Department of Economics
and
Institute for Policy Analysis
University of Toronto
150 St. George St.
Toronto, Canada
M5S 1A1

WORKING PAPER
NUMBER UT-ECIPA-PETERS-02

A REVELATION PRINCIPLE FOR COMPETING MECHANISMS

Larry G. Epstein Mike Peters*

December 9, 1996

Department of Economics
University of Toronto
ISSN 0829-4909
Author's email: peters@chass.utoronto.ca
On-line Version: http://www.epas.utoronto.ca:8080/ecipa/wpa.html

Abstract

In modelling competition among mechanism designers, it is necessary to specify the set of feasible mechanisms. These specifications are often borrowed from the optimal mechanism design literature and exclude mechanisms that are natural in a competitive environment; for example, mechanisms that depend on the mechanisms chosen by competitors. This paper constructs a set of mechanisms that is *universal* in that any specific model of the feasible set can be embedded in it. An equilibrium for a specific model is robust if and only if it is an equilibrium also for the universal set of mechanisms. A key to the construction is a language for describing mechanisms that is not tied to any preconceived notions of the nature of competition.

Journal of Economic Literature Classification Numbers: C72,D82

Department of Economics, University of Toronto, 150 St. George Street, Toronto, Canada, M5S 3G7, epstein@chass.utoronto.ca, peters@chass.utoronto.ca. We gratefully acknowledge the financial support of the Social Sciences Humanities Research Council of Canada and the comments of Jim Peck and seminar participants at Columbia, Yale and U. Montreal.

1. INTRODUCTION

There are many situations where sellers engaged in negotiations might want to know what promises competitors have made. For example, a seller might have incomplete information about the distribution of prices offered by his competitors. In a typical wage negotiation, the firm might want to offer a different wage depending on what wage the worker has been offered by the firm's competitors. A homeowner soliciting bids from contractors may be worried about a contractor's interest in completing a renovation project quickly. For this reason the homeowner might be interested in knowing about the contractor's other projects.

In each of these situations, the seller or firm or homeowner plausibly has less information about available alternatives than the trader with whom he is bargaining. Standard competitive models typically ignore this. For example, in wage search models ([10]) firms simply make take it or leave it wage offers without making any attempt to learn what other offers the worker thinks are available. A more sophisticated approach is to allow sellers to compete by offering alternative auctions [16], or even special kinds of *direct* mechanisms ([8, 15]) in which buyers are asked to report their willingness to pay. These models assume that sellers attempt to make outcomes depend on buyers' private information about their preferences, but stop short of allowing sellers to try to make outcomes depend on buyers' private information about what is happening in the rest of the market.

Such models are open to the criticism that they are overly restrictive. For example, a model of wage negotiations that ignores the seller's ability to attempt to extract information about alternate offers may not give accurate predictions about outcomes. The usual remedy for this is to invoke the *revelation principle*. The argument is based on the following idea. No matter what selling scheme the seller uses, the allocations associated with the scheme can be *replicated* using a *direct mechanism* where the buyer is asked simply to report his preferences and the seller assigns different outcomes to different reports. Any model of the set of feasible selling schemes (such as assuming that firms must make take it or leave it wage offers) amounts to an assumption that firms must choose from a subset of the set of direct mechanisms. The model is then sensible or *robust* if the seller's optimal mechanism within this subset is also an optimal mechanism with respect to the entire set of direct mechanisms.

In the multi-firm competitive setting, however, the class of direct mechanisms in which allocations depend on buyers' reports about their preferences is not sufficiently rich to allow any *equilibrium* to be replicated. That is because the actions undertaken in equilibrium at firm 1 will in general depend not only on buyers' preferences, but also on the other firm's mechanism and the latter may vary within an equilibrium. One reason for such dependence on the other firm's mechanism is that firm 1 may design his mechanism so that it asks buyers explicitly to report on the other firm; for example, because firm 1 may want to match or undercut his competitor's price as in [17]. Alternatively, even if the mechanisms are not explicitly interrelated, equilibrium buyer behavior at firm 1 may vary with firm 2's mechanism, because, for example, the latter may serve as a correlating device. Peck's [14] example of the inadequacy of valuations-based direct mechanisms can be interpreted in this way. We add that Peck's example helped to stimulate our interest and subsequent thinking on this topic.

In principle, it is clear how to deal with this. A buyer's information about other firms' mechanisms should be modelled as a part of his *type* in the sense of Harsanyi [5]. If we assume that the competitor's mechanism can be described in some concrete way, for instance, by describing the price that the other firm has offered, then this type is unambiguous; it is simply the combination of the description of the buyer's preference (for example, his willingness to pay) and his observation of the other seller's price. Given that the buyer will report the competing price, then when the first seller designs his mechanism, he can choose an allocation that depends on this second firm's price (assuming there is some incentive compatible way to learn about it) just as easily as he can make it depend on the buyer's valuation. But if it pays him to take advantage of this, then it will presumably also pay the other firm to do so. This means that if we want to take advantage of the revelation principle to avoid ad hoc restrictions on mechanisms in a *competitive* environment, we need to have the buyer describe not only the other firm's price, but also the way that the other firm's price depends on the price of the first firm, and on the way that the other firm's price depends on the way that the first firm's price depends ... and so on.

This raises two difficulties. First, there is a problem of infinite regress. The question arises whether there exists a *types* space having the property that a type from this space conveys all the information that the seller might ever want to know about the buyer and about the other seller's mechanism. The second difficulty seems even more challenging. Describing a firm through its price, the way that price depends on price, price depends on the way price depends...and so on, presumes that firms are employing a limited class of mechanisms. But they might be using more complex mechanisms like auctions. The rationale behind the revelation principle suggests that the types space should not depend on ad hoc properties associated with specific indirect mechanisms. Thus it is desirable to describe a space, or a language for describing mechanisms, that is independent of any preconceived notion of what competition is. In particular, this *universal language* should be sufficiently rich to permit descriptions of simple price setting, price matching, auctions or any other mechanisms in a large class.

One purpose of this paper is to propose a workable model of a types space for this competitive environment that resolves this problem of infinite regress and provides the needed universality. Our broader purpose is to provide a revelation principle that will work in a competitive environment. This we do by constructing a 'universal' class of mechanisms. The mechanisms in this class all have the property that buyers are asked to report a type from the universal type space described above to characterize both their preferences and the mechanism offered by the other firm. The equilibrium allocations associated with any specific competitive model satisfying some regularity conditions can be replicated by having sellers offer mechanisms in this universal class.

We want to know whether the equilibrium relative to some specific model of competition is an equilibrium also if sellers can deviate to a broader class of mechanisms than permitted by the specific (ad hoc) model of competition. It is only equilibria that are *robust* in this sense that are sensible bases for predicting behaviour. Our revelation principle provides a way to determine such robustness, in principle. As with the standard revelation principle, the equilibrium strategy of each seller in a specific competitive model can be converted into a randomization relative to the universal class of mechanisms. Thus we can consider the profitability of deviations by sellers to any mechanism in the universal class, taking the corresponding randomization of the other sellers to be fixed. Because of the universality of our class of mechanisms, we can conclude that, in a precise formal sense and subject to suitable assumptions as stated in Theorem 5.3, equilibria relative to our universal class are the *only* robust equilibria.

At a technical level, our analysis is related to the modelling of 'beliefs about beliefs' as in [9] and [2], for example. The associated infinite regress problem is resolved there through the analysis of appropriate hierarchies of probability measures. We adapt these analyses, more particularly their extension in [4], in order to resolve the problem of infinite regress that arises in our context. Because we choose to describe mechanisms via their implied expected payoff functions, and these are assumed to be upper semi-continuous, we are led to study hierarchies of upper semi-continuous functions. But our primary contribution is *not* to the literature on problems of infinite regress and their resolution through appropriate hierarchies (see also [6] and [19]). Rather, we feel that the innovative feature of our analysis is the demonstration that expected payoff functions can provide a universal language for describing mechanisms and that this can be achieved in such a way as to provide a revelation principle as described above. The problem of infinite regress is only one aspect of this analysis.

It may be useful to elaborate briefly on another aspect. It is clear that the analytical power of the standard revelation principle comes from the fact that the distribution of buyers' types facing a given seller is assumed to be independent of the mechanism offered by that seller (though see [15] for a model in which this distribution varies in a systematic way with the seller's mechanism). In a competitive environment, this property is difficult to attain since mechanisms are inherently *self-referential*. To see this, suppose we want to use the allocations associated with mechanisms to describe them. As pointed out above, the allocation at firm 1 will generally depend on the mechanism or allocation available at firm 2. Symmetrically, the allocation at firm 2 that describes it's mechanism must depend on the mechanism or allocation available with firm 1. Thus it would seem that as firm 1 varies it's mechanism, the allocation describing firm 2 must change, which will generally change the type of any buyer who selects firm 1. A similar problem arises if one describes mechanisms by their expected payoff functions, because the expected payoff at firm 1 depends in general on the mechanism employed at firm 2. However, this interdependence can be overcome by an appropriate (re)definition of 'expected payoff function' and this is one of our methodological contributions.

Finally, we mention that Vassilakis [20] shares with us the objective of endogenizing features of strategic situations that are typically taken to be exogenous. Given an initial game, with specified rules such as strategy sets for each player, he defines an induced *universal game* "whose strategy spaces include the strategy spaces of the original game and, in addition, all possible proposals to modify the rules, including the rules of the universal game" (p.2). In the mechanism design context, the universal game is intended to endogenize the identity of the agents that can propose mechanisms, the degree of collusion between participants in a mechanism (co-ordinated lying) and other features that in our analysis, as in the received literature, are taken to be exogenous. While we are sympathetic towards such an ambitious goal, we feel that our more modest objective, and the added structure that it permits, may prove more successful in producing concrete applied results. Finally, we note that Vassilakis is not explicitly concerned with the questions of primary interest here, namely the modelling of competing mechanisms and an appropriate revelation principle.

A partial guide to the reader may be helpful. In order to minimize the notational and technical hurdles to reaching our statement of the revelation principle, we have divided the latter into two theorems. The first, Theorem 4.1, though incomplete is nevertheless informative and has the added virtue of being easily accessible. A number of examples follow (Section 3). Readers who are sufficiently motivated by this material may then proceed to the completion of the revelation principle (Theorem 5.2) and to our robustness result (Theorem 5.3). Technical details are confined as much as possible to Section 6 and the appendices. The former is intended to outline the essential mathematical structure underlying our revelation principle and to clarify our universal language for describing mechanisms. Proofs are completed in the appendices, particularly in Appendix C. Section 7 concludes with further discussion, some extensions and questions for future research.

2. INDIRECT MECHANISMS

2.1. Primitives

Throughout the paper, where we refer to a set X as a 'space', the intention is that X is a *compact metric* space. Where only a weaker structure is needed, that will be made explicit. Where a measurable structure is needed, the corresponding Borel σ-algebra, denoted $\mathcal{B}(X)$, is used.

For notational simplicity, we deal with the case of two buyers and two sellers or firms. The trading process begins when the sellers simultaneously announce the mechanisms they plan to use. We assume that neither seller observes directly the offer made by the other seller. Buyers, on the other hand, can observe the offers made by both the sellers. After seeing them, each buyer selects one and only one of the mechanisms. Once the buyers have made their choices and these choices have been revealed to the sellers, the sellers' mechanisms are played out with any participating buyers.

The environment that we have in mind is similar to the one that prevails in a simple wage search model. Buyers are responsible for searching out the market, and in our model do so with 0 cost. This search results in each buyer receiving a pair of offers, one from each seller. The offers are not simple prices, but contracts that describe how the buyer will be treated if he accepts the offer. Once the buyer accepts the offer, the contract determines how the final allocation will be decided.

To accommodate the participation choice of buyers, let $P = \{0, 1\}$, where the intention is that $p_i = 1$ if and only if buyer i participates at the firm under consideration.

The primitives for our model are:

$\mathcal{A}_0$: space of 'pure' actions

Ω: valuations space (including the 'usual type' of a buyer)

F: cdf according to which buyers' valuations are drawn (independently)

The independence assumption is made purely to simplify notation; correlated types can be accommodated as long as the distribution of types conditional on a realized own type ω varies continuously with ω.

Sellers may randomize and condition their choice of actions on the participation decisions of buyers. Thus we are led to consider the space $(\Delta(\mathcal{A}_0))^{P^2}$ of participation contingent actions.[1] The value to the seller of any participation contingent plan a_c depends on the participation probabilities of each buyer. In other words, the seller is concerned with the 'full' action (a_c, π, π'), consisting of the contingent action and the probabilities with which each buyer participates in the seller's mechanism. Thus we are led to the actions space $\mathcal{A} = (\Delta(\mathcal{A}_0))^{P^2} \times [0,1]^2$.

Now we can define:

$u : \mathcal{A} \times \Omega \longrightarrow [0,1]$, buyers' utility functions

We assume that buyers who do not participate in either mechanism get 0 utility and that there is an action $\underline{a} \in \mathcal{A}$ such that $u(\underline{a}, \cdot) = 0$. In several examples below, the action $\underline{a}$ corresponds to 'no trade', implying the default utility level of 0 regardless of valuation. For example, a seller might choose a price that he knows no one could afford to pay. Notice that buyers always do at least as well by participating in one of the mechanisms as they would by staying out of the process.

Assume that each buyer is an expected utility maximizer with continuous vNM utility function $u^0 : \mathcal{A}_0 \times P^2 \times \Omega \longrightarrow [0,1]$, that is, the value of u^0 depends on the firm's chosen pure action, on participation decisions and on the buyer's valuation. To complete the description of utility, adopt the following additional notation. For any contingent action $a_c \in (\Delta(\mathcal{A}_0))^{P^2}$, denote by $a_c(p,p')$ the marginal on

[1] $\Delta(\mathcal{A}_0)$ denotes the space of Borel probability measures on $\mathcal{A}_0$, endowed with the standard topology of weak convergence, which makes it compact metric.

the (p, p')-component space $\mathcal{A}_0$. Then the utility of the action $a = (a_c, \pi, \pi') \in \mathcal{A}$ for a buyer with valuation ω is given by

$$u(a, \omega) = \left[\pi' \int_{\mathcal{A}_0} u^0(\cdot, 1, 1, \omega)\, da_c(1,1)(\cdot) + (1 - \pi') \int_{\mathcal{A}_0} u^0(\cdot, 1, 0, \omega)\, da_c(1,0)(\cdot)\right], \tag{2.1}$$

Turn finally to sellers' payoffs. Each seller has a vNM index $v : \mathcal{A} \times \Omega^2 \longrightarrow [0,1]$, where the dependence of $v(a, \omega, \omega')$ on the valuations of the two buyers allows us to interpret each contingent action as an *option* to trade at a specified price. The value of such an option depends on whether or not the buyer decides to exercise the option and this depends on his valuation. Assume also that $v(\underline{a}, \cdot) = 0$.

2.2. Standard Model of Competition

The usual way to model behavior in the above setting proceeds by specifying a set Γ of feasible indirect mechanisms from which sellers may choose. We outline this modelling approach here.

To define indirect mechanisms, fix a space of messages C that is used by both firms. The message space C is perfectly general and in particular need not contain any communication about the mechanisms of competing firms. If such communication is permitted by C, then the nature of the description of the other mechanisms is unspecified. An indirect mechanism γ assigns an action to each of the messages that might be communicated by buyers, that is, γ is a measurable map from C^2 into $\mathcal{A}$. Write $\gamma = (\gamma_c, \gamma_{\pi_1}, \gamma_{\pi_2})$, where

$$\gamma_c : C^2 \longrightarrow (\Delta(\mathcal{A}_0))^{P^2} \text{ and } (\gamma_{\pi_1}, \gamma_{\pi_2}) : C^2 \longrightarrow [0,1]^2. \tag{2.2}$$

Thus, $\gamma_c(\cdot)$ describes the contingent action and $(\gamma_{\pi_1}(\cdot), \gamma_{\pi_2}(\cdot))$ describes the 'rest' of the action prescribed by the mechanism $\gamma(\cdot)$.[2]

This formulation admits two possible interpretations with respect to the timing of communication. The set of feasible mechanisms Γ may or may not allow the outcome that the seller specifies when *only* buyer 1 chooses his mechanism to depend on the message sent by buyer 2. Such dependence occurs in models

[2]It may seem unnatural to have the range of γ lie in $\mathcal{A}$ instead of in $(\Delta(\mathcal{A}_0))^{P^2}$, since an element in $(\Delta(\mathcal{A}_0))^{P^2}$ already specifies a set of participation contingent actions. Our chosen formulation anticipates the discussion of direct mechanisms below where we wish to interpret a mechanism as a recommendation about how to choose among the sellers, and a specification of how buyers will be treated when they do choose the seller. In that setting, the component $\gamma_r(\cdot)$ will represent recommended participation probabilities.

where buyers communicate with sellers before committing themselves to one of the mechanisms. The alternative and common assumption ([8], for example) is that buyers communicate after committing themselves. This assumption can be accommodated within our formalism by restricting mechanisms so that the action prescribed when only buyer i participates is independent of buyer j's message. (See Section 7.3 for further discussion.) The examples that we give below all fall into this latter category.

