

A Revelation Principle for Competing Mechanisms

Epstein, Larry G. and Michael Peters

Working Paper No. 452
July 1998

University of
Rochester

A Revelation Principle for Competing Mechanisms

Larry G. Epstein, Michael Peters

Rochester Center for Economic Research
Working Paper No. 452

July 1998

A REVELATION PRINCIPLE FOR COMPETING MECHANISMS

Larry G. Epstein Michael Peters*

July 10, 1998

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.

1. INTRODUCTION

Mechanism design problems are solved by restricting mechanism designers to *direct mechanisms* that assign outcomes to agents' reports about their private information. This approach is based on the *revelation principle* which states that for

*Epstein is at University of Toronto and Department of Economics, Harkness Hall, University of Rochester, Rochester, NY 14627-0156, USA - email epstein@chass.utoronto.ca. Peters is at Department of Economics, University of Toronto, 150 St. George Street, Toronto, Canada, M5S 3G7, peters@chass.utoronto.ca. We gratefully acknowledge the financial support of the Social Sciences Humanities Research Council of Canada and the comments of John Geanakoplos, Tom Gresik, Aviad Heifetz, Jim Peck and seminar participants at Columbia, U. Montreal, Rochester and Yale.

every indirect mechanism, there exists a direct mechanism that (induces truthful reporting and) produces the same outcomes. In other words, the class of direct mechanisms forms a *universal class*. Despite the important insights that the revelation principle generates into problems with asymmetric information, questions have been raised about its usefulness in environments where there are multiple principals (sometimes referred to as sellers in what follows) competing for one or more agents (sometimes called buyers). A series of examples presented in Peck [20] and in Martimort and Stole [13] illustrate apparent failures of the standard revelation principle. In their examples, allocations that are supported as equilibria with indirect mechanisms are not supported when sellers are restricted to using direct mechanisms where buyers report only their private valuations and do so truthfully. In addition, [13] provides an instance of an equilibrium relative to such direct mechanisms that is *not robust* to the possibility that sellers might deviate to more complicated mechanisms, illustrating another limitation of direct mechanisms that is specific to the competitive setting.

The reason for such failures stems from the fact, pointed out in [14] and [11], that in a multi-principal environment agents possess private information not only about their own preferences or valuations, but also about what different principals are doing, that is, about what is happening in the market. Moreover, it is important that such market information be included in the agent's type. When the principal attempts to make use of this market information, he is essentially designing his mechanism in a way that makes it responsive to what other mechanism designers are doing.

An analogous problem arises in the discussion of 'meet the competition clauses' [23] in the industrial organization literature. In a typical Bertrand price competition between two firms, they bid down the price until it equals marginal cost. This changes if firms are allowed to offer prices along with promises to match a competitor's price if the latter is lower. In that case, the monopoly price for both firms is an equilibrium because a firm considering deviating by lowering price realizes that this will simply force the other firm to cut price as well, resulting in no new customers. This 'meet the competition' argument illustrates the essential problem. The monopoly outcome cannot be supported when firms are restricted to direct mechanisms where buyers report only private valuations, because these rule out the possibility that firms might write contracts that make their price offers respond to what other firms are doing.

In principle it is clear how to deal with this - simply incorporate market information into an agent's type. However, there are some serious obstacles to doing

so that we now outline.

An obvious way for a seller to learn his competitor's price is to ask buyers to report it at the same time that they report their preference information. However, limiting the seller to this specific form of price matching is restrictive. For example, the seller might wish to make his price depend also on whether or not his opponent has made his price depend on . . . and so on. Market information seems to involve an infinite regress that must be resolved.

The problem of infinite regress associated with a type is now familiar from the work of [15] or [4]. The hierarchy of dependencies that arises here is outwardly similar to the hierarchies of beliefs that they study. The infinite regress is basic, but it is only part of the problem in our setting. Logically prior is the question of how to describe the competitor's mechanism. A description based on the fact that price depends on whether price depends on whether . . . is inadequate because restricting the language to prices is itself an ad hoc restriction. We are seeking a *universal language*, one that is sufficiently rich to permit descriptions of mechanisms in a large class that is not limited by preconceived notions of the nature of competition. In contrast, probability measures provide the obvious tool for describing beliefs, which form the essence of a type in the setting of [15] and [4].

In the absence of any obvious way to deal with these problems, the literature has responded by imposing ad hoc restrictions on the set of indirect mechanisms from which sellers can choose. The literature on competing mechanisms [14, 21, 22] restricts sellers to direct mechanisms in which buyers report only private information about their preferences. Competition in price schedules is the common assumption in the financial literature [7, 2, 12] and in the industrial organization literature [24, 3]. At first glance this does not seem unreasonable. It is natural to model sellers as competing in price when it is prices that are actually observed. However a complete positive theory needs to explain *why* sellers compete the way that they do despite the fact that more imaginative mechanisms are available to them. In some cases, it might be argued that institutional constraints justify the *a priori* restrictions on feasible mechanisms. However, even when there is a law that explicitly restricts the set of mechanisms that sellers can use, it is impossible to evaluate the impact of such a law without knowing what would happen without it.

This leads finally to the contribution of this paper. We construct a language for describing mechanisms that provides a way to incorporate private market information into an agent's type. This language is the key to the specification of a class of mechanisms having the property that any well behaved set of indirect

mechanisms can be embedded within it. In this sense any ad hoc model of competition among mechanism designers can be viewed as a model that restricts sellers to offering mechanisms that lie in a subset of this universal class. This provides a natural way of thinking about the apparent restrictiveness of the usual sort of direct mechanisms, since they constitute a relatively small subset of the universal class. The nonrobustness of equilibria in direct mechanisms and the failures of the standard revelation principle illustrated by the examples in [20] and [13] (as well as the examples that we give below) then become transparent.

Furthermore, the existence of this universal class of mechanisms makes it possible to show the sense in which the revelation principle *does* hold. We show that equilibria relative to the universal class are robust in the sense that there are no profitable deviations to more complicated mechanisms, and that all robust equilibria can be represented as equilibria relative to the universal class. Thus equilibria relative to the set of universal mechanisms can never give rise to the problems identified by Martimort and Stole [13] or Peck [20]. Mechanisms in the universal class ask buyers to report their type (including market information), and in this sense they are ‘direct’ mechanisms. We show that there is no loss of generality in restricting sellers to this universal class of direct mechanisms.

At a less formal level, we can describe our result as follows. The set of robust equilibrium allocations constitute all feasible allocations when sellers compete with one another and are free to communicate and contract with the buyers who visit them. One way to interpret this is to think of sellers’ activities as uncontractable moral hazard variables as in Myerson [16]. A central planner might try to make incentive compatible recommendations to sellers after collecting buyer private (preference) information. Incentive compatibility means that no seller has any incentive to deviate from the recommendations provided that he expects the other seller to conform. If the planner recommends only ‘actions’ to the sellers, then the revelation principle can fail in two ways - there will be allocations that appear incentive compatible to the planner that will not be, and there will be equilibrium allocations that do not look incentive compatible from the planner’s point of view. These difficulties are resolved and the revelation principle holds in its usual form if the planner recommends mechanisms from the universal set of mechanisms to the sellers instead of simple actions.

To illustrate some of the above, fix ideas and clarify later notation, we begin by presenting two finite examples. (An economically more meaningful example, involving competition in price matching mechanisms, is provided in Appendix A.) Our formulation of the competing mechanisms problem is contained in Section 3.

Section 4 provides the statement of our main result (Theorem 4.1), which verifies the existence of a universal set of mechanisms. A discussion of robustness and the revelation principle follows in Section 5. Section 6 provides some intuition for the nature of types, that is, the language for describing mechanisms; this is accompanied by an intuitive sketch of the proof of Theorem 4.1. The complete proof may be found in a series of appendices. Section 7 concludes with an outline of some extensions.

2. FINITE EXAMPLES

As mentioned, the difficulties that arise in a multi-principal environment are illustrated by examples in [13] and [20]. But the former deals with a problem of competing agency, which differs from the problem of competing mechanisms as explained later; and Peck's example, while it is closer to our framework, is quite complicated. Therefore, we present a pair of examples that illustrate the special difficulties for the setting of competing mechanisms. They serve also to clarify later notation and terminology.

We begin with a description of the canonical 'competing mechanism' problem. This description applies to the examples given in this section as well as the formal model described later on. There are two sellers and two buyers. Sellers simultaneously offer mechanisms to buyers in the first stage of the game. Each mechanism consists of a message space and a mapping that converts messages sent by buyers into actions (later called 'simple actions'). In the second stage, buyers send messages to sellers, and these determine the action that each seller takes. In the third stage, which occurs before buyers observe the actions that result from their communications or the messages sent by the other buyer, each buyer chooses one of the sellers.¹

Sellers are restricted to two actions named A and B . Preferences are common knowledge. Therefore, a 'direct mechanism' of the usual sort is simply an action.

¹It may seem more natural to have buyers communicate with sellers after committing themselves and to allow sellers to vary the simple action taken with the number of participating buyers. The examples could be modified to include these features, but at the cost of additional complexity. Our formal analysis permits the latter feature; for the former see Section 7.2.

2.1. Example 1: Simple Price Competition and Nonrobustness

The first example illustrates the notion of 'robustness' that is central later in the paper. The idea is that equilibria in direct mechanisms may have the property that sellers can profitably deviate to more complicated mechanisms. The example illustrates the way that sellers can exploit buyers 'market' information.

In this example, sellers compete in price. The seller can offer a high price (A) or a low price (B). If more than one buyer selects the same seller, the seller chooses one of the buyers at random as his trading partner. Payoffs depend on both the action taken and on the number of participating buyers, as summarized in the following table:

| | | | |
|-----|---|------|--------|
| | 0 | 1 | 2 |
| A | 0 | 2, 1 | 2, 1/2 |
| B | 0 | 1, 2 | 1, 1 |

The columns differ according to the number of buyers who offer to trade with the seller. From the first column, sellers receive 0 payoff if there are no buyers. For other columns, the two entries in any cell give the seller's payoff (first entry) and the payoff to each participating buyer. For example, if the seller uses the high price (A), he gets the payoff 2 if at least one buyer offers to trade. Buyers who offer to trade with the seller get payoff 1 if they are alone and payoff 1/2 if the other buyer also participates, because there is only probability 1/2 to be chosen as the seller's trading partner.

Suppose that sellers compete in direct mechanisms (as in [14]), which here means that each seller offers either the high or the low price. To proceed, we need to specify the 'continuation equilibrium' associated with these mechanisms, that is, we need to specify how buyers select between sellers for each pair of options (A or B) on offer. Focussing on a symmetric equilibrium where each buyer uses the same selection strategy, it is easy to show that there is a unique continuation equilibrium and that it involves buyers randomizing equally between sellers when they offer the same price and choosing the low priced seller with probability 1 when the prices offered are different.

This buyer behavior generates the following table of payoffs:

| | | |
|-----|---------------|---------------|
| | A | B |
| A | 3/2, 3/2, 3/4 | 0, 1, 1 |
| B | 1, 0, 1 | 3/4, 3/4, 3/2 |

The first entry in each cell is the payoff to the seller who chooses the row, the

second entry is the payoff to the seller who chooses the column and the final payoff is the common payoff to each of the buyers. Ignoring the third entry for a moment, this box represents the normal form game that is played by sellers when they choose direct mechanisms in the first stage of the game. They understand the way that changes in their mechanisms affect buyers' behavior and incorporate this into their calculations.

