Toward a Strategic Foundation for Rational Expectations Equilibrium*

Philip J. Reny
Department of Economics
University of Chicago

Motty Perry
Department of Economics
Hebrew University of Jerusalem, and
University of Chicago

February 2006

Abstract

A step toward a strategic foundation for rational expectations equilibrium is taken by considering a double auction with n buyers and m sellers with interdependent values and affiliated private information. If there are sufficiently many buyers and sellers, and their bids are restricted to a sufficiently fine discrete set of prices, then, generically, there is an equilibrium in nondecreasing bidding functions that is arbitrarily close to the unique fully revealing rational expectations equilibrium of the limit market with unrestricted bids and a continuum of agents. In particular, the large double auction equilibrium is almost efficient and almost fully aggregates the agents' information.

1. Introduction

The Rational Expectations Equilibrium (REE) concept has had a profound effect on economic theory. However, there is widespread agreement that an adequate foundation has yet to be provided. This is evidenced by the various paradoxes that accompany virtually any presentation of the idea.

The desire to provide a foundation for REE is not new. Hellwig (1980) notes that the rational expectations hypothesis typically requires traders to act rationally with respect to information, yet fail to recognize their influence on the price. Hellwig presents a

*We wish to thank four referees and the editor for providing numerous insightful comments and suggestions. We also wish to thank George Mailath and Jeroen Swinkels for helpful conversations. Both authors gratefully acknowledge financial support from the National Science Foundation (SES-9905599, SES-0214421, and SES-0001744).

model with both competitive traders and noise traders. Competitive traders do not take into account their effect on the price, but do take into account the information conveyed by the price. On the other hand, noise traders' demands are exogenous, random, and independent of fundamentals. Their presence is required to keep the market from collapsing. In an attempt to justify his assumption that agents are price-takers, Hellwig considers a sequence of economies in which the number of competitive and noise traders grow. He shows that the market equilibrium converges to an REE. However, due to the presence of noise traders this equilibrium is neither fully revealing nor efficient. While this impressive result marked an important step, explaining rational expectations through the presence of irrational noise traders and nonstrategic price-taking competitive traders would not be the end of the story.[1]

Building upon Wilson's (1977) striking information aggregation result in a bidding model, Milgrom (1981) provides a seminal contribution toward a foundation for REE. This remarkable paper initiates what might be called the strategic bidding approach to REE. Milgrom demonstrates that a strategic bidding model, with its explicit price formation process and well-defined order of moves, is capable of resolving all of the paradoxes typically associated with REE. On the other hand, Milgrom shows that obtaining a *fully-revealing* REE in his model requires strong assumptions, even as the number of bidders grows.[2] In general, Milgrom's limiting equilibrium is an REE, but it typically fails to be fully revealing, thereby resulting in an inefficient outcome.

In an important recent contribution, Pesendorfer and Swinkels (1997) reveal why the limiting equilibria in Milgrom (1981) often fail to be fully revealing. Pesendorfer and Swinkels note that Milgrom's limit result focuses on a market with an arbitrarily large number of bidders but with a *fixed finite* number of units of the good for sale. In contrast, they demonstrate that when *both* the number of bidders *and* the number of units for sale grow large, the unknown state of nature typically *is* eventually revealed by the equilibrium price. That is, a fully-revealing REE typically does arise in Milgrom's model when there are sufficiently many bidders and sufficiently many units for sale.[3]

[1] A strategic version of Hellwig's model is considered in Kyle (1989), where the competitive traders, but not the noise traders, take into account their effect on the price. Still, without noise traders the market breaks down. Consequently, as in Hellwig (1980), the limiting equilibrium is necessarily a nonfully-revealing, inefficient REE. A related analysis, but with two-dimensional uncertainty and a continuum of (price-taking) traders, can be found in Messner and Vives (2001). Their multi-dimensional uncertainty precludes the one-dimensional price from acheiving a fully revealing REE, a possibility that is consistent with the work of Allen (1981).

[2] These assumptions are built into Wilson's model.

[3] The Cournot literature provides a related model with strategic agents (sellers) on just one side of the market. This literature too contains both positive (Palfrey (1985)) and negative (Vives (1988)) results on information aggregation. On the other hand, this literature is not precisely in the same vein as the auction literature because submitting a bid in an auction, given single-unit demands, is equivalent to submitting a demand schedule. In contrast, the Cournot literature permits multi-unit supplies but firms are restricted to submitting quantities, not supply schedules. For the same reason,

The bidding models employed in both Milgrom (1981) and Pesendorfer and Swinkels (1997) are Vickrey-type auctions in which, on the one hand, buyers behave strategically based on their private information. On the other hand, the sellers are passive, simply showing up at auction with a fixed number of units available for sale at any nonnegative price. One can view these sellers as playing a role similar to that played by the noise traders in the models of Hellwig and Kyle. Not only does their presence serve to ensure that trade takes place, but their number must grow without limit to ensure that information is aggregated.

Thus, while all of the above models have produced important insights, they share the undesirable feature that their markets would collapse were it not for the presence of uninformed or irrational or nonstrategic agents.

A fully satisfactory foundation for REE should permit all traders, buyers and sellers, to be informed, rational, and strategic. It is the purpose of the present paper to move in this direction. The natural bidding model with which to study the resulting two-sided market is a double auction.

Wilson (1985) introduces the many-buyer/many-seller double-auction model and shows that it is incentive-efficient in large markets with strategic traders.[4] This pioneering and significant work focuses on settings in which traders have independent private values.[5,6] Our analysis covers the independent private value setting as well as settings in which the good's value is partially common and partially private and where the agents possess affiliated private information. For our purposes, the inclusion of a common value component is critical, since this is the hallmark of environments in which rational expectations are required. On the other hand, we do not cover the pure common value case.[7]

---

the Cournot literature is not directy related to the literature on rational expectations, wherein consumers are imagined to submit excess demand schedules to a Walrasian auctioneer.

[4] Chaterjee and Samuelson (1983) introduced the one-buyer/one-seller double auction in their work on bargaining with two-sided incomplete information.

[5] For the buyer's-bid double-auction, Satterthwaite and Williams (1989) establish a $1/n$ rate of convergence to price-taking behavior. In addition to this but in a k-double auction, Rustichini, Satterthwaite and Williams (1994) establish a $1/n^2$ rate of convergence to efficiency. In both cases, n is the equal number of buyers and sellers.

[6] When values are private, double auction equilibria have been shown to be nearly efficient under quite general distributional assumptions, so long as the market is sufficiently large. For example, Swinkels (2001) considers the case in which value distributions are independent but not identical, while Cripps and Swinkels (2002) and Fudenberg et. al. (2003) permit value distributions that are neither independent nor identical. Cripps and Swinkels allow mixed equilibria (relying on Jackson and Swinkels (2002) for existence), while Fudenberg et. al. restrict attention to, but also prove the existence of, monotone pure strategy equilibria in large markets.

[7] Indeed, our proof technique does not work under pure common values. Note however that given the risk-neutrality of our agents, the pure common value case is not particularly interesting. This is because every allocation is then ex-ante efficient and so Milgrom and Stokey's (1982) no-trade theorem applies.

The double-auction model possesses at least two important features. First, in contrast to the one-sided Vickrey auction models employed in Milgrom (1981) and Pesendorfer and Swinkels (1997), rational traders in a double auction act strategically to manipulate the market-clearing price in their favor. This leads to some efficient trades going unrealized with positive probability. It is therefore possible to investigate the extent to which such manipulation and inefficiency vanishes, and the extent to which price-taking behavior approximates strategic behavior as the market becomes large. Second, the double auction mimics the workings of actual markets in use. Specifically, it operates like a call market which itself is used to conduct and price trades in many financial markets, such as, for example, the overnight market on the NYSE. Consequently, the analysis of double auctions with nonprivate values has the potential to provide insights into whether such financial markets are able to aggregate the participants' information and generate efficient outcomes.

We put together the two strands of the above literature — the one-sided (buyer-only) markets with interdependent-values and the two-sided markets with private values — by studying a single two-sided market in which ex-ante symmetric buyers and sellers with single-unit demands and possibly interdependent values participate in a double auction.

Our main result is as follows. Suppose bids and offers must be submitted in discrete units from a sufficiently fine price grid. If the market contains sufficiently many buyers and sellers, then generically in the agents' value functions and the fineness of the grid, the double auction possesses a nontrivial Bayes-Nash equilibrium in pure monotone bidding functions.[8,9] Further, the equilibrium outcome is arbitrarily close to an efficient fully-revealing REE, and equilibrium behavior is arbitrarily close to price-taking behavior. Hence in this sense, *for the very simple two-sided market studied here*, we provide a strategic foundation for fully-revealing REE.

Notably, we do not establish the existence of a nontrivial pure bidding equilibrium when the market is small. Our proof technique is designed to take advantage of the economics of large markets. The structure of our proof is as follows. We begin by analyzing an idealized limit market with a continuum of buyers and sellers, and in which each agent's bid can be any nonnegative real number. The continuum of agents precludes any single agent from affecting the price and the continuum of feasible bids implies that, in equilibrium, ties in bids will not occur. Together these properties allow us to construct, straightforwardly, a symmetric equilibrium of the double auction in pure strictly increasing bidding functions. This limit market equilibrium is, furthermore, a fully-revealing REE.

The heart of the proof establishes a continuity property for the fully-revealing double-

[8] Trivial "no trade" equilibria always exist.

[9] The genericity result alluded to here is unrelated to that in Allen (1981) since our limit economy always possesses a fully-revealing REE.

auction equilibrium of the limit market. The "nearby" market with large but finite numbers of buyers and sellers, and a sufficiently fine grid of bids, possesses a nearby (and so almost efficient and almost fully-revealing) equilibrium. Demonstrating this is not straightforward because, when there are finitely many traders, buyers and sellers can each affect the price but wish to affect it in opposite directions. Hence, buyers and sellers will bid differently. Furthermore, standard equilibrium existence proofs in bidding models rely heavily on affiliation properties of order statistics, and these properties fail here precisely because buyers and sellers employ distinct bidding functions.[10]

Our main insight is that establishing the existence of a monotone equilibrium poses serious difficulties only when individual agents can have a significant impact on the price. In sufficiently large markets, where strategic price manipulation is negligible, we demonstrate that a monotone pure strategy equilibrium generically exists. Our proof technique also exploits the ex-ante symmetry, modulo endowments, of the agents in our model. However, we believe that our existence result can be pushed through even without symmetry. On the other hand, the symmetry of the agents *is* used here to establish the efficiency of the limit equilibrium.

For technical reasons, our proof technique requires us to drop the somewhat unrealistic, yet often very convenient, assumption that agents can submit bids from a continuum. Instead, we assume that bids are restricted to a sufficiently fine discrete grid of prices. This in turn forces us to deal head on with the issue of rationing, which has received little attention in the literature.

Important avenues for future research include extending the present model to settings with asymmetric and risk-averse agents, multi-unit demands, multiple markets, and multi-dimensional private information. While each of these extensions is likely to present non-trivial challenges, we are hopeful that each new result, including the present, will provide a useful point from which to begin thinking about those that remain.

The remainder of the paper is organized as follows. Section 2 describes the basic setup. In Section 3 the limit model with a continuum of agents and a continuum of prices is studied. The finite agent model with a discrete grid of prices is described in Section 4, and Section 5 illustrates the difficulties in obtaining monotone equilibria when either the number of agents is small or the grid of prices is insufficiently fine. Section 6 contains the main result and a detailed sketch of the proof. The supplement to the present paper, Reny and Perry (2006), contains the complete proof as well as additional related material, and will henceforth be referred to as RP-s.

[10] Evidently, this is why in the strategic models cited above (Kyle (1989), Milgrom (1981), and Pesendorfer and Swinkels (1997)) all strategic traders are symmetric.

2. The Basic Setup

There are N agents, each of whom desires at most one unit of a single indivisible good. The "state of the good" is a random variable, $\tilde{\omega}$, taking values in $[0, 1]$. The realization, ω, of $\tilde{\omega}$ is unknown to all agents.[11] The density from which $\tilde{\omega}$ is drawn is $g(\omega)$, defined on $[0, 1]$. Conditional on the realization of the state of the good, ω, each agent i receives an i.i.d. signal, $\tilde{x}_i$, taking values in $[0, 1]$ according to the density $f(x|\omega)$, defined on $[0, 1]^2$. We make the following assumptions.

A.1 On their domains, g is C^1 and f is C^2 and both are strictly positive.

A.2 $\partial^2 \ln f(x|\omega)/\partial x \partial \omega > 0$ for all $(x, \omega) \in [0, 1]^2$.

Assumption A.2 captures the idea that a high signal is good news about the state of the good by requiring $f(x|\omega)$ to satisfy the strict monotone likelihood ratio property. Equivalently, for each i, agent i's signal, $\tilde{x}_i$, and the state of the good, $\tilde{\omega}$, are strictly affiliated. Because the N signals $\tilde{x}_1, ..., \tilde{x}_N$ are i.i.d. conditional on $\tilde{\omega}$, this implies that the $N + 1$ random variables $\tilde{x}_1, ...\tilde{x}_N, \tilde{\omega}$ are strictly affiliated as well.

When his signal is $x \in [0, 1]$ and the state of the good is $\omega \in [0, 1]$, an agent's von Neumann-Morgenstern utility is $v(x, \omega) \cdot \chi - p$, where $\chi = 0$ if he ends up with zero units of the good, $\chi = 1$ if he ends up with at least one unit of the good, and $p \in \mathbb{R}$ is the price paid. Hence, each agent wants at most one unit of the good. We make the following assumptions on $v(\cdot, \cdot)$.

A.3 On $[0, 1]^2$, $v(\cdot, \cdot) \geq 0$ is C^1, $v_x(\cdot, \cdot) > 0$, and $v_\omega(\cdot, \cdot) \geq 0$.[12]

A.4 $v_\omega(x, 0) = v_\omega(x, 1) = 0$ for every $x \in [0, 1]$.

Assumption A.3 says that each agent's value is nonnegative, weakly increases if ω increases, and strictly increases if x increases; so there is a strictly positive, although perhaps arbitrarily small, private value component. Thus, while our model includes the pure private value model as a special case, the pure common value model is excluded. Assumption A.4 says that values are almost private when the state is extreme. Its essentially technical role is explained in Remark 2 below. Note that any function satisfying A.3 but not A.4 can be approximated arbitrarily closely by one that satisfies both A.3 and A.4.[13]

[11] Letters with tildes denote random variables, while letters without tildes represent their realizations.

[12] Subscripts denote partial derivatives.

