action $a^* \in \mathcal{A}^*$. Together, the principals' strategies δ, the continuation strategies $(\tilde{c}, \tilde{\pi})$and the agent's type ω induce a joint distribution over $(\mathcal{A}^*)^n \times E$.

Definition 7 *For any $D \subset \Gamma^n$ such that the $\delta(D) > 0$ define $b(\omega, D)$ to be the joint distribution on $(\mathcal{A}^*)^n \times E$ induced by the equilibrium $(\Gamma, \tilde{c}, \tilde{\pi}, \delta)$* conditional *on the principals' mechanisms being drawn from the subset $D \subset \Gamma^n$.*[20] *Write $b^{\mathcal{A}^*}(\omega, D)$ as the marginal distribution of b on the space of actions, $(\mathcal{A}^*)^n$.*

Write (a, q) as a typical element of $\mathcal{A} \times \triangle(\mathcal{A})$. Since the agent's posterior beliefs are defined before he is forced to choose an effort, the utility associated with any particular array of actions $a^* \in (\mathcal{A}^*)^n$ is given by

$$y(a^*, \omega) \equiv$$

$$\int \cdots \int \left\{ \max_e \sum_{a_1} \cdots \sum_{a_n} u(a_1 \ldots a_n, e, \omega) q_1(a_1) \ldots q_n(a_n) \right\} da_1^*(a_1, q_1) \ldots da_n^*(a_n, q_n)$$

With this notation, the agent's payoff in any continuation equilibrium can be written as

$$U_\Gamma(\gamma, \tilde{c}, \tilde{\pi}, \omega) =$$

$$\int \cdots \int u(a_1 \ldots a_n, e, \omega) \, dm_{\tilde{c},\tilde{\pi}}(\omega, \gamma) =$$

$$\int \cdots \int \{y(a_1^*, \ldots a_n^*, \omega)\} \, db^{\mathcal{A}^*}(\omega, \gamma) =$$

$$\max\left\{y(a_1^*, \ldots a_n^*, \omega) : a_j^* \in \gamma_j(C) \, \forall j = 1 \ldots n\right\} \quad (5)$$

The first equality follows from the fact that every effort level that the agent takes with positive probability must maximize his interim payoff. The final equality follows from the fact that b is the distribution associated with a continuation equilibrium following γ. The use of max instead of sup in the final operation results from the assumption that a continuation equilibrium always exists.

[20]The fact that this distribution depends on $\tilde{c}, \tilde{\pi}$ and δ is suppressed to ease notation.

8.1 Proof of Theorem 2

Proof. To begin, the map ψ is described, and the strategies that principals and the agent use in the continuation equilibrium relative to the new space of mechanisms are described. Define the transformation ψ_j from Γ into Γ^* as

$$\psi_j : \gamma_j \longmapsto \gamma_j^* (\cdot) = \beta (\gamma_j (C)) (\cdot)$$

This transformation simply converts a mechanism γ_j into the menu of alternatives in $\mathcal{A}^*$ that it provides. Let δ_j^* be the measure induced by δ_j on Γ^* by this map.

ψ_j will generally be a many to one mapping with inverse correspondence $\psi_j^{-1} (\cdot)$. Choose an arbitrary selection $\bar{\psi}_j^{-1} (\cdot)$ from this correspondence. For any collection $\{\gamma_1^* \dots \gamma_n^*\} \subset \Gamma^n$ define the sets

$$B_j \left(\gamma_j^*\right) = \begin{cases} \psi_j^{-1} \left(\gamma_j^*\right) \cap \operatorname{supp} \delta_j & \text{if } \psi_j^{-1} \left(\gamma_j^*\right) \cap \operatorname{supp} \delta_j \neq \emptyset \\ \bar{\psi}_j^{-1} \left(\gamma_j^*\right) & \text{otherwise} \end{cases} \tag{6}$$

We can now specify the continuation equilibrium relative to Γ^*. Begin with the case where $\gamma^* \notin \psi (\Gamma)^n$. By the definition of Γ^*, The agent must choose a distribution over actions from each principal that lies in a closed subset of $\mathcal{A}^*$, as well as an effort from a finite set. Since the agent's payoffs are linear (by expected utility) with respect to the distribution of actions that principals choose, they are continuous. It follows that a best choice-action pair from each menu exists for the agent. So it will always be possible to assign some continuation equilibrium to the array of mechanisms. Since arrays like this will play no role in what follows, the continuation equilibrium can be assigned in any convenient fashion.