There are two pure strategy equilibria for this game in which sellers offer the same price and buyers randomize between them. When sellers use the high price, a deviating seller realizes that he will capture the whole market by cutting price, but he does not find this profitable. When they use the low price, a deviating seller realizes that he will lose all his customers by raising price, which makes this unprofitable.

At this point, we are more interested in the mixed strategy equilibrium where each seller offers the high price with probability $3/5$ and receives the equilibrium payoff $9/10$. We can show that this equilibrium is *nonrobust* in the sense that seller 1 can improve his payoff strictly by a unilateral deviation to a more complicated mechanism that works as follows: Each buyer must send one of two messages, labelled s and t . If both buyers send s , then the seller charges the low price B . In response to any other combination of messages, the high price A is charged. The expected payoff from use of this mechanism depends on buyer behavior, or the 'continuation equilibrium'. The latter must specify both a participation strategy that describes the likelihoods with which buyers select each seller, and a communication strategy, describing the messages that buyers send.

Consider the following continuation equilibrium: If seller 2 offers the mechanism A , both buyers report t to seller 1 and randomize equally between the two sellers. If seller 2 offers B , buyers both report s to seller 1 and randomize equally between sellers. These strategies are best responses for either buyer given that they are used by the other buyer.² Given this buyer behavior, seller 1 charges the high (low) price when seller 2 does and this produces the payoff $6/5 > 9/10$. Thus the deviation is strictly profitable for seller 1 and the mixed strategy equilibrium allocation is supported only because of an ad hoc restriction on the set of contracts feasible for sellers. It is evident that the added profit is attained by making the choice of action depend on the action of seller 2. Seller 1 learns about

²If seller 2 takes action A , then neither buyer can unilaterally induce seller 1 to switch to B . If seller 2 offers B , both buyers report s to seller 1 who will offer B . Either buyer can induce seller 1 to switch to A by altering his message to t . But this deviation for a buyer only raises the price charged by seller 1.

the action of seller 2 by asking buyers about it.

2.2. Example 2: More complex mechanisms may generate new equilibrium allocations

The example in this section is simply a counterexample to the naive revelation principle, and is closely related to the examples in Peck [20] and Martimort and Stole [13]. We provide it to show that these counterexamples themselves are not the result of ad hoc restrictions on the set of feasible mechanisms. As we show later, the equilibrium in indirect mechanisms that cannot be reproduced when firms are only allowed to use direct mechanisms is a 'robust' equilibrium.

Payoffs are given in the following table (read as in the previous example):

| | | | |
|----------|-------|-------|-----------|
| | 0 | 1 | 2 |
| <i>A</i> | -1, 2 | -1, 1 | 3 + 2ε, 1 |
| <i>B</i> | -2, 0 | 0, 2 | 2, 0 |

Here ε is a small positive number. This problem is slightly richer than simple price competition. Mechanism *A* is good (bad) for the seller (buyer) when only one buyer chooses to participate, but bad (good) when both buyers participate. Mechanism *B* has the opposite property - the seller prefers both buyers to participate, while each buyer prefers to be the sole participant.

Suppose first that sellers are restricted to employing direct mechanisms, that is, to choosing an action *A* or *B*. The following table describes one possible continuation equilibrium (consisting exclusively of participation probabilities).

| | | |
|----------|-------------|-------------|
| | <i>A</i> | <i>B</i> |
| <i>A</i> | mix equally | both to 2 |
| <i>B</i> | both to 1 | mix equally |

This buyer behavior implies the following (expected) payoffs for each pair of actions:

| | | |
|----------|-------------|----------|
| | <i>A</i> | <i>B</i> |
| <i>A</i> | 1, 1, 1 + ε | -1, 2, 0 |
| <i>B</i> | 2, -1, 0 | 0, 0, 1 |

Focussing once again on the first two payoffs in each cell, there is a unique equilibrium in direct mechanisms where each firm offers the direct mechanism

B .³ However, the outcome corresponding to each seller using A can be supported through the use of more complex mechanisms. To see this, suppose that sellers are free to use indirect mechanisms in which buyers communicate messages from $\{s, t\}$. We will show that there is an equilibrium for this game in which each seller uses the mechanism m , where m specifies action A if both buyers report t and action B otherwise.

The continuation equilibrium (buyer behavior) that supports AA as an equilibrium play is given as follows: If both sellers offer m , then buyers report t to both sellers and randomize equally between them. (There is no incentive for either buyer to deviate from this by sending the message s to either seller. Such a message would change the action taken by the seller from A to B but would have no effect on the participation strategy used by the other buyer. Thus the new message would generate a payoff $1 < 1 + \epsilon$ from any seller who changes to action B .)

Suppose next that seller 2 offers m and seller 1 deviates to m' . If there exists a message pair (c, c') in $\{s, t\}^2$ for which m' prescribes the action A , then let the buyers communicate c and c' respectively to the deviator and each send t to the nondeviating seller 1. This leads to the action A by each seller. Consequently, buyers randomize equally between them. (As in the earlier continuation equilibrium, neither buyer has an incentive to induce the action B by either seller.) Finally, if m' delivers only the action B , let buyers send the message s to the nondeviator and randomize equally between sellers. This generates the payoff 0 for each seller and 1 for each buyer.

This description of buyer behavior is sufficient to permit examination of the profitability of a unilateral deviation from m . Any deviation generates either 1 or 0 as the deviator's payoff, neither of which exceeds the payoff received in the putative equilibrium outcome. We conclude that AA is supported as an equilibrium outcome by this continuation equilibrium. Whenever the deviator tries to induce the action B using some indirect mechanism in the feasible set, buyers send messages to the nondeviator that induce him to change actions as well, making a profitable deviation impossible. In this way, AA is supported along an equilibrium path even though it is not supported with competition in direct mechanisms. Thus, restricting mechanism designers to direct mechanisms

³One caution is in order here. From the table it appears that offering the direct mechanism A is a dominated strategy. This interpretation is not appropriate since the table represents a reduced version of a two-stage game that presumes a particular continuation equilibrium. There exist alternative choice strategies leading to A dominating B for the seller.

not only generates invalid predictions (in the sense that equilibria may not be robust, as in the previous example), it also makes it impossible to understand the equilibrium behavior that does occur. Our goal is to provide a space of mechanisms broad enough to ensure that these problems cannot occur.

3. INDIRECT MECHANISMS

3.1. 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 sellers simultaneously announce the mechanisms they plan to use. As is common in the search literature, we assume that buyers search out the market beforehand and consequently have better information than sellers. More particularly, neither seller observes directly the mechanism chosen by the other seller but buyers can observe both mechanisms.⁴ After seeing them, each buyer selects one of the firms. Once buyers have made their choices and these have been revealed to the sellers, then sellers’ mechanisms are played out with any participating buyers.

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 ‘simple’ actions

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

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

A simple action is a complete description of the allocation, including possibly randomization. In our example of price matching (Appendix A), a simple action

⁴Our model applies without change to the situation where the sellers can see each others’ offers but cannot write binding contracts on offers made by others. See Section 7.3 for the case where such binding contracts are feasible.

is a lottery over buyers and the option to buy at a specified price to be offered to the buyer that is ultimately selected. In an auction environment, a simple action might be a set of (randomized) transfers paid to and received from each bidder along with a specification of the probability with which each bidder is allocated the commodity. The independence assumption is made 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 condition their choice of simple actions on the participation decisions of buyers. Thus we are led to consider the space $(\mathcal{A}_0)^{P^2}$ of participation contingent simple actions. 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 $\mathcal{A} = (\mathcal{A}_0)^{P^2} \times [0, 1]^2$.

Sellers’ payoffs are represented by $v : \mathcal{A} \times \Omega^2 \rightarrow [0, 1]$, 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.

For buyers, payoffs are represented by the function $u : \mathcal{A} \times \Omega \rightarrow [0, 1]$. Interpret $u(a, \omega)$ as the expected payoff to a buyer (say buyer 1) with valuation ω who is participating at a given firm where action a is taken. It is computed prior to his learning if the other buyer is also participating there. By the definition of actions, each a in \mathcal{A} has the form $a = (a_c, \pi, \pi')$, where π and π' represent the respective probabilities with which buyers 1 and 2 choose the seller. Because u represents 1’s utility conditional on his already having chosen the seller, we assume that $u(a, \omega)$ is independent of the π -component of the action a . But it will in general depend on π' , because 1’s payoff ex post may depend not only on the simple action chosen but also on whether or not the other buyer is participating. Thus the likelihood of such participation is important ex ante.

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 $u(\underline{a}, \cdot) = 0$. The action \underline{a} may correspond to ‘no trade’, implying the default utility level 0 regardless of valuation. For example, a seller might choose a price that he knows no one could afford to pay. Because we assume that $u(\cdot) \geq 0$, the utility obtained in the absence of participation, it follows 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

also that $v(\underline{a}, \cdot) = 0$.

At this stage it is useful to point out the difference between our formulation and the better known problem of common agency [3], involving two (or more) sellers dealing with a single buyer whose payoffs depend on the actions of both sellers. In particular, the buyer's ranking of alternatives offered by one seller depend on the action selected by the other seller. This externality makes it possible to improve upon simple direct mechanisms in the common agency environment. In our formulation, the payoff that a buyer gets from one seller is independent of the other seller's action, but it depends on whether or not the other buyer chooses to participate with the same seller. The probability with which this occurs depends on the action taken by the other seller. This indirect dependence gives rise to the same sort of contractual externality that appears in common agency - the buyer's ranking of a menu of alternatives depends on the action taken by the other seller. The added complexity that arises in the competing mechanism problem is that this ranking of alternatives and the nature of the externality are not unique (as in common agency), because they may vary with the continuation equilibrium describing buyer behavior.

3.2. 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 is perfectly general in the sense of the degree and nature of the communication about competing mechanisms that it permits. 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

$$\gamma_c : C^2 \longrightarrow (\mathcal{A}_0)^{P^2} \quad \text{and} \quad (\gamma_{\pi_1}, \gamma_{\pi_2}) : C^2 \longrightarrow [0, 1]^2. \quad (3.1)$$

Thus, $\gamma_c(\cdot)$ describes the contingent simple action and $(\gamma_{\pi_1}(\cdot), \gamma_{\pi_2}(\cdot))$ describes the 'rest' of the action prescribed by the mechanism $\gamma(\cdot)$. The components $\gamma_{\pi_i}(\cdot)$ can be interpreted as the seller's recommended participation probabilities for buyer i .

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 where buyers communicate with sellers before committing themselves to one of the mechanisms. The alternative and common assumption ([14], 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 7.2 for further discussion.)

Denote by Γ the set of feasible indirect mechanisms, endowed with some topology. Unless specified otherwise, we assume below that Γ is compact metric.

Turn to behavior. A *communication strategy* \tilde{c} is a measurable mapping from $\Omega \times \Gamma^2$ into C , with the interpretation that $\tilde{c}(\omega, \gamma, \gamma')$ 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 $\tilde{\pi} : \Omega \times \Gamma^2 \rightarrow [0, 1]$, where $\tilde{\pi}(\omega, \gamma, \gamma')$ 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 $(\tilde{c}, \tilde{\pi})$ is a *continuation equilibrium* if no buyer has any incentive to deviate from either the reporting strategy \tilde{c} or the selection strategy $\tilde{\pi}$, for any of his valuations and for any pair of mechanisms offered by the two sellers. We assume the existence of continuation equilibria. We view this assumption as completely innocuous. Of course it is not difficult to construct models of indirect competition where continuation equilibrium do not exist (one mechanism might be - I will trade with the buyer who names the largest integer). It is also easy to think of models of indirect competition where sellers can offer mechanisms that do not make sense (each seller offers a price equal to the price offered by the other seller). There is no need to worry about whether such models are good descriptions of competition between sellers - it is immediately apparent that they are not. Thus in the discussion that follows we restrict attention to indirect models that possess all the usual desiderata of predictive models.

