Learning and Price Discovery in a Search Market*

Stephan Lauermann† Wolfram Merzyn‡ Gábor Virág§

May 2016

Abstract

We develop a steady-state 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: Overlapping generations of traders gradually learn about the state of the market through a sequence of multilateral bargaining rounds. We construct an intuitive equilibrium where sellers accept all offers when the frictions are small. We characterize equilibrium trading, and belief patterns. We show that equilibrium outcomes are approximately competitive when the frictions are small given an intuitive refinement on off-equilibrium beliefs.

1 Introduction

Uncertainty about supply and demand is common, especially in decentralized markets. Labor markets, over-the-counter asset markets, and housing markets are examples where the market conditions may not be fully known to the market participants. This uncertainty has implications for agents' behavior: the agents may spend time experimenting with offers unlikely to be accepted, they may employ strategies to learn from market outcomes, and they may be more accommodating after being in the market for a longer time period without trading. These patterns of experimentation, learning, and accommodation over time—while common in real-life markets—cannot be captured by search models in which the aggregate supply and demand conditions are known by the agents in the market.

We introduce a model in which no individual trader knows the relative scarcity of the good being traded. Our model builds on standard frictional search-and-bargaining models in

*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 Boston College, Concordia University, Michigan State, NYU Stern, Penn State University, the University of Toronto, the Canadian Economic Theory Conference, the Decentralization Conference at Washington University, and the Midwest Theory Conference. We thank Isaac Sorkin and Qinggong Wu for research assistance. Stephan Lauermann acknowledges the National Science Foundation for financial support under Grant SES-1061831.

†University of Bonn, s.lauermann@uni-bonn.de.

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

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

the tradition of Mortensen (1982). These models assume that the agents know the market conditions, while we assume that the agents do not know the supply conditions. First, we analyze the trade and learning patterns that emerge under market-level uncertainty. Second, we ask whether the traders eventually learn the aggregate characteristics and whether the prices accurately reflect relative scarcity when the frictions are small.

Studying those questions, we make contributions to equilibrium search theory, a branch of the literature in which both sides of the market make strategic decisions. First, we construct a "full-trade" equilibrium, a construction that may help analyze a wide range of decentralized markets where the market conditions are not fully known. In our model, the identified equilibrium allows for a detailed analysis of trading strategies under market uncertainty. We relate these strategies to the well-known winner's curse and loser's curse that feature prominently in auction theory.

Second, we study the conditions under which the equilibrium price reflects market conditions, that is, conditions under which information is aggregated. The literature on dynamic matching and bargaining games, pioneered by Rubinstein and Wolinsky (1985) and by Gale (1987), addresses the question of how prices are formed in decentralized markets and whether these prices are Walrasian. Existing models, however, assume that market demand and supply are common knowledge among traders. This assumption is restrictive because markets have been advocated over central planning precisely on the grounds of the markets' supposed ability to "discover" the equilibrium prices by eliciting and aggregating information that is dispersed in the economy; see Hayek (1945). Beyond the theoretical importance, the question whether markets aggregate information has implications for policies that are intended to increase the transparency of markets, to remove insider information, to provide innovation incentives, or to centralize markets to decrease informational frictions.[1]

As an illustration, consider a bidder on eBay who seeks to purchase a consumer good. eBay is a decentralized market but has a well-defined trading mechanism that is close to our model. When bidding on a particular item, a bidder shades her bid below her true valuation to account for the expected continuation value (waiting for another auction). Waiting costs give the current seller market power to increase his reserve prices. Conventional search theory reflects these considerations.

However, there is evidence that is not consistent with buyers knowing aggregate market conditions. Juda and Parkes (2006) study the eBay market for a specific Dell monitor. They find that "[a]mong the 508 bidders that won exactly one monitor and participated in multiple auctions, 201 (40%) paid more than $10 more than the closing price of another auction *in which they bid*, paying on average $35 more (standard deviation $21) than the closing price

[1] For informationally opaque over-the-counter asset markets and their performace, see Duffie (2003).

of the cheapest auction *in which they bid* but did not win" (emphasis theirs). This piece of evidence shows not only that there is considerable amount of price dispersion in that market but also that buyers do not always know the "going price" of the object. In fact, a bidder may later regret not having been more aggressive before. This is reflected in dynamic bid patterns as well: "A simple regression analysis shows that bidders tend to submit maximal bids to an auction that are $1.22 higher after spending twice as much time in the system, as well as bids that are $0.27 higher in each subsequent auction."

These pieces of evidence are consistent with buyers experimenting, learning, and being more accommodating over time, which are hard to reconcile with models of conventional search theory where buyers fully know market conditions from the outset.[2] However, these observations are consistent with an environment with aggregate uncertainty where buyers need to learn market conditions.

We develop a model to study markets with aggregate uncertainty. The model combines elements of Wolinsky (1990) and Satterthwaite and Shneyerov (2008). The matching technology and the bargaining protocol are adopted from Satterthwaite and Shneyerov (2008). 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 who are matched with each seller. Each seller conducts a first-price sealed-bid auction with a secret reserve price. Successful buyers and sellers leave the economy, and unsuccessful traders leave the market with some exogenous exit rate; otherwise, they remain in the market to be rematched in the next period. The possibility of exiting makes waiting costly and is interpreted as the "friction" of trade.

There are two advantages of using auctions instead of the more standard bilateral meeting (and bargaining) technology in our model. First, we intend to capture the effect of local competition, which is an important part of many real-world markets. Second, we believe that purely bilateral matching represents a very special environment and it may fail to aggregate information even when most other settings do.[3]

The defining feature of our model is uncertainty about a binary state of nature (high or low) similarly to Wolinsky (1990). The realized state is unknown to the traders and does not change over time. The state of nature determines the relative scarcity of the good. The mass of incoming buyers is larger in the high state and smaller in the low state, whereas the mass of incoming sellers is independent of the state of nature. The larger the mass of entering buyers relative to the mass of entering sellers, the scarcer the good. Every buyer receives a noisy signal about the state of the world upon birth. Moreover, after every auction, the

[2] In fact, based on insights from conventional search theory Juda and Parkes (2006) conclude that these observations imply that buyers make mistakes.

[3] For a discussion, see Section 7.1.

losing buyers obtain additional information regarding the state, because they are able to draw an inference from the fact that their respective bids lost; the losing bidders do not observe other buyers' bids. We concentrate on steady-state equilibria that are (weakly) monotone in the beliefs of the agents: A monotone strategy is such that a buyer who attaches a higher probability to the high state bids more, and a seller who attaches a higher probability to the high state sets a higher reserve price. This assumption is consistent with the idea that the high state favors sellers (more buyers enter than in the low state), and therefore, if the high state is more likely, then intuitively the buyers should compete more vigorously, and the sellers should demand a higher price.

We show that the buyers shade their bids to account for the opportunity cost of foregone continuation payoffs. Moreover, although 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 and introduces an endogenous common value element. The resulting winner's curse leads to further bid shading: Winning an auction implies that, on average, fewer bidders participate and that the participating bidders are more optimistic about their continuation payoff.[4] Therefore, the winner's curse implies a lower value for winning the good than expected before winning. Offsetting the winner's curse is the "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 participate and that the participating bidders are more pessimistic about their continuation payoffs. The loser's curse implies that bidders become more pessimistic and raise their bids after repeated losses over time.[5]

The behavior of the sellers is quite different: After a seller has not been able to transact for a short period of time, he lowers his reserve price believing that the state is most likely to be low. The reason for the quick concession is that the pivotal conditioning event (when calculating the winner's curse effect) is the event when the largest bid is *equal* to a certain reserve price. Conditioning on this event tends to decrease the probability of the high state. An interesting illustration of this observation comes when we identify *full-trade* equilibria where the sellers accept all equilibrium bids. In such equilibria, the pivotal event of the highest bid being equal to $b < v$ indicates that the state is low almost surely in the limit. This explains why sellers are happy accepting any offers that are made in (a full-trade) equilibrium.

[4] Conditional on winning, the buyers also learn from the fact that the seller's reservation price was lower than the accepted bid. This learning pattern reinforces the winner's curse effect as sellers set low reserve prices when the state is likely to be low.

[5] Thus, 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*.

We also study whether the equilibrium allocation reflects the scarcity of the object, that is, the conditions under which information is aggregated properly. For this, we characterize the equilibria 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 in many ways but not fully. First, we show that the equilibrium trading probabilities are competitive in the limit; that is, the short side of the market trades almost surely in the limit. Moreover, 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; that is, the price is competitive in the limit. However, if in the low state the mass of incoming buyers is smaller than the mass of incoming sellers and vice versa in the high state, then the equilibrium price may be higher than the seller's costs in the low state; that is, there are equilibria with noncompetitive limits.

We introduce an intuitive refinement, the refinement of monotone beliefs in Section 6. This refinement requires off-equilibrium beliefs to satisfy that a higher winning bid indicates that the state of the world is high with a larger probability. Under this refinement, we show that the equilibrium price is competitive in the limit even if sellers outnumber buyers.

In Section 2, we discuss our contribution to the literature. In Section 3, we introduce the model and the equilibrium concept. In Section 4, we show existence of equilibrium by constructing a full-trade equilibrium. We also discuss the strategies of agents and how learning about the state affects equilibrium behavior. Our main results for information aggregation are stated in Sections 5 and 6. In Section 7, we provide a discussion of the role of the matching protocol in information aggregation, as well as of the speed of trade in calendar time. The Appendix contains the proofs. A supplementary online Appendix presents the construction of full-trade equilibria.

2 Contribution to the Literature

We contribute to the 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 fully specified "decentralized" trading institution leads to outcomes that are competitive when trade frictions 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

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

important for *positive* theory in order to understand under which conditions markets can and cannot be well approximated with competitive analysis and for *normative* theory in order to understand which trading institutions are able to decentralize desirable allocations.

In existing dynamic matching and bargaining models, market demand and supply are known. Thus, each market participant can individually compute the market-clearing price before trading. In the early matching and bargaining literature, the preferences and endowments of each individual trader were typically assumed to be observable. Satterthwaite and Shneyerov (2007, 2008) introduce a model with private information. Because they assume a continuum of agents, the realized distribution of preferences is known in their model; thus, there is idiosyncratic but no aggregate uncertainty.

The absence of aggregate uncertainty in existing dynamic matching and bargaining games is a substantial restriction, and considering uncertainty is important for at least two reasons. First, assuming that aggregate market conditions are known to all participants is unrealistic in many markets. Our model allows us to study those markets in which this assumption is not met. Second, as previously argued, price discovery has been highlighted as an integral function of markets. Our model allows us to investigate whether and under which conditions decentralized markets can serve this function.

Modeling uncertainty about market demand and supply leads to novel conceptual challenges. First, we need to characterize the endogenous distribution of beliefs in the economy. This is a key difficulty that has blocked progress on equilibrium learning in search models thus far. Second, when buyers and sellers bargain over the price, their common outside option is to continue searching for other partners. This introduces an endogenous common value element into the bargaining game. Bargaining with common values is known to cause difficulties because of the multiplicity-of-equilibrium problem. We propose a tractable model that minimizes the impact of these problems, while still allows us to study the important questions.

A recent contribution by Majumdar, Shneyerov, and Xie (2015) is the only other paper that considers a dynamic matching and bargaining game in which market demand and supply depend on an unknown state of nature. Their key assumption is that traders from each market side are subjectively certain that the state of nature is in their favor. This assumption allows them to prove the existence of equilibrium and provide a full characterization of equilibrium even when the traders are heterogenous. Given the assumption, however, there is no uncertainty from the viewpoint of each trader. Therefore, issues at the heart of this paper such as private learning and the endogenous distribution of beliefs, as well as the induced value interdependence, are absent.

There is also related literature on consumer search, some of which considers uncertainty on the buyers' market side; see, for example, Benabou (1993), Dana (1994), and Janssen and Shelegia (2015). Different from our model, in that literature the other market side (sellers) know the state. These contributions typically focus on "reserve-price equilibria." An exception is Janssen, Parakhonyak, and Parakhonyak (2014). In our model, no trader knows the state, we consider the full set of equilibria, and we use our model to study information aggregation with small frictions. Our companion paper, Lauermann and Virag (2012), studies the case in which the losers of an auction can participate in a large market and seek their best outside options.