On the other hand if $\gamma^* = \{\gamma_1^*, \dots \gamma_n^*\} \in \psi (\Gamma)^n$ define

$$c^* (\omega, \gamma^*) = b^{\mathcal{A}^*} (\omega, \{B_1 (\gamma_1^*), \dots B_n (\gamma_n^*)\}) \tag{7}$$

where $b^{\mathcal{A}^*} (\omega, \{B_1 (\gamma_1^*), \dots B_n (\gamma_n^*)\})$ denotes the marginal distribution on $(\mathcal{A}^*)^n$ induced by $b (\omega, \{B_1 (\gamma_1^*), \dots B_n (\gamma_n^*)\})$.

To define the agent's effort choice, let $a^* = \{a_1^*, \dots a_n^*\}$ denote the agent's actual choice from the menu offered by each principal, and let $q \in \triangle (\mathcal{A})^n$ be the array of messages that the agent receives. Note that the agent's posterior beliefs (after seeing the principals' messages) also lie in $\triangle (\mathcal{A})^n$. The agent chooses his effort level according to

$$\pi^* (\omega, \gamma^*, q) = b^{\mathcal{E}} (\omega, \{B_1 (\gamma_1^*), \dots B_n (\gamma_n^*)\} | a^*, q) \tag{8}$$

where $b^{\mathcal{E}}\left(\omega,\left\{B_1\left(\gamma_1^*\right),\dots B_n\left(\gamma_n^*\right)\right\}|a^*,q\right)$ is the distribution of effort induced by $b\left(\omega,\left\{B_1\left(\gamma_1^*\right),\dots B_n\left(\gamma_n^*\right)\right\}\right)$ conditional on the principals' actions being given by a^* and the agent's posterior beliefs being given by q.

The next step is to show that this strategy constitutes a continuation equilibrium, and that the agent's payoff when he follows this continuation strategy in the transformed problem is the same as his payoff in the original problem. To begin, consider the payoff that the agent receives in the continuation equilibrium relative to (γ^*, c^*, π^*) after choosing the array a^* from the various principals menus, and receiving the array of messages q from the principals. From (8), this is given by

$$\int\left\{\sum_{a_1}\cdots\sum_{a_n}u\left(a_1\dots a_n,e,\omega\right)q_1\left(a_1\right)\dots q_n\left(a_n\right)\right\}db^{\mathcal{E}}\left(\omega,\left\{B_1\left(\gamma_1^*\right),\dots B_n\left(\gamma_n^*\right)\right\}|a^*,q\right) \tag{9}$$

Consider any action a_j^* in the menu offered by principal j. By construction, each such action must coincide with a joint distribution of simple actions and posterior beliefs that the agent can induce (by sending appropriate messages) to the principal in the continuation equilibrium relative to $(\gamma_j, \gamma_{-j}, \tilde{c}, \tilde{\pi})$ for any $\gamma_j \in B_j\left(\gamma_j^*\right)$. It follows that any message q_j that the agent could receive when he selects a_j^* from principal $j's$ menu must be equal to the correct posterior for a conditional on that message. Then from the properties of the continuation equilibrium $(\tilde{c}, \tilde{\pi})$ every effort level in the support of $b^{\mathcal{E}}\left(\omega,\left\{B_1\left(\gamma_1^*\right),\dots B_n\left(\gamma_n^*\right)\right\}|a^*,q\right)$ must maximize the agent's expected payoff conditional on the posterior belief q. Thus (9) can be simplified to

$$\max_e\sum_{a_1}\cdots\sum_{a_n}u\left(a_1\dots a_n,e,\omega\right)q_1\left(a_1\right)\dots q_n\left(a_n\right)$$

