Learning and Price Discovery in a Search Model*

Stephan Lauermann † Wolfram Merzyn † Gábor Virág § ${\rm August}\ 20,\ 2011$

Abstract

We develop a dynamic matching and bargaining game with aggregate uncertainty about the relative scarcity of a commodity. We use our model to study price discovery in a decentralized exchange economy: Traders gradually learn about the state of the market through a sequence of multilateral bargaining rounds. We characterize the resulting equilibrium trading patterns. We show that equilibrium outcomes are approximately competitive when frictions are small. Therefore, prices aggregate information about the scarcity of the traded commodity; that is, prices correctly reflect the commodity's economic value.

JEL CLASSIFICATIONS: D44, D82, D83

KEYWORDS: Dynamic Matching and Bargaining, Decentralized Markets, Foundations of Competitive Equilibrium, Search Theory, Learning

^{*}We have benefited from helpful comments by Georg Nöldeke, Tilman Börgers, Jan Eeckhout, Chris House, Philipp Kircher, Benny Moldovanu, Tymofiy Mylovanov, and Art Shneyerov, as well as seminar audiences at NYU Stern, Concordia University, the 2009 Canadian Economic Theory Conference, the 2009 Decentralization Conference at Washington University in St. Louis and the 2009 Midwest Theory conference at Penn State. We thank Isaac Sorkin for research assistance.

[†]University of Michigan, slauerma@umich.edu.

[‡]Deutsche Bahn AG, wolfram.merzyn@gmx.de.

[§]University of Toronto, gabor.virag@utoronto.ca.

1 Introduction

General equilibrium theory famously states that—under certain assumptions—there exists a vector of prices such that all markets clear. It fails, however, to explain how the market-clearing ("Walrasian") price vector comes about. The literature on dynamic matching and bargaining games, pioneered by Rubinstein and Wolinsky (1985) and by Gale (1987), aims to fill this gap in the foundations for general equilibrium theory. It addresses the question of how prices are formed in decentralized markets and whether these prices are Walrasian. Existing models of dynamic matching and bargaining, however, assume that market demand and supply—and, hence, the market-clearing price—are common knowledge among traders. This assumption is restrictive because markets have been advocated over central planning precisely on the grounds of their supposed ability to "discover" the equilibrium prices by eliciting and aggregating information that is dispersed in the economy; see Hayek (1945). By construction, existing models that take market-clearing prices to be common knowledge remain silent about whether this argument is correct and whether markets can indeed solve the price discovery problem.

We develop a dynamic matching and bargaining game to study price discovery in a decentralized market. We relax the standard assumption that the aggregate state of the market is common knowledge. In our model, individual traders are uncertain about market demand and supply. No individual trader knows the relative scarcity of the good being traded. We analyze the resulting patterns of trade and learning that emerge in equilibrium. We ask whether traders eventually learn the relevant aggregate characteristics and whether prices accurately reflect relative scarcity when frictions are small.

Our model combines elements from Satterthwaite and Shneyerov (2008) and Wolinsky (1990). Specifically, the matching technology and the bargaining protocol are adapted from Satterthwaite and Shneyerov. Time is discrete. In every period, a continuum of buyers and sellers arrives at the market. All buyers are randomly matched to the sellers, resulting in a random number of buyers that are matched with each seller. Each seller conducts a second-price sealed-bid auction with no reserve price. At the end of each round, successful buyers and sellers leave the economy. Unsuccessful traders either leave the market without trading with some exogenous exit probability or remain in the market. The exogenous exit rate makes waiting costly and is interpreted as the "friction" of trade.¹

The defining feature of our model is aggregate uncertainty, which we model similar to Wolinsky (1990). There is a binary state of nature. The realized state is unknown

¹The presence of an exit rate is the main difference between Satterthwaite and Shneyerov (2008) and Satterthwaite and Shneyerov (2007). The exit rate ensures that strategy profiles determine steady-state outcomes uniquely; see Nöldeke and Tröger (2009).

to the traders and does not change over time. For each state of nature, we consider the corresponding steady state of the market. In Wolinsky (1990), the state determines the common value of the good being traded. We consider a private value setting instead. In our model, the state of nature determines the relative scarcity of the good. Depending on the state of nature, the mass of incoming buyers is either large or small, whereas the mass of incoming sellers is independent of the state of nature. The larger the mass of entering buyers is relative to the mass of entering sellers, the scarcer is the good. Every buyer receives a noisy signal upon birth. Moreover, after every auction, the losing buyers obtain additional information regarding the state. In our model, losing bidders do not observe other buyers' bids. Nevertheless, they are able to draw an inference from the fact that their respective bids lost.²

We characterize equilibrium learning and bargaining strategies. The buyers shade their bids to account for the opportunity cost of foregone continuation payoffs. Moreover, despite the fact that the consumption value of the good is known, the fact that continuation payoffs depend on the unknown common state of nature makes the buyers' preferences interdependent. A resulting winner's curse leads to further bid shading: Winning an auction implies that on average fewer bidders are participating and that the participating bidders are more optimistic about their continuation payoff. Both of these facts imply a lower value of winning the good than expected prior to winning. Countervailing the winner's curse is a "loser's curse." The role of the loser's curse for information aggregation in large double auctions was identified by Pesendorfer and Swinkels (1997). In our model, losing an auction implies that on average more bidders are participating and that the participating bidders are more pessimistic about their continuation payoff. The loser's curse implies that bidders become more pessimistic and raise their bids after repeated losses over time.³

We are particularly interested in the characterization of the equilibrium when the exogenous exit rate is small, which is interpreted as the frictionless limit of the decentralized market. Our main result shows that the limit outcome approximates the Walrasian outcome relative to the realized aggregate state of the market. If the realized state is such that the mass of incoming buyers exceeds the mass of incoming sellers, the resulting limit price at which trade takes place is equal to the buyers' willingness to pay; if the realized state is such that the mass of incoming buyers is smaller than the mass of incoming sellers, the price is equal to the seller's costs. Therefore, prices

²Our model differs from the two cited papers. It differs from Satterthwaite and Shneyerov (2008), because in our model buyers and sellers have homogeneous cost and valuations—but heterogeneous beliefs—and because sellers cannot post secret reservation prices. Unlike Wolinsky (1990), in our model valuations and costs are private and bargaining is according to an auction with a continuum of possible bids, rather than bilateral price posting.

³Thus, the loser's curse refers to the effect of the learning dynamics *over time*, whereas the winner's curse refers to bid shading *in each period*.

aggregate information about the scarcity of the traded commodity; that is, prices reflect the commodity's economic value.

Our analysis reveals how the winner's curse and the loser's curse shape equilibrium outcomes and information aggregation. We show that as the exit rate vanishes, entrants' initial bids are dominated by the winner's curse and the buyers bid for an increasing number of periods as if they are certain that the continuation value is maximal. Eventually, however, the loser's curse becomes strong enough so that those buyers who have lost in a sufficiently large number of periods raise their bids over time. Specifically, suppose that the market-clearing price can be either high or low depending on realized market demand. Then the number of periods in which buyers bid the low price diverges to infinity. However, the number of periods in which buyers bid low grows more slowly than the rate at which the exit rate goes to zero. Therefore, bids become high sufficiently fast relative to the exit rate. This ascending bid pattern ensures that actual transaction prices are equal to the sellers' costs if the realized state is such that the buyers are on the short side of the market and actual transaction prices are equal to the buyers' valuations if the buyers are on the long side of the market.

In the following section, we discuss our contribution to the literature. In Section 3 we introduce the model. We provide existence and uniqueness results for steady-state equilibria in monotone bidding strategies in Section 4. We also provide some preliminary characterization of equilibrium. Proving the existence of equilibrium is a non-trivial problem in a search model with aggregate uncertainty because of the endogeneity of the distribution of population characteristics (beliefs in our model); see Smith (2011). Some techniques that we develop might be useful more generally. For example, we demonstrate an interesting failure of the monotone likelihood ratio property in auctions with a random number of bidders.⁴ Our main result on convergence to the competitive outcome is stated in Section 5. We provide an extension to an economy with heterogeneous buyers in Section 6. Section 7 provides a discussion of extensions and conclusion. Some technical results about the steady-state stock and the proof of existence of equilibrium are relegated to a supplementary online appendix.⁵

2 Contribution to the Literature

We contribute to a body of research that studies the foundations for general equilibrium through the analysis of dynamic matching and bargaining games, which was initiated by Rubinstein and Wolinsky (1985) and Gale (1987).⁶ A central question is whether a

⁴This is discussed in Section 1.3 of the online appendix.

⁵The online appendix is included in this submission but it is not intended for publication.

⁶For recent contributions, see for example Satterthwaite and Shneyerov (2008), Shneyerov and Wong (2010), and Kunimoto and Serrano (2004) and the references therein.

fully specified "decentralized" trading institution leads to outcomes that are competitive when frictions of trade are small. Well-known negative results by Diamond (1971) and Rubinstein and Wolinsky (1985) have demonstrated that this question is not trivial. Studying foundations is important for *positive* theory in order to understand under which conditions markets can and cannot be well approximated by competitive analysis, and for *normative* theory in order to understand what trading institutions are able to decentralize desirable allocations.⁷

In the early matching and bargaining literature, the preferences and endowments of each individual trader were assumed to be observable. Thus, the problems of price discovery and learning were absent. Satterthwaite and Shneyerov (2007, 2008) introduce a model with private information, proving convergence to the competitive outcome. However, in their model the preferences are independently distributed among a continuum of traders, and so the realized distribution of preferences—and, hence, market demand and supply—is commonly known. Thus, in their model, there is idiosyncratic uncertainty but there is no aggregate uncertainty.

Our paper is related to work on matching and bargaining with common values by Wolinsky (1990) and later work by Blouin and Serrano (2001).⁸ This work on search with common values provides negative convergence results. It uncovers a fundamental problem of information aggregation through search: As frictions vanish, traders can search and experiment at lower costs. This might seem to make information aggregation simple. However, it also implies that traders increasingly insist on favorable terms—the buyers on low prices and the sellers on high prices—turning the search market into "a vast war of attrition" (Blouin and Serrano (2001, p. 324)). This insistence on extreme positions makes information aggregation difficult even when search frictions are small. In our model, the winner's curse implies a similar effect: When the exit rate vanishes, the buyers bid low and insist on a price equal to the sellers' cost for an increasingly large number of periods. Yet, in our setting, this "insistence problem" is overcome by the opposing loser's curse, as discussed before.

There are two important differences from our paper, which might also explain the divergent results. First, in models with common values, preferences depend on an unknown state, and, consequently, these models are used to study foundations for Rational Expectations Equilibria. By contrast, we study the microfoundations for competitive equilibrium in a standard exchange economy (albeit a stylized, quasilinear one). Second, in Wolinsky (1990) and Blouin and Serrano (2001) traders can choose only between

⁷For an excellent discussion of strategic foundations for general equilibrium, see the introductory chapter in Gale (2000).

 $^{^8}$ Serrano and Yosha (1993) consider a related problem with one-sided private information and Gottardi and Serrano (2005) consider a "hybrid" model of decentralized and centralized trading.

⁹In our model, a common value component arises as well. However, it arises *endogenously* through the equilibrium continuation payoffs.

two price offers (bargaining postures). It is an open question whether trading outcomes in these models are competitive if the frictions are small and if the restriction on the set of prices is not imposed.¹⁰

Golosov, Lorenzoni and Tsyvinski (2011) consider a related model of search with common values without this restriction on the set of price offers. They consider an economy in which two divisible assets of unknown common value are traded. The friction in their model is an exogenous probability that trading stops in any given period. They show that equilibrium outcomes approximate ex-post efficient outcomes in the event that the game has not stopped for a sufficiently large number of periods. However, Golosov et al. study the trading outcome with a fixed, positive stopping probability. They do not study the question whether outcomes become competitive in the "frictionless" limit when the stopping probability is small, which is the question that motivates our paper.¹¹

Shneyerov, Majumdar and Xie (2011) is the only other paper that considers a dynamic matching and bargaining game with uncertainty about market demand and supply. Like us, Shneyerov et al. study the frictionless limit and show that, in the limit, outcomes are approximately competitive. The main difference between the papers is the dimension along which heterogeneity is introduced and how learning is modeled. We consider buyers and sellers who have homogeneous preferences but who can hold a large set of heterogeneous beliefs. Shneyerov et al. allow for heterogeneous preferences, but they restrict the set of possible beliefs to place either probability one or probability zero on each of two possible states. Specifically, they consider a model of "optimism": Buyers and sellers start with diametrically opposed priors, each believing that the state is in their favor with probability one. Learning takes place by traders switching away from probability one beliefs after observing events that have subjective probability zero.

There is a large body of related work on information aggregation in centralized institutions in which all traders simultaneously interact directly (see, e.g., the work on large double auctions by Reny and Perry (2006) and Pesendorfer and Swinkels (1997, 2000)) and on the behavior of traders in financial markets (e.g., Kyle (1989), Ostrovsky (2011) and Rostek and Weretka (2011)). The assumption of a central price formation mechanism distinguishes this literature from dynamic matching and bargaining games in which prices are determined in a decentralized manner through bargaining.

Finally, our work is related to the literature on social learning (Banerjee and Fudenberg (2004)), the recent work on information percolation in networks (Golub and Jackson (2010)), and information percolation with random matching (Duffie and Manso (2007)). In the latter model, agents who are matched observe each other's information.

¹⁰See the concluding discussion of the trading procedure in Blouin and Serrano (2001).

¹¹As explained in Golosov et al., ex-post efficiency of the outcome in the event that the game does not stop does not imply that this is the rational expectations equilibrium relative to the initial allocation.

In our model, the amount of information that one bidder learns from other traders is endogenous and depends on the action (bid) that they choose.

3 Model and Equilibrium

3.1 Setup

There are a continuum of buyers and a continuum of sellers present in the market. In periods $t \in \{..., -1, 0, 1, ...\}$, these traders exchange an indivisible, homogeneous good. Each buyer demands one unit, and the buyers have a common valuation v for the good. Each seller has one unit to trade. The common cost of selling is c = 0. Trading at price p yields payoffs v - p and p - c, respectively. The valuation exceeds the cost, so there are gains from trade. Buyers and sellers maximize expected payoffs.

Similar to Wolinsky (1990), there are two states of nature, a high state and a low state $w \in \{H, L\}$. Both states are equally likely. The realized state of nature is fixed throughout and unknown to the traders. For each realization of the state of nature, we consider the corresponding steady-state outcome, indexed by w. The state of nature determines the constant and exogenous number of new traders who enter the market (the flow), and, indirectly, it also determines the constant and endogenous number of traders in the market (the stock). In the low state, the mass of buyers entering each period is d^L , and, in the high state, it is d^H . More buyers enter in the high state, $d^H > d^L$. The mass of sellers who enter each period is the same in both states and is equal to one.

The buyers are characterized by their beliefs $\theta \in [0,1]$, the probability that they assign to the high state. In the following, we often refer to θ as the *type* of a buyer. Each buyer who enters the market privately observes a noisy signal and forms a posterior based on Bayesian updating. In state w, the posteriors of the entering buyers are assumed to be distributed on the support $[\underline{\theta}, \overline{\theta}]$, with cumulative distribution functions G^H and G^L , respectively. The distributions are continuous and admit continuous probability density functions, g^H and g^L . Notice that using Bayes' rule the distributions must be such that $\theta = \frac{d^H g^H(\theta)}{d^H g^H(\theta) + d^L g^L(\theta)}$, or, equivalently, the likelihood ratio satisfies

$$\frac{\theta}{1-\theta} = \frac{d^{H}}{d^{L}} \frac{g^{H}\left(\theta\right)}{g^{L}\left(\theta\right)}.$$

For a buyer, the mere fact of entering the market contains news because the inflow is larger in the high state. Conditional on entering the market, a buyer is pessimistic and believes that the high state is more likely than the low state. This is expressed by the likelihood ratio $d^H/d^L > 1$.¹²

¹²To formally define updating based on entering the market, suppose that there is a *potential* set of buyers of mass d, $d \ge d^H$. In state w, a mass d^w of the potential buyers actually enters the market.