Our paper is also related to work on matching and bargaining with exogenously assumed *common values*. In these models, preferences depend on an unknown state, and consequently, these models are used to study the foundations for rational expectations equilibria. In contrast, we study the foundations for competitive equilibrium in a standard exchange economy. Particularly prominent contributions to search with common values are Wolinsky (1990) and later work by Blouin and Serrano (2001).[7] These contributions provide negative convergence results and uncover 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 the search frictions are small. In our model, the winner's curse implies a similar effect: When the exit rate vanishes, the buyers bid low for an increasingly large number of periods. Yet in our setting, this "insistence problem" is overcome by the opposing loser's curse.

Wolinsky (1990) and Blouin and Serrano (2001) assume that traders can choose between only two price offers (bargaining postures). This assumption makes their models tractable. It is an open question whether the trading outcomes in these models are competitive if the frictions are small and if a restriction is not imposed on the set of prices. Golosov, Lorenzoni, and Tsyvinski (2011) consider a related search model with common values in which the traded good is divisible. They do not impose a restriction on the set of price offers. The friction in their model is an exogenous probability that trading stops in any given period. They show that the equilibrium outcomes approximate ex-post efficient outcomes in the event that the game has not stopped for a sufficiently large number of periods. Golosov, Lorenzoni, and Tsyvinski study the trading outcome with a fixed, positive stopping probability, and they

[7] 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. Lauermann and Wolinsky (2016) study information aggregation if a single, privately informed buyer searches among many sellers.

do not study the question whether outcomes become competitive in the "frictionless" limit when the stopping probability is small.[8]

There is a large body of related work on the foundation for rational expectation equilibrium 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 the work on information aggregation in Cournot models, summarized in Vives (2010)) and on the behavior of traders in financial markets (e.g., Kyle (1989), Ostrovsky (2011), and Rostek and Weretka (2012)). 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. 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 Set-up

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 $v > 0$ for the good. Each seller has one unit to trade. The common cost of selling is $c \in [0, v)$. Trading at price p yields payoffs $v - p$ and $p - c$, respectively. A trader who exits the market without trading has a payoff of zero. The valuation exceeds the cost, so there are gains from trading. Buyers and sellers maximize the expected payoffs.

Similar to Wolinsky (1990), there are two states of nature, a high state and a low state $w \in \{h, \ell\}$. Both states are equally likely. The realized state of nature is fixed throughout and is 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, the state of nature 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

[8] As explained in Golosov, Lorenzoni, and Tsyvinski, the 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 the high state, it is d^h. More buyers enter in the high state, $d^h > d^\ell$. The mass of sellers who enter each period is the same in both states and is equal to s.[9] We are most interested in the case where $d^h > s > d^\ell$, so that there is uncertainty about who is on the long side of the market.

The buyers and the sellers 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 or a seller. 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 $[\underline{\theta}^B, \bar{\theta}^B]$, with cumulative distribution functions $G^B(\theta \mid h)$ and $G^B(\theta \mid \ell)$, respectively. The distributions are continuous and admit continuous probability density functions, $g^B(\theta \mid h)$ and $g^B(\theta \mid \ell)$. Notice that using Bayes' rule the distributions must be such that $\theta = \frac{d^h g^B(\theta|h)}{d^h g^B(\theta|h)+d^\ell g^B(\theta|\ell)}$, which is the no-introspection condition of Smith (2011) for entering buyers.[10] Equivalently, the likelihood ratio is

$$\frac{\theta}{1-\theta} = \frac{d^h g^B(\theta \mid h)}{d^\ell g^B(\theta \mid \ell)}.$$

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 $d^h/d^\ell > 1$.[11]

The sellers' side is analogous. Each seller 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 sellers are assumed to be distributed on the support $[\underline{\theta}^S, \bar{\theta}^S]$, with cumulative distribution functions $G^S(\theta \mid h)$ and $G^S(\theta \mid \ell)$, respectively. The distributions are continuous and admit continuous probability density functions, $g^S(\theta \mid h)$ and $g^S(\theta \mid \ell)$. Notice that using Bayes' rule the distributions must be such that $\theta = \frac{d^h g^S(\theta|h)}{d^h g^S(\theta|h)+d^\ell g^S(\theta|\ell)}$, or, equivalently, the likelihood ratio satisfies

$$\frac{\theta}{1-\theta} = \frac{d^h g^S(\theta \mid h)}{d^\ell g^S(\theta \mid \ell)}. \tag{1}$$

It is assumed that $\bar{\theta}^B < 1$; that is, the agents' signals are boundedly informative. To avoid

[9] The assumption that $s^h = s^\ell$ is a normalization that does not change any of the results.

[10] The no-introspection condition states that an agent should not be able to update his belief just by knowing what his belief is. For more on this condition, please see the more general formula (23) and the discussion there.

[11] To formally define updating based on entering the market, suppose that there is a *potential* set of buyers of mass d, with $d \geq d(h)$. In state w, a mass $d(w)$ of the potential buyers actually enters the market. Alternatively, one can simply interpret $d(h)/d(\ell)$ as the prior of an entering buyer. For games with population uncertainty and updating about an unknown state of nature, see Myerson (1998) and, especially, Milchtaich (2004).

technical difficulties, and to simplify exposition, we assume that[12]

$$\underline{\theta}^B > 1/2. \tag{2}$$

Each period unfolds as follows:

1. Entry occurs (the *"inflow"*): A mass s of sellers and a mass d^w of buyers enter the market. The buyers and sellers privately observe signals, as previously described.

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 $n = 0, 1, 2, ...$ buyers is Poisson distributed[13] and is equal to $e^{-\mu}\mu^n/n!$, where $\mu(w) = D(w)/S(w)$ is the endogenous ratio of the mass of buyers to the mass of sellers in the stock as described below. The expected number of buyers who are matched with each seller is equal to $\mu(w)$, of course. We sometimes refer to $\mu(w)$ as a measure of market "tightness."

3. Each seller runs a sealed-bid auction with a secret reserve price r. The buyers do not observe r or how many other buyers are matched with the same seller. The bids are not revealed ex post, so the buyers learn only whether they have won with their submitted bid. Moreover, for simplicity, the seller observes only whether the highest bid is above r or not.

4. A seller leaves the market if his good is sold; otherwise, the seller stays in the stock with probability $\delta \in [0, 1)$ to offer his good in the next period. A winning buyer pays his 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.

5. Upon losing, the remaining buyers update their beliefs based on the information gained from losing with their submitted bids. Upon not trading, the remaining sellers update their beliefs based on the information gained from not trading with their chosen reserve price r. 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.

[12] By (1), it follows that if all buyers had the same signal upon entry, then $\theta = d^h/(d^h + d^\ell)$. Therefore, the assumption that $\underline{\theta}^B > 1/2$ means that agents do not receive too precise information upon entry. A similar assumption cannot be made for the sellers as without different signals all entering sellers would hold a belief $d^\ell/(d^h + d^\ell) < 1/2$.

[13] This 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.

On the individual level, the exit rate $1-\delta$ acts similarly to a discount rate: Not trading today creates a risk of losing all trading opportunities with probability $1-\delta$. On the aggregate level, the exit rate ensures that a steady state exists for all strategy profiles. Traders do not discount future payoffs beyond the implicit discounting of the exit rate.

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

The matching technology and the bargaining protocol are adapted from Satterthwaite and Shneyerov (2008). In fact, our model is a special case of theirs if the state is known, with the remaining difference being that they allow for heterogeneous values and costs. We discuss the case in which the state is known now.

3.2 Full Information Benchmark

Suppose that $d(h) = d(\ell) = d$. In this case, the state of the market is known, and we can drop the beliefs from the description of the model. The analysis of the steady-state of this case is standard. In particular, let V^S and V^B denote the steady-state equilibrium payoffs of the sellers and the buyers, respectively. A standard perfection requirement implies that the sellers' reserve price satisfies

$$r - c = \delta V^S. \tag{3}$$

The sellers accept the highest price offer if trading at that price yields higher payoffs than continuing the search and rejects the price offer otherwise. By standard arguments, the buyers' bidding strategies must be fully mixed without atoms and no gaps on some interval $[r, \bar{p}]$, where $\bar{p}$ is determined by the buyers' indifference between the end-points,

$$v - \bar{p} = e^{-\mu}(v - r) + \delta\left(1 - e^{-\mu}\right)V^B,$$

where $e^{-\mu}$ is the probability that the seller is matched with no buyer. Here, $v - \bar{p}$ is the payoff from bidding $\bar{p}$ as the highest bid in the support is certain to win. The bid r wins only if there is no other bidder, which happens with probability $e^{-\mu}$.[14] The lowest bid in the offer distribution must be equal to r exactly because it wins only if there is no other bidder. By using a similar indifference condition for intermediate bids, one can characterize the whole distribution of the bids as a function of V^B and V^S (determining r). Then, a fixed-point argument implies the existence of an equilibrium and allows to determine V^S and V^B as well.

This full-information benchmark is essentially identical to Burdett and Judd (1983), one of the earliest contributions to equilibrium search theory with a non-degenerate distribution

[14] The Poisson distribution implies that $e^{-\mu}$ is also the probability that there is no other buyer from the perspective of a buyer who is matched with the seller.

of prices. It is a (very) special case of Satterthwaite and Shneyerov (2008), who allow for heterogeneous values and costs in order to study market clearing with private information. In our model, the heterogeneity of beliefs takes a somewhat similar role as the heterogeneity of preferences in theirs. In particular, this heterogeneity "purifies" the mixed offer strategy.

3.3 Steady-state Equilibrium

We now return to our base model in which $d(h) > d(\ell)$. A steady-state equilibrium specifies the strategies and the endogenous stocks (the masses of the buyers and sellers and the distribution of their beliefs). We restrict our attention to pure strategy equilibria where the bid is a (weakly) increasing function of the belief of the buyer and the reserve price is (weakly) increasing in the belief of the seller. In addition, the distribution of beliefs is sufficiently "smooth."

Formally, the masses of buyers and sellers in the stock are $D(w)$ and $S(w)$. The distributions of beliefs are given by cumulative distribution functions $\Gamma^j(\cdot \mid w)$ for $j = B, S$ and $w = \ell, h$. We assume that each function Γ^j is absolutely continuous.[15] Furthermore, we assume that Γ^j is piecewise twice continuously differentiable,[16] and that we can choose a density, denoted $\gamma^{w,j}$, that is right continuous on $[0, 1)$. Finally, we assume that $\Gamma^j(\cdot|h)$ and $\Gamma^j(\cdot|\ell)$ are mutually absolutely continuous. Thus, the distributions have common support.[17]

The bidding strategy β is a weakly increasing function and maps beliefs from $[0, 1]$ to $[c, v]$. We often use the generalized inverse of β, given by $\beta^{-1}(b) = \inf\{\theta | \beta(\theta) \geq b\}$, where $\beta^{-1}(b) = 1$ if $\beta(\theta) < b$ for all θ. Moreover, we assume that β is strictly increasing in the interior of the support of Γ^B so there are no ties. The reserve price strategy ρ is a weakly increasing function and maps beliefs from $[0, 1]$ to $[c, v]$.

Finally, $\theta^B_+(\theta, b)$ is the posterior of a buyer with initial belief θ conditional on not trading with bid b, and $\theta^S_+(\theta, r)$ is the posterior of a seller with initial belief θ conditional on not trading with reserve price r.

We characterize the requirements for the equilibrium objects. Let $\theta^B_{(1)}$ denote the first-order statistic of the buyers' beliefs in any given match. We set $\theta^B_{(1)} = 0$ if there is no buyer present. Let $\Gamma^B_{(1)}(x \mid w)$ denote the probability that the highest belief in the auction is below

[15] Note that this is a restriction on the kind of steady-state equilibria that we study.

[16] A 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.

[17] Mutual absolute continuity can be shown to be implied by each Γ^j being absolutely continuous, but it is convenient to assume it right away.

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 having no buyer present is $\Gamma^B_{(1)}(0 \mid w)$ by our assumption that there is no atom in the distribution of beliefs at zero. The Poisson distribution implies $\Gamma^B_{(1)}(0 \mid w) = e^{-\mu(w)}$, where $\mu(w) = D(w)/S(w)$ as defined before. In general, the first-order statistic of the distribution of beliefs is given by

$$\Gamma^B_{(1)}(x \mid w) = e^{-\mu(w)(1-\Gamma^B(x|w))}. \tag{4}$$

Here, $\mu(w)(1 - \Gamma^B(\theta \mid w))$ is the ratio of the mass of buyers who have a belief above θ to the mass of sellers, and $e^{-\mu(w)(1-\Gamma^B(\theta|w))}$ is the probability that the seller is matched with no buyer who has such a belief.