Then the ex ante payoff associated with the choice a^* is just

$$\int\cdots\int\left\{\max_e\sum_{a_1}\cdots\sum_{a_n}u\left(a_1\dots a_n,e,\omega\right)q_1\left(a_1\right)\dots q_n\left(a_n\right)\right\}da_1^*\left(a_1,q_1\right)\dots da_n^*\left(a_n,q_n\right)$$

$$=y\left(a^*,\omega\right)$$

The agent selects according to the distribution $b^{\mathcal{A}^*}\left(\omega,\left\{B_1\left(\gamma_1^*\right),\dots B_n\left(\gamma_n^*\right)\right\}\right)$. So this strategy yields the payoff

$$\int y\left(a^*,\omega\right)db^{\mathcal{A}^*}\left(\omega,\left\{B_1\left(\gamma_1^*\right),\dots B_n\left(\gamma_n^*\right)\right\}\right)$$

$$= \mathbf{E}_{B(\gamma^*)} \int y\left(a^*, \omega\right) db^{\mathcal{A}^*}\left(\omega, \{\gamma_1, \ldots \gamma_n\}\right) \quad \text{(by definition)}$$

$$= \mathbf{E}_{B(\gamma^*)} \max\left\{y\left(a_1^*, \ldots a_n^*, \omega\right) : a_j^* \in \cup b^j\left(\cdot, \gamma_j, \cdot\right) \forall j = 1 \ldots n\right\} \quad \text{(from (5))}$$

$$= \max\left\{y\left(a_1^*, \ldots a_n^*, \omega\right) : a_j^* \in \cup b^j\left(\cdot, \gamma_j, \cdot\right) \forall j = 1 \ldots n\right\}$$

(because of the fact that $\cup b^j\left(\cdot, \gamma_j, \cdot\right)$ is the same for all $\gamma_j \in B_j\left(\gamma_j^*\right)$). Since the choice sets in this last expression are the ones offered in the menus associated with the mechanisms in γ^*, this line verifies that the continuation strategy specified for the agents is a continuation equilibrium. Finally applying (5) once again gives this payoff equal to

$$= U_\Gamma\left(\gamma, \tilde{c}, \tilde{\pi}, \omega\right)$$

which establishes the first property in the theorem.

Finally, it remains to show that

$$V_{\Gamma^*}(\psi\left(\gamma_j\right); \delta_{-j}^*, c^*, \pi^*) = V_\Gamma(\gamma_j; \delta_{-j}, \widetilde{c}, \tilde{\pi})$$

for all $\gamma_j \in \operatorname{supp} \delta_j$.

For any $\left(\gamma_j^*, \gamma_{-j}^*\right)$

$$\bar{v}\left(\gamma_j^*, \gamma_{-j}^*\right) =$$

$$\int\int v\left(a_1 \ldots a_n, e\right) db\left(\omega, \{B_1\left(\gamma_1^*\right), \ldots B_n\left(\gamma_n^*\right)\}\right) dF\left(\omega\right)$$

$$= \mathbf{E}_{\gamma \in B(\gamma^*)} \int\int v\left(a_1 \ldots a_n, e\right) dm_{\tilde{c}, \tilde{\pi}}\left(\omega, \left(\gamma_j, \gamma_{-j}\right)\right) dF\left(\omega\right)$$

Integrating this using the distribution δ_{-j}^* gives

$$V_{\Gamma^*}(\gamma_j^*; \delta_{-j}^*, c^*, \pi^*)$$

$$\mathbf{E}_{\gamma_j' \in B_j\left(\gamma_j^*\right)} \int \left\{\int\int v\left(a_1 \ldots a_n, e\right) dm_{\tilde{c}, \tilde{\pi}}\left(\omega, \left(\gamma_j', \gamma_{-j}\right)\right) dF\left(\omega\right)\right\} d\delta_{-j}$$

$$= \mathbf{E}_{\gamma_j' \in B_j\left(\gamma_j^*\right)} V_\Gamma(\gamma_j'; \delta_{-j}, \widetilde{c}, \tilde{\pi})$$

