

Common Agency and the Revelation Principle

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

January 19, 1999

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.

The set of direct mechanisms form a universal class of mechanisms for problems in which a single principal tries to control the incentives of one or more agents. These mechanisms are universal in the sense that the payoffs (or allocations) generated by any indirect mechanism can be reproduced by

^aThanks are due to Barry 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.

¹Early references include Gibbard [5], Green and Laffont [6], and Myerson [10].

an appropriate direct mechanism. An immediate implication is that the most profitable direct mechanism is always at least as profitable as any indirect mechanism. In this sense, optimal direct mechanisms are robust in the sense that sellers will continue to use them even if the set of feasible mechanisms is enlarged.

It has been known for some time that 'direct' mechanisms of the usual sort do not have these properties when there are multiple principals competing against one another.² The difficulty is that agents have market information at the time that they communicate with principals and 'naive' direct mechanisms do not allow sellers to use this information. Alternatively, a complete description of the agent's preferences (and beliefs) is not enough to reveal the agent's 'type' in the sense of Harsanyi. Examples in the literature show that equilibrium indirect mechanisms are not robust against the possibility that the mechanism designers might deviate outside this class of mechanisms³ and that there are classes of indirect mechanisms whose equilibrium allocations cannot be replicated when sellers are compelled to use only direct mechanisms⁴.

At a conceptual level, this problem has been resolved by Epstein and Peters [4] who construct a universal set of mechanisms based, in part, on a refined notion of an agent's type. This set of mechanisms has both the desirable properties that the set of direct mechanisms possess in the usual single principle problem. First, any indirect model of competition among mechanism designers gives rise to a set of feasible mechanisms that can be embedded into the universal set of mechanisms. Secondly, equilibria relative to the universal set of mechanisms are robust in the sense that equilibria will continue as equilibria when the set of feasible mechanisms is extended.

The set of types involved in such a construction are complex. First of all, since type contains market information it is endogenous. Second, even in single agent environments where the principals' beliefs are common knowledge,

² The basic problem with direct mechanisms in competing mechanism design problems was first observed by Myerson [9] and Katz [7]. Examples illustrating the difficulties that arise are given in [8], [14] and [4].

³ An example is given in Martimort and Stole [8].

⁴ Martimort and Stole [8] give an equilibrium in which sellers use reasonable looking indirect mechanisms. Peck [14] gives an example of a non-truthful equilibrium in direct mechanisms whose allocations cannot be replicated in any equilibrium in which sellers use direct mechanisms.

the principals' mechanisms can depend on one another in complex ways. For example, seller 1's mechanism may depend on whether seller 2's mechanism depends on whether seller 1's mechanism depends... and so on. So the set of agent's types must come from an infinite dimensional space even when the set of agent's types is finite in the conventional sense. So despite the fact that a universal set of mechanisms is known to exist, it is difficult to use this result in practice.

The purpose of this paper is to ask whether the construction of a universal set of mechanisms is simpler in a more specialized environment. The natural candidate for such an investigation is that class of problems where multiple principals attempt to control the incentive of a single agent. This model has been the focus of the emerging literature on common agency.⁵ One of the methods that has been used to analyze problems in this framework is to imagine that the principal offers the agent menus of alternatives from which he can choose. In the single mechanism designer problem with a single agent who has private information but takes no unobservable action, menus of alternatives are clearly equivalent to direct mechanisms.⁶ This is true no matter how complicated the type space is, any contract that asks the agent to report a type is offering the agent a menu of alternatives of some kind. The purpose of this paper is simply to evaluate menus as an appropriate model of competition among mechanism designers when there is only a single agent.

As will become apparent, menus are not universal in the sense of Stein and Peters[4].⁷ The universal set of mechanisms has the property that any arbitrary class of direct mechanisms can be embedded into it. In this sense, the universal set extends any arbitrary set of direct mechanisms. It creates new deviations that can destroy equilibrium relative to some arbitrarily restricted set of direct mechanisms. It also creates mechanisms that can be used to support equilibrium outcomes that cannot be supported relative to arbitrarily restricted sets of direct mechanisms. The embeddings (in this universal class) associated with menus and alternative sets of direct mech-

⁵ For example, Bernheim and Wilson [1], Hartman and Stole [8] or Biais, Hartman and Rached [2].

⁶ It has been suggested to me that this is referred to as the 'taxation principle' although I am unsure of the source of this term.

⁷ For the same reason, menus are not universal in the way that direct mechanisms are universal in the single mechanism designer problem. Arbitrary indirect mechanisms cannot be uniquely associated with some specific menu.

arisms will generally intersect, without any associated proper containment.

The difficulty is that the same menu offered in different ways (that is to say, offered as alternative indirect mechanisms) can induce different continuation equilibria. To account for this the universal set of mechanisms must be at least as large the set of all menus plus the set of all continuation equilibria associated with those menus. Menus alone cannot be used to describe more general indirect mechanisms.⁸

However, menus possess properties weaker than universality in common agency problems that are almost as good. The first theorem presented below shows that any equilibrium relative to some set of indirect mechanisms can be translated into an equilibrium relative to some subset of the set of menus.⁹ This translation does not preserve payoffs in the way that translations into the universal set of mechanisms does. However the translations of indirect mechanisms played along the equilibrium path do preserve payoffs so that the equilibrium payoffs relative to some set of feasible indirect mechanisms can be 'represented' as equilibrium payoffs relative to some restricted set of menus. This means that if there are models of indirect competition that have desirable properties like robustness [4], then it must be possible to characterize the associated equilibrium payoffs as equilibrium relative to the set of menus.

This result does not imply that every equilibrium in menus is necessarily interesting. The set of menus is a strict subset of the universal set of mechanisms so it is likely that some equilibrium relative to the set of menus will not be robust to the possibility that principals deviate to more complicated mechanisms. This property is the key property of the universal set of mechanisms discussed in Epstein and Peters [4]. Equilibrium relative to the universal set of mechanisms will continue as equilibria no matter how the set of feasible mechanisms is extended. The second result in the paper is concerned with this issue.

⁸It might seem that all that this requires is that the seller offer a menu along with a suggestion about how the agent should choose from the menu (following Myerson). The difficulty with this approach is that the suggestion that the seller makes to the agent should depend on the agent's true type. Since this type depends on the market information that the agent has, this immediately makes the problem as difficult as trying to describe a universal set of mechanisms.

⁹As will be seen, the menus that are needed are more complex than those that have traditionally been used in the literature.

Unfortunately menus do not have the strong robustness properties of equilibria relative to the universal set of mechanisms. However, they do have a weaker robustness property in the sense that for any extension of the set of mechanisms, it will always be possible to construct a continuation equilibrium for the extension such that the payoffs from the original equilibrium are preserved.