We derive the posterior of a buyer upon not trading with bid b. Given the assumption that the bidding strategies are strictly increasing, losing with a bid b implies that there was some bidder in the match with a belief above $x = \beta^{-1}(b)$ or that the seller set a reserve price higher than r. Then let

$$q^B(b \mid w) := \Gamma^B_{(1)}\left(\beta^{-1}(b) \mid w\right) \Gamma^S(\rho^{-1}(b) \mid w)$$

denote the probability that a buyer wins with bid b in state w. We often refer to q^B as the per-period trading probability. Denote with $q^B(\theta, b) = \theta q^B(b \mid h) + (1-\theta) q^B(b \mid \ell)$ the probability of winning if the probability of the high state is θ. Bayes' rule for the posterior upon not trading with bid b requires that

$$\theta^B_+(\theta, b) = \frac{\theta(1 - q^B(1, b))}{1 - q^B(\theta, b)}, \tag{5}$$

if the denominator is positive. Otherwise, it can be chosen in any way with the constraint (introduced and explained later) that θ^B_+ is monotone in b.

Let $q^S(r \mid w) = 1 - \Gamma^B_{(1)}\left(\beta^{-1}(r) \mid w\right)$ denote the probability that a seller trades with reserve price r in state w, and $q^S(\theta, r) = \theta q^S(r \mid h) + (1-\theta) q^S(r \mid \ell)$ the probability of trading if the probability of the high state is θ. The belief of a seller who did not trade with reserve price r is given as

$$\theta^S_+(\theta, r) = \frac{\theta(1 - q^S(1, r))}{1 - q^S(\theta, r)}. \tag{6}$$

For some later discussions, it is useful to define the "tieing posterior". Let $\mathcal{B}$ denote the Borel sigma-algebra defined on $[c, v]$, and let $\theta^S_0 : [0,1] \times \mathcal{B} \to [0,1]$, with $\theta^S_0(\theta, A) = \Pr\left(h | b_{(1)} \in A, \theta\right)$, be the posterior probability of h conditional on the highest bid being in a set A. Whenever the conditioning event $b_{(1)} \in A \subset \mathcal{B}$ has positive probability,

$$\theta^S_0(\theta, A) = \frac{\theta \Pr(b_{(1)} \in A | h)}{\theta \Pr(b_{(1)} \in A | h) + (1-\theta) \Pr(b_{(1)} \in A | \ell)}. \tag{7}$$

Taking the conditioning event A to be a single bid b, we obtain $\theta_0^S(\theta, b) = \Pr\left(h | b_{(1)} = b, \theta\right)$.

The steady-state stock needs to satisfy

$$S(w)\Gamma^S(\theta \mid w) = \tag{8}$$
$$sG^S(\theta \mid w) + \delta S(w) \int_{\left\{\tau: \theta_+^S(\tau, \rho(\tau)) \leq \theta\right\}} \Gamma_{(1)}^B\left(\beta^{-1}(r) \mid w\right) d\Gamma^S(\theta \mid w).$$

To explain this steady-state condition for the stock, suppose that the mass of sellers is $S(w)$ today. The steady-state mass of sellers in the stock who have a type below θ is equal to $S(w)\Gamma^S(\theta \mid w)$. This mass has to be equal to the mass of sellers in the inflow with type less than θ (the first term on the right-hand side) plus the mass of sellers who lose, survive, and update to some type less than θ (the second term). In steady state, these two populations must be identical.

We follow Fudenberg and Levine (1993) who study a steady-state model with stationary equilibrium beliefs. Despite focusing on a steady-state distribution of beliefs in the entire economy, individual agents change their beliefs over time in our model-as in Fudenberg and Levine (1993). Sellers who trade quickly infer that the state is likely to be high, while buyers who trade quickly infer that the state is likely to be low. These agents do not continue in the pool; therefore, continuing agents become pessimistic on both sides (buyers attach higher probabilities to the high state, while sellers attach higher probabilities to the low state upon not being able to trade). This is counterbalanced by new entrants who are more optimistic; thus, the composition maintains a steady-state distribution of beliefs about the state of the economy.

The inflow of buyers with type less than θ is $d^w G^B(\theta \mid w)$. The steady-state condition is

$$D(w)\Gamma^B(\theta \mid w) = d^w G^B(\theta \mid w) + \tag{9}$$
$$+\delta D(w) \int_{\left\{\tau: \theta_+^B(\tau, \beta(\tau)) \leq \theta\right\}} \left(1 - \Gamma_{(1)}^B\left(\beta^{-1}(b) \mid w\right) \Gamma^S(\rho^{-1}(b) \mid w)\right) d\Gamma^B(\tau \mid w).$$

The steady-state mass of buyers in the stock who have a type below θ is equal to $D(w)\Gamma^B(\theta \mid w)$. This mass has to be equal to the mass of 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]

[18] For the purpose of this paper, the steady-state model is *defined* by (8) and (9). 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 Duffie and Sun (2007).

Let $V^B(\theta)$ denote the value function, which is equal to[19]

$$\max_b \, (v-b)q^B(\theta,b) + \delta(1-q^B(\theta,b))V^B\left(\theta_+^B(\theta,b)\right). \tag{10}$$

A bidding strategy β is optimal if $b=\beta(\theta)$ solves problem (10) for every θ.

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

$$\max_r q^S(\theta,r)(r-c) + (1-q^S(\theta,r))V^S\left(\theta_+^S(\theta,r)\right). \tag{11}$$

A reserve price strategy ρ is optimal if $r=\rho(\theta)$ solves the problem (11) 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, $S(h), D(h), S(\ell), D(\ell)$, and distribution functions Γ^B, Γ^S such that the steady-state conditions (8) and (9) hold for all θ; (ii) updating functions θ_+^B, θ_+^S that are consistent with Bayes' rule (5), (6); and (iii) strictly increasing functions β and ρ that are optimal (solve (10) and (11), respectively).

An equilibrium determines the value functions $V^j(\theta)$, $j\in\{B,S\}$, (lifetime) trading probabilities $Q^j(\theta|w)$, expected trading prices $P^j(\theta|w)$, and payoffs $EU^j(\theta|w)$ with $V^j(\theta) = \theta EU^j(\theta|h) + (1-\theta)EU^j(\theta|\ell)$. We denote with (b,β) the bidding sequence of b today and then following the strategy β from tomorrow onward (given the updated beliefs). We abuse notation and write $Q^B((\beta,\theta)|w)$, $Q^S((\rho,\theta)|w)$, $P^B((\beta,\theta)|w)$, etc., for the trading probabilities and the (expected) transaction prices for the allocation of type θ who follows strategy β or ρ. Let $Q^B(b|w)$, $Q^S(r|w)$, $P^B(b|w)$, $Q^S(r|w)$ denote the trading probability and price variables for the constant bidding strategy $\beta\equiv b$ and $\rho\equiv r$. For later reference, we note the following connection[20] between the per-period and lifetime probabilities for the same bid b:

$$Q^B(b|w) = \frac{q^B(b|w)}{1-\delta+\delta q^B(b|w)}, \tag{12}$$

and similarly for the sellers:

$$Q^S(r|w) = \frac{q^S(r|w)}{1-\delta+\delta q^S(r|w)}. \tag{13}$$

We impose a notion of perfectness on the sellers' equilibrium strategies. In particular, we require the following undominatedness condition:

$$\rho(1)-c=\delta V^S(1) \qquad \text{and} \qquad \rho(0)-c=\delta V^S(0). \tag{14}$$

[19] Recall that θ_+^B is the posterior of a buyer who lost with b and had an initial belief of θ.

[20] This follows from straightforward algebra.

Thus, a seller who knows the state accepts prices if and only if accepting the price yields larger payoffs than continued searching. This requirement is consistent with the perfectness notion (3) that was required in the full-information benchmark.

The main bite of (14) comes when we combine it with the assumption that equilibrium strategies are monotone. Note that, as indicated earlier, the assumption that ρ is non-decreasing also means that we assume that $V^S(1) \geq V^S(0)$. In addition, (14) rules out "no-trade" equilibria such as $\beta(1) = 0$ and $\rho(0) = 1$, as as —using that $\beta(\theta) \leq v$ for all θ by construction—$\rho(1) \leq \delta(v - c)$. Therefore, the main way we utilize the condition is by using that for all $\theta \in (0, 1)$,

$$\delta V^S(0) \leq \rho(\theta) - c \leq \delta V^S(1). \tag{15}$$

3.4 Information Aggregation

We study the equilibrium trading outcome in each state as the exit rate becomes small. In particular, we ask whether trade between buyers and sellers takes place at the "correct," market-clearing prices. We recall the definition of *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 as $Q^B(\theta \mid w)$ and $P^B(\theta \mid w)$, respectively. For a seller, the trading outcome consists of the probability of being able to sell the good and the expected price received, denoted as $Q^S(\theta \mid w)$ and $P^S(\theta \mid w)$. 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 *competitive outcome* relative to the economy defined by the inflow if the prices and trading probabilities are as follows. If $d^w < s$, then $P^B(\theta \mid w) = P^S(\theta \mid w) = c$, $Q^B(\theta \mid w) = 1$, and $Q^S(\theta \mid w) = d^w$. If $d^w > s$, then $P^B(\theta \mid w) = P^S(\theta \mid w) = v$, $Q^B(\theta \mid w) = s/d^w$, and $Q^S(\theta \mid w) = 1$. We omit the non-generic case of $d^w = s$.

The competitive allocation has three important features. First, the law of one price holds; that is, all transactions take place at the same price. Second, this price clears the market. Third, the short side of the market trades almost surely.

We consider whether the trading outcome becomes competitive 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. We know that an equilibrium exists for each δ_k that is large enough, as we show it in the online Appendix. Pick any such equilibrium and denote the corresponding equilibrium magnitudes with β_k,ρ_k, Γ_k^h, Γ_k^ℓ, D_k^h,P_k^j, Q_k^j, 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. In this case, we say that the information is aggregated in the economy.

4 Full-trade Equilibria

4.1 What are Full-trade Equilibria, and Why are They Useful?

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). To keep our model tractable, we made several simplifying assumptions. Our first simplification is to use a tractable auction model with a random number of bidders as our trading protocol. Even in this relatively simple model, difficulties for our analysis arise as the sellers' costs and the buyers' values depend (through endogenous outside options) on a state of the world about which the agents have private information. In other words, our auction is an interdependent value auction, but with a random number of bidders and endogenous outside options. Our second simplification is to assume that traders learn nothing beyond losing with their bids (or reserve price), which abstracts away from other channels of learning. Even with those modeling shortcuts, two sets of technical problems remain when proving existence of a tractable and intuitive equilibrium in which agents update monotonically.

The Affiliation Problem with Asymmetric Traders. As Reny and Perry (2011) show, the first order statistic of affiliated random variables that are mapped by two different functions (reserve prices and bidding strategies) does not have to be affiliated. We solve this problem by first studying full-trade equilibria that preserve tractability.

Random Number of Bidders. With a Poisson distributed number of bidders (even if the means are the same), the affiliation of the first-order statistic (of the bidders' type) with the state is lost. We "solve" this problem by placing a lower bound on the bidders' types (in the inflow). The formal condition is (2), which implies that belief updating is monotone upward in any *full-trade* equilibrium.

To illustrate the problem, suppose that the expected number of bidders is the same in the two states. Then, there is no winner's curse at the bottom (the bidder wins if no other bidder is there), and there is no winner's curse at the top. However, there is a potential winner's curse in the middle, meaning that the expected value conditional on winning is not monotone.

Given these complications, our analysis proceeds in steps. Before providing general characterization results that apply to all equilibria, we **construct equilibria of a particularly**

simple form where sellers accept all equilibrium bids. A formal definition follows below, but for now, imagine a profile of strategies where all sellers post the same reserve price r_0, and (consequently) all buyers bid at least r_0. Therefore, all matches result in trades, and thus, we use the term full-trade equilibrium for an equilibrium of this form. Constructing such full-trade equilibria solves the problem of equilibrium existence. To follow this constructive approach, we need to derive the distribution of equilibrium beliefs when all matches result in trades. To derive the belief distribution, we show that the steady-state distribution of beliefs can be constructed using an intuitive recursive algorithm when all matches result in trades.

Beyond proving the existence of equilibrium, there are several additional reasons to study full-trade equilibria. First, such equilibria are standard to study in the literature. See Satterthwaite and Shneyerov (2007) and the consumer search literature discussed in our literature review for examples where full-trade equilibria are used.