If $\gamma_j^* = \psi\left(\gamma_j\right)$ for some $\gamma_j \in \operatorname{supp} \delta_j$, then by the definition, each $\gamma_j' \in B_j\left(\gamma_j^*\right)$ is an element of the support of δ_j. Since δ_j is an equilibrium distribution for principal j, every γ_j' in its support gives the same payoff. Thus the integrand in this last expression is constant and equal to

$$V_\Gamma(\gamma_j; \delta_{-j}, \widetilde{c}, \tilde{\pi})$$

as required.

Alternatively if $\psi^{-1}\left(\gamma_j^*\right) = \bar{\gamma} \notin \operatorname{supp} \delta_j$ then

$$\mathbf{E}_{\gamma_j' \in B_j\left(\gamma_j^*\right)} V_\Gamma(\gamma_j'; \delta_{-j}, \widetilde{c}, \tilde{\pi}) = V_\Gamma(\bar{\gamma}; \delta_{-j}, \widetilde{c}, \tilde{\pi}) \leq V_\Gamma(\delta_j; \delta_{-j}, \widetilde{c}, \tilde{\pi})$$

which proves the theorem. ■

8.2 Proof of theorem 6

Proof. Let $(\Gamma, \tilde{c}, \tilde{\pi})$ be an extension of (Γ^*, c^*, π^*) with embedding given by α. The method of the proof is to convert deviations that lie outside of the range of α into the menus that they do provide, and then change the continuation equilibrium associated with those menus to coincide with the original equilibrium.

Use the transformation ψ from Theorem 2 to associate with each mechanism $\gamma_j \in \Gamma$ the corresponding menu $\gamma_j^*(\cdot) = \beta\left(\gamma_j(C)\right)(\cdot)$. Now let $\mathcal{D}\left(\gamma_1, \ldots \gamma_n\right) \subset \triangle\left(\mathcal{A}^n \times E\right)$ be the set of joint distributions over simple actions and efforts that can be induced with *some* distribution of messages and some effort choice strategy, given the mechanisms $\{\gamma_1, \ldots \gamma_n\}$. If the agent if offered an array of mechanisms $\{\gamma_1, \ldots \gamma_n\}$ in the continuation equilibrium relative to Γ, he can induce any array of actions $\{a_1, \ldots a_n\}$ in the set $\gamma_1(C) \times \cdots \times \gamma_n(C) \equiv \gamma_1^*(\cdot) \times \ldots \gamma_n^*(\cdot) \subset \left(\mathcal{A}^*\right)^n$ by sending appropriate messages to the principals. It follows that $\mathcal{D}\left(\gamma_1, \ldots \gamma_n\right) = \mathcal{D}\left(\gamma_1^*, \ldots \gamma_n^*\right)$. Since (c^*, π^*) is a continuation equilibrium relative to Γ^*, $m_{c^*,\pi^*}\left(\omega, \gamma_1^*, \ldots \gamma_n^*\right)$ maximizes the expected utility over $\mathcal{D}\left(\gamma_1^*, \ldots \gamma_n^*\right)$ of a agent of type ω. It follows immediately that this distribution maximizes the agent ω's expected utility over the set $\mathcal{D}\left(\gamma_1, \ldots \gamma_n\right)$. Choose any continuation strategy $\tilde{c}^*$ and $\tilde{\pi}^*$ relative to Γ that satisfies

$$m_{\tilde{c}^*,\tilde{\pi}^*}\left(\omega, \gamma_1, \ldots \gamma_n\right) = m_{c^*,\pi^*}\left(\omega, \psi\left(\gamma_1\right), \ldots \psi\left(\gamma_n\right)\right)$$

(this strategy exists by the definition of $\mathcal{D}$). Since $\mathcal{D}(\gamma_1, \ldots \gamma_n) = \mathcal{D}(\psi(\gamma_1), \ldots \psi(\gamma_n))$ this continuation strategy must be a continuation equilibrium.