Denote by Γ the set of feasible indirect mechanisms, endowed with some topology. Frequently, we will assume that Γ is compact metric.

Turn to behavior. A *communication strategy* $\widetilde{c}$ is a measurable mapping from $\Omega \times \Gamma^2$ into C, with the interpretation that $\widetilde{c}(\omega, \gamma, \gamma')$ is the message sent to the firm using γ by a buyer of valuation ω when the other firm is using γ'. Similarly, a *participation strategy* is a measurable function $\widetilde{\pi} : \Omega \times \Gamma^2 \longrightarrow [0, 1]$, where $\widetilde{\pi}(\omega, \gamma, \gamma')$ is the probability of participating only at the firm using γ by a buyer of valuation ω when the other firm is using γ'.

Say that the strategy pair $(\widetilde{c}, \widetilde{\pi})$ is a *continuation equilibrium* if no buyer has any incentive to deviate from either the reporting strategy $\tilde{c}$ or the selection strategy $\widetilde{\pi}$, for any of his valuations and for any pair of mechanisms offered by the two sellers. We assume the existence of continuation equilibria. When we want to emphasize the underlying set of indirect mechanisms Γ, we refer to $(\widetilde{c}, \widetilde{\pi})$ as a continuation equilibrium *relative to* Γ or we refer to the triple $(\Gamma, \widetilde{c}, \widetilde{\pi})$ as a continuation equilibrium. When we wish to emphasize a particular pair of mechanisms, we will refer to $(\tilde{c}(\cdot, \gamma, \gamma'), \widetilde{\pi}(\cdot, \gamma, \gamma'))$ as a continuation equilibrium relative to (γ, γ').

The key to the standard (one firm) revelation principle, is that composing a mechanism with buyers' strategies yields a mapping from pairs of valuations into actions, or in other words, a 'direct mechanism'. A corresponding composition plays an important role in the present setting. To be precise, given γ, each communication and participation strategy $(\tilde{c}, \widetilde{\pi})$ induces the mapping $m_\gamma : \Omega^2 \times \Gamma^2 \longrightarrow \mathcal{A}$, where

$$m_\gamma(\omega, \omega', \gamma', \gamma'') = (\gamma_c(\widetilde{c}(\omega, \gamma, \gamma'), \widetilde{c}(\omega', \gamma, \gamma'')), \widetilde{\pi}(\omega, \gamma, \gamma'), \widetilde{\pi}(\omega', \gamma, \gamma'')). \quad (2.3)$$

Each m_γ can be thought of as a direct mechanism (see Section 5). The expression $m_\gamma(\omega, \omega', \gamma', \gamma'')$ describes the action forthcoming at the firm employing γ, in the given continuation equilibrium, if the ω-valuation buyer acts as though the other firm is employing γ' and the ω'-valuation buyer acts as though the other firm is employing γ''. In equilibrium, $\gamma' = \gamma''$ and both equal the mechanism actually chosen by the other firm, but allowing $\gamma' \neq \gamma''$ in principle will permit us later to express the appropriate incentive compatibility restrictions on direct mechanisms. The dependence on the other firm's mechanism of the action chosen differentiates our setting from the more familiar single firm setting, where valuations alone matter.

The preceding definition also simplifies the description of seller behavior. Suppose that the competing firm chooses the randomization $\delta' \in \Delta(\Gamma)$ and that buyer behavior is described by the strategy pair $(\widetilde{c}, \widetilde{\pi})$. Then the seller who chooses the randomization δ receives the payoff

$$V(\delta; \delta', \widetilde{c}, \widetilde{\pi}) = \int v\left(m_{\gamma}(\omega, \omega', \gamma', \gamma'), \omega, \omega'\right) \, dF\left(\omega\right) \, dF\left(\omega'\right) \, d\delta'(\gamma') \, d\delta(\gamma). \quad (2.4)$$

Consequently, say that the triple $(\widetilde{c}, \widetilde{\pi}, \delta^*)$ is a *(symmetric) equilibrium* relative to Γ, or simply that $(\Gamma, \widetilde{c}, \widetilde{\pi}, \delta^*)$ is an equilibrium, if: $(\widetilde{c}, \widetilde{\pi})$ is a continuation equilibrium and

$$\delta^* \in \arg \max_{\delta \in \Delta(\Gamma)} V(\delta; \delta^*, \widetilde{c}, \widetilde{\pi}).$$

A serious difficulty with using any such equilibrium to model behavior is that the set of equilibria depends on the set of feasible mechanisms Γ, including C, and these are usually specified for convenience and tractability rather than for compelling economic reasons. As described in the introduction, this ad hocery is a major motivation for our paper. In Sections 4 and 5, we describe a universal class of mechanisms that provides a way to avoid such ad hocery. First, we present some illustrative and motivating examples.

3. EXAMPLES

A simple competitive environment is one in which each of the two sellers has a single indivisible unit of output to sell and each of the buyers wishes to acquire exactly one unit if the price is low enough.

In the usual modelling approach, a feasible set Γ of indirect mechanisms is specified and an associated equilibrium is used to describe behaviour. Here (and in Appendix A) we describe three alternative specifications for Γ . Though each is in some sense natural or at least familiar, each specification is ad hoc and the predictions that they deliver are therefore suspect. On the other hand, we will show later that each specification for Γ can be identified with a subset of the universal class of indirect mechanisms delivered by our Theorems 4.1 and 5.2, illustrating the unifying power of our framework. Therefore, an equilibrium delivered by any of these models Γ is a sensible basis for explanantion or prediction if it is robust to deviations by sellers to mechanisms in the universal class. It is beyond the scope of this paper to determine such robustness or to derive directly an equilibrium relative to the universal class of indirect mechanisms. Our contribution is to provide a framework within which the question of robustness can be meaningfully posed.

Two forms of price competition are described next. See Appendix A for a model where sellers compete in auctions. The environment for all three cases is as follows: Buyers' valuations are independently drawn from the interval $\Omega = [0,1]$ using the continuous probability distribution function F. The set of pure actions $\mathcal{A}_0 = [0,1] \times \Delta^2$, where $\Delta^2 = \left\{\mu \in \mathbf{R}^2_+ : \mu_1 + \mu_2 \leq 1\right\}$. The generic action $a_0 = (q, \mu)$ indicates that the seller chooses buyer i with probability μ_i and offers him the option to trade at price q. A buyer with valuation ω who trades at a price p receives utility $\omega - p$ while the seller in the same situation receives payoff p. More formally, if buyer 1 has valuation $\omega \in [0,1]$, then his utility function is

$$u^0(a_0, p, p', \omega) = \begin{cases} \mu_1(\omega - q) & \text{if } p = 1 \text{ and } \omega \geq q \\ 0 & \text{otherwise,} \end{cases}$$

where $a_0 = (q, \mu)$. Buyer 2's utility function is defined symmetrically. Sellers' utility functions are described formally below.

3.1. Simple Price Competition

The standard approach to this problem in the search literature is to allow each seller to set a take it or leave it price. Buyers search out these price offers and then choose one and only one seller as a potential trading partner. If both buyers select the same seller, the seller can provide the unit only to one of them, so one of the buyers will be rationed, and, by assumption will be unable to trade at all[3]. We model this by assuming that each seller announces a price and promises to choose each of the buyers who select him with equal probability. The buyer who is chosen is then given the *option* to trade at the seller's announced price.

[3]Of course, this rationed buyer should have the opportunity to go back to the other seller. It seems reasonable that this return will be costly, perhaps in the sense that trade is delayed. For convenience we simply assume that the cost of returning to the other seller is prohibitively high.

The rules of the fixed price mechanism do not generally permit the seller to randomize over the price that he offers, though he needs to randomize his choice over the buyers who select him in a specific way. Thus the only relevant actions, elements of $\mathcal{A}$, are tuples of the form

$$a = \left\{ \left(q, 1, 0\right), \left(q, 0, 1\right), \left(q, \ .5, \ .5\right), \pi, \pi' \right\}. \tag{3.1}$$

(To clarify, $(q, \ .5, \ .5)$ indicates that if both buyers participate $((p, p') = (1, 1))$, then each buyer receives, with probability $1/2$, the option to buy the good at price q. The other two triples describe the pure action undertaken if $(p, p') = (0, 1)$ or $(1, 0)$.) This gives

$$u\left(a, \omega\right) = \left(\pi' \max\left[\omega - q, 0\right]/2 \ + \ \left(1 - \pi'\right)\max\left[\omega - q, 0\right]\right) \tag{3.2}$$

and for the seller

$$v\left(a, \omega, \omega'\right) = \begin{cases} \left(1 - \left(1 - \pi\right)\left(1 - \pi'\right)\right) q & \text{if } \omega, \omega' \geq q \\ \pi q & \text{if } \omega \geq q, \, \omega' < q \\ \pi' q & \text{if } \omega < q, \, \omega' \geq q. \end{cases}$$

The fixed price mechanism requires that the action that the seller chooses be independent of buyers' private information. (Formally, the message space C may be taken to be a singleton.) This means that the set of indirect mechanisms is $\Gamma = \{\gamma_q : q \in [0, 1]\}$, where $\gamma_q \in \mathcal{A}$ is defined by

$$\gamma_q = \left(\left(q, 1, 0\right), \left(q, 0, 1\right), \left(q, \ .5, \ .5\right), \bar{\pi}, \bar{\pi}\right).$$

Because firms make no attempt to influence buyers' choices, we assign a recommendation $\bar{\pi}$ arbitrarily. Henceforth, denote γ_q simply by q.

Focus on a continuation equilibrium. Because there is nothing for buyers to report to sellers, the continuation equilibrium consists only of a choice strategy. It is straightforward to show that the continuation equilibrium choice strategy is characterized by a cutoff valuation $\omega^*(q, q')$ with the property that buyers whose valuations are below $\omega^*(q, q')$ choose the lower priced seller, while buyers with valuations in the other interval choose the higher priced seller. That is,

$$\widetilde{\pi}\left(\omega, q, q'\right) = \begin{cases} 1 & \text{if } \omega \leq \omega^*\left(q, q'\right) \ \text{ and } q \leq q' \\ 1 & \text{if } \omega > \omega^*\left(q, q'\right) \text{ and } q > q' \\ 0 & \text{otherwise.} \end{cases} \tag{3.3}$$

It is readily verified that the cutoff valuation $\omega^*(q,q')$ is given by the solution ω to

$$\frac{\omega-\min[q,q']}{\omega-\max[q,q']}=\frac{1+F(\omega)}{2-F(\omega)+F(\min[q,q'])}, \tag{3.4}$$

whenever this equation has a solution in the interval $[\max[q,q'],1]$. Otherwise, the cutoff is equal to 1 and all buyers choose the lower priced firm. As a result, $\omega^*(\cdot)$ is a continuous function.

The direct mechanism m_q corresponding to this equilibrium as in (2.3), (and imposing $q'=q''$), is given by $m_q(\omega,\omega',q',q')=$

$$((q,1,0),(q,0,1),(q,.5,.5),\widetilde{\pi}(\omega,q,q'),\widetilde{\pi}(\omega',q,q')). \tag{3.5}$$

The seller's payoff for any pair of mechanisms $q\leq q'$ is given by

$$V(q;q',\widetilde{c},\widetilde{\pi})=\int v\left(m_q(\omega,\omega',q',q'),\omega,\omega'\right)\,dF(\omega)\,dF(\omega')$$

$$=q-q\left(1-F\left(\omega^*(q,q')\right)\right)^2.$$

From (3.4), this payoff function is continuous.[4] Because Γ is compact, there exists an equilibrium in mixed strategies.

3.2. Price Matching

The model of simple price competition is rich enough to illustrate many of the properties that we discuss below. Yet it is obviously inadequate as a description of equilibrium mechanisms. For example, it is well known that sellers can profitably deviate by holding auctions. Alternatively, they may wish to deviate by occasionally matching the other seller's price. We focus on such price matching mechanisms here.

The class of feasible indirect mechanisms is described as follows. Sellers initially announce a price. Buyers are asked to tell the seller what price the other seller has offered by naming a price from the message space $C=[0,1]$. There are then two possibilities. If only one of the buyers selects the seller, his report about the other firm is ignored and that buyer is offered the option to trade at

[4]Though there is price competition here, the seller's payoff function is not discontinuous as in the usual Bertrand problem. A small cut in price increases continuously the probability with which each buyer selects the seller. This is because of the possibility of rationing when both buyers choose the same seller.

the price that the seller announced. If two buyers select the seller, the seller takes the maximum of the two reported prices. If this maximum exceeds the price that the seller has offered, the seller ignores the reports. He chooses one of the buyers randomly and offers him the opportunity to trade at the price that the seller originally announced. If the maximum of the prices reported by the two buyers is below the price that the seller has announced, the seller picks one of the buyers at random and offers him the option to trade at this maximum price. The only real option that the seller has in this class of mechanisms is the price that he sets. However the use of buyers' messages makes it possible to vary the price charged according to the price offered by the other firm.[5]

Again, all relevant actions have the form in (3.1). However, the action that the firm chooses is allowed to depend on the messages that the seller receives from the buyers. Formally $\Gamma = \{\gamma_q : q \in [0,1]\}$, where $\gamma_q : C^2 \to \mathcal{A}$ is defined by

$$\gamma_q(c, c') = ((q, 1, 0), (q, 0, 1), (\min[q, \max[c, c']], .5, .5), \bar{\pi}, \bar{\pi}).$$

As in the case of simple price competition, we assign a recommendation $\bar{\pi}$ arbitrarily and refer to γ_q simply as q.

Turn to continuation equilibria. There will evidently be many continuation equilibria. We focus on the one in which buyers report the other firm's price truthfully, that is,

$$\tilde{c}(\omega, q, q') = q'.$$

It is straightforward to show that this is an equilibrium: The buyer's report does not affect the price he pays when he is the only buyer to visit a seller, thus the best report that the buyer can make is the one that is best when the other buyer selects the same seller and reports truthfully. In that case, he cannot affect the price paid if he reports a lower price for the other firm than the true one. If he reports a price above the true price he will either have no effect on the price or raise it.

For determination of the choice strategy, the problem is similar in structure to that with simple price competition. The price that the seller announces is the price at which trade will occur when there is only one buyer at the firm. Since the high priced seller ends up trading at the lower priced seller's price in many situations, there is less advantage in this situation to being the lower priced firm.

[5]This example differs from the model of Salop [17] in two ways. First, Salop assumes that firms observe the prices set by other firms. Here sellers can learn about the other seller's price only by asking buyers about it. Secondly, the sellers in this example face a capacity constraint that Salop's firms do not.

The continuation equilibrium choice strategy is again characterized by a cutoff valuation $\omega^*(q,q')$ as in (3.3). Here, however, the cutoff valuation must satisfy

$$\frac{1}{2F(\omega)} = \frac{(\omega - \max[q,q'])}{(\omega - \min[q,q'])}, \tag{3.6}$$

when this equation has a solution. Otherwise the cutoff is equal to 1 and all buyers select the lower priced seller.

The direct mechanism m_q corresponding to this equilibrium as in (2.3), (and imposing $q' = q''$), is given by $m_q(\omega,\omega',q',q') =$

$$((q,1,0),(q,0,1),(\min[q,q'],.5,.5),\widetilde{\pi}(\omega,q,q'),\widetilde{\pi}(\omega',q,q')). \tag{3.7}$$

The seller's payoff for any pair of mechanisms $q \le q'$ is given by $V(q;q',\widetilde{c},\widetilde{\pi}) =$

$$\begin{cases} \left[1-(1-F(\omega^*(q,q')))^2\right] q & \text{if } q \le q' \\ 2(1-F(\omega^*(q,q')))F(\omega^*(q,q'))q + (1-F(\omega^*(q,q')))^2 q' & \text{otherwise.} \end{cases}$$

From (3.6), this payoff function is continuous.[6] Because Γ is compact, there exists an equilibrium in mixed strategies.

4. A REVELATION PRINCIPLE

4.1. Additional Assumptions

We have seen that the actions forthcoming in any continuation equilibrium can be represented by functions of the form $m_\gamma(\cdot)$, defined in (2.3), that assign actions directly to reports about valuations and the mechanism of the other firm. As mentioned, a report contingent action may be interpreted as including a contingent recommendation about the way the buyers should choose among the sellers. Thus, each such m_γ can be interpreted informally as a direct mechanism in which the firm makes a recommendation to buyers about how they should choose among the two competing sellers, and specifies how the buyers will be treated as functions of their reports and subsequent participation decisions.

Focus on a continuation equilibrium $(\Gamma,\widetilde{c},\widetilde{\pi})$. We need to restrict such equilibria in two ways. To express these restrictions, introduce the payoff functions

[6] An 'explanation' similar to that provided for model of simple price competition applies.

induced by mechanisms. To be precise, denote the expected utility of a buyer facing γ by

$$U(\omega, \gamma'; \gamma) = \int u(m_\gamma(\omega, \omega', \gamma', \gamma'), \omega)\, dF(\omega'), \tag{4.1}$$

if the buyer has valuation ω and the other firm is using the mechanism γ'.