1 Basics

There are n sellers dealing with a single buyer. The terms buyer and seller need not be interpreted literally. This terminology is common in the literature on competing mechanisms ([9], [15], or [4]). The literature on common agency has been more creative in generating examples, for example multiple lenders dealing with a single borrower (Biais, Martimort and Rochet [2] and Parlour and Rajan [13]), multiple regulators overseeing the operation of single firm (Berreheim and Whinston [1]) or multiple lobbyists trying to influence government policy (Dixit and Grossman [3]). The trading process begins when sellers simultaneously announce the mechanisms they plan to use. An important assumption is that the buyer observes market information completely while no seller can directly observe the mechanism that has been offered by any other seller. After seeing the contracts offered by each seller, the buyer communicates with sellers and take actions that affect the payoffs of all the sellers.

The primitives of the model are

- A: space of 'simple' actions available to each seller
- E: space of actions that the buyer carries out for the sellers
- : valuations space (including the 'usual type' of a buyer)
- F : a mapping to which (all sellers believe) the buyer's valuation is raw in

To avoid measure theoretic considerations it will be assumed that the sets A, E and - are all finite. An element of the set A is a complete description of the contractible part of the relationship between the seller and the buyer. This would include any transfer between buyer and seller and any action

taken by the seller. Actions for the seller might include contracts that specify contractible actions that the buyer has to take. An important assumption is that sellers cannot write contracts contingent on the contracts or other actions of the other sellers. The set E describes the set of uncontractable actions available to the buyer. The set E might consist of a vector of contracts that the buyer undertakes for each seller. Alternatively, in a competing mechanism problem, E is just a list of all the sellers. The buyer's choice simply reflects the mechanism that the buyer chooses to participate in.

Buyers and sellers have expected utility preferences. Sellers' payoffs are represented by $v : A^n \times E \rightarrow [0; 1]$. For the buyer, payoffs are represented by the function $u : A^n \times E \times \{1, 2, \dots\} \rightarrow [0; 1]$. To simplify some of the arguments below, it is assumed that there is an action a_j that each seller j can take such that for any $a_j \in A$

$$u(a; a_{-j}; e; 1) \geq u(a_j; a_{-j}; e; 1)$$

for all $a_{-j} \in A^{n-1}$; $e \in E$ and for any $1 \leq j \leq n$. This action is used at various points to prevent the buyer from sending messages the seller does not want to hear under any circumstances.

1.1 Standard Model of Competition

For any measurable set X , let $\mathcal{P}(X)$ denote the set of probability measures on X . When a topology on $\mathcal{P}(X)$ is needed, the topology of weak convergence is assumed. For any $x \in \mathcal{P}(X)$, let $\text{supp } x$ be the support of the distribution x . Fix a measurable space of messages C that the buyer can send to a seller. For simplicity it is assumed that this set of messages that the buyer can send to each seller is the same. Let R be the (measurable) set of responses that sellers can send back to the buyer. Both these message spaces are assumed compact metric, but beyond that are perfectly general in the degree and nature of the communication about the other sellers' mechanisms that they permit. For example, the messages could allow the buyer to communicate the mechanisms being used by other sellers. The messages could be contingent plans describing the way the buyer will communicate under different circumstances.

An indirect mechanism ϕ_j for seller j specifies feasible messages C and R , as well as the probability distribution over actions and messages that the

seller will use to respond to any message that the buyer sends him. So \circ_j is a measurable map from C into $4(A \times R)$. With a slight abuse of notation let $\circ = (\circ_1, \dots, \circ_n)$ refer to the entire array of mechanisms offered by the sellers. For any particular element of $4(A \times R)$, each report $r \in R$ that is sent with positive probability induces a unique posterior belief about the probability with which the seller has chosen each action. This posterior probability lies in $4(A)$. Since Bayes rule involves a measurable map from $R \rightarrow A$, each mechanism can be associated with a measurable map from C into $4(A \times 4(A))$. In words, each message induces a joint distribution on actions and posterior beliefs. It will be convenient below to think of indirect mechanisms in this way. It will be assumed in the sequel that the image of each feasible \circ_j is closed in $4(A \times 4(A))$.

Denote by \mathcal{J} the set of feasible indirect mechanisms endowed with some topology. The goal of this paper is to try to characterize 'good' models of indirect competition i.e., to specify 'good' choices for \mathcal{J} . Good models of indirect equilibrium should be such that equilibrium exists so for the purposes of this paper, it is innocuous to assume that \mathcal{J} has the 'usual' desirable properties for example that it is compact metric.

The allocations are then determined by the following process. First, sellers simultaneously and publicly announce the mechanisms that they plan to use to determine their allocations. The buyer then communicates with each of the sellers by sending each of them one of the messages in the sets C_j that they specify. The sellers respond and choose actions according to the \circ_j that they have committed themselves to. Once the buyer sees the sellers' response to his message, he chooses his sort.

The simplest example might be price competition. Suppose that each seller has a single unit of output to offer and that the buyer wishes to purchase exactly one unit. Each seller announces a price. The buyer chooses one and only one seller as a potential trading partner after seeing the prices that the sellers have to offer. In this case there is no communication so that C and R are empty for every seller. The buyer's sort is simply the seller that he selects so that E is simply a list of all the sellers.

Buyer behavior in each mechanism depends on the buyer's evaluation and on the mechanisms that he observes being offered by the other sellers. A communications strategy is a measurable mapping $\epsilon: \mathbb{E}^n \rightarrow 4(C^n)$ that describes the (probability distribution over) messages that the buyer will send to the sellers as a function of the buyer's type and the array of mechanisms

that he is offered by the sellers. A decision strategy ψ : $E_j^n \times C^n \times R^n \rightarrow \Delta(E)$ is a measurable mapping that describes the probability distribution the buyer 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 sellers. The pair (ϵ, ψ) together constitute a continuation strategy for the buyer. Say that the continuation strategy (ϵ, ψ) is a continuation equilibrium if the buyer has no incentive to deviate from this strategy for any of his valuations for any array of mechanisms offered by the sellers or for any array of messages that he receives from the sellers.

Again since j is intended to represent a reasonable model of the competition among the sellers it is natural to assume that continuation equilibria exist. It is not difficult to construct models of indirect competition where continuation equilibria do not exist (one mechanism might be - I will trade with the buyer who names the largest integer).¹⁰ It is also easy to think of models of indirect competition where sellers can offer mechanisms that do not make sense (each seller offers a price equal to the price offered by the other seller). There is no need to worry about whether such models are good descriptions of competition between sellers - it is immediately apparent that they are not. Thus in the discussion that follows we restrict attention to indirect models that possess all the usual desirata of predictive models.