[13] For example, fix $\varepsilon \in (0, 1)$ and let $\lambda_e : [0, 1] \to [0, 1]$ be any nondecreasing differentiable function that is 0 at $\omega = 0$ and 1 at $\omega = \varepsilon$ (hence $\lambda'_\varepsilon(\omega) = 0$ for $\omega \geq \varepsilon$). If $v(x, \omega)$ satisfies A.3, then $\tilde{v}(x, \omega) = (1 - \lambda_e(\omega))v(x, 0) + \lambda_e(\omega)\lambda_e(1 - \omega)v(x, \omega) + (1 - \lambda_e(1 - \omega))v(x, 1)$ satisfies A.3 and A.4 and converges uniformly to $v(x, \omega)$ as $\varepsilon \to 0$.

Our main result (Theorem 6.1) will be shown to hold for a generic subset of the space of value functions satisfying A.3 and A.4. But because this space is not finite dimensional, we cannot define "generic" in terms of Lebesgue measure. We therefore adopt a standard topological approach.

Let V denote the subset of functions, $v(x,\omega)$, satisfying A.3 and A.4, and define a norm on V by $\|v\| = \max_{x,\omega} v(x,\omega)$, thereby inducing a topology on V.[14] Genericity in V can now be defined in terms of so-called residual sets.

Definition 2.1. *A subset of a topological space is residual if it contains a countable intersection of open dense sets.*[15]

A standard interpretation is that if a property holds on a residual set, then the property holds "generically." We will demonstrate that our main result holds on a residual set of value functions satisfying A.3 and A.4.[16]

Up to this point, all agents are ex-ante symmetric. We now break this symmetry by endowing $m < N$ agents with one unit of the good. Because each agent desires at most one unit, these agents, should they trade, will necessarily give up their units. Hence, we shall refer to these m agents as sellers. The remaining $n = N - m$ agents will be referred to as buyers. That is, there are n buyers and m sellers.

We now provide an example that is consistent with assumptions A.1-A.3.[17] In the example, an agent's income is a private signal that influences the expected value of the good in two ways: directly through its effect on private incentives (providing a private value component) and indirectly via an informational effect (providing a common value component).

> **Example.** There are three periods. At date zero there is a market for an indivisible asset. The asset's expected value depends on an unknown state of the economy, $\tilde{\omega}$, distributed uniformly on $[0,1]$. For simplicity assume that, given ω, the asset's expected value is $1+\omega$. An agent's income, $\tilde{x}$, is strictly affiliated with the state of the economy. In particular, $\tilde{x}$ has density $f(x|\omega)$ satisfying assumption A.2. Hence, both A.1 and A.2 are satisfied.

[14] Any topology on V that is at least this strong and for which linear combinations of elements of V are continuous in the coefficients will do.

[15] Recall that a set is dense if arbitrarily close to every point not in the set is a point in the set.

[16] In RP-s it is shown that the main result of the present paper also holds when one adopts the measure-motivated approach of Christensen (1974) or Hunt Sauer and York (1996) for infinite-dimensional spaces. In their approach, a set, A, is "generic" if there exists a probability measure, μ, on the ambient space such that $\mu(A^c + x) = 0$ for every x in the ambient space. When the ambient space is $\mathbb{R}^n$, this implies that A^c has Lebesgue measure zero, which is an attractive feature of this approach.

[17] As mentioned above, it is always possible to perburb such an example so that the technical assumption A.4 also holds.

Agents participate in the date-zero market after learning their own income. The asset's return is realized at date one but is available at date two. At date one, each agent is liquidity constrained with probability $q(x)$, and, if he owns the asset, he will find it optimal to collect the return early at a cost of $c > 0$.

The higher is the agent's income, the less likely it is that he will be liquidity constrained. Hence, $q'(x)$ is strictly negative. Therefore, conditional on ω and x, the expected value of the asset at date zero is $v(x, \omega) = 1 + \omega - q(x)c$. If $q'(x)$ is continuous and $c < 1$, then $v(\cdot, \cdot)$ is C^1 and non negative and so assumption A.3 is satisfied.

Our main interest lies in studying the outcome of strategic bidding behavior among the n buyers and m sellers in a double-auction market when n and m are large. But before doing so, we first analyze the double auction in a setting in which no agent can affect the price. This limit market is relevant because it's unique fully-revealing REE will turn out to be approximated arbitrarily well by an equilibrium of a double-auction market with sufficiently many buyer and sellers.

3. The Limit Model

Let us alter the basic setup by supposing that there are a continuum of buyers and sellers. Specifically, suppose that there is a unit mass of agents, of whom $\alpha \in (0, 1)$ are buyers and $1 - \alpha$ are sellers. Consequently, if the state of the good is ω, then for every $x \in [0, 1]$, the mass of agents with signals below x is $F(x|\omega)$. Let $\mathcal{E}(\alpha, v, f, g)$ denote this continuum economy.[18]

We now analyze this continuum economy from two perspectives. First, we view the situation as a pure exchange economy and compute a fully-revealing rational expectations equilibrium. It will be shown to be unique among fully-revealing REE. Second, we study the agents' behavior in this continuum setting when they participate in a double auction. A Bayes–Nash equilibrium in bidding strategies will be constructed, and we will demonstrate that it induces, for every state of the good, the same price of the good and the same allocation of the units among the agents as the unique fully-revealing

[18] There is no need to appeal to a continuum version of the law of large numbers. The continuum-economy equilibrium price function and bidding strategies we derive here are equilibria given the *definitions* of the continuum-economy payoff functions we employ. The payoff functions as well as the equilibria are significant because, as we will show, they are the limits of payoffs, equilibrium price functions and equilibrium bidding strategies in double auctions when the finite number of agents grows arbitrarily large. Put somewhat differently, our convergence results in large finite double auctions provide the justification for carrying out computations in the continuum economy in the manner that we do here.

REE. Proposition 3.1 below gives a general statement of this result, but its essence can be conveyed through an example, which we now provide.

Suppose that half the agents are buyers and half are sellers. Suppose further that given a signal, x, and the state ω, an agent's value for the good is $v(x, \omega) = x + \omega$. Finally, suppose that $\tilde{\omega}$ is uniform on $[0, 1]$, and that, given $\tilde{\omega} = \omega$, each agent's signal is drawn uniformly from $[0, \omega]$.[19,20]

Let us first view this as a pure exchange economy and attempt to find a fully-revealing REE. Since the only relevant uncertainty is the state of the good, $\tilde{\omega}$, a fully-revealing REE is a one-to-one map, $P : [0, 1] \to \mathbb{R}$, from the state of the good into the price of the good.[21]

If $P(\cdot)$ is fully revealing, then all agents can infer the state, ω, when deciding whether to buy or sell at price $P(\omega)$. Consequently, an agent with signal x will wish to leave the market *with* a unit if $x + \omega > P(\omega)$ and will wish to leave the market without a unit if $x + \omega < P(\omega)$. Because there are half as many units as agents, market clearing requires exactly half the agents to have a signal $x > P(\omega) - \omega$. Now, because each agent's signal is independently drawn uniformly from $[0, \omega]$, exactly half the agents have a signal above $\omega/2$. Consequently, we must have $P(\omega) - \omega = \omega/2$, that is, $P(\omega) = 3\omega/2$. Evidently, this is the only fully-revealing REE.

Given this REE, one might ask, "How does an agent's information get into the price?" As Milgrom (1981) does for finite one-sided markets using a uniform-price Vickrey auction, we provide an answer to this question by considering next the strategic bidding behavior of the agents in the continuum economy when they participate in a double auction.

So, put aside, for the moment, the above REE. Recall that in a double auction, the order of moves is as follows. Agents receive their signals, then simultaneously submit bids to the auctioneer. Each buyer submits a sealed bid indicating the maximum price he is willing to pay for the (single) unit he desires and each seller submits a sealed offer indicating the minimum price at which he is willing to sell the (single) unit he owns.

[19] This conditional density does not satisfy our full support and differentiability assumptions, nor does the value function satisfy assumption A.4. But these are not substantive issues. The particular value function and densities in this example merely permit a simple illustration of this section's main ideas.

[20] The value function and signal density employed here are taken from Milgrom (1981), who takes advantage of their special features to show that the price in a single-unit second-price auction can, for the winning bidder, be a sufficient statistic for the private information held by all competitors. In contrast, we employ these particular value and density functions for simplicity of calculation only. For our results, a sufficient statistic is unnecessary, and this example is representative of the general case.

[21] One might instead reserve the term "fully-revealing" for an REE that reveals not only ω but also reveals the distribution of signals accross agents. (Note that ω alone is enough to infer the distribution from which signals are drawn, but it is not enough to determine which signal any particular agent receives.) However, once ω is revealed, such additional information is redundant (as the argument below shows), and so our definition of a fully-revealing REE as one that reveals ω is the natural definition in this context.

The auctioneer then chooses a market-clearing price, p, and an agent leaves the market with one unit if his bid is above p, and leaves the market with zero units if his bid is below p, where all trade occurs at the price p.[22] Note that each agent submits a bid based solely on his private information. The market-clearing price is determined only *after* the bids have been submitted, at which time it is too late to change one's bid. Consequently, agents cannot condition their bids *directly* on the market price. On the other hand, in equilibrium, they are aware of the others' bidding strategies and take them into account.

We claim that it is a Bayesian–Nash equilibrium for all buyers to employ the bidding strategy, $b(x)$, and all sellers to employ the bidding strategy, $s(x)$, where $b(x) = s(x) = 3x$. That is, every agent submits a bid equal to three times the signal he receives.[23]

To see this, suppose that all agents do indeed employ the above bidding strategy. Now, because there are half as many units as agents, the double-auction market-clearing price, $P(\omega)$, in state ω will be determined by the median bid. Because all agents employ the same strictly increasing bidding function, the median bid is submitted by the agent with the median signal, namely, $x = \omega/2$ (recall that signals are uniform on $[0, \omega]$). The median bid is therefore $3\omega/2$, because all agents bid three times their signal. Consequently, $P(\omega) = 3\omega/2$ is the resulting double-auction price in state ω given the strategies $b(x) = s(x) = 3x$. So, in every state ω, the market price is the same as that in the fully-revealing REE.

To complete the Bayes–Nash equilibrium argument, consider an agent, say a buyer, whose signal is x. We must argue that this agent can do no better than to submit the bid $b(x) = 3x$, given that all other agents bid three times their signals. Now, while this buyer does not know the state, ω, he can reason as follows. *If* the state were ω, then, given the others' strategies, the median bid would be $3\omega/2$ regardless of this buyer's bid (given the continuum of agents). Hence, as above, the double-auction price will be $P(\omega) = 3\omega/2$. Now, by bidding $3x$, this buyer will be allocated a unit at price $3\omega/2$ if and only if his bid is above that price, or equivalently, if and only if $x > \omega/2$. But this is equivalent to $x + \omega > \omega/2 + \omega$, which itself is equivalent to $v(x, \omega) > P(\omega)$. That is, by bidding $3x$, this buyer, regardless of the state of the good ω, will end up purchasing a unit at price $P(\omega)$ if and only if his ex-post value, $v(x, \omega)$, exceeds $P(\omega)$. Clearly, this buyer cannot possibly improve upon this. A similar argument for sellers establishes the desired result. No agent can improve his payoff by bidding other than three times his signal.

Hence, the strategies $b(x) = s(x) = 3x$ form a Bayes–Nash equilibrium of the double auction and the resulting price, in every state ω, coincides with the price in the fully-revealing REE. Finally, observe that the Bayes–Nash allocation of units also corresponds

[22] We postpone consideration of bids that are equal to the price set by the auctioneer. Such bids will arise with probability zero in the present context and so, for now, we may safely ignore them.

[23] There are other equivalent equilibria. See Remark 1 below.

to the efficient outcome of the fully-revealing REE. This is because, in this equilibrium of the double auction, the units are allocated to the agents with the highest signals, and hence to the agents with the highest values. Therefore, the above bidding equilibrium of the double auction is the strategic equivalent of the fully-revealing REE.

The above Bayes-Nash equilibrium is symmetric in that buyers and sellers employ the same bidding function. This is to be expected, because, in the present continuum agent setting, no single agent's bid can affect the price. Consequently, the only asymmetry in buyer and seller preferences, namely that buyers prefer lower prices and sellers prefer higher prices, is irrelevant insofar as their optimal bidding behavior is concerned. Consequently, buyers and sellers have identical preferences over bids.[24]

The example generalizes. To prepare for the general result, define the αth percentile of $F(\cdot|\omega)$ as that signal, x, satisfying $F(x|\omega) = \alpha$, and denote it by $x(\omega)$. Let $\omega(x)$ denote the state of the good, ω, in which the αth percentile of $F(\cdot|\omega)$ is closest to x. Hence, $\omega(x(\omega_0)) = \omega_0$ for all $\omega_0 \in [0, 1]$, and if $x < x(0)$ or $x > x(1)$, then $\omega(x) = 0$ or 1, respectively.

Proposition 3.1. *Given $\alpha \in (0, 1)$, suppose that the continuum economy $\mathcal{E}(\alpha, v, f, g)$, with α buyers and $1 - \alpha$ sellers, satisfies A.1-A.3 Then,*

(i) there is a unique fully-revealing and efficient REE, namely $P(\omega) = v(x(\omega), \omega)$, and

(ii) the double auction possesses a Bayes–Nash equilibrium in symmetric, nondecreasing bidding strategies. In this equilibrium, each agent, buyer or seller, with signal x, submits a bid equal to $v(x, \omega(x))$. This Bayes–Nash equilibrium induces, in every state $\omega \in [0, 1]$, the same price and allocation as in the unique fully-revealing REE.[25]

The proof of Proposition 3.1 proceeds as in the example above and is omitted.

Remark 1. *It can be shown that, among fully-revealing Bayes–Nash equilibria, the equilibrium displayed in Proposition 3.1 is essentially unique. In particular, the strategies are uniquely determined over the range of signals, $x \in [x(0), x(1)]$. All nondecreasing functions that agree with $v(x, \omega(x))$ on this range also form a Bayes–Nash equilibrium,*

[24] To see this, note that because an agent's bid has no effect on the price, a buyer's bidding incentives would be unchanged if he were forced to buy the good, regardless of the market-clearing price, and his bid were treated as if he were a seller. Hence, a buyer with signal x prefers one bid over another if and only if a seller with signal x does.