Of course this assumption does impose a restriction on the models Γ of indirect competition to which the analysis in this paper applies. The restriction is analogous to the restriction associated with the usual revelation principle in single mechanism designer problems, since indirect mechanisms in which do not possess equilibria cannot be replaced by direct mechanisms.

It should also be noticed that when we assign a particular continuation equilibrium $\tilde{c}(\cdot, \gamma, \gamma')$, $\tilde{\pi}(\cdot, \gamma, \gamma')$ to a pair of mechanisms, we are not assuming that it is unique. We view the value of a particular mechanism to be partly determined by the continuation equilibrium that it delivers. Thus the continuation equilib-

rium is part of the model of competition that we wish to understand. If there are multiple continuation equilibria, these will generate new models for which the set of indirect mechanisms can once again be embedded in our universal set of mechanisms.

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

The key to the standard (one principal) 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 $(\tilde{c}, \tilde{\pi})$ induces the mapping

$$m_\gamma : \Omega^2 \times \Gamma^2 \longrightarrow \mathcal{A}, \text{ where}$$

$$m_\gamma(\omega, \omega', \gamma', \gamma'') = (\gamma_c(\tilde{c}(\omega, \gamma, \gamma'), \tilde{c}(\omega', \gamma, \gamma'')), \tilde{\pi}(\omega, \gamma, \gamma'), \tilde{\pi}(\omega', \gamma, \gamma'')). \quad (3.2)$$

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 appropriate incentive compatibility restrictions on direct mechanisms. The dependence of the action chosen on the other firm's mechanism differentiates our setting from the more familiar single seller 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 $(\tilde{c}, \tilde{\pi})$.⁵ Then the seller who chooses the randomization δ receives the payoff

$$V(\delta; \delta', \tilde{c}, \tilde{\pi}) = \int v(m_\gamma(\omega, \omega', \gamma', \gamma''), \omega, \omega') dF(\omega) dF(\omega') d\delta'(\gamma') d\delta(\gamma). \quad (3.3)$$

⁵ $\Delta(\Gamma)$ denotes the space of Borel probability measures on Γ , endowed with the standard topology of weak convergence.

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

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

We impose symmetry on the strategies that buyers and sellers use in equilibrium purely for the sake of the notational simplification that symmetry permits.

Clearly equilibria of this kind depend on the specification of Γ including the message space C . Typically Γ and C are selected for reasons of tractability, both mathematical and economic. If one has data on prices, it is natural to want to formulate a model in which firms compete in prices. We are interested in analyzing the exact sense in which this might be restrictive.

4. A UNIVERSAL CLASS OF MECHANISMS

4.1. Additional Assumptions

Our objective is to show that any set Γ of indirect mechanisms (with a given continuation equilibrium) can be embedded in the class of universal mechanisms. We do this by showing that the actions produced by any pair of mechanisms in Γ can be replicated by an appropriate combination of mechanisms in the universal class. The continuation equilibrium for the latter features agents reporting their private information truthfully and obeying all participation recommendations made to them by sellers.

We impose two additional assumptions on continuation equilibria. Focus on a continuation equilibrium $(\Gamma, \tilde{c}, \tilde{\pi})$ and the corresponding function m_γ , defined by (3.2), that summarizes the actions produced by $\gamma \in \Gamma$. To express the assumptions on $(\Gamma, \tilde{c}, \tilde{\pi})$, introduce the payoff functions induced by mechanisms. To be precise, denote the expected utility of a buyer facing γ by

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

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

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

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

Upper semi-continuity (in fact continuity) of $U(\cdot; \gamma)$ in valuation alone is implied by a continuation equilibrium (this is well known - [25] for example). It follows that the condition (i) of payoff usc is innocuous if Γ is finite. More generally, it can be shown that a sufficient condition for payoff usc, including part (ii), is that $U(\cdot)$ be continuous on $\Omega \times \Gamma^2$.

The second restriction on continuation equilibria (called non-redundancy) is more difficult to explain. We provide a formal (and possibly impenetrable) definition of the property here and defer interpretation until Section 6, after we have shown what the assumption of non-redundancy delivers.

The formal definition follows: Given a continuation equilibrium $(\Gamma, \tilde{c}, \tilde{\pi})$, 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 $\gamma \mapsto \sup\{U(\omega, \gamma'; \gamma) : (\omega, \gamma') \in E\}$, where E varies over $\mathcal{B}(\Omega) \times \Sigma_0$, and $\Sigma_{n+1} = \sigma$ -algebra generated by the mappings $\gamma \mapsto \sup\{U(\omega, \gamma'; \gamma) : (\omega, \gamma') \in E\}$, 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 $k > n$. Say that $(\Gamma, \tilde{c}, \tilde{\pi})$ is *non-redundant* if any pair of distinct points in Γ can be separated by some Σ_n .⁶

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 simply “technical matter”.

4.2. The Main Result

Our main result is presented here. First, we explain some notation and terminology used in the theorem.

Mechanisms in the universal class resemble the usual sorts of direct mechanisms in that buyers are asked to report their private information directly. To do this, buyers must be able to describe the mechanism that is being used by the

⁶Our definition of non-redundancy is adapted from that in [15] and [9]. In the price matching example (Appendix A), the separation required by non-redundancy is achieved by the first-order σ -algebra Σ_1 .

other seller. This description must be adequate to describe every order in the hierarchy of dependencies built into the mechanism. It must also be free of ad hoc terminology, like price, since the mechanisms being described may not involve simple price offers. A major contribution of the theorem is to provide a suitable language, in the form of the set T . Buyers report their preference information by using an element of the set Ω and they describe their market information by using an element from T .

To clarify the sense in which T constitutes a language, denote by $\mathcal{A}^{\Omega^2 \times T^2}$ the set of all measurable maps $m : \Omega^2 \times T^2 \rightarrow \mathcal{A}$. 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. Since T is a language that can be used to describe such mechanisms, there is a one to one map $\psi : T \rightarrow \mathcal{A}^{\Omega^2 \times T^2}$. Interpret $\psi(t)$ as the direct mechanism that is *described* by $t \in T$. Thus T constitutes a language for describing direct mechanisms that have as inputs reports from this same language.

We have defined actions to include recommended probabilities. Thus the action $\psi(t)(\omega, \omega', t', t'')$ includes a recommended participation probability to the buyer with valuation ω , given that the other buyer has valuation ω' and that both buyers report the type t' for the other seller. Denote that recommended probability by $\psi(t)_\pi(\omega, \omega', t', t'')$, paralleling the notation in (3.1).

The set $\psi(T)$ can be viewed also as a set of *indirect* mechanisms, that is, a particular specification of Γ and one for which the message space C is $\Omega \times T$. This interpretation for $\psi(T)$ gives meaning to the theorem's reference to $(\psi(T), c^*, \pi^*)$, a continuation equilibrium relative to $\psi(T)$.⁷

We can now state our main result.⁸

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

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

$$\gamma(\tilde{c}(\omega, \gamma, \gamma'), \tilde{c}(\omega', \gamma, \gamma')) = m(c^*(\omega, e(\gamma), e(\gamma')), c^*(\omega', e(\gamma), e(\gamma'))) \text{ and}$$

⁷We make the obvious modification in previous formalism whereby strategies, including c^* and π^* , are defined and measurable on T , rather than on $\Gamma = \psi(T)$.

⁸Say that $e : \Gamma \rightarrow T$ 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 .

$$\gamma'(\tilde{c}(\omega, \gamma', \gamma), \tilde{c}(\omega', \gamma', \gamma)) = m'(c^*(\omega, e(\gamma'), e(\gamma)), c^*(\omega', e(\gamma'), e(\gamma))).$$

(b) For all $(\omega, \omega', t, t') \in \Omega^2 \times T^2$,

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

$$\pi^*(\omega, t, t') = \psi(t)_\pi(\omega, \omega', t', t').$$

We have explained the sense in which the space T constitutes a language. The theorem establishes the universality of that language, in that, under the conditions stated, indirect mechanisms in any given feasible set Γ can be described in terms of T by means of the translation represented by e . In particular, the same T applies for any continuation equilibrium $(\Gamma, \tilde{c}, \tilde{\pi})$ satisfying payoff usc and non-redundancy.⁹

The theorem also provides a continuation equilibrium (c^*, π^*) relative to the set $\psi(T)$ of indirect mechanisms. By part (a), the actions forthcoming in this equilibrium replicate those in the given equilibrium $(\Gamma, \tilde{c}, \tilde{\pi})$. (Because actions have been defined to include participation probabilities, the latter are also replicated.) This establishes that $\psi(T)$ is a sufficiently rich set of mechanisms. Part (b) states that the continuation equilibrium (c^*, π^*) has two natural properties - c^* involves truthful reporting of the other seller's type t' and the probability $\pi^*(\omega, t, t')$ with which the ω -buyer chooses the seller of type t when the other seller has type t' , coincides with the recommendation of the type t seller.¹⁰ A consequence is that any pair of mechanisms in $\psi(T)$, here viewed as a pair of direct mechanisms, one for each firm, can be implemented by the continuation equilibrium (c^*, π^*) . The parallel with the standard single-firm setting is apparent - $\psi(T)$ is the counterpart for our competitive setting of the familiar class of incentive compatible direct mechanisms based on reports about valuations alone that is the key to the standard revelation principle. (See Appendix C for more on the nature of the mechanisms in $\psi(T)$.)

We elaborate on the significance of the theorem in Section 5 and provide an intuitive outline of its proof in Section 6. Here we offer a brief comparison with

⁹For a given indirect mechanism γ , its translation $e(\gamma)$ is not unique in the following sense. If γ lies in both Γ_1 and Γ_2 and if the conditions of the theorem are satisfied so that there is an embedding e_i of Γ_i into T for each i , then $e_1(\gamma) \neq e_2(\gamma)$, in general. This is natural because the 'nature' of the mechanism γ depends on the context. Similarly, it is natural that, for given Γ , e depend on the continuation equilibrium $(\tilde{c}, \tilde{\pi})$ that is being considered.

¹⁰It follows from this equality that $\psi(t)_\pi(\omega, \omega', t', t')$ must be independent of ω' . See Appendix C for details.

the case of a single mechanism designer where the revelation principle appears tautological - ‘direct mechanisms’ can be constructed in a straightforward way by composing equilibrium reporting strategies with the rule that assigns actions to reports. Theorem 4.1 differs substantially from the single agent result because we cannot assume that the types describing buyers’ private information lie in some known and well behaved set. Instead this set of types must be constructed from ‘scratch’. Moreover, this construction is complicated by the special nature of the multi-principal setting. In the single mechanism designer problem, the belief hierarchy is a natural candidate for a ‘universal’ description of buyers’ private information only because it is independent of the modeler’s notion of what indirect mechanisms are available to the mechanism designer. Such independence is not given in our setting, because private information includes market information and this is expressed in terms of the modeler’s conception of the nature of competition. Independence from the modeler’s conception is restored by use of the universal language T , making its construction novel and a central contribution. This perspective may also help to explain why we need to adopt the non-redundancy restriction, though it has no counterpart in the construction of types based on beliefs.