We require first that the continuation equilibrium satisfy the following two-faceted continuity property. For any space S, $\mathcal{U}(S)$ denotes the set of upper-semicontinuous (usc) functions from S into $[0,1]$ and endowed with the topology described in Section 6.1 below.

Definition: Say that the *continuation equilibrium* $(\Gamma, \widetilde{c}, \widetilde{\pi})$ *is payoff upper- semi-continuous* if: (i) $U(\cdot; \gamma)$ is usc on $\Omega \times \Gamma$; and (ii) the mapping $\gamma \longmapsto U(\cdot; \gamma) \in \mathcal{U}(\Omega \times \Gamma)$ is continuous.

Upper-semicontinuity of $U(\cdot; \gamma)$ in valuation alone is implied by a continuation equilibrium (this is well known - [21] for example). It follows that payoff usc of the continuation equilibrium is innocuous if Γ is finite.[7] But it is restrictive more generally. On the other hand, it is satisfied by all our examples (see Appendix A), showing that payoff usc is not overly restrictive. Another justification for this assumption is pragmatic - it works. Upper semi-continuity plays a large role in the technical analysis.

The second restriction on continuation equilibria (called non-redundancy) is much more difficult to explain. We provide a formal (and possibly impenetrable) definition of the property here and defer interpretation and justification until after Theorem 4.1, which shows what the assumption of non-redundancy delivers. At this point, we add only that non-redundancy is satisfied by each of the continuation equilibria described for the two forms of price competition and the model of competing auctions.

The formal definition follows: Given a continuation equilibrium $(\Gamma, \widetilde{c}, \widetilde{\pi})$, define a sequence $\{\Sigma_n\}$ of σ-algebras on Γ, each contained in the Borel σ-algebra. Let $\Sigma_0 = \{\emptyset, \Gamma\}$, $\Sigma_1 = \sigma$-algebra generated by the mappings $\gamma \longmapsto \sup\{U(\omega, \gamma'; \gamma) : (\omega, \gamma') \in E\}$, where E varies over $\mathcal{B}(\Omega) \times \Sigma_0$, and $\Sigma_{n+1} = \sigma$-algebra generated by the mappings $\gamma \longmapsto \sup\{U(\omega, \gamma'; \gamma) : (\omega, \gamma') \in E\}$, where E varies over $\mathcal{B}(\Omega) \times \Sigma_n$. Observe that $\Sigma_n \nearrow_n$ and that if $\Sigma_n = \Sigma_{n+1}$ for some n, then $\Sigma_n = \Sigma_k$ for all

[7]This case is not without economic interest. It is common in the literature to assume that sellers have recourse to only one or two alternative mechanisms, for example, sellers might choose to bargain or hold an auction.

$k > n$. Say that $(\Gamma, \tilde{c}, \tilde{\pi})$ is *non-redundant* (NR) if any pair of distinct points in Γ can be separated by some Σ_n.[8]

The statement and interpretation of non-redundancy are simpler when Ω is finite (or countable). In that case, the σ-algebras defined above are unchanged if, for all n, E is restricted to vary only over $\{\{\omega\} \times \Sigma_n : \omega \in \Omega\}$. The complicating need to rely on nonsingleton subsets of Ω in the infinite case appears to be a purely "technical matter".

4.2. The Main Result

Our main result is presented here. It provides the sought after universal language T for describing mechanisms. First, we explain some notation and terminology used in the theorem.

The set $\mathcal{A}^{\Omega^2 \times T^2}$ denotes all measurable maps $m : \Omega^2 \times T^2 \longrightarrow \mathcal{A}$. Each such m can be viewed as a direct mechanism employing message space $\Omega \times T$ (for each buyer), that assigns action $m(\omega, \omega', t', t'')$ directly to reports (ω, t') and (ω', t'') by the two buyers. Associated with T is a map $\psi : T \longrightarrow \mathcal{A}^{\Omega^2 \times T^2}$; interpret $\psi(t)$ as the direct mechanism that is described by $t \in T$. Given such a pair (T, ψ), we can also view the image $\psi(T)$ as a set of indirect mechanisms using message space $C = \Omega \times T$. This identification gives meaning to the theorem's reference to $(\psi(T), c^*, \pi^*)$, a continuation equilibrium relative to $\psi(T)$.

The theorem also refers to the mapping $\Psi : T \longrightarrow \mathcal{U}(\Omega \times T)$, that takes any t into $U(\cdot; \psi(t))$, where, similar to (4.1),

$$U(\omega, t'; \psi(t)) = \int u(\psi(t)(\omega, \omega', t', t'), \omega) \, dF(\omega').$$

Denote by $\theta_{(\Gamma, \tilde{c}, \tilde{\pi})}$, or simply by θ, the map from Γ into $\mathcal{A}^{\Omega^2 \times \Gamma^2}$ taking γ into m_γ (see (2.3)). Non-redundancy implies that this map is one-to-one.

We can now state our main result (see Theorem 5.2 for an extension):[9]

Theorem 4.1. *There exist a separable metric space T and a one-to-one map $\psi : T \longrightarrow \mathcal{A}^{\Omega^2 \times T^2}$ such that for any payoff usc and non-redundant continuation equilibrium $(\Gamma, \tilde{c}, \tilde{\pi})$, with Γ compact metric, there exist:*

---

[8] Our definition of non-redundancy is adapted from that in [6]. The definition in [9] is analogous.

[9] We say that $e : \Gamma \longrightarrow T$ is an embedding if it is continuous and one-to-one. When Γ is compact Hausdorff, this is equivalent to e being a homeomorphism into T.

(a) An embedding $e : \Gamma \longrightarrow T$ and a mapping $\xi : \mathcal{A}^{\Omega^2 \times \Gamma^2} \longrightarrow \mathcal{A}^{\Omega^2 \times T^2}$ that render the following diagram commutative:

$$\begin{array}{lll} \Gamma & \longrightarrow^{e} & T \\ \downarrow_{\theta} & & \downarrow_{\psi} \\ \mathcal{A}^{\Omega^2 \times \Gamma^2} & \longrightarrow_{\xi} & \mathcal{A}^{\Omega^2 \times T^2} \end{array}$$

(b) A payoff usc and non-redundant continuation equilibrium $(\psi(T), c^*, \pi^*)$ satisfying

$$m(\omega, \omega', t', t') = m\left(c^*(\omega, t, t'), c^*(\omega', t, t')\right), \text{ and}$$

$$m'(\omega, \omega', t, t) = m'\left(c^*(\omega, t', t), c^*(\omega', t', t)\right),$$

for all $m = \psi(t)$ and $m' = \psi(t')$ in $\psi(T)$ and for all $(\omega, \omega') \in \Omega^2$.

(c) Moreover, the continuation equilibrium $(\psi(T), c^*, \pi^*)$ satisfies

$$\gamma\left(\widetilde{c}(\omega, \gamma, \gamma'), \widetilde{c}(\omega', \gamma, \gamma')\right) = m\left(c^*(\omega, e(\gamma), e(\gamma')), c^*(\omega', e(\gamma), e(\gamma'))\right),$$

for all $(\omega, \omega', \gamma, \gamma') \in \Omega^2 \times \Gamma^2$ and $m = \psi(e(\gamma))$.

(d) Finally, the mapping $\Psi : T \longrightarrow \mathcal{U}(\Omega \times T)$ is an embedding, where

$$\Psi(t)(\cdot) \equiv U(\cdot; \psi(t)). \tag{4.2}$$

This theorem constitutes the crux of our contribution and thus warrants thorough interpretation. First, we point out that in order to present the simplest possible statement that communicates the essence of our result, we have deleted reference to continuity properties of $\theta_{(\Gamma, \widetilde{c}, \widetilde{\pi})}$, ψ and ξ and to further details regarding $\psi(T)$. Such details are provided in Section 5.

The theorem provides a pair (T, ψ), where each $t \in T$ is the type or description of the direct mechanism $\psi(t)$. The pair is universal in the sense of being independent of the particular continuation equilibrium $(\Gamma, \widetilde{c}, \widetilde{\pi})$. The types space T is universal also in that any set of indirect mechanisms Γ satisfying the specified assumptions can be viewed as a subspace of T through the embedding e. As noted earlier, elements of $\mathcal{A}^{\Omega^2 \times \Gamma^2}$ can be interpreted as direct mechanisms that use reports about valuations and the mechanism, an element of Γ, of the other firm. The map ξ translates any such direct mechanism into one in which the report about the other firm is expressed in terms of the universal language represented by T. In particular, this is true for the direct mechanism $\theta_{(\Gamma, \widetilde{c}, \widetilde{\pi})}(\gamma) = m_\gamma$ that is derived from the continuation equilibrium $(\Gamma, \widetilde{c}, \widetilde{\pi})$ as in (2.3). Thus, the actions at a given firm implied by any (payoff usc and NR) continuation equilibrium can be viewed as functions of the two valuations and the type $t' \in T$ of the other firm. In other words, the set $\mathcal{A}^{\Omega^2 \times T^2}$ of direct mechanisms is large enough to replicate the outcomes of any suitable continuation equilibrium. Indeed, by commutativity of the diagram, the subset $\psi(T)$ suffices.

The above corresponds to the observation in the single firm setting that direct mechanisms based on valuations are enough to replicate the actions implied in a buyer equilibrium associated with any indirect mechanism. In that standard setting, any direct mechanism that is incentive compatible can be implemented, implying that the set of incentive compatible direct mechanisms is a universal class for modelling the mechanism choice problem of a single mechanism designer. This constitutes the received revelation principle. Parts (b) and (c) of the theorem establish that in a similar sense $\psi(T)$ is a universal class of indirect mechanisms for modelling mechanism choice in the setting of multiple competing mechanism designers. This constitutes our version of the revelation principle. In particular, (b) states that the actions prescribed by mechanisms in $\psi(T)$ can be *implemented* or achieved in the continuation equilibrium provided. Part (c) is an immediate implication of (a) and (b). It states explicitly that the continuation equilibrium outcomes implied by any given $(\Gamma, \tilde{c}, \tilde{\pi})$ are *replicated* by $(\psi(T), c^*, \pi^*)$.

Applications of the single firm revelation principle rely on the explicit characterization of the universal class of mechanisms - mechanisms in $\mathcal{A}^{\Omega^2}$ that satisfy incentive compatibility. Fruitful applications of our theorem presumably would require further information about $\psi(T)$ and *a fortiori* about T. Fortunately, the proof of the theorem is constructive; each of T, ψ and (c^*, π^*) is constructed explicitly. In particular, the constraints on mechanisms in $\mathcal{A}^{\Omega^2 \times T^2}$ that define the subset $\psi(T)$ are described, providing counterparts of the familiar incentive compatibility conditions (see also Section 5). Admittedly, the constructions and the constructs that they deliver are much more complicated than in the single firm setting. However, this added complexity seems to us to be a reflection of the added complexity of a competitive environment and our desire to at least reduce the degree of ad hocery in modelling. (Our restriction to payoff usc and non-redundant continuation equilibria is ad hoc, but seems considerably less so than a restriction to a particular set of indirect mechanisms such as take-it-or leave-it prices or specified auctions.)

Part (d) provides some information about the nature of types and the sense in which they describe mechanisms. Each type t can be identified with the expected payoff function $U(\cdot; \psi(t))$ corresponding to the direct mechanism $\psi(t)$ that it describes. In that sense, descriptions in the universal language are in the form of payoff functions.

4.3. Non-Redundancy

The intuition surrounding non-redundancy begins with the observation that any direct mechanism m_γ employs as an input the description of the other firm. This description γ' is expressed in the 'language' of Γ, which is adequate for reporting on mechanisms lying in Γ, but not more generally. To overcome the ad hocery in the specification of Γ, at the very least we need a language that permits descriptions that make no use of the specific Γ being considered. This universal language for describing mechanisms should also permit descriptions of mechanisms lying outside of Γ. Thus it must involve a rich message space T. In particular, T should be rich enough so that any pair of distinct mechanisms (in a suitably large class) may be differentiated by descriptions in T. Non-redundancy is designed to achieve such invertibility.

To motivate and clarify non-redundancy, we offer a heuristic outline. Given two mechanisms γ_1 and γ_2 in Γ, we seek a way to distinguish between them that is independent of the language used in Γ. One natural approach is to consider the way the payoffs or actions associated with these two mechanisms vary across Ω within the given continuation equilibrium $(\widetilde{c}, \widetilde{\pi})$. This amounts to asking whether we can identify with each γ_i a 'marginal' mechanism that employs the message space Ω^2; or, because mechanisms matter primarily through their payoff functions, the task is to associate with each $U(\cdot, \cdot; \gamma_i)$, a function on $\Omega \times \Gamma$, a 'marginal' function that is defined on Ω.

This is somewhat analogous to associating with each joint probability measure a suitable marginal measure. Unlike the case for probability measures, however, there is no compelling and uncontentious notion of marginal for this setting. Our choice is to define the Ω-marginal to be $\sup_{\gamma' \in \Gamma} U(\cdot, \gamma'; \gamma_i)$. Though other definitions seem as plausible, (using the inf is one example), it remains to be seen whether any of them leads to successful completion of analyses paralleling ours.[10]

To continue with the heuristic outline, distinguish between γ_1 and γ_2 if

$$\sup_{\gamma' \in \Gamma} U(\cdot, \gamma'; \gamma_1) \neq \sup_{\gamma' \in \Gamma} U(\cdot, \gamma'; \gamma_2).$$

[10] Another possibility is to use the correspondence $\omega \longmapsto U(\omega, \Gamma; \gamma')$ to represent the dependence of expected payoffs on valuation. That would require a substantial modification of the technical analysis to deal with hierarchies of correspondences.

This simple criteria will generally fail to distinguish among all mechanisms in Γ. Even so, Γ can be partitioned into equivalence classes such that all mechanisms within an equivalence class are indistinguishable according to this sup criterion.

Now refine the criterion and calculate the supremum for γ_1 and γ_2 on each one of these equivalence classes instead of on all of Γ. Distinguish between γ_1 and γ_2 if there is some equivalence class on which the sup calculations above are different. This sharper test will define a new and finer partition into equivalence classes. This refinement procedure can be repeated as often as we like. The continuation equilibrium $(\Gamma, \widetilde{c}, \widetilde{\pi})$ is non-redundant (NR) if, for any pair of mechanisms in the class, some finite repetition of this procedure places the two mechanisms into separate equivalence classes.

Non-redundancy is satisfied by *each* of our examples, namely the two forms of price competition and competing auctions (see Appendix A for details). To illustrate, consider the price matching example. From the expression (A.2) for buyers' payoffs, one obtains that

$$\sup_{q'} U\left(\omega, q'; q\right) = U\left(\omega, 0; q\right) =$$

$$F\left(\omega^*\left(q, 0\right)\right) \max\left[\omega - q, 0\right] + \left(1 - F\left(\omega^*\left(q, 0\right)\right)\right) \frac{\omega}{2}.$$

That is, the best that any buyer can do at the first seller occurs when the other seller sets a price of 0. (This is apparent by noting that (3.4) implies that the cutoff valuation is inversely related to the price q' set by the lower priced seller. Thus, as seller 2 cuts his price, the buyer who selects seller 1 will be more likely to be chosen by seller 1 as a trading partner, and will pay a lower price in the event that the other buyer also chooses seller 1.) It is a routine matter (see Lemma A.1) to show that

$$U(\cdot, 0; q_1) = U(\cdot, 0; q_2) \implies q_1 = q_2,$$

proving that the first level σ-algebra Σ_1 is sufficient to separate distinct mechanisms.

Nevertheless, the assumption of non-redundancy is restrictive. One reason it may not hold is that there are two distinct actions (or perhaps the same action labelled in different ways) that yield every buyer type the same expected payoff. For example, if Ω consists of a single valuation, then there may be many randomizations over the set of pure actions that yield the buyer the same expected payoff. Our approach is to think of all these mechanisms as being equivalent.

This might seem odd. For example, suppose that Ω is a singleton representing a single risk averse buyer type who is trying to buy an insurance policy from the risk neutral seller. The set of pure actions is then the set of outcome contingent transfers, and there are clearly many distinct transfer functions that will yield the buyer the same expected utility. In our formal statement of non-redundancy, only the buyers payoffs are taken into account. More generally, it would be natural to include sellers' payoffs when trying to decide when two pure actions are equivalent.