[25] It is interesting to note that the proposition remains valid even in the pure common value case that is ruled out by our assumptions. Indeed, if $v(x, \omega) = v(\omega)$ and $v'(\omega) > 0$, then $P(\omega) = v(\omega)$ is a fully revealing and efficient REE, and it is supported by the Bayes-Nash equilibrium $b(x) = v(\omega(x))$ of the double auction. Indeed, given this bidding behavior, in any state ω a fraction α of all bids are below $b(x(\omega))$ and the remaining fraction are above. Hence, the equilibrium price in state ω is $P(\omega) = b(x(\omega)) = v(\omega)$. Consequently, all bidders are indifferent between winning and losing and these bidding strategies therefore constitute an equilibrium.

are interchangeable,[26] *and lead to the same price and allocation in every state of the good. For instance, the equilibrium, $b(x) = s(x) = 3x$, given in the above example is simply one among the set of equilibria: $b(x) = s(x) = 3x$, for $x \in [0, 1/2)$, and $b(x), s(x) \geq 3/2$ for $x \in [1/2, 1]$.*

Remark 2. *We can now offer some insight into the role of assumption A.4. Although the equilibrium of the continuum economy is essentially arbitrary outside the interval $[x(0), x(1)]$, our analysis of the finite economy draws particular attention to the continuum economy equilibrium in which $b(x) = v(x, 0)$ for $x < x(0)$, $b(x) = v(x, \omega(x))$ for $x \in [x(0), x(1)]$, and $b(x) = v(x, 1)$ for $x > x(1)$. Indeed, the equilibrium whose existence we ultimately establish for large finite double-auctions converges to this particular equilibrium, $b(\cdot)$, of the continuum economy, and our existence proof, which is partly constructive, makes use of the differentiability of $b(\cdot)$ at $x(0)$ and $x(1)$. Assumption A.4 ensures that $b(\cdot)$ is differentiable at these points independently of the density f which determines them. Loosely, differentiability at $x(0)$ is helpful because it permits us to employ knowledge of the rate at which equilibrium bids rise to the right of $x(0)$ in the continuum economy, where the equilibrium is uniquely determined, to construct the equilibrium to the left of $x(0)$ in the large finite economy; and similarly for $x(1)$. See, in particular, Lemma 1.8 in RP-s and its use in the proof of Theorem 1.25 there.*

Thus, the fully-revealing REE of this *idealized* market can be supported by a Bayes–Nash equilibrium in bidding functions. As we shall show, the double auction in the finite economy with sufficiently many agents and a sufficiently fine grid of prices generically possesses a Bayes-Nash equilibrium that approximates arbitrarily well the Bayes-Nash equilibrium described in Proposition 3.1. Consequently, this Bayes-Nash equilibrium for the finite economy will be arbitrarily close to a fully-revealing and efficient REE.

4. The Finite Economy with a Grid of Prices

We now return to the finite agent setting described in Section 2 in which there are n buyers and m sellers. The $n + m$ agents participate in a double auction. From now on, unless noted otherwise, we restrict the agents' bids/offers to the discrete set of prices, $\mathcal{P} = \{0, \Delta, 2\Delta, ..., \}$, where $\Delta > 0$ is the fineness of the grid. We now fully describe the double-auction rules.

4.1. Double Auction Rules

As described in Section 3, each buyer submits a bid indicating the maximum he is willing to pay for the good and each seller submits an offer indicating the minimum he

[26] That is, any combination of them are in equilibrium.

is willing to accept for the good. The auctioneer then sets a market-clearing price, P. However, because bids and offers must be chosen from the discrete grid of prices, $\mathcal{P}$, the market-clearing price will often coincide with one or more agent's bid. Hence, unlike Section 3, we must now specify precisely how the market clears in this case.

Given the $n+m$ bids and offers, let us first discuss carefully the choice of the market-clearing price. Because there are m units of the good in all, the market-clearing price must ensure that exactly m agents end up with one unit of the good. Since any agent with a bid or offer above the price leaves the market with one unit, and any agent with a bid or offer below the price leaves the market with zero units, a price can clear the market only if it lies between the mth and $(m+1)$st highest among all bids and offers.

Hereafter, it will be more convenient to refer to the sellers' offers as "bids." Consequently, each agent, whether a buyer or a seller, submits a bid to the auctioneer. Let b_k denote the kth highest among all $n+m$ bids. Hence, $b_1 \geq ... \geq b_{n+m}$, and the market-clearing price, P, must lie between b_m and b_{m+1}.

When $b_{m+1} < b_m$, the auction rules must specify which of the continuum of prices in $[b_{m+1}, b_m]$ is the market-clearing price. For our purposes, the particular choice is unimportant. The market-clearing price, P, can be determined by any prespecified nondecreasing function $\rho : \mathbb{R}^{m+n} \to \mathbb{R}_+$,[27] such that for all vectors of bids, $b_1 \geq ... \geq b_{m+n}$,

$$b_{m+1} \leq \rho(b_1, ..., b_{m+n}) \leq b_m.$$

Note that we do not require, although we permit, the range of the price function, $\rho(\cdot)$, to be restricted to the grid of prices $\mathcal{P}$. Thus, given □ $\in [0, 1]$, the □-double-auction pricing rule, $P =$ □$b_m + (1 -$ □$)b_{m+1}$, is a special case of our general formulation. From now on, we assume that some prespecified measurable function $\rho(\cdot)$ satisfying the above inequality is employed to determine the market-clearing price.

Having taken care of the market-clearing price, we now fully describe the trading rules. As already mentioned, agents with bids above the market-clearing price, P, end up with the good and agents with bids below P do not. Among the remaining agents with bids equal to P, precisely that number of them required to clear the market (i.e., (#goods)−(#bids above P)) are randomly (uniformly) chosen to end up with the good.[28] All trades are conducted at the market-clearing price P. This completes the

[27] The nondecreasing requirement permits us to place an upper bound, namely $v(1, 1) + \Delta$, on bids when searching for an equilibrium.

[28] According to this rationing rule, all agents, buyers and sellers, who bid the market price are treated symmetrically. This differs from the rationing rule described in Wilson (1985) and Rustichini, Satterthwaite and Williams (1994), where only the long side of the market is rationed equiprobably. The latter rule maximizes trade, but introduces an asymmetry. For example, when there are twice as many buyers as sellers, buyers are more likely to be rationed than sellers if all agents employ the same bidding function.

description of the double-auction rules.

Note, finally, that agents can affect the market price. For example, if for some vector of bids, $b_{m+1} < \rho(b_1, ..., b_{m+n}) < b_m$, then a buyer who bid b_m or higher can, by lowering his bid to just below $\rho(b_1, ..., b_{m+n})$, decrease the price at which he buys the unit and, similarly, a seller who bid b_{m+1} or lower can, by raising his bid to just above ρ, increase the price at which he sells his unit. Consequently, it is in the interest of buyers to "underbid," and of sellers to "overbid," in an attempt to affect the price to their advantage. As a result, some efficient trades may go unrealized. Our main result implies that such inefficiencies disappear as the number of agents grows and as the price grid becomes sufficiently fine.

4.2. Equilibrium

Let $\mathcal{E}(n, m, v, f, g, \Delta)$ denote the finite economy with n buyers and m sellers described in Section 2, but where the price grid is now $\mathcal{P} = \{0, \Delta, 2\Delta, ...\}$ instead of $\mathbb{R}_+$. The value function $v(x, \omega)$ and density functions $f(x|\omega)$ and $g(\omega)$ are as in Section 2, and should henceforth be assumed to satisfy assumptions A.1-A.4.

Suppose that $b : [0, 1] \to \mathcal{P}$ and $s : [0, 1] \to \mathcal{P}$ are nondecreasing bidding functions. Let $u^\beta(p, x|b(\cdot), s(\cdot))$ denote the double-auction expected payoff of a buyer whose signal is x, when he bids the price $p \in \mathcal{P}$, and all other $n-1$ buyers employ the bidding function $b(\cdot)$, and all m sellers employ the bidding function $s(\cdot)$. Similarly, let $u^\sigma(p, x|b(\cdot), s(\cdot))$ denote the double-auction expected payoff of a seller whose signal is x, when he bids the price $p \in \mathcal{P}$, and all n buyers employ $b(\cdot)$ and all other $m-1$ sellers employ $s(\cdot)$.[29]

The pair of nondecreasing bidding functions, $(\hat{b}(\cdot), \hat{s}(\cdot))$, constitutes a double-auction equilibrium of $\mathcal{E}(n, m, v, f, g, \Delta)$ if for every $x \in [0, 1]$,

$$\hat{b}(x) \text{ solves } \max_{p \in \mathcal{P}} u^\beta(p, x|\hat{b}(\cdot), \hat{s}(\cdot)),$$

and

$$\hat{s}(x) \text{ solves } \max_{p \in \mathcal{P}} u^\sigma(p, x|\hat{b}(\cdot), \hat{s}(\cdot)).$$

Hence, we restrict attention to equilibria in which all buyers employ the same bidding function and all sellers employ the same bidding function. A double-auction equilibrium is *nontrivial* if trade occurs with strictly positive probability.[30]

[29] The dependence of the payoff functions upon n and m is suppressed.

[30] There always exists a trivial no-trade equilibrium in which, for example, buyers always bid zero and sellers always bid $v(1, 1) + 1$.

5. The Failure of Monotone Best Replies

The double-auction equilibrium whose existence we ultimately establish is pure and nondecreasing. Standard existence results (e.g., Athey (2001)) do not apply because, in the present setting, agents need not possess monotone best replies when all other agents employ monotone strategies. There are two underlying reasons for this. The first is that, owing to the finite grid, an agent might be rationed with positive probability. The second is related to the fact that, in a finite economy, buyers and sellers are asymmetric: each can affect the price but wish to affect it in opposite directions. We now discuss these in more detail, beginning with the effect of rationing.

5.1. The Rationing/Grid-Size Effect

The reason that rationing can lead to the nonexistence of monotone best replies can be understood through a simple example. To isolate the effect of rationing, suppose that there are a continuum of agents, half of whom are buyers. This implies that no agent can affect the price and hence that buyers and sellers are symmetric. Consequently, the second effect mentioned above is not present.

Suppose the agents have private values. That is, suppose $v(x, \omega) = v(x)$. Consider a nondecreasing bidding function such that $b(x) = p_1$ when $x < 1/2$ and $b(x) = p_2 > p_1$ when $x \geq 1/2$. Suppose that all agents employ $b(\cdot)$, that p_1 and p_2 are consecutive grid-prices, and that $p_1 < v(x) < p_2$ for $x \in (\underline{x}, \bar{x})$.

Whenever a buyer's signal is in the interval $(\underline{x}, \bar{x})$, he strictly prefers to trade when the price is p_1 and strictly prefers not to trade when the price is p_2. Because ω and x are affiliated, it can happen that, when the buyer's signal is low (i.e., just above $\underline{x}$), the state is also low and so it may be very likely that the median signal is less than $1/2$. Hence, it is very likely that the market-price will be p_1, and very unlikely that it will be p_2. In this event, if the buyer bids p_1 he will be rationed with positive probability and consequently he strictly prefers to bid p_2 since this guarantees that he trades when the price is p_1 while it is very unlikely that he will have to trade at p_2.

On the other hand, when his signal increases to just below $\bar{x}$, it may be very likely that the state is high and so very likely that the median signal is above $1/2$, and consequently that the market-price will be p_2. The buyer will then strictly prefer to bid p_1 to avoid trading at p_2. Hence, the buyer's best reply is nonmonotone.[31]

The key feature of the example is that the steps of the bid function employed by all but the one agent are large. In Section 6 we show that when there are a continuum of agents and the grid of prices is sufficiently fine, this difficulty can be overcome. The idea is to first restrict the agents to monotone bidding functions and show that

[31] It should be noted that the rationing effect described here is present even when there are finitely many agents. The example demonstrates that this effect does not disappear in large markets.

an equilibrium bidding function in this restricted setting must have sufficiently narrow steps due to the fineness of the grid. When these steps are sufficiently narrow, an agent's signal has almost no effect on his assessment of the relative probability of one price versus an adjacent one. Hence, the negative effect of increasing one's signal, as described above, is outweighed by the positive effect of the increase in one's value; it is here where our assumption A.3, that there is a private value component to preferences, is crucial. Consequently, agents do in fact possess a monotone best reply and so the restriction to monotone bidding functions is ultimately not binding.

5.2. The Strategic/Small-Numbers Effect

For fixed bidding functions of the others, an agent's optimal bid for a given signal depends upon his assessment of the resulting distribution of the market-clearing price, which must lie between the mth and $m+1$st highest bid. Hence, when an agent's signal increases, the change in his optimal bid will depend upon how this increase affects his assessment of the distribution of the order statistics of the others' bids.

In general, one cannot guarantee that an agent will increase his bid in response to an increase in his signal. A standard method used to obtain the desired monotonicity is to first establish that an agent's signal is affiliated with the order statistics of the others' bids, since these help to determine the price. Now, if the other agents employ nondecreasing bidding functions, then because all the agents' *signals* are affiliated, the others' *bids* will be affiliated with any one agent's signal.[32] Based upon this, it would be entirely reasonable to conjecture the following.

($*$) *Each agent's signal is affiliated with the order statistics of the others' bids.*

In fact, ($*$) always holds in one-sided markets with symmetric buyers and symmetric affiliated signals (as in the models of Milgrom (1981), Milgrom and Weber (1982), and Pesendorfer and Swinkels (1997)), regardless of the market's size. This is because symmetric buyers can, in equilibrium, employ the same nondecreasing bidding function and hence, by results from Milgrom and Weber (1982), their bids will be symmetric and affiliated and so the order statistics of their bids will be affiliated with their signals, as required.

However, in our two-sided market, when there are finitely many agents, even though their signals are symmetrically distributed, buyers and sellers have different incentives and so will typically employ *distinct* bidding functions. Consequently, standard results do not apply and one cannot, in general, conclude that the order statistics of their bids will be affiliated with their signals. Indeed, as we now show, this need not hold (i.e., ($*$) can fail).[33]

[32] See Milgrom and Weber (1982).

[33] We thank Jeroen Swinkels for helping us develop a closely related example.

Figure 5.1: Failure of Affiliation

Consider a market with two buyers and one seller. The state of the good, $\tilde{\omega}$, is uniform on $[0, 1]$. If $\tilde{\omega} = \omega < 1/2$, the agents' signals are independent and uniform on $[0, 2/3]$, while if $\tilde{\omega} = \omega \geq 1/2$ their signals are independent and uniform on $[1/3, 1]$.[34]

Figure 5.1 depicts two nondecreasing bidding functions, $b(\cdot)$ for buyer 1 who receives the signal $\tilde{x}_1$, and $s(\cdot)$ for the seller, whose signal is $\tilde{x}_2$. The remaining "undecided" buyer receives the signal $\tilde{x}_3$, and is considering his bid given the strategies of the other two agents.