4.3. Finitely Many Types

A possible concern with Theorem 4.1 is tractability. In problems with a single mechanism designer, the set of types is theoretically very complex, an infinite series of beliefs about beliefs to higher and higher orders, paralleling the complexity of the types space T upon which our universal class of mechanisms is based. Normative applications of the revelation principle usually come from making assumptions that make the types space simple. For example, buyers might have high or low marginal utility, or valuation information may be expressed as an interval on the real line.

In order to bolster confidence that our approach may prove useful in simple applied models of competing mechanism designers, we describe conditions on primitives that are sufficient to deliver the finiteness of T . Given the length of this paper, we content ourselves with illustrating the potential for simplification in plausibly interesting environments, rather than attempting to provide a general result. Thus we proceed under the assumption that there is no private information, that is, Ω is a singleton which can be suppressed in the notation. We also continue to assume that buyers behave symmetrically.

Consider the following natural specialization of buyers' payoff functions u . Let $u_0 : \mathcal{A}_0 \times P \rightarrow [0, 1]$, where $u_0(a_0, p')$ gives the payoff to a buyer participating at a mechanism that has produced simple action a_0 and where the other buyer's participation status is given by p' . Given any action $a = (a_c, \pi, \pi')$ in $\mathcal{A} = (\mathcal{A}_0)^{P^2} \times [0, 1]^2$, write

$$a_c = (a_c^{(1,0)}, a_c^{(0,1)}, a_c^{(1,1)});$$

$a_c^{(1,1)}$ denotes the simple action prescribed by the plan a_c if both buyers participate (that is, $(p, p') = (1, 1)$) and so on. Suppose finally that $u(a)$ is given by

$$u(a) = \pi' u_0(a_c^{(1,1)}, 1) + (1 - \pi') u_0(a_c^{(1,0)}, 0), \quad (4.2)$$

the expected payoff to the participating buyer when π' is the probability of the other buyer also participating.

Finiteness of T is implied if $u_0(\mathcal{A}_0 \times P)$ is finite (*a fortiori* if \mathcal{A}_0 is finite) and if we assume that for all simple actions a_0 and b_0 ,

$$u_0(a_0, 1) \neq u_0(b_0, 0). \quad (4.3)$$

Theorem 4.2. *Suppose that there is no private information, that buyers' payoff functions satisfy (4.2) and (4.3) and that $u_0(\mathcal{A}_0 \times P)$ is finite. Then the set of types T provided by Theorem 4.1 is finite.*

The proof is given at the end of Section 6. One drawback to finite action spaces is that they do not permit sellers to use randomized actions (though randomized *strategies* are permitted). This assumption may appear innocuous, but randomized actions have strong incentive effects when buyers are risk averse, making them desirable to sellers.

5. ROBUSTNESS AND THE REVELATION PRINCIPLE

If the restrictions imposed on the seller's ability to offer mechanisms are unreasonable, then the predictions 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 modeler. 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 principal 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.

Given a continuation equilibrium (Γ, c, π) and $\gamma \in \Gamma$, denote by m_γ the function defined in (3.2). Similarly, denote by $m_{\gamma_1}^1$ the function corresponding to the continuation equilibrium (Γ_1, c_1, π_1) , where γ_1 is an arbitrary mechanism in Γ_1 .

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

$$m_\gamma(\cdot, \gamma', \gamma') = m_{\alpha(\gamma)}^1(\cdot, \alpha(\gamma'), \alpha(\gamma')) \text{ on } \Omega^2. \quad (5.1)$$

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 .¹¹ As an example, if Γ is compact metric, then the continuation equilibrium $(\psi(T), c^*, \pi^*)$ provided by Theorem 4.1 extends (Γ, c, π) , with embedding $\alpha = \psi \circ e$.

Say that an equilibrium (Γ, c, π, δ) is *robust* if for any extension (Γ_1, c_1, π_1) of (Γ, c, π) , where Γ_1 is compact metric, 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 5.1. (a) *If the equilibrium (Γ, c, π, δ) is robust, where Γ is compact metric, 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.*

Proof. (a) If not, then there exists $m \in \psi(T)$ that is a profitable unilateral deviation by a seller. Define $\Gamma_1 = \psi(e(\Gamma)) \cup \{m\}$. (Because ψ is a homeomorphism, $\psi(e(\Gamma))$ is compact metric. Addition of the discrete point m leaves Γ_1 compact metric, as required by our definitions of ‘extension’ and ‘robustness’.) Further, (Γ_1, c^*, π^*) extends (Γ, c, π) , (take the restriction of $\psi \circ e$ as the required embedding α), and $(\Gamma_1, c^*, \pi^*, \psi \circ e[\delta])$ is not an equilibrium, contradicting robustness.

(b) Observe that the continuation equilibrium $(\psi(T), c, \pi)$ need not feature truthful reporting. Let (Γ_1, c_1, π_1) extend $(\psi(T), c, \pi)$, with embedding α . The

¹¹Both here and in the definition of robustness, we do not require that Γ be compact. That permits us to apply the term ‘extension’ also to the case where the set of indirect mechanisms Γ is $\psi(T)$.

appropriate form of (5.1) is

$$m_{\psi(t)}(\cdot, \psi(t'), \psi(t')) = m_{\alpha(\psi(t))}^1(\cdot, \alpha(\psi(t')), \alpha(\psi(t'))) \text{ on } \Omega^2.$$

By Theorem 4.1, 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

$$m_{\gamma_1}^1(\cdot, \gamma'_1, \gamma'_1) = \psi(e_1(\gamma_1))(\cdot, e_1(\gamma'_1), e_1(\gamma'_1)) \text{ on } \Omega^2.$$

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

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

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. ■

Robust equilibrium allocations (allocations supported by equilibria relative to $\psi(T)$) constitute the primary normative contribution of our analysis. However, in general, neither T nor $\psi(T)$ is compact, raising questions about the existence of robust equilibria. Comparison with the single principal context provides a useful perspective. In the standard setting, existence of an optimal mechanism is proven after imposing additional structure corresponding to specific applied problems. We suspect that such a procedure could succeed here as well. It is beyond the scope of this already lengthy paper to pursue this much further, but we offer some supporting comments.

First note that if the set of simple actions is finite, there is no private information, and preferences satisfy 4.2 and 4.3, then T is finite by Theorem 4.2. This implies that the universal set of mechanisms is finite. Then by Nash's theorem there exists an equilibrium (possibly in mixed strategies) relative to $(\psi(T), c^*, \pi^*)$. By Theorem 5.1, this equilibrium is robust, which guarantees that there are robust equilibrium allocations for such problems. There is as yet no proof of existence of robust equilibria when there is asymmetric information about buyers' preferences. However, it is the asymmetric information about the market, not asymmetric information about preferences, that creates all the conceptual problems in this

paper. Theorem 4.2 allows this asymmetric information about market outcomes. Thus it already provides the key existence result.

Each of the examples discussed in Section 2 satisfies 4.2 and 4.3 (this is readily checked by looking at the payoff matrices given in section 2) and therefore has a robust equilibrium allocation. By Theorem 4.2 it is sufficient to think of each seller offering buyers a mechanism consisting of a *finite* set of messages T . As there is no private preference information, a naive 'direct' mechanism is one that assigns the same action to each of these signals.

In the price competition game described in section 2.1, a seller who deviates from the set of naive direct mechanisms can choose a mechanism from the universal set in which the price that he charges depends on the messages from T that the buyers send him. Buyers will correctly anticipate this price and respond with a selection strategy that is a best reply to the anticipated price and the fixed price of the other seller. Since the symmetric equilibrium selection strategy associated with any pair of simple actions is unique, it will be the same strategy that buyers would have used in response to a direct deviation to this action in the original game. It follows that the deviation in the complicated message space generates the same profits as the corresponding deviation in the original game. Because we are starting from an equilibrium allocation for the original game, the deviation cannot be profitable.

The price matching example is not very interesting because the robust equilibrium allocations are simply the pure strategy equilibrium allocations with direct mechanisms. Example 2.2 includes a robust equilibrium allocation that is not an equilibrium in direct mechanisms. The proof of this, which appears in Appendix A, is more complicated because equilibrium path mechanisms have the property that by sending appropriate messages, buyers could induce sellers to use different actions from the ones used on the equilibrium path. Deviations by one seller in a more complicated message space could then be used to induce the other seller to switch actions, or to induce buyers to change their selection strategies. The payoff functions in that problem also support multiple symmetric continuation equilibria for any pair of simple actions offered by sellers. Deviations in a more complicated space of mechanisms might simply trigger continuation equilibria that the deviating seller prefers even when the actions that sellers take do not change. To verify robustness, it is necessary to show that the deviating seller cannot profitably exploit any of these opportunities.

6. The Nature of T and Non-Redundancy

This section is devoted to providing some intuition for the proof of Theorem 4.1, focussing primarily on the nature of T and the meaning of non-redundancy. At the end, we provide a proof of Theorem 4.2 (finiteness of T). For the technical details supporting this section the reader is referred to Appendices B-D.

Fix a continuation equilibrium $(\Gamma, \tilde{c}, \tilde{\pi})$ and consider the problem of trying to describe mechanisms in Γ in a way that is not tied to the specific view of competition embodied in Γ . This is the heart of our problem. An initial intuition is to use the payoff function generated by any mechanism as a way to describe that mechanism. This approach, which is the one we adopt, seems promising as a route to universality because mechanisms of all sorts deliver payoff functions.

To be more precise, consider using the buyers' payoff function $U(\cdot; \gamma)$, defined by (4.1), to describe the mechanism γ used by firm 1. A difficulty with doing so is that one of its arguments is the mechanism γ' in Γ used by the other firm. Thus the above payoff function is tied by its very definition to the given class Γ , (that is, its domain is $\Omega \times \Gamma$), contrary to the desired universality. The latter can be achieved, however, if we confine our description of γ_1 to the way in which its payoffs vary with valuations, a primitive of the model. The task, therefore, is to associate each $U(\cdot, \cdot; \gamma)$, a function on $\Omega \times \Gamma$, with a 'marginal' function that is defined on Ω . This is somewhat analogous to associating with each joint probability measure a suitable marginal measure, though there is no compelling and uncontentious notion of marginal for our setting. Our choice is to define the Ω -marginal to be $\sup_{\gamma' \in \Gamma} U(\cdot, \gamma'; \gamma)$.¹²

We arrive at an initial description of γ by means of $\Phi_0(\gamma)$, the function on Ω defined by

$$\Phi_0(\gamma)(\cdot) = \sup_{\gamma' \in \Gamma} U(\cdot, \gamma'; \gamma). \quad (6.1)$$

In words, our 0-level description of γ is given by the best valuation-contingent payoff that γ delivers, where 'best' is over all feasible mechanisms γ' for the other firm. The latter supremum evidently makes this a coarse description of γ and thus we proceed to refine it. This is possible because the 0-level description can be applied also to describe mechanisms used by the other firm. Thus we can refine (6.1) by computing the best valuation-contingent payoff that γ delivers, where 'best' is now over all feasible mechanisms γ' for the other firm that have

¹²Though 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.

a given 0-level description. In other words, we arrive at a level 1 description in terms of the function $\Phi_1(\gamma)(\cdot)$ defined by

$$\Phi_1(\gamma)(\cdot, h'_0) = \sup_{\gamma' \in \Gamma} \{U(\cdot, \gamma'; \gamma) : \Phi_0(\gamma') = h'_0\}, \quad (6.2)$$

where h'_0 varies over all possible 0-level descriptions. Proceeding inductively in the obvious way, one obtains a sequence of progressively finer descriptions $\Phi_n(\gamma)(\cdot)$ of γ , $n \geq 0$. The complete description of γ is provided by the infinite sequence of all n^{th} order descriptions.¹³

Thus we describe γ by means of its ‘type’

$$e(\gamma) = (\Phi_n(\gamma)(\cdot))_{n=0}^{\infty}. \quad (6.3)$$

The space T consists of all sequences of descriptions that can be constructed in this way, varying over all possible continuation equilibria $(\Gamma, \tilde{c}, \tilde{\pi})$.¹⁴

Turn to the meaning of non-redundancy of $(\Gamma, \tilde{c}, \tilde{\pi})$. It is very ‘close’ to the assumption that distinct mechanisms γ_1 and γ_2 in Γ have distinct descriptions of the sort just outlined, that is, e defined in (6.3) is one-to-one.¹⁵ Some violations of this assumption are not troubling. For example, non-redundancy is violated if there exist two distinct mechanisms in Γ that are effectively identical, but one employs communication in English while the other employs French. Our approach is to think of these mechanisms as being equivalent.