It should also be noted that when a particular continuation equilibrium is assigned, this does not involve or require that this continuation equilibrium be unique. A mechanism is of interest to a seller partly because of the continuation equilibrium that it delivers. Alternative equilibrium selections will be interpreted as alternative models of indirect competition in this paper.

When we want to emphasize the underlying set of indirect mechanisms; we refer to (ϵ, ψ) as a continuation equilibrium relative to j . When we wish to emphasize a particular array of mechanisms we refer to $(\epsilon(\Phi^j), \psi(\Phi^j))$ as a continuation equilibrium relative to j .

The key to the standard (one seller) revelation principle, is that composing a mechanism with buyers' strategies yields a mapping from valuations into actions or in other words a 'direct mechanism'. In the context of this

¹⁰ In fact Myerson [12] gives an example in which the sellers equilibrium does not exist when sellers are constrained to use naive direct mechanisms in which the buyers report only their types.

paper, there is a corresponding composition. Let

$$m_{\epsilon^k} := E_i^n | \! - A^n E E, \quad (1)$$

be the joint distribution of actions and efforts induced by a particular continuation equilibrium (ϵ^k).

These definitions make it possible to describe the buyer's and seller's payoffs. Suppose that the competing firms choose the randomizations $\pm = f_{\pm_1} \dots f_{\pm_n}$ with $\pm_k \in \{1, 2, 4\}$ for all $j = 1 \dots n$, and that buyer behavior is described by the continuation equilibrium strategy (ϵ^k). The buyer's payoff depends on the mechanisms that sellers offer. Let $\circ = f^0_1 \dots f^0_n$ be any array of mechanisms. Then the payoff that the buyer gets is

$$\begin{aligned} U_i(\circ; \epsilon^k; !) = \\ \sum_{a_1 \dots a_n} \sum_{e} u(a_1 \dots a_n; e; !) m_{\epsilon^k}(!; \circ) \end{aligned} \quad (2)$$

The seller who chooses the randomization \pm_j expects the payoff

$$\begin{aligned} V_i(\pm_j; \pm_{i,j}; \epsilon^k) = \\ \sum_{a_1 \dots a_n} \sum_{e} v(a_1 \dots a_n; e) m_{\epsilon^k}(!; (\circ_j; \circ_{i,j})) F(!) \end{aligned} \quad (3)$$

where

$$F(\circ_j; \circ_{i,j}) = \sum_{a_1 \dots a_n} v(a_1 \dots a_n; e) m_{\epsilon^k}(!; (\circ_j; \circ_{i,j})) dF(!)$$

In an obvious notation $d_{\pm_{i,j}}(\circ_{i,j}) = f_{\pm_1} \dots d_{\pm_{j-1}} d_{\pm_j} \dots d_{\pm_n}$.

Say that $(\epsilon^k; \pm_1 \dots \pm_n)$ is an equilibrium relative to $(M; i)$, or simply that $(i; \epsilon^k; \pm_1 \dots \pm_n)$ is an equilibrium for M , if (ϵ^k) is a continuation equilibrium and for each $j = 1 \dots n$ and

$$\pm_j \in \arg \max_{\pm_{i,j} \in \{1, 2, 4\}} V_i(\pm^0, \pm_{i,j}; \epsilon^k).$$

It should be apparent from this that the actions that any particular seller takes and the effort level that is induced depend not only on the buyer's valuation but also on the mechanisms that have been offered by the other firms. It might seem reasonable to model common agency in the familiar way by having sellers choose allocations in response to the buyer's reports.

about preferences. Actions and efforts will generally depend on mechanisms offered by other firms. It should be clear that the allocations and effort induced by some model of direct competition will not be reproducible with this kind of direct mechanism. This was illustrated initially by the example in Peck [14]. A complete series of examples illustrating the pitfalls of this approach are given in Martimort and Stole [8] for the common agency case.

2 Menus

In the simplest single seller single buyer 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.

In its simplest form, a menu is simply a mechanism with a message space equal to the set of actions A along with a mapping $\circ^{\alpha} : A^{\alpha} \rightarrow A$. A satisfying \circ^{α} (a) $\in A$ for all $a \in A^{\alpha}$. The interpretation is that the mechanism \circ^{α} either gives the buyer the action that he requests or gives the buyer a . The seller's message space R is set equal to the empty set. This is the approach taken in [2] who allow the sellers to compete in nonlinear pricing schemes. The action that the seller takes in that case consists of the monetary transfer and an asset trade.

In general, mechanisms formed by offering menus of this kind will be restrictive since sellers might want to exploit buyer risk aversion by offering random mechanisms in which the buyer is offered a distribution over the set of outcomes for some of his types. If the seller does this he might also want to send a random reply to the buyer that is correlated with the action that he chooses in order to maintain some connection between his action and the effort that the seller makes. Menus of the simple kind described above can not be used to reproduce this kind of behavior.

So to model menus in this more general setting suppose the seller uses the message space $R = \{A\}$ and let $A^{\alpha} = \{A \in \mathcal{P}(A)\}$. The message that the seller sends in this case should be interpreted as a recommendation

to the buyer about the posterior probability with which he should believe that the seller has taken the various actions. A action for the seller is then a joint distribution over simple actions and signals. In the formulation that follows the seller will offer the buyer the choice from a menu of actions in A^x . For this reason the set of messages C available to the buyer will be equal to A^x . Now define \circ^x to be the set of all mappings $\circ^x : A^x \rightarrow A^x$ satisfying

$$\circ^x(a) = \begin{cases} a & a \in P \\ \underline{a}^0 & a \notin P \end{cases}$$

for some closed subset $P \subseteq A^x$, where \underline{a}^0 is a degenerate measure that assigns all probability weight to the outcome a and sends the degenerate message \underline{a} with probability 1.

The following theorem shows that arbitrary indirect mechanisms can be replaced with mechanisms in which sellers offer the buyer a menu of alternatives (from A^x). This new menu is simply the set of alternatives that the buyer could have induced in the original direct mechanism by sending appropriate messages to the seller.