Consider two possible signals for the undecided buyer, $\underline{x}_3 \in (0, 1/3)$ and $\bar{x}_3 \in (2/3, 1)$, shown in Figure 5.1. Note that when $\tilde{x}_3 = \underline{x}_3$, the undecided buyer knows that the others' signals are uniform on $[0, 2/3]$, while when $\tilde{x}_3 = \bar{x}_3$ he knows that their signals are uniform on $[1/3, 1]$.

We now show that, even though both agents employ nondecreasing bidding functions, the undecided buyer's signal, $\tilde{x}_3$, is *not* affiliated with the first order statistic (i.e., the maximum) of their bids. Hence, $(*)$ fails.

Let $\tilde{Z} = \max(b(\tilde{x}_1), s(\tilde{x}_2))$. Affiliation requires, in particular, that the ratio

$$\frac{\Pr(\tilde{Z} = p_2 \mid \tilde{x}_3 = x_3)}{\Pr(\tilde{Z} = p_1 \mid \tilde{x}_3 = x_3)}$$

[34] The implied conditional density of x given ω is not strictly positive as required by A.1, and satisfies the MLRP weakly, but not strictly as in A.2. This is for simplicity only. This density can be approximated by one statisfying A.1 and A.2 while preserving the essential features of the example.

be nondecreasing in x_3. However, it is easy to see that

$$\frac{\Pr(\tilde{Z} = p_2 \mid \tilde{x}_3 = \underline{x}_3)}{\Pr(\tilde{Z} = p_1 \mid \tilde{x}_3 = \underline{x}_3)} \to 1 \quad \text{as } \varepsilon \to 0,$$

while

$$\frac{\Pr(\tilde{Z} = p_2 \mid \tilde{x}_3 = \bar{x}_3)}{\Pr(\tilde{Z} = p_1 \mid \tilde{x}_3 = \bar{x}_3)} \to 0 \quad \text{as } \varepsilon \to 0.$$

So, for ε small enough, the undecided buyer's signal is not affiliated with the first order statistic of the others' bids.

The lesson here is that in two-sided markets, wherein buyers and sellers should be expected to employ distinct strategies, the key affiliation condition, $(*)$, does not generally hold. Because of this, and because each agent can affect the price, one can construct examples in which the rationing effect described above is *not* present, and yet the only best reply for an agent against the monotone bidding functions of his opponents is nonmonotone.[35]

Fortunately, our interest lies primarily with large markets, i.e., those in which both the number of buyers, n, and the number of sellers, m, is large. As the market grows, the difference between the other agents' $m-1$st and $m+1$st highest bids converges to zero almost surely, and with it the strategic incentive to manipulate the price also converges to zero for both buyers and sellers. Thus, the strategic effect that is required for the present failure of monotone best replies vanishes.

6. The Main Result: Existence, Information Aggregation, Efficiency, and Price-Taking Behavior

Recall from Section 3 that $x(\omega)$ is the αth percentile of $F(\cdot|\omega)$, and that $\omega(x)$ is the state, ω, in which the αth percentile of $F(\cdot|\omega)$ is closest to x. Also, recall that V denotes the set of value functions satisfying A.3 and A.4. Finally, observe that if a property holds on a residual subset of a topological space (see Definition 2.1), the property is commonly interpreted as holding generically. Our main result is as follows.[36]

Theorem 6.1. *Fix any $\alpha \in (0,1)$. For every v in a residual subset of V and for every $\varepsilon > 0$, there exists $\bar{\Delta} > 0$ such that, for all Δ in a residual subset of $(0,\bar{\Delta})$ and for*

[35] Such an example involving two buyers and seven sellers is available from the authors upon request. The strategies in the example do not constitute an equilibrium. Whether there is an equilibrium in which at least one agent's best reply must be non-monotone is an interesting open question. We wish to thank Oren Rigbi for carrying out the necessary programming.

[36] The topology on the space of value functions is defined in Section 2, and the reals are endowed with their usual topology.

all $n_r, m_r \to_r \infty$ such that $n_r/(n_r + m_r) \to_r \alpha$, there is a sequence of buyer-seller nondecreasing bidding functions, $\{(b_r(\cdot), s_r(\cdot))\}_r$, such that for all large enough r,

(1) *$(b_r(\cdot), s_r(\cdot))$ is a nontrivial double-auction equilibrium of $\mathcal{E}(n_r, m_r, v, f, g, \Delta)$,*

(2) *$|b_r(x) - v(x, \omega(x))| < \varepsilon$ and $|s_r(x) - v(x, \omega(x))| < \varepsilon$, for all $x \in [0, 1]$,*

(3) *$\Pr(\left|\tilde{P}_r - v(x(\tilde{\omega}), \tilde{\omega})\right| < \varepsilon) > 1 - \varepsilon$, and*

(4) *$\Pr(\tilde{\beta}_r < \varepsilon) > 1 - \varepsilon$,*

where $\tilde{P}_r$ is the random market-clearing price and $\tilde{\beta}_r$ is the random fraction of agents who inefficiently end up with the good given the buyer-seller bidding functions $(b_r(\cdot), s_r(\cdot))$.

A detailed sketch of the proof is provided below. The complete proof can be found in Reny and Perry (2005).

Remark 3. *The existence result is provided by (1). The inequalities in (2) express a price-taking result, namely, that the strategic bidding behavior of the agents in these double-auction equilibria is approximated arbitrarily well by the rational price-taking behavior of the agents in the continuum economy of Section 3 (see Proposition 3.1 (ii)). The limit in (3) is an information aggregation result. It says that for almost every sufficiently fine grid of prices, and whenever the market is sufficiently large, the double auction possesses an equilibrium in which the market-clearing price is, with probability arbitrarily close to one, arbitrarily close to the fully-revealing REE price, $P(\omega) = v(x(\omega), \omega)$, of the continuum economy of Section 3 (see Proposition 3.1 (i)). The limit in (4) is an efficiency result. It says that such equilibria also have the property that the fraction of agents inefficiently ending up with the good is arbitrarily close to zero with probability arbitrarily close to one.*

Remark 4. *The order of limits is important. The size of the market must grow faster than the grid size shrinks to zero. This is because our demonstration that bidders possess monotone best replies to monotone bidding functions relies upon a key single-crossing property which we can establish for sufficiently large markets only when the price grid is sufficiently fine.*

6.1. Sketch of the Proof

A complete proof of Theorem 6.1 can be found in RP-s. We shall provide here a detailed sketch only, beginning with a broad outline.

6.1.1. Overview

A main ingredient of the proof is the use of a discrete grid of prices. Its purpose is to permit the application of a fixed point technique developed by Athey (2001), in which the jump points of the agents' monotone best-reply step-functions play an important role. But there is a major difficulty here. Even though players are symmetric in terms of information and values, they are necessarily asymmetric by virtue of their role as either a buyer or a seller. This, together with the ability of agents to manipulate the price, can preclude the existence of monotone best replies (see Section 5), rendering Athey's techniques inapplicable. The essence of the proof is to show that, when there are sufficiently many agents, the ability of any single agent to affect the price becomes weak, and, in combination with a sufficiently fine grid of prices, the monotone best reply property is restored. Once monotonicity is established, existence of equilibrium, information aggregation, and efficiency follow relatively easily.

The proof is broken into four parts, A-D. Part A shows that the continuum economy with a sufficiently fine grid of prices possesses a symmetric equilibrium in which all agents employ the same monotone step function (RP-s Proposition 1.3), that the steps grow narrow as the grid size, Δ, shrinks (RP-s Lemma 1.2), and that as Δ shrinks, the outcome converges to a fully revealing rational expectations equilibrium outcome (RP-s Proposition 1.7). It is here (see the proof of Lemma 1.1 of RP-s) where we make substantive use of the private value component of the agents' preferences.

Part B establishes an important property of symmetric equilibria of the continuum economy. Generically, except for bids that win with probability zero or one, all nonequilibrium bids are strictly suboptimal for all signals. This result is established by carefully considering the incentives of an agent whose signal makes him indifferent between one bid and another.

Part C establishes a symmetry property of equilibria for the continuum economy. Except for no-trade equilibria, the continuum economy has no monotone equilibria in which buyers and sellers employ distinct bidding functions.

Finally, part D considers the large but finite economy for a sufficiently small Δ. Here, we focus upon buyer-seller bidding functions that are fixed points of the convex-hull of a correspondence derived from the double-auction, but where buyers and sellers are restricted to monotone step function strategies that are forced to be near one another. The objective is to show that neither of these restrictions (monotonicity and closeness) is binding when there are sufficiently many buyers and sellers. The fixed points are then equilibria of the double auction.

Part D considers the limit of the above fixed points as the number of agents grows and the economy tends to the continuum economy. The effect of the limit is to eliminate the asymmetry between buyers and sellers and also to eliminate any agent's strategic effect on the price. The limit of the fixed points is shown to be an equilibrium, possibly asymmetric, of the continuum economy. But note that because the step functions of buyers and sellers are sufficiently close, they induce trade with probability bounded away from zero. Consequently, as shown in part C, the limit equilibrium involves buyers and sellers employing the same bidding function. Hence, far enough along the sequence, the restriction that strategies are sufficiently close ceases to be binding. Further, once the strategies are sufficiently close, the results from parts A-C show that the single-crossing condition holds and so the restriction to increasing strategies also ceases to bind. (A key part of this argument involves part B as follows. Because prices that are unused at the limit are strictly suboptimal, they remain strictly suboptimal far enough along the sequence. Consequently, no equilibrium price has a positive but vanishingly small probability of occurring as the number of agents grows. Such prices could be problematic because, conditional upon them, an agent's ability to affect the price need not vanish, and so the effective asymmetry between buyers and sellers also need not vanish. When both of these effects are present, single-crossing and monotonicity might fail.) Thus, far enough along the sequence, neither the restriction that buyer and seller strategies must be monotone nor the restriction that they must be close is binding, and so the fixed points along the sequence constitute an equilibrium of the double auction when there are sufficiently many buyers and sellers. This completes the overview of the proof. We next provide a detailed sketch.

6.1.2. Detailed Sketch

We begin by studying the limit economy, denoted by $\mathcal{E}(\alpha, v, f, g, \Delta)$, in which there are a continuum of agents, α of whom are buyers, $1-\alpha$ of whom are sellers, and where the grid of prices is $\mathcal{P} = \{0, \Delta, 2\Delta, ...\}$. As in Section 3, buyers and sellers have identical preferences here owing to the continuum of agents. One way to see this is to note that it is equivalent for a seller to submit a bid of b and sell when the market-clearing price is above b, or to first commit to selling his unit at the (currently unknown) market-clearing price, and submit the bid b as if he were a buyer, whereupon he would buy back his unit when the market-clearing price is below his bid. Buyers and sellers thus have the same preferences over bids.[37] Only the presence of the discrete grid of prices, $\mathcal{P}$, differentiates the present economy $\mathcal{E}(\alpha, v, f, g, \Delta)$ with the economy $\mathcal{E}(\alpha, v, f, g)$ of Section 3.

Part A of the proof establishes the existence of a nondecreasing bidding function, $b : [0, 1] \to \mathcal{P}$, that constitutes a double-auction equilibrium for $\mathcal{E}(\alpha, v, f, g, \Delta)$ when employed by all agents, so long as the grid of prices, $\Delta > 0$, is sufficiently small. All such equilibria are shown to be fixed points of a particular correspondence; this correspondence will play a role later on in the proof. In addition, it is shown that as the

[37] Note the reliance of this argument on the fact that, because of the continuum of agents, the seller's bid cannot affect the market-clearing price.

grid size tends to zero, all such equilibria converge to the fully revealing and efficient rational expectations equilibrium derived in Section 3.

Part A Recall that $x(\omega)$ is the αth percentile of $F(\cdot|\omega)$. That is, $F(x(\omega)|\omega) = \alpha$ for every $\omega \in [0,1]$. If every agent employs the same nondecreasing bidding function, $b : [0,1] \to \mathcal{P}$, then the market-clearing price in state $\omega \in [0,1]$ is $P(\omega) = b(x(\omega))$. This holds even though $b(\cdot)$ will not be strictly increasing owing to the grid of prices. The presence of the grid implies that ties in bids, and hence rationing, will occur with positive probability. The rationing probabilities are spelled out in Section 1.1 of RP-s, the precise details being unnecessary for this proof sketch.[38]

We may restrict attention to the finite price grid $\bar{\mathcal{P}} = \{0, \Delta, 2\Delta, ..., K\Delta\}$, where $(K-1)\Delta < v(1,1) \leq K\Delta$. Since our focus is on nondecreasing bidding functions and the distribution of signals is atomless, it is without additional loss to suppose that this bidding function is right continuous on $[0,1]$ and continuous at 1. Such nondecreasing functions taking values in the finite grid of prices $\bar{\mathcal{P}}$ are characterized by their jump points. Hence (following Athey (2001)), we may denote a bidding function by a nondecreasing vector, $\mathbf{x} = (x_1, ..., x_K) \in [0,1]^K$, where the agent bids $k\Delta$ for every $x \in [x_k, x_{k+1})$, and where $x_0 = 0$ and $x_{K+1} = 1$. Note that some of these intervals may be empty and so not all prices in the grid need be employed.

For example, consider the step function depicted in Figure 6.1, including the dotted steps (ignore the curved functions for the moment as well as the distinction between solid and dotted steps). There, $\bar{\mathcal{P}} = \{0, p_1, ..., p_{10}\}$, where $p_k = k\Delta$, and $\hat{x}_k$ is the signal at which the bid function jumps from p_{k-1} to p_k. Because $\hat{x}_6 = \hat{x}_7$, this function does not assume the value p_6. Suppose now that all agents employ this bidding function. Because the αth percentile of the distribution of signals always lies between $x(0)$ and $x(1)$, the market-clearing price, $P(\cdot) \equiv b_{\hat{\mathbf{x}}}(x(\cdot))$, will always lie between p_3 and p_9. Moreover, $P(\omega)$ is never p_6 since no agent ever submits this bid. Thus, the range of $P(\cdot)$ is $\{p_3, p_4, p_5, p_7, p_8, p_9\}$, which is indicated by the solid line portion of the bid function. The dotted line portion of the bidding function denotes bids that are made with positive probability by each agent (and hence submitted by a positive fraction of agents), but never occur as market-clearing prices. We will return to Figure 6.1 later on.