Second, such equilibria are easy to work with and allow for more explicit characterization results. For example, we are able to derive the distribution of beliefs explicitly so the model is amenable to numerical analysis. Third, as we argue later, in our model most of the important characterization results extend to other types of equilibria; see the discussion at the end of Section 4.2.3. Finally, in Section 4.2.1 we argue that any full-trade equilibrium identified above would remain equilibrium if the seller observed the winning bid before setting his reserve.

4.2 Full-trade Equilibria with State-dependent Competitive Allocation

4.2.1 Equilibrium Construction

In this section, we concentrate on the leading case where the competitive allocation depends on the state (that is, where $d^{\ell} < s < d^{h}$). The other cases are studied in Section 5 where general characterization results are provided. All proofs are given in the online Appendix.

For our construction, fix any $r_0 \in (c, v)$ and assume that all sellers set reserve price r_0. The sellers' strategies then induce a game among buyers who take the common reserve prices as fixed. A **steady-state bidding equilibrium** given r_0 is then a combination of stocks and strategies such that the stocks satisfy the steady-state conditions given the strategies, the buyers' strategy is optimal, and the sellers' strategy is $\rho(\theta) \equiv r_0$; that is, we simply drop the sellers' optimality condition. The analysis is substantially simplified. In particular, it is immediate that the buyers always bid at least r_0, and thus, in equilibrium all bids are accepted by all sellers. The following result shows that the equilibrium in the buyers' game is unique by applying the contraction-mapping theorem (see the online Appendix).

Proposition 1 *For any $r_0 \in [c, v]$, there exists a unique steady-state bidding equilibrium.*

We show that when δ_k is high, then for any r_0 there exists a full-trade equilibrium of the overall game where all the sellers set reserve price r_0 and the buyers bid according to the unique bidding equilibrium identified in Proposition 1.

Proposition 2 *Let $d^\ell < s < d^h$. Let $r_0 \in (c, v)$. If δ_k is sufficiently large, there exists a steady-state equilibrium of the original game with $\rho(\theta) \equiv r_0$ for all $\theta \in (0, 1)$.*

The proof is in the online Appendix. The proof establishes two results:

1. Setting a reserve price above r_0 is not optimal for any seller with any belief.
2. The equilibrium satisfies undominatedness.

For point 1, the proof shows that for any $b \in (r_0, v)$, the probability of the low state conditional on the highest bid being equal to b converges to 1 as δ_k goes to 1. The starting point of the proof is that in a full-trade equilibrium the buyer-seller ratio in the high state ($\mu_k(h)$) converges to infinity at rate $1/(1-\delta_k)$, and the buyer-seller ratio in the low state ($\mu_k(\ell)$) converges to zero at the rate of $1-\delta_k$.[21] Take any bid $b \in (r_0, v)$. The proof shows that the density of the highest bid at any such b is proportional to $1-\delta_k$ in the low state, and proportional to $e^{-1/(1-\delta_k)}$ in the high state.[22] Given that $e^{-1/(1-\delta_k)}$ tends to zero much faster than $1-\delta_k$ does, it follows that the probability of the low state conditional on the highest bid being equal to b converges to 1 as δ_k goes to 1.

Given this observation, to establish point 1, it is sufficient to argue that the seller's revenue is lower than r_0 in the low state in the limit. The argument has two parts. First, as there are more sellers than buyers entering each period in the low state, the mechanics of any (full-trade or not) equilibrium implies that the trading probability for the seller in the low state is strictly less than one in the limit. Second, to avoid falling victim to the winner's curse, each buyer depresses his bid substantially and bids close to r_0 until the bidder is almost sure that the state is high. Consequently, the expected transaction price in the low state in the limit converges to r_0 in the low state.[23] These two observations imply that the seller's expected revenue is lower than r_0 in the low state in the limit.

This argument also implies that if the seller *observed* the winning bid b before setting a reserve price, then the full-trade equilibrium constructed above would remain an equilibrium for small frictions. It directly follows that observing any $b < v$, the seller would accept such

[21] The mass of buyers in state ℓ is roughly d^ℓ, and the mass of sellers is roughly $\left(1-d^\ell\right))/(1-\delta_k)$. In state h, the mass of buyers is roughly $(d^h - s)/(1-\delta_k)$, and the mass of sellers is roughly s.

[22] See (4) to build intuition. In particular, letting $x_k = \beta_k^{-1}(b)$, the density of the highest type at x_k can be calculated upon taking a derivative of (4) to obtain $\gamma^B_{(1),k}(x_k \mid w) = \mu_k(w)\gamma^B_k(x_k \mid w)\, e^{-\mu^w_k(1-\Gamma^B_k(x_k|w))}$. Then $\lim \mu^h_k(1-\delta_k) = A$ for some $A > 0$, and $\lim \mu^\ell_k/(1-\delta_k) = B$ for some $B > 0$ imply that $\lim \gamma^B_{(1),k}(x_k \mid h)/\gamma^B_{(1),k}(x_k \mid \ell)$ is of the order $\frac{e^{-1/(1-\delta_k)}}{(1-\delta_k)^2}$, which converges to zero.

[23] This argument is spelled out in more details in Section 4.2.2.

a highest bid because such a highest bid indicates that the state is likely to be low. The argument can also be extended to any bid close to v.[24]

To prove that the equilibrium satisfies undominatedness (point 2), we need to verify that for any belief $\theta \in (0,1)$ the reserve price set lies between the lowest and the highest possible values of the outside option. For any given δ and θ, this requirement boils down to $\delta V^S(0) \leq \rho(\theta) - c \leq \delta V^S(1)$ as it was observed in (15). Therefore, undominatedness follows for our construction when $\delta_k \to 1$ if for all $r_0 \in (c, v)$,

$$\lim V_k^S(0) \leq r_0 - c \leq \lim V_k^S(1).$$

The fact that these inequalities hold follows from two observations. First, in the high state the expected price converges to v, and the trading probability converges to 1 for the sellers; thus, $\lim V_k^S(1) = v - c$. Second, in the low state the expected price converges to r_0, and the trading probability converges to less than 1 for the sellers; thus, $\lim V_k^S(0) < r_0 - c$.

4.2.2 Learning, Equilibrium Strategies, and the Winner's Curse

In the equilibrium constructed, the buyers and the sellers behave very differently. The sellers set the same reserve price r_0 regardless of how long ago they entered the market. The buyers, however, increase their bids every time they have not traded. A buyer who has entered the market a long enough time ago bids close to v in the full-trade equilibrium constructed.

To explain the difference in the dynamics of the buyers' bidding and sellers' reserve price strategies, we consider two important effects. These two effects are useful for analyzing the buyers' behavior. First, the winner's curse means that a buyer may bid low even if he is almost certain that the state favors the sellers. Second, the loser's curse means that if a buyer has not traded enough times, then he increases his bid and tries to beat the competition. These two opposing effects are described in what follows using the full-trade equilibria for illustration.

Recall that $\theta_+^B(\theta, b)$ is the posterior of a buyer who starts with belief θ and learns that his bid b did not win (either there was a higher bidder or the seller had a higher reserve price set). Let $\gamma_{(1)}^h = \gamma_{(1)}^B(\theta \mid h)$, $\gamma_{(1)}^\ell = \gamma_{(1)}^B(\theta \mid \ell)$ describe the density function of the highest bidder type who participates in the auction in state $\omega = \ell, h$, and $W_+^h = EU^B(\theta_+^B(\theta, \beta(\theta)) \mid h)$, $W_+^\ell = EU^B(\theta_+^B(\theta, \beta(\theta)) \mid \ell)$ the relevant continuation values of type θ upon not winning the

[24] First, any highest bid $b \leq \delta_k v$ can be shown to imply that the state is likely to be low upon a minor modification of the proof. Second, any bid $b > \delta_k v$ is accepted by the seller because the seller's outside option is not more than $\delta_k v$. Combining these observations yields that accepting any highest bid (in the support of equilibrium bids) is in the interest of the seller.

auction. The differential equation, which describes β, can be written[25] as

$$\beta'(\theta) = \frac{\theta\gamma^h_{(1)}[v - \beta(\theta) - \delta W^h_+] + (1-\theta)\gamma^\ell_{(1)}[v - \beta(\theta) - \delta W^\ell_+]}{\theta\Gamma^B_{(1)}(\theta \mid h) + (1-\theta)\Gamma^B_{(1)}(\theta \mid \ell)}.$$

The winner's curse effect can be formally captured by the observation that even if a bidder is very certain that the state is high (θ is close to 1), the bid function may still not respond much (β' is close to zero) if, conditional on θ being the highest type in the auction, the belief $\gamma^h_{(1)}$ remains low. Intuitively, in the pivotal event of tieing at the top, the state can be low with a very high probability even if the prior entering the auction prescribes a very high probability for the high state. This leads to rational bidders depressing their bids even if the bidders attach a very high probability to the high state (and thus a high value of the object). This phenomenon is called the winner's curse in the auction literature. In the context of our dynamic model, the winner's curse implies that buyers may keep bidding low even if they already hold the belief that the state is very likely to be high after they have not been able to trade a large number of times.

We can also describe the loser's curse in full-trade equilibria. In particular, buyers who have not traded for a long enough time know that the state likely favors the sellers and increase their bids. Formally, the loser's curse can be captured by the observation that buyers update their beliefs upon not trading such that $\theta^B_+(\theta, b) \leq \theta$ for any $b \geq r_0$ where r_0 is the reserve price set by all sellers.

The balance of these two effects means that after a sufficient number of losses buyers bid close to v, but in the limit, there are buyers with arbitrarily high beliefs who still bid arbitrarily close to the reserve price r_0 set by the sellers.

The winner's curse is absent for the sellers when the frictions are low (δ_k is high enough) because when the state is high many buyers compete, and thus, the value of the reserve price is not important. In other words, the bidder competition reduces the potential winner's curse problem arising for the sellers. On a more technical level, the pivotal conditioning event for the sellers is that the highest bid is equal to the equilibrium reserve price r_0. In the full-trade equilibrium constructed, this pivotal event indicates that the state is likely to be the low state (when δ_k is high). Therefore, the winner's curse effect does not arise for the sellers, and the sellers are happy to accepts all bids, which are made in equilibrium. Given that the sellers accept all bids from the beginning, the scope for the loser's curse is limited as well. In the equilibrium constructed, the sellers do not decrease their reserve prices over time at all.

[25] See Lemma 17 in the online Appendix.

4.2.3 Limiting Properties of Full-trade Equilibria and Discussion

The above result shows that there are equilibria where the limiting price is non-competitive in the low state. In particular, in every equilibrium the expected price is at least $r_0 > c$ in the low state, while the competitive price is c in the low state. All the other properties of the competitive equilibrium can be shown to hold[26]:

i) *The short side trades for sure in each state in the limit in any full-trade equilibrium.*

This result is immediate in the high state for the sellers who trade every time they meet a buyer, which is with positive probability every period in the limit. The result for buyers in the low state follows from the fact that in each period a buyer is alone with the seller with a positive probability, and thus, he trades with a positive probability each period.

ii) *Almost all trades take place at prices close to v in the high state in the limit.*

We observed above that for any $b < v$, the probability of the low state conditional on the highest bid being equal to b converges to 1 as δ_k goes to 1. This observation implies the result in point ii).

iii) *Almost all trades take place at prices close to r_0 in the low state in the limit.*

Bidding in the low state is very non-competitive: Any bid $r_0 + \varepsilon$ guarantees sure winning in the low state in the limit. Therefore, buyers make sure to bid close to r_0 until they have a belief that attaches a positive probability to the low state. Thus, almost all transactions take place at a price close to r_0 in the low state in the limit.

One of the main results, Proposition 3 shows that any sequence of equilibria, and not just sequences of full-trade equilibria, satisfies properties i) and ii). In addition, there is a price $p^\ell \in [c, v]$ such that property iii) is satisfied with price p^ℓ instead of r_0.

5 Limiting Allocations for All Equilibria

We argued above that for any sequence of *full-trade* equilibria, the limiting allocation is competitive with one exception: The limiting transaction price may be above the competitive price in the low state when there is uncertainty about the competitive allocation ($d^\ell < s < d^h$). In the Appendix, we prove that these results hold for *all* undominated equilibria, and not just for full-trade equilibria:

Proposition 3 *For almost all $\theta^B \in [\underline{\theta}^B, \bar{\theta}^B]$ and $\theta^S \in [\underline{\theta}^S, \bar{\theta}^S]$, it holds that equilibrium trading probabilities are competitive in the limit: If $d^w < s$, then*

$$\lim Q_k^S(\theta^S \mid w) = d^w, \text{ and } \lim Q_k^B(\theta^B \mid w) = 1.$$

[26] Separate proofs are not added, as these results all follow from Proposition 3.