We assume that the lower and upper bounds of the support $[\underline{\theta}, \overline{\theta}]$ are such that $1/2 \leq \underline{\theta} < \overline{\theta} < 1$. The assumption that $1/2 \leq \underline{\theta}$ is needed to ensure monotonicity of a certain posterior; see the remarks following Lemma 3. Substantively, this assumption is consistent with signals being sufficiently noisy, so that even the most favorable signal $\underline{\theta}$ is not strong enough to overturn the initial pessimism of an entering buyer. The assumption $\overline{\theta} < 1$ implies that signals are boundedly informative.¹³ Each period consists of several steps:

- 1. Entry occurs (the "inflow"): A mass one of sellers and a mass d^w of buyers enter the market. The buyers privately observe signals, as described before.
- 2. Each buyer in the market (the "stock") is randomly matched with one seller. A seller is matched with a random number of buyers. The probability that a seller is matched with k = 0, 1, 2, ... buyers is Poisson distributed¹⁴ and equal to $e^{-\mu}\mu^k/k!$, where μ is the endogenous ratio of the mass of buyers to the mass of sellers in the stock. We sometimes refer to μ as a measure of market "tightness." The expected number of buyers who are matched with each seller is equal to μ , of course.
- 3. Each seller runs a sealed-bid second-price auction with no reserve price. The buyers do not observe how many other buyers are matched with the same seller. The bids are not revealed ex post, so bidders learn only whether they have won with their submitted bid.
- 4. A seller leaves the market if its good is sold; otherwise, the seller stays in the stock with probability $\delta \in [0,1)$ to offer its good in the next period. A winning buyer pays the second highest bid, obtains the good, and leaves the market. A losing buyer stays in the stock with probability δ and is matched with another seller in the next period. Those who do not stay exit the market permanently. A trader who exits the market without trading has a payoff of zero.
- 5. Upon losing, the remaining buyers update their beliefs based on the information gained from losing with their submitted bids. The remaining buyers and sellers who neither traded nor exited stay in the market. Together with the inflow, these traders make up the stock for the next period.

On the individual level, the exit rate $1 - \delta$ acts similar to a discount rate: Not trading today creates a risk of losing trading opportunities with probability $1 - \delta$. On the

Alternatively, one can simply interpret d^H/d^L as the prior of an entering buyer.

¹³We endow buyers with initial signals to ensure existence of a pure strategy equilibrium. Our results hold even if the interval $[\underline{\theta}, \overline{\theta}]$ is arbitrarily small and even if the initial signal is fully uninformative.

¹⁴This distribution is consistent with the idea that there are a large number of buyers who are independently matched with sellers. The resulting distribution of the number of buyers matched with a seller is binomial. When the number of buyers and sellers is large, the binomial distribution is approximated by the Poisson distribution.

aggregate level, the exit rate ensures that a steady state exists for all strategy profiles; see Nöldeke and Tröger (2009). Traders do not discount future payoffs beyond the implicit discounting implied by the exit rate.

3.2 Steady-State Equilibrium

We study steady-state equilibria in stationary strategies so that the distribution of bids depends only on the state and not on (calendar) time. An immediate consequence is that in any period the set of optimal bids of a buyer depends only on the current belief about the likelihood of being in the high state.

We restrict attention to symmetric and pure strategy equilibria where the bid is a strictly increasing function of the belief of the buyer and where the distribution of beliefs is sufficiently "smooth," as defined below. A symmetric steady-state equilibrium is a vector $(\Gamma^H, \Gamma^L, S^H, D^H, S^L, D^L, \beta, \theta^+)$. Next, we describe each of these components. First, the distributions of beliefs are given by atomless cumulative distribution functions Γ^w . Each function Γ^w is absolutely continuous and nondecreasing. Furthermore, we assume that Γ^w is piecewise twice continuously differentiable.¹⁵ These assumptions ensure that we can choose a density, denoted γ^w , that is right continuous on [0, 1).¹⁶

The masses of buyers and sellers in the stock are D^w and S^w . The bidding strategy β is a strictly increasing function and maps beliefs from [0,1] to nonnegative bids. We often use the generalized inverse of β , given by $\beta^{-1}(b) = \inf \{\theta | \beta(\theta) \ge b\}$, where $\beta^{-1}(b) = 1$ if $\beta(\theta) < b$ for all θ . Finally, $\theta^+(x,\theta)$ is the posterior of a buyer with initial belief θ conditional on losing against buyers with beliefs above x.

We characterize the equilibrium requirements for these objects; a formal definition of equilibrium follows at the end of this section. Let $\theta_{(1)}$ denote the first-order statistic of beliefs in any given match. We set $\theta_{(1)} = 0$ if there is no bidder present. $\Gamma_{(1)}^w$ denotes the c.d.f. of the first-order statistic in state w; that is, $\Gamma_{(1)}^w(x)$ is the probability that the highest belief in the auction is below x. The event in which all the buyers have a belief below x includes the event in which there are no buyers present at all. The probability of this event is $\Gamma_{(1)}^w(0)$ by our assumption that there is no atom in the distribution of beliefs at zero. The Poisson distribution implies $\Gamma_{(1)}^w(0) = e^{-\mu^w}$, where $\mu^w = D^w/S^w$ as defined before. The fact that this probability is positive implies that the buyers must have positive expected payoffs because any buyer has some probability of being the sole

¹⁵A function is piecewise twice continuously differentiable on [0, 1] if there is a partition of [0, 1] into a countable collection of open intervals and points such that the function is twice continuously differentiable on each open interval. Moreover, we require that the set of non-differentiable points has no accumulation point except at one. Smoothness ensures that we can work conveniently with densities.

¹⁶We believe that these restrictions are without loss of generality. The restriction to symmetric and pure strategies is without loss of generality by the uniqueness of the optimal bids, whenever belief distributions are atomless and bidding strategies are strictly increasing. However, we have not been able to show that all equilibria have these latter properties.

bidder and receiving the good at a price of zero. In general, the first-order statistic of the distribution of beliefs is given by

$$\Gamma_{(1)}^{w}(\theta) = e^{-\mu^{w}(1-\Gamma^{w}(\theta))}.$$
(1)

Intuitively, $\mu^w(1-\Gamma^w(\theta))$ is the ratio of the mass of buyers having belief above θ to the mass of sellers, and $e^{-\mu^w(1-\Gamma^w(\theta))}$ is the probability that the seller is matched with no buyer having such belief. Let $\Gamma^{\theta}_{(1)}(x) = \theta \Gamma^{H}_{(1)}(x) + (1-\theta)\Gamma^{L}_{(1)}(x)$ be the unconditional probability that the highest belief is below x if the probability of the high state is θ .

We derive the posterior upon losing. Given the assumption that bidding strategies are strictly increasing, losing with a bid b implies that there was some bidder in the match with a belief above $x = \beta^{-1}(b)$. Bayes' rule for the posterior θ^+ requires that

$$\theta^{+}(x,\theta) = \frac{\theta(1 - \Gamma_{(1)}^{H}(x))}{1 - \Gamma_{(1)}^{\theta}(x)}$$
(2)

if $1 - \Gamma_{(1)}^{\theta}(x) > 0$. Otherwise, we set $\theta^{+}(x,\theta) \equiv \sup\{(\theta^{+}(x',\theta)) | x' : 1 - \Gamma_{(1)}^{\theta}(x') > 0\}$, which is well defined by monotonicity of $\Gamma_{(1)}^{w}$. This particular choice of the "off-equilibrium" belief does not affect our analysis.¹⁷

To derive the steady-state conditions for the stock, suppose that the mass of sellers is S^w today. A seller trades if and only if matched with at least one buyer. Tomorrow's population of sellers therefore consists of the union of those sellers who were not matched with any buyer and the newly entering sellers. In steady state, these two populations must be identical, requiring

$$S^{w} = 1 + \delta \Gamma_{(1)}^{w} (0) S^{w}. \tag{3}$$

The inflow of buyers having type less than θ is $d^w G^w(\theta)$. The stationarity condition is

$$D^{w}\Gamma^{w}\left(\theta\right) = d^{w}G^{w}(\theta) + \delta D^{w} \int_{\left\{\tau:\theta^{+}\left(\tau,\tau\right)\leq\theta\right\}} \left(1 - \Gamma_{(1)}^{w}\left(\tau\right)\right) d\Gamma^{w}\left(\tau\right). \tag{4}$$

The steady-state mass of buyers in the stock having a type below θ is equal to $D^w\Gamma^w(\theta)$. This mass has to be equal to the mass of the buyers in the inflow with type less than θ (the first term on the right-hand side) plus the mass of buyers who lose, survive, and update to some type less than θ (the second term).¹⁸

¹⁷When restricting beliefs to the two states of nature, we implicitly assume that following an off-equilibrium event—the only such event is losing with a bid above the highest equilibrium bid—a buyer continues to believe that all other buyers play according to their equilibrium strategies.

¹⁸For the purpose of this paper, the steady-state model is *defined* by (3) and (4). Formally, these equations are taken as the primitives of our analysis and they are not derived from some stochastic matching process. This allows us to avoid well-known measure theoretic problems with a continuum of random variables. These problems can be solved, however, at the cost of additional complexity; see

Let $V(\theta)$ denote the value function, which is equal to

$$\max_{b} v \Gamma_{(1)}^{\theta}(0) + \int_{0+}^{\beta^{-1}(b)} (v - \beta(\tau)) d\Gamma_{(1)}^{\theta}(\tau) + \delta(1 - \Gamma_{(1)}^{\theta}(\beta^{-1}(b))) V(\theta^{+}), \quad (5)$$

where $\theta^{+} = \theta^{+}(\theta, \beta^{-1}(b))$. A bidding strategy β is optimal if $b = \beta(\theta)$ solves the maximization problem (5) for every θ .

A steady-state equilibrium in symmetric, strictly increasing bidding strategies with an atomless distribution of types (an equilibrium from now on) consists of (i) masses of buyers and sellers, S^H, D^H, S^L, D^L , and distribution functions Γ^H , Γ^L such that the steady-state conditions (3) and (4) hold for all θ ; (ii) an updating function θ^+ that is consistent with Bayes' rule (2); (iii) a strictly increasing bidding function β that is optimal (maximizes (5)).

4 Characterization and Existence of Equilibrium

4.1 The Equilibrium Stock

The following lemmas establish necessary implications of equilibrium for the steady-state stock. Generally, characterizing stocks in equilibrium search models is difficult because of an intricate feedback between stocks and strategies, which requires determining these two objects simultaneously. In our model, however, we can "decouple" the stock from the strategies. This is because the bidding strategy is strictly increasing: The identity of the winning bidder as well as the updated belief is the same for all strictly increasing bidding strategies. We now describe properties of the stock, assuming (and verifying later) that a monotone equilibrium exists. All proofs of the results from this section are in the supplementary online appendix, with the exception of the proof of the following lemma.

Lemma 1 (Unique Masses.) For each state w, there are unique masses of buyers D^w and sellers S^w that satisfy the steady-state conditions. The market is tighter in the high state; that is, $\frac{D^H}{S^H} > \frac{D^L}{S^L}$.

The lemma is intuitive: The larger mass of buyers in the high state implies that more buyers stay in the market because each buyer has a smaller chance to transact. Moreover, each seller has a higher chance to transact in the high state, so the sellers leave the market more quickly, and there are fewer sellers on the market in the high state.

Duffie and Sun (2007).

A distribution of beliefs is said to have the no-introspection property if

$$\frac{\theta}{1-\theta} = \frac{D^H}{D^L} \frac{\gamma^H(\theta)}{\gamma^L(\theta)} \tag{6}$$

for all $\theta < 1$ with $\gamma^L(\theta) > 0$. The condition implies that a buyer does not update based merely on observing its own belief ("introspection"). The following lemma follows from the steady-state conditions.

Lemma 2 (No-Introspection.) If Γ^w is an atomless and piecewise twice continuously differentiable c.d.f. and if Γ^w satisfies the steady-state conditions given the steady-state masses S^w , D^w , then Γ^w and D^w have the no-introspection property.

The distribution of beliefs satisfies the "monotone likelihood ratio property" (MLRP) if $\gamma^H(\theta'') \gamma^L(\theta') \geq \gamma^H(\theta') \gamma^L(\theta'')$ whenever $\theta'' \geq \theta'$. The no-introspection condition implies the MLRP. Intuitively, observing a buyer with a higher belief makes the high state more likely. The no-introspection condition also implies that the distributions Γ^H and Γ^L have identical support, given that there are no atoms at zero or one.

We use the MLRP to characterize updating. Suppose that $0 < \Gamma^L(\theta) < 1$. The MLRP implies that $1 - \Gamma^H(\theta) > 1 - \Gamma^L(\theta)$. By Lemma 1, $\mu^H > \mu^L$. Therefore, $\mu^H\left(1 - \Gamma^H(\theta)\right) > \mu^L\left(1 - \Gamma^L(\theta)\right)$, the expected number of buyers with belief above θ who are matched with a seller is higher in the high state. From the definition of $\Gamma^w_{(1)}$, $1 - \Gamma^H_{(1)}(\theta) > 1 - \Gamma^L_{(1)}(\theta)$, the likelihood of losing is higher in the high state for any θ . Hence, "losing is bad news," and the posterior conditional on losing satisfies $\theta^+(x,\theta) > \theta$ whenever $0 \le \Gamma^L(x) < 1$; see Definition (2). An implication is that all the buyers in the stock must have beliefs above the most optimistic type in the inflow $\underline{\theta}$: All of those buyers who have just entered hold beliefs above the cutoff $\underline{\theta}$. For all other buyers in the stock who have entered at least one period before, the finding that $\theta^+(x,\theta) > \theta$ for all θ implies that their beliefs are above $\underline{\theta}$ as well. Therefore, all beliefs in the stock are above the cutoff $\underline{\theta}$.

We also need the posterior conditional on being tied when characterizing optimal bidding. Conditional on state w, the density of the first-order statistic is $\gamma_{(1)}^w = \frac{D^w}{S^w} \gamma^w \Gamma_{(1)}^w$. The unconditional density is $\gamma_{(1)}^\theta(x) = \theta \gamma_{(1)}^H(x) + (1-\theta)\gamma_{(1)}^L(x)$. The posterior of type θ after tying with a buyer with belief x at the top spot is

$$\theta^{0}(x,\theta) = \frac{\theta \gamma_{(1)}^{H}(x)}{\gamma_{(1)}^{\theta}(x)} \tag{7}$$

¹⁹This informal argument is verified in the proof of Lemma 4.

if $\gamma_{(1)}^{\theta}(x) > 0.20$ The next lemma establishes that updating is monotone.

Lemma 3 (Monotonicity of Posteriors.) Suppose that Γ^w is an atomless and piecewise twice continuously differentiable c.d.f., (i) the monotone likelihood ratio property holds, and (ii) $\mu^H \geq \mu^L > 0$. Then, the posterior upon losing, $\theta^+(x,\theta)$, is nondecreasing in x. If, in addition, (iii) $\gamma^H(\theta) \geq \gamma^L(\theta) \frac{\mu^L}{\mu^H}$ for all θ , then the posterior upon being tied, $\theta^0(x,\theta)$, is nondecreasing in x on [0,1].