Let X_K denote the nonempty, compact, convex set of nondecreasing vectors of jump points $\mathbf{x} \in [0,1]^K$, and let $b_{\mathbf{x}}(\cdot)$ denote the monotone bidding function uniquely determined by $\mathbf{x} \in X_K$.[39] It will be convenient to sometimes refer to $\mathbf{x}$ as an agent's

[38] A useful feature of the tie-break rule we employ is that it renders buyers and sellers precisely symmetric in this continuum-agent setting, whereas the standard tie-break rule, which maximizes the number of trades, would not. Symmetry is useful in establishing the existence of an equilibrium, although we suspect that, with some effort, it could be dispensed with.

[39] That is, $b_{\mathbf{x}} : [0,1] \to \bar{\mathcal{P}}$ is the unique nondecreasing right-continuous function that is continuous at

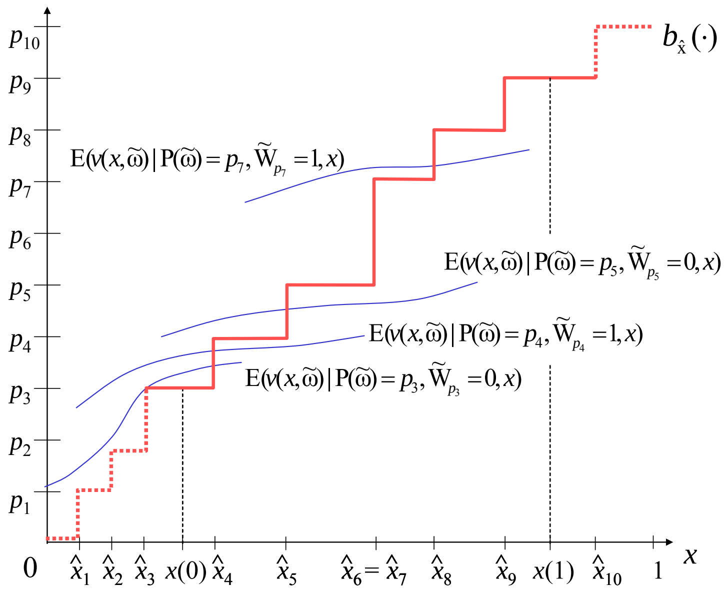

Figure 6.1: Continuum Economy Equilibrium

strategy.

For $\mathbf{x} \in X_K$, let $u(p, x|\mathbf{x})$ denote the double-auction payoff of an agent (buyer or seller) in $\mathcal{E}(\alpha, v, f, g, \Delta)$, when the agent's signal is x and he bids the price p, and all other agents employ the strategy $\mathbf{x}$. Given $\hat{\mathbf{x}} \in X_K$, the monotone bidding function $b_{\hat{\mathbf{x}}}(\cdot)$ is a double-auction equilibrium for $\mathcal{E}(\alpha, v, f, g, \Delta)$ when, for every $x \in [0, 1]$,

$$b_{\hat{\mathbf{x}}}(x) \text{ solves } \max_{p \in \mathcal{P}} u(p, x|\hat{\mathbf{x}}).$$

A related and very useful maximization problem is the following.

$$\max_{\mathbf{y} \in X_K} \int_0^1 u(b_{\mathbf{y}}(x), x|\mathbf{x}) f(x) dx, \tag{6.1}$$

where $f(x) = \int_0^1 f(x|\omega) g(\omega) d\omega$ is the ex-ante density over the agent's signal. In this problem, the agent chooses his bidding function ex-ante, before finding out his signal. Moreover, the agent is restricted to choosing a *nondecreasing* bidding function since $\mathbf{y} \in X_K$ produces the nondecreasing function $b_{\mathbf{y}}(\cdot)$. Given $\mathbf{x} \in X_K$, let $B(\mathbf{x})$ denote the set of $\mathbf{y} \in X_K$ solving (6.1).

$x = 1$ and jumps from $(k-1)\Delta$ to $k\Delta$ at x_k.

Note that $b_{\hat{\mathbf{x}}}(\cdot)$ is a double-auction equilibrium if $\hat{\mathbf{x}}$ is a fixed point of $B(\cdot)$ *and* if restricting the agent's ex-ante bidding function to be nondecreasing is not binding.

It is not difficult to show that $B(\cdot)$ is nonempty-valued and upper hemicontinuous, but it need not be convex-valued. However, Kakutani's theorem guarantees the existence of a fixed point of the correspondence, $coB(\cdot)$, whose value for any $\mathbf{x} \in X_K$ is the convex hull, $coB(\mathbf{x})$, of the set $B(\mathbf{x})$.

The first important result from part A of the proof is the following.

RP-s Proposition 1.3. *There exists $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$, the following statements are equivalent.*
(a) $b_{\hat{\mathbf{x}}}(\cdot)$ *is a double-auction equilibrium for* $\mathcal{E}(\alpha, v, f, g, \Delta)$.
(b) $\hat{\mathbf{x}} \in B(\hat{\mathbf{x}})$.
(c) $\hat{\mathbf{x}} \in coB(\hat{\mathbf{x}})$.

The proposition states that when the grid of prices is sufficiently fine, every fixed point, $\hat{\mathbf{x}}$, of $coB(\cdot)$ is a fixed point of $B(\cdot)$ and $b_{\hat{\mathbf{x}}}(\cdot)$ is a double-auction equilibrium. We now sketch its proof.

Because the implications (a) $\Longrightarrow$ (b) $\Longrightarrow$ (c) are obvious, we need only argue that (c) $\Longrightarrow$ (a). This follows from Lemmas 1.1 and 1.2 in RP-s. Recall that $P(\omega) \equiv b_{\hat{\mathbf{x}}}(x(\omega))$ is the market-clearing price when the state is ω and all agents employ $b_{\hat{\mathbf{x}}}(\cdot)$. The lemmas are as follows.

RP-s Lemma 1.1. There is a step-size, $\eta > 0$, such that for all K and all $\mathbf{x} \in X_K$, if each step of $P(\cdot) = b_{\mathbf{x}}(x(\cdot))$ has length less than η, then (i) $B(\mathbf{x})$ is convex, and (ii) if all other agents employ $b_{\mathbf{x}}(\cdot)$, then for some $\mathbf{y} \in X_K$ the nondecreasing bidding function $b_{\mathbf{y}}(\cdot)$ maximizes the agent's ex-ante (and interim) payoff among *all* measurable bidding functions, nondecreasing or not.

RP-s Lemma 1.2. If Δ is sufficiently small and $\hat{\mathbf{x}} \in coB(\hat{\mathbf{x}})$, the function $P(\cdot) \equiv b_{\hat{\mathbf{x}}}(x(\cdot))$ has arbitrarily narrow steps.

Together, the lemmas yield (c) $\Longrightarrow$ (a). To see this, note that for sufficiently small Δ, Lemma 1.2 implies that $P(\cdot) \equiv b_{\hat{\mathbf{x}}}(x(\cdot))$ has arbitrarily narrow steps. Consequently, by Lemma 1.1 (i), $\hat{\mathbf{x}} \in coB(\hat{\mathbf{x}})$ implies $\hat{\mathbf{x}} \in B(\hat{\mathbf{x}})$, so that $b_{\hat{\mathbf{x}}}(\cdot)$ maximizes the agent's ex-ante payoff among nondecreasing bidding functions when all others employ $b_{\hat{\mathbf{x}}}(\cdot)$. But then $b_{\hat{\mathbf{x}}}(\cdot)$ is an equilibrium because, by Lemma 1.1 (ii), such a nondecreasing function is a best reply.

We now sketch the proofs of the lemmas.

Proof of RP-s Lemma 1.1 (sketch). By Athey (2001), to prove (i) and (ii) it suffices to show that $u(p, x|\mathbf{x})$ satisfies single crossing in (p, x).

Because there are a continuum of agents, the strategic/small-numbers effect discussed in Section 5.1 is not present here. Hence, any failure of single crossing is due to the first effect discussed in Section 5.1, namely, the rationing/grid-size effect. This effect, we remind the reader, can lead to a failure of single-crossing as follows. Consider a buyer with signal x for whom it is profitable to buy the good if the price is p_0 but not if the price is the next higher grid-price, p_1. If the likelihood that the market-clearing price is p_1 versus p_0 is not too high, the buyer might be better off bidding p_1, so as to guarantee winning the good when the market-clearing price is p_0, than bidding p_0 and being rationed with positive probability. But if the likelihood that the price is p_1 increases when his signal increases, it might become better for him to bid the lower price p_0 at the higher signal, violating single-crossing. As we now discuss, this difficulty vanishes when the market-clearing price function, $P(\cdot)$, has sufficiently narrow steps.

For example, suppose $I_0 = [\bar{\omega} - \varepsilon, \bar{\omega}]$ and $I_1 = [\bar{\omega}, \bar{\omega} + \varepsilon]$ are adjacent intervals of states and that $P(\omega) = p_i$ precisely on I_i for $i = 0$ and 1. If the length of each I_i, namely $\varepsilon > 0$, is sufficiently small, then given x, the relative likelihood of the event $P(\tilde{\omega}) = p_1$ versus $P(\tilde{\omega}) = p_0$, namely

$$\frac{\int_{I_1} f(x|\omega)g(\omega)d\omega}{\int_{I_0} f(x|\omega)g(\omega)d\omega},$$

is virtually independent of the agent's signal, x. Indeed, this ratio converges to unity for every x as ε tends to zero.[40] Consequently, because the agent's value has a strict private value component, and because x and ω are affiliated, an increase in the agent's signal leads to a strict increase in his value for the good conditional upon the occurrence of either price, p_1 or p_0, and this strict increase outweighs the arbitrarily small negative effect of the increased likelihood of the higher price. Hence, the difference in his payoff from bidding p_1 versus p_0 strictly increases, implying that single crossing, and hence also (i) and (ii), hold. ■

Proof of RP-s Lemma 1.2 (sketch). If the lemma fails, then Δ can be arbitrarily small and yet some price $p^\Delta \in [0, v(1,1) + \Delta]$ is assumed by $P^\Delta(\cdot) \equiv b_{\hat{\mathbf{x}}^\Delta}(x(\cdot))$ on an interval of states whose length is bounded away from zero as $\Delta \to 0$. Since $\hat{\mathbf{x}}^\Delta \in coB(\hat{\mathbf{x}}^\Delta)$, $\hat{\mathbf{x}}^\Delta$ is the convex combination of finitely many members of $B(\hat{\mathbf{x}}^\Delta)$ and therefore for some $\mathbf{y}^\Delta \in B(\hat{\mathbf{x}}^\Delta)$ given positive weight in the convex combination, $b_{\mathbf{y}^\Delta}(\cdot)$ assumes the value p^Δ on an interval of signals, $[\underline{y}^\Delta, \bar{y}^\Delta)$, whose length is bounded away from zero.

Because $\mathbf{y}^\Delta \in B(\hat{\mathbf{x}}^\Delta)$, there can be no nondecreasing bidding function that is strictly better than $b_{\mathbf{y}^\Delta}(\cdot)$. Consequently, (i) an agent must weakly prefer bidding p^Δ to bidding the next higher grid-price p_+^Δ when his signal is $\bar{y}^\Delta$ and (ii) he must weakly prefer bidding

[40] In general, the rates at which the intervals shrink need not be the same. The complete proof in RP-s takes this into account.

p^Δ to bidding the next lower grid-price p_-^Δ when his signal is $\underline{y}^\Delta$, where we shall employ the convention that $p_-^\Delta = p^\Delta = 0$ if $p^\Delta = 0$.

Now, if an agent bids p_+^Δ instead of p^Δ when his value is $\bar{y}^\Delta$, there are two effects. First, he now ends up with the good when the price is p^Δ and he would have been rationed had he bid p^Δ. Second, he now ends up with the good when the price is p_+^Δ and he is not rationed given his bid of p_+^Δ. In the first case his expected value of the good is

$$E(v(\bar{y}^\Delta, \tilde{\omega})|P^\Delta(\tilde{\omega}) = p^\Delta, \bar{y}^\Delta, \text{ and lose at } p^\Delta) \geq E(v(\bar{y}^\Delta, \tilde{\omega})|P^\Delta(\tilde{\omega}) = p^\Delta, \bar{y}^\Delta), \quad (6.2)$$

where the inequality follows because being rationed is good news about the state.[41] In the second case, his expected value of the good is

$$E(v(\bar{y}^\Delta, \tilde{\omega})|P^\Delta(\tilde{\omega}) = p_+^\Delta, \bar{y}^\Delta, \text{ and win at } p_+^\Delta) \geq E(v(\bar{y}^\Delta, \tilde{\omega})|P^\Delta(\tilde{\omega}) = p^\Delta, \bar{y}^\Delta), \quad (6.3)$$

where the inequality follows because every state in which the price is p_+^Δ is higher than every state in which the price is p^Δ.

Now, if the the left-hand side of (6.2) is greater than p^Δ and the left-hand side of (6.3) is greater than p_+^Δ, then each of the two effects results in a net increase in the agent's payoff. In particular then, if $E(v(\bar{y}^\Delta, \tilde{\omega})|P^\Delta(\tilde{\omega}) = p^\Delta, \bar{y}^\Delta) > p_+^\Delta$ the agent is strictly better off bidding p_+^Δ than p^Δ, contradicting (i). Hence, an implication of (i) is,

$$E(v(\bar{y}^\Delta, \tilde{\omega})|P^\Delta(\tilde{\omega}) = p^\Delta, \bar{y}^\Delta) \leq p_+^\Delta.$$

By a similar argument, an implication of (ii) is,

$$E(v(\underline{y}^\Delta, \tilde{\omega})|P^\Delta(\tilde{\omega}) = p^\Delta, \underline{y}^\Delta) \geq p_-^\Delta.$$

Since p_+^Δ and p_-^Δ differ by at most two grid points, we obtain

$$E(v(\bar{y}^\Delta, \tilde{\omega})|P^\Delta(\tilde{\omega}) = p^\Delta, \bar{y}^\Delta) \leq E(v(\underline{y}^\Delta, \tilde{\omega})|P^\Delta(\tilde{\omega}) = p^\Delta, \underline{y}^\Delta) + 2\Delta,$$

which is impossible for Δ sufficiently small since $\bar{y}^\Delta - \underline{y}^\Delta$ is positive and bounded away from zero, and $v_x(\cdot, \cdot) > 0$ and continuous. ■

Recall from Section 3 that $\omega(x)$ is the state, ω, in which the αth percentile of $F(\cdot|\omega)$ is closest to x. The second important result from part A of the proof is the following.

[41] Given the uniform rationing rule and conditional upon one's bid being equal to the market-clearing price, an agent is more likely to end up with a unit when the state is low than when it is high. Thus, after observing the market price, an agent whose bid coincides with it increases his value assessment upon finding out that he will not end up with the good, and reduces his value assessment upon finding out that he will end up with the good. This is a form of the winner's curse.