Theorem 1 Suppose that \mathcal{J} is some space of indirect mechanisms let $(\epsilon^{\frac{1}{n}}; \pm_1; \dots; \pm_n)$ be an equilibrium relative to \mathcal{J} . Then there is a map $\tilde{A} : \mathcal{J} \rightarrow \mathcal{J}^x$, an array of randomizations $(\pm_1^x; \dots; \pm_n^x)$ over \mathcal{J}^x and a continuation equilibrium $(C^x; \frac{1}{n})$ relative to \mathcal{J}^x such that

1.

$$U_{i^x}(\tilde{A}(\circ^x); C^x; \frac{1}{n}; \epsilon^{\frac{1}{n}}) = U_i(\circ^x; \epsilon^{\frac{1}{n}}; \epsilon^{\frac{1}{n}})$$

for all i and for all $\circ^x \in \mathcal{J}^x$; and

2. for each $\circ_j^x \in \text{supp } \pm_j$

$$V_i(\tilde{A}(\circ_j^x); \pm_{i,j}^x; C^x; \frac{1}{n}) = V_i(\circ_j^x; \pm_{i,j}^x; \epsilon^{\frac{1}{n}})$$

otherwise if $\circ_j^x \notin \text{supp } \pm_j$

$$V_i(\tilde{A}(\circ_j^x); \pm_{i,j}^x; C^x; \frac{1}{n}) = V_i(\pm_j^x; \pm_{i,j}^x; \epsilon^{\frac{1}{n}})$$

Every equilibrium relative to j has a payoff equivalent representation in i^* where the representation is given by the mapping \tilde{A} . The representations are equivalent in the sense if all sellers' mechanisms are translated into i^* the buyer gets the same payoff in the continuation equilibrium relative to these translated mechanisms as he does in the original continuation equilibrium. Furthermore sellers get the same payoff from translated mechanisms as they receive in the initial equilibrium provided that they play one of their equilibrium strategies. Unprofitable deviations for sellers in j correspond after translation into unprofitable deviations relative to i^* .

The translation \tilde{A} is straightforward. It simply takes all of the actions that the buyer could have induced by sending some message to the seller, $\hat{\sigma}_j(C)$, and offers this to the buyer as a menu of alternatives that he need sto choose from. The buyer's continuation strategy is then constructed to reproduce his behavior in the original continuation equilibrium. This preserves the payoffs of the various players.

The proof is mathematically trivial, but complex because of the fact (mentioned previously) that many different indirect mechanisms might translate into the same menu of alternatives. This would not be a problem, except for the fact that the buyer could conceivably choose from the menu and select an effort in a way that depends on the manner in which the menu is presented. An example is presented to illustrate this below.

This is the reason that payoff equilibrium path mechanisms cannot be translated in a way that preserves the seller's payoffs.¹¹ If the same menu is offered in two different ways in the original game, and the buyer chooses different effort levels in response to each, then the seller's payoff will generally differ in the original continuation equilibrium depending on which of the two (menu equivalent) indirect mechanisms that he offers. The translation into menus will require that the buyer respond with only one effort to both mechanisms. If the two indirect mechanisms both lie on the equilibrium path, this won't be a problem because they will generate the same payoff for the seller anyway. All that can be guaranteed of the equilibrium path is that the translated mechanism will yield a payoff that is no larger than the equilibrium path payoff in the original game.

The reader might wonder whether this problem could be resolved if the

¹¹ It bears mention here that the universal class of mechanisms described by Epstein and Peters [4] generates a transformation that preserves these payoffs of the equilibrium path.

seller simply extends the menu by offering a message instructing the buyer how to select his effort (and how to choose from the menu). This is the approach that Myerson [11] used for the single principal problem. Indeed, this is why the revelation principle will normally require the principal to send a message to his agents recommending the actions they should take. This approach is not pursued here, partly because it is unnecessary for the purpose of this paper, and partly because it raises much more 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 agent's full type including all his market information. To construct the appropriate messages mechanism need to deal with the conceptually complex considerations associated with universal set of mechanisms as discussed in [4] - precisely the considerations that this paper is trying to avoid.

The second point to keep in mind is that equilibria for sellers relative to j will not generally translate into equilibria relative to j^x . The reason is that j^x will typically incorporate various implicit restrictions on the mechanisms that sellers are allowed to use. This is expressed formally by observing that $\tilde{A}(j) \neq \tilde{A}(j^x)$. 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 j is 'small' enough in j^x the equilibrium outcomes relative to j^x will generate payoffs that cannot be attained by any combination of mechanisms in j . This is the basis of the example given in Peck [14]. Secondly, even if the payoffs in some equilibrium relative to j^x are attainable by combinations of mechanisms in j , these payoffs may not be supportable as equilibria relative to j . This is the basis of the example given in Martimort and Stole [8].

The third point to note is that the theorem given here is considerably weaker than the one in Epstein and Peters [4]. In that paper conditions are given under which every mechanism in j (or on the equilibrium path) is translated into a payoff equivalent mechanism in j^x (again the translation here is not payoff equivalent on the equilibrium path).

Despite the fact that menus in A^x are relatively tractable, they are considerably more complicated than the menus that have been used in the literature. Nonlinear pricing schemes, for example, involve menus of simple actions without any randomization of any messages from the seller to attempt to influence buyer behavior. The nonlinear pricing problem with a single agent is simpler than the problem considered here if the buyer takes

no uncontractable action on behalf of the seller. In this case, there is no need for the seller to send messages to the buyer to influence his effort, and the set of menus of elements from 4 (A) will be sufficient to prove a revelation principle similar to the one above.

The required randomization means that this set of menus will still be more complicated than the simple nonlinear pricing schemes that have been used in the literature. Optimal incentive schemes with risk aversion will often require the use of random outcomes. The typical approach is to assume that randomization is impossible. A previous version of this paper assumed that contractible randomization was difficult enough to keep the set of feasible actions (of either a random or nonrandom kind) finite. In this case, a version of the Theorem 1 can be proved in which sellers offer only menus of simple actions.

3 Example

This section develops a simple example to illustrate the revelation principle proved above. There are two sellers and a single buyer who can have one of two possible types μ_1 or μ_2 with equal probability. The sellers have three available actions called A , B , and C . The buyer does not make any effort and it is assumed impossible for the seller to write contracts that commit him to any specific randomization over these three outcomes.¹² The payoffs to seller 1, seller 2 and the buyer respectively are listed in the following table.

			μ_1			μ_2		
			A	B	C	A	B	C
A	0;0;0	2;0;2	1;1;5;3			3;3;3	0;0;0	0;0;0
B	0;2;2	1;1;1	0;0;0			0;0;0	1;1;1	0;0;0
C	5;1;3	0;0;0	0;0;0			0;0;0	0;0;0	0;0;0

The first box is a variant of the example presented in Hartmanis and Stole [8].¹³ There is an equilibrium for this problem in which both sellers offer the

¹²This does not stop him, of course, from using a random strategy to choose among the three actions.