A standard sufficient condition for monotonicity of the posteriors would be that the first-order statistic $\theta_{(1)}$ inherits the monotone likelihood ratio property of the parent distribution of θ . However, in contrast to standard auction settings, the MLRP is not inherited here because the first-order statistic is taken from a *random* number of random variables.²¹

Given a symmetric equilibrium, the posterior conditional on losing is $\theta^+(\theta,\theta)$. An implication of Lemma 3 is that this posterior is strictly increasing in θ . The same holds for the posterior conditional on tying, $\theta^0(\theta,\theta)$. This follows from the fact that conditions (i)—(iii) hold in equilibrium: Conditions (i) and (ii) follow from Lemma 2 and Lemma 1, respectively. For Condition (iii), note the following: The support of Γ^w is a subset of $[\underline{\theta},1]$; see the previous remark following $\theta^+(x,\theta)>\theta$. Therefore, the no-introspection property from Lemma 2 and $S^H \leq S^L$ from Lemma 1 together imply $\frac{\gamma^H(\theta)}{\gamma^L(\theta)} \frac{\mu^H}{\mu^L} \geq \frac{\gamma^H(\theta)}{\gamma^L(\theta)} \frac{D^H}{D^L} \geq \frac{\theta}{1-\theta}$. Finally, the assumption that $\underline{\theta} \geq 1/2$ implies $\underline{\theta} \geq 1$; that is, (iii) holds.

We show that the distribution of beliefs of buyers in the stock is unique. Together with the previous finding that the mass of buyers and sellers is unique, the lemma implies that there exists a unique steady-state stock.

Lemma 4 (Uniqueness of the Steady-State Distributions.) There exists a unique absolutely continuous and piecewise twice continuously differentiable distribution Γ^w that satisfies the steady-state conditions.

We describe the basic idea of the proof and some of the complications here. To illustrate the construction and the uniqueness argument, let us suppose momentarily that we have found some stock Γ^w, D^w, S^w that satisfies the steady-state conditions and suppose further that the interval of initial beliefs $[\underline{\theta}, \overline{\theta}]$ is sufficiently small such that upon updating, $\theta^+(\underline{\theta}, \underline{\theta}) > \overline{\theta}$. Therefore, the set of beliefs of buyers who have lost once is above

²⁰We extend the definition of the posterior to all types: If $\min \{\Gamma^L(x), \Gamma^H(x)\} < 1$, we set $\theta^0(x,\theta) = \inf\{(\theta^0(x',\theta)) | x' \geq x \text{ and } \gamma^{\theta}_{(1)}(x) > 0\}$ and if $\Gamma^L(x) = \Gamma^H(x) = 1$, we set $\theta^0(x,\theta) = \sup\{(\theta^0(x',\theta)) | x' \leq x \text{ and } \gamma^{\theta}_{(1)}(x) > 0\}$. Bidders do not observe whether they are tied, and the particular choice of the extension of Bayes' formula does not affect our analysis.

²¹We provide a detailed discussion of the failure of the MLRP of the first-order statistic with a random number of bidders in our supplementary online appendix in Section 1.3.

the interval of the initial beliefs. Consequently, the mass of buyers with beliefs below any $\theta' \in [\underline{\theta}, \overline{\theta}]$ is just the mass of such buyers in the inflow; that is, $D^w \Gamma\left(\theta'\right) = d^w G^w\left(\theta'\right)$. Moreover, the mass of buyers with beliefs above θ' is $D^w - d^w G^w\left(\theta'\right)$. Therefore, the probability that a buyer with belief θ' loses is $e^{-(D^w - d^w G^w(\theta'))/S^w}$ in each state w. This determines the posterior of θ' after losing once, $\theta^+\left(\theta',\theta'\right)$. Conversely, for any θ'' from the set of beliefs who lost once— $\left[\theta^+\left(\underline{\theta},\underline{\theta}\right),\theta^+\left(\overline{\theta},\overline{\theta}\right)\right]$ —we can find the prior $\hat{\theta}$ such that $\theta^+(\hat{\theta},\hat{\theta}) = \theta''$. Taking our observations together, the distribution of beliefs of buyers who have lost once is given by

$$D^{w}\Gamma^{w}\left(\theta''\right) = d^{w} + \delta d^{w} \int_{\underline{\theta}}^{\hat{\theta}} e^{-(D^{w} - G^{w}(\tau))/S^{w}} dG^{w}\left(\tau\right) \qquad \forall \theta'' \in \left[\theta^{+}\left(\underline{\theta}, \underline{\theta}\right), \theta^{+}\left(\overline{\theta}, \overline{\theta}\right)\right].$$

The above reasoning suggests that we can construct the population of buyers inductively, starting with the distribution of initial beliefs and then proceeding to the distribution of beliefs of buyers who have lost once, twice, ... and so on. Furthermore, the above arguments suggest that the construction would yield a unique candidate for an equilibrium steady-state stock.

The existence proof is based on induction, following the line of reasoning laid out before. There are two difficulties with the argument, however. First, we have assumed that intervals of beliefs of successive generations of buyers do not overlap. This does not need to be the case. To take care of this problem, we use the fact that the losing probabilities of the lowest type $\underline{\theta}$ are determined by the total masses D^w and S^w , which are unique by Lemma 1. This determines the posterior $\theta^+(\underline{\theta},\underline{\theta})$ and implies that the set of buyers with beliefs in $[\underline{\theta}, \min\{\theta^+(\underline{\theta},\underline{\theta}), \overline{\theta}\}]$ is given by the inflow. We then apply similar arguments successively. Second, the construction above uses the fact that posteriors after losing are monotone in priors. However, the argument following Lemma 3 for monotonicity of θ^+ presupposes the steady-state conditions to conclude that no-introspection holds. When we prove the existence of a steady-state stock, we need to directly ensure that the conditions of Lemma 3 hold, which is done in the main technical lemma of the proof, Lemma 12.

4.2 Characterization of Bidding and Existence of Equilibrium

We characterize the equilibrium bidding strategy. Let

$$EU(\theta,\beta|w) = v\Gamma_{(1)}^{w}(0) + \int_{\underline{\theta}}^{\theta} (v - \beta(\tau)) d\Gamma_{(1)}^{w}(\tau) + \delta\left(1 - \Gamma_{(1)}^{w}(\theta)\right) EU(\theta^{+}(\theta,\theta),\beta|w)$$

denote the expected utility of a bidder with belief θ given a symmetric bidding strategy β , conditional on state w. The unconditional expected payoff is $EU(\theta, \beta|\hat{\theta}) = \hat{\theta}EU(\theta, \beta|H) + (1-\hat{\theta})EU(\theta, \beta|L)$. The function $EU(\theta, \beta|\hat{\theta})$ can be interpreted as type

 $\hat{\theta}$'s expected (off-equilibrium) payoff from bidding like type θ .

We prove that equilibrium bids must be

$$\beta(\theta) = v - \delta EU\left(\theta^{+}\left(\theta, \theta\right), \beta | \theta^{0}\left(\theta, \theta\right)\right). \tag{8}$$

An intuition for this bidding strategy is as follows. By standard reasoning about bidding in second-price auctions, the bid must be "truthful" and equal to the expected payoff from winning conditional on being tied. Here, the expected payoff from winning is equal to the valuation v minus the relevant continuation payoff. For the relevant continuation payoff, note that the strategy adopted from tomorrow onwards is the optimal strategy given the updated belief conditional on having lost, $\theta^+(\theta,\theta)$. We need to evaluate the expected value of that strategy using the posterior conditional on being tied (the "pivotal event"). Therefore, the expected continuation payoff is calculated by evaluating the utility derived from the future bidding sequence of a bidder with belief $\theta^+(\theta,\theta)$, given the posterior probability of the high state conditional on being tied, $\theta^0(\theta,\theta)$. Thus, the relevant continuation payoff is $\delta EU\left(\theta^+(\theta,\theta),\beta|\theta^0(\theta,\theta)\right)$, and buyers optimally "shade" their bids by this amount.

We provide some auxiliary observations. First, the value function is convex in beliefs: Optimal bidding is a decision problem under uncertainty, implying a convex value function by standard arguments from information economics. Second, the envelope theorem dictates a simple relation between EU, V, and the derivative V'.

Lemma 5 (Characterizing the Value Function.) The value function $V(\theta)$ is convex. At all interior differentiable points of the value function, $V'(\theta) = \frac{\partial}{\partial \hat{\theta}} EU(\theta, \beta | \hat{\theta})|_{\hat{\theta} = \theta}$, and

$$EU(\theta, \beta | \hat{\theta}) = V(\theta) + (\hat{\theta} - \theta)V'(\theta). \tag{9}$$

The following Lemma establishes a unique candidate for the equilibrium bidding function for given continuation payoffs. The lemma follows from rewriting the necessary first-order condition for optimal bids; that is, we determine the derivative of the objective function (5) with respect to b and set it equal to zero.

Lemma 6 (Equilibrium Candidate.) For almost all types in the support of the distribution of beliefs, in equilibrium

$$\beta(\theta) = v - \delta V \left(\theta^{+} \left(\theta, \theta \right) \right) + \delta V' \left(\theta^{+} \left(\theta, \theta \right) \right) \left(\theta^{+} \left(\theta, \theta \right) - \theta^{0} \left(\theta, \theta \right) \right). \tag{10}$$

We can use Lemma 5 to substitute EU for V' and V in equation (10). After the substitution, the expression for the bidding strategy is as claimed in the equation (8).

We have identified a unique candidate for the equilibrium bidding strategy for given continuation payoffs in this section. We have also proven that there exists a unique steady-state stock in Section 4.1. The following proposition shows that there exists an equilibrium. The exogenous parameters— δ , d^H , and d^L —determine the market outcome in an essentially unique way. The proof is in the online appendix.

Proposition 1 (Existence and Uniqueness of Equilibrium.) There exists a steady-state equilibrium in strictly increasing strategies. The equilibrium distribution of beliefs and the value function $V(\theta)$ are unique. For almost all types in the support of the distribution of beliefs, the bidding function is $\beta = v - \delta EU(\theta^+(\theta, \theta), \beta|\theta^0(\theta, \theta))$.

5 Price Discovery with Small Frictions

We state and prove our main result: as the exit rate becomes small, the equilibrium trading outcome becomes competitive in each state. In particular, all trade between buyers and sellers takes place at the "correct," market-clearing prices.

We define trading outcomes. For buyers, the trading outcome in state w consists of the equilibrium probability of winning in an auction (instead of being forced to exit) and the expected price paid conditional on winning, denoted $q^w(\theta)$ and $p^w(\theta)$, respectively. For a seller, the trading outcome consists of a probability of being able to sell the good and the expected price received, denoted $q^w(S)$ and $p^w(S)$. The inflow defines a large quasilinear economy, where the mass of buyers is d^w and the mass of sellers is independent of w and equal to one. A trading outcome is said to be a (perfectly) competitive outcome (or Walrasian outcome) relative to the economy defined by the inflow if prices and trading probabilities are as follows. If $d^w < 1$ (i.e., if buyers are on the short side of the market), then $p^w(\theta) = p^w(S) = 0$, $q^w(\theta) = 1$, and $q^w(S) = d^w$. If $d^w > 1$ (i.e., if buyers are on the long side of the market), then $p^w(\theta) = p^w(S) = v$, $q^w(\theta) = 1/d^w$, and $q^w(S) = 1$. We do not characterize the competitive outcome in the case in which both market sides have equal size, $d^w = 1$. If an outcome is competitive, it is necessarily an efficient outcome relative to the economy defined by the inflow.

We consider the trading outcome when the exit rate is small. Let $\{\delta_k\}_{k=1}^{\infty}$ be a sequence such that the exit rate converges to zero, $\lim (1 - \delta_k) = 0$. Intuitively, a smaller exit rate corresponds to a smaller cost of searching. To interpret our results, it might be helpful to observe that decreasing the exit rate is equivalent to increasing the speed of matching.²² We know that an equilibrium exists for each δ_k . Pick any such equilibrium

 $^{^{22}}$ Let Δ_k denote the length between time periods and let d denote the (fixed) exit probability per unit of time. With $1 - \delta_k = \Delta_k d$, one can interpret a decrease in the exit rate $1 - \delta_k$ as a decrease in Δ_k . In this interpretation, the market friction $1 - \delta_k$ arises because it takes time Δ_k to come back to the market after a loss. As this time lag Δ_k goes to zero, the friction vanishes.

and denote the corresponding equilibrium magnitudes by $\beta_k, \Gamma_k^H, \Gamma_k^L, D_k^H, p_k^w, q_k^w$ and so on. A sequence of trading outcomes converges to the competitive outcome relative to the economy defined by the inflow in state w if the sequence of outcomes converges pointwise for all θ and for S.

Proposition 2 (Price Discovery with Small Frictions.) For any sequence of vanishing exit rates and for any sequence of corresponding steady-state equilibria in strictly increasing strategies the sequence of trading outcomes converges to the competitive outcome for each state of nature.

We illustrate the proposition through a few observations. First, we restate the implications of the proposition in terms of the limit of the value function.

Corollary 1 (Limit Payoffs.) For any sequence of equilibria for $\lim_{k\to\infty} (1-\delta_k) = 0$: $\lim_{k\to\infty} V_k(\theta) \equiv v$ if $d^L < d^H < 1$; $\lim_{k\to\infty} V_k(\theta) \equiv 0$ if $1 < d^L < d^H$; and $\lim_{k\to\infty} V_k(\theta) = (1-\theta)v + \theta 0$ if $d^L < 1 < d^H$.

The corollary is immediate and the proof is omitted. Intuitively, the short side of the market captures the surplus from trading. Moreover, the corollary states that the value function is no longer convex but linear in the limit. Information loses its value when the friction of trade is small.

The following result is the main intermediate step towards proving the Proposition. The lemma illustrates some of the main forces at work.

Lemma 7 (Limit Market Population.) For any sequence of equilibria for $\lim_{k\to\infty} (1-\delta_k) = 0$ the following statements hold: (i)

$$\lim_{k \to \infty} \frac{D_k^w}{S_k^w} = \begin{cases} 0 & \text{if } d^w < 1\\ \infty & \text{if } d^w > 1. \end{cases}$$

(iia) If $d^w < 1$, the probability of being the sole bidder becomes one, $\lim_{k \to \infty} e^{-D_k^w/S_k^w} = 1$. (iib) If $d^w > 1$, the probability of being the sole bidder converges to zero and it converges to zero faster than the exit probability, $\lim_{k \to \infty} \frac{e^{-D_k^w/S_k^w}}{1-\delta_k} = 0$.

In the following, we describe equilibrium when the exit rate is small. We consider the case in which $d^L < 1 < d^H$. In that case, buyers are on the short side in the low state and on the long side in the high state. The cases in which $d^H < 1$ or $d^L > 1$ are less interesting because in these cases it is known whether the buyers or the sellers are on the short side of the market.

As stated in the previous lemma, the difference between the sizes of the market sides in the inflow is magnified in the stock. Therefore, in the low state, the number of buyers per seller vanishes to zero, and a buyer is almost sure to be the sole bidder. In the high state, the number of buyers per seller diverges to infinity, and there is almost never only a single bidder.

The fact that a buyer becomes sure to be the sole bidder in the low state has an immediate implication: If a buyer is the sole bidder, the buyer wins and pays nothing. Therefore, in the low state, payoffs must converge to v for all buyers. Consequently, the characterization of the limit trading outcome conditional on the low state is relatively straightforward. The corresponding fact for the high state has no immediate implication, however. Even though a bidder becomes less and less likely to be the sole bidder, the bidder also becomes more and more patient. If the bidder becomes patient fast enough, the bidder could just wait to be the sole bidder and receive the good for free too. In the second part of (iib), we show that this is not the case. Relative to the buyer's increasing patience, the probability of being the sole bidder converges to zero even faster. Therefore, in the limit, a buyer who uses the strategy of always bidding lowest would almost surely have to exit the market before being able to trade.

Let us discuss bidding. Lemma 7 implies that conditional on the high state buyers learn that the state is high very quickly after entry, after losing only once. This is because losing is very unlikely in the low state, but it is very likely in the high state. However, optimal bids depend on the buyer's belief conditional on being pivotal (tied at the top). Therefore, we characterize beliefs conditional on being pivotal. As the next result shows, being pivotal is very positive news, indicating that the low state is very likely. Intuitively, the growing imbalance of the two sides of the market implies a very strong winner's curse.