If $d^w > s$, then

$$\lim Q_k^S(\theta^S \mid w) = 1, \text{ and } \lim\ Q_k^B(\theta^B \mid w) = s/d^w.$$

The law of one price holds if $d^w < s$ or if $d^w > s$. If $d^w > s$, then the trading price is competitive in the limit; that is, the trading price converges to v in probability in state w. If $d^w < s$, then consider a subsequence where there exists $p^w \in [c, v]$ such that for some $\theta^B \in (0,1)$, $\lim P_k^B(\theta^B \mid w) = p^w$. Then the transaction price of almost all types θ^B, θ^S converges to p^w in probability. Moreover, if $d^h < s$, then the limit is competitive in both states; that is, $p^h = p^\ell = c$.

The main difficulty is showing that trading probabilities become competitive in all undominated equilibria in the limit; that is, the short side trades almost surely in the limit. First, we discuss the most interesting case where $d^\ell < s < d^h$. For the high state, the key step is to show (see the proof of Lemma 4) that any reserve price posted in a monotone equilibrium allows the seller to trade with a positive probability each period (and thus with probability 1 over the lifetime) in the high state in the limit.[27] The proof of the lemma exploits the undominatedness refinement in a simple manner. Let $\overline{\rho}_k$ be the highest value in the support of the equilibrium reserve prices, and suppose that buyers bidding $\overline{\rho}_k$ would trade for sure (per period) in the limit. Then they would not bid much higher than $\overline{\rho}_k$, and it would be dominated for the sellers to set a reserve price as high as $\overline{\rho}_k$.

The key step for the low state is Lemma 5, which shows that if $\liminf D_k(\ell)/S_k(\ell) > 0$ and $\liminf \beta_k(0) < v$, then in the limit buyers trade for sure in the low state.[28] The proof first shows that there is an atom in the distribution of bids at $\beta_k(0) \approx V_k^S(0) + c$; otherwise, the seller would be able to sell with probability 1 in the low state in the limit at a price above $\lim \beta_k(0)$, which would violate the optimality conditions for the seller. Due to the atom at $\beta_k(0)$, bidding slightly above $\beta_k(0) \approx V_k^S(0) + c$ yields an increase in the lifetime trading probabilities; that is, for all $\varepsilon > 0$, $\limsup(Q_k^B(V_k^S(0) + c + \varepsilon \mid \ell)\ /\ Q_k^B(\beta_k(0) \mid \ell)) > 1$. This implies that overtaking the atom at $\beta_k(0) \approx V_k^S(0) + c$ is profitable for the buyers in the low state, which contradicts the optimality of bid $\beta_k(0)$.

In the other cases, there is no uncertainty about the competitive allocation. When $d^h, d^\ell > s$, both states feature a competitive allocation where the sellers trade for sure at price v. When $d^h, d^\ell < s$, in both states the buyers trade for sure at price c. While all the other results from the case of $d^\ell < s < d^h$ continue to hold, the requirement of undominatedness is now also enough to rule out equilibria with non-competitive limits. Therefore, the limit is competitive in these two cases.

[27] See Lemma 4, where it is easy to rule out $\lim \beta_k(1) = c$, when $d^\ell < s < d^h$.

[28] Since it is easy to rule out $\lim \beta_k(0) = v$ when $d^\ell < s < d^h$, it follows that $\lim D_k(\ell)/S_k(\ell) = 0$, which implies that the buyers trade for sure in the low state in the limit.

The main idea is that in any equilibrium the reserve price r_0 converges to the competitive price in this case. Take the case where $d^h, d^\ell < s$, and suppose that the reserve price set by the sellers r_0 remains higher than c in the limit. It can be shown that in each state almost all transactions take place at a price that converges to the reserve price r_0 because of the lack of buyer competition. Therefore, each seller sells at a price near r_0 with a probability less than 1 in both states in the limit, and the continuation value of the sellers is less than r_0 in both states in the limit. Therefore, it is dominated for the sellers to reject a bid of $r_0 - \varepsilon$; thus, a reserve price of $r_0 > c$ cannot be part of an undominated equilibrium in the limit. Similarly, in the case where $d^h, d^\ell > s$, it is dominated for the sellers to accept any bid below v in the limit; thus a reserve price of $r_0 < v$ cannot be part of an undominated equilibrium in the limit.

This argument not only implies that the limit is competitive in these two cases but also that the equilibria constructed for the case of $d^\ell < s < d^h$ do not pass the test of undominatedness in the other cases. To show that an equilibrium in undominated strategies still exists, we need to modify our construction in Section 4. We provide the modified construction in the online Appendix when we prove Proposition 5. The main novelty is that sellers no longer post the exact same reserve prices, but the property that all bids are accepted in equilibrium still holds. Moreover, the reserve prices set by all sellers converge to c if $d^h, d^\ell < s$, and to v if $d^h, d^\ell > s$.

6 Competitive Limit with Monotone Beliefs

6.1 Off-equilibrium Beliefs and Information Aggregation

Studying the set of full-trade equilibria, we showed that equilibrium trading prices may be non-competitive in the limit if $d^\ell < s < d^h$. The sellers all post a non-competitive reserve price $r_0 > c$ even as the frictions become low. As a result, no buyer bids below r_0 in equilibrium, and reducing the reserve price does not provide any benefits for the sellers. Now suppose that the seller considers the possibility of buyers making mistakes and bidding less than r_0. If it is more likely that such mistakes are made in the low state than in the high state, then reducing the reserve price from r_0 is profitable for the seller, and thus, setting a reserve price r_0 is not rationalizable with such off-equilibrium beliefs.

However, if the seller attaches a high probability to the high state in the off-equilibrium event where the winning bid is less than r_0, then those off-equilibrium beliefs rationalize the seller setting a reserve price of $r_0 > c$. Attaching a high probability to the high state if a bid below r_0 is made in error, and is the winning bid, is counterintuitive. Such a belief implies that the lack of a bid that is higher than r_0 induces sellers to hold beliefs according to which it is likely that many buyers enter each period.

Another reason to believe that such beliefs are counterintuitive is due to Lemma 2. In Lemma 2, we show that for bids in the support of equilibrium such beliefs cannot be sustained. In particular, that Lemma shows that if $b' > b$, and both bids are in equilibrium support, then conditioning on the highest bid being b' induces a higher belief for the seller than conditioning on the highest bid being equal to b.

6.2 Equilibrium Refinement and Competitive Limit

Given the above discussion, we formally introduce a refinement that requires monotonicity of beliefs for off-equilibrium beliefs. We show that under this refinement all sequences of undominated equilibria have competitive limits; that is, they aggregate information efficiently. Recall that $\theta_0^S(\theta, A) = \Pr\left(h|b_{(1)} \in A, \theta\right)$ is the posterior probability of h conditional on the highest bid being in a set $A \subset \mathcal{B}$; that is, it satisfies (7) if the conditioning event $b_{(1)} \in A$ has a positive probability. We adopt the following refinement:

Refinement of Monotone Beliefs. *An equilibrium satisfies the refinement of monotone beliefs if there exists a belief system $\{\theta_0^S(\theta, \cdot)\}$ such that*
i) $\theta_0^S(\theta, [b_1, b_2])$ is weakly increasing in b_1 and b_2;
ii) for all $b_1 < b_2$ it holds that if $\rho(\theta) = b_2$, then

$$\delta V^S(\theta_0^S(\theta, [b_1, b_2])) \geq b_1 - c. \tag{16}$$

Condition i) states that beliefs need to satisfy an intuitive monotonicity refinement, as discussed above. Condition ii) states that if a seller sets a reserve price of b_2, then decreasing his reserve price to $b_1 < b_2$ is not profitable, given his beliefs $\theta_0^S(\theta, \cdot)$. To see this, suppose that the seller switches his strategy from reserve price b_2 to reserve price b_1. If the highest bid is not on interval $[b_1, b_2]$, then there is no change in the seller's payoff. Therefore, when changing his strategy, the seller can assume that the highest bid belongs to interval $[b_1, b_2]$, and the seller's belief conditional on this event is $\theta_0^S(\theta, [b_1, b_2])$. If he does not accept bids on $[b_1, b_2]$, then his continuation utility is $\delta V^S(\theta_0^S(\theta, [b_1, b_2]))$. If he accepts those bids, then his profit conditional on $b_{(1)} \in [b_1, b_2]$ is at least b_1, as all the bids are at least b_1.[29] Given this discussion, (16) means that taking the belief θ_0^S as given, the seller does not reject all bids on $[b_1, b_2]$ if doing so yields a lower continuation payoff than the worst outcome (b_1) that can be guaranteed by accepting those bids.

[29] The actual profit conditional on this event is

$$E[b_{(1)} \mid b_{(1)} \in [b_1, b_2], \theta],$$

that is, the expected value of the highest bid given the prior θ, and the event that $b_{(1)} \in [b_1, b_2]$.

The equilibria featured in Proposition 2 do not satisfy the refinement.[30] The reason is that as we discussed, if $b_1, b_2 \to r_0$, then $\theta_0^S(\theta, [b_1, b_2]) \to 0$. Given that in the limit trade occurs with probability d^ℓ in the low state at a price close to r_0, it follows that $\delta V^S(\theta_0^S(\theta, [b_1, b_2])) \to d^\ell(r_0 - c)$ and thus, the refinement fails because the right-hand side of (16) converges to $r_0 - c > d^\ell(r_0 - c)$.

We are ready to prove our main result: Under the refinement of monotone beliefs, all equilibria have competitive limits.

Proposition 4 *The allocation provided by any sequence of equilibria that satisfy the refinement of monotone beliefs converges to the competitive limit.*

Proof. See Appendix 2.

Our convergence proof exploits the monotonicity refinement to identify a bid b where for all $\theta < 1$ it holds that $\lim \theta_{0,k}^S(\theta, b) = 0$. We prove the following important result for this case:

Lemma 1 *Take any sequence of equilibria that satisfy the refinement of monotone beliefs and suppose that* $\lim \theta_{0,k}^S(\theta_k, b_k) = 0$ *for some sequence* $b_k \in supp\beta_k$, *and* $b_k = \rho_k(\theta_k)$ *and* $b_k \to b \in [c, v]$. *Then*

$$\lim V_k^S(0) + c \geq b.$$

Proof. See Appendix 2 for the proof.

This result states that if a sequence of winning bids b_k implies that the state is almost surely low in the limit, then the seller sets such reserve prices in the limit only if that reserve price is not higher than the outside option in the low state in the limit. In particular, it is not the case that the seller sets a higher reserve price than the outside option in the low state even if he believes that the state is very likely to be low conditional on the highest bid being equal to the reserve price.

The other key step of the proof of Proposition 4 (Step 1) is that all equilibrium bids below p^ℓ (as defined in Proposition 3) make the sellers believe that the state is almost surely low in the limit because such a bid comes from a buyer with a low belief. Therefore, any such winning bid *in the support* of equilibrium bids signals that the state is very likely to be ℓ. Therefore, Lemma 1 implies that no seller sets a reserve $\rho \in (\lim V_k^S(0) + c, p^\ell)$ for a high enough k. Then it can be shown (see Steps 2 and 3 of the proof) that bidding $V_k^S(0) + c + \varepsilon$ is

[30] In the proof of Proposition 5, we construct an equilibrium that satisfies the refinement.

better than bidding p^ℓ in the limit, which contradicts with how p^ℓ was defined in Proposition 3.

Without the refinement, it would not be true that *any* bid on $(\lim V_k^S(0)+c, p^\ell)$ is accepted for a high enough k. Although all such low bids in the support of the equilibrium indicate that the state is likely to be low, the seller may believe that the state is likely to be high if the winning bid is below p^ℓ but is *not* in the support of the equilibrium bid distribution.[31]

6.3 Supporting Technical Results

6.3.1 Monotone Equilibrium Updating

To motivate the refinement of monotone beliefs, we continue by showing that $\theta_0^S(\theta, b) = \lim_{\varepsilon\to 0} \theta_0^S(\theta, [b, b+\varepsilon])$ is increasing in b on the equilibrium path. This result implies that a higher winning bid in the support of equilibrium bids makes the high state more likely.