RP-s Proposition 1.7. *For each $\Delta > 0$, let $b^{\Delta}(\cdot)$ be a double-auction equilibrium for $\mathcal{E}(\alpha, v, f, g, \Delta)$. Then, as $\Delta \to 0$, $b^{\Delta}(x) \to v(x, \omega(x))$ uniformly on $(x(0), x(1))$ and $P^{\Delta}(\omega) \to v(x(\omega), \omega)$ uniformly on $(0, 1)$. Consequently, the market-clearing price function converges uniformly to the unique fully-revealing rational expectations equilibrium of the limit economy $\mathcal{E}(\alpha, v, f, g)$ with a continuum of agents and prices.*

Thus, as the grid of prices becomes arbitrarily fine, limits of double-auction equilibria for $\mathcal{E}(\alpha, v, f, g, \Delta)$ converge uniformly to the essentially unique fully-revealing and efficient REE of the continuum economy with a continuum of prices from Section 3. The logic behind this result is straightforward, as we now demonstrate.

Proof of RP-s Proposition 1.7 (sketch). Fix a signal $\bar{x} = x(\bar{\omega})$ for some $\bar{\omega} \in (0, 1)$. For simplicity, we will argue only that

$$\lim_{\Delta \to 0} b^{\Delta}(\bar{x}) = v(\bar{x}, \bar{\omega}),$$

leaving the uniformity of the convergence to RP-s.

Let $p^{\Delta} = b^{\Delta}(\bar{x})$. Because $p^{\Delta} = b^{\Delta}(\bar{x}) = b^{\Delta}(x(\bar{\omega})) = P^{\Delta}(\bar{\omega})$ and $\bar{\omega} \in (0, 1)$, the interval of states, I^{Δ}, on which $P^{\Delta}(\cdot)$ is p^{Δ} is nonempty and contains $\bar{\omega}$. As we have seen above, I^{Δ} shrinks to a point as $\Delta \to 0$. Since $\bar{\omega} \in I^{\Delta}$ for all Δ, I^{Δ} must shrink to $\bar{\omega}$.

Now, suppose, by way of contradiction, that $\lim_{\Delta \to 0} b^{\Delta}(\bar{x}) < v(\bar{x}, \bar{\omega})$. Then, when the agent's signal is $\bar{x}$, his bid of $p^{\Delta} = b^{\Delta}(\bar{x})$ loses whenever the market-clearing price is above p^{Δ}. But $P^{\Delta}(\omega)$ is above p^{Δ} only when $\omega > \bar{\omega}$. Hence, the agent strictly prefers to win when $P^{\Delta}(\omega)$ is between p^{Δ} and $v(\bar{x}, \bar{\omega})$. Moreover, because I^{Δ} shrinks to $\bar{\omega}$, for Δ small the agent knows that the state must be very close to $\bar{\omega}$ when $P(\omega) = p^{\Delta}$. Hence, the agent strictly prefers to win when $P(\omega) = p^{\Delta}$ as well. Consequently, increasing his bid from p^{Δ} to any price between p^{Δ} and $v(\bar{x}, \bar{\omega})$, strictly increases his payoff, a contradiction. Hence, $\lim_{\Delta \to 0} b^{\Delta}(\bar{x}) \geq v(\bar{x}, \bar{\omega})$. A similar argument establishes the opposite inequality. ■

This concludes part A of the proof. However, to further clarify part A, and to prepare for part B, it will be helpful to reconsider Figure 6.1. The step function depicted there, $b_{\hat{\mathbf{x}}}(\cdot)$, is a double-auction equilibrium for $\mathcal{E}(\alpha, v, f, g, \Delta)$. Also, $\bar{\mathcal{P}} = \{0, p_1, ..., p_{10}\}$, where $p_k = k\Delta$, and $\hat{x}_k$ is the signal at which the bid function jumps from p_{k-1} to p_k. For the moment, ignore the curved lines and focus only upon the step function. As we have already observed, the range of $P(\cdot)$ is $\{p_3, p_4, p_5, p_7, p_8, p_9\}$, as indicated by the solid line portion of the bid function.

Now, because $P(\cdot)$ is never below p_3, any bid below p_3 is a losing bid in the sense that the agent is sure to end up without the good. Similarly, any bid above p_9 is a winning bid. Consequently, changing $b_{\hat{\mathbf{x}}}(\cdot)$ by replacing any bid below p_3 with any other bid below p_3 and replacing any bid above p_9 with any other bid above p_9 (even bids above p_{10}) yields a distinct bidding function that is also an equilibrium. Indeed, this new equilibrium is outcome equivalent to $b_{\hat{\mathbf{x}}}(\cdot)$. Thus, the dotted portion of the bidding function displayed in the figure is not uniquely determined. On the other hand, as we will argue later, the dotted portion of the bidding function is significant because it happens to be the limit of the sequence of large finite economy equilibria that will eventually be shown to exist. Along this sequence of equilibria, the dotted line prices arise as market-clearing prices with *positive, but vanishingly small*, probability. More on this later.

The jump points, $\hat{x}_k$, in the figure are signals at which the agent is indifferent between bidding p_{k-1} and p_k. Given that all agents employ $b_{\hat{\mathbf{x}}}(\cdot)$, it is important to understand the effects on an agent's payoff of increasing his bid from p_{k-1} to p_k.

Fix an agent. For every p in the range of $P(\cdot) \equiv b_{\hat{\mathbf{x}}}(x(\cdot))$, let $\tilde{W}_p$ denote the zero-one random variable indicating whether the agent ends up with a unit of the good (i.e., $W_p = 1$) conditional upon his bid being p and conditional upon the market-clearing price also being p. Because rationing is uniform, $\tilde{W}_p$ assumes both values, zero and one, with positive probability for each p in the range of $P(\cdot)$.

Recall from the above sketch of the proof of RP-s Lemma 1.2, when an agent with signal x increases his bid from $\underline{p}$ to $\bar{p}$ there are two changes to his payoff.

I. First, when $P(\omega) = \underline{p}$, he now ends up with the good (i.e., wins) with probability one instead of with probability strictly less than one. Hence, conditional on these additional wins, the change in his payoff is

$$E(v(x,\tilde{\omega})|P(\tilde{\omega}) = \underline{p}, \widetilde{W_{\underline{p}}} = 0, x) - \underline{p}. \tag{6.4}$$

II. Second, when $P(\omega) = \bar{p}$ he now ends up with the good with positive probability instead of with probability zero. Hence, conditional on these additional wins, the change in his payoff is

$$E(v(x,\tilde{\omega})|P(\tilde{\omega}) = \bar{p}, \widetilde{W_{\bar{p}}} = 1, x) - \bar{p}. \tag{6.5}$$

The overall change in the agent's payoff is a weighted sum of these two changes, where the weights are the probabilities of the additional winning events given the agent's signal, x. Now, if this agent is indifferent between $\underline{p}$ and $\bar{p}$, the overall change in his payoff, (i.e., the weighted sum of (6.4) and (6.5)), must be zero. Consequently, (6.4) and (6.5) must have opposite signs, as can be seen in Figure 6.1. (Let us ignore, for the moment, the possibility that both are zero.)

For example, when $\underline{p} = p_3$, $\bar{p} = p_4$ and $x = \hat{x}_4$, (6.4) is positive because (see the figure)

$$E(v(\hat{x}_4,\tilde{\omega})|P(\tilde{\omega}) = p_3, \widetilde{W_{p_3}} = 0, \hat{x}_4) > p_3,$$

and (6.5) is negative because

$$E(v(\hat{x}_4,\tilde{\omega})|P(\tilde{\omega})=p_4,\widetilde{W}_{p_4}=1,\hat{x}_4)<p_4.$$

Hence, when the agent's signal is $\hat{x}_4$, he strictly prefers to win when the price is p_3 and strictly prefers to lose when the price is p_4, but on balance he is indifferent between the two bids. Note then that if there were an unused price (like p_6) between p_3 and p_4, the agent with signal $\hat{x}_4$ could strictly improve his payoff by submitting this bid rather than p_3 since he would then win when the price is p_3, which strictly increases his payoff, and he would still lose when the price is p_4. The general conclusion is this:

> When an equilibrium bid function jumps from $\underline{p}$ to $\bar{p}$ at signal x and (6.4) is positive and (6.5) is negative, $\underline{p}$ and $\bar{p}$ must be consecutive prices in $\mathcal{P}$. (6.6)

On the other hand, when $\underline{p}=p_5$, $\bar{p}=p_7$ and $x=\hat{x}_6=\hat{x}_7$, (6.4) is negative and (6.5) is positive (see the figure). Hence, when the agent's signal is $\hat{x}_6=\hat{x}_7$, he strictly prefers to lose when the price is p_5, and he strictly prefers to win when the price is p_7. But again, on balance, he is just indifferent between the two bids. In this case, there is an unused price, namely p_6, between p_5 and p_7. This is consistent with equilibrium because, if the agent bids p_6 instead of p_5, he would then win for sure when the price is p_5, strictly decreasing his payoff because (6.4) is negative, while he would still lose when the price is p_7. The general conclusion is this:

> When an equilibrium bid function jumps from $\underline{p}$ to $\bar{p}$ at signal x and (6.4) is negative and (6.5) is positive, any price strictly between $\underline{p}$ and $\bar{p}$ is strictly suboptimal. (6.7)

Note that, by affiliation and the strict private value component, the functions of x in (6.4) and (6.5) are strictly increasing. Consequently, the change in the agent's payoff from increasing his bid from $\underline{p}$ to $\bar{p}$ is a weighted sum of two strictly increasing functions of his signal. But, as shown in part A of the proof, when the grid is sufficiently fine, the relative weights on the two functions, which are essentially the relative weights of the two prices $\underline{p}$ and $\bar{p}$, are virtually independent of the agent's signal. Hence, the difference in the agent's payoff from bidding high versus low is strictly increasing.

Consequently, assuming a sufficiently fine grid, apart from the indifference at jump points, it should now be evident from the above discussion that in the equilibrium of Figure 6.1, the agent is strictly optimizing by employing $b_{\hat{\mathbf{x}}}(x)$ for all signals $x \in (\hat{x}_3,\hat{x}_{10})$. Hence, except for the indeterminacy of bids for signals below $\hat{x}_3$ and above $\hat{x}_{10}$, the equilibrium bidding function is an essentially unique best reply against itself. This uniqueness will be important later on.

Finally, we return to the presumption underlying the above discussion that (6.4) and (6.5) are not zero. Part B of the proof establishes that, generically, (6.4) and (6.5) are indeed *not* zero, and so the equilibrium strategy, $b_{\hat{\mathbf{x}}}(\cdot)$, is a unique best reply against itself.

Part B Part B of the proof establishes that when the price grid is sufficiently fine, the following statement holds *generically.*

(†) In every double-auction equilibrium $b_{\hat{\mathbf{x}}}(\cdot)$ for $\mathcal{E}(\alpha, v, f, g, \Delta)$, if $\underline{p}$ and $\bar{p}$ are in the range of the market-clearing price function $P(\cdot)$, and $b_{\hat{\mathbf{x}}}(\cdot)$ jumps up from $\underline{p}$ to $\bar{p}$ at the signal x, then (6.4) and (6.5) have *strictly* opposite signs. In particular, neither can be zero.

As we have already seen, Figure 6.1 has been drawn so that (†) holds true. Also, note that in order for (†) to fail, *both* (6.4) and (6.5) would have to be zero because the indifference of the agent whose signal is the jump point requires a positively weighted sum of (6.4) and (6.5) to be zero. Part B of the proof establishes, using techniques from differential topology (i.e., Sard's theorem), that it would require an unlikely coincidence in the choice of the grid size, Δ, and the value function, $v(x, \omega)$, in order for an equilibrium to exhibit equality in both (6.4) and (6.5) at any one of its jump points. In particular, were such a coincidence to occur, it would be possible to eliminate it by perturbing Δ and $v(\cdot, \cdot)$ ever so slightly. It is shown that such coincidences can occur only on a non-generic set (i.e., the complement of a residual set) of grid sizes and value functions.

Condition (†) is important because, together with the results of part C, it permits us to establish the single-crossing property in a range of prices and for certain sequences of buyer-seller bidding functions in large double-auctions. This will be explained in more detail following part C.

Part C Given the symmetry of the agents in the continuum economy, up to now we have focused on equilibria in which *all* agents, buyers and sellers, employ the same bidding function. When there are finitely many agents, buyers and sellers are not symmetric and so they will employ distinct equilibrium bidding functions. On the other hand, when the finite number of agents is *large* we expect the difference in buyer and seller strategies to be small.

Hence, we are led to reconsider the continuum economy, now allowing all buyers to employ the same bidding function, $b(\cdot)$ say, and all sellers to employ the same bidding function, $s(\cdot)$ say, but where $b(\cdot)$ and $s(\cdot)$ need not be the same. A pair of nondecreasing buyer-seller bidding functions, $(b(\cdot), s(\cdot))$, is defined to be a double-auction equilibrium for $\mathcal{E}(\alpha, v, f, g, \Delta)$ if when all buyers employ $b(\cdot)$ and all sellers employ $s(\cdot)$, each agent's strategy is a best reply against the induced market-clearing price function and given the double-auction rationing rule.[42] Such equilibria are *nontrivial* if trade occurs with positive probability. When $(b(\cdot), b(\cdot))$ is a double-auction equilibrium, we will sometimes simply say that $b(\cdot)$ is a double-auction equilibrium, maintaining consistency with our previous terminology. For emphasis, we may also say that $b(\cdot)$ is a symmetric double-auction equilibrium.

Part C shows that permitting buyers and sellers to employ distinct strategies adds essentially no new equilibria except those in which there is no trade. Part C establishes that all nontrivial nondecreasing equilibria of the continuum economy are outcome-equivalent to an equilibrium in which all agents employ the same bidding function. We now provide a sketch of the details.

Observe that a pair of buyer-seller bidding functions will induce trade with probability zero only if for some pair of prices in the grid, $\Delta k < \Delta k'$ say, the buyers never bid above Δk and the sellers never bid below $\Delta k'$. In terms of jump point vectors, this implies that the kth coordinate of a buyer's jump point vector is $x_k = 1$, and the k'th coordinate of a seller's jump point vector is $y_{k'} = 0$. But because $k < k'$, this implies that $x_{k'} = 1$ and $y_k = 0$. Hence, $y_k - x_k = y_{k'} - x_{k'} = 1$. Therefore, trade occurs with probability zero only if the distance between some coordinate of the buyer's jump–point vector and the same coordinate of the seller's jump point vector is one. With this in mind, let us say that two step functions are ε-close when their jump-point vectors are within ε of one another coordinate by coordinate. Finally, note that if $\varepsilon \in [0, 1)$, then the set of ε-close pairs of buyer-seller bidding functions induce trade with positive probability bounded away from zero.