Let $\theta_k^t(\theta)$ denote the posterior of a buyer who has entered the market with a prior θ and who has lost t times. The following proposition characterizes $\theta_k^0 \left(\theta_k^t \left(\theta \right), \theta_k^t \left(\theta \right) \right)$, the posterior conditional on being pivotal (tied at the top) after having lost t times before.

Proposition 3 (Time Pattern on Bids.) Suppose that $d^L < 1 < d^H$ and suppose that

$$\lim_{k \to \infty} (1 - \delta_k) = 0. \text{ Let } \theta_k^t = \theta_k^t(\theta) \text{ for some } \theta \in [\underline{\theta}, \overline{\theta}].$$
(i) For any t , (a) $\lim_{k \to \infty} \theta_k^0(\theta_k^t, \theta_k^t) = 0$, and (b) $\lim_{k \to \infty} \beta_k(\theta_k^t) = 0$.
(ii) Let $t_k \equiv \frac{-0.5}{(1 - \delta_k) \ln(1 - \delta_k)}$. Then, (a) $\lim_{k \to \infty} \theta_k^0(\theta_k^{t_k}, \theta_k^{t_k}) = 1$, and (b) $\lim_{k \to \infty} \beta_k(\theta_k^{t_k}) = v$.

The proposition is not part of the proof of Proposition 2. In fact, we use findings from the proof of Proposition 2 to prove it. We state the proposition because we believe it provides some interesting insights into bidding when the exit rate is small.

As discussed previously, $\lim \theta_k^t(\theta) = 1$ as $k \to \infty$ for all $\theta \in [\underline{\theta}, \overline{\theta}]$ and for all $t \ge 1$; that is, bidders who have lost at least once have a posterior that puts probability close to

one on w = H. The first part of the proposition states, however, that the event of being tied conditional on having belief θ_k^t is sufficiently "good news" such that the posterior switches to putting probability zero on the high state. Therefore, conditional on being tied, a buyer believes that the continuation payoff is close to v and, consequently, bids zero. Two related and immediate consequences of Part (i) of Proposition 3 are that when $(1 - \delta_k) \to 0$, (a), fixing any time t, the bid of a buyer who has entered t periods ago decreases to zero when $(1 - \delta_k)$ decreases, and, (b), buyers bid close to zero for an increasingly long time. This observation illustrates why proving convergence to the competitive outcome is not immediate from the fact that the imbalance of the masses of buyers and sellers in the stock explodes in the high state.

Of course, the finding from Proposition 2 requires that buyers stop bidding zero at some time and bid close to v eventually. This is reflected in Part (ii) of Proposition 3. The significance of that part is that the number t_k is chosen so that $\lim (\delta_k)^{t_k} = 1$ as $k \to \infty$: the probability of exogenous exit within t_k periods is vanishing to zero. Part (ii) states that after at most t_k periods buyers are eventually sufficiently pessimistic that they bid high. The fact that $\lim (\delta_k)^{t_k} = 1$ can be interpreted as saying that buyers start bidding high "quickly" relative to the exit rate $(1 - \delta_k)$.

Proposition 3 illustrates the combined effect of the winner's and the loser's curse. Initial bids are predominantly shaped by the winner's curse (buyers bid cautiously low to avoid winning in the low state). Eventually, however, the loser's curse is sufficiently strong so that after losing at most t_k number of periods, buyers bid close to their maximum willingness to pay.

Proposition 3 illustrates how the presence of aggregate uncertainty affects bidding, and, consequently, expected prices. We have conducted a preliminary analysis to compare the expected price conditional on the high state in the current model with the expected price in a model in which the state is known. Our analysis suggests that the relative distance from the competitive price is much higher with aggregate uncertainty than without. The reason is precisely the fact that with uncertainty the initial bids are very low in the high state. However, we chose not to provide a complete analysis of the relative speed of convergence because we believe this to be beyond the scope of this paper.

6 Extension: Heterogeneous Buyers

Buyers in our model have one-dimensional types (beliefs). This is due to our assumptions that there is a binary state of nature and that the buyers have homogeneous preferences. Assuming one-dimensional types makes our model tractable and explicitly solvable. Specifically, this assumption enables us to provide an explicit characterization

of the endogenous population of traders, and it allows us to use standard techniques from auction theory to characterize bidding.

We now generalize our characterization result to a setting in which buyers have heterogeneous preferences. While we cannot prove existence of equilibrium, this extension gives rise to a somewhat richer economic environment. For example, in our base model, whether or not prices aggregate information has no consequence for welfare. Our model shares this feature with many of the standard models of information aggregation in large common value auctions; see, e.g., Milgrom (1979) and Pesendorfer and Swinkels (1997). With heterogeneous preferences, information aggregation is consequential for efficiency.²³

We extend our model as follows. The mass of buyers who enter the market is either d^L or d^H , where $d^L < d^H$, as before. However, buyers' valuations are now drawn from a finite set $\mathcal{V} \subset [0,1]$. The share of buyers who have valuation $v \in \mathcal{V}$ is f(v) and $\sum_{v \in \mathcal{V}} f(v) = 1$. Each entering buyer privately observes a signal. The distribution of posteriors induced by the signal is given by G^w , as before. Thus, we assume that valuations and signals are independently and identically distributed conditional on w. The mass of the entering sellers is one in either state. We consider only the case where $d^L > 1$. We assume that the distribution f is non-degenerate in the following sense. There are two marginal types v_*^L and v_*^H and rationing variables r^L and r^H , $0 < r^w < 1$, such that

$$d^{w}(r^{w}f(v_{*}^{w}) + \sum_{v > v_{*}^{w}} f(v)) = 1, \qquad w \in \{L, H\}.$$
(11)

In the competitive outcome (allocation) for the quasilinear economy defined by f and d^w , buyers with valuations above v_*^w receive the good. Buyers with valuation v_*^w are rationed and only a share r^w of these buyers receives the good. The competitive price is $p_*^w = v_*^w$. We assume $p_*^H > p_*^L$. We omit the degenerate case in which (11) holds for some v_*^w with r = 0. In that case, the competitive price need not be uniquely determined by market clearing.

We consider monotone steady-state equilibria. Equilibria consist of a bidding strategy $\beta: \mathcal{V} \times [0,1] \to [0,1]$, an updating rule $\theta^+(\theta,b)$, and a steady-state population, characterized by the mass of buyers and sellers, D^w and S^w , and a probability measure on types $\mathcal{V} \times [0,1]$ (valuations and beliefs), denoted by Φ^w . We restrict attention to monotone equilibria, by which we mean that $\beta(v,\theta)$ is strictly increasing in v and θ , and that θ^+ is strictly increasing in θ and weakly increasing in θ . Furthermore, we assume that the marginal distribution of beliefs, $\Phi^w(\theta|v)$, is atomless and we assume that equilibrium bids are given by $\beta(v,\theta) = v - \delta EU\left(v,\theta^+(\theta,\beta(v,\theta))|\theta^0(\theta,\beta(v,\theta))\right)$, where $\theta^0(\theta,\beta(v,\theta))$ is the posterior conditional on being tied. We conjecture that there

²³For example, Pesendorfer and Swinkels (2000). study efficiency of a large double auction with heterogeneous preferences.

is an essentially unique steady-state equilibrium satisfying these requirements. However, we have not been able to prove either existence of equilibrium or uniqueness.

Let $q^w(v,\theta)$ denote the lifetime trading probability and let $p^w(v,\theta)$ denote the expected price conditional on trading. Similarly, $q^w(v) = \int q^w(v,\theta) dG^w(\theta)$ and $p^w(v) = \int p^w(v,\theta) dG^w(\theta)$. The monotonicity of the equilibrium implies that $q^w(v,\theta)$ is nondecreasing in v and v, and v

In the competitive outcome the buyers' trading probabilities are $q^w(v) = 0$ for $v < v_*^w$, $q^w(v_*^w) = r^w$, and $q^w(v) = 1$ for $v > v_*^w$, the sellers' trading probabilities are $q^w(S) = 1$, and the price is $p^w(v) = p_*^w$ for $v \ge v_*^w$. Given a sequence of vanishing exit rates $\{1 - \delta_k\} \to 0$, suppose there is a sequence of monotone steady-state equilibria with trading probabilities q_k^w and prices p_k^w . We say that the sequence of outcomes converges to the competitive outcome if the trading probabilities and prices converge to the competitive outcome.

Proposition 4 (Revealing Prices with Heterogeneous Buyers.) Consider an economy with heterogeneous buyers. For any sequence of vanishing exit rates and for any sequence of corresponding monotone steady-state equilibria, the sequence of trading outcomes converges to the competitive outcome for each state of nature.

The proof is relegated to our online appendix. The proof involves two important observations. First,

$$\frac{e^{-D_k^w \Phi_k^w(\{(v,\theta)|v \ge v_*^w\})/S_k^w}}{1 - \delta_k} \to 0; \tag{12}$$

that is, the probability that a buyer ends up participating some time in an auction in which there is no bidder present who has a valuation weakly above v_*^w vanishes to zero. Intuitively, it becomes almost common knowledge that all bidders have valuations of at least v_*^w . This observation is an important step towards showing that buyers bid up the price to at least v_*^w ; that is, in state w, the price is at least p_*^w . Equation (12) follows from the fact that there are more buyers with valuations weakly above v_*^w entering the market than there are sellers entering. Thus, buyers having valuations $v \geq v_w^*$ accumulate in the stock. Note the similarity to the previous finding in Lemma 7, Part (ii). Second,

$$\frac{e^{-D_k^w \Phi_k^w(\{(v,\theta)|\beta_k((v,\theta)) \ge b'\})/S_k^w}}{1 - \delta_k} \to \infty, \qquad \forall b' > v_*^w; \tag{13}$$

that is, when a buyer keeps bidding $b' = v_*^w + \varepsilon$, the buyer trades eventually with a probability converging to one at a price that is at most b'. The second observation implies that buyers pay at most p_*^w . Equation (13) follows from the fact that (i) only buyers having valuation strictly above v_*^w bid higher than b' and that (ii) there are fewer buyers having such values who enter the market than there are sellers.

7 Discussion and Conclusion

7.1 Discussion of Assumptions and Extensions

Bid Disclosure. In our model, learning is "minimal": losing buyers learn nothing except that they lost. In our companion paper, Lauermann and Virag (2011), we ask whether such nontransparent auctions would arise if each seller could individually choose the auction format. We show that sellers have an incentive to hide information from the buyers because of a "continuation value effect": If bidders receive information when losing, then they can refine their future bids, which raises the expected value of their outside options. This leads to less aggressive bidding and lower revenues for the seller. Countervailing the continuation value effect is the well-known linkage principle effect for common value auctions. We study how these two effects determine the sellers' preferences for information disclosure. For example, we show that the sellers do not have an incentive to reveal any information about the submitted bids after the auction.

Multiple States. The assumption that there are two states of nature is a standard method to ensure that learning is tractable. As explained in Section 6, the combination of a binary state of nature and homogeneous preferences is what allows us to provide an explicit characterization of the equilibrium. While we cannot extend our explicit equilibrium characterization, we believe that our characterization of the limit trading outcome extends to a variation in which there are multiple states of nature. Intuitively, what matters is only whether or not buyers are on the long side of the market in the inflow. Suppose there are M states and the mass of entering buyers is d^m in state m, ordered such that $d^m < d^{m+1}$. The mass of entering sellers is one. It is straightforward to show that the outcome will be competitive in the limit for all those states where $d^m < 1$. For those state with $d^m > 1$, we believe that the proof of Proposition 2 extends too; that is, buyers bid up prices to v. However, we have not formally verified this conjecture.

Reservation Prices. In our model, the sellers do not set a reservation price. More generally, the sellers take no actions.²⁴ The absence of reservation prices lets us avoid two modeling problems, Diamond's paradox and multiplicity of equilibrium that is due to freedom in assigning off-equilibrium beliefs.

Diamond's paradox: Diamond (1971) found—roughly speaking—that if the sellers have all the market power in a simple search model then the sellers can charge monopoly

²⁴Note that in our model individual sellers do not have a strong incentive to set reservation prices when the exit rate is small, because the expected winning bid is close to the buyers' maximum willingness to pay in each state, given the continuation payoffs. For example, suppose we modify our model by giving sellers the option to set a positive reservation price at a small cost. Then, for a sufficiently small exit rate, it would be an equilibrium for sellers to not use that option. Given that no seller sets a positive reservation price, the resulting equilibrium outcome would be equivalent to the equilibrium outcome of our original model. This argument depends, of course, on the cost of setting a reservation price being large relative to the exit rate.

prices. In our model, if the sellers can set reservation prices, a version of Diamond's paradox emerges, and in the unique equilibrium the bids and the reservation prices are equal to the buyers' valuations. However, it is well known that Diamond's paradox can be avoided in richer search models. For example, Satterthwaite and Shneyerov (2008) have shown that the paradox can be avoided by assuming that buyers have heterogeneous valuations. We conjecture that this would be the case here, too; that is, there would be nontrivial equilibria in a variation with reservation prices and heterogeneous valuations.

Multiplicity of equilibria: If the reservation price is observable, it becomes a signal of the seller's beliefs about the state of the market, and there is freedom in assigning beliefs following off-equilibrium reservation prices. This freedom can then be used to support multiple equilibria, possibly also equilibria that are not competitive in the limit. One might be able to impose known refinements for signaling games and show the implications of some reasonable refinement of beliefs for limit outcomes. The advantage of the current model is that we do not need such refinements.

In summary, it might be possible to introduce reservation prices once the model is further extended to heterogeneous buyers as in Section 6 (to avoid Diamond's paradox) and once some refinement is imposed (to select among the multiple equilibria that might arise otherwise). However, including reservation prices makes the model considerably more complex. We therefore chose not to include reservation prices.

7.2 Conclusion

We provide a framework to study price discovery through trading in a decentralized market. In our model, buyers learn about the relative scarcity of a good through repeated bidding in auctions. In particular, individual traders never observe the whole market and they directly interact only with small groups of traders. We characterize the resulting distribution of beliefs in the population, the learning process, and the bidding behavior of buyers. Despite the fact that there is no centralized price formation mechanism, we show that the equilibrium trading outcome is approximately Walrasian when the exit rate is small and search becomes cheap. Thus, prices reveal aggregate scarcity and correctly reflect economic value. We discussed possible extensions of our model. A particularly interesting question for further research might be the effect of aggregate uncertainty on the speed of convergence. Because we provide an explicit characterization of equilibrium, our model might be well suited to study this question. We conjecture that whenever there is uncertainty about the market-clearing price, the outcome is further away from the competitive outcome for small frictions than when the market-clearing price is commonly known ex ante.

We have emphasized the analysis of trading outcomes with small frictions and related our work to research that studies foundations for general equilibrium. Nevertheless, we also provide a complete characterization of equilibrium bidding and learning for all level of frictions. We believe that our model is a first step towards a tractable framework to study the search behavior of agents who learn about the economy in an equilibrium analysis. Such a framework has been lacking so far. Search theory has proven to be a remarkably useful tool for studying important decentralized markets, such as those for housing, certain financial assets, and labor. However, little progress has been made in incorporating the possibility of learning, which has limited the applicability of search theory. Rothschild (1974) criticized early on that "the results [from search theory] depend on the untenable assumption that searchers know the probability distribution from which they are searching," (p. 689) and he writes further that "it seems absurd to suppose that consumers know them [the price distributions] with any reasonable degree of accuracy" (p692).²⁵ Beyond their conceptual importance, equilibrium models of search with learning may also be helpful in explaining some empirical findings. For example, standard models of search assume that searchers know the actual distribution that they are sampling and from have difficulties explaining observed large dispersions of accepted prices and wages. Equilibrium models of search with learning provide a potential explanation because searchers with different beliefs about the prevailing market conditions have different acceptance behavior, adding a novel source of heterogeneity.

²⁵Rothschild (1974) himself and subsequent work has characterized optimal rules for sampling from unknown distributions. However, these are models of single-person decision problems where the distribution of prices is exogenous instead of equilibrium search models. This restriction arose presumably because of the difficulties associated with developing tractable model of equilibrium search with learning.