Lemma 2 *Assume that $d^h > s > d^\ell$. There exists a $\underline{\delta} < 1$ such that for all $\delta > \underline{\delta}$ the following holds. For all $b' > b \geq c$ such that $b, b' \in supp(\beta)$,*

$$\theta_0^S(\theta, b) > \theta_0^S(\theta, b'). \tag{17}$$

Moreover, let $\theta_0^S(\theta, \varnothing)$ be the posterior if no bid is received. Then for all $b \geq c$ such that $b \in supp(\beta)$

$$\theta_0^S(\theta, b) \geq \theta_0^S(\theta, \varnothing). \tag{18}$$

Let us discuss the above Lemma briefly. By Proposition 3, trading probabilities converge to the competitive trading probabilities if $d^\ell < s < d^h$. Therefore, in the limit all monotone equilibria in undominated strategies feature equilibrium stocks that are competitive, and thus, $S(\ell) > S(h)$, $D(h) > D(\ell)$ for low enough frictions. Using this result, Lemma 2 shows that $\theta_0^S(\theta, b)$ is monotone in b for all bids that are in the support of equilibrium strategies. The intuition for this is that if $S(\ell) > S(h)$, $D(h) > D(\ell)$, then there is more competition in the high state than in the low state; thus, the winning bid tends to be higher in the high state than in the low state.[32]

6.3.2 Equilibrium Existence

In the online Appendix, we prove that an equilibrium satisfying the refinement exists:

[31] In the equilibria constructed in Section 4.2.1 with $r_0 > c$, such bids below $p^\ell = r_0$ would all be rejected by the sellers (if made).

[32] Although we conjecture that this result holds for the other cases, we have not been able to establish it when frictions are high or if $d^\ell < s < d^h$ does not hold. The main complication is that given the endogenous nature of equilibrium stocks, it is difficult to characterize stocks away from the limit.

Proposition 5 *For a threshold $\underline{\delta} < 1$, there exists a full-trade equilibrium for all $\delta_k > \underline{\delta}$, which satisfies the refinement of monotone beliefs.*

Proposition 2 already provided a construction such that the sellers do not find it optimal to reject any bids made in equilibrium. For Proposition 5, we modify that construction to make sure that for some off-equilibrium beliefs that satisfy the monotonicity refinement, the seller does not have an incentive to *decrease* his reserve price below the equilibrium reserve price. We adopt off-equilibrium beliefs for the sellers that prescribe probability one to the lowest bidder type $\underline{\theta}^B$, if the bid is less than $\beta_k(\underline{\theta}^B)$. Then the key idea is to require that the seller with the highest belief $(\bar{\theta}^S)$ is exactly indifferent between accepting bid $\beta_k(\underline{\theta}^B)$ or rejecting such a bid, and such a seller sets a reserve price $\rho_k(\bar{\theta}^S) = \beta_k(\underline{\theta}^B)$. This indifference condition pins down a sequence of reserve prices $\{\rho_k(\bar{\theta}^S)\}$.[33] Sellers with beliefs below $\bar{\theta}^S$ strictly prefer accepting a bid $\beta_k(\underline{\theta}^B)$, and thus, to satisfy the refinement, they need to set a reserve price strictly below $\beta_k(\underline{\theta}^B)$.

In the most interesting case where $d^h > s > d^\ell$, the indifference condition implies that $\lim \rho_k(\bar{\theta}^S) = c$, so when there is uncertainty about the competitive price, all sellers set a reserve price close to c in the limit. It follows from Proposition 3 and the discussion afterward that in any equilibrium that satisfies undominatedness (but not necessarily the refinement of monotone beliefs) the equilibrium reserve prices already converge to c if $d^h, d^\ell < s$ and to v if $d^h, d^\ell > s$. The construction for Proposition 5 is very easy to verify in this case.

7 Discussion

In the previous section, we used the refinement of monotone beliefs to obtain that all sequences of equilibria have competitive limits under this refinement. In this section, we drop this refinement and revisit equilibria that feature non-competitive limits. First, we discuss what features of the matching and bargaining protocol are conducive to information aggregation in our model. Second, we show that such equilibria feature fast concessions from the sellers and, thus, are similar to full-trade equilibria.

7.1 Matching and Bargaining Protocols, Aggregate Uncertainty, and Non-competitive Limits

The analysis in Section 6 implies that under full information about market conditions, any undominated equilibrium has a competitive limit. Therefore, in Proposition 3 the possibility of non-competitive limits arises due to the combination of incomplete information about the

[33] We believe that the construction identifies the unique full-trade equilibrium, which satisfies the refinement of monotone beliefs.

economy and the particular matching and bargaining protocol. It is then natural to ask whether maintaining the assumption of incomplete information about the economy alternative protocols would deliver competitive limits, and thus, information aggregation. We start the discussion by observing that in our model the buyers compete directly by bidding while there is no such competition between the sellers. This helps explain why the price is competitive when frictions are low when there are many buyers but the price may stay above the competitive level when there are more sellers than buyers. In particular, the sellers do not compete away their non-competitive rents as they face no direct competition.

In this section, we argue that when both buyers and sellers face direct competition, the price becomes competitive in both states. Moreover, the amount of direct competition does not have to be large in absolute value, only large compared to the level of frictions. First, consider a modification where with probability ε each buyer is matched with two sellers. Then the sellers compete with each other, and it is our strong conjecture that the limit must be competitive even in the low state in this case. The reason is that each seller has an incentive to undercut "rival" seller(s) to make sure that the buyer chooses his offer. This is sufficient to force the reserve prices down to the cost of the sellers when there are more sellers than buyers in the economy.

Second, consider a setting where buyer competition is limited in that for any given seller at most two bids go through for the seller to choose from. We conjecture that this more limited buyer competition is sufficient for the price to be competitive in the high state as buyers face a Bertrand-type head-on competition with probability one in the high state in the limit. Therefore, a competitive limit arises as long as there is some but not necessarily overly fierce competition to trade with an agent on the short side. However, the popular bilateral matching protocol will not necessarily deliver a competitive price in the limit, as all direct competition is eliminated in this protocol.

7.2 Belief Updating and Trading Strategies in Calendar Time

So far, we considered the market outcome as the exit rate becomes small. Now, we consider an equivalent interpretation. We fix the exit probability per unit time and increase the speed of trading. This reinterpretation makes it easier to interpret how quickly agents can trade in the limit. Formally, fix some $r > 0$, and suppose that the probability of not exiting the market is $\delta_k = e^{-r\Delta_k}$ where Δ_k is the length of time between periods of trading. We consider $\Delta_k \to 0$, so that trading becomes fast and $\delta_k \to 1$. Take any even sequence of events that take place in period τ_k and assume that $\lim (\delta_k)^{\tau_k} = 1$. This implies that $\lim \Delta_k \tau_k = 0$, which means that the calendar time that elapses before period τ_k converges to zero.

We consider how beliefs and bids or reserve prices change in calendar time after entry.

First, belief updating happens quickly in calendar time in the limit. Proposition 4 implies that almost all buyers trade in zero calendar time in the low state in the limit. Consequently, almost all entering buyers update to a belief close to one in zero calendar time in the limit. Similarly, Proposition 4 implies that almost all sellers trade in zero calendar time in the high state in the limit, and thus, almost all entering sellers update to a belief close to zero in zero calendar time in the limit.

Now, we characterize equilibria with non-competitive limits and show that in all equilibria the sellers concede quickly as frictions vanish.

Proposition 6 *Assume that $d^\ell < s < d^h$. For any sequence of undominated equilibria with $p^\ell > c$, there exists a positive integer t_k with $\lim t_k(1-\delta_k)^\alpha = 0$ for all $\alpha > 0$, such that all sellers who entered at least t_k periods ago set a reserve price r that satisfies*

$$\lim \Pr(b_{(1),k} < r \mid b_{(1),k} \geq c, \ell) = 0.$$

Proof. See Appendix 2.

With the interpretation of calendar time as set out above, Proposition 6 means that the sellers match almost all equilibrium bids in zero calendar time in the limit if the sequence of equilibria does not converge to the competitive allocation (that is, if $p^\ell > c$). From a more symmetric point of view, all buyers and sellers choose trading strategies that concede quickly.[34]

When the limit is competitive, the sellers may not concede quickly (in zero calendar time) in the limit. There is a possibility for a war of attrition in which the sellers set high reserve prices for a positive calendar time, and the buyers bid close to $p^\ell = c$ for a long time as well. In this case, the sellers may not have incentives to accept those low bids because those offers are close to the sellers' outside options, $\lim V_k^S(0) + c = p^\ell = c$ in the limit.[35]

8 Conclusion

We study search markets where the market conditions are not known by the agents. We provide three contributions to the literature. First, we study the combined effects of search

[34] It does not directly follow from Proposition 3 that both sides concede quickly. Proposition 3 would allow that sellers do not trade in the low state in the limit for a positive amount of calendar time, while buyers concede quickly.

[35] In particular, there may exist an equilibrium where sellers set a reserve $r \in (c, v)$ for t_k periods with $\lim \delta_k^{t_k} < 1$ and then set a low reserve price τ_k such that $\lim \tau_k = c$. However, if we focus on equilibria that satisfy the refinement of monotone beliefs, then the proof of Proposition 6 implies that the pivotal belief $\theta_{k,0}^S$ converges to zero in zero calendar time. Then the monotonicity refinement implies that a reserve price close to $c + \lim V_k^S(0) = c$ is set in zero calendar time in the limit. In this sense, the sellers concede quickly in the limit.

and learning in a novel two-sided equilibrium model in order to capture realistic uncertainty about market conditions. Second, we identify tractable "full trade equilibria" that provide insights into the economics of decentralized markets. In particular, we study bidding behavior and learning over time in a search market with aggregate uncertainty. Third, we test the hypothesis that even if decentralized, markets can nevertheless aggregate information that is dispersed among its participants. When frictions are small, trade takes place at the correct market clearing prices "as if" the demand and supply conditions were known to the participants.

9 Appendix A

9.1 Proof of Proposition 3

Restatement of Proposition 3. *For almost all $\theta^B \in [\underline{\theta}^B, \bar{\theta}^B]$ and $\theta^S \in [\underline{\theta}^S, \bar{\theta}^S]$, it holds that equilibrium trading probabilities are competitive in the limit: If $d^w < s$, then*

$$\lim Q_k^S(\theta^S \mid w) = d^w, \text{ and } \lim Q_k^B(\theta^B \mid w) = 1.$$

If $d^w > s$, then

$$\lim Q_k^S(\theta^S \mid w) = 1, \text{ and } \lim Q_k^B(\theta^B \mid w) = s/d^w.$$

The law of one price holds if $d^w < s$ or if $d^w > s$. If $d^w > s$, then the trading price is competitive in the limit, that is, it converges to v in probability in state w. If $d^w < s$, then consider a subsequence where there exists, $p^w \in [c, v]$ such that for some $\theta^B \in (0, 1)$, $\lim P_k^B(\theta^B \mid w) = p^w$. Then the transaction price of almost all types θ^B, θ^S converges to p^w in probability. Moreover, if $d^h < s$ then the limit is competitive in both states, that is $p^h = p^\ell = c$.

We start the proof with the following Lemma, which captures the fact that experimentation becomes cheap when frictions vanish. For simplicity, we only state the result for the sellers' side, a similar result hold for the buyers.

Lemma 3 *Take any equilibrium in monotone strategies. If for some $\theta > 0$ it holds that $\lim Q_k^S(\theta \mid h) = 1$, then the realized transaction price of a seller with belief θ converges in distribution to a single price $p^h(\theta) \in [c, v]$ in the high state. If $\lim Q_k^S(\theta' \mid h) = \lim Q_k^S(\theta'' \mid h) = 1$ for some $1 \geq \theta', \theta'' > 0$, then $p^h(\theta') = p^h(\theta'') = p^h$. Moreover, if $\lim Q_k^S(\theta \mid h) = 1$ for some $\theta < 1$, then $\lim W_k^S(\theta \mid \ell) = \lim V_k^S(0)$.*

Proof. Fix $\theta, \theta' > 0$ such that $\lim Q_k^S(\theta \mid h) = \lim Q_k^S(\theta' \mid h) = 1$ and let $p^h(\theta) = \lim P_k^S(\theta \mid h)$.[36]

[36] If the original price sequence P_k^S did not converge, then the argument is applied to a converging subsequence without loss of generality.

Step 1. We show that the actual trading price of type θ converges to $p^h(\theta)$ in probability in the high state.