This can be done. Denote by $V(\gamma';\gamma)$ the expected payoff to a seller using γ when the other seller is using γ'. If $V(\cdot)$ is incorporated into the definition of non-redundancy in the natural way described in Section 7, then the resulting notion is weaker and can be violated only if there exist distinct γ_1 and γ_2 satisfying *both*

$$\sup_{\gamma'\in\Gamma} U(\cdot,\gamma';\gamma_1) = \sup_{\gamma'\in\Gamma} U(\cdot,\gamma';\gamma_2) \quad \textit{and} \tag{4.3}$$

$$\sup_{\gamma'\in\Gamma} V(\gamma';\gamma_1) = \sup_{\gamma'\in\Gamma} V(\gamma';\gamma_2). \tag{4.4}$$

The 'cost' of this generalization is added notational complexity because of the need to differentiate throughout between the payoff functions of buyers and sellers. For this reason, we have chosen to focus on the notationally simpler version and to provide an outline of the generalization (Section 7).

Finally, we present one example of a redundant set of mechanisms, where redundancy is defined using buyers' payoffs only. Suppose, in the price matching example given above, that a seller who is chosen by only a single buyer is required to offer that buyer the opportunity to trade at a price 0, instead of the announced price. When two buyers select the seller, the seller price matches in the manner described earlier. Then it is straightforward to show that $\sup_{\gamma'\in\Gamma} U(\omega,\gamma';\gamma_1) = U(\omega,0;q) = 3\omega/4$, which is independent of q.[11] Thus (4.3) holds for all γ_1 and γ_2. On the other hand, (4.4) fails, implying that the weaker form of non-redundancy that includes sellers' payoffs *is* satisfied. We have been unable to find any interesting examples violating this weaker form. In addition, our revelation principle is readily generalized to accommodate the weaker notion of non-redundancy (see Section 7).

[11] To see this result, note that the price that a buyer will be offered at each seller is the same, either 0 or the lower of the two prices that were offered. Thus buyers are indifferent and choose each seller with probability 1/2 for any valuation.

5. DIRECT MECHANISMS

5.1. A Characterization

Theorem 4.1 proves that the class $\psi(T)$ of indirect mechanisms employing the message space $\Omega \times T$ is sufficiently rich to replicate the actions forthcoming in any (suitable) continuation equilibrium. Thus $\psi(T)$ is robust in the sense that if we could find an equilibrium for the seller's game relative to this class, then we would not have to worry about the possibility that some seller might profitably deviate outside the class. Moreover, because of the particular message space $\Omega \times T$ employed by mechanisms in $\psi(T)$, any such mechanism can be viewed also as a direct mechanism. The corresponding class of mechanisms in the single firm setting is usually described as consisting of all incentive compatible valuations-based direct mechanisms. Here we provide a parallel characterization of the direct mechanisms in $\psi(T)$.

A complication in achieving such a characterization of $\psi(T)$ arises from the possibility that there are actions that yield the same payoffs for all buyer types. The requirement that continuation equilibria be non-redundant means that we need to be able to recover the actions generated by the mechanism by looking at payoffs. In order to do this, we need to adopt some convention for choosing among payoff equivalent actions.

For example, consider a mechanism that specifies that both buyers choose the seller with probability 0 no matter what the other seller offers and no matter what the buyers' types.[12] Then the action that the mechanism specifies for the case when both buyers select the seller is irrelevant. We treat every mechanism of this kind as identical, which means that (in the specification of Γ) we adopt some convention about what the seller offers to do when both buyers choose him. Non-redundancy then implies that when we recover the mechanism from it's description in the universal space T we recover a mechanism that adheres to this convention.

However, if we characterize direct mechanisms by their incentive properties alone, then there will be nothing in the specification to describe this kind of convention. Thus mechanisms that are incentive compatible and otherwise look like they have been derived from direct mechanisms will not lie in $\psi(T)$ unless they adhere to the convention used in constructing it. However, even if this is the case, they will be *payoff equivalent* to mechanisms that do lie in $\psi(T)$, that is,

[12]Such a mechanism may not be relevant economically, but it may well be feasible. It is much like setting a take it or leave it price equal to the highest valuation in the support of F.

to mechanisms that do adhere to our convention. Thus we show that incentive compatibility characterizes $\psi(T)$ up to payoff equivalence.

We saw above that the actions forthcoming in a continuation equilibrium can be represented by functions of the form m_γ, defined in (2.3), that assign actions directly to valuations and descriptions of the competing firm. Because of the definition of actions, each m_γ can be interpreted as a direct mechanism through which the firm makes a recommendation to buyers about how to choose among the competing sellers, and also specifies how buyers will be treated as a function of their types and subsequent participation decisions.[13]

We can express both the initial set Γ of indirect mechanisms and the above mapping $\gamma \longmapsto m_\gamma$ via the pair $(\Gamma, \theta_{(\Gamma,\widetilde{c},\widetilde{\pi})})$, where $\theta_{(\Gamma,\widetilde{c},\widetilde{\pi})} : \Gamma \longrightarrow \mathcal{A}^{\Omega^2 \times \Gamma^2}$. Since each γ is associated with a unique direct mechanism in $\mathcal{A}^{\Omega^2 \times \Gamma^2}$, we could alternatively view Γ as a language used to *describe* direct mechanisms which have as inputs reports from this same language. As Theorem 4.1 shows, this is exactly the role that the universal message space T plays. Each $t \in T$ describes the direct mechanism $\psi(t)$ which takes as an input a report from T about the other firm's mechanism.

More generally, consider pairs of the form (Γ_d, θ), where Γ_d is a (not necessarily compact metric) topological space and $\theta : \Gamma_d \longrightarrow \mathcal{A}^{\Omega^2 \times \Gamma_d^2}$. Here Γ_d is a language to be used to *describe* direct mechanisms which use as inputs, reports in this same language. Because of this circularity, refer to any such pair (Γ_d, θ) as a *self-referential class of direct mechanisms.* We have already provided a concrete example of such a self referential class when we described the set of price matching mechanisms. Each price describes a mechanism that uses the other seller's price as an input.

Direct mechanisms written in any language are of interest only if they possess all the properties delivered by indirect mechanisms. Therefore, we impose on self-referential classes the essential properties of pairs $(\Gamma, \theta_{(\Gamma,\widetilde{c},\widetilde{\pi})})$ that are delivered by indirect mechanisms. To express these, note that by the nature of the space $(\Delta(\mathcal{A}_0))^{P^2} \times [0,1]^2$ of actions, we can describe any direct mechanism $m \in \theta(\Gamma_d) \subset \mathcal{A}^{\Omega^2 \times \Gamma_d^2}$ in the form

$$m(\cdot) = (m_{10}(\cdot), m_{01}(\cdot), m_{11}(\cdot), m_{\pi_1}(\cdot), m_{\pi_2}(\cdot))$$

where $m_{10}(\cdot)$ describes the distribution over pure actions specified when buyer 1 participates and buyer 2 does not, $m_{01}(\cdot)$ and $m_{11}(\cdot)$ have similar meanings and

[13]Observe that even though buyers' valuations are independent, their types are correlated because both buyers observe the same mechanism at the other seller.

$m_{\pi_i}(\cdot)$ describes the probability with which the firm recommends that buyer i choose to participate.

The recommended choice probabilities must satisfy three constraints. First, in order that they represent symmetric continuation equilibria, require that

$$m_{\pi_1}(\cdot) = m_{\pi_2}(\cdot) \equiv m_\pi(\cdot). \tag{5.1}$$

Buyer one does not learn the valuation of buyer 2 before he makes his choice. Therefore, require that

$$m_\pi(\omega, \omega', \gamma', \gamma'') = m_\pi(\omega, \overline{\omega}', \gamma', \overline{\gamma}'') \quad \forall \omega, \omega', \overline{\omega}', \gamma', \gamma'', \overline{\gamma}''; \tag{5.2}$$

Therefore, we can write simply $m_\pi(\omega, \gamma')$.

Finally, since we assume that with probability 1 each buyer selects a seller, m should satisfy: If $m = \theta(\gamma)$, then for each $m' = \theta(\gamma')$ in $\theta(\Gamma)$,

$$m_\pi(\omega, \gamma') = 1 - m'_\pi(\omega, \gamma). \tag{5.3}$$

We refer to the preceding restrictions, (5.1) through (5.3) as *strong measurability constraints.*[14]

Secondly, any mechanism $m \in \theta(\Gamma_d) \subset \mathcal{A}^{\Omega^2 \times \Gamma_d^2}$ should satisfy the following *incentive compatibility constraints*:

(i) for each ω, ω' and ω^r in Ω and for each triple $\gamma, \gamma', \gamma'' \in \Gamma_d$,

$$\int u(m(\omega, \omega', \gamma'', \gamma''), \omega)\, dF(\omega') \geq \int u(m(\omega^r, \omega', \gamma', \gamma''), \omega)\, dF(\omega'); \tag{5.4}$$

If $m = \theta(\gamma)$, then for each $\omega, \omega' \in \Omega$ and $m' = \theta(\gamma')$,

$$m_\pi(\omega, \gamma') \int u(m(\omega, \omega', \gamma', \gamma'), \omega)\, dF(\omega')$$

$$+ [1 - m_\pi(\omega, \gamma')] \int u(m'(\omega, \omega', \gamma, \gamma), \omega)\, dF(\omega') \geq$$

$$\max\left[\int u(m(\omega, \omega', \gamma', \gamma'), \omega)\, dF(\omega'), \int u(m'(\omega, \omega', \gamma, \gamma), \omega)\, dF(\omega')\right] \tag{5.5}$$

[14] The adjective 'strong' is used to avoid confusion with the measurability of each m as a function, which is built into the definition of $\mathcal{A}^{\Omega^2 \times \Gamma_d^2}$.

Constraint (5.4) says that buyers will not have an incentive to lie to either seller about their own type, or the mechanism that has been offered by the other seller, provided that they expect the other buyer not to lie. This constraint is standard. Constraint (5.5) imposes that buyers have no incentive to deviate from the recommended choice probability announced by the seller. More than any other, this constraint reflects the multiplicity of firms; note, in particular, the role played by θ.

Two self explanatory 'technical' conditions follow.

USC $m \in \mathcal{A}^{\Omega^2 \times \Gamma_d^2}$ is *USC* if $U(\cdot; m)$ is usc on $\Omega \times \Gamma_d$, where

$$U(\omega, \gamma'; m) = \int u\left(m(\omega, \omega', \gamma', \gamma'), \omega\right) \, dF(\omega').$$

Compact Support $m \in \mathcal{A}^{\Omega^2 \times \Gamma_d^2}$ has *compact support* if there exists compact $Y \subset \Gamma_d$ such that $U(\cdot, \gamma'; m) = 0$ for $\gamma' \in \Gamma_d \setminus Y$.

Finally, for any (Γ_d, θ), define $\mathcal{M}(\Omega^2 \times \Gamma_d^2)$ as the set of all direct mechanisms $m \in \mathcal{A}^{\Omega^2 \times \Gamma_d^2}$ satisfying strong measurability, incentive compatibility, USC and having compact support. This is our candidate for the subclass of direct mechanisms that can 'implement' any (suitable) continuation equilibrium.

The following is immediate:

Lemma 5.1. *Let $(\Gamma, \widetilde{c}, \widetilde{\pi})$ be a payoff usc continuation equilibrium, with Γ compact metric, and consider the self-referential class $(\Gamma, \theta_{(\Gamma, \widetilde{c}, \widetilde{\pi})})$. Then*

$$\theta_{(\Gamma, \widetilde{c}, \widetilde{\pi})}(\Gamma) \subset \mathcal{M}(\Omega^2 \times \Gamma^2). \tag{5.6}$$

Proof. The measurability and continuity restrictions are immediate. The incentive compatibility restrictions follow in the usual way by observing that a deviation to another type in the direct mechanism is equivalent to using a strategy that is best for another type in the indirect equilibrium. Thus incentive compatibility follows from the definition of equilibrium. ■

To further establish the special role of the restrictions embodied in $\mathcal{M}(\cdot)$, we state the following theorem that extends the earlier revelation principle (Theorem 4.1). A proof is provided in Section 6 and Appendix C.

Theorem 5.2. *Adopt the notation and assumptions of Theorem 4.1 and let (T, ψ) be the self-referential pair provided there. Then*

$$\psi(T) \subset \mathcal{M}(\Omega^2 \times T^2), \text{ and} \tag{5.7}$$

$$cep\,(\psi(T)) = \left\{U(\cdot; m) : \, m \in \mathcal{M}(\Omega^2 \times T^2)\right\}. \tag{5.8}$$

From Theorem 4.1, we know that each type in T describes the direct mechanism $\psi(t) \in \mathcal{A}^{\Omega^2 \times T^2}$. The set inclusion (5.7), paralleling (5.6), states that each such direct mechanism is strongly measurable, incentive compatible, USC and has compact support. Furthermore, (5.8) provides a sense in which these properties characterize $\psi(T)$. The set equality in (5.8) implies that if we identify mechanisms that deliver the same expected payoff functions, then $\psi(T)$ 'equals' $\mathcal{M}(\Omega^2 \times T^2)$. The set $\psi(T)$ is the counterpart for our setting of the class of all incentive compatible valuations-based mechanisms familiar from the single firm setting. Given the types space T, our characterization seems as tractable and potentially useful as the latter has proven to be in the case of a single firm. Admittedly, T is complicated, but it may be constructed by countable recursive methods as described in Section 6.

The formal statement does not include topological details, but we add some here (the proof of the theorem confirms them). We need topologies for $\mathcal{M}(\Omega^2 \times \Gamma^2)$ and $\mathcal{M}(\Omega^2 \times T^2)$, one difference being that Γ is compact metric but T need not be. In the former case, adopt the weak topology induced by $cep : \mathcal{M}(\Omega^2 \times \Gamma^2) \longrightarrow \mathcal{U}(\Omega \times \Gamma)$; that is, two mechanisms m_1 and m_2 are 'close' if their payoff functions $U(\cdot; m_1)$ and $U(\cdot; m_2)$ are 'close' as elements of $\mathcal{U}(\Omega \times \Gamma)$. This topology is not Hausdorff - two mechanisms that imply the same payoff functions for all buyers and sellers cannot be separated. This reflects the view that there is no reason to distinguish between such mechanisms. For $\mathcal{M}(\Omega^2 \times T^2)$, use the fact, shown in the proof, that there exists a compact metric $\mathcal{T}$ containing T. Define $\widehat{cep} : \mathcal{M}(\Omega^2 \times T^2) \longrightarrow \mathcal{U}(\Omega \times \mathcal{T})$ by

$$\widehat{cep}(m)\,(\cdot, t') = \begin{cases} U(\cdot, t'; m) & \text{if } t' \in Y \\ 0 & \text{otherwise.} \end{cases}$$

where $Y \subset T$ is *any* compact set such that $U(\cdot, t'; m) = 0$ if $t' \in T \backslash Y$. (Such a set Y exists by the assumption of compact support.) For $\mathcal{M}(\Omega^2 \times T^2)$, adopt the weak topology induced by this map.[15] Observe that were T compact, this construction would coincide with that defined above for the compact Γ (use Theorem 6.2(a)).

[15]We remind the reader that the topological space $\mathcal{U}(S)$ is defined (in Section 6.1) only for compact S. This explains our resorting to the ambient compact space $\mathcal{T}$.

With these topologies, ψ and ξ are both continuous and $\psi(T)$ is separable metric (even though its ambient space $\mathcal{M}(\Omega^2 \times T^2)$ is not even Hausdorff). While T and $\psi(T)$ are not generally compact, Γ and $\theta_{(\Gamma,\tilde{c},\tilde{\pi})}(\Gamma)$ are, the first by assumption and the second because $\theta_{(\Gamma,\tilde{c},\tilde{\pi})}$ is continuous. This asymmetry constitutes a gap; ideally, our candidate $\psi(T)$ for a universal space of indirect mechanisms should have similar topological properties to the primitive spaces Γ.

5.2. Robustness

If the restrictions that are imposed on the seller's ability to offer mechanisms to buyers are unreasonable, then the predictions that are forthcoming from a model of indirect competition will be unreliable. For this reason, we are interested in knowing when equilibria in particular models of indirect competition will survive the possibility that sellers might invent mechanisms that are not considered possible by the modeller. We have suggested that the universal class of mechanisms $\psi(T)$ provides an appropriate framework for examining such robustness of equilibria. Here we provide a formal result confirming this suggestion. In the single firm setting, the revelation principle shows that a mechanism that is optimal in the class of incentive compatible direct mechanisms is also optimal in an unrestricted sense. The theorem to follow may be thought of as a counterpart result for the present setting of competing mechanism designers.

Say that the payoff usc and NR continuation equilibrium (Γ_1, c_1, π_1) *extends* (Γ, c, π) if there exists an embedding $\alpha : \Gamma \longrightarrow \Gamma_1$ such that: For all γ, γ' in Γ,

$$(\theta\gamma)(\cdot, \gamma', \gamma') = \theta_1(\alpha(\gamma))(\cdot, \alpha(\gamma'), \alpha(\gamma')) \text{ on } \Omega^2. \tag{5.9}$$

In words, the actions implied by any pair of mechanisms γ and γ' in Γ are replicated by their translations $\alpha(\gamma)$ and $\alpha(\gamma')$, mechanisms in Γ_1. As an example, the continuation equilibrium $(\psi(T), c^*, \pi^*)$ provided by our revelation principle extends (Γ, c, π), with embedding $\alpha = \psi \circ e$. This is Theorem 4.1(c).