¹³They describe a menu equilibrium for this box in which each seller offers the menu $\{B; C\}$ to the buyer who then chooses the option B from both. This is not an equilibrium if contractible randomizations are possible. If they are either seller can profitably deviate

action C to the buyer. This can be interpreted as an equilibrium in direct mechanism where each seller offers a direct mechanism in which the action chosen is independent of the buyer's report. This is also an equilibrium in menus (in which each seller offers a degenerate menu).

If the competition is modelled by allowing sellers to offer menus of simple actions to the buyer, there is a relatively simple equilibrium in which each seller offers the buyer the menu $f_A; C g$. The buyer chooses A from both sellers when his type is μ_2 and randomizes choosing A from one seller and C from the other with equal probability when his type is μ_1 . The buyer's payoff in this case is 3 whatever his type. Each seller gets 3 when the buyer is type μ_2 and 2 (that is, $\frac{1}{2}5 + \frac{1}{2}1$) when the buyer is type μ_1 .

To show that this is an equilibrium in menus it is necessary to check that neither seller can profitably deviate by modifying the menu that he offers. For example, consider the payoff that the seller can get by deleting the option A from his menu. If the buyer has type μ_1 he will choose A from the other seller for sure and the deviating seller will get payoff 5. The downside is that the seller now receives payoff 0 when the buyer is type μ_2 . This gives the seller's expected payoff as $5 \cdot \frac{1}{2} + 0 \cdot \frac{1}{2} = 2.5$, exactly what it was in the initial equilibrium. It is straightforward to check that deviations to alternative menus are also unprofitable.¹⁴

Alternatively, considering modelling the competition by allowing sellers to choose from a slightly larger set of contracts consisting of all possible menus and all direct mechanisms (meaning that the buyer reports either μ_1 or μ_2 and the seller responds with one of the actions A, B, or C). Refer to the menu $f_A; C g$ as the mechanism \circ^m in this space. Then no deviation that the seller might consider is to offer a direct mechanism \circ^d which chooses the action C if the buyer reports μ_1 and A if the buyer reports μ_2 . Suppose that the buyer responds to this deviation by reporting μ_2 to the deviator no matter what his

¹⁴ Alternatively, consider a random mechanism (that is, a degenerate menu from 4 (A)) that chooses action A with small enough probability. I am indebted to Sergei Severinov for pointing this out to me.

¹⁴ If contractible randomizations are possible, this menu will not be an equilibrium. A profitable deviation would involve offering a menu consisting of the outcome C for sure, or a random outcome that provides A with high probability and C with complementary probability. If the buyer has type μ_1 he can no longer randomize, but is forced instead to choose the outcome C for sure with the deviator. Since the option A is still available, the buyer will choose it when he has type μ_2 .

type is (so that the deviator always chooses A). Then when the buyer's type is μ_1 he chooses the action C from the other seller's menu. This continuation equilibrium makes the deviation strictly unprofitable since the deviator loses any chance at the payoff 5 that he gets with positive probability along the equilibrium path.

To translate σ^d into the corresponding menu observe that if the other seller offers the degenerate menu consisting only of the action A than the buyer will report truthfully under σ^d . Thus σ^d is just equal to the menu consisting of the options fA; C g, obviously the same as σ^m . The same menu of options can be offered in two distinct ways and the buyer responds to these offers differently. This difference makes the sellers' profits strictly lower when he offers σ^d than when he offers σ^m in the original game. Then when the mechanism σ^d gets translated into j^x the payoff that the seller gets from the translation is strictly higher than the payoff he got from the deviation in the original equilibrium. Nonetheless the payoff is no larger than his equilibrium payoff associated by Theorem 1.

This example has another property. Turn the argument around and imagine starting with the set of menus and extending the set of mechanisms available to the seller to include direct mechanisms. Whether the sellers will continue to offer the menu fA; C g depends on how they expect buyers to respond.. If the buyer is expected to report truthfully when the seller deviates to the direct mechanism σ^d , then the equilibrium in menus will fall apart since the deviating seller will then ensure the payoff 5 for himself when the buyer has type μ_1 . This illustrates that the equilibrium in menus is not robust in the sense of Epstein and Peters [4].

This leads nicely into the second theorem in the paper. There is a continuation equilibrium that could be assigned to the deviation that will break up the equilibrium, but there is also one that will preserve it. The theorem in the next section shows that this is generally true with common agency.

4 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 sellers 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. This is illustrated in the example given above. These same examples suggest an alternative, since the ability of the seller 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 there always seem to be an alternative assignment that does not increase the seller's profits. 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 sellers 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 direct mechanisms, \mathcal{M}_i and \mathcal{M}_i^1 . Say that $\mathcal{M}_i^1 \subset \mathcal{M}_i$ (\mathcal{M}_i^1 is bigger than \mathcal{M}_i) if there exists an embedding $\Phi^0: \mathcal{M}_i \rightarrow \mathcal{M}_i^1$. This idea is natural, there are more mechanisms in \mathcal{M}_i^1 than in \mathcal{M}_i . Now take two models of competition with feasible sets of mechanisms \mathcal{M}_i and \mathcal{M}_i^1 such that $\mathcal{M}_i^1 \subset \mathcal{M}_i$. Let the associated continuation equilibria be (ϵ_i) and (ϵ_i^1) respectively. As before, let $m_{\epsilon_i}(\Phi^0)$ and $m_{\epsilon_i^1}(\Phi^0)$ be the joint distributions on actions and effort induced by these continuation equilibria when the array of mechanisms is ϵ^0 and ϵ^1 respectively. The continuation equilibrium $(\mathcal{M}_i^1; \epsilon_i^1)$ is said to extend $(\mathcal{M}_i; \epsilon_i)$ if there is an embedding $\Phi^0: \mathcal{M}_i \rightarrow \mathcal{M}_i^1$ such that for all ϵ^0 in \mathcal{M}_i

$$m_{\epsilon_i}(\Phi^0) = m_{\epsilon_i^1}(\Phi^0(\epsilon^0))$$

The generalization of the idea behind direct mechanisms in the single mechanism designer problem is to imagine sellers exploring more complex mechanisms (in \mathcal{M}_i^1) that are not envisioned by the 'ad hoc' model of competition specified in the economic model of competition (i.e., \mathcal{M}_i). 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 sellers' feasible sets.

Definition 2 (Strong Robustness) An equilibrium $(\mathcal{M}_i; \epsilon_i; \pm)$ is said to be strongly robust if for any extension $(\mathcal{M}_i^1; \epsilon_i^1)$ for which \mathcal{M}_i^1 is compact metric, $(\mathcal{M}_i^1; \epsilon_i^1; \Phi^0[\pm])$ is an equilibrium, where $\Phi^0[\pm]$ is the randomization on \mathcal{M}_i^1 induced by \pm and Φ^0 .