8 Appendix

This appendix contains the proofs of our main results from Section 5 (Price Discovery with Small Frictions). The proofs of the results from Section 4 (Characterization and Existence of Equilibrium) and Section 6 (Extension: Heterogeneous Buyers) are contained in an **Online Appendix**. The exceptions are the proofs of Lemma 1 (Unique Masses), Lemma 5 (Envelope Theorem) and Lemma 6 (Characterizing the Equilibrium Bidding) which we give here, too. We keep these three proofs in the regular appendix because we use elements from these proofs when showing convergence and because the characterization of the equilibrium bids is a cornerstone of our convergence proof.

8.1 Proof of Lemma 1 (Uniqueness of Masses)

We show that the steady-state conditions for the stocks can be written as:

$$d^{w} = (1 - \delta) D^{w} + \delta S^{w} \left(1 - e^{-D^{w}/S^{w}} \right)$$
(14)

$$1 = (1 - \delta) S^{w} + \delta S^{w} \left(1 - e^{-D^{w}/S^{w}} \right). \tag{15}$$

These conditions have a simple interpretation: the left-hand side is the inflow for each market side. The right-hand side is outflow from each market side, that is, the sum of the number of traders who exit through discouragement and the number of traders who exit through trade. The number of traders who exit through trade is $S^w \left(1 - e^{-D^w/S^w}\right)$, which is equal for both market sides.

Rewriting the steady-state condition for buyers, (4),

$$D^{w} = d^{w} + \delta D^{w} \int_{0}^{1} \left(1 - e^{-D^{w}(1 - \Gamma^{w}(\theta))/S^{w}}\right) d\Gamma^{w}(\theta)$$

$$= d^{w} + \delta D^{w} - \delta D^{w} \int_{0}^{1} \frac{\partial}{\partial \theta} \left(S^{w} e^{-D^{w}(1 - \Gamma^{w}(\theta))/S^{w}}\right) d\theta$$

$$= d^{w} + \delta D^{w} - \delta S^{w} \left(1 - e^{-D^{w}/S^{w}}\right).$$

Recall the steady-state condition for sellers, (3), $S^w = 1 + \delta S^w e^{-D^w/S^w}$. Rewriting the steady-state conditions further yields (14) and (15).

A solution to the steady-state conditions exists, and the solution is unique. The difference of (14) and (15),

$$d^{w} - 1 = (1 - \delta) (D^{w} - S^{w}), \tag{16}$$

defines D^w as a function of S^w , $J(S^w)$. We can write (15) as a function of S^w only,

$$1 = (1 - \delta) S^{w} + \delta S^{w} \left(1 - e^{-J(S^{w})/S^{w}} \right).$$

This equation has a solution by the intermediate value theorem. At $S^w \to 0$, the right-hand side becomes zero, while for $S^w \to \infty$, the right-hand side becomes infinite (recall that $(1 - e^{-J(S^w)/S^w}) \in [0,1]$). A solution exists in between.

The solution is unique. Let S'^w, D'^w and S''^w, D''^w be two solutions, and suppose that $D''^w \geq D'^w$. Then, by (16), $S''^w \geq S'^w$. We can show that (15) and $S''^w > S'^w$ leads to a contradiction. Hence, it must be that $S''^w = S'^w$, which implies $D''^w = D'^w$ by (16). The contradiction arises as follows. The first term of (15) is trivially strictly increasing in S^w . The second term (15) is also increasing in S^w and in D^w , which can be seen by inspection of the derivatives, S^{26}

$$\frac{\partial}{\partial S^w} \left(S^w \left(1 - e^{-D^w/S^w} \right) \right) = \left(1 - e^{-D^w/S^w} \right) - \frac{D^w}{S^w} e^{-D^w/S^w} \ge 0$$

$$\frac{\partial}{\partial D^w} \left(S^w \left(1 - e^{-D^w/S^w} \right) \right) = e^{-D^w/S^w} > 0.$$

Thus, if (15) holds for S'^w , it cannot also hold for $S''^w > S'^w$. Intuitively, if the number of buyers and sellers is higher, then (i) more sellers exit due to discouragement (the first term) and (ii) more sellers trade (the second term) because there are simply more sellers (the first derivative is positive) and, in addition, the number of buyers increases and so less sellers have no bidders (the second derivative is positive).

By assumption, $d^H > d^L$. We show that this implies $D^H > D^L$ and $S^L < S^H$. First, it cannot be that both the number of sellers and the number of buyers increases (or decreases) when the state is changed from L to H. By our earlier observation, if both, the number of sellers and buyers increases, then the right-hand side of (15) would strictly increase, leading to a failure of the equation. Similarly, the right-hand side of (15) would strictly decrease if both market sides shrink. Because $d^H > d^L$, inspection of (16) shows that the difference $(D^H - S^H) > (D^L - S^L)$. Hence, it cannot be that the buyers' market side weakly decreases while the sellers' market side weakly increases. Therefore, $D^H > D^L$ and $S^L < S^H$, as claimed. QED.

The first inequality follows if $\frac{1}{x}(1-e^{-x}) \ge e^{-x}$. This inequality holds if $e^x - 1 \ge x$, which is true, since, at zero, $e^0 - 1 \ge 0$, and for x > 0, $(e^x - 1 - x)' = e^x - 1 \ge 0$.

8.2 Proof of Lemma 5 (Envelope Theorem)

Convexity follows from a standard argument. By definition, $V(\theta) = EU(\theta, \beta|\theta)$. With $\theta^{\alpha} = (\alpha\theta' + (1-\alpha)\theta'')$,

$$V(\alpha \theta' + (1 - \alpha) \theta'') = \alpha EU(\theta^{\alpha}, \beta | \theta') + (1 - \alpha) EU(\theta^{\alpha}, \beta | \theta'')$$

$$\leq \alpha EU(\theta', \beta | \theta') + (1 - \alpha) EU(\theta'', \beta | \theta'') = \alpha V(\theta') + (1 - \alpha) V(\theta'').$$

The equalities follow by definition of EU and V and linearity of EU. The inequality follows from optimality of β .

The envelope formula follows from standard arguments as well: (i) optimality requires $EU(\theta,\beta|\hat{\theta}) \leq EU(\hat{\theta},\beta|\hat{\theta})$ for all $\theta,\hat{\theta}$ and (ii) $EU(\theta,\beta|\hat{\theta})$ is differentiable everywhere in $\hat{\theta}$. Hence, Theorem 1 by Milgrom and Segal (2002) implies $V'(\theta) = \frac{\partial}{\partial \hat{\theta}} EU(\theta,\beta|\hat{\theta})|_{\hat{\theta}=\theta}$ at interior differentiable points of V. Linearity of $EU(\theta,\beta|\hat{\theta})$ in $\hat{\theta}$ implies that $EU(\theta,\beta|\hat{\theta}) = EU(\theta,\beta|\theta) + (\hat{\theta}-\theta)\frac{\partial EU(\theta,\beta|\hat{\theta})}{\partial \hat{\theta}}|_{\hat{\theta}=\theta}$. Together, (9) follows. QED.

8.3 Proof of Lemma 6 (Equilibrium Candidate).

The derivative of the objective function (5) is

$$\beta^{-1\prime}(\beta(x))(\gamma_{(1)}^{\theta}(v-\beta(x)-\delta V\left(\theta^{+}\left(x,\theta\right)\right))+\delta(1-\Gamma_{(1)}^{\theta}\left(x\right))V'\left(\theta^{+}\left(x,\theta\right)\right)\frac{\partial\theta^{+}\left(x,\theta\right)}{\partial x}). \tag{17}$$

The derivative exists for almost every type in the support of Γ^w . The optimal bid for almost all types is characterized by the first-order condition (17)= 0.

Note that

$$\frac{\partial \theta^{+}(x,\theta)}{\partial x} = \frac{-\theta \gamma_{(1)}^{H}(x) (1 - \Gamma_{(1)}^{\theta}(x)) + \gamma_{(1)}^{\theta}(x) \theta (1 - \Gamma_{(1)}^{H}(x))}{(1 - \Gamma_{(1)}^{\theta}(x))^{2}}.$$
 (18)

Further,

$$\frac{\gamma_{(1)}^{\theta}(x)\theta(1-\Gamma_{(1)}^{H}(x))}{1-\Gamma_{(1)}^{\theta}(x)}-\theta\gamma_{(1)}^{H}(x)=\gamma_{(1)}^{\theta}(x)(\theta^{+}(x,\theta)-\theta^{0}(x,\theta))$$
(19)

by the definitions of θ^+ and $\theta^0(x,\theta)$. Using (18) and (19), the necessary first-order condition (17)= 0 can be rewritten as (10). *QED*

8.4 Proof of Proposition 2: Price Discovery

The proposition follows from a sequence of lemmas. We start by proving **Lemma 7**. Recall Equation (16),

$$S_k^w = D_k^w - \frac{d^w - 1}{1 - \delta_k}. (20)$$

Substituting (20) into (15) yields

$$1 - e^{-D_k^w/S_k^w} = \frac{1 - (1 - \delta_k) S_k^w}{\delta_k S_k^w} = 1 + \frac{1 - S_k^w}{\delta_k S_k^w}.$$
 (21)

We can solve this equation for D_k^w to obtain

$$D_k^w = -S_k^w \ln \frac{S_k^w - 1}{\delta_k S_k^w}. (22)$$

Case 1: $d^w < 1$. In this case, $\lim \frac{d^w - 1}{1 - \delta_k} = -\infty$, so that (20) implies $\lim (S_k^w - D_k^w) = \infty$; hence, $\lim S_k^w = \infty$. This implies that the right-most side of (21) converges to 0. Therefore, the limit of the left-most side $\lim \left(1 - e^{-D_k^w/S_k^w}\right) = 0$, that is,

$$\lim_{k \to \infty} \mu_k^w = \lim_{k \to \infty} D_k^w / S_k^w = 0,$$

as claimed. From the steady-state condition for buyers, (4), $D_k^w \ge d^w$. Reordering terms and evaluating the integral on the right-hand side of (4) at zero,

$$d^{w} \ge (1 - \delta_{k}) D_{k}^{w} + \delta_{k} D_{k}^{w} \Gamma_{k,(1)}^{w}(0) = (1 - \delta_{k}) D_{k}^{w} + \delta_{k} D_{k}^{w} e^{-\mu_{k}}.$$

Taking limits on the last two inequalities implies

$$\lim_{k \to \infty} D_k^w \ge d^w \ge \lim_{k \to \infty} D_k^w e^{-\mu_k} = \lim_{k \to \infty} D_k^w.$$

Hence, $\lim D_k^w = d^w$, as claimed. Therefore, from (20) it follows that

$$\lim_{k \to \infty} (1 - \delta_k) S_k^w = 1 - d^w. \tag{23}$$

Letting $\mu_k^w = D_k^w/S_k^w$, it follows from (22) and (23) that

$$\lim_{k \to \infty} \frac{\mu_k^w}{1 - \delta_k} = \frac{d^w}{1 - d^w}.$$
 (24)

Case 2: $d^w > 1$. From (20), it follows that $\lim D_k^w - S_k^w = \infty$; thus, $\lim D_k^w = \infty$. Then, (22) implies that $\lim S_k^w = 1$.²⁷ This implies that $\mu_k^w \to \infty$. The rest of the proof

To see this, note that if $\lim S_k^w > 1$, but finite was true, then (22) would imply that $\lim D_k^w < \infty$, a contradiction with what we have already established above. If $\lim S_k^w = \infty$, then $\log \frac{S_k^w - 1}{\delta_k S_k^w} \to 0$, and therefore, by (22), $\lim \mu_k^w = \lim D_k^w / S_k^w = 0$, which contradicts $\lim D_k^w - S_k^w = \infty > 0$. Therefore, $\lim S_k^w = 1$ must hold.

establishes that $\lim_{k \to -\frac{1}{2}} \frac{1-\delta_k}{e^{-\mu_k^w}} = \infty$. Using (20) and $\lim_{k \to \infty} S_k^w = 1$ yields that

$$\lim_{k \to \infty} (1 - \delta_k) D_k^w = d^w - 1. \tag{25}$$

Formula (25) and $\lim S_k^w = 1$ imply that

$$\lim_{k \to \infty} (1 - \delta_k) \mu_k^w = d^w - 1. \tag{26}$$

Finally, using (26) implies that $\lim_{k \to -\mu_k^w} \frac{1-\delta_k}{e^{-\mu_k^w}} = \lim_{k \to -\mu_k^w} (1-\delta_k)e^{\mu_k^w} = (d^w-1)\lim_{k \to -\mu_k^w} \frac{e^{\mu_k^w}}{\mu_k^w} = \infty$. QED.

Let $q_k^w\left(\theta\right)$ denote the lifetime trading probability of a buyer having type θ ,

$$q_k^w\left(\theta\right) = \left(1 - \xi_k^w\left(\theta\right)\right) + \delta_k \xi_k^w\left(\theta\right) \left(1 - \xi_k^w\left(\theta_k^1\left(\theta\right)\right)\right) + \delta_k^2 \xi_k^w\left(\theta\right) \xi_k^w\left(\theta_k^1\left(\theta\right)\right) \left(1 - \xi_k^w\left(\theta_k^2\left(\theta\right)\right)\right) + \cdots$$

where $1 - \xi_k^w(\theta)$ denotes the probability that a buyer with type θ trades in any given period, and $\theta_k^t(\theta)$ denotes the posterior of a buyer with a prior θ who has lost t times. The steady-state conditions imply a bound on the average expected trading probability:

Lemma 8 The average expected lifetime trading probability is bounded by the ratio of the number of entering sellers to the number of entering buyers, $\int_{\underline{\theta}}^{\overline{\theta}} q_k^w(\theta) dG^w(\theta) \leq 1/d^w$. Moreover, if $1 < d^w$, then the average expected trading probability $\lim_{k \to \infty} \int_{\underline{\theta}}^{\overline{\theta}} q_k^w(\theta) dG^w(\theta) = 1/d^w$.