Say that an equilibrium (Γ, c, π, δ) is *robust* if for any extension (Γ_1, c_1, π_1) of (Γ, c, π), then $(\Gamma_1, c_1, \pi_1, \alpha[\delta])$ is an equilibrium, where $\alpha[\delta]$ is the randomization on Γ_1 induced by δ and α.

Theorem 5.3. *(a) If the equilibrium (Γ, c, π, δ) is robust, then $(\psi(T), c^*, \pi^*, \psi \circ e[\delta])$ is an equilibrium, where c^*, π^*, e and ψ are defined in Theorem 4.1.*

(b) If $(\psi(T), c, \pi, \delta)$ is an equilibrium, then the equilibrium is robust.[16]

[16]Because $\psi(T)$ need not be compact, this statement requires the obvious extension of the formal definition of robustness.

Proof. (a) If not, then there exists $m \in \psi(T)$ that is a profitable unilateral deviation by a seller. Define $\Gamma_1 = \psi(e(\Gamma)) \cup \{m\}$, making Γ_1 compact metric. Then (Γ_1, c^*, π^*) extends (Γ, c, π), (take the restriction of $\psi \circ e$ as the required emebedding α), and $(\Gamma_1, c^*, \pi^*, \psi \circ e[\delta])$ is not an equilibrium, contradicting robustness.

(b) Denote by θ the mapping $\psi(t) \longmapsto m_{\psi(t)}$ corresponding to the continuation equilibrium $(\psi(T), c, \pi)$; observe that the latter need not feature truthful reporting. Let (Γ_1, c_1, π_1) extend $(\psi(T), c, \pi)$, with embedding α. The appropriate form of (5.9) is

$$\theta(\psi(t))(\cdot, \psi(t'), \psi(t')) = \theta_1(\alpha(\psi(t)))(\cdot, \alpha(\psi(t')), \alpha(\psi(t'))) \text{ on } \Omega^2.$$

By the revelation principle, Γ_1 may be embedded into $\psi(T)$ by $\psi \circ e_1$, with associated truth-telling continuation equilibrium $(\psi(T), c^*, \pi^*)$, such that

$$(\theta_1\gamma_1)(\cdot, \gamma_1', \gamma_1') = \psi(e_1(\gamma_1))(\cdot, e_1(\gamma_1'), e_1(\gamma_1')) \text{ on } \Omega^2.$$

Consequently, $\psi \circ e_1 \circ \alpha$ embeds $\psi(T)$ into itself and

$$\theta(\psi(t))(\cdot, \psi(t'), \psi(t')) = \psi(e_1\alpha(\psi(t)))(\cdot, e_1\alpha(\psi(t')), e_1\alpha(\psi(t'))) \text{ on } \Omega^2.$$

This identity states that the two continuation equilibria $(\psi(T), c, \pi)$ and $(\psi(T), c^*, \pi^*)$ imply the same valuation and report contingent actions, after suitable translation by the embedding $\psi \circ e_1 \circ \alpha$ of $\psi(T)$ into itself. We are given that $(\psi(T), c, \pi, \delta)$ is an equilibrium. It follows that so is $(\psi(T), c^*, \pi^*, \psi \circ e_1 \circ \alpha[\delta])$.

Suppose that $(\Gamma_1, c_1, \pi_1, \alpha[\delta])$ is not an equilibrium. Then there exists a profitable unilateral deviation to some $\gamma \in \Gamma_1$ not in the support of $\alpha[\delta]$. But then the deviation to $\psi(e_1(\gamma))$ is profitable, contradicting the fact that $(\psi(T), c^*, \pi^*, \psi \circ e_1 \circ \alpha[\delta])$ is an equilibrium. ■

6. OUTLINE OF PROOF OF REVELATION PRINCIPLE

6.1. Upper Semi-Continuous Functions

The first step is to specify the space of usc functions precisely in order to make even the statements of Theorems 4.1 and 5.2 formally meaningful.

For any topological space S, denote by $\mathcal{U}(S)$ the set of upper semi-continuous (usc) functions from S into $[0, 1]$. Adopt the topology τ for $\mathcal{U}(S)$ that is generated by the following subbasis:

$$\{g : \exists s \in G,\ g(s) > \kappa\},\ \{g : \forall s \in K,\ g(s) < \kappa\}, \tag{6.1}$$

where G and K vary over the open and compact subsets of S and where κ varies over $[0,1]$. This is the weakest topology such that the mapping $g \longmapsto \sup_{s\in A} g(s)$ is lower semi-continuous (lsc) for each open A and usc for each compact A.

To illuminate the topology τ, we observe that it is consistent with topologies that have been widely employed. Denote by $\mathcal{F}(S)$ the set of all closed subsets of S endowed with the closed convergence topology [7, pp. 18-21]. Each closed set F can be identified with its indicator function, an usc function. Secondly, let $\mathcal{C}ap(S)$ denote the collection of all (regular Borel) capacities on S, endowed with the vague topology.[17] Each usc g can be associated uniquely with the capacity ν_g,

$$\nu_g(A) = \sup_{s\in A} g(s), \tag{6.2}$$

for Borel measurable set A. Therefore, we have

$$\mathcal{F}(S) \subset \ \mathcal{U}(S) \subset \mathcal{C}ap(S). \tag{6.3}$$

More to the point, the three topologies are consistent with one another in the sense that these set inclusions are topological embeddings.[18]

The association in (6.2) of any usc function with a set function may be helpful for clarifying the intuition underlying some of the analysis to follow. This is particularly true if the non-additive nature of ν_g is overlooked and if the reader thinks in terms of additive measures. Such an association should help to interpret what follows in more familiar terms such as marginals of measures, the weak convergence topology, and so on. For example, the meaning of Theorems 6.1 and 6.2 is clearer if one thinks of their widely known counterparts for measures. Similar counterparts exist for most key steps in our analysis.

The following two theorems describe a number of attractive and useful features of the topology.

Theorem 6.1. *If S is compact Hausdorff, then so is $\mathcal{U}(S)$. If S is also metric, then so is $\mathcal{U}(S)$.*

[17]See [18] for applications of capacities, also termed 'non-additive probabilities', to decision theory.

[18]The capacity ν_g is called a *sup measure*; see [13] for details regarding these non-additive measures and the assertions in (6.3). See also [11, p. 285] for more on the connection between $\mathcal{U}(S)$ and other topological spaces. For example, the mapping $g \longmapsto \{(s,t) : 0 < t \leq g(s)\}$, (the latter hypograph of g being a closed subset of $S \times (0,1]$), is a homeomorphism of $\mathcal{U}(S)$ into the space of closed subsets of $S \times (0,1]$, endowed with the Fell topology.

Proof. See [13, Theorems 2.2-3, Cor. 2.5]. For the second claim, use also [3, Theorem XI.4.1]. ■

Remark 1. *τ is the* unique *compact Hausdorff topology on $\mathcal{U}(S)$ [12, p. 19].*

Theorem 6.2. *Let S be compact Hausdorff. Then:*

(a) If $e : S' \longrightarrow S$ is continuous for some space S' and if $\widehat{e}$ is defined by

$$(\widehat{e}g)(s) = \sup\{g(s') : e(s') = s\}, \tag{6.4}$$

and the sup *is understood to equal zero where the constraint set is empty, then $\widehat{e} : \mathcal{U}(S') \longrightarrow \mathcal{U}(S)$ is continuous. In the special case where $S' = A \subset S$, then $\widehat{e}$ takes g into $\widehat{e}g$, where $\widehat{e}g(s) = g(s)$ if $s \in A$ and 0 otherwise, and $\widehat{e}$ is a homeomorphism of $\mathcal{U}(A)$ onto $\mathcal{U}(S \mid A) \equiv \{h \in \mathcal{U}(S) : h = 0 \text{ on } S \backslash A\}$.*

(b) Let $S = S_1 \times S_2 \times S_3$ and let $\xi : S_3 \rightsquigarrow S_2$ be a compact-valued correspondence such that $\xi(A_3)$ is open (closed) for every open (closed) set $A_3 \subset S_3$. Then $g \longmapsto g_1$, $g_1(s_1, s_3) = \max_{s_2 \in \xi(s_3)} g(s_1, s_2)$, is a continuous map from $\mathcal{U}(S_1 \times S_2)$ into $\mathcal{U}(S_1 \times S_3)$.

(c) The mapping $(s, g) \longmapsto g(s)$ is usc on $S \times \mathcal{U}(S)$.

See Appendix B for a proof. These properties will be exploited heavily below. For (b), two special cases are exploited. In the special case where $\xi(\cdot) = S_2$, (b) can be rewritten to say that the mapping from $\mathcal{U}(S_1 \times S_2)$ into $\mathcal{U}(S_1)$ taking g into g_1, $g_1(s_1) = \max_{s_2 \in S_2} g(s_1, s_2)$, is continuous. It is useful to think of g_1 as the S_1-marginal of g. With this terminology, the operation of taking a marginal is continuous, as it is in the more familiar setting of measures with the weak convergence topology. For the second special case, let $f : S_2 \longrightarrow S_3$ be a continuous function and $\xi = f^{-1}$. The hypothesized properties for ξ are implied by the continuity of f. These hypotheses cannot be deleted in (b). For example, if $\xi(s_3) = \{\overline{s}_2\}$ for all s_3, then $g_1(s_1, s_3) = g(s_1, \overline{s}_2)$, but the mapping from g into the restriction $g(\cdot, \overline{s}_2)$ is generally not continuous (see [13, Theorem 4.5] for related results).

6.2. Hierarchies

In this subsection, we begin the construction of the universal types space T. Because mechanisms are described through the expected payoff functions that they deliver, and these functions are usc, we are led to study hierarchies of usc functions. These are studied here with only informal connection to mechanisms. This connection is made precise and formal in the following subsection. As mentioned in the introduction, the analysis of hierarchies parallels the analysis of hierarchies of probability measures (see [9] and [2], for example) and is a special case of the analysis in [4].[19] For the convenience of the reader, the construction of hierarchies of usc functions is briefly outlined here.

Let the two firms be i and j and consider the message spaces that each might use in designing a mechanism. To describe these spaces, define

$$C_0^i = \Omega,\ C_1^i = \Omega \times \mathcal{U}(\Omega),\ C_n^i = C_{n-1}^i \times \mathcal{U}(C_{n-1}^j),\ n \geq 1,$$

and similarly for j. By above, each C_n^i is a well-defined compact metric space. Use of Ω as the message space would have buyers report only their valuations. However, it may be advantageous for i to have buyers report also a description of the mechanism used by j, which description would correspond to a payoff function $g \in \mathcal{U}(C_0^j)$. This leads to adoption of C_1^i as the message space, but it is also *incomplete* in that a message in C_1^i does not include a description of j's mechanism based on C_1^j. Similarly, each C_n^i and C_n^j is an incomplete message space. We show, however, that there exists a well-defined message space $C = C^i = C^j$ that is rich enough to include descriptions of mechanisms over C. Such *completeness* corresponds to the homeomorphism

$$C^i \sim_{hmeo} \Omega \times\ \mathcal{U}(C^j).$$

Write

$$C^i = \Omega \times \mathcal{T}^i, \text{ and } C^j = \Omega \times \mathcal{T}^j,$$

corresponding to the hypothesis that a message consists of a report of both valuation and 'type', namely everything else. Then types spaces must satisfy

$$\mathcal{T}^i \sim_{hmeo}\ \mathcal{U}(\Omega \times \mathcal{T}^j) \text{ and } \mathcal{T}^j \sim_{hmeo}\ \mathcal{U}(\Omega \times \mathcal{T}^i).$$

Suppress firm specific superscripts. We employ the above sequence of incomplete message spaces to construct a types space $\mathcal{T}$ satisfying

$$\mathcal{T}\ \sim_{hmeo}\ \mathcal{U}(\Omega \times \mathcal{T}); \tag{6.5}$$

the associated complete message space is $C = \Omega \times \mathcal{T}$.

[19] It is only Theorem 6.3 that is a special case. The proof of our main result, Theorem 4.1, requires additional effort (see the next subsection and Appendix C).

Define $\mathcal{T}_0 = \Pi_0^\infty \mathcal{U}(C_n)$ with generic element $t = (g_0, g_1, ..., g_n, ...)$. Whereas this type is interpreted as a complete description reported to i of the mechanism used by j, each component g_n provides only a partial description of j's mechanism. Think of $g_n(c_n)$ as the maximum expected utility for buyer 1 if the message $c_n \in C_n$ is sent to j, where 'maximum' is relevant because j's mechanism uses a richer message space than C_n. Not all such elements t are sensible, however, because the descriptions g_n may contradict one another. To clarify, each g_{n-1} in $\mathcal{U}(C_{n-1})$ may be identified with an element of $\mathcal{U}(C_n) = \mathcal{U}(C_{n-1} \times \mathcal{U}(C_{n-1}))$, by viewing g_{n-1} as a function on $C_{n-1} \times \mathcal{U}(C_{n-1})$ that does not depend on the second argument. If the interpretation 'maximum expected utility' is to be warranted, then we would expect that for all $n \geq 1$, $c_{n-1} \in C_{n-1}$,

$$\max_{\mathcal{U}(C_{n-1})} g_n(c_{n-1}, \cdot) = g_{n-1}(c_{n-1}). \tag{6.6}$$

Refer to a type $t = (g_0, ..., g_n, ...)$ satisfying this restriction as *coherent*. The subspace of $\mathcal{T}_0$ consisting of coherent types is denoted $\mathcal{T}_1$.

An important first step in the construction is to note that $\mathcal{T}_1$ is homeomorphic to $\mathcal{U}(\Omega \times \mathcal{T}_0)$.[20] Let $\Psi : \mathcal{T}_1 \longrightarrow \mathcal{U}(\Omega \times \mathcal{T}_0)$ denote the homeomorphism. Each coherent type t includes a description $\Psi(t)$ of the mechanism used by firm j were j to use the message space $\Omega \times \mathcal{T}_0$. But if types are coherent, then j would 'know' this and in principle dismiss incoherent types by giving them low payoffs. This suggests consideration of the decreasing sequence of types spaces$\{\mathcal{T}_k\}$, where

$$\mathcal{T}_2 = \{t \in \mathcal{T}_1 : \Psi(t) = 0 \text{ on } \Omega \times (\mathcal{T} \backslash \mathcal{T}_1)\},$$

$$\mathcal{T} = \{t \in \mathcal{T}_1 : \Psi(t) = 0 \text{ on } \Omega \times (\mathcal{T}_0 \backslash \mathcal{T}_{k-1})\}, \quad k \geq 2.$$

Finally, define

$$\mathcal{T} = \cap \mathcal{T}_k.$$

Theorem 6.3. *The space $\mathcal{T}$ is nonempty, compact metric and satisfies (6.5).*

Proof. We observed above that $\mathcal{U}(C)$ is homeomorphic to a subspace of capacities. Hierarchies of capacities are a special case of the class of hierarchies studied in [4]. To be precise, the proof follows from Theorem 6.1 above and from Theorems 4.2, 4.3 and 6.1 of [4].

[20]As explained in the proof of Theorem 6.3 below, this and other assertions to follow are corollaries of [4].

We outline only two details in the proof that will be needed later. The first, that logically precedes the statement of the theorem, is the asserted homeomorphism $\Psi : \mathcal{T}_1 \longrightarrow \mathcal{U}(\Omega \times \mathcal{T}_0)$. Given a coherent type $t = (g_0, ..., g_n, ...)$, $g_n \in \mathcal{U}(C_n)$, define $\Psi(t)$ as follows: For any $z \in \Omega \times \mathcal{T}_0 = \Omega \times \Pi_{n=0}^{\infty}\mathcal{U}(C_n)$, let z^N be the projection of z onto $\Omega \times \Pi_{n=0}^{N-1}\mathcal{U}(C_n) = C_N$ and define

$$\Psi(t)(z) = \inf_N g_N(z^N). \tag{6.7}$$

Then $\Psi(t)(\cdot)$ is usc because it is the pointwise limit of a decreasing sequence of usc functions. Further, Ψ is the desired homeomorphism. (This is readily proven directly. Alternatively, it follows by translation to the present setting of the arguments in [4] cited above.)

To prove (6.5), observe first that

$$\mathcal{T} = \cap \mathcal{T}_k = \{t \in \mathcal{T}_1 : \ \Psi(t) = 0 \text{ on } \Omega \times (\mathcal{T}_0 \backslash \mathcal{T})\}.$$

The latter set is homeomorphic to $\mathcal{U}(\Omega \times \mathcal{T})$; see Theorem 6.2. ■

Remark 2. *Note that $\Psi(t)(\cdot)$, defined by (6.7), is not necessarily continuous on $\Omega \times \mathcal{T}$ even if each g_N is a continuous function. This is one reason that we cannot replace upper semi-continuity by continuity as the basic regularity condition.*

The space $\mathcal{T}$ is a large and complicated space even if Ω is finite; in that case, even $\mathcal{U}(\Omega)$ has the cardinality of the continuum because it equals $[0,1]^\Omega$. Some examples of elements of $\mathcal{T}$ may help to clarify its nature.