However, there exist violations that are serious. For example, suppose that Ω is a singleton representing a single risk averse buyer type who is trying to buy an insurance policy from a risk neutral seller. The set of simple 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, and in our description of mechanisms, only buyers’

¹³In some cases, only finitely many orders of description are ‘nontrivial’. For instance, in the price matching example (Appendix A) distinct mechanisms have distinct level 1 descriptions. Thus higher level descriptions are redundant (think of the level 2 counterpart to (6.2)).

¹⁴Our formal proof of the existence of T is constructive but is less intuitive than the argument sketched here. See [8] for a discussion of the relative merits of these two approaches to proving the existence of a types space in the context of types as beliefs.

¹⁵Lemma D.5 shows how non-redundancy yields that e is one-to-one. Note that if we used this invertibility as the (alternative) definition of non-redundancy, then Theorem 4.1 remains valid. In fact, we could then also drop the assumption of metrizability for ‘spaces’ if we simultaneously dropped the claim that T is separable metric. (This reveals the limited purpose of the assumption of metrizability, namely to permit a simpler statement of non-redundancy.)

payoffs are used. It seems natural to include sellers' payoffs also when describing and distinguishing between mechanisms.

This can be done by formally viewing sellers as buyers that have an artificial valuation $\bar{\omega}$ lying in the expanded space $\Omega \cup \{\bar{\omega}\}$. Details are provided in Section 7. The resulting form of non-redundancy is weaker because it is easier to distinguish between mechanisms. For example, it can be violated only if there exist distinct γ_1 and γ_2 satisfying *both*

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

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

where $V(\gamma'; \gamma)$ denotes the expected payoff to a seller using γ when the other firm is using γ' . We have been unable to find any interesting examples violating this notion of non-redundancy and our revelation principle is readily generalized to accommodate it. 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 in Section 7.

Turn to other aspects of Theorem 4.1 and its proof.¹⁶ Their explication requires that we provide some additional formal detail regarding T . Level 0 descriptions are functions of valuation and thus are elements of $\mathcal{U}(\Omega)$. Level 1 descriptions are functions of valuation and level 0 descriptions and thus lie in $\mathcal{U}(\Omega \times \mathcal{U}(\Omega))$. Thus if one defines the sequence $\{C_n\}$ inductively by

$$C_0 = \Omega, C_1 = \Omega \times \mathcal{U}(\Omega), C_n = C_{n-1} \times \mathcal{U}(C_{n-1}), n \geq 1, \quad (6.4)$$

then level n descriptions are elements of $\mathcal{U}(C_n)$ and

$$T \subset \prod_{n=0}^{\infty} \mathcal{U}(C_n).$$

Consequently, if $e(\gamma) = t = (h_n)_{n=0}^{\infty}$ is the type of the indirect mechanism γ , then its level n description $h_n \in \mathcal{U}(C_n)$ gives a buyer's expected payoff from γ as a function of $(\omega, h'_0, \dots, h'_{n-1})$, the buyer's valuation and all lower level descriptions of the other seller's mechanism. The problem of infinite regress mentioned in the introduction takes the following form: Given that we are describing a mechanism

¹⁶We emphasize that what follows is intended to provide intuition rather than a literal outline of the proof.

by the sequence t of all its finite level descriptions, does such a description *uniquely* determine a buyer's expected payoff from γ as a function of valuation and the sequence t' of all finite level descriptions of the other firm? The answer is 'yes' and the *unique* function that does the job is

$$\Psi(t)(\omega, t') = \inf_n h_n(\omega, h'_0, \dots, h'_{n-1}),$$

where $t = (h_n)_{n=0}^\infty$ and $t' = (h'_n)_{n=0}^\infty$. This positive result relies heavily on upper semi-continuity (see Appendix B).

The theorem asserts also that each type t may be associated with $\psi(t)$, a direct mechanism using message space $\Omega^2 \times T^2$. To see how this mapping is constructed, suppose that t is the type of some $\gamma \in \Gamma$ for the continuation equilibrium $(\Gamma, \tilde{c}, \tilde{\pi})$. Then application of the translation of Γ into T provided by e yields (recalling the notation m_γ defined in (3.2))

$$\psi(t)(\omega, \omega', t', t'') = \begin{cases} m_\gamma(\omega, \omega', e^{-1}(t'), e^{-1}(t'')) & \text{if } t', t'' \in e(\Gamma) \\ \underline{a} & \text{otherwise.} \end{cases} \quad (6.5)$$

By way of interpretation, only types in $e(\Gamma)$ are feasible given the continuation equilibrium $(\Gamma, \tilde{c}, \tilde{\pi})$. Consequently, reports of types outside $e(\Gamma)$ lead to the 'no trade' action \underline{a} . There remains the question "what if the same type t is associated with an indirect mechanism γ_1 coming from a different continuation equilibrium $(\Gamma_1, \tilde{c}_1, \tilde{\pi}_1)$?" In that case, because γ_1 has the same type t , using γ_1 in place of γ as above would yield a direct mechanism with the identical buyer's expected payoff function $\Psi(t)(\cdot)$. Thus $\psi(t)$ is well-defined up to 'payoff equivalence' and that suffices for our purposes.

The basis for the remaining claims in the theorem is now clear. Because the direct mechanisms $\psi(t)$ are constructed as above from some continuation equilibrium in indirect mechanisms, they embody incentives for truthful reporting of valuation and the other firm's type, as well as agreement with the 'recommended' choice probabilities. This ensures implementation (part (b)). Replication (part (a)) follows from the construction (6.5).

Turn finally to the proof of Theorem 4.2. Consider the continuation equilibrium $(\psi(T), c^*, \pi^*)$ provided by Theorem 4.1 and fix the types t and t' for sellers 1 and 2. In the absence of private information, both buyers communicate the identical message $c^*(t, t') = t'$ to the seller of type t and visit him with the common probability $\pi^*(t, t')$; similarly for behavior vis a vis the seller of type t' . Buyers

know the types t and t' and the communication strategy c^* and therefore can foresee the simple actions that will be taken at each seller, contingent on how many buyers participate. Thus in choosing where to participate, buyers play a game $G(a_0, b_0, a'_0, b'_0)$ of the form

| | | |
|---|-----------------------------|------------------------------|
| | 1 | 2 |
| 1 | $u_0(a_0, 1), u_0(a_0, 1)$ | $u_0(b_0, 0), u_0(b'_0, 0)$ |
| 2 | $u_0(b'_0, 0), u_0(b_0, 0)$ | $u_0(a'_0, 1), u_0(a'_0, 1)$ |

The interpretation is that the row buyer and column buyer choose whether to participate at firm 1, where simple action a_0 is taken if both buyers appear and b_0 is taken if only one appears, or at firm 2, where simple actions a'_0 and b'_0 are taken in corresponding circumstances. These simple actions are those prescribed by the mechanisms $\psi(t)$ at seller 1 and $\psi(t')$ at seller 2 for each possible participation status of the two buyers. Because in the continuation equilibrium, π^* is a best response to the other buyer's use of π^* , it follows that $(\pi^*(t, t'), \pi^*(t, t'))$ is a Nash equilibrium for the above game. It is also important to observe that, in terms of the notation introduced earlier in this section, $U(t'; \psi(t))$ is the expected payoff to a buyer from choosing seller 1 given that the other buyer chooses that seller with probability $\pi^*(t, t')$; similarly for $U(t; \psi(t'))$.

As outlined earlier, mechanisms are described by the payoffs that they deliver. Moreover, non-redundancy of $(\psi(T), c^*, \pi^*)$ means that payoffs must be sufficiently diverse to permit any two distinct mechanisms in $\psi(T)$, or equivalently any two distinct types in T , to be distinguished. On the other hand, our assumptions, including the finiteness of $u_0(\mathcal{A}_0 \times P)$, imply that only finitely many types can be distinguished. Thus T must be finite.

A more detailed argument is as follows: If $\pi^*(t, t') = 0$ or 1 , then $U(t'; \psi(t))$ is equal to one of the payoffs to the row buyer in the first row of $G(a_0, b_0, a'_0, b'_0)$. Hence it lies in $u_0(\mathcal{A}_0 \times P)$. Similarly for $U(t; \psi(t'))$. In the other case, where $0 < \pi^*(t, t') < 1$, then one can compute directly, using (4.3), that $\pi^*(t, t')$ is unique.¹⁷ This implies finiteness of the set of all Nash equilibrium payoffs for any game $G(a_0, b_0, a'_0, b'_0)$, of which there are only finitely many. Because $U(t'; \psi(t))$ is a Nash equilibrium payoff, it must lie in a finite set that is independent of the particular t and t' .

This proves that $\{U(t'; \psi(t)) : (t', t) \in T^2\}$ is a finite set, with cardinality $\#U$. For any n , the number of distinct descriptions of level n cannot exceed

¹⁷It equals $(u_0(a'_0, 1) - u_0(b_0, 0)) [u_0(a_0, 1) + u_0(a'_0, 1) - u_0(b'_0, 0) - u_0(b_0, 0)]^{-1}$.

$\#U$. Non-redundancy of $(\psi(T), c^*, \pi^*)$ requires that any two distinct mechanisms can be distinguished by descriptions of some level. Finally, descriptions become progressively finer the higher the level - mechanisms that have distinct descriptions of level n also have distinct descriptions of all higher levels. The conclusion is that there can be at most $\#U$ types.

7. EXTENSIONS

7.1. Sellers' Payoffs

We mentioned in Section 6 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

$$V(\gamma'; \gamma) = \int v(m_\gamma(\omega, \omega', \gamma', \gamma), \omega, \omega') dF dF. \quad (7.1)$$

In order to unify the notation for dealing with buyers and sellers, add a discrete point $\{\bar{\omega}\}$ to Ω and define $\bar{\Omega} = \Omega \cup \{\bar{\omega}\}$. Then proceed roughly as before with $\bar{\Omega}$ replacing Ω as the primitive valuations space. In order to do so, define $U : \bar{\Omega} \times \Gamma^2 \rightarrow [0, 1]$ so that for $\omega \in \Omega$, $U(\omega, \gamma'; \gamma)$ retains the interpretation as the buyer's payoff function, while $U(\bar{\omega}, \gamma'; \gamma) \equiv V(\gamma'; \gamma)$, representing the seller's payoff function. Expand the domain of the direct mechanism (3.2) m_γ so that

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

reflecting the fact that only valuations in Ω are conceivable for buyers. Then the assumptions payoff usc and non-redundancy, expressed relative to the expanded valuations space $\bar{\Omega}$, are both meaningful and appropriate. Theorem 4.1, with $\bar{\Omega}$ 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 $\bar{\omega}$), and the other seller's type, now expanded to include payoffs to sellers. This reformulation of non-redundancy 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.