Strongly robust equilibria are analyzed in Epstein and Peters [4]. 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 3 (Weak Robustness) An equilibrium $(j^*; \epsilon^*; \pm^*)$ is said to be weakly robust if for every compact metric $j^1 < j^*$ there exists an extension $(j^1; \epsilon^1; \pm^1)$ such that $(j^1; \epsilon^1; \pm^1)$ is an equilibrium, where \pm^1 is the randomization on j^1 induced by \pm^* .

This leads to our second result

Theorem 4 In the common agency problem with a single agent, let $(j^*; C^*; \epsilon^*; \pm^*)$ be an equilibrium. Then $(j^*; C^*; \epsilon^*; \pm^*)$ is weakly robust.

To see the idea behind the proof, it might help to consider the case where there is only a single mechanism designer. For any arbitrary mechanism \circ that this seller offers from j^* , the image $\circ(C)$ is the 'menu' of alternative actions that the buyer can choose from. Typically there will be one (or more) mechanisms \circ in j^* that offer the buyer the same menu of alternatives as $\circ^* 2 j^*$. Typically these will yield the seller a different payoff than the payoff he gets from \circ^* . The reason for this is that the continuation equilibrium will induce the buyer to send different messages and take different actions than he does when he sees \circ^* . This is precisely why the seller might like the larger set of mechanisms. Yet since the menu of actions and efforts available to the buyer is the same when he is offered \circ as it is when he is offered \circ^* , 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 buyer choose from the 'menu' associated with \circ in exactly the same way as he does when he chooses from the menu \circ^* . Then in this new continuation equilibrium, the seller must get the same payoff when he offers \circ as he does when he offers \circ^* . With this construction every payoff that the seller can attain in j^* coincides with a payoff that he could have attained in j^* . It follows immediately that whatever was optimal for the seller relative to j^* will then be optimal relative to j^* . The proof simply extends this reasoning to the case where there are many sellers.

5 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 1 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 his choice varies with the way the choice set is 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 agent's 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 impossible to deal with this (see [4]) but the appropriate set of mechanisms is complex.

6 Appendix - Proofs

We begin with some basic definitions and relationships. For any subset $P \subseteq A^x$ define $(P) : A^x \rightarrow A^x$ such that

$$(P)(a) = \begin{cases} a & \text{if } a \in P \\ a^0 & \text{otherwise} \end{cases} \quad (4)$$

In word $\zeta^-(P)$ is a function that respondsto any joint distribution (action) in P with the same joint distribution (over actions and messages), otherwise it responds with a^0 (recall this is a degenerate distribution that assigns all probability weight to the pair $(a; a)$). With this understanding we will refer to $\zeta^-(P)$ as a menu of alternatives from A^α .

Each mechanism ζ_j and report ζ_j produces a joint distribution over simple actions and messages for the seller. Each message that the buyer receives induces some posterior belief in $\zeta_j(A)$ about the simple action that the seller has undertaken. Thus each report that the buyer makes to the seller induces a pure action in $A^\alpha \setminus A \in \zeta_j(A)$. With a slight abuse of notation denote the set of actions ($\zeta_j(A)$) that the buyer can induce by sending some message to the seller as $\zeta_j(C) \subset A^\alpha$. Recall that this set is closed in A^α by assumption.

The buyer chooses his report after observing the message associated with the action $a^\alpha \in A^\alpha$. Together, the sellers strategies \pm , the continuation strategies ($\epsilon^{1/4}$) and the buyer's type ι induce a joint distribution over $(A^\alpha)^n \in E$.

Definition 5 For any $D \subset \iota^n$ such that $\zeta(D) > 0$ define $b(\iota; D)$ to be the joint distribution on $(A^\alpha)^n \in E$ induced by the equilibrium $(\iota; \epsilon^{1/4}; \pm)$ conditional on the sellers mechanisms being drawn from the subset $D \subset \iota^n$.¹⁵ Write $b^{A^\alpha}(\iota; D)$ as the marginal distribution of b on the space of actions $(A^\alpha)^n$.

Write $(a; q)$ as a typical element of $A \in \zeta(A)$. Since the buyer's posterior beliefs are defined before he is forced to choose a report, the utility associated with any particular array of actions $a^\alpha \in (A^\alpha)^n$ is given by

$$y(a^\alpha; \iota) = \max_{\epsilon} \max_{a_1, \dots, a_n} \max_{q_1, \dots, q_n} u(a_1, \dots, a_n; \epsilon; \iota) q_1(a_1) \dots q_n(a_n) d a_1^\alpha(a_1; q_1) \dots d a_n^\alpha(a_n; q_n)$$

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

$$U_i(\iota; \epsilon^{1/4}; \iota) =$$

¹⁵The fact that this distribution depends on $\epsilon^{1/4}$ and \pm is suppressed to ease notation.

$$\begin{aligned}
 & z \quad z \\
 & \text{defn } u(a_1, \dots, a_n; e; \cdot) \text{dm}_{\epsilon^k}(\cdot; \cdot) = \\
 & z \quad z \\
 & \text{defn } fy(a_1^{\alpha}, \dots, a_n^{\alpha}; \cdot) g \text{db}^A(\cdot; \cdot) = \\
 & \max_{j=1}^n y(a_1^{\alpha}, \dots, a_n^{\alpha}; \cdot) : a_j^{\alpha} \geq \bar{a}_j(C) \quad (5)
 \end{aligned}$$

The first equality follows from the fact that every effort level that the buyer 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 \cdot . The use of \max instead of \sup in the final operation results from the assumption that a continuation equilibrium always exists.

6.1 Proof of Theorem 1

Proof. To begin the map \tilde{A} is described, and the strategies that sellers and the buyer use in the continuation equilibrium relative to the new space of mechanisms are described. Define the transformation \tilde{A}_j from \mathcal{J} into \mathcal{I}^{α} as

$$\tilde{A}_j : \mathcal{J} \ni \cdot_j^{\alpha} (\Phi) = \cdot_j(C)(\Phi)$$

This transformation simply converts a mechanism \cdot_j^{α} into the menu of alternatives in A^{α} that it provides. Let \pm_j^{α} be the measure induced by \cdot_j^{α} on \mathcal{I}^{α} by this map.