Examples of Types: Fix $N \geq 0$ and a closed set of functions $\mathcal{V}_N \subset \mathcal{U}(C_N)$. Define a set of types $Y \subset \mathcal{T}$ as follows: For each $g_N \in \mathcal{V}_N$, construct a type $t = (g_n)_0^\infty$, where $g_n \in \mathcal{U}(C_n) = \mathcal{U}(\Omega \times \Pi_0^{n-1}\mathcal{U}(C_k))$. First, for $n \leq N$, let g_n be the 'marginal' of g_N, that is,

$$g_n(\omega, h_0, ..., h_{n-1}) = \sup \{g_N(\omega, h_0, ..., h_{n-1}, ..., h_{N-1}) : \ h_k \in \mathcal{U}(C_k), \ n \leq k < N\}.$$

For $n > N$, let

$$g_n(\omega, h_0, ..., h_{n-1}) = \begin{cases} g_{n-1}(\omega, h_0, ..., h_{n-2}) & \text{if } h_{n-1} \in \mathcal{V}_{n-1}, \\ 0 & \text{otherwise,} \end{cases}$$

where the (closed) sets $\mathcal{V}_k \subset \mathcal{U}(C_k)$ are also defined inductively; $\mathcal{V}_{N+1}$ equals the set of all functions g_{N+1} obtained as above beginning by using some g'_N in $\mathcal{V}_N$, and so on. Then $t = (g_n)$ lies in $\mathcal{T}$. We compute that if $t' = (h'_n)$, then

$$\Psi(t)(\omega, t') = \begin{cases} g_N(\omega, h'_0, ..., h'_{N-1}) & \text{if } t' \in Y, \\ 0 & \text{otherwise.} \end{cases}$$

Consequently, t describes a mechanism in which the expected payoff equals g_N as a function of valuation and 'level N' descriptions of the other firm, as long as the buyer reports that the other firm has a type in Y. Any other report is penalized by the assignment of zero utility. Further interpretation is afforded by examining the structure of the g_n's. The assignment of zero utility by g_{N+1} to the report $(h_0, ..., h_N) \notin \mathcal{V}_N$ reflects 'first-order knowledge' that the payoff function at the other firm is an element of $\mathcal{V}_N$. Similarly, the definition of g_{N+2} reflects 'second-order knowledge' that the other firm has the first order knowledge described above. Similarly for higher order functions. Therefore, if each firm is characterized by a type in Y, then the expected payoff as a function of valuation at each firm is given by some function in $\mathcal{V}_N$ and this is 'common knowledge'.

While $Y \subset \mathcal{T}$, it is only types in $T \subset \mathcal{T}$ that describe mechanisms. Therefore, it is of interest to have an example where $Y \subset T$. This can be provided by using the three examples of Γ's (two forms of price competition and competing auctions) and the associated continuation equilibria described in Section 3 and Appendix A. For concreteness, consider the price matching example. The set of types associated with price-matching indirect mechanisms may be described as above if we take $N = 1$ and define $\mathcal{V}_1$ as follows: Let $\Phi_0 : [0,1] \longrightarrow \mathcal{U}(\Omega)$, where

$$\Phi_0(q)(\cdot) = \sup_{q' \in [0,1]} U(\cdot, q'; q) = U(\cdot, 0; q);$$

and let $\Phi_1 : [0,1] \longrightarrow \mathcal{U}(\Omega \times \mathcal{U}(\Omega))$, where

$$\Phi_1(q)(\cdot, h_0) = \sup_{q'} \{U(\cdot, q'; q) : \Phi_0(q)(\cdot) = h_0(\cdot)\}.$$

(The supremum equals 0 if the constraint set is empty.) Then take $\mathcal{V}_1 = \{\Phi_1(q)(\cdot) : q \in [0,1]\}$. Similar constructions apply to the simple price competition and auctions examples. The fact that $N = 1$ works reflects the common feature of these models whereby non-redundancy is delivered by the separating property of the first level σ-algebra Σ_1. In each case, the implied set of types Y is a subset of T (this follows from Theorem 4.1 and its proof).

6.3. Proof Outline

The construction of $\mathcal{T}$ dealt with arbitrary usc functions without any formal reference to mechanisms. Thus the relevance of $\mathcal{T}$ to mechanisms may not be evident. However, we will identify a subset $T \subset \mathcal{T}$ that satisfies the claims made in the revelation principle (Theorems 4.1 and 5.2).

In the rest of this section, we outline the construction of T and ψ. Remaining details are provided in Appendix C.

One obvious restriction defining the sought after subspace T of $\mathcal{T}$ is that any type t in T must be such that $\Psi(t)$ has the form

$$\Psi(t)(\omega, t') = U(\omega, t'; m) = \int u(m(\omega, \omega', t', t'), \omega)\, dF(\omega'), \tag{6.8}$$

for some $m \in \mathcal{M}(\Omega^2 \times T^2)$. Only in this case can one think of t as describing the expected payoff of a suitable direct mechanism employing message space $\Omega \times T$. If we denote m by $\psi(t)$, then (T, ψ) will be the desired pair.

But there are also other restrictions that are natural. For example, if only mechanisms in $\mathcal{M}(\Omega^2 \times T^2)$ are relevant, then any mechanism should 'know' this and penalize (by assigning zero utility) a reported type for the other firm that contradicts this. Accordingly, T (and ψ) should satisfy: For each $t \in T$,

$$\Psi(t)(\cdot) = U(\cdot; m) \text{ for some } m \in \mathcal{M}(\Omega^2 \times T^2) \text{ and} \tag{6.9}$$

$$\Psi(t)(\cdot, t') = 0 \text{ if } t' \in \mathcal{T} \setminus T. \tag{6.10}$$

In the above informal sense of 'knowledge', this condition can be interpreted as saying that types in T express the 'common knowledge' that mechanisms lie in $\mathcal{M}(\Omega^2 \times T^2)$.

An iterative construction of the (largest) set of types modelling the noted common knowledge will now be described. As in Section 6.1, if $B \subset S$ is compact, then $\mathcal{U}(S \mid B) = \{g \in \mathcal{U}(S) : g = 0 \text{ on } S \backslash B\}$. For arbitrary (possibly nonmeasurable) A, define

$$\mathcal{U}(S \mid A) = \cup \{\mathcal{U}(S \mid B) : B \subset A \text{ compact}\}.$$

If g is in $\mathcal{U}(S \mid A)$, say that g *knows* A.

Construct Y^k recursively. Let

$$Y^0 = \Psi^{-1}\left(\{U(\cdot; m) \in \mathcal{U}(\Omega \times \mathcal{T}) :\, m \in \mathcal{A}^{\Omega^2 \times \mathcal{T}^2}\}\right).$$

Let $\mathcal{M}^0(\Omega^2 \times \mathcal{T}^2)$ consist of those mechanisms m in $\mathcal{A}^{\Omega^2 \times \mathcal{T}^2}$ such that m satisfies the strong measurability constraints (5.1)-(5.2), incentive compatible (in the sense of (5.4)), USC, has compact support and satisfies the following variations of (5.3) and (5.5) respectively:[21] For each $m' \in \mathcal{A}^{\Omega^2 \times \mathcal{T}^2}$ and $t' \in Y^0$, if $\Psi(t) = U(\cdot; m)$ and $\Psi(t') = U(\cdot; m')$, then: (a) The equality (5.3) is satsified with with (t, t') replacing (γ, γ'); and (b) the inequality (5.5) is satisfied with (t, t') replacing (γ, γ').

Then let

$$Y^{k+1} = \Psi^{-1}\left[\mathcal{U}(\Omega \times \mathcal{T} \mid \Omega \times Y^k) \cap cep\left(\mathcal{M}^0(\Omega^2 \times \mathcal{T}^2)\right)\right].$$

Finally, define

$$T = \cap_{k=0}^{\infty} Y^k. \tag{6.11}$$

Informally, the above recursive construction 'suggests' the limiting property that

$$T = \Psi^{-1}\left[\mathcal{U}(\Omega \times \mathcal{T} \mid \Omega \times T) \ \cap \ cep\left(\mathcal{M}^0(\Omega^2 \times \mathcal{T}^2)\right)\right]. \tag{6.12}$$

(This can be verified as follows: That T contains the set on the right is immediate. For the converse, let $t \in \cap_{k=0}^{\infty} Y^k$. Then, $\Psi(t)$ knows each Y^k. It follows that $\Psi(t)$ knows their intersection, which is T. In addition, $\Psi(t) = U(\cdot; m^0)$ for some $m^0 \in \mathcal{M}^0(\Omega^2 \times \mathcal{T}^2)$.) Because $T \subset \mathcal{T}$ and the latter is compact metric, conclude that T is separable metric [3, pp. 176, 233].

Having thus defined and characterized T, turn to the construction of $\psi : T \longrightarrow \mathcal{M}(\Omega^2 \times T^2)$. From (6.12), it follows that for any $t \in T$, there exists $m \in \mathcal{M}^0(\Omega^2 \times \mathcal{T}^2)$ such that $\Psi(t) = U(\cdot; m)$ and $\Psi(t)$ knows $Y \subset T$ for some compact subset Y. Define $\psi(t)$ to equal the restriction of m to $\Omega^2 \times T^2$. Then (T, ψ) is a self-referential class. Moreover, it follows from the construction and (6.12) that

$$\psi(T) \subset \mathcal{M}(\Omega^2 \times T^2) \text{ and}$$
$$cep\left(\psi(T)\right) = \left\{g_{|\Omega \times T} : g \in \Psi(T)\right\} = \left\{U(\cdot; m) : \ m \in \mathcal{M}(\Omega^2 \times T^2)\right\}, \tag{6.13}$$

both of which are asserted in Theorem 5.2. Finally, for each $t \in T$,

$$\Psi(t)(\cdot) = U(\cdot; \psi(t)) \text{ on } \Omega \times T,$$

as asserted in Theorem 4.1. Because Ψ is one-to-one, so is ψ.[22]

[21] These conditions require special treatment because they involve the mapping θ, which here is not given. Rather, it must be constructed simultaneously with T; eventually, ψ will serve as the appropriate map θ.

[22] ψ is continuous as described following Theorem 5.2. It follows from [3, p. 175] and the separability of T that $\psi(T)$ is separable.

See Appendix C for a proof that (T, ψ) satisfies the other assertions in the theorems.

7. EXTENSIONS AND POTENTIAL APPLICATIONS

7.1. Non-metrizable Spaces

While compact Hausdorff is essential for our arguments, the assumption of metrizability is used only for one purpose - to *simplify* the statement of the assumption of non-redundancy (Section 4.1). Lemma C.4 shows the nature of this simplification, while Lemma C.5 reveals the purpose of non-redundancy, namely to prove that e is one-to-one. An alternative would be to use this invertibility as the *definition* of NR, in which case metrizability is unnecessary. Theorems 4.1 and 5.2 remain valid with the natural exception that T and $\psi(T)$ need not be separable metric. However, they are completely regular because each can be embedded in a compact Hausdorff topological space [3, XI.8].

A related topological issue is the asymmetry inherent in the failure of T (and $\psi(T)$) to be compact even though the primitive spaces Γ of indirect mechanisms are compact by assumption. This asymmetry is a technical issue that requires further study.

7.2. Sellers' Payoffs

We mentioned in Section 4.1 that it was both feasible and desirable to extend our analysis so that sellers' payoffs, in addition to buyers' payoffs, are used to describe mechanisms. Here we outline how such an extension can be accomplished.

The first step is to reformulate the assumptions for a continuation equilibrium $(\Gamma, \tilde{c}, \tilde{\pi})$. A seller using γ and whose competitor is using γ' receives expected payoff

$$V(\gamma'; \gamma) = \int v\left(m_\gamma(\omega, \omega', \gamma', \gamma'), \omega, \omega'\right)\, dF\, dF. \tag{7.1}$$

In order to unify the notation for dealing with buyers and sellers, add a discrete point $\{\varpi\}$ to Ω and define $\overline{\Omega} = \Omega \cup \{\varpi\}$. Then proceed roughly as before with $\overline{\Omega}$ replacing Ω as the primitive valuations space. In order to do so, define $U : \overline{\Omega} \times \Gamma^2 \longrightarrow [0, 1]$ so that for $\omega \in \Omega$, $U(\omega, \gamma'; \gamma)$ retains the interpretation as the buyer's payoff function, while $U(\varpi, \gamma'; \gamma) \equiv V(\gamma'; \gamma)$, representing the seller's payoff function. Expand the domain of the direct mechanism (2.3) m_γ so that

$$m_\gamma(\omega, \omega', \cdot) = \underline{a} \quad \text{if } (\omega, \omega') \notin \Omega^2, \tag{7.2}$$

reflecting the fact that only valuations in Ω are conceivable for buyers. Then the assumptions payoff usc and NR, expressed relative to the expanded valuations space $\overline{\Omega}$, are both meaningful and appropriate. Theorem 4.1, with $\overline{\Omega}$ replacing Ω, is valid; the existing proof requires only trivial notational modifications. In the universal class of mechanisms, buyers report a valuation in Ω, (because it doesn't pay to report ϖ), and the other firm's type, now expanded to include sellers payoffs. For Theorem 5.2, the definition of $\mathcal{M}(\overline{\Omega}^2 \times T^2)$ must be modified in obvious ways, for example, by adding a restriction paralleling (7.2). Note that this reformulation of NR delivers a *weaker* assumption than that used previously because the use of sellers payoffs makes it easier to distinguish between distinct mechanisms. In that sense, the reformulated revelation principle generalizes the one stated in the text.[23]

7.3. More Specialized Models of Competition

It has been our intention to assume as little as possible *a priori* about the nature of competition, because of the view that such assumptions are invariably ad hoc. However, there may be situations where additional restrictions would be acceptable to some readers. Therefore, we outline the implications for our analysis of such additional restrictions.

Our indirect mechanisms assume that the reporting strategy $\tilde{c}$ is a function of the buyer's type and the mechanism that the buyer has seen offered at the other firm. One alternative is to allow the buyer to observe whether or not another buyer has selected the seller before choosing his report, so that the reporting strategy would be a pair consisting of a report when he is alone and a report when he is competing with another buyer. Our formalism allows for this possibility. Simply let $C = C_0^{P^2}$, in which case the reporting strategy $\tilde{c}$ is a plan that specifies a message that depends on whether or not the other buyer is there. The implication for direct mechanisms is that the payoff to a buyer who is the only participant is independent of his report about the other firm.

The existing theorem still applies and thus $\mathcal{M}(\Omega^2 \times T^2)$ is still a sufficiently large class to replicate any continuation equilibrium, but it is 'too large' in general because some mechanisms in $\mathcal{M}(\Omega^2 \times T^2)$ will generate payoffs that may vary with the report about the other firm provided by a lone participant. Therefore, viewed as indirect mechanisms over the message space $C = \Omega \times T$, they will not belong

[23] On the other hand, the current assumption that each seller's payoff is usc in the other firm's mechanism was not needed in the earlier version.

to the newly restricted admissible class. To improve upon $\mathcal{M}(\Omega^2 \times T^2)$, suppose that the modified model of competition implies restrictions on direct mechanisms that can be modelled by a correspondence $\mathcal{M}^*(\cdot) \subset \mathcal{M}(\cdot)$. Then, for some specifications of $\mathcal{M}^*(\cdot)$, the revelation principle may be strengthened to assert that there exists $T^* \subset T$, such that the assertions in Theorems 4.1 and 5.2 remain valid of T is replaced everywhere by T^*. In this way, one can obtain a revelation principle for a number of alternative specializations of our model.

7.4. Decentralization

An important application of the standard revelation principle is to answer the question "which (type-contingent) allocations are feasible?" The typical answer, of course, is "all incentive compatible allocations." From the perspective of this paper, this answer presumes the existence of a single mechanism designer or central planner and must be reconsidered in a decentralized environment. Here we elaborate briefly, with particular attention to the meaning of 'decentralization'.[24]

An *allocation* in our setting is a function $\alpha : \Omega^2 \to \mathcal{A}^2$, where $\alpha(\omega, \omega')$ represents the actions (including recommended choice probabilities) forthcoming at the two firms given the pair (ω, ω') of valuations. The allocations that can be implemented by a central planner are those satisfying the strong measurability and incentive constraints (5.1)-(5.5), and the USC condition, where these are reformulated in the natural way. Such allocations can be implemented by having the buyers report their valuations to a central planner who then instructs each seller to choose the action specified by α.