Proof: Let $(\theta_k^t)^{-1}(\theta')$ be the generalized inverse, $(\theta_k^t)^{-1}(\theta') = \sup\{\theta|\theta_k^t(\theta) \leq \theta'\}$. The steady-state conditions require that for every θ

$$D_{k}^{w}\Gamma_{k}^{w}(\theta) = d^{w}\left(\int_{0}^{\theta} dG^{w}(\tau) + \delta_{k} \int_{0}^{(\theta_{k}^{1})^{-1}(\theta)} \xi_{k}^{w}(\tau) dG^{w}(\tau) + \delta_{k}^{2} \int_{0}^{(\theta_{k}^{2})^{-1}(\theta)} \xi_{k}^{w}(\tau) \xi_{k}^{w}(\theta_{k}^{1}(\tau)) dG^{w}(\tau) + \cdots \right)$$
(27)

By the fundamental theorem of calculus for Lebesgue integration, we can multiply the above identity with $1 - \xi_k^w(\tau)$ point-by-point, which yields

$$\begin{split} & D_{k}^{w} \int_{0}^{\theta} \left(1 - \xi_{k}^{w}\left(\tau\right)\right) d\Gamma_{k}^{w}(\tau) \\ &= d^{w} \left(\int_{0}^{\theta} \left(1 - \xi_{k}^{w}\left(\tau\right)\right) dG^{w}\left(\tau\right) + \delta_{k} \int_{0}^{\left(\theta_{k}^{1}\right)^{-1}(\theta)} \left(1 - \xi_{k}^{w}\left(\theta_{k}^{1}\left(\tau\right)\right)\right) \xi_{k}^{w}\left(\tau\right) dG^{w}\left(\tau\right) + \dots \right) \end{split}$$

Evaluating at $\theta = 1$ and using the definition of $q_k^w(\theta)$ to simplify the right-hand side

$$D_{k}^{w} \int_{0}^{1} (1 - \xi_{k}^{w}(\tau)) d\Gamma_{k}^{w}(\tau) = d^{w} \left(\int_{0}^{1} ((1 - \xi_{k}^{w}(\tau)) + \delta_{k} \left(1 - \xi_{k}^{w} \left(\theta_{k}^{1}(\tau) \right) \right) \xi_{k}^{w}(\tau) + \delta_{k}^{2} \left(1 - \xi_{k}^{w} \left(\theta_{k}^{2}(\tau) \right) \right) \xi_{k}^{w}(\tau) \xi_{k}^{w} \left(\theta_{k}^{1}(\tau) \right) + \dots \right) dG^{w}(\tau)$$

$$= d^{w} \int_{0}^{1} q_{k}^{w}(\theta) dG^{w}(\theta).$$

As shown in the proof of Lemma 1, $D_k^w \int_0^1 (1 - \xi_k^w(\tau)) d\Gamma_k^w(\tau) = S_k^w (1 - e^{-D_k^w/S_k^w})$, the total mass of buyers who trade in any period is equal to the total mass of sellers who trade. Rewriting the steady-state condition for the sellers, (3), implies $1 = S_k^w (1 - \delta_k e^{-D_k^w/S_k^w})$. Because $1 \ge \delta_k$, $1 \ge S_k^w (1 - e^{-D_k^w/S_k^w})$. Taken together, we have shown the following chain of (in-)equalities, which proves the first claim of the lemma:

$$1 \ge S_k^w \left(1 - e^{-D_k^w / S_k^w} \right) = D_k^w \int_0^1 \left(1 - \xi_k^w \left(\tau \right) \right) d\Gamma_k^w (\tau) = d^w \int_0^1 q_k^w \left(\theta \right) dG^w \left(\theta \right). \tag{28}$$

Equation (15) implies that if $d^w > 1$, then $S_k^w \left(1 - e^{-D_k^w/S_k^w}\right) \to 1$. Taking limits on the last three equalities in (28) implies the second claim of the lemma,

$$1 = \lim_{k \to \infty} S_k^w \left(1 - e^{-D_k^w / S_k^w} \right) = \lim_{k \to \infty} d^w \int_0^1 q_k^w \left(\theta \right) dG^w \left(\theta \right)$$
 QED.

The following Lemma strengthens the finding of Lemma 8 for a special case.

Lemma 9 Suppose that $d^H > 1$ and $d^L < 1$. Then, for all $\theta \in [\underline{\theta}, \overline{\theta}]$

$$\lim_{k \to \infty} q_k^H(\theta) = 1/d^H.$$

Proof: First, the trading probability is monotone in the type: Let $\theta_l < \theta_h$ and let $\left\{\theta_l^t\right\}_{t=0}^{\infty}$ and $\left\{\theta_h^t\right\}_{t=0}^{\infty}$ be the sequence of updated beliefs after losing t times. By monotonicity of θ_k^t , $\theta_l^t < \theta_h^t$ for all t; hence, $\beta_k\left(\theta_l^t\right) < \beta\left(\theta_h^t\right)$ for all t. Therefore, the probability of winning in any given period after having lost t times, $1 - \xi_k^w\left(\theta_l^t\right) < 1 - \xi_k^w\left(\theta_h^t\right)$ for all t; hence,

$$q_{k}^{w}(\theta_{l}) = \left(1 - \xi_{k}^{w}(\theta_{l}^{0})\right) + \delta_{k}\xi_{k}^{w}(\theta_{l}^{0})\left(1 - \xi_{k}^{w}(\theta_{l}^{1})\right) + \delta_{k}^{2}\xi_{k}^{w}(\theta_{l}^{1})\xi_{k}^{w}((\theta_{l}^{0})) \dots \right.$$

$$< q_{k}^{w}(\theta_{h}) = \left(1 - \xi_{k}^{w}(\theta_{h}^{0})\right) + \delta_{k}\xi_{k}^{w}(\theta_{h}^{0})\left(1 - \xi_{k}^{w}(\theta_{h}^{1})\right) + \delta_{k}^{2}\xi_{k}^{w}(\theta_{h}^{1})\xi_{k}^{w}((\theta_{h}^{0})) \dots \right.$$

The posterior of the most optimistic new buyer after losing once becomes one, $\theta_k^1(\underline{\theta}) = \theta_k^+(\underline{\theta},\underline{\theta}) \to 1$, since the likelihood ratio of losing $\frac{1-e^{-D_k^H/S_k^H}}{1-e^{-D_k^L/S_k^L}} \to \infty$ by Lemma 7. This implies that $\lim \theta_k^+(\underline{\theta},\underline{\theta}) > \overline{\theta}$. Hence, by the monotonicity of the trading probability, we

can "sandwich" the trading probability of all $\theta \in [\underline{\theta}, \overline{\theta}]$ for sufficiently large k,

$$q_k^H\left(\theta_k^+\left(\underline{\theta},\underline{\theta}\right)\right) \ge q_k^H(\theta) \ge q_k^H(\underline{\theta}) \quad \forall \quad \theta \in [\underline{\theta},\overline{\theta}], \ k \text{ large.}$$
 (29)

Using (29) and Lemma 8,

$$\lim_{k \to \infty} \inf \int_{\underline{\theta}}^{\overline{\theta}} q_k^H \left(\theta_k^+ (\underline{\theta}, \underline{\theta}) \right) dG^H (\theta)$$

$$\geq \lim_{k \to \infty} \int_{\theta}^{\overline{\theta}} q_k^H (\theta) dG^H (\theta) = 1/d^H \geq \lim_{k \to \infty} \sup \int_{\theta}^{\overline{\theta}} q_k^H (\underline{\theta}) dG^H (\theta) .$$
(30)

By construction, $q_k^H(\underline{\theta}) \geq \delta_k q_k^H\left(\theta_k^+(\underline{\theta},\underline{\theta})\right)$. By monotonicity of θ^+ and monotonicity of q_k^H , $q_k^H(\underline{\theta}) \leq q_k^H\left(\theta_k^+(\underline{\theta},\underline{\theta})\right)$. Therefore, the difference $q_k^H\left(\theta_k^+(\underline{\theta},\underline{\theta})\right) - q_k^H(\underline{\theta}) \in [0,1-\delta_k]$. When $\delta_k \to 1$, $\lim \left(q_k^H\left(\theta_k^+(\underline{\theta},\underline{\theta})\right) - q_k^H(\underline{\theta})\right) = 0$ (the expected trading probability with the initial type $\underline{\theta}$ and the expected trading probability after updating once become the same). Hence, $\liminf q_k^H\left(\theta_k^+(\underline{\theta},\underline{\theta})\right) \leq \limsup q_k^H(\underline{\theta})$. This inequality together with the inequalities (30) implies

$$\lim_{k \to \infty} q_k^H \left(\theta_k^+ \left(\underline{\theta}, \underline{\theta} \right) \right) = 1/d^H = \lim_{k \to \infty} q_k^H \left(\underline{\theta} \right);$$

(recall, $\int_{\underline{\theta}}^{\overline{\theta}} dG^H(\theta) = 1$). Hence, (29) implies $\lim q_k^H(\theta) = 1/d^H$ for all $\theta \in [\underline{\theta}, \overline{\theta}]$. QED. We prove that trading probabilities satisfy the conditions of Proposition 2.

Lemma 10 Trading probabilities satisfy:

$$\lim_{k\to\infty}q_k^w\left(\theta\right)=\left\{\begin{array}{ll} 1 & \text{if } d^w<1\\ \frac{1}{d^w} & \text{if } d^w>1 \end{array}\right. \qquad and \qquad \lim_{k\to\infty}q_k^w\left(S\right)=\left\{\begin{array}{ll} d^w & \text{if } d^w<1\\ 1 & \text{if } d^w>1. \end{array}\right.$$

Proof: For buyers: If $d^w < 1$, then $\lim_{k \to \infty} q_k^w(\theta) = 1$ is immediate from Lemma 7. We have argued that $\lim_{k \to \infty} q_k^w(\theta) = \frac{1}{d^w}$ if $d^w > 1$ for the case w = H and $d^L < 1$ in Lemma 9. The case in which in both states $d^w > 1$ follows from the steady-state conditions along similar lines. We omit the proof of that case.

For sellers: The trading probability is recursively defined as $q_k^w(S) = 1 - e^{-D_k^w/S_k^w} + \delta_k e^{-D_k^w/S_k^w} q_k^w(S)$. If $d^w > 1$, then $\lim_{k \to \infty} q_k^w(S) = 1$ follows from $D_k^w/S_k^w \to \infty$, shown in Lemma 7.

If $d^w < 1$, then (23) implies $(1 - \delta_k)S_k^w = 1 - d^w$. From the steady-state condition, (15),

$$1 = (1 - \delta_k) S_k^w + \delta_k S_k^w \left(1 - e^{-D_k^w / S_k^w} \right).$$

Rewriting the definition of $q_k^w(S)$, $1 - e^{-D_k^w/S_k^w} = \frac{(1-\delta_k)q_k^w(S)}{1-\delta_kq_k^w(S)}$, substituting into the

steady-state condition, and taking limits,

$$1 = \lim (1 - \delta_k) S_k^w + \lim \delta_k S_k^w \frac{(1 - \delta_k) q_k^w (S)}{1 - \delta_k q_k^w (S)} = 1 - d^w + (1 - d^w) \lim \frac{q_k^w (S)}{1 - q_k^w (S)},$$

from which $\lim q_k^w\left(S\right) = d^w$ follows, as claimed. *QED*.

Let $q_k^w(\theta, b')$ denote the probability that a type θ eventually ends up trading at a price $p \leq b'$, and let θ_k denote the highest type in the stock who bids below b', $\theta_k = \sup\{\theta | \beta_k \leq b', \theta \in \operatorname{supp}\Gamma_k^w\}$ (if there is no such type, $\theta_k = 0$). The probability $q_k^w(\theta, b')$ is defined as

$$q_{k}^{w}\left(\theta,b'\right) = \left(1 - \xi_{k}^{w}\left(\min\left\{\theta_{k},\theta\right\}\right)\right) + \delta_{k}\left(1 - \xi_{k}^{w}\left(\theta\right)\right)\xi_{k}^{w}\left(\min\left\{\theta_{k},\theta_{k}^{1}\left(\theta\right)\right\}\right) + \delta_{k}^{2}\left(1 - \xi_{k}^{w}\left(\theta\right)\right)\left(1 - \xi_{k}^{w}\left(\theta_{k}^{1}\left(\theta\right)\right)\right)\xi_{k}^{w}\left(\min\left\{\theta_{k},\theta_{k}^{2}\left(\theta\right)\right\}\right) + \dots,$$

where we need to use $(\min\{\theta_k, \theta\})$ because $\beta_k(\theta)$ might be below b'; that is, $\theta < \theta_k$.

Let $\overline{V} = \lim V_k(1)$. We show that the expected price conditional on winning becomes equal to $v - \overline{V}$. To show this, we want to prove that if a bidder wins with some bid with positive probability, the second highest bidder bids almost surely $\beta = v - \overline{V}$ (no bidder would bid higher). The proof works as follows: If there is a positive chance to win against a buyer with a belief θ' , then the posterior of this buyer conditional on being tied must converge to one; formally,

$$\forall \left\{ \theta_{k}^{\prime} \right\} : \qquad \lim q_{k}^{H} \left(\theta, \beta_{k} \left(\theta_{k}^{\prime} \right) \right) > 0 \Rightarrow \lim \theta_{k}^{0} \left(\theta_{k}^{\prime}, \theta_{k}^{\prime} \right) = 1.$$

Because $\theta_k^0(\theta_k', \theta_k') \to 1$, the bid $\beta_k(\theta_k')$ is shown to converge to $\lim V_k(1) = \bar{V}$.

Lemma 11 Suppose that $d^H > 1$ and $d^L < 1$. Let $\overline{V} = \lim_{k \to \infty} V_k(1)$. For any type $\theta \in [\underline{\theta}, \overline{\theta}]$ the expected price conditional on winning in the high state converges to $v - \overline{V}$, $\lim_{k \to \infty} E_k[p, \theta|H] = v - \overline{V}$.

Proof: Suppose that there are some b' and $\theta^* \in [\underline{\theta}, \overline{\theta}]$ such that $q_k^H(\theta^*, b')$ converges to some positive number along some subsequence,

$$d^{H} \lim_{k \to \infty} \inf q_{k}^{H} \left(\theta^{*}, b' \right) = \varepsilon > 0. \tag{31}$$

We prove that this implies $b' \geq v - \bar{V}$. This implies the claim.

Let $\theta_k \equiv \sup \{\theta | \beta_k(\theta) \leq b'\}$. We prove that (31) implies

$$\lim_{k \to \infty} \inf(1 - \xi_k^H(\theta_k)) D_k^H \ge \varepsilon. \tag{32}$$

We are done if $\liminf \xi_k^H(\theta_k) < 1$. So, suppose that $\xi_k^H(\theta_k) \to 1$. The inequality now

follows from the following chain of equations. For k sufficiently large,

$$\varepsilon \leq d^{H}(1 - \xi_{k}^{H}(\theta_{k})) + \delta_{k}d^{H}(1 - \xi_{k}^{H}(\theta_{k})) \xi_{k}^{H}(\theta_{k}^{1}(\theta)) + \delta_{k}^{2}(1 - \xi_{k}^{H}(\theta_{k})) \dots \\
= d^{H}(1 - \xi_{k}^{H}(\theta_{k})) + \delta_{k}d^{H}(1 - \xi_{k}^{H}(\theta_{k})) \left[\xi_{k}^{H}(\theta_{k}^{1}(\theta)) + \delta_{k}\xi_{k}^{H}(\theta_{k}^{1}(\theta)) \xi_{k}^{H}(\theta_{k}^{2}(\theta)) + \dots\right] \\
\leq d^{H}(1 - \xi_{k}^{H}(\theta_{k})) + \delta_{k}d^{H}(1 - \xi_{k}^{H}(\theta_{k})) \left[\xi_{k}^{H}(\bar{\theta}) + \delta_{k}\xi_{k}^{H}(\theta_{k}^{1}(\bar{\theta})) \xi_{k}^{H}(\bar{\theta}) + \dots\right],$$

where the first inequality comes from the definition of q_k^H and the second inequality comes from $\theta_k^1(\theta) \to 1 > \overline{\theta}$ and ξ_k^H being nonincreasing. Integrating both sides with respect to G^H , taking limits with $k \to \infty$, and noting that $\xi_k^H(\theta_k) \to 1$, we rewrite further

$$\liminf \left(1 - \xi_k^H(\theta_k)\right) d^H \int_0^1 \left[\xi_k^H(\bar{\theta}) + \delta_k \xi_k^H(\theta_k^1(\bar{\theta})) \xi_k^H(\bar{\theta}) + \dots\right] dG^H$$

$$\leq \liminf \left(1 - \xi_k^H(\theta_k)\right) d^H \int_0^1 \left[1 + \delta_k \xi_k^H(\tau) + \dots\right] dG^H(\tau)$$

$$= \lim \left(1 - \xi_k^H(\theta_k)\right) D_k^H.$$

where we used that ξ_k^H is nonincreasing, $\xi_k^H(\bar{\theta}) \to 1$, and the steady-state conditions. Together, the two displayed chains of equations imply the desired inequality (32).