\tilde{A}_j will generally be a many-to-one mapping with inverse correspondence $\tilde{A}_j^{-1}(\Phi)$. Choose an arbitrary selection $\tilde{A}_j^{-1}(\Phi)$ from this correspondence. For any collection $f^{\alpha}_1, \dots, f^{\alpha}_n$ define the sets

$$B_j^{\alpha} = \begin{cases} \tilde{A}_j^{-1}(\cdot_j^{\alpha}) \setminus \text{supp } \pm_j^{\alpha} & \text{if } \tilde{A}_j^{-1}(\cdot_j^{\alpha}) \setminus \text{supp } \pm_j^{\alpha} \neq \emptyset; \\ \tilde{A}_j^{-1}(\cdot_j^{\alpha}) & \text{otherwise} \end{cases} \quad (6)$$

We can now specify the continuation equilibrium relative to \cdot_j^{α} . Begin with the case where $\cdot_j^{\alpha} \in \tilde{A}_j(\mathcal{J})^n$. By the definition of \cdot_j^{α} , the buyer must choose a distribution over actions from each seller that lies in a closed subset of A^{α} , as well as an effort from a finite set. Since the buyer's payoffs are linear (by expected utility) with respect to the distribution of actions that sellers choose, they are continuous. It follows that a best choice-action pair from each menu exists for the buyer. 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 $\hat{C}^{\alpha} = f^{\alpha_1} \cdots f^{\alpha_n} g^{\alpha}$, then

$$C^{\alpha}(\cdot; \cdot) = b^{A^{\alpha}}(\cdot; fB_1(\cdot_1); \dots; B_n(\cdot_n))g^{\alpha} \quad (7)$$

where $b^{A^{\alpha}}(\cdot; fB_1(\cdot_1); \dots; B_n(\cdot_n))g^{\alpha}$ denotes the marginal distribution on $(A^{\alpha})^n$ induced by $b(\cdot; fB_1(\cdot_1); \dots; B_n(\cdot_n))g^{\alpha}$.

To define the buyer's best choice, let $a^{\alpha} = fa_1^{\alpha} \cdots a_n^{\alpha}$ denote the buyer's actual choice from the menu offered by each seller, and let $q \in \Delta(A)^n$ be the array of messages that the buyer receives. Note that the buyer's posterior beliefs (after seeing the sellers' messages) also lie in $\Delta(A)^n$. The buyer chooses his best level according to

$$\hat{V}^{\alpha}(\cdot; \cdot; q) = b^E(\cdot; fB_1(\cdot_1); \dots; B_n(\cdot_n))g^{\alpha}; q \quad (8)$$

where $b^E(\cdot; fB_1(\cdot_1); \dots; B_n(\cdot_n))g^{\alpha}; q$ is the distribution of e^{α} induced by $b(\cdot; fB_1(\cdot_1); \dots; B_n(\cdot_n))g^{\alpha}$ conditional on the sellers' actions being given by a^{α} and the buyer's posterior beliefs being given by q .

The next step is to show that this strategy constitutes a continuation equilibrium, and that the buyer'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 buyer receives in the continuation equilibrium relative to $(\cdot^{\alpha}; C^{\alpha}; V^{\alpha})$ after choosing the array a^{α} from the various sellers' menus and receiving the array of messages q from the sellers. From (8), this is given by

$$z^{\alpha} = \max_{a_1, \dots, a_n} u(a_1, \dots, a_n; e; \cdot) q_1(a_1), \dots, q_n(a_n) \quad b^E(\cdot; fB_1(\cdot_1); \dots; B_n(\cdot_n))g^{\alpha}; q \quad (9)$$

Consider any action a_j^{α} in the menu offered by seller j . By construction each such action must coincide with a joint distribution of simple actions and posterior beliefs that the buyer can induce (by sending appropriate messages) to the seller in the continuation equilibrium relative to $(\cdot_j; \cdot_{-j}; e^{\alpha})$ for any $\cdot_j \in B_j$. It follows that any message q_j that the buyer could receive when he selects a_j^{α} from seller j 's menu must be equal to the correct posterior for a conditional on that message. Then from the properties

of the continuation equilibrium ($\epsilon \frac{1}{4}$) every effort level in the support of $b^E(\cdot; fB_1(\cdot_1); \dots; B_n(\cdot_n)g; j; q)$ must maximize the buyer's expected payoff conditional on the posterior belief q . Thus (9) can be simplified to

$$\max_e \max_{a_1} \max_{a_n} u(a_1; \dots; a_n; e; \cdot) q_1(a_1) \dots q_n(a_n)$$

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

$$\begin{aligned} z &= z' \\ &= \max_e \max_{a_1} \max_{a_n} u(a_1; \dots; a_n; e; \cdot) q_1(a_1) \dots q_n(a_n) \cdot a_1^\alpha(a_1; q_1) \dots a_n^\alpha(a_n; q_n) \\ &= y(a^\alpha; \cdot) \end{aligned}$$

The buyer selects according to the distribution $b^A(\cdot; fB_1(\cdot_1); \dots; B_n(\cdot_n)g)$. So this strategy yields the payoff

$$\begin{aligned} z &= y(a^\alpha; \cdot) \cdot b^A(\cdot; fB_1(\cdot_1); \dots; B_n(\cdot_n)g) \\ &= E_{B(\cdot)} y(a^\alpha; \cdot) \cdot b^A(\cdot; fB_1(\cdot_1); \dots; B_n(\cdot_n)g) \quad (\text{by definition}) \\ &= E_{B(\cdot)} \max_{\mathbf{a}} y(a_1^\alpha; \dots; a_n^\alpha; \cdot) : a_j^\alpha \geq [b(\Phi_j; \Psi_j = 1; \dots; n)]_j \quad (\text{from (5)}) \\ &= \max_{\mathbf{a}} y(a_1^\alpha; \dots; a_n^\alpha; \cdot) : a_j^\alpha \geq [b(\Phi_j; \Psi_j = 1; \dots; n)]_j \end{aligned}$$

(because of the fact that $[b(\Phi_j; \Psi_j = 1; \dots; n)]_j$ is the same for all $a_j \in B_j$). Since the choice sets in this last expression are the ones ordered in the menus associated with the mechanisms in \cdot^α , this line verifies that the continuation strategy specified for the buyers is a continuation equilibrium. Finally applying (5) once again gives this payoff equal to

$$= U_i(\cdot; \epsilon \frac{1}{4}; \cdot)$$

which establishes the first property in the theorem.

Finally, it remains to show that

$$V_i(\tilde{A}(\cdot_j); \pm_j^\alpha; C^\alpha; \epsilon \frac{1}{4}) = V_i(\cdot_j; \pm_j; \epsilon \frac{1}{4})$$

for all $\cdot_j \in \text{supp } \pm_j$.