But what if we require that the implementation be *decentralized* in the sense that the allocation be delivered by an equilibrium of the sellers' mechanism design game? The answer, naturally, depends on the power retained by the central planner. For example, suppose that the central planner can specify the set Γ of indirect mechanisms available to sellers. Then, any allocation that is implementable in the above sense can also delivered by an equilibrium relative to some Γ. To see this, define $\Gamma = \{\gamma^*\}$ where γ^* is a mechanism having message space $C = \Omega$ and satisfying

$$\gamma(\omega, \omega') = \alpha(\omega, \omega').$$

Because the allocation α is incentive compatible, there is a continuation equilibrium for Γ in which buyers report their valuations to sellers truthfully. The

[24]This discussion was stimulated by a question posed by John Geanakoplos.

equilibrium allocation associated with Γ is clearly the same as that associated with the allocation α.

However, the analysis differs if sellers are free to employ any mechanism that they view as profitable. The corresponding notion of decentralization is that the allocation be delivered by a robust (in the sense of Section 5.2) equilibrium, or equivalently, by an equilibrium relative to our universal set of mechanisms $\psi(T)$. The set of allocations that can be implemented in a robust way ought to be a strict subset of those implementable by an omnipotent central planner. However, it remains an open problem to characterize this set.

A. APPENDIX: Auctions and Added Details

A.1. Details for Price Competition Models

We verify some assertions regarding the two models of price competition that were made in Section 4.1. They deal with buyers' payoff functions.

For simple price competition, it follows from (3.2) and (3.5) that $U(\omega, q'; q) =$

$$\begin{cases} F(\omega^*(q,q')) \frac{\max[\omega-q,0]}{2} + (1 - F(\omega^*(q,q'))) \max[\omega - q, 0] & \text{if } q \leq q' \\ F(\omega^*(q,q')) \max[\omega - q, 0] + (1 - F(\omega^*(q,q'))) \frac{\max[\omega-q,0]}{2} & \text{otherwise.} \end{cases} \tag{A.1}$$

By the continuity of $\omega^*(q,q')$, this is evidently continuous in (ω, q, q'), which can be shown to imply payoff upper-semicontinuity.

Proceed similarly for the price matching model to derive the payoff function $U(\omega, q'; q) =$

$$\begin{cases} F(\omega^*(q,q')) \frac{\max[\omega-q,0]}{2} + (1 - F(\omega^*(q,q'))) \max[\omega - q, 0] & \text{if } q \leq q' \\ F(\omega^*(q,q')) \max[\omega - q, 0] + (1 - F(\omega^*(q,q'))) \frac{\max[\omega-q',0]}{2} & \text{otherwise.} \end{cases} \tag{A.2}$$

Joint continuity in (ω, q, q') implies payoff upper-semicontinuity for the underlying continuation equilibrium.

Lemma A.1. *Non-redundancy is satisfied for each of the price competition models and the continuation equilibria described in Section 3.*

Proof. In each case $\sup_{q'} U(\omega, q'; q)$ is achieved at $q' = 0$. Therefore, it suffices to show that

$$q_1 \neq q_2 \implies U(\cdot, 0; q_1) \neq U(\omega, 0; q_2).$$

For simple price competition,

$$U(\omega,0;q) = F\left(\omega^{*}(q,0)\right)(\omega-q) + \left[1-F\left(\omega^{*}(q,0)\right)\right](\omega-q)/2. \tag{A.3}$$

Thus $U(\cdot,0;q_1) = U(\cdot,0;q_2)$ would imply (by equating coefficients of ω) that

$$F\left(\omega^{*}(q_1,0)\right) = F\left(\omega^{*}(q_2,0)\right). \tag{A.4}$$

This is impossible if (3.4) determines these two cutoffs. The only other possibility is that both cutoffs equal 1 and so the common value in (A.4) also equals 1. But then (A.3) for each of q_1 and q_2 implies that $q_1 = q_2$.

For price matching, suppose that $q_2 > q_1$. Then

$$\begin{gathered} U\left(\cdot,0;q_1\right) = U\left(\cdot,0;q_2\right) \text{ on } \left(q_1,q_2\right) \implies \\ F\left(\omega^{*}\left(q_1,0\right)\right)\left(\omega-q_1\right) + \left(1-F\left(\omega^{*}\left(q_1,0\right)\right)\right)\frac{\omega}{2} = \\ \left(1-F\left(\omega^{*}(q_2,0)\right)\right)\frac{\omega}{2} \end{gathered}$$

on the same interval, implying

$$F\left(\omega^{*}\left(q_1,0\right)\right) + \frac{\left(1-F\left(\omega^{*}\left(q_1,0\right)\right)\right)}{2} = \frac{\left(1-F\left(\omega^{*}\left(q_2,0\right)\right)\right)}{2},$$

and hence that $F\left(\omega^{*}(q_1,0)\right) = F\left(\omega^{*}(q_2,0)\right) = 0$. This contradicts the definition of the cutoff values, whereby either (3.6) or $\omega^{*}(q_1,0) = \omega^{*}(q_2,0) = 1$. ■

A.2. Auctions

Here we discuss a class of indirect mechanisms consisting of auctions. This class is of interest to us for two primary reasons. First, this class of indirect mechanisms permits actions to vary with buyers' valuations in a relatively complete way. Secondly, it further demonstrates the relevance of our assumptions of payoff usc and non-redundancy for continuation equilibria.

The environment is identical to that in the price competition models. In particular, *pure* actions consist of pairs (q,μ) with $\mu = (\mu_1,\mu_2)$ such that $\mu_1+\mu_2 \leq 1$. The interpretation is that with probability μ_i buyer i is given the *option* of trading at the price q. Other notation is also adopted from the earlier examples. Here we assume that the distribution F from which the buyers' valuations are drawn is atomless and has a strictly positive density.

When the seller's compete by offering alternative auctions, each seller announces a reserve price for the auction. Buyers see the two reserve prices, then select one and only one seller and submit a bid to that seller. The seller chooses the buyer who submits the higher bid and offers that buyer the option to trade at the pre-announced reserve price, or the bid of the other buyer, whichever is higher. If only one buyer selects the seller, the seller offers him the opportunity to trade at the reserve price.

Buyers report valuations, which are elements of $[0,1]$; thus the message space C equals $[0,1]$. Let $\Gamma = \{\gamma_r : r \in [0,1]\}$ denote the set of indirect mechanisms where $\gamma_r : C^2 \to \mathcal{A}$ is defined by

$$\gamma_r(c, c') = \begin{cases} \{(r,1,0), (r,0,1), (c',1,0), \pi, \pi\} & \text{if } c \geq c' \geq r \\ \{(r,1,0), (r,0,1), (c,0,1), \pi, \pi\} & \text{if } c' > c \geq r \\ \{(r,1,0), (r,0,1), (r,1,0), \pi, \pi\} & \text{if } c \geq r > c' \\ \{(r,1,0), (r,0,1), (r,0,1), \pi, \pi\} & \text{if } c' \geq r \geq c \\ \{(r,1,0), (r,0,1), (r,0,1), \pi, \pi\} & \text{if } r > c' \geq c \\ \{(r,1,0), (r,0,1), (r,1,0), \pi, \pi\} & \text{if } r > c > c' \end{cases}$$

To clarify this notation, consider the action prescribed by the mechanism γ_r if $c \geq c' \geq r$. That action is $\{(r,1,0), (r,0,1), (c',1,0), \pi, \pi\}$, according to which buyer 1 is offered the good at the reserve price r if only he selects the firm, the second triple is interpreted similarly for buyer 2, and $(c', 1, 0)$ indicates that if both buyers select the firm, then the first buyer is offered the good at a price equal to the other buyer's reported valuation. As in the case of simple price competition, the seller makes no attempt to influence the choice probabilities of the buyers, so we set these arbitrarily at π.

Now consider continuation equilibria. It is well known that bidding the true valuation is at least a weakly dominant strategy in the second price auction independent of the buyer's beliefs about his opponent. This implies that $\tilde{c}(\omega, \gamma, \gamma') = \omega$ constitutes an equilibrium reporting strategy. The equilibrium choice strategy for this problem was originally described by [16]. It is characterized by a single cutoff valuation $\omega^*(\gamma, \gamma')$. Buyers whose valuations are below the cutoff choose the seller with the lower reserve price with probability 1. Buyers whose valuations are above the cutoff choose each of the two sellers with equal probability.

Standard arguments (developed in [16]) can be used to show that the cutoff valuation is given by the solution to

$$\left[1 - \frac{1 - F(y)}{2}\right][y - r']$$

$$= \int_r^y \left[1 - (F(y) - F(s)) - \frac{1 - F(y)}{2}\right] ds \tag{A.5}$$

assuming that $r \leq r'$ whenever this solution exists on the interior of the interval $[0,1]$, and the cutoff is given by 1 otherwise. The solution to this equation is monotonically increasing in r' and monotonically decreasing in r. To see this simply note that under our assumptions, both sides of the equation are monotonically increasing, but the derivative of the left hand side is strictly larger than the derivative of the right hand side. Since both sides are monotonic functions of r or r', the result follows.

Again using the results in [16], if $r \leq r'$ then the buyer's utility is given by $U(\omega, r'; r) = 0$ when $\omega \leq r$;

$$U(\omega, r'; r) =$$

$$\int_r^\omega \left[1 - (F(y(r,r')) - F(s)) - \frac{1 - F(y(r,r'))}{2}\right] ds \tag{A.6}$$

when $r \leq \omega < y(r,r')$; and

$$\int_r^{y(r,r')} \left[1 - (F(y(r,r')) - F(s)) - \frac{1 - F(y(r,r'))}{2}\right] ds$$

$$+ \int_{y(r,r')}^\omega \left[1 - \frac{1 - F(s)}{2}\right] ds \tag{A.7}$$

when $\omega \geq y(r,r')$.

For $r > r'$, we have: $U(\omega, r'; r) = 0$ when $\omega < r$;

$$U(\omega, r'; r) = \left[1 - \frac{1 - F(y(r,r'))}{2}\right] [\omega - r] \tag{A.8}$$

when $r' \leq \omega < y(r,r')$; and

$$U(\omega, r'; r) = \left[1 - \frac{1 - F(y(r,r'))}{2}\right] [y(r,r') - r]$$

$$+ \int_{y(r,r')}^\omega \left[1 - \frac{1 - F(s)}{2}\right] ds \tag{A.9}$$

otherwise.

Despite the algebraic complexity of these expressions, they all involve sums and products of continuous functions. Thus we conclude that $U(\omega, r'; r)$ is jointly continuous. Payoff usc follows.

We now wish to show that this class of mechanisms is non-redundant.

Claim : $U(\omega, r'; r)$ is weakly decreasing in r' for all (ω, r).

Proof. Take the case where $r < r'$. The argument for the other case is analogous. Let $r'' > r'$. Then $y(r, r') < y(r, r'')$ and for $\omega \in (r, y(r, r'))$

$$U(\omega, r'; r) - U(\omega, r''; r) =$$

$$\frac{F(y(r, r'')) - F(y(r, r'))}{2}[\omega - r] \geq 0$$

using (A.6). When $\omega > y(r, r'')$ the corresponding expression is

$$U(\omega, r'; r) - U(\omega, r''; r) =$$

$$\frac{F(y(r, r'')) - F(y(r, r'))}{2}[\omega - r]$$

$$+ \int_{y(r,r')}^{y(r,r'')} \frac{F(y(r, r'')) - F(s)}{2} ds \geq 0.$$

A similar expression applies to the case where $y(r, r') \leq \omega \leq y(r, r'')$. ■

Then $\sup_{r'} U(\omega, r'; r) = U(\omega, 0; r) =$

$$\left[1 - \frac{1 - F(y(r, 0))}{2}\right][y(r, 0) - r] + \int_{y(r,0)}^{\omega} \left[1 - \frac{1 - F(s)}{2}\right] ds.$$

For $r^1 > r$, $U(\omega, 0, r) - U(\omega, 0; r^1) =$

$$\left[1 - \frac{1 - F(y(r, 0))}{2}\right][y(r, 0) - r] - \left[1 - \frac{1 - F(y(r^1, 0))}{2}\right]\left[y\left(r^1, 0\right) - r^1\right]$$

$$+ \int_{y(r,0)}^{y\left(r^1,0\right)} \left[1 - \frac{1 - F(s)}{2}\right] ds.$$

To see that this expression is positive, note that the first and third term together yield the payoff of a buyer of type $y(r^1, 0)$ when the reserve price chosen by the seller is r (and the other seller offers the reserve price 0). In the continuation equilibrium, this is equal to $U(y(r^1, 0), r; 0)$ because this buyer must be just indifferent between the two options. The second term is $U(y(r^1, 0), 0; r^1) = U(y(r^1, 0), r^1; 0)$ again, since the buyer must be indifferent between the two options in a continuation equilibrium. Thus the expression is equal to

$$U\left(y\left(r^1, 0\right), r; 0\right) - U\left(y\left(r^1, 0\right), r^1; 0\right),$$

which is strictly positive by Claim A.2.

Thus the payoff functions $U(\cdot, 0; r)$ and $U(\cdot, 0; r^1)$ are distinct for any distinct pair r and r^1. As in the other examples of indirect mechanisms, any two auctions can be separated by the 'first level' σ-algebra Σ_1, that is, by looking at suprema of payoff functions with respect to the other firm's mechanism. In particular, non-redundancy is proven.

B. APPENDIX: Proof for 'Space of USC Functions'

Proof of Theorem 6.2: (a) Routine.
(b) Denote the indicated mapping by Φ. The Maximum Theorem [1, Lemma 14.29] implies that $\Phi(g)$ is usc. For the continuity of Φ, suppose that

$$\Phi(g) \in \mathcal{N}_1 \equiv \{h_1 \in \mathcal{U}(C_1 \times C_3) : \sup_{G_1 \times G_3} h_1 > \kappa\},$$

for open sets $G_1 \subset C_1$ and $G_3 \subset C_3$. Then

$$g \in \mathcal{N} \equiv \{h \in \mathcal{U}(C_1 \times C_2) : \sup_G h > \kappa\},$$

where $G = G_1 \times \xi(G_3)$, an open set. Moreover, $\Phi(\mathcal{N}) \subset \mathcal{N}_1$.

If the neighborhood of $\Phi(g)$ is of the form

$$\Phi(g) \in \mathcal{N}_1 \equiv \{h_1 \in \mathcal{U}(C_1 \times C_3) : \sup_{K_1 \times K_3} h_1 < \kappa\},$$

for compact sets K_1 and K_3, then $\Phi(\mathcal{N}) \subset \mathcal{N}_1$ where

$$g \in \mathcal{N} \equiv \{h \in \mathcal{U}(C_1 \times C_2) : \sup_{K_1 \times K_2} h < \kappa\},$$

and $K_2 = \xi(K_3)$, a compact set.
(c) Let $g^\alpha \longrightarrow g$, $C^\alpha \longrightarrow C$, $g^\alpha(C^\alpha) \geq \kappa$ all α and show that $g(C) \geq \kappa$. For any open neighborhood G of C, $\exists \alpha_0$, $C^\alpha \in G$ for $\alpha > \alpha_0$. Let $\overline{G}$ denote the closure. Then $\forall \alpha > \alpha_0$, $\sup_{\overline{G}} g^\alpha \geq g^\alpha(C^\alpha) \geq \kappa$. By the nature of convergence in $\mathcal{U}(C)$, it follows that $\sup_{\overline{G}} g \geq \kappa$. Any compact Hausdorff space is normal. Therefore, the relatively compact neighborhoods G of C define a directed set D such that $C^G \longrightarrow C$ and $g(C^G) = \sup_{\overline{G}} g \geq \kappa$. Conclude that $g(C) \geq \kappa$ because g is usc. ■

C. APPENDIX: Proof of Revelation Principle

Prove Theorems 4.1 and 5.2 simultaneously. Notation from Section 6 is used freely. In particular, recall the definition in Section 6.2 of the message spaces C_n.

First, we elaborate on the nature of $\mathcal{T}$ and Ψ by describing $\Psi^{-1}(\cdot) : \mathcal{U}(\Omega \times \mathcal{T}) \longrightarrow \mathcal{T}$.

Lemma C.1. *For each* $V \in \mathcal{U}(\Omega \times \mathcal{T})$, $\Psi^{-1}(V) = (\Theta_n(V)) \in \Pi_0^\infty \mathcal{U}(C_n)$, *where*

$$\Theta_0(V)(\omega) = \sup\{V(\omega, t') : t' \in \mathcal{T}\},$$

$$\Theta_n(V)(\omega, g_0', ..., g_{n-1}') = \sup_{t'}\{V(\omega, t') : \Theta_k(\Psi(t')) = g_k',\ k < n\}.$$

Proof. In light of (6.3), this is a special case of [4, Theorem D.1]. An interpretation is that $\Theta_n(V)$ is a projection of V onto $\Pi_0^n \mathcal{U}(C_i)$. ■

Lemma C.2. *(a)* $\{U(\cdot; m) : m \in \theta(\Gamma)\} \subset \mathcal{U}(\Omega \times \Gamma)$ *is compact; (b)* $\theta(\Gamma)$ *is compact; (c)* $\{\sup_{\gamma'} U(\cdot, \gamma'; m) : m \in \theta(\Gamma)\}$ *is compact in* $\mathcal{U}(\Omega)$.