We expand (32) using the definition of ξ_k^w ,

$$\liminf_{k \to \infty} D_k^H e^{-D_k^H (1 - \Gamma_k^H(\theta_k))/S_k^H} \ge \varepsilon. \tag{33}$$

Equation (33) implies that $\lim_{k\to\infty}\Gamma_k^H(\theta_k)=1$: Otherwise, if $\limsup \sup \Gamma_k^H(\theta_k)<1$ were true, $S_k^H\to 1$ and $D_k^H\to \infty$ would imply that $\liminf D_k^H e^{-D_k^H(1-\Gamma_k^H(\theta_k))/S_k^H}=0$ by l'Hospital's rule, ²⁸ contradicting (33). Using $D_k^L\to d^L<1$, we obtain that

$$\limsup D_k^L \Gamma_k^L(\theta_k) e^{-D_k^L(1-\Gamma_k^L(\theta_k)/S_k^L)} \le 1.$$

Hence,

$$\liminf_{k \to \infty} \frac{D_k^H \Gamma_k^H \left(\theta_k\right) e^{-D_k^H \left(1 - \Gamma_k^H \left(\theta_k\right) / S_k^H}\right)}{D_k^L \Gamma_k^L \left(\theta_k\right) e^{-D_k^L \left(1 - \Gamma_k^L \left(\theta_k\right) / S_k^L}\right)} \ge \varepsilon.$$
(34)

The likelihood ratio of tying satisfies

$$\frac{\theta_{k}^{0}(\theta_{k})}{1 - \theta_{k}^{0}(\theta_{k})} = \frac{\theta_{k}}{1 - \theta_{k}} \frac{\gamma_{k}^{H}(\theta_{k})}{\gamma_{k}^{L}(\theta_{k})} \frac{D_{k}^{H}}{D_{k}^{L}} \frac{S_{k}^{L}}{S_{k}^{H}} \frac{e^{-D_{k}^{H}(1 - \Gamma_{k}^{H}(\theta_{k})/S_{k}^{H}}}{e^{-D_{k}^{L}(1 - \Gamma_{k}^{L}(\theta_{k}))/S_{k}^{L}}} \\
\geq \frac{\theta_{k}}{1 - \theta_{k}} \frac{\Gamma_{k}^{H}(\theta_{k})}{\Gamma_{k}^{L}(\theta_{k})} \frac{D_{k}^{H}}{D_{k}^{L}} \frac{S_{k}^{L}}{S_{k}^{H}} \frac{e^{-D_{k}^{H}(1 - \Gamma_{k}^{H}(\theta_{k}))/S_{k}^{H}}}{e^{-D_{k}^{L}(1 - \Gamma_{k}^{L}(\theta_{k}))/S_{k}^{L}}}.$$

The large state $x_k = D_k^H$ and $c_k = (1 - \Gamma_k^H(\theta_k))/S_k^H$. Using l'Hospital's rule, $\lim \inf x_k e^{-x_k c_k} = 0$ if $x_k \to \infty$ and $\lim \inf c_k > 0$. Hence, $c_k = (1 - \Gamma_k^H(\theta_k))/S_k^H$ must converge to 0. From before, $S_k^H \to 1$, so this requires $\Gamma_k^H(\theta_k) \to 1$.

The inequality follows from the MLRP of Γ^w , $\frac{\gamma_k^H(\theta_k)}{\gamma_k^L(\theta_k)} \ge \frac{\Gamma_k^H(\theta_k)}{\Gamma_k^L(\theta_k)}$. Using (34) (for the first inequality) and using $\frac{S_k^L}{S_k^H} \to \infty$ (for the second equality),

$$\lim_{k \to \infty} \inf \frac{\theta_k^0(\theta_k)}{1 - \theta_k^0(\theta_k)} \ge \varepsilon \lim_{k \to \infty} \inf \frac{\theta_k}{1 - \theta_k} \frac{S_k^L}{S_k^H} = \infty.$$
 (35)

Therefore, the posterior $\theta_k^0(\theta_k, \theta_k) \to 1$.

We show that $\theta_k^0(\theta_k, \theta_k) \to 1$ implies that $\beta_k(\theta_k) \to v - \bar{V}$; that is, $b' = v - \bar{V}$, as claimed in the beginning. From Lemma 15, $\beta_k(\theta_k) = v - \delta_k EU\left(\theta_k^+(\theta_k, \theta_k), \beta_k | \theta_k^0(\theta_k, \theta_k)\right)$. By $\theta_k^+(\theta_k, \theta_k) \to 1$ and because the sequence of payoffs $EU\left(\theta, \beta_k | \theta'\right)$ is Lipschitz continuous in θ' with a uniform Lipschitz constant, we can pass the limit through; that is $\lim EU\left(\theta_k^+(\theta_k, \theta_k), \beta_k | H\right) = \lim EU\left(1, \beta_k | H\right) = \bar{V}$. Hence,

$$\lim \beta_{k}(\theta_{k})$$

$$= v - \lim \delta_{k} EU\left(\theta_{k}^{+}(\theta_{k}, \theta_{k}), \beta_{k} | \theta_{k}^{0}(\theta_{k}, \theta_{k})\right)$$

$$= v - \lim \underbrace{\theta_{k}^{0}(\theta_{k}, \theta_{k})}_{\rightarrow 1} \underbrace{EU\left(\theta_{k}^{+}(\theta_{k}, \theta_{k}), \beta_{k} | H\right)}_{\rightarrow \bar{V}} - \underbrace{\left(1 - \theta_{k}^{0}(\theta_{k}, \theta_{k})\right)}_{\rightarrow 0} EU\left(\theta_{k}^{+}(\theta_{k}, \theta_{k}), \beta_{k} | L\right)$$

$$= v - \bar{V} = b'.$$

Thus, $\lim q_k^H(\theta, p) = 0$ for all $p < v - \bar{V}$, $\theta \in [\underline{\theta}, \overline{\theta}]$: If not, then $\liminf q_k^H(\theta', p') > 0$ for some $p' < v - \bar{V}$ and θ' . As we have shown before, $\liminf q_k^H(\theta', p') > 0$ implies $p' = v - \bar{V}$, a contradiction.

From $\beta_k(\theta) \leq \beta_k(1) = v - \delta V_k(1)$ for all θ and k, it follows that $\lim q_k^H(\theta, p) = 1$ for all $p > v - \bar{V}$, $\theta \in [\underline{\theta}, \overline{\theta}]$. Hence,

$$\lim_{k \to \infty} E_k[p, \theta | H] = \lim_{k \to \infty} \frac{1}{q_k^w(\theta)} \int_0^1 (1 - q_k^w(\theta, b)) db = v - \bar{V}. \ QED.$$

Proof of Proposition 2.

We have shown that the trading probabilities become competitive in Lemma 10. We now show that prices become competitive, too.

Case 1: $d^L < d^H < 1$. By Lemma 1, in both states, the probability of being the sole bidder becomes one, $e^{-D_k^w/S_k^w} \to 1$, which implies $V_k(\theta) \to v$ for all θ by inspection of the payoffs. In particular, $V_k(1) \to v$. Hence, the bidding strategy $\beta_k(\theta) \to 0$ for all θ . Because the expected price is smaller than $\beta_k(1)$ by definition and the monotonicity of β_k and because $\beta_k(1) \to 0$, $\lim_{k \to \infty} p_k^w(\theta) = \lim_{k \to \infty} p_k^w(S) = 0$ follows.

Case 2: $d^L > 1$ and $d^H > 1$. We show $EU_k [\beta_k, \theta = 0|L] \to 0$. By monotonicity of the bidding strategy, $\beta_k (0) \le \beta_k (\theta)$ for all θ ; hence, a buyer with belief $\theta = 0$ wins only

as the sole bidder,

$$\begin{split} EU_k\left[\beta_k,\theta=0|L\right] &= e^{-D_k^L/S_k^L}\left(v\right) + \delta_k\left(1 - e^{-D_k^L/S_k^L}\right)EU_k\left[\beta_k,\theta=0|L\right] \\ \Leftrightarrow EU_k\left[\beta_k,\theta=0|L\right] &= \frac{e^{-D_k^L/S_k^L}}{\frac{1-\delta_k}{e^{-D_k^L/S_k^L}} + \delta_k}v. \end{split}$$

From Lemma 1, $\frac{1-\delta_k}{e^{-D_k^L/S_k^L}} \to \infty$, while $e^{-D_k^L/S_k^L} \to 0$. Therefore,

$$\lim_{k \to \infty} EU_k \left[\beta_k, \theta = 0 | L \right] = 0.$$

Since $EU_k\left[\beta_k, \theta = 0 \middle| L\right] = V_k\left(0\right)$ and $V_k\left(0\right) \to 0$, we have $\lim \beta_k\left(\theta\right) = v$ for all θ . Because a bidder is never a sole bidder in the limit, $\lim_{k \to \infty} p_k^w\left(\theta\right) = \lim_{k \to \infty} p_k^w\left(S\right) = v$ follows.

Case 3: $d^H > 1$ and $d^L < 1$.

As in Case 1, by Lemma 1, $\lim_{k\to\infty} p_k^L(\theta) = \lim_{k\to\infty} p_k^L(S) = 0$. We now argue w = H. From before, $\theta_k^1(\underline{\theta}) \to 1$. Since, in the high state, the first cohort has a vanishing winning probability, and since there is no exogenous exit in the limit either, it follows that

$$\lim_{k \to \infty} V_k^H(\underline{\theta}) = \lim_{k \to \infty} V_k^H(\theta_k^1(\underline{\theta})) = \bar{V}.$$

From Lemma 9, $q_k^H\left(\underline{\theta}\right) \to 1/d^H$. From Lemma 11 $p_k^H\left(\underline{\theta}\right) = v - \bar{V}$. Together,

$$\lim_{k \to \infty} V_k^H(\underline{\theta}) = \frac{v - (v - \overline{V})}{d^H} = \frac{\overline{V}}{d^H}.$$

Since $d^{H} > 1$, $\bar{V} = \frac{\bar{V}}{d^{H}}$ implies that $\bar{V} = 0$. Thus, the expected price $\lim_{k \to \infty} p_{k}^{H}(\theta) = \lim_{k \to \infty} p_{k}^{H}(S) = v - \bar{V} = v$. QED.

8.5 Proof of Proposition 3: Time Pattern of Bids

Proof. Step 1: We derive some auxiliary observations. By definition, $\theta_k^t(\theta)$ satisfies

$$\frac{\theta_{k}^{t}\left(\theta\right)}{1-\theta_{k}^{t}\left(\theta\right)} = \frac{\theta_{k}^{t-1}\left(\theta\right)}{1-\theta_{k}^{t-1}\left(\theta\right)} \frac{1-\Gamma_{(1),k}^{H}(\theta_{k}^{t-1}\left(\theta\right))}{1-\Gamma_{(1),k}^{L}(\theta_{k}^{t-1}\left(\theta\right))}.$$

The belief $\theta_k^t(\theta)$ is strictly increasing in t. From the proof of Proposition 2, $\lim_{k\to\infty} \Gamma^H_{(1),k}(\theta) = \lim_{k\to\infty} 1 - \Gamma^L_{(1),k}(\theta) = 0$ for all $\theta \in [\underline{\theta}, \overline{\theta}]$. Therefore, $\lim_{k\to\infty} \theta_k^1(\theta) = 1$; thus, by monotonicity, $\lim_{k\to\infty} \theta_k^t(\theta) = 1$. The posterior after losing t times and then being tied satisfies (if

densities are positive)

$$\frac{\theta_{k}^{0}\left(\theta_{k}^{t}\left(\theta\right),\theta_{k}^{t}\left(\theta\right)\right)}{1-\theta_{k}^{0}\left(\theta_{k}^{t}\left(\theta\right),\theta_{k}^{t}\left(\theta\right)\right)} = \frac{\theta_{k}^{t}\left(\theta\right)}{1-\theta_{k}^{t}\left(\theta\right)} \frac{\gamma_{(1),k}^{H}(\theta_{k}^{t}\left(\theta\right))}{\gamma_{(1),k}^{L}(\theta_{k}^{t}\left(\theta\right))}.$$
(36)

By definition, $\gamma_{(1),k}^w = \mu_k^w \gamma_k^w \Gamma_{(1),k}^w$. Using the no-introspection condition to substitute for γ_k^w ,

$$\frac{\gamma_{(1),k}^{H}(\theta)}{\gamma_{(1),k}^{L}(\theta)} = \frac{\theta}{1-\theta} \frac{S_k^L}{S_k^H} \frac{\Gamma_{(1),k}^{H}(\theta)}{\Gamma_{(1),k}^{L}(\theta)}.$$
 (37)

Substituting iteratively into (36),

$$\frac{\theta_{k}^{0}\left(\theta_{k}^{t}\left(\theta\right),\theta_{k}^{t}\left(\theta\right)\right)}{1-\theta_{k}^{0}\left(\theta_{k}^{t}\left(\theta\right),\theta_{k}^{t}\left(\theta\right)\right)} = \frac{\theta_{k}^{t-1}}{1-\theta_{k}^{t-1}}\frac{1-\Gamma_{(1),k}^{H}(\theta_{k}^{t-1})}{1-\Gamma_{(1),k}^{L}(\theta_{k}^{t-1})}\frac{\theta_{k}^{t}}{1-\theta_{k}^{t}}\frac{S_{k}^{L}}{S_{k}^{H}}\frac{\Gamma_{(1),k}^{H}(\theta_{k}^{t})}{\Gamma_{(1),k}^{L}(\theta_{k}^{t})} = \\ = \left(\frac{\theta_{k}^{t-1}}{1-\theta_{k}^{t-1}}\frac{1-\Gamma_{(1),k}^{H}(\theta_{k}^{t-1})}{1-\Gamma_{(1),k}^{L}(\theta_{k}^{t-1})}\right)^{2}\frac{S_{k}^{L}}{S_{k}^{H}}\frac{\Gamma_{(1),k}^{H}(\theta_{k}^{t})}{\Gamma_{(1),k}^{L}(\theta_{k}^{t})} = \\ = \left(\frac{\theta_{k}^{t-2}}{1-\theta_{k}^{t-2}}\frac{1-\Gamma_{(1),k}^{H}(\theta_{k}^{t-2})}{1-\Gamma_{(1),k}^{L}(\theta_{k}^{t-2})}\frac{1-\Gamma_{(1),k}^{H}(\theta_{k}^{t-1})}{1-\Gamma_{(1),k}^{L}(\theta_{k}^{t-1})}\right)^{2}\frac{S_{k}^{L}}{S_{k}^{H}}\frac{\Gamma_{(1),k}^{H}(\theta_{k}^{t})}{\Gamma_{(1),k}^{L}(\theta_{k}^{t})} = \\ = \left(\frac{\theta}{1-\theta}\frac{1-\Gamma_{(1),k}^{H}(\theta)}{1-\Gamma_{(1),k}^{L}(\theta)}\cdots\frac{1-\Gamma_{(1),k}^{H}(\theta_{k}^{t-1})}{1-\Gamma_{(1),k}^{L}(\theta_{k}^{t-1})}\right)^{2}\frac{S_{k}^{L}}{S_{k}^{H}}\frac{\Gamma_{(1),k}^{H}(\theta_{k}^{t})}{\Gamma_{(1),k}^{L}(\theta_{k}^{t})} (38)$$

Recall from the proof of Proposition 2 that

$$\lim_{k \to \infty} \mu_k^L = 0 \text{ and } \lim_{k \to \infty} \frac{\mu_k^L}{1 - \delta_k} = \frac{d^L}{1 - d^L}.$$
 (39)

L'Hospital's rule implies that $\lim_{x\to 0} \frac{1-e^{-x}}{x} = 1$. Hence,

$$\lim_{k \to \infty} \frac{1 - e^{-\mu_k^L}}{1 - \delta_k} = \frac{d^L}{1 - d^L}.$$
 (40)

Step 2: Proof of Statement (i) of Proposition 3.

From the previous step, $\lim \theta_k^1(\theta) > \overline{\theta}$ for all $\theta \in [\underline{\theta}, \overline{\theta}]$. Thus, monotonicity of beliefs implies that for sufficiently large k,

$$\theta_k^0\left(\theta_k^t(\underline{\theta}), \theta_k^t(\underline{\theta})\right) \ge \theta_k^0\left(\theta_k^{t-1}(\theta), \theta_k^{t-1}(\theta)\right)$$

for all $\theta \in [\underline{\theta}, \overline{\theta}]^{29}$ Therefore, it is sufficient to prove statement (i) at $\theta = \underline{\theta}$. In the

The tying posterior θ_k^0 is defined at $\underline{\theta}$ by definition of γ_k^w and g_k^w ($\underline{\theta}$) > 0, which ensure $\gamma_{(1)}^w$ (θ_k^t ($\underline{\theta}$)) > 0 for all t.

following, we simplify

$$\theta_k^t \equiv \theta_k^t \left(\underline{\theta}\right)$$
.