Now we begin with an equilibrium $(\delta_1^*, \ldots \delta_n^*)$relative to Γ^* with $\alpha[\delta^*]$ as the induced distribution relative to Γ. The payoff to some principal who unilaterally deviates to a mechanism γ' outside of $\alpha[\Gamma^*]$ while all other principals offer mechanisms γ_{-j} in the support of δ^* is given by

$$\int\int \cdots \int v(a_1 \ldots a_n, e)\, dm_{\tilde{c}^*, \tilde{\pi}^*}(\omega, (\gamma', \gamma_{-j}))\, dF(\omega)$$

$$= \int\int \cdots \int v(a_1 \ldots a_n, e)\, dm_{c^*, \pi^*}(\omega, (\psi(\gamma'), \psi(\gamma_{-j})))\, dF(\omega) \leq$$

$$V_\Gamma(\delta_j, \delta_{-j}, c^*, \pi^*) =$$

$$V_\Gamma(\alpha[\delta_j], \alpha[\delta_{-j}], \tilde{c}^*, \tilde{\pi}^*)$$

which implies that the equilibrium $(\Gamma^*, \delta, c^*, \pi^*)$ is weakly robust. ■

References

Bernheim, B. D., and M. D. Whinston (1986): "Common Agency," *Econometrica*, 54(4), 923–942.

Bernheim, D., and M. Whinston (1987): "Menu Auctions, Resource Allocations and Economics Influence," *Quarterly Journal of Economics*, 101, 1–31.

Biais, B., D. Martimort, and J.-C. Rochet (1997): "Competing Mechanisms in a Common Value Environment," Discussion paper, Universite de Toulouse.

Burrill, C. (1972): *Measure Integration and Probability.* McGraw Hill.

Dixit, A., G. Grossman, and E. Helpman (1997): "Common agency and coordination: General theory and application to government policy-making," pp. 752–769.

Epstein, L., and M. Peters (1997): “A Revelation Principle for Competing Mechanisms,” University of Toronto.

Grossman, G. M., and E. Helpman (1994): “Protection for Sale,” *American Economic Review*, 84(4), 833–850.

Guesnerie, R. (1995): *A Contribution to the Pure Theory of Taxation.* Cambridge University Press.

Katz, M. (1987): “Game Playing Agents: Unobservable Contracts as Precommitments,” mimeo, University of California at Berkeley.

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

Martimort, D., and L. Stole (1997): “Communications Spaces, Equilibria Sets and the Revelation Principle under Common Agency,” University of Chicago unublished manuscript.

——— (1999a): “Contractual Externalities and Common Agency Equilibria,” University of Chicago unpublished manuscript.

Martimort, D., and L. Stole (1999b): “The Revelation and Taxation Principles in Common Agency Games,” mimeo University of Chicago.

McAfee, P. (1993): “Mechanism Design by Competing Sellers,” *Econometrica*, 61(6), 1281–1312.

Myerson, R. (1979): “Incentive Compatibility and the Bargaining Problem,” 47, 61–73.

——— (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 55–73.

Parlour, C. A., and U. Rajan (1997): “Competition in Loan Contracts,” Discussion paper, Carnegie Mellon University.

Peck, J. (1994): “Competition in Transactions Mechanisms: The Emergence of Competition,” Ohio State University mimeo.

——— (1995): “Competing Mechanisms and the Revelation Principle,” *Ohio State University Manuscript*, p. 12.

Rochet, J. (1986): “Le Controle Des Equations Aux Derivees Partielles Issues de la Theories Des Incitations,” Ph.D. thesis, Universite Paris IX.