Given $\lim Q_k^S(\theta \mid h) = 1$, it follows that there exists a sequence t_k with $\lim \delta_k^{t_k} = 1$ such that the seller trades almost surely by period t_k in the high state in the limit if he employs his equilibrium strategy. Let $\chi_k(p)$ denote the probability of selling at price p or above by period t_k if the equilibrium strategy is used by type θ. Suppose that for some $p > p^h(\theta)$, $\lim \chi_k(p) > 0$, otherwise convergence to $p^h(\theta)$ in probability follows.[37] We propose a strategy $\Sigma(p)$ with two properties:

P1. the seller trades at or above price p with a probability converging to 1 in the high state,

P2. the payoff achieved in the low state converges to the (full information) optimal payoff in the low state, $\lim V_k^S(0)$.

The strategy $\Sigma(p)$ has the following form: the seller sets a reserve price p for z_k periods, where $z_k = \alpha_k t_k$. After z_k periods, if the seller has not traded he adopts a strategy that is optimal in the low state. We require that $\lim \alpha_k = \infty$, and that $\lim \delta_k^{z_k} = 1$.[38] This strategy satisfies properties P1 and P2. To see that P1 is satisfied by $\Sigma(p)$, we show that the probability of trading by period z_k (at or above price p) in the high state converges to one. Let q_k denote the probability of trading in the high state by period t_k if $\Sigma(p)$ is used. It holds that $q_k \geq \chi_k(p)$,[39] and thus $\lim \chi_k(p) > 0$ implies $q = \lim q_k > 0$. Using, that $\lim \delta_k^{z_k} = 1$, the trading probability offered by strategy $\Sigma(p)$ by period $z_k = \alpha_k t_k$ in the high state in the limit is

$$\lim q + (1-q)q + ...(1-q)^{\alpha_k - 1} q = \lim \frac{q(1-q^{\alpha_k})}{1-(1-q)} = 1.$$

To see that P2 is satisfied, note that the exogenous exit probability by period z_k converges to zero as $\lim \delta_k^{z_k} = 1$. Therefore, the seller drops out with a zero probability in the limit before he switches to a strategy that is optimal in the low state (at period z_k), and thus his payoff in the low state converges to the full information limiting payoff $\lim V_k^S(0)$.[40]

It is immediate from P1 and P2 that the original strategy was not optimal, a contradiction with our starting assumption.

---

[37] Because then the seller trades almost surely at the expected price p^h in the limit as the mean trading price is p^h.

[38] Formally, $\lim \delta_k^{t_k} = 1$ is equivalent to $\lim t_k(1-\delta_k) = 0$. We need to have $\lim \delta_k^{z_k} = 1$ or $\lim z_k(1-\delta_k) = 0$. Let $(z_k(1-\delta_k))^2 = t_k(1-\delta_k)$, which clearly implies that $\lim z_k(1-\delta_k) = 0$ if $\lim t_k(1-\delta_k) = 0$. Then $\alpha_k^2 t_k^2 (1-\delta_k)^2 = t_k(1-\delta_k)$ or $\alpha_k^2 t_k (1-\delta_k) = 1$. Then $\lim t_k(1-\delta_k) = 0$ implies that $\lim \alpha_k = \infty$, which then provides the appropriate construction.

[39] This holds because if an arbitrary reserve price strategy trades at a price p or above in a period, then strategy $\Sigma(p)$ trades at a price p or above in that same period or before.

[40] This argument assumes that the trading probability by period z_k stays zero in the low state in the limit. But otherwise, a sure trading probability (over the entire lifetime) at price p or above could be guaranteed in the low state. But then in both states a revenue of p or above could be guaranteed and thus setting a reserve below p and achieving a revenue of p^h cannot be optimal in the high state.

Step 2: Step 1 implies that a price $p^h(\theta')$ is achieved almost surely by type θ' in the equilibrium in the high state if $\lim Q_k^S(\theta' \mid h) = 1$. If $p^h(\theta') < p^h(\theta)$ then type θ' could adopt strategy $\Sigma(p^h(\theta)-\varepsilon)$ for ε arbitrarily small, which would improve his payoff in state h to $p^h(\theta)$, and attain the optimal payoff in state ℓ.[41] Therefore, $p^h(\theta') = p^h(\theta)$ must hold in equilibrium. Finally, the full information payoff can be achieved by such a strategy $\Sigma(p^h(\theta) - \varepsilon)$ in the low state in the limit (as outlined in Step 1) regardless of the value of $p^h(\theta)$ and $\varepsilon > 0$, which concludes the proof. *Q.E.D.*

Lemma 4 *If* $\lim \beta_k(1) > c$, *then*

$$\lim Q_k^S(1|h) = 1.$$

Proof. Suppose $\lim Q_k^S(1|h) < 1$. Then, $\lim 1 - q_k^S(\rho_k(1)|h) = \lim q_k^B(\rho_k(1)|h) = 1$ (almost no buyer bids above $\rho_k(1)$). Hence, $\lim Q_k^B(\rho_k(1)|h) = 1$, which implies $\lim \beta_k(1) \leq \lim \rho_k(1)$. From undominatedness (14) and $V_k^S(1) \leq \beta_k(1) - c$, we have $\rho_k(1) - c \leq \delta_k(\beta_k(1) - c)$ and hence $\rho_k \leq \beta_k(1)$ for all all k. Hence, $\lim \beta_k(1) = \lim \rho_k(1)$. It follows from monotonicity of β_k that $\lim P_k^S(1|h) = \lim \rho_k(1)$. From the definition of V_k^S and (14), for all k large enough,

$$\rho_k(1) - c = \delta_k V_k^S(1) = \delta_k Q_k^S(1|h)\left(P_k^S(1|h) - c\right).$$

From the displayed equation, the hypothesis $\lim Q_k^S(1|h) < 1$, and the previous observation $\lim P_k^S(1|h) = \lim \rho_k(1)$, it follows that $\lim \rho_k(1) = c$. Since we already showed $\lim \beta_k(1) = \lim \rho_k(1)$, the hypothesis $\lim Q_k^S(1|h) < 1$ implies that $\lim \beta_k(1) = c$. Hence, if $\lim \beta_k(1) > c$, then $\lim Q_k^S(1|h) = 1$, as claimed. *Q.E.D.*

Lemma 5 *Suppose that* $\liminf D_k(\ell)/S_k(\ell) > 0$. *Then, either* $\lim Q_k^B(\beta_k(0) \mid \ell) = 1$ *or* $\lim \beta_k(0) = v$ *(or both).*

Proof. Suppose that $\liminf D_k(\ell)/S_k(\ell) > 0$. We show that

$$\liminf Q_k^B(\beta_k(0) \mid \ell) < 1 \Rightarrow \lim \beta_k(0) = v,$$

which proves the lemma.

First, $\beta_k(0) \geq \rho_k(0)$. Otherwise, $V_k^B(0) = 0$, in contradiction to $\rho_k(1) - c \leq \delta_k(v - c)$ implying strictly positive profits when bidding $b = \rho_k(1)$. Hence, by monotonicity of β_k and the hypothesis $\liminf D_k(\ell)/S_k(\ell) = \liminf \mu_k^\ell > 0$, we have

$$\lim q_k^S(\rho_k(0)|\ell) = \lim 1 - e^{-\mu_k^\ell} > 0,$$

[41] Again, the argument assumes that the trading probability by period z_k stays zero in the low state in the limit. This assumtpion is without loss of generality, see the footnote above for further discussion.

which implies that

$$\lim Q_k^S\left(\theta=0|\ell\right)=1. \tag{19}$$

Let $p^\ell := \lim P_k^S\left(0|\ell\right)$. From (14) and (19), $\lim \rho_k\left(0\right)=p^\ell$. Of course,

$$\lim Q_k^S\left(\rho_k\left(0\right)+\varepsilon|\ell\right)<1, \tag{20}$$

for any $\varepsilon>0$, by optimality of setting reserve price $\rho_k\left(0\right)$ (if $\rho_k\left(0\right)\to v$ such ε does not exist but then the lemma follows directly from $\beta_k\left(0\right)\geq\rho_k\left(0\right)$). We have

$$\lim q_k^S\left(b_k+\varepsilon|\ell\right)=\lim 1-e^{-\mu_k^\ell\left(1-\Gamma_k^B\left(\beta^{-1}(b_k)|\ell\right)\right)}=0,$$

where the first equality follows from definition of q_k^S and the second follows from (20). So, $\lim e^{-\mu_k^\ell}<1$ and $\lim e^{-\mu_k^\ell\left(1-\Gamma_k^B\left(\beta^{-1}(b_k)|\ell\right)\right)}=1$. Hence,

$$\begin{aligned}\lim\frac{q_k^B\left(b_k+\varepsilon|\ell\right)}{q_k^B\left(b_k|\ell\right)}:=\alpha &= \lim\frac{e^{-\mu_k^\ell\left(1-\Gamma_k^B\left(\beta^{-1}(b_k)|\ell\right)\right)}\Gamma_k^S\left(\rho_k^{-1}\left(b_k+\varepsilon\right)\mid\ell\right)}{e^{-\mu_k^\ell}\Gamma_k^S\left(\rho_k^{-1}\left(b_k\right)\mid\ell\right)}\\ &\geq \lim\frac{e^{-\mu_k^\ell\left(1-\Gamma_k^B\left(\beta^{-1}(b_k)|\ell\right)\right)}}{e^{-\mu_k^\ell}}>1,\end{aligned} \tag{21}$$

where the first equality is from the definition of q_k^B, the first inequality from $\Gamma_k^S\left(\rho_k^{-1}\left(b_k\right)\mid\ell\right)\leq\Gamma_k^S\left(\rho_k^{-1}\left(b_k+\varepsilon\right)\mid\ell\right)$ and the second inequality from the previous findings. Recall

$$Q_k^B\left(b_k|\ell\right)=\frac{\frac{q_k^B(b_k|\ell)}{1-\delta_k}}{1+\delta_k\frac{q_k^B(b_k|\ell)}{1-\delta_k}}$$

and so the hypothesis $\lim Q_k^B(\beta_k(0)\mid\ell)<1$ implies $\lim\frac{q_k^B(b_k|\ell)}{1-\delta_k}:=z<1$. Therefore,

$$\frac{Q_k^B\left(b_k+\varepsilon|\ell\right)}{Q_k^B\left(b_k|\ell\right)}=\frac{q_k^B\left(b_k+\varepsilon|\ell\right)}{q_k^B\left(b_k|\ell\right)}\frac{1+\delta_k\frac{q_k^B(b_k|\ell)}{1-\delta_k}}{1+\delta_k\frac{q_k^B(b_k+\varepsilon|\ell)}{1-\delta_k}}\to\alpha\frac{1+z}{1+z\alpha}>1.$$

Thus, if $\lim\beta_k(0)<v$ we can choose ε small enough such that the ratio of the profits at $b_k+\varepsilon$ and b_k, respectively, satisfies

$$\lim\frac{Q_k^B\left(b_k+\varepsilon|\ell\right)}{Q_k^B\left(b_k|\ell\right)}\frac{v-b_k-\varepsilon}{v-b_k}>0. \tag{22}$$

Hence, assuming $\lim\beta_k(0)<v$ implies a contradiction to the optimality of $b_k=\beta_k\left(0\right)$. *Q.E.D.*

Remark: The critical observation for the Lemma is that the following holds. If $\lim D_k(\ell)/S_k(\ell)>0$, then either $\lim Q_k^B(\beta_k(0)\mid\ell)=1$ or $\lim\beta_k\left(0\right)=v$. The argument is that when $\lim D_k(\ell)/S_k(\ell)>0$, then (i) there would have to be a mass point at p^ℓ and (ii) a buyer having type 0 would have an incentive to overbid that mass point, unless either $\lim Q_k^B(\beta_k(0) \mid \ell) = 1$ (she is trading already for sure) or $\lim \beta_k(0) = v$ (increasing the trading probability does not increase payoffs).

Lemma 6 *Suppose $d^\ell < s < d^h$. Then: Trading probabilities are competitive and the law-of-one-price holds, with trade taking place at prices p^ℓ and p^h. For almost all $\theta^B \in \left[\underline{\theta}^B, \bar{\theta}^B\right]$, and $\theta^S \in \left[\underline{\theta}^S, \bar{\theta}^S\right]$ the payoffs are*

$$\begin{aligned} \lim EU^S\left(\theta^S|w\right) &= \frac{\min\{s, d^w\}}{s}\left(p^w - c\right), \\ \lim EU^B\left(\theta^B|w\right) &= \frac{\min\{s, d^w\}}{d^w}\left(v - p^w\right). \end{aligned}$$

Proof. Consider $w = h$.