Part C defines, for $\varepsilon \in [0, 1)$, a continuum-economy best-response-like correspondence, $\Psi_\varepsilon(\cdot)$, from pairs of ε-close buyer-seller bidding functions, $(b(\cdot), s(\cdot))$, into subsets of them. The correspondence $\Psi_\varepsilon(\cdot)$ has the property that if $b(\cdot)$ and $s(\cdot)$ are ε-close and $(b(\cdot), s(\cdot))$ is a double-auction equilibrium for the continuum economy $\mathcal{E}(\alpha, v, f, g, \Delta)$, then $(b(\cdot), s(\cdot))$ is a fixed-point of $\Psi_\varepsilon(\cdot)$. Part C demonstrates that, generically, and when Δ is sufficiently small, *every* fixed point, $(\bar{b}(\cdot), \bar{s}(\cdot))$, of $co\Psi_\varepsilon(\cdot)$ is such that both $(\bar{b}(\cdot), \bar{b}(\cdot))$ and $(\bar{s}(\cdot), \bar{s}(\cdot))$ are double-auction equilibria of $\mathcal{E}(\alpha, v, f, g, \Delta)$, i.e., that both $\bar{b}(\cdot)$ and $\bar{s}(\cdot)$ are equilibria. Further, these latter two equilibria are outcome-equivalent in the sense that they induce the same market-clearing price function, and for every price p in its range, the interval on which $\bar{b}(\cdot)$ is p coincides with that on which $\bar{s}(\cdot)$ is p. Hence, any double-auction equilibrium for $\mathcal{E}(\alpha, v, f, g, \Delta)$ for which trade occurs with positive probability is outcome-equivalent to some symmetric equilibrium in which buyers and sellers employ the same bidding function.

The intuition for this is as follows. When Δ is sufficiently small, the arguments employed in Part A can also be employed here to show that fixed points, $(\bar{b}(\cdot), \bar{s}(\cdot))$, of

[42] Part C of the proof in RP shows explicitly how such a pair of strategies induces a market-clearing price function.

$\Psi_\varepsilon(\cdot,\cdot)$ are nontrivial double-auction equilibria for $\mathcal{E}(\alpha, v, f, g, \Delta)$. (Nontriviality follows because $\bar{b}(\cdot)$ and $\bar{s}(\cdot)$ are sufficiently close.) In fact, the reasoning from part A shows that strict single-crossing must hold between any pair of prices having a market-clearing price weakly between them. Consequently, if p is a market-clearing price and if I is the nondegenerate interval on which $\bar{b}(\cdot)$ say, is p, then p is *strictly* optimal for buyers for all but perhaps one signal, x, in I. But because buyers and sellers have identical preferences in the continuum economy this implies that $\bar{s}(\cdot)$ must be p on I as well. It follows that, except for finitely many signals, $\bar{b}(\cdot)$ and $\bar{s}(\cdot)$ coincide on all market-clearing prices. Also, by the genericity result from part B, every unused price between the highest and lowest market-clearing price is strictly suboptimal. Hence, no such price can be in the range of either $\bar{b}(\cdot)$ or $\bar{s}(\cdot)$. We conclude that $\bar{b}(\cdot)$ and $\bar{s}(\cdot)$ are outcome-equivalent. But from this it follows easily that both $\bar{b}(\cdot)$ and $\bar{s}(\cdot)$ are outcome-equivalent double-auction equilibria for $\mathcal{E}(\alpha, v, f, g, \Delta)$.

The Import of Parts B and C We now discuss how parts B and C work together to help establish single-crossing on a range of prices in large double-auctions and for certain buyer-seller strategies. Call a strategy *monotone-optimal* if it is a best-reply subject to the constraint that it is monotone. Consider then a sequence of double-auctions where the number of agents converges to infinity, and consider also a corresponding sequence of monotone-optimal buyer-seller bidding functions that converge to an equilibrium of the continuum economy. By part C, the continuum-economy equilibrium is essentially symmetric and so the results of part A apply. In particular, given a sufficiently fine price grid, single-crossing holds at the continuum-economy equilibrium for prices between the highest and lowest market-clearing price. Now, ideally, we would like to invoke a continuity argument to conclude that single-crossing must also hold far enough along the sequence of monotone-optimal buyer-seller bidding functions in the large finite double-auctions. This would imply that the strategies are fully optimal far enough along the sequence because monotone best replies exist when single-crossing holds. But there is a potential problem with such a continuity argument.

For market-clearing prices that occur with probability bounded away from zero along the sequence, continuity arguments (provided in part D of the formal proof) do permit the application of the results from part A and single-crossing can be straightforwardly established for such prices. However, in principle, as the market grows, there may be market-clearing prices receiving positive, but vanishingly small, weight; the weight on these market-clearing prices may vanish in the limit if the length of the intervals of signals over which they are bid by any agent vanishes. Single-crossing may well fail to hold for such prices because, conditional upon their occurrence, the number of agents who submit them as bids can be small with high probability, even though there are many agents altogether. Consequently, an agent's incentives when contemplating any such bid are effectively the same as his incentives when there are a small number of agents, where we have already seen that single-crossing can fail (see Section 5.2). One way to solve this problem is to rule out the possibility that the weight on these prices vanishes in the limit. As we shall eventually see, the weight on some prices necessarily vanishes and for these prices we must face the problem head on; this is done in part D of the proof. On the other hand, for an important range of prices, condition (†) from part B ensures that their weight cannot vanish in the limit; their weight is eventually either bounded away from zero or equal to zero.

To see this last point, suppose that the limit equilibrium is $b_{\hat{\mathbf{x}}}(\cdot)$, and suppose that (†) holds. Then, the situation depicted in Figure 6.1 is in effect. In particular, recalling the conclusions reached in (6.6) and (6.7), every price between the highest and lowest market-clearing price that is not in the range of $b_{\hat{\mathbf{x}}}(\cdot)$ is *strictly suboptimal regardless of the agent's signal.* Consequently, by a continuity argument (also provided in part D of the formal proof), such prices must eventually be strictly suboptimal in all nearby strategies of large finite economies. But then such prices cannot be employed with positive probability in any nearby monotone-optimal strategy, because one can remove weight from the suboptimal price to an adjacent, utility-increasing, price without violating monotonicity. Hence, every price between the highest and lowest market-clearing price that is not in the range of $b_{\hat{\mathbf{x}}}(\cdot)$ is eventually employed with probability zero, while all prices in the range of $b_{\hat{\mathbf{x}}}(\cdot)$ are eventually employed with probability bounded away from zero (otherwise they would not be in the range of $b_{\hat{\mathbf{x}}}(\cdot)$).

So, parts B and C (together with some limit results from part D) establish that, if the number of agents increases in a double-auction, then far enough along a sequence of monotone-optimal buyer-seller bidding functions, single-crossing holds for all prices between the highest and lowest market-clearing price of the limit equilibrium of the continuum economy.

This establishes single-crossing over an important range of prices and so goes a long way toward establishing that the monotone-optimal strategies are in fact fully optimal and hence in equilibrium for the large double-auction. The remaining work centers around establishing single-crossing for prices that are below the lowest or above the highest market-clearing price of the limit equilibrium (like the prices 0, p_1, p_2, and p_{10}, indicated by the dotted lines in Figure 6.1). These prices are handled in part D.

Part D The final part of the proof considers sequences of large finite economies. Given unbounded sequences of natural numbers $\{n_r\}, \{m_r\}$ such that $n_r/(n_r + m_r) \to \alpha \in (0,1)$ as $r \to \infty$, consider the sequence of finite economies $\mathcal{E}(n_r, m_r, v, f, g, \Delta)$ with n_r buyers and m_r sellers, and the price grid $\mathcal{P} = \{0, \Delta, 2\Delta, ...\}$.

For $\mathbf{x}, \mathbf{y} \in X_K$, let $u_r^\beta(p, x|\mathbf{x}, \mathbf{y})$ denote the double-auction expected payoff of a buyer in $\mathcal{E}(n_r, m_r, v, f, g, \Delta)$ whose signal is x, when he bids the price p and all other $n_r - 1$

buyers employ the bidding function $b_{\mathbf{x}}(\cdot)$ and all m_r sellers employ the bidding function $b_{\mathbf{y}}(\cdot)$. Similarly, let $u_r^{\sigma}(p, x|\mathbf{x}, \mathbf{y})$ denote the double-auction expected payoff of a seller whose signal is x, when he bids the price p and all n_r buyers employ $b_{\mathbf{x}}(\cdot)$ and all other $m_r - 1$ sellers employ $b_{\mathbf{y}}(\cdot)$.

Hence, $(b_{\hat{\mathbf{x}}}(\cdot), b_{\hat{\mathbf{y}}}(\cdot))$ constitutes a double-auction equilibrium for $\mathcal{E}(n_r, m_r, v, f, g, \Delta)$ when, for every $x \in [0, 1]$,

$$b_{\hat{\mathbf{x}}}(x) \text{ solves } \max_{p \in \bar{\mathcal{P}}} u_r^{\beta}(p, x|\hat{\mathbf{x}}, \hat{\mathbf{y}}),$$

and

$$b_{\hat{\mathbf{y}}}(x) \text{ solves } \max_{p \in \bar{\mathcal{P}}} u_r^{\sigma}(p, x|\hat{\mathbf{x}}, \hat{\mathbf{y}}).$$

We shall place two restrictions upon buyer-seller jump-point vector pairs $(\mathbf{x}, \mathbf{y})$, and consider equilibria subject to these restrictions. Most of the work then entails showing that, generically, when the economy is sufficiently large and Δ is sufficiently small, these restrictions are not binding.

The first restriction on $(\mathbf{x}, \mathbf{y})$ is designed to ensure that the equilibrium is nontrivial. The vectors $\mathbf{x}$ and $\mathbf{y}$ will be required to be sufficiently close to one another, ensuring that trade occurs with probability bounded away from zero.

The second restriction on $(\mathbf{x}, \mathbf{y})$ is related to the use of jump-points as strategies, which implicitly restricts the bidding functions of buyers and sellers to be nondecreasing. To ensure that this strategy restriction is not binding, we must establish the single-crossing property. The second restriction on $(\mathbf{x}, \mathbf{y})$ is designed to establish the single-crossing property over the range of prices for which the single-crossing argument from parts B and C does not apply. Recall that the part B and C single-crossing argument does not apply to prices that may arise as market-clearing prices with positive but vanishing probability along a sequence of monotone-optimal buyer-seller strategies in double-auctions whose finite number of agents converges to infinity. Let us refer to these as "vanishing prices."

Consider now the conditions under which vanishing prices can occur. As already mentioned in part B above, a vanishing price can occur when the length of the intervals of signals over which buyers and sellers bid this price shrinks to zero. But a price can vanish even if the lengths of these intervals do not shrink to zero. For example, suppose that there are n buyers and $m = n$ sellers, each of whom employs the bidding function $b_{\hat{\mathbf{x}}}(\cdot)$ depicted in Figure 6.1.[43] Then, for every n, every price, $0, p_1, p_2, ..., p_{10}$, in the range of $b_{\hat{\mathbf{x}}}(\cdot)$ occurs with positive probability as a market-clearing price of the double-auction. However, in the limit as $n \to \infty$, the probabilities that the prices $0, p_1, p_2$ and

[43] In general, buyers and sellers will employ distinct bidding functions and this function will not be independent of the number of agents. The use here of a common bidding function that does not depend on the number of agents is for simplicity of exposition only.

p_{10} (the dotted bids in Figure 6.1) occur as market-clearing prices converge to zero, even though the intervals of signals over which they are bid remain bounded away from zero. To see why these prices vanish, note that with probability approaching one as $n \to \infty$, it is the median bid (because $n = m$) that determines the market-clearing price. Consequently, the market price is determined by the bid made at the median signal. However, in the limit as the number of agents grows, the median signal will lie between $x(0)$ and $x(1)$ with probability approaching one because $x(0)$ (resp., $x(1)$) is the median signal when the state ω takes on its lowest (resp., highest) possible value. Hence, in the limit, even though a positive and bounded away from zero fraction of agents submit the bids $0, p_1$, p_2 and p_{10}, the *median bid* will lie strictly between p_2 and p_{10} with probability approaching one. The market-clearing prices $0, p_1, p_2$ and p_{10} therefore vanish in the limit.

Establishing the single-crossing property for vanishing prices precludes the use of straightforward continuity arguments based upon the law of large numbers. Our technique for dealing with these prices is to control, to some extent, the rates at which the probabilities that they occur as market-clearing prices can vanish. We do this by ensuring that the step sizes of the agents' bid functions on these prices do *not* vanish. Hence, as in the example discussed in the previous paragraph, the probability that such prices clear the market can vanish only because, in the limit, the signals at which agents submit them are either less than $x(0)$ or above $x(1)$.

More formally, the second restriction is as follows. We choose a small $\varepsilon > 0$, and restrict $(\mathbf{x}, \mathbf{y})$ so that the length of each interval on which the step functions $b_{\mathbf{x}}(\cdot)$ and $b_{\mathbf{y}}(\cdot)$ assume each grid price in $(v(0,0) - \Delta, v(x(0), 0))$ and $(v(x(1), 1), v(1, 1) + \Delta)$ is at least ε.[44] We repeat that, *even with this restriction,* the probability that such prices arise as *market-clearing prices* can tend to zero as the economy grows since agents may bid such prices only when their signal is less than (resp., greater than) the lowest (resp., highest) possible αth percentile, $x(0)$ (resp., $x(1)$). Let us refer to this as the ε step-size restriction on vanishing prices.

In part D of the proof, C^0_ε denotes the subset of joint strategies $(\mathbf{x}, \mathbf{y}) \in X_K \times X_K$ such that $|x_k - y_k| \le \varepsilon^2$ for all k, and the ε step-size restriction on vanishing prices is satisfied. The choice of ε^2 in the one place and ε in the other is helpful in establishing certain bounds on probabilities involving order statistics. Some $\varepsilon > 0$ and sufficiently small is chosen and fixed.