For any $\hat{\psi}_{j,i}^{(0)}$

$$\hat{\psi}_{j,i}^{(0)} = \mathbb{E}_{\hat{\pi}_j} \left[v(a_1, \dots, a_n; e) \right]$$

$$= \mathbb{E}_{\hat{\pi}_j} \left[v(a_1, \dots, a_n; e) \mid m_{e \in \mathcal{M}} \neq \hat{\pi}_j \right]$$

Integrating this using the distribution $\hat{\pi}_j$ gives

$$\mathbb{E}_{\hat{\pi}_j} \left[v(a_1, \dots, a_n; e) \mid m_{e \in \mathcal{M}} \neq \hat{\pi}_j \right] = \mathbb{E}_{\hat{\pi}_j} \left[V_i(\hat{\pi}_j; e) \right]$$

If $\hat{\pi}_j = \tilde{\pi}_j$ for some $\tilde{\pi}_j \in \text{supp } \hat{\pi}_j$, then by the definition each $\hat{\pi}_j \in \text{supp } \hat{\pi}_j$ is an element of the support of $\tilde{\pi}_j$. Since $\tilde{\pi}_j$ is an equilibrium distribution for seller j , every $\hat{\pi}_j$ in its support gives the same payoff. Thus the integral in this last expression is constant and equal to

$$V_i(\hat{\pi}_j; e)$$

as required.

Alternatively if $\hat{\pi}_j \neq \tilde{\pi}_j \in \text{supp } \hat{\pi}_j$ then

$$\mathbb{E}_{\hat{\pi}_j} \left[V_i(\hat{\pi}_j; e) \right] = V_i(\tilde{\pi}_j; e) + V_i(\hat{\pi}_j; e)$$

which proves the theorem. ■

6.2 Proof of theorem 4

Proof. Let $(j; e)$ be an extension of $(j; C; \frac{1}{n})$ with embedding given by Φ . The method of the proof is to convert deviations that lie outside of the range of Φ 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 $\tilde{\pi}_j$ from Theorem 1 to associate with each mechanism $\tilde{\pi}_j$ the corresponding menu $\tilde{\pi}_j(\Phi = \tilde{\pi}_j(C))$. Now let

$D(\circ_1; \dots; \circ_n) \subseteq A^n \times E$ 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 $f^o_1; \dots; f^o_n$. If the buyer offered an array of mechanisms $f^o_1; \dots; f^o_n$ in the continuation equilibrium relative to j , he can induce any array of actions $a_1; \dots; a_n$ in the set $\circ_1(C) \times \dots \times \circ_n(C)$ by sending appropriate messages to the sellers. It follows that $D(\circ_1; \dots; \circ_n) = D(\circ_1^*; \dots; \circ_n^*)$. Since $(C^*; \frac{1}{n})$ is a continuation equilibrium relative to j^* , $m_{C^*; \frac{1}{n}}(l; \circ_1^*; \dots; \circ_n^*)$ maximizes the expected utility over $D(\circ_1^*; \dots; \circ_n^*)$ of a buyer of type l . It follows immediately that this distribution maximizes the buyer's expected utility over the set $D(\circ_1; \dots; \circ_n)$. Choose any continuation strategy C^* and $\frac{1}{n}$ relative to j that satisfies

$$m_{C^*; \frac{1}{n}}(l; \circ_1; \dots; \circ_n) = m_{C^*; \frac{1}{n}}(l; \bar{A}(\circ_1); \dots; \bar{A}(\circ_n))$$

(this strategy exists by the definition of D). Since $D(\circ_1; \dots; \circ_n) = D(\bar{A}(\circ_1); \dots; \bar{A}(\circ_n))$ this continuation strategy must be a continuation equilibrium.

Now we begin with an equilibrium $(\pm^*; \dots; \pm^*)$ relative to j^* with $\mathbb{R}[\pm^*]$ as the induced distribution relative to j . The payoff to some seller who unilaterally deviates to a mechanism \circ^o outside of $\mathbb{R}[\pm^*]$ while all other sellers offer mechanisms \circ_{-j} in the support of \pm^* is given by

$$\begin{aligned} & \mathbb{E}_{\pm^*} v(a_1; \dots; a_n; e) - m_{C^*; \frac{1}{n}}(l; (\circ^o, \circ_{-j})) \mathbb{E}(l) \\ &= \mathbb{E}_{\pm^*} v(a_1; \dots; a_n; e) - m_{C^*; \frac{1}{n}}(l; (\bar{A}(\circ^o), \bar{A}(\circ_{-j}))) \mathbb{E}(l). \\ & V_i(\pm_j; \pm_{-j}; C^*; \frac{1}{n}) = \\ & V_i(\mathbb{R}[\pm_j]; \mathbb{R}[\pm_{-j}]; C^*; \frac{1}{n}) \end{aligned}$$

which implies that the equilibrium $(j^*; \pm; C^*; \frac{1}{n})$ is weakly robust. ■

References

- [1] D. Bernheim and M. Whinston. Mechanisms, resource allocations and economic incentives. *Quarterly Journal of Economics* 101:131, 1987.

- [2] Bruno Biais, David Martimort, and Jean-Charles Rochet. Competing mechanisms in a common value environment. Technical report, Université de Toulouse, 1997.
- [3] Avinash Dixit, Gene Grossman, and Ehud Helpman. Common agency and coordination: General theory and application to government policymaking. *jpe*, pages 75–76, 1997.
- [4] L. Epstein and M. Peters. A revelation principle for competing mechanisms. University of Toronto, 1997.
- [5] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–611, 1973.
- [6] J. Green and J.-J. Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45:427–438, 1977.
- [7] Michael Katz. Game playing agents. Unobservable contracts as precommitments. mimeo, University of California at Berkeley, 1987.
- [8] D. Martimort and L. Stole. Communication spaces, equilibria sets and the revelation principle under common agency. mimeo University of Chicago, 1997.
- [9] Preston McAfee. Mechanism design by competing sellers. *Econometrica*, 61(6):1281–1312, November 1993.
- [10] Roger Myerson. Incentive compatibility and the bargaining problem. *Econometrica*, 47:61–73, 1979.
- [11] Roger Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:55–73, 1981.
- [12] Roger Myerson. Optimal coordination mechanisms in generalized principal agent problems. *Journal of Mathematical Economics*, 10:67–81, 1982.
- [13] Christine A. Parlour and Uday Rajan. Competition in loan contracts. Technical report, Carnegie Mellon University, 1997.

- [14] James Peck. Competing mechanisms and the revelation principle. Ohio State University Manuscript, page 12, 1995.
- [15] Michael Peters. On the equivalence of Walrasian and non-Walrasian equilibria in contract markets. *Review of Economic Studies*, pages 241–265, 1997.