Proof. By payoff usc, $\gamma \longmapsto U(\cdot; \theta(\gamma)) \in \mathcal{U}(\Omega \times \Gamma)$ is continuous. Thus (a) follows from the compactness of Γ. Part (c) follows from Theorem 6.2. Part (b) is a consequence of the weak topology used for the space of mechanisms (Section 5). ■

Let $(\Gamma, \tilde{c}, \tilde{\pi})$ be given and construct a suitable $e : \Gamma \longrightarrow T$. This is done in several stages.

Define $\Phi : \theta(\Gamma) \longrightarrow \Pi_0^\infty \mathcal{U}(C_n)$, where $\Phi = (\Phi_n)_0^\infty$, by

$$\Phi_0(m)(\omega) = \sup\{U(\omega, \gamma'; m) : \gamma' \in \Gamma\},$$

$$\Phi_n(m)(\omega, g_0', ..., g_{n-1}') = \sup\{U(\omega, \gamma'; m) : \exists m', \theta(\gamma') = m', \Phi_k(m') = g_k', k < n\}.$$

Lemma C.3. $\Phi : \theta(\Gamma) \longrightarrow \mathcal{T}$ *and* Φ *is continuous.*

Proof. Recall that $T \subset \mathcal{T} \subset \Pi_0^\infty \mathcal{U}(C_n)$. Prove by induction that

$$\Phi_n : \theta(\Gamma) \longrightarrow \mathcal{U}(C_n) \text{ continuously.}$$

Theorem 6.2 is used repeatedly.

$n = 0$: The only complication is in the proof that Φ_0 is continuous: Let $\{m^\alpha\}_{\alpha \in I}$ be a net in $\theta(\Gamma)$ converging to m. Given the weak topology on $\theta(\Gamma) \subset \mathcal{M}(\Omega^2 \times \Gamma)$, this yields

$$U(\cdot; m^\alpha) \longrightarrow U(\cdot; m) \text{ in } \mathcal{U}(\Omega \times \Gamma).$$

This implies (by Theorem 6.2(b)) that

$$\sup_{\gamma'} U(\cdot, \gamma'; m^\alpha) \longrightarrow \sup_{\gamma'} U(\cdot, \gamma'; m) \text{ in } \mathcal{U}(\Omega),$$

proving continuity of Φ_0.
$n > 0$: Assuming that Φ_k is continuous for each $k < n$, Theorem 6.2(a) and (b) deliver the desired conclusion for Φ_n.
Claim: $t = (\Phi_n(m))_0^\infty \in \mathcal{T}$ for any $m = \theta(\gamma)$: Recall the outline of the proof of Theorem 6.3. It suffices to show that

$$\Psi(t) = 0 \text{ on } \Omega \times (\mathcal{T}_0 \backslash \mathcal{T}), \tag{C.1}$$

because then $\Psi(t) \in \mathcal{U}(\Omega \times \mathcal{T})$ and thus $t \in \mathcal{T}$.

From (6.7), $\Psi(t)(z) = \inf_N \Phi_N(m)(z^N)$, where $z \in \Omega \times \Pi_{n=0}^\infty \mathcal{U}(C_n)$ and $z^N \equiv \pi^N(z)$ equals the projection of z onto $\Omega \times \Pi_{n=0}^{N-1} \mathcal{U}(C_n)$. It follows, therefore, from the recursive definition of the functions Φ_n, that

$$\Psi(t)(z) = 0 \text{ if } z \notin \cap_{N=0}^\infty \left(\Omega \times \pi^N(\mathcal{T}_N) \times \Pi_{N+1}^\infty \mathcal{U}(C_n)\right).$$

But

$$\cap_{N=0}^\infty \left(\Omega \times \pi^N(\mathcal{T}_N) \times \Pi_{N+1}^\infty \mathcal{U}(C_n)\right) \subset \Omega \times \cap_{N=0}^\infty \mathcal{T}_N \;=\; \Omega \times \mathcal{T}, \tag{C.2}$$

proving (C.1). (The routine proof of (C.2) exploits the fact that $\mathcal{T}_N$ is compact and declines with N. The weak set inclusion is actually an equality.) ■

The following lemma is a special case (by (6.3)) of [13, p. 55]:

Lemma C.4. *Let $\varphi : (X, \mathcal{B}(X)) \longrightarrow (\mathcal{U}(S), \mathcal{B}(\mathcal{U}(S)))$, where X is an arbitrary topological space and S is locally compact and separable (e.g., compact metric). Then φ is measurable if and only if the map from X into $[0,1]$ defined by*

$$x \longmapsto \sup_{s \in A} \varphi(x)(s),$$

is measurable for each $A \in \mathcal{B}(S)$.

Lemma C.5. *Define e as the composition $\Phi \circ \theta$. Then e is continuous, one-to-one and $e(\Gamma) \subset T$.*

Proof. Φ is continuous on $e(\Gamma)$ by a previous lemma. θ is continuous because the continuation equilibrium is payoff usc and the weak topology is used for the space of direct mechanisms (see the discussion following Theorem 5.2).

Let $\{\Sigma_n\}$ be the σ-algebras on Γ appearing in the definition of NR. Recall again that $\mathcal{T} \subset \Pi_0^\infty \mathcal{U}(C_n)$, where $C_0 = \Omega$ and $C_n = C_{n-1} \times \mathcal{U}(C_{n-1})$. Denote by $\mathcal{B}_n$ the Borel σ-algebra on $\mathcal{U}(C_n)$. By Lemma C.4, Σ_n is the weakest σ-algebra such that $\Phi_{n-1} \circ \theta : \Gamma \longrightarrow (\mathcal{U}(C_{n-1}), \mathcal{B}_{n-1})$is measurable. Therefore, non-redundancy implies that $e = \Phi \circ \theta : \Gamma \longrightarrow \Pi_0^\infty \mathcal{U}(C_n)$ is one-to-one.

Define $\overline{\xi} : \mathcal{A}^{\Omega^2 \times \Gamma^2} \longrightarrow \mathcal{A}^{\Omega^2 \times \mathcal{T}^2}$ by

$$(\overline{\xi} m)(\omega, \omega', t', t'') \equiv \begin{cases} m(\omega, \omega', e^{-1}(t'), e^{-1}(t'')) & \text{if } t', t'' \in e(\Gamma) \\ \underline{a} & \text{otherwise.} \end{cases}$$

Claim: $\Psi\left(\Phi \circ \theta(\gamma)\right) = U(\cdot; \overline{\xi} \circ \theta(\gamma)), \quad \gamma \in \Gamma$: By Lemma C.1, it suffices to show that $\Phi \circ \theta(\gamma) = \Theta\left(U(\cdot; \overline{\xi}\theta(\gamma)\right)$. But this is readily verified by applying the equality $U(\cdot, \gamma'; \theta\gamma) = U(\cdot, e\gamma'; \overline{\xi}\theta\gamma)$, $\gamma' \in \Gamma$.

Show that $e(\Gamma) \subset T$: Recall the definition (6.11), whereby $T = \cap Y^k$. Let $\gamma \in \Gamma$. Then, by the claim, $\Psi(e\gamma) = \Psi\left(\Phi \circ \theta(\gamma)\right) = U(\cdot; \overline{\xi} \circ \theta(\gamma))$, implying that

$$e(\Gamma) \subset Y^0.$$

Show next that $e(\gamma) \in Y^1$: First, $\overline{\xi}\theta(\gamma) \in \mathcal{M}^0(\Omega \times \mathcal{T})$, because incentive compatibility and the other constraints that define the latter set of mechanisms are inherited from the corresponding properties of $\theta\gamma$, a direct mechanism over $\Omega \times \Gamma$. (The fact that $e(\Gamma) \subset Y^0$ is also relevant here.) Second, $\Psi(e\gamma) \in \mathcal{U}(\Omega \times \mathcal{T} \mid \Omega \times Y^0)$ if $\Psi(e\gamma)(\cdot) = 0$ on $\Omega \times (\mathcal{T} \setminus Y)$ for some compact $Y \subset Y^0$. Let $Y = e(\Gamma)$. Then Y is compact and, by above, $Y \subset Y^0$. Moreover, $t' \in \mathcal{T} \setminus Y \implies \overline{\xi}(\theta\gamma)(\cdot, t') = \underline{a} \implies \Psi(e\gamma)(\cdot, t') = 0$. Conclude that $e(\gamma) \in Y^1$. The proof may be completed by induction. ■

Turn to the remaining assertions in the theorems.

For the map ξ in Theorem 4.1, exploit the fact that $e(\Gamma) \subset T$, the domain of ψ, and let

$$\xi\left(\theta(\gamma)\right) = \psi\left(e(\gamma)\right), \gamma \in \Gamma; \tag{C.3}$$

ξ can be defined arbitrarily on $\mathcal{A}^{\Omega^2 \times \Gamma^2} \setminus \theta(\Gamma)$. Then the diagram in Theorem 4.1 commutes by construction. (Note that ξ generally differs from $\overline{\xi}$ as defined in the preceding proof, though $\xi(\theta\gamma)$ and $\overline{\xi}(\theta\gamma)$ imply the same payoff functions. Note also that ξ is one-to-one on $\psi(T)$.)

$\psi(T)$ is metric: Because Ψ and ψ are one-to-one, so also is the mapping from $\psi(T)$ into $\mathcal{U}(\Omega \times \mathcal{T} \mid \Omega \times T)$ that takes $m = \psi(t)$ into $\Psi(t)$. Using the weak topology for mechanisms defined following Theorem 5.2, this map is a homeomorphism. Finally, $\mathcal{U}(\Omega \times \mathcal{T} \mid \Omega \times T) \subset \mathcal{U}(\Omega \times \mathcal{T})$ is metric.

Turn to the continuation equilibrium described in Theorem 4.1(b)-(c). View $\psi(T)$ as a feasible set of indirect mechanisms using message space $C = \Omega \times T$. We constructed T and ψ so that $\psi(T) \subset \mathcal{M}(\Omega^2 \times T^2)$ (see Section 6.3). Thus every mechanism in $\psi(T)$ is payoff usc, strongly measurable and incentive compatible. Adopt the notation in (2.2) and express any indirect mechanism $\gamma = \psi(t)$ in the form

$$\gamma(\cdot) = (\gamma_c(\cdot), \gamma_{\pi_1}(\cdot), \gamma_{\pi_2}(\cdot)).$$

By strong measurability, the recommended probabilities satisfy $\gamma_{\pi_1}(\cdot) = \gamma_{\pi_2}(\cdot) = \gamma_\pi(\cdot)$ and the latter can be viewed as a function of (ω, t'). The candidate continuation equilibrium strategies are defined by:

$$c^*(\omega, t, t') = (\omega, t') \text{ and } \pi^*(\omega, t, t') = \gamma_\pi(\omega, t'),$$

where $\gamma = \psi(t)$.

We show that this defines a suitable continuation equilibrium. For any deviation $c^r = (\omega^r, t^r)$,

$$\int u\left(\gamma_c\left(c^r, c^*(\omega', t, t')\right), \pi^*(\omega, t, t'), \pi^*(\omega', t, t'), \omega\right) dF(\omega') =$$

(because $u(\cdot, \pi_1, \pi_2)$ is independent of π_1)

$$\int u\left(\gamma_c\left(c^r, c^*(\omega', t, t')\right), \pi^*(\omega^r, t, t^r), \pi^*(\omega', t, t'), \omega\right) dF(\omega') =$$

(by definition of π^*)

$$\int u\left(\gamma_c\left(c^r, c^*(\omega', t, t')\right), \gamma_{\pi_1}(\omega^r, t^r), \gamma_{\pi_2}(\omega', t'), \omega\right) dF(\omega') =$$

(because the message space is $\Omega \times T$)

$$\int u\left(\gamma\left(\omega^r, t^r, \omega', t'\right), \omega\right) dF(\omega') \leq \int u\left(\gamma\left(\omega, t', \omega', t'\right), \omega\right) dF(\omega').$$

The last inequality follows from the incentive compatibility of the mechanism γ.

To see that the choice strategy constitutes a continuation equilibrium relative to (γ, γ'), where $\gamma = \psi(t)$ and $\gamma' = \psi(t')$, observe that

$$\pi^*(\omega, t, t') \int u\left(\left\{\gamma_c\left(c^*(\omega, t, t'), c^*(\omega', t, t')\right), \pi^*(\omega, t, t'), \pi^*(\omega', t, t')\right\}, \omega\right) dF(\omega') +$$

$$[1 - \pi^*(\omega, t, t')] \int u\left(\left\{\gamma_c\left(c^*(\omega, t', t), c^*(\omega', t', t)\right), \pi^*(\omega, t', t), \pi^*(\omega', t', t)\right\}, \omega\right) dF(\omega')$$

$$= \gamma_{\pi_1}(\omega, t') \int u\left(\gamma\left(\omega, \omega', t', t'\right), \omega\right) dF(\omega') +$$

$$(1 - \gamma_{\pi_1}(\omega, t')) \int u\left(\gamma'\left(\omega, \omega', t, t\right), \omega\right) dF(\omega') \geq$$

$$\max\left[\int u\left(\gamma\left(\omega, \omega', t', t'\right), \omega\right) dF(\omega'), \int u\left(\gamma'\left(\omega, \omega', t, t\right), \omega\right) dF(\omega')\right],$$

by incentive compatibility for γ. The optimality of the choice strategy can be seen by expanding the expression involving the maximum of the two functions.

Finally, prove that $(\psi(T), c^*, \pi^*)$ is non-redundant: Let $\{\Sigma_n\}$ be the sequence of σ-algebras on $\psi(T)$ as in the definition of NR. We have to show that they separate any distinct $\psi(t_1)$ and $\psi(t_2)$. But $\psi(t_1) \neq \psi(t_2) \Longrightarrow t_1 = (g_n^1) \neq (g_n^2) = t_2 \Longrightarrow$ (by Lemma C.1) $\exists N$ such that $\Theta_N(U(\cdot; \psi(t_1))(\cdot) \neq \Theta_N(U(\cdot; \psi(t_2))(\cdot)$. If $N = 0$, then $\psi(t_1)$ and $\psi(t_2)$ are separated by $\Sigma_1' \subset \Sigma_1$ defined as the weakest σ-algebra such that $\psi(t) \longmapsto \Theta_0(U(\cdot; \psi(t))(\cdot)$ is measurable. (In other words, modify the measurability constraint in the definition of NR to consider singleton sets $E \subset \Omega$ rather than all Borel measurable subsets.) Proceed by induction on N.

References

[1] C. Aliprantis and K. Border, *Infinite Dimensional Analysis*, Springer-Verlag (1994).

[2] A. Brandenburger and E. Dekel, Hierarchies of beliefs and common knowledge, *J. Econ. Theory* **59** (1993), 189–198.

[3] J. Dugundji, *Topology*, Allyn and Bacon (1966).

[4] L. Epstein and T. Wang, 'Beliefs about beliefs' without probabilities, *Econometrica* **64** (1996), 1343-1374.

[5] J. Harsanyi, Games with incomplete information played by 'Bayesian' players, *Management Science* **14** (1967), 159–182.

[6] A. Heifetz and D. Samet, Universal partition structures, Working Paper 26/93, Tel Aviv, (1993).

[7] W. Hildenbrand, *Core and Equilibria of a Large Economy*, Princeton University Press (1974).

[8] P. McAfee, Mechanism design by competing sellers, *Econometrica* **61** (1993), 1281–1312.

[9] J. Mertens and S. Zamir, Formulation of Bayesian analysis for games with incomplete information, *Int. J. Game Theory* **14** (1985), 1–29.

[10] D. Mortensen, *Job Search and Labor Market Analysis,* North Holland, Amsterdam (1986).

[11] T. Norberg, Random capacities and their distributions, *Probab. Th. Rel. Fields* **73** (1986), 281–297.

[12] T. Norberg and W. Vervaat, Capacities on non-Hausdorff spaces, mimeo, (1989).

[13] G. O'Brien and W. Vervaat, *Capacities, Large Deviations and LogLog Laws*, Birkhauser (1991).

[14] J. Peck, Competing mechanisms and the revelation principle, *Ohio State University Manuscript* (1995).

[15] M. Peters, A competitive distribution of auctions, University of Toronto manuscript, to appear in *R. Econ. Studies*(1994).

[16] M. Peters and S. Severinov, Limit equilibrium concepts for problems with competing mechanism designers, to appear in *J. Econ. Theory*, (1993).

[17] S. Salop, *Practises that Credibly Facilitate Ologopoly Coordination*, Cambridge, MIT Press (1986).

[18] D. Schmeidler, Subjective probability and expected utility without additivity, *Econometrica* **57** (1989), 571–587.

[19] S. Vassilakis, Some economic applications of Scott domains, *Math Soc. Sciences* **24** (1992), 173-208.

[20] S. Vassilakis, Rules for changing the rules, Working Paper 265, U. Pittsburgh (1990).

[21] R. Wilson, Incentive Efficiency of Double Auctions,*Econometrica* **53** (1985), 11012–1117