7.2. More Specialized Models of Competition

It has been our intention to assume as little as possible *a priori* about the nature of competition. However, there may be situations where additional restrictions would be acceptable to some readers. Here we indicate briefly how these may be handled in our formal framework.

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 that was mentioned in Section 3.2 is to allow communication only after the buyer has chosen where to participate. Such a restriction on communication can be imposed using our formalism. Let $C = C_0^P$, in which case a message in C is a pair $c = (c_0, c_1)$ in C_0^2 , representing communication in the event of nonparticipation ($p = 0$) and participation ($p = 1$). Then restrict mechanisms so that only the messages of participants affect the action that is taken. To express this formally, use the notation in (3.1) and denote by $\gamma_c^{(p,p')}(c, c')$ the simple action prescribed by the indirect mechanism γ when p and p' describe the participation status of buyers 1 and 2. Then assume that $\gamma_c^{(1,0)}(c_0, c_1, c'_0, c'_1)$ depends only on c_1 , $\gamma_c^{(1,1)}(c_0, c_1, c'_0, c'_1)$ depends only on (c_1, c'_1) and so on. The resulting formal model specializes ours both by restricting message spaces C to the special form indicated and by restricting indirect mechanisms as just described.

Theorem 4.1 is valid also for this specialization of our model. The proof is similar, though the types space \hat{T} is smaller, because there are fewer mechanisms that require descriptions. Similarly, other specializations of our model lead to suitable subspaces of T as the relevant types space for describing mechanisms, with the revelation principle intact.¹⁸ We view this generality and flexibility as attractive features of our analysis.

7.3. More Information for Sellers

We have assumed that sellers are unable to observe the mechanisms offered by their competitors, or, if they can observe these mechanisms, that they are unable to write contracts that are explicitly contingent on them. It is well known that explicit contracts of this kind are powerful collusive devices. Our framework is easily modified to allow for such explicit contracts. A mechanism from our universal class asks the buyer to report his valuation along with a report about the

¹⁸ Another specialization that might be of interest and that can be accommodated is to allow a buyer to communicate with a seller only after observing whether or not the other buyer has selected the same seller.

mechanism being used by the other seller. If this latter mechanism is observable (and verifiable) then the report about the other seller’s mechanism can simply be reinterpreted as the observation of that mechanism. Because the incentive compatibility condition associated with reporting the competing mechanism can be dropped, this leads to straightforward changes in the construction of the universal language T and the corresponding class of mechanisms.

This environment is not very interesting because collusive outcomes can be supported in the usual way, that is, each seller writes a contract promising to punish deviations from the most collusive contract with counter deviations as in the original argument in [23]. The only difficulty is the precise form of this contract. A contract that responds only to deviations in price is not complete because it does not specify how sellers respond to deviations to more complicated contracts. Our formulation suggests that complete contracts of this kind should come from our universal class of mechanisms.

A. APPENDIX: Price Matching and Added Details

A.1. Proof of Robustness for Example 2.2

The outcome AA in example 2.2 can be supported by a robust equilibrium. From Theorem 4.2, the universal set of mechanisms consists of all mechanisms mapping from some finite set T into the two actions available to the seller. Thus every deviation in the universal set of mechanisms simply involves an alternative assignment of the elements of T to the two actions A and B . Since T is finite, the set of all such assignments is finite. Such deviations by seller 1, say, induce two potentially profitable outcomes. They might induce buyers in the continuation equilibrium to change the messages they send to seller 2 in a way that changes the action that seller 2 chooses, and they may alter the selection strategies that buyers use in the continuation equilibrium. In the example, the set of alternative pairs of actions for the sellers and the set of continuation equilibrium selection strategies associated with these is small, so we can check for robustness by exhaustively checking the potential outcomes associated with a deviation.

The various possibilities are described in the following table. The first two columns list the possible actions that the sellers might choose along the continuation equilibrium path associated with the deviation. The third column gives the payoff to seller 1 (the deviator) if the continuation equilibrium specifies that both buyers choose seller 1 along the continuation equilibrium path (“*” indicates that

there is no continuation equilibrium of this kind). The fourth column gives the payoff to seller 1 when both buyers select seller 2 on the continuation equilibrium path, while the final column gives seller 1's payoff in the case where buyers use a mixed selection strategy on the continuation equilibrium path.

| seller 1's action | seller 2's action | seller 1 | seller 2 | mixed |
|-------------------|-------------------|----------|----------|--|
| <i>A</i> | <i>A</i> | 1 | -1 | 1 |
| <i>B</i> | <i>B</i> | * | * | 0 |
| <i>A</i> | <i>B</i> | 1 | -1 | $-\frac{-1-\epsilon+\epsilon^2}{(1+\epsilon)^2} < 1$ |
| <i>B</i> | <i>A</i> | 2 | -2 | $2\frac{\epsilon}{1+\epsilon}$ |

In the final two rows of the matrix, the probability with which the seller whose action is *A* is chosen equals $\frac{1}{2(1+\epsilon)}$.^{19,20}

As is apparent from the table, the only way for the deviating seller to increase his profits (they are equal to 1 in the equilibrium constructed in example 2.2) is if the new mechanism induces seller 2 to play *A* while he plays *B*. This is profitable for seller 1 only if both buyers choose him. This result, however, cannot be part of any continuation equilibrium path. The reason is that seller 2 takes action *A* only when both buyers report *t* to him. The mechanism that seller 2 is offering is such that either buyer can unilaterally induce him to switch to *B* by sending the message *s*. Either buyer can improve upon the 0 payoff received along the equilibrium path by sending the message *s* to seller 2 and then choosing him with probability 1. Since the payoff a buyer gets from *B* when he is alone is 2, this deviation is profitable.

This rules out the possibility that seller 1 can achieve this payoff by deviating. Now from the table, each of the other possible outcomes that seller 1 could achieve by deviating yield him profits that are no higher than the profit (equal to 1) achieved in the original equilibrium. Thus there is no deviation in the set of universal mechanisms that will improve upon the original mechanism, making *AA* a robust equilibrium allocation. By the definition of robustness there will be no extension of the set of mechanisms described in example 2.2 within which seller 1 can raise his payoff if seller 2 sticks to his original mechanism.

¹⁹To understand inequalities relating to the last column, recall that ϵ is small.

²⁰Some outcomes listed in the table may be inconsistent with equilibrium. For example, it may not be an equilibrium for buyers to send messages that lead them to expect the deviator to use *A* and the non-deviator to use *B* when the buyers expect each other to mix.

A.2. Price Matching

In order to clarify notation and other aspects of our formalism, consider a simple competitive environment in which each of the two sellers has a single indivisible unit of output to sell and each buyer wishes to acquire exactly one unit if the price is low enough. We suppose that the feasible set Γ of indirect mechanisms consists of price matching mechanisms and then describe a continuation equilibrium that satisfies the assumptions required by Theorem 4.1. (We have conducted a similar exercise and verified the assumptions of Theorem 4.1 also for an alternative specification of Γ in which sellers compete in auctions. Details are available from the authors upon request.)

The environment is more completely described as follows: Buyers' valuations are independently drawn from the interval $\Omega = [0, 1]$ using the continuous probability distribution function F . The set of simple actions $\mathcal{A}_0 = [0, 1] \times \Delta^2$, where $\Delta^2 = \{\mu \in \mathbf{R}_+^2 : \mu_1 + \mu_2 \leq 1\}$. The generic simple 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 . Examples of 'full' actions, elements a of \mathcal{A} , include tuples of the form

$$a = ((q, 1, 0), (q, 0, 1), (q, .5, .5), \pi, \pi'). \quad (\text{A.1})$$

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 simple action undertaken if $(p, p') = (1, 0)$ or $(0, 1)$.

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 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. By using buyers' messages, it is possible to vary the price charged

according to the price offered by the other firm.²¹ However, the price cannot vary in response to the number of participants. Thus the actions described in (A.1) are the only ones in $\mathcal{A} = (\mathcal{A}_0)^{P^2} \times [0, 1]^2$ that are relevant.

Turn to payoffs. A buyer with valuation ω who trades at a price q receives utility $\omega - q$, while the seller in the same situation receives payoff q . This leads to the following specification of payoff functions (a is given by (A.1)):

$$u(a, \omega) = \pi' \max[\omega - q, 0] / 2 + (1 - \pi') \max[\omega - q, 0] \quad \text{and} \quad (\text{A.2})$$

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

The action that the seller chooses is allowed to depend on the messages received from buyers. Formally $\Gamma = \{\gamma_q : q \in [0, 1]\}$, where $\gamma_q : C^2 \rightarrow \mathcal{A}$ is defined by

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

Because firms make no attempt to influence buyers' participation choices, 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: When he is the only one to visit a seller, the buyer's report does not affect the price. 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.

The continuation equilibrium participation strategy is characterized by a cut-off 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

²¹This example differs from the model of Salop [23] 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.

interval choose the higher priced seller. That is,

$$\tilde{\pi}(\omega, q, q') = \begin{cases} 1 & \text{if } \omega \leq \omega^*(q, q') \text{ and } q \leq q' \\ 1 & \text{if } \omega > \omega^*(q, q') \text{ and } q > q' \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.3})$$

The cutoff valuation $\omega^*(q, q')$ is given by the solution ω to

$$\frac{1}{2F(\omega)} = \frac{(\omega - \max[q, q'])}{(\omega - \min[q, q'])}, \quad (\text{A.4})$$

when this equation has a solution. Otherwise the cutoff is equal to 1 and all buyers select the lower priced seller. As a result, $\omega^*(\cdot)$ is a continuous function.

The direct mechanism m_q corresponding to this equilibrium as in (3.2), (and imposing $q' = q''$), is given by $m_q(\omega, \omega', q', q') =$

$$((q, 1, 0), (q, 0, 1), (\min[q, q'], .5, .5), \tilde{\pi}(\omega, q, q'), \tilde{\pi}(\omega', q, q')). \quad (\text{A.5})$$

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

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

From (A.4), this payoff function is continuous.²² Because Γ is compact, there exists an equilibrium in mixed strategies.

Next we verify some properties of the continuation equilibrium $(\Gamma, \tilde{c}, \tilde{\pi})$. It follows from (A.2) and (A.5) that $U(\omega, q'; q) =$

$$\begin{cases} F(\omega^*(q, q')) \frac{\max[\omega - q, 0]}{2} + (1 - F(\omega^*(q, q'))) \max[\omega - q, 0] & \text{if } q \leq q' \\ F(\omega^*(q, q')) \max[\omega - q, 0] + (1 - F(\omega^*(q, q'))) \frac{\max[\omega - q', 0]}{2} & \text{otherwise.} \end{cases} \quad (\text{A.6})$$

By the continuity of $\omega^*(q, q')$, this is evidently jointly continuous in (ω, q, q') and this implies payoff upper-semicontinuity for $(\Gamma, \tilde{c}, \tilde{\pi})$.

²²Unlike in the usual Bertrand pricing problem, here sellers face capacity constraints that ensure that the low priced seller cannot serve the whole market. Buyers are thus reluctant to switch to the low priced seller for fear of being rationed. Thus when a seller cuts price he raises *continuously* the probability with which each buyer comes to him.