Given buyer-seller strategies $(\mathbf{x}, \mathbf{y}) \in C^0_\varepsilon$, consider the maximization problem

$$\max_{(\mathbf{x}', \mathbf{y}') \in C^0_\varepsilon} \left[\int_0^1 u^\beta_r(b_{\mathbf{x}'}(x), x|\mathbf{x}, \mathbf{y}) f(x) dx + \int_0^1 u^\sigma_r(b_{\mathbf{y}'}(x), x|\mathbf{x}, \mathbf{y}) f(x) dx \right], \tag{6.8}$$

and let $\Psi^r_\varepsilon(\mathbf{x}, \mathbf{y})$ denote the set of solutions.

[44] Our proof ensires that vanishing prices cannot occur outside these sets.

Note that if $(b_{\mathbf{x}^r}(\cdot), b_{\mathbf{y}^r}(\cdot))$ is a double-auction equilibrium for $\mathcal{E}(n_r, m_r, v, f, g, \Delta)$ and $(\mathbf{x}^r, \mathbf{y}^r) \in C_\varepsilon^0$, then $(\mathbf{x}^r, \mathbf{y}^r) \in \Psi_\varepsilon^r(\mathbf{x}^r, \mathbf{y}^r)$. The opposite implication, however, need not be true because of the two restrictions embodied in the constraint set C_ε^0. But there is another reason as well. Recall the third restriction that is implicit in the formulation of the above maximization problem. The agents are implicitly restricted to choosing nondecreasing bidding functions. Hence, it must be shown that all three restrictions are not binding when r, which indexes the size of the economy, is sufficiently large.

The objective function in (6.8) is continuous in $\mathbf{x}, \mathbf{x}', \mathbf{y}, \mathbf{y}'$. Hence, $\Psi_\varepsilon^r(\cdot, \cdot)$ is nonempty valued and upper hemicontinuous, but it need not be convex valued. Of course, Kakutani's theorem guarantees the existence of $(\mathbf{x}^r, \mathbf{y}^r) \in co\Psi_\varepsilon^r(\mathbf{x}^r, \mathbf{y}^r)$. We may assume without loss that $(\mathbf{x}^r, \mathbf{y}^r) \to (\hat{\mathbf{x}}, \hat{\mathbf{y}})$ as $r \to \infty$. The remainder of the argument establishes that, under the conditions stated above, $(\mathbf{x}^r, \mathbf{y}^r)$ defines a double-auction equilibrium when r is large enough.

Part D includes several limit results showing that an agent's payoff in the finite economy converges to his payoff in the continuum economy. Consequently, upper hemicontinuity arguments can be employed to show that $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is a fixed point of the correspondence defined in part C. Hence, as argued there, $b_{\hat{\mathbf{x}}}(\cdot)$ and $b_{\hat{\mathbf{y}}}(\cdot)$ are outcome-equivalent and each one is a double-auction equilibrium for the continuum economy $\mathcal{E}(\alpha, v, f, g, \Delta)$.

As we already know, the results from part B can then be employed to argue that, generically, vanishing prices cannot occur for prices between the lowest and highest price in the range of $P(\cdot) \equiv b_{\hat{\mathbf{x}}}(x(\cdot)) = b_{\hat{\mathbf{y}}}(x(\cdot))$, and so strict single-crossing must hold for all prices in the range of $P(\cdot)$ not only at the limit, but also far enough along the sequence of finite economies. Hence, for prices in this range and for r sufficiently large, buyers and sellers each have *unique* (but possibly distinct) best replies. We will eventually show that the same holds true for the remaining prices as well.

We next wish to argue that $\hat{\mathbf{x}} = \hat{\mathbf{y}}$, i.e., that $b_{\hat{\mathbf{x}}}(\cdot) = b_{\hat{\mathbf{y}}}(\cdot)$. Suppose that $b_{\hat{\mathbf{x}}}(\cdot)$ is as shown in Figure 6.1. Now, the equality $b_{\hat{\mathbf{x}}}(x(\cdot)) = b_{\hat{\mathbf{y}}}(x(\cdot))$ established in the previous paragraph implies that $b_{\hat{\mathbf{x}}}(x) = b_{\hat{\mathbf{y}}}(x)$ for all $x \in [x(0), x(1)]$. It remains to show that $b_{\hat{\mathbf{x}}}(x) = b_{\hat{\mathbf{y}}}(x)$ for $x < x(0)$ and $x > x(1)$. We shall be content to show this for the lower interval $x < x(0)$.

Referring to Figure 6.1, the lower dotted portion of the bidding function $b_{\hat{\mathbf{x}}}(\cdot)$ is determined as follows. The proof in RP-s ensures that the prices $0, p_1, p_2$, and p_3 are consecutive prices in the grid. The jump points $\hat{x}_1$ and $\hat{x}_2$ are then uniquely determined by $v(\hat{x}_1, 0) = p_1$ and $v(\hat{x}_2, 0) = p_2$.

Because $b_{\hat{\mathbf{x}}}(x) = b_{\hat{\mathbf{y}}}(x)$ for all $x \in [x(0), x(1)]$ and $(\mathbf{x}^r, \mathbf{y}^r) \to (\hat{\mathbf{x}}, \hat{\mathbf{y}})$, we have that $b_{\hat{\mathbf{x}}^r}(x)$ and $b_{\hat{\mathbf{y}}^r}(x)$ each converge to $b_{\hat{\mathbf{x}}}(x) = b_{\hat{\mathbf{y}}}(x)$ for each $x \in (x(0), x(1))$ and so the prices $p_3, ..., p_9$ each occur as market-clearing prices with positive probability for large enough r. Furthermore, the step-size restriction ensures that each of the prices $0, p_1, p_2$, and p_{10} occur as a market-clearing price with positive probability for each r. Consequently, for large enough r, any one of the prices $0, p_1, ..., p_{10}$ is a market-clearing price with positive probability.

As $r \to \infty$, it becomes more and more likely that the αth percentile of the agents' signals is between $x(0)$ and $x(1)$ and it is straightforward to show, using the law of large numbers, that the relative likelihood that the market-clearing price is p_2 versus p_3 tends to zero. Slightly less obvious, but nonetheless true (and shown in Part D of RP-s), is that the relative likelihood that the market-clearing price is p_1 versus p_2 or 0 versus p_1 both tend to zero, a result that hinges on our step-size restriction ensuring that the lengths of the intervals on which p_1,and p_2 are bid are bounded away from zero. Indeed, establishing these relative likelihood limits is the role of the step-size restriction. We now flesh out some of the implications.

We now claim that for sufficiently large r, buyers and sellers with signals below $\hat{x}_3$ optimize uniquely by employing bidding strategies near the dotted strategy in Figure 6.1. For example, consider a buyer with signal $x' \in (\hat{x}_1, \hat{x}_2)$. By the definition of $\hat{x}_1$ and $\hat{x}_2$, $p_1 < v(x', 0) < p_2$. For large r, strict single-crossing holds for prices between p_3 and p_{10}. Hence, the buyer's optimal bid is less than p_3 because $x' < \hat{x}_3$. To see that p_2 is not optimal, consider the difference in his payoff from bidding p_2 versus p_1. This difference depends only upon the events in which the market price is either p_2 or p_1. But, as we have already argued, the market price is infinitely more likely to be p_2 than p_1. Furthermore, conditional on a market-price of p_2, the state of the good is almost surely $\omega = 0$, since any higher state makes the already unlikely event that the market price is p_2 infinitely less likely. Hence, conditional on a market-price of p_1 or p_2, a buyer with signal x' is almost certain that his value is $v(x', 0) < p_2$ and that the market-clearing price is p_2. Hence, bidding p_2 is strictly suboptimal. Similarly, bidding $p_0 = 0$ is strictly suboptimal since he then runs the risk of being rationed when the market price is zero, while bidding p_1 guarantees that he wins when the price is zero and gives him a chance of winning when the price is $p_1 < v(x', 0)$. Consequently, for a given $x' \in (\hat{x}_1, \hat{x}_2)$ and for all large enough r, the buyer's optimal bid is p_1, as indicated by the dotted bidding strategy in the figure.

But exactly the same argument applies to sellers and so in the limit both the buyers' and the sellers' strategies converge to the dotted bidding function for $x < x(0)$. A similar argument for $x > x(1)$ therefore establishes that $b_{\hat{\mathbf{x}}}(x) = b_{\hat{\mathbf{y}}}(x)$ for all x and hence that $\hat{\mathbf{x}} = \hat{\mathbf{y}}$.

A further implication of the above argument is that agents with signals below $\hat{x}_3$ behave almost as if they are certain that the state is $\omega = 0$ and so as if they have private values. But then (as proven in RP-s) strict single-crossing holds and the agents have unique best replies close to the dotted strategy in the figure. We conclude that best replies are unique at the prices $0, p_1$, and p_2; and similarly for p_{10}. Since we already know that best replies are unique at the other prices, buyers and sellers each have unique best replies against $(\mathbf{x}^r, \mathbf{y}^r)$ for large r.

Note that because for large r a strategy close to the dotted one is optimal, the step-size constraint on an agent's strategy is not binding. Further, because $\mathbf{x}^r$ and $\mathbf{y}^r$ converge to the *same* vector $\hat{\mathbf{x}} = \hat{\mathbf{y}}$, the constraint that they must be close to one another is eventually not binding. Finally, because single-crossing eventually holds, the restriction to nondecreasing bidding functions is also eventually not binding. Hence, for large r, every member of $\Psi^r_\varepsilon(\mathbf{x}^r, \mathbf{y}^r)$ is an *unconstrained* best reply. But because there is a unique best reply, $\Psi^r_\varepsilon(\mathbf{x}^r, \mathbf{y}^r)$ must be a singleton. Hence, because $(\mathbf{x}^r, \mathbf{y}^r) \in co\Psi^r_\varepsilon(\mathbf{x}^r, \mathbf{y}^r)$ and $\Psi^r_\varepsilon(\mathbf{x}^r, \mathbf{y}^r)$ is a singleton, it must in fact be the case that $(\mathbf{x}^r, \mathbf{y}^r) \in \Psi^r_\varepsilon(\mathbf{x}^r, \mathbf{y}^r)$, and so $(\mathbf{x}^r, \mathbf{y}^r)$ is an unconstrained best reply to itself; it is an equilibrium.

Hence, for r large enough, $(b_{\mathbf{x}^r}(\cdot), b_{\mathbf{y}^r}(\cdot))$ is a double-auction equilibrium of $\mathcal{E}(n_r, m_r, v, f, g, \Delta)$. It is nontrivial because $\mathbf{x}^r$ and $\mathbf{y}^r$ were restricted so that this was so. This completes the proof of (1) of Theorem 6.1.

Because $b_{\mathbf{x}^r}(\cdot)$ and $b_{\mathbf{y}^r}(\cdot)$ each converge uniformly to $b_{\hat{\mathbf{x}}}(\cdot) = b_{\hat{\mathbf{y}}}(\cdot)$, a double-auction equilibrium of $\mathcal{E}(\alpha, v, f, g, \Delta)$, an appeal to Proposition 1.7 proves (2) of Theorem 6.1, and several applications of the law of large numbers proves (3) and (4). RP-s contains the details.

7. References

Allen, B. (1981): "Generic Existence of Completely Revealing Equilibria for Economies with Uncertainty when Prices Convey Information," *Econometrica*, 49, 1173-1199.

Athey, S. (2001): "Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games with Incomplete Information," *Econometrica*, 69, 861-890.

Chatterjee, K., and W. Samuelson (1983): "Bargaining Under Incomplete Information," *Operations research,* 31, 835-851.

Christensen, J.P.R. (1974): *Topology and Borel Structures*, North-Holland Mathematics Studies 10, Elsevier, New York.

Cripps, M. and J. M. Swinkels (2003): "Depth and Efficiency in Large Double Auctions with Private Values," mimeo, Olin School of Business, Washington University.

Fudenberg, D., M. M. Mobius, and A. Szeidl (2003): "Existence of Equilibrium in Large Double Auctions," mimeo, Harvard University.

Hellwig, M. R. (1980): "On the Aggregation of Information in Competitive Markets," *Journal of Economic Theory,* 22, 477-498.

Hunt, B. R., T. Sauer, and J. A. Yorke (1996): "Prevalence: A Translation Invariant 'Almost Every' on Infinite-Dimensional Spaces," *Bulletin (New Series) of the American Mathematical Society*, 27, 217-238.

Jackson, M. O., and J. M. Swinkels (2002): "Existence of Equilibrium in Single and Double Private Value Auctions," mimeo, Washington University.

Kyle, A. S. (1989): "Informed Speculation with Imperfect Competition," *The Review of Economic Studies,* 56, 317-355.

Messner, S., and X. Vives (2001): "Allocative and Productive Efficiency in REE with Asymmetric Information," Center for Economic Policy Research," Discussion paper No. 2678.

Milgrom P. R. (1981): "Rational Expectations, Information Acquisition, and Competitive Bidding," *Econometrica,* 49, 921-943.

Milgrom P. R. and R. J. Weber (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica,* 50, 1089-1122.

Milgrom P. R. and N. Stokey (1982): "Information, Trade and Common Knowledge", *Journal of Economic Theory,* 26, 17-27.

Oxtoby, J. C. (1980): *Measure and Category*, Second Edition, Springer-Verlag, New York.

Palfrey, T. R. (1985): "Uncertainty Resolution, Private Information Aggregation and the Cournot Competitive Limit," *Review of Economic Studies,* 52, 69-83.

Pesendorfer W. and J. M. Swinkels (1997): "The Loser's Curse and Information Aggregation in Common Value Auctions," *Econometrica,* 65, 1247-1281.

Reny, P. and M. Perry (2003): "Toward a Strategic Foundation for Rational Expectations Equilibrium," mimeo, University of Chicago.

Reny, P. and M. Perry (2005): "Supplement to 'Toward a Strategic Foundation for Rational Expectations Equilibrium'," mimeo, University of Chicago.

Rustichini A., M. A. Satterthwaite and S. R. Williams (1994): "Convergence to Efficiency in a Simple Market with Incomplete Information," *Econometrica,* 62, 1041-1063.

Satterthwaite, M.A., and S.R. Williams (1989): "The Rate of Convergence to Efficiency in the Buyer's Bid Double Auction as the Market Becomes Large," *The Review of Economic Studies,* 56, 477-498.

Swinkels, J.M. (2001): "Efficiency of Large Private Value Auctions," *Econometrica,* 69, 37-68.

Vives, X. (1988): “Aggregation of Information in Large Cournot Markets,” *Econometrica*, 56, 851-876.

Wilson R. (1977): “A Bidding Model of Perfect Competition,” *The Review of Economic Studies,* Vol. 44, 511-518.

Wilson R. (1985): “Incentive Efficiency of Double Auctions,” *Econometrica,* 53, 1101-1116.