For sufficiently large k such that $\theta_k^1 > \overline{\theta}$, the steady-state conditions imply

$$D_{k}^{w}\Gamma_{k}^{w}(\theta_{k}^{t})=d^{w}+\int_{\theta}^{\bar{\theta}}\delta_{k}\xi_{k}^{w}\left(\tau\right)dG^{w}\left(\tau\right)+...+\int_{\theta}^{\bar{\theta}}\delta_{k}^{t-1}\xi_{k}^{w}\left(\tau\right)...\xi_{k}^{w}\left(\theta_{k}^{t-1}\left(\tau\right)\right)dG^{w}\left(\tau\right)\leq td^{w}.$$

From the proof of Proposition 2, $D_k^H \to \infty$. Hence,

$$\lim_{k \to \infty} \frac{\Gamma_{(1),k}^{H}(\theta_k^t)}{e^{-\mu_k^H}} = \lim_{k \to \infty} \frac{e^{-\mu_k^H \left(1 - \frac{ta^H}{D_k^H}\right)}}{e^{-\mu_k^H}} = 1.$$
(41)

Moreover, the steady-state conditions imply that

$$D_k^w \left(1 - \Gamma_k^w(\theta_k^t) \right) \le \frac{d^L \left(1 - e^{-\mu_k^L} \right)^t}{1 - \delta_L + e^{-\mu_k^L}} \tag{42}$$

because $\xi_k^w\left(\theta\right) \leq \xi_k^w\left(\underline{\theta}\right) = 1 - e^{-\mu_k^L}$ for all $\theta \geq \underline{\theta}$ and

$$\begin{split} D_k^w \left(1 - \Gamma_k^w(\theta_k^t) \right) &= \int_{\underline{\theta}}^{\overline{\theta}} \delta_k^t \xi_k^w \left(\tau \right) ... \xi_k^w \left(\theta_k^t \left(\tau \right) \right) dG^w \left(\tau \right) \\ &+ \int_{\underline{\theta}}^{\overline{\theta}} \delta_k^{t+1} \xi_k^w \left(\tau \right) ... \xi_k^w \left(\theta_k^{t+1} \left(\tau \right) \right) dG^w \left(\tau \right) + \\ &\leq \delta_k^t \left(1 - e^{-\mu_k^L} \right)^t d^L + \delta_k^{t+1} \left(1 - e^{-\mu_k^L} \right)^{t+1} d^L + \\ &= \frac{1}{1 - \delta_k + e^{-\mu_k^L}} \delta_k^t \left(1 - e^{-\mu_k^L} \right)^t d^L. \end{split}$$

From (40) and (41),

$$\lim_{k\to\infty}\frac{1-\Gamma_{(1),k}^L(\underline{\theta})}{1-\delta_k}=\lim_{k\to\infty}\frac{1-e^{-\mu_k^L}}{1-\delta_k}=\frac{d^L}{1-d^L}.$$

Hence, (42) implies

$$\lim_{k \to \infty} \sup \frac{D_k^L \left(1 - \Gamma_k^L(\boldsymbol{\theta}_k^t)\right)}{\left(1 - \delta_k\right)^t} \le d^L \left(\frac{d^L}{1 - d^L}\right)^t.$$

This implies

$$\lim_{k \to \infty} \sup \frac{1 - \Gamma_{(1),k}^{L}(\theta_k^t)}{(1 - \delta_k)^{t+1}} = \lim_{k \to \infty} \sup \frac{1 - e^{-\mu_k^L (1 - \Gamma_k^L(\theta_k^t))}}{(1 - \delta_k)^{t+1}} \le \left(\frac{d^L}{1 - d^L}\right)^{t+1}, \tag{43}$$

since,

$$\lim_{k \to \infty} \sup \frac{\mu_k^L \left(1 - \Gamma_k^L(\theta_k^t)\right)}{\left(1 - \delta_k\right)^{t+1}} \leq \lim_{k \to \infty} \sup \frac{1}{\left(1 - \delta_k\right) S_k^L} \left(\frac{D_k^L \left(1 - \Gamma_k^L(\theta_k^t)\right)}{\left(1 - \delta_k\right)^t}\right) = \frac{d^L}{1 - d^L} \left(\frac{d^L}{1 - d^L}\right)^t.$$

Note that (23), (26), and (41) imply

$$\lim S_k^L \Gamma_{(1),k}^H(\theta_k^t) = \lim \frac{1 - d^L}{1 - \delta_k} e^{-\frac{1 - d^H}{1 - \delta_k}}.$$
(44)

Taking limits on (38) and ignoring all terms that are finite and non-zero,

$$\lim_{k \to \infty} \frac{\theta_k^0 \left(\theta_k^t \left(\bar{\theta}\right), \theta_k^t \left(\bar{\theta}\right)\right)}{1 - \theta_k^0 \left(\theta_k^t \left(\bar{\theta}\right), \theta_k^t \left(\bar{\theta}\right)\right)}$$

$$= \lim_{k \to \infty} \left(\frac{1}{1 - \Gamma_{(1),k}^L(\underline{\theta})} ... \frac{1}{1 - \Gamma_{(1),k}^L(\theta_k^{t-2})} \frac{1}{1 - \Gamma_{(1),k}^L(\theta_k^{t-1})}\right)^2 S_k^L \Gamma_{(1),k}^H(\theta_k^t)$$

$$= \lim_{k \to \infty} \frac{\frac{1}{1 - \delta_k} e^{-\left(1 - d^H\right)/(1 - \delta_k)}}{(1 - \delta_k)(1 - \delta_k)^2 ... (1 - \delta_k)^t} = 0,$$

where we used (44) and (43) for the second and l'Hospital's rule for the final equality.

From the proof of Proposition 2, $\lim EU(\theta, \beta_k|0) = v$ for all θ . Because the sequence of payoffs $EU(\theta, \beta_k|\theta')$ is Lipschitz continuous in θ' with a uniform Lipschitz constant, we can pass the limit through:

$$\begin{split} \lim_{k \to \infty} \beta_k(\theta_k^t) &= v - \lim_{k \to \infty} \delta_k EU\left(\theta_k^+ \left(\theta_k^t, \theta_k^t\right), \beta_k | \theta_k^0 \left(\theta_k^t, \theta_k^t\right)\right) \\ &= v - \lim_{k \to \infty} EU\left(1, \beta_k | 0\right) = 0. \end{split}$$

This concludes the proof of statement (i).

Step 3: Proof of Statement (ii) of Proposition 3.

Let $t_k \equiv \frac{-(0.5)}{(1-\delta)\ln(1-\delta)}$. Note that Lemma 3 implies that the likelihood ratio of losing, $\frac{1-\Gamma_{(1),k}^H(\theta)}{1-\Gamma_{(1),k}^L(\theta)}$, is a nondecreasing function of θ on the support of $\Gamma_{(1),k}^w$. Therefore, (38) implies for $\theta \in [\underline{\theta}, \overline{\theta}]$:

$$\frac{\theta_k^0\left(\theta_k^t\left(\theta\right),\theta_k^t\left(\theta\right)\right)}{1-\theta_k^0\left(\theta_k^t\left(\theta\right),\theta_k^t\left(\theta\right)\right)} \ge \left(\frac{\theta}{1-\theta}\right)^2 \left(\frac{1-\Gamma_{(1),k}^H(\theta)}{1-\Gamma_{(1),k}^L(\theta)}\right)^{2t_k} \frac{S_k^L}{S_k^H} \frac{\Gamma_{(1),k}^H(\theta_k^t)}{\Gamma_{(1),k}^L(\theta_k^t)}.$$

Evaluating the limit of terms on the right-hand side:

$$\lim_{k \to \infty} \inf \left(\frac{1 - \Gamma_{(1),k}^H(\theta)}{1 - \Gamma_{(1),k}^L(\theta)} \right)^{2t_k} S_k^L \Gamma_{(1),k}^H(\theta_k^t)$$

$$= \lim_{k \to \infty} \inf \left(\frac{1}{1 - \Gamma_{(1),k}^L(\theta)} \right)^{2t_k} \frac{1 - d^L}{1 - \delta_k} e^{-\frac{1 - d^H}{1 - \delta_k}}$$

$$= \lim_{k \to \infty} \inf \left(\frac{d^L}{1 - d^L} (1 - \delta_k) \right)^{-2t_k} \frac{1 - d^L}{1 - \delta_k} e^{-\frac{1 - d^H}{1 - \delta_k}} = \infty,$$

where we used $\Gamma_{(1),k}^H(\theta) \to 0$ and (44) for the second line, (40) for the first equality of the third line, and the following observation for the second equality of the third line:

$$\lim_{(1-\delta)\to 0} \frac{\frac{1-d^L}{1-\delta}e^{-\frac{1-d^H}{1-\delta}}}{\left(\frac{d_L}{1-d_L}\left(1-\delta\right)\right)^{-\frac{1}{(1-\delta)\ln(1-\delta)}}} = \lim_{(1-\delta)\to 0} \frac{1-d^L}{1-\delta}\left(\left(\frac{d_L}{1-d_L}\right)^{\frac{1}{\ln(1-\delta)}}\frac{e^{1-d_H}}{e}\right)^{\frac{-1}{1-\delta}} = \infty.$$

This follows from $\lim \left(\frac{d_L}{1-d_L}\right)^{\frac{1}{\ln(1-\delta)}} = 1$ and $\lim \left(\frac{d_L}{1-d_L}\right)^{\frac{1}{\ln(1-\delta)}} \frac{e^{1-d_H}}{e} = e^{-d_H} < 1$, so that $\lim_{(1-\delta)\to 0} \left(\left(\frac{d_L}{1-d_L}\right)^{\frac{1}{\ln(1-\delta)}} \frac{e^{1-d_H}}{e}\right)^{\frac{-1}{1-\delta}} = \infty$.

Thus, $\lim_{k\to\infty}\theta_k^0\left(\theta_k^t\left(\theta\right),\theta_k^t\left(\theta\right)\right)=1$, as claimed. From the proof of Proposition 2, $\lim EU\left(1,\beta_k|1\right)=0$. Together,

$$\begin{split} \lim_{k \to \infty} & \beta_k(\theta_k^{t_k}) &= v - \lim_{k \to \infty} \delta_k EU\left(\theta_k^+ \left(\theta_k^{t_k}, \theta_k^{t_k}\right), \beta_k | \theta_k^0 \left(\theta_k^{t_k}, \theta_k^{t_k}\right)\right) \\ &= v - \lim_{k \to \infty} \delta_k EU\left(1, \beta_k | 1\right) = v. \end{split} \qquad QED.$$

References

Bajari, Patrick and Ali Hortacsu, "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions," *RAND Journal of Economics*, 2003, 34 (2), 329–55.

Banerjee, Abhijit and Drew Fudenberg, "Word-of-mouth learning," Games and Economic Behavior, 2004, 46 (1), 1–22.

Blouin, Max R and Roberto Serrano, "A Decentralized Market with Common Values Uncertainty: Non-Steady States," *Review of Economic Studies*, 2001, 68 (2), 323–46.

Diamond, Peter A., "A Model of Price Adjustment," *Journal of Economic Theory*, 1971, 3 (2), 156–168.

- **Duffie, Darrell and Gustavo Manso**, "Information Percolation in Large Markets," *American Economic Review*, 2007, 97 (2), 203–209.
- _ and Yeneng. Sun, "Existence of independent random matching," The Annals of Applied Probability, 2007, 17 (1), 386-419.
- Gale, Douglas, "Limit Theorems for Markets with Sequential Bargaining," Journal of Economic Theory, 1987, 43 (1), 20–54.
- _ , Strategic Foundations of General Equilibrium, Cambridge University Press, 2000.
- Golosov, Michael, Guido Lorenzoni, and Aleh Tsyvinski, "Decentralized Trading with Private Information," February 2011. Mimeo.
- Golub, Benjamin and Matthew O. Jackson, "Naïve Learning in Social Networks and the Wisdom of Crowds," American Economic Journal: Microeconomics, 2010, 2 (1), 112–49.
- Gottardi, Piero and Roberto Serrano, "Market Power And Information Revelation In Dynamic Trading," Journal of the European Economic Association, 2005, 3 (6), 1279–1317.
- **Hayek, F. A.**, "The Use of Knowledge in Society," *The American Economic Review*, 1945, 35 (4), pp. 519–530.
- Kunimoto, Takashi and Roberto Serrano, "Bargaining and competition revisited," Journal of Economic Theory, 2004, 115 (1), 78–88.
- **Kyle, Albert S.**, "Informed Speculation with Imperfect Competition," Review of Economic Studies, 1989, 56 (3), 317–55.
- Lauermann, Stephan and Gabor Virag, "Auctions in Markets: Common Outside Options and the Continuation Value Effect," May 2011. Mimeo.
- Milgrom, Paul, "A Convergence Theorem for Competitive Bidding with Differential Information," *Econometrica*, 1979, 47 (3), 679–88.
- _ and Ilya Segal, "Envelope Theorems for Arbitrary Choice Sets," Econometrica, 2002, 70 (2), 583-601.
- Milgrom, Paul R. and Robert J. Weber, "A Theory of Auctions and Competitive Bidding," *Econometrica*, 1982, 50 (5), pp. 1089–1122.
- Nöldeke, Georg and Thomas Tröger, "Matching Heterogeneous Agents with a Linear Search Technology," January 2009. Mimeo.

- Ostrovsky, Michael, "Information Aggregation in Dynamic Markets with Strategic Traders," March 2011. Mimeo.
- **Pesendorfer, Wolfgang and Jeroen Swinkels**, "The Loser's Curse and Information Aggregation in Common Value Auctions," *Econometrica*, 1997, 65 (6), 1247–1282.
- and _ , "Efficiency and Information Aggregation in Auctions," American Economic Review, 2000, 90 (3), 499–525.
- Reny, Philip J and Motty Perry, "Toward a Strategic Foundation for Rational Expectations Equilibrium," *Econometrica*, 2006, 74 (5), 1231–1269.
- Rostek, Marzena and Marek Weretka, "Price Inference in Small Markets," September 2011. Mimeo.
- Rothschild, Michael, "Searching for the Lowest Price When the Distribution of Prices Is Unknown," *Journal of Political Economy*, 1974, 82 (4), 689–711.
- **Rubinstein, Ariel and Asher Wolinsky**, "Equilibrium in a Market with Sequential Bargaining," *Econometrica*, 1985, 53 (5), 1133–50.
- Satterthwaite, Mark and Artyom Shneyerov, "Dynamic Matching, Two-Sided Incomplete Information, and Participation Costs: Existence and Convergence to Perfect Competition," *Econometrica*, 2007, 75 (1), 155–200.
- _ and _ , "Convergence to Perfect Competition of a Dynamic Matching and Bargaining Market with Two-sided Incomplete Information and Exogenous Exit Rate," Games and Economic Behavior, 2008, 63 (2), 435–467.
- Serrano, Roberto and Oved Yosha, "Information Revelation in a Market with Pairwise Meetings: The One Sided Information Case," *Economic Theory*, 1993, 3 (3), 481–99.
- Shneyerov, Artyom and Chi Leung Wong, "The rate of convergence to perfect competition of matching and bargaining mechanisms," *Journal of Economic Theory*, 2010, 145 (3), 1164–1187.
- _ , Dip Majumdar, and Huan Xie, "An Optimistic Search Equilibrium," May 2011. Mimeo.
- Smith, Lones, "Frictional Matching Models," Annual Review of Economics, 2011, 3 (1), 319–38.
- Wolinsky, Asher, "Information Revelation in a Market with Pairwise Meetings," Econometrica, January 1990, 58 (1), 1–23.