We prove non-redundancy as follows: Evidently, $\sup_{q'} U(\omega, q'; q)$ is achieved at $q' = 0$. Therefore, it suffices to show that $q_1 \neq q_2 \implies U(\cdot, 0; q_1) \neq U(\cdot, 0; q_2)$. If $q_2 > q_1$, then

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

on the same interval, implying

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

and hence that $F(\omega^*(q_1, 0)) = F(\omega^*(q_2, 0)) = 0$. This contradicts the definition of the cutoff values, whereby either (A.4) or $\omega^*(q_1, 0) = \omega^*(q_2, 0) = 1$.

B. APPENDIX: USC Functions and Hierarchies

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:

$$\{g : \exists s \in G, g(s) > \kappa\}, \{g : \forall s \in K, g(s) < \kappa\}, \quad (\text{B.1})$$

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 $g \mapsto \sup_{s \in A} g(s)$ is lower semi-continuous (lsc) for each open A and usc for each compact A .

The topology τ 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 [10, pp. 18-21]. Each closed set F can be identified with its indicator function, an usc function. Secondly, let $\mathcal{Cap}(S)$ denote the collection of all (regular Borel) capacities on S , endowed with the vague topology. Each usc g can be associated uniquely with the capacity ν_g ,

$$\nu_g(A) = \sup_{s \in A} g(s), \quad (\text{B.2})$$

for Borel measurable set A . Therefore, we have

$$\mathcal{F}(S) \subset \mathcal{U}(S) \subset \mathcal{Cap}(S). \quad (\text{B.3})$$

More to the point, the three topologies are consistent with one another in the sense that these set inclusions are topological embeddings.²³

The association in (B.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 B.1 and B.2 is clearer if one thinks of their well known counterparts for measures.

Theorem B.1. *If S is compact Hausdorff, then so is $\mathcal{U}(S)$. If S is also metric, then so is $\mathcal{U}(S)$.*

Proof. See [19, Theorems 2.2-3, Cor. 2.5]. For the second claim, use also [5, Theorem XI.4.1]. ■

Theorem B.2. *Let S be compact Hausdorff. Then:*

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

$$(\hat{e}g)(s) = \sup\{g(s') : e(s') = s\}, \quad (\text{B.4})$$

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

(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 $g \mapsto g_1$, $g_1(s_1, s_3) = \max_{s_2 \in \xi(s_3)} g(s_1, s_2)$, is a continuous map from $\mathcal{U}(S_1 \times S_2)$ into $\mathcal{U}(S_1 \times S_3)$.*

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

These properties are exploited heavily. 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

²³The capacity ν_g is called a *sup measure*; see [19] for details regarding these non-additive measures and the assertions in (B.3). See also [17, p. 285] for more on the connection between $\mathcal{U}(S)$ and other topological spaces.

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 $f : S_2 \rightarrow S_3$ 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 $\xi(s_3) = \{\bar{s}_2\}$ for all s_3 , then $g_1(s_1, s_3) = g(s_1, \bar{s}_2)$, but the mapping from g into the restriction $g(\cdot, \bar{s}_2)$ is generally not continuous (see [19, Theorem 4.5] for related results).

Proof of Theorem B.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

$$\Phi(g) \in \mathcal{N}_1 \equiv \{h_1 \in \mathcal{U}(S_1 \times S_3) : \sup_{G_1 \times G_3} h_1 > \kappa\},$$

for open sets $G_1 \subset S_1$ and $G_3 \subset S_3$. Then

$$g \in \mathcal{N} \equiv \{h \in \mathcal{U}(S_1 \times S_2) : \sup_G h > \kappa\},$$

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

$$\Phi(g) \in \mathcal{N}_1 \equiv \{h_1 \in \mathcal{U}(S_1 \times S_3) : \sup_{K_1 \times K_3} h_1 < \kappa\},$$

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

$$g \in \mathcal{N} \equiv \{h \in \mathcal{U}(S_1 \times S_2) : \sup_{K_1 \times K_2} h < \kappa\},$$

and $K_2 = \xi(K_3)$, a compact set.

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

Turn to hierarchies of usc functions, which provide the first step in the construction of T from Theorem 4.1. The following analysis parallels the analysis of hierarchies of probability measures (see [15] and [4], for example) and is a special case of the analysis in [6].

Theorem B.3. *There exists a non-empty compact metric space \mathcal{T} satisfying*

$$\mathcal{T} \sim_{\text{hmeo}} \mathcal{U}(\Omega \times \mathcal{T}). \quad (\text{B.5})$$

Proof. We observed above that $\mathcal{U}(S)$ is homeomorphic to a subspace of capacities. Hierarchies of capacities are a special case of the class of hierarchies studied in [6]. To be precise, the proof follows from Theorem B.1 above and from Theorems 4.2, 4.3 and 6.1 of [6]. ■

Some details from the proof, which is constructive, will be useful and so are outlined here. Define by (6.4) the spaces $\{C_n\}$, thought of as successively richer message spaces. Let $\mathcal{T}_0 = \Pi_0^\infty \mathcal{U}(C_n)$ with generic element $t = (g_0, g_1, \dots, g_n, \dots)$. Refer to the type t as *coherent* if

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

for all $n \geq 1$ and $c_{n-1} \in C_{n-1}$. The subspace of \mathcal{T}_0 consisting of coherent types is denoted \mathcal{T}_1 .

An important first step in the construction of \mathcal{T} is to note that \mathcal{T}_1 is homeomorphic to $\mathcal{U}(\Omega \times \mathcal{T}_0)$, with homeomorphism $\Psi : \mathcal{T}_1 \rightarrow \mathcal{U}(\Omega \times \mathcal{T}_0)$ constructed as follows: Let $t = (g_0, \dots, g_n, \dots) \in \Pi_{n=0}^\infty \mathcal{U}(C_n)$ be a coherent type. 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

$$\Psi(t)(z) = \inf_N g_N(z^N). \quad (\text{B.6})$$

Next consider the decreasing sequence of types spaces $\{\mathcal{T}_k\}$, where

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

Finally, define $\mathcal{T} = \cap \mathcal{T}_k$. To prove (B.5), observe first that

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

The latter set is homeomorphic to $\mathcal{U}(\Omega \times \mathcal{T})$; see Theorem B.2. ■

C. APPENDIX: Construction of \mathcal{T} and ψ

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 Theorem 4.1. In this appendix, we define T and ψ . Remaining assertions in the Theorem are proven in the next appendix.

There are some obvious restrictions on the pair (T, ψ) . Recall the interpretation whereby t describes a direct mechanism, denoted $\psi(t)$, that assigns an action to reports by buyers of valuations and the type of the other firm's mechanism. Formally, $\psi(t) \in \mathcal{A}^{\Omega^2 \times T^2}$, the set of measurable maps from $\Omega^2 \times T^2$ into \mathcal{A} . Direct mechanisms are of interest only if they possess all the properties (incentive compatibility, for example) possessed by indirect mechanisms. Therefore, (T, ψ) should deliver all such properties. To express these, note that by the nature of the space $\mathcal{A}_0^{P^2} \times [0, 1]^2$ of actions, we can describe any direct mechanism $m \in \psi(T) \subset \mathcal{A}^{\Omega^2 \times T^2}$ in the form

$$m(\cdot) = (m_c(\cdot), m_{\pi_1}(\cdot), m_{\pi_2}(\cdot))$$

where $m_c(\cdot)$ describes the participation contingent simple action and $m_{\pi_i}(\cdot)$ 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

$$m_{\pi_1}(\cdot) = m_{\pi_2}(\cdot) \equiv m_{\pi}(\cdot). \quad (\text{C.1})$$

Buyer 1 does not learn the valuation of buyer 2 before he makes his choice. Therefore, require that

$$m_{\pi}(\omega, \omega', t', t'') = m_{\pi}(\omega, \bar{\omega}', t', t'') \quad \forall \omega, \omega', \bar{\omega}', t', t'', t''' \quad (\text{C.2})$$

and we can write simply $m_{\pi}(\omega, t')$.

Finally, since we assume that with probability 1 each buyer selects a seller, m should satisfy: If $m = \psi(t)$, then for each $m' = \psi(t')$ in $\psi(T)$,

$$m_{\pi}(\omega, t') = 1 - m'_{\pi}(\omega, t). \quad (\text{C.3})$$

We refer to the preceding restrictions, (C.1) through (C.3) as *strong measurability constraints*.²⁴

Secondly, any direct mechanism should satisfy the following *incentive compatibility constraints*:

²⁴The adjective 'strong' is used to avoid confusion with the measurability of each m as a function, which is built into the definition of $\mathcal{A}^{\Omega^2 \times T^2}$.

(i) For each ω, ω' and ω^r in Ω and for each triple $t, t', t'' \in T$,

$$\int u(m(\omega, \omega', t'', t'), \omega) dF(\omega') \geq \int u(m(\omega^r, \omega', t', t''), \omega) dF(\omega'). \quad (\text{C.4})$$

If $m = \psi(t)$, then for each $\omega, \omega' \in \Omega$ and $m' = \psi(t')$,

$$\begin{aligned} & m_\pi(\omega, t') \int u(m(\omega, \omega', t', t'), \omega) dF(\omega') \\ & + [1 - m_\pi(\omega, t')] \int u(m'(\omega, \omega', t, t), \omega) dF(\omega') \geq \\ & \max \left[\int u(m(\omega, \omega', t', t'), \omega) dF(\omega'), \int u(m'(\omega, \omega', t, t), \omega) dF(\omega') \right] \end{aligned} \quad (\text{C.5})$$

Constraint (C.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 (C.5) imposes that buyers have no incentive to deviate from the recommended choice probability announced by the seller.

Two self explanatory ‘technical’ conditions follow.

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

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

Compact Support $m \in \mathcal{A}^{\Omega^2 \times T^2}$ has *compact support* if there exists compact $Y \subset T$ such that $U(\cdot, t'; m) = 0$ for $t' \in T \setminus Y$.

Assuming for the moment that (T, ψ) has been constructed, denote by $\mathcal{M}(\Omega^2 \times T^2)$ the set of all direct mechanisms $m \in \mathcal{A}^{\Omega^2 \times T^2}$ satisfying strong measurability, incentive compatibility, USC and having compact support. To this point we have argued that it is natural to require that

$$\psi(T) \subset \mathcal{M}(\Omega^2 \times T^2). \quad (\text{C.6})$$

We turn to an iterative construction of a suitable pair (T, ψ) .²⁵

²⁵As in Appendix B, if $B \subset S$ is compact, then $\mathcal{U}(S | B) = \{g \in \mathcal{U}(S) : g = 0 \text{ on } S \setminus B\}$. For arbitrary (possibly nonmeasurable) A , define $\mathcal{U}(S | A) = \cup \{\mathcal{U}(S | B) : B \subset A \text{ compact}\}$. If g is in $\mathcal{U}(S | A)$, say that g *knows* A .

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 (C.1)-(C.2), incentive compatible (in the sense of (C.4)), USC, has compact support and satisfies the variations of (C.3) and (C.5) for which the qualifiers ' $m = \psi(t)$ ' and ' $m' = \psi(t')$ ' are replaced by ' $m' \in \mathcal{A}^{\Omega^2 \times \mathcal{T}^2}$, $t' \in Y^0$, $\Psi(t) = U(\cdot; m)$ and $\Psi(t') = U(\cdot; m')$ '.

Then for each $k \geq 0$, let

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

Finally, define

$$T = \bigcap_{k=0}^{\infty} Y^k. \quad (\text{C.7})$$

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

$$T = \Psi^{-1} \left[\mathcal{U}(\Omega \times \mathcal{T} \mid \Omega \times T) \cap \text{cep} \left(\mathcal{M}^0(\Omega^2 \times \mathcal{T}^2) \right) \right]. \quad (\text{C.8})$$

(This can be verified as follows: That T contains the set on the right is immediate. For the converse, let $t \in \bigcap_{k=0}^{\infty} Y^k$. 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 [5, pp. 176, 233].

Having thus defined T , turn to the definition of ψ . From (C.8), 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 $\psi(t)$ as the restriction of m to $\Omega^2 \times T^2$. Then (C.6) follows. Because Ψ is one-to-one, so is ψ .

The following additional property is worth noting:

$$\{U(\cdot; m) : m \in \psi(T)\} = \{U(\cdot; m) : m \in \mathcal{M}(\Omega^2 \times T^2)\},$$

indicating that if we identify mechanisms that deliver the same expected payoff functions, then $\psi(T)$ 'equals' $\mathcal{M}(\Omega^2 \times T^2)$.

D. APPENDIX: Proof of Embedding

Let $(\Gamma, \tilde{c}, \tilde{\pi})$ be a continuation equilibrium as in Theorem 4.1. Here we construct the embedding e and the continuation equilibrium $(\psi(T), c^*, \pi^*)$. Notation introduced in the preceding appendices, including the message spaces C_n defined in (6.4), is used freely.

Denote by θ the map taking each $\gamma \in \Gamma$ into $m_\gamma \in \mathcal{A}^{\Omega^2 \times \Gamma^2}$ (see (3.2)). Endow $\theta(\Gamma)$ with the weak topology induced by the map into $\mathcal{U}(\Omega \times \Gamma)$ that takes γ into $U(\cdot; \theta(\gamma))$. Thus 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 cannot be separated. This reflects the view that there is no reason to distinguish between such mechanisms. Observe that this topology makes θ a continuous map, because of the assumption that $(\Gamma, \tilde{c}, \tilde{\pi})$ is payoff usc.

The following commutative diagram may provide a useful guide. Use *cep* (conditional expected payoff) to denote the map taking $m \in \psi(T)$ into $U(\cdot; m)$. The maps θ , ψ and Ψ have already been defined, while Φ and e will be defined here.

$$\begin{array}{ccccc} \Gamma & \xrightarrow{e} & T \subset \mathcal{T} & \xrightarrow{\Psi} & \mathcal{U}(\Omega \times T) \\ \downarrow \theta & \nearrow \Phi & \downarrow \psi & \nearrow_{cep} & \\ \theta(\Gamma) \subset \mathcal{A}^{\Omega^2 \times \Gamma^2} & & \psi(T) \subset \mathcal{A}^{\Omega^2 \times T^2} & & \end{array}$$

Lemma D.1. For each $V \in \mathcal{U}(\Omega \times \mathcal{T})$, $\Psi^{-1}(V) = (\Theta_n(V))_0^\infty \in \Pi_0^\infty \mathcal{U}(C_n)$, where

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

$$\Theta_n(V)(\omega, g'_0, \dots, g'_{n-1}) = \sup_{t'} \{V(\omega, t') : \Theta_k(\Psi(t')) = g'_k, k < n\}.$$

Proof. In light of (B.3), this is a special case of [6, Theorem D.1]. An interpretation is that $\Theta_n(V)$ is a projection of V onto $\Pi_0^\infty \mathcal{U}(C_i)$. ■

Lemma D.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, $\gamma \mapsto U(\cdot; \theta(\gamma)) \in \mathcal{U}(\Omega \times \Gamma)$ is continuous. Thus (a) follows from the compactness of Γ . Part (c) follows from Theorem B.2. Part (b) is a consequence of the weak topology used for the space of mechanisms. ■

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

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

$$\Phi_n(m)(\omega, g'_0, \dots, g'_{n-1}) = \sup \{U(\omega, \gamma'; m) : \exists m', \theta(\gamma') = m', \Phi_k(m') = g'_k, k < n\}.$$

Lemma D.3. $\Phi : \theta(\Gamma) \rightarrow \mathcal{T}$ 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 B.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 $\theta(\Gamma)$, this yields

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

This implies (by Theorem B.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 B.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 B.3. It suffices to show that

$$\Psi(t) = 0 \text{ on } \Omega \times (\mathcal{T}_0 \setminus \mathcal{T}), \quad (\text{D.1})$$

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

From (B.6), $\Psi(t)(z) = \inf_N \Phi_N(m)(z^N)$, 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

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

But

$$\bigcap_{N=0}^\infty \left(\Omega \times \pi^N(\mathcal{T}_N) \times \Pi_{N+1}^\infty \mathcal{U}(C_n) \right) \subset \Omega \times \bigcap_{N=0}^\infty \mathcal{T}_N = \Omega \times \mathcal{T}, \quad (\text{D.2})$$

proving (D.1). (The routine proof of (D.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 (B.3)) of [19, p. 55]:

Lemma D.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 D.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 and θ is continuous.

Let $\{\Sigma_n\}$ be the σ -algebras on Γ appearing in the definition of non-redundancy. 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 D.4, Σ_n is the weakest σ -algebra such that $\Phi_{n-1} \circ \theta : \Gamma \rightarrow (\mathcal{U}(C_{n-1}), \mathcal{B}_{n-1})$ is measurable. Therefore, non-redundancy implies that $e = \Phi \circ \theta : \Gamma \rightarrow \Pi_0^\infty \mathcal{U}(C_n)$ is one-to-one.

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

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

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

Show that $e(\Gamma) \subset T$: Recall the definition (C.7), whereby $T = \cap Y^k$. Let $\gamma \in \Gamma$. Then, by the claim, $\Psi(e\gamma) = \Psi(\Phi \circ \theta(\gamma)) = U(\cdot; \xi \circ \theta(\gamma))$, implying that

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

Show next that $e(\gamma) \in Y^1$: First, $\xi \theta(\gamma) \in \mathcal{M}^0(\Omega \times \mathcal{T})$, 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, $t' \in \mathcal{T} \setminus Y \implies \xi(\theta \gamma)(\cdot, t') = \underline{a} \implies \Psi(e\gamma)(\cdot, t') = 0$. Conclude that $e(\gamma) \in Y^1$. The proof may be completed by induction. ■

Turn to the continuation equilibrium described in Theorem 4.1. View $\psi(T)$ as a feasible set of *indirect* mechanisms using message space $C = \Omega \times T$. We constructed T and ψ to satisfy (C.6). Thus every mechanism in $\psi(T)$ is payoff usc, strongly measurable and incentive compatible. Adopt the notation in (3.1) 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 $c^r = (\omega^r, t^r)$,

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

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

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

(by definition of π^*)

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

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

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

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

$$\begin{aligned} & \pi^*(\omega, t, t') \int u(\{\gamma_c(c^*(\omega, t, t'), c^*(\omega', t, t')), \pi^*(\omega, t, t'), \pi^*(\omega', t, t')\}, \omega) dF(\omega') + \\ & [1 - \pi^*(\omega, t, t')] \int u(\{\gamma_c(c^*(\omega, t', t), c^*(\omega', t', t)), \pi^*(\omega, t', t), \pi^*(\omega', t', t)\}, \omega) dF(\omega') \\ & \quad = \gamma_{\pi_1}(\omega, t') \int u(\gamma(\omega, \omega', t', t'), \omega) dF(\omega') + \\ & \quad (1 - \gamma_{\pi_1}(\omega, t')) \int u(\gamma'(\omega, \omega', t, t), \omega) dF(\omega') \geq \\ & \max \left[\int u(\gamma(\omega, \omega', t', t'), \omega) dF(\omega'), \int u(\gamma'(\omega, \omega', t, t), \omega) dF(\omega') \right], \end{aligned}$$

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 non-redundancy. We have to show that they separate any distinct $\psi(t_1)$ and $\psi(t_2)$. But $\psi(t_1) \neq \psi(t_2) \implies t_1 = (g_n^1) \neq (g_n^2) = t_2 \implies$ (by Lemma D.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 $\psi(t) \mapsto \Theta_0(U(\cdot; \psi(t))(\cdot))$ is measurable. (In other words, modify the measurability constraint in the definition of non-redundancy 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] D. Bernhardt and E. Hughson, Splittling orders, *Rev. Fin. Stud.* **10** (1997), 69–102.
- [3] D. Bernheim and M. Whinston, Menu auctions, resource allocation and economic influence, *Quart. J. Econ.* **101** (1987), 1–31.
- [4] A. Brandenburger and E. Dekel, Hierarchies of beliefs and common knowledge, *J. Econ. Theory* **59** (1993), 189–198.
- [5] J. Dugundji, *Topology*, Allyn and Bacon (1966).
- [6] L. Epstein and T. Wang, ‘Beliefs about beliefs’ without probabilities, *Econometrica* **64** (1996), 1343–1374.
- [7] L. Glosten, Insider trading, liquidity and the role of the monopolistic specialist, *J. Business* **62** (1989), 211–236.
- [8] A. Heifetz and D. Samet, Topology-free typology of beliefs, Working Paper 26/96, Tel Aviv, (1996).
- [9] A. Heifetz and D. Samet, Universal partition structures, Working Paper 26/93, Tel Aviv, (1993).

- [10] W. Hildenbrand, *Core and Equilibria of a Large Economy*, Princeton University Press (1974).
- [11] M. Katz, Game playing agents: unobservable contracts as precommitments, *Rand J. Econ.* (1991), 307–329.
- [12] A. Kyle, Continuous auctions and insider trading, *Econometrica* **53** (1985), 1335–1355.
- [13] D. Martimort and L. Stole, Communication spaces, equilibria sets and the revelation principle under common agency, mimeo University of Chicago, (1997).
- [14] P. McAfee, Mechanism design by competing sellers, *Econometrica* **61** (1993), 1281–1312.
- [15] J. Mertens and S. Zamir, Formulation of Bayesian analysis for games with incomplete information, *Int. J. Game Theory* (1985), 1–29.
- [16] R. Myerson, Optimal coordination mechanisms in generalized principal agent problems, *J. Math. Econ.* **10** (1982), 67–81.
- [17] T. Norberg, Random capacities and their distributions, *Probab. Th. Rel. Fields* **73** (1986), 281–297.
- [18] T. Norberg and W. Vervaat, Capacities on non-hausdorff spaces, mimeo, (1989).
- [19] G. O’Brien and W. Vervaat, *Capacities, Large Deviations and LogLog Laws*, Birkhauser (1991).
- [20] J. Peck, Competing mechanisms and the revelation principle, *Ohio State University Manuscript* (1995), 12.
- [21] M. Peters, On the equivalence of Walrasian and non-Walrasian equilibria in contract markets, *Rev. Econ. Stud.* (1997), 241–265.
- [22] M. Peters and S. Severinov, Competition among sellers who offer auctions instead of prices, *J. Econ. Theory* **75** (1997), 141–179.
- [23] S. Salop, *Practices that Credibly Facilitate Oligopoly Coordination*, Cambridge, MIT Press (1986).

- [24] R. Wilson, Auctions of shares, *Quart. J. Econ.* **93** (1979), 675–689.
- [25] R. Wilson, Incentive efficiency of double auctions, *Econometrica* **53** (1985), 1101–1115.