Suppose $\lim \beta_k(1) > c$. Then, the monotonicity of ρ_k and Lemma 4 imply that $\lim Q_k^S(\theta|h) = 1$ for all $\theta \in \left[\underline{\theta}^S, \bar{\theta}^S\right]$, as claimed. Suppose $\lim \beta_k(1) = c$. We show that this implies a contradiction. As in the proof of Lemma 4, $\beta_k(1) \geq \rho_k(1)$, and so $\lim \rho_k(1) = c$. But if $\lim \rho_k(1) = \lim \beta_k(1) = c$, then the monotonicity of ρ_k and β_k implies that given any ε, a bid $c+\varepsilon$ wins for sure when k is large enough. Thus, $\lim V_k^B(\theta) = v - c$ for all θ. This requires $\lim Q_k^B(\theta|h) = 1$ for almost all $\in \left[\underline{\theta}^B, \bar{\theta}^B\right]$—in contradiction to mass balance given $d^h > s$. Finally, mass balance requires that $\lim d^h \int_{\underline{\theta}^B}^{\bar{\theta}} Q_k^B(\theta|h)\, d\theta = s$. Thus, trading probabilities are competitive if $w = h$.

Take some $\theta \in \left[\underline{\theta}^S, \bar{\theta}^S\right]$ with $\lim Q_k^S(\theta|h) = 1$ and let $p^h := \lim P_k^S(\theta|h)$. Then, almost all sellers trade at p^h by $\lim Q_k^S(\theta|h) = 1$ for almost all $\theta \in \left[\underline{\theta}^S, \bar{\theta}^S\right]$ and Lemma 3 (the distribution of trading prices collapses). Thus, the law of one price holds if $w = h$.

Consider $w = \ell$.

Case 1: $\limsup D_k(\ell)/S_k(\ell) > 0$. Then, from Lemma 5, either $\lim Q_k^B(\beta_k(0) \mid \ell) = 1$ or $\lim \beta_k(0) = v$ (or both). If $\lim Q_k^B(\beta_k(0) \mid \ell) = 1$, then by monotonicity of β_k, $\lim Q_k^B(\theta \mid \ell) = 1$ for all θ, so the law-of-one-price holds. If $\lim \beta_k(0) = v$, then this and $\liminf D_k(\ell)/S_k(\ell) > 0$ implies that sellers trade for sure when setting $r = v - \varepsilon$ for any $\varepsilon > 0$, $\lim Q_k^S(v - \varepsilon|w) = 1$ for all ε and $w \in \{\ell, h\}$. Thus, $\lim V_k^S(\theta) = v - c$ for all θ. But this would require $\lim Q_k^S(\theta) = 1$ for all θ, violating mass balance given $d^\ell < s$. Thus, trading probabilities are competitive in case 1.

Case 2: $\limsup D_k(\ell)/S_k(\ell) = 0$. Then, $\lim Q_k^B(\theta \mid \ell) = 1$ for almost all θ. To see why, suppose otherwise and suppose that a positive fraction ϕ of the entering cohort has belief

$\theta^B \in [\underline{\theta}^B, \bar{\theta}^B]$ such that $\liminf Q_k^B(\theta^B \mid \ell) < 1$. Then there exists a sequence t_k with $\liminf \delta_k^{t_k} < 1$ such that in a steady state equilibrium there is a positive mass of buyers still on the market who entered at least t_k periods ago. Noting that $\liminf \delta_k^{t_k} < 1$ is equivalent to $\liminf t_k(1-\delta_k) > 0$, we obtain that $\limsup D_k(\ell)(1-\delta_k) > 0$. Contradiction. Thus, $\lim Q_k^B(\theta \mid \ell) = 1$ for almost all θ in case 2. Therefore, trading probabilities are competitive in both possible cases.

Take some $\theta \in [\underline{\theta}^B, \bar{\theta}]$ with $\lim Q_k^B(\theta|\ell) = 1$ and let $p^\ell := \lim P_k^\ell(\theta|\ell)$. Then, almost all buyers trade at p^ℓ by Lemma 3. Thus, the law of one price holds if $w = \ell$.

Characterization of Payoffs. Consider $w = h$. From competitive trading probabilities and the law of one price, sellers' expected payoffs for almost all $\theta \in [\underline{\theta}^S, \bar{\theta}]$ are $\lim EU^S(\theta|h) = (p^h - c)$. Consider a buyer having type $\theta \in [\underline{\theta}^B, \bar{\theta}]$ such that $\lim Q_k^B(\theta|\ell) = 1$. From Lemma 3 and $\lim Q_k^B(\theta|\ell) = 1$ for almost all $\theta \in [\underline{\theta}^B, \bar{\theta}]$, we have $\lim EU^B(\theta|h) = \lim V_k^B(1)$ for almost all $\theta \in [\underline{\theta}^B, \bar{\theta}]$. Moreover, for almost all θ, the price conditional on trading is p^ℓ. Take some θ with $\lim EU^B(\theta|h) = \lim V_k^B(1)$ and $P_k^B(\theta) \to p^\ell$ and let $\bar{Q} = \lim Q_k^B(\theta|h)$. Then, $\lim V_k^B(1) = \bar{Q}(v - p^\ell)$. Thus, for almost all θ, we have $\bar{Q} = \lim Q_k^B(\theta|h)$. Finally, from the mass balance requirement $\lim d^h \int_{\underline{\theta}^B}^{\bar{\theta}} Q_k^B(\theta|h)\, d\theta = s$, we have $\lim Q_k^B(\theta|h) = \frac{d^h}{s}$ for almost all θ, proving $\lim EU^B(\theta|h) = \frac{d^h}{s}(v - p^h)$.

An analogous argument establishes the same for $w = \ell$. Together, this proves the characterization of payoffs.

Q.E.D.

Proof of Proposition 3.

Proof of Proposition 3 for the case where $d^\ell < s < d^h$.

Lemma 6 proved the law of one price and that trading probabilities are competitive. It remains to show that $p^h = v$, with p^h as defined in Lemma 6.

Suppose, $\lim \beta_k(1) = c$. As argued before in the proof of Lemma 6, if $\lim \beta_k(1) = c$ then $\lim V_k^B(\theta) = v - c$ for all θ. This requires $\lim Q_k^B(\theta|w) = 1$ for almost all $\theta \in [\underline{\theta}^B, \bar{\theta}]$ and $w \in \{\ell, h\}$. This violates mass balance since $s < d^h$. Thus, $\lim \beta_k(1) > c$. Then, Lemma 4 implies that $\lim Q_k^S(1|h) = 1$. So, Lemma 3 requires $P_k^S(1|h) = P_k^S(\theta|h)$ for all θ and from Lemma 6, $P_k^S(1|h) = p^h$. Thus, by Lemma 3, the probability that some buyer bids higher than $p^h + \varepsilon$ must vanish for every ε. In addition, $\lim \rho_k(1) = p^h$ from (14), $P_k^S(1|h) \to p^h$ and $\lim Q_k^S(1|h) = 1$. Hence, a bid $p^h + \varepsilon$ wins with probability converging to one for every $\varepsilon > 0$. Thus, $\lim V_k^B(1) \geq v - p^h$. From Lemma 6 and its proof, $\lim V_k^B(1) \leq \frac{d^h}{s}(v - p^h)$. Since $\frac{d^h}{s} < 1$, this requires $p^h = v$.

Proof of Proposition 3 for the case where $s < d^{\ell} < d^h$.

Consider $w = \ell$. From $d^{\ell} > s$, mass balance requires $\liminf D_k(\ell)/S_k(\ell) > 0$. Thus, from Lemma 5, either $\lim Q_k^B(\beta_k(0) \mid \ell) = 1$ or $\lim \beta_k(0) = v$ (or both). Mass balance and $s < d^{\ell}$ prohibits $\lim Q_k^B(\beta_k(0) \mid \ell) = 1$ (since then $\lim Q_k^B(\theta \mid \ell) = 1$ for all θ). Thus, $\lim \beta_k(0) = v$. As argued before in the proof of Lemma 6: This and $\liminf D_k(\ell)/S_k(\ell) > 0$ implies that sellers trade for sure when setting $r = v - \varepsilon$ for any $\varepsilon > 0$, that is, $\lim Q_k^S(v - \varepsilon | w) = 1$ for all ε and $w \in \{\ell, h\}$. Thus, $\lim V_k^S(\theta) = v - c$ for all θ. This implies $\lim Q_k^S(\theta) = 1$ (trading probabilities are competitive), the law of one price with $p^h = p^{\ell} = v$, and the characterization of payoffs.

Proof of Proposition 3 for the case where $s > d^{\ell} > d^h$.

Suppose $\lim \beta_k(1) > c$. Then, the monotonicity of ρ_k and Lemma 4 imply that $\lim Q_k^S(\theta|h) = 1$ for all $\theta \in [\underline{\theta}^S, \bar{\theta}]$. But this violates mass balance since $s < d^h$.

Thus, $\lim \beta_k(1) = c$. As argued before in the proof of Lemma 6, when $\lim \beta_k(1) = c$ then $\lim V_k^B(\theta) = v - c$ for all θ. This requires $\lim Q_k^B(\theta|w) = 1$ for almost all $\theta \in [\underline{\theta}^B, \bar{\theta}]$ and $w \in \{\ell, h\}$. Thus trading probabilities are competitive. Of course, if $\lim \beta_k(1) = c$, then $\lim P_k^B(\theta|w) = \lim P_k^S(\theta|w) = c$ (the law of one price holds). The characterization of payoffs follows from $p^h = p^{\ell} = c$.

This completes the proof of Proposition 3. *Q.E.D.*

9.2 Monotone Updating

In this Appendix, we prove monotonicity of certain posteriors for the case where $d^h > s > d^{\ell}$. We start the analysis with the following result:

Lemma 7 *If δ is large enough, then $\frac{\Gamma^B_{(1)}(\theta|h)}{\Gamma^B_{(1)}(\theta|\ell)}$ and $\frac{\gamma^B_{(1)}(\theta|h)}{\gamma^B_{(1)}(\theta|\ell)}$ are strictly increasing in θ for all $\theta \geq \underline{\theta}^B$.*

Proof. Recall that $\Gamma^B_{(1)}(\theta|w) = e^{-\frac{D(w)}{S(w)}(1-\Gamma^B(\theta|w))}$, and thus

$$\gamma^B_{(1)}(\theta|w) = \frac{D(w)}{S(w)}\gamma^B(\theta|w)\,\Gamma^B_{(1)}(\theta|w).$$

The likelihood ratio can be written as

$$\frac{\gamma^B_{(1)}(\theta|h)}{\gamma^B_{(1)}(\theta|\ell)} = \frac{\frac{D(h)}{S(h)}\gamma^B(\theta|h)\,\Gamma^B_{(1)}(\theta|h)}{\frac{D(\ell)}{S(\ell)}\gamma^B(\theta|\ell)\,\Gamma^B_{(1)}(\theta|\ell)}.$$

The no-introspection condition implies that

$$\frac{\gamma^B(\theta|h)}{\gamma^B(\theta|\ell)} = \frac{D(\ell)\theta}{D(h)(1-\theta)}, \tag{23}$$

and thus

$$\frac{\gamma_{(1)}^{B}(\theta|h)}{\gamma_{(1)}^{B}(\theta|\ell)} = \frac{S(\ell)}{S(h)}\frac{\theta}{1-\theta}\frac{\Gamma_{(1)}^{B}(\theta|h)}{\Gamma_{(1)}^{B}(\theta|\ell)}. \tag{24}$$

Suppose that for all $\theta \in supp\Gamma^B$ it holds that $\frac{S(\ell)}{S(h)}\frac{\theta}{1-\theta} > 1$. Then $\frac{\gamma_{(1)}^{B}(\theta|h)}{\gamma_{(1)}^{B}(\theta|\ell)} > \frac{\Gamma_{(1)}^{B}(\theta|h)}{\Gamma_{(1)}^{B}(\theta|\ell)}$, so $\frac{\Gamma_{(1)}^{B}(\theta|h)}{\Gamma_{(1)}^{B}(\theta|\ell)}$ is increasing in θ because

$$\begin{aligned}\left(\frac{\Gamma_{(1)}^{B}(\theta|h)}{\Gamma_{(1)}^{B}(\theta|\ell)}\right)' &= \frac{\gamma_{(1)}^{B}(\theta|h)\,\Gamma_{(1)}^{B}(\theta|\ell) - \gamma_{(1)}^{B}(\theta|\ell)\,\Gamma_{(1)}^{B}(\theta|h)}{\left(\Gamma_{(1)}^{B}(\theta|\ell)\right)^2} \\ &= \frac{\gamma_{(1)}^{B}(\theta|\ell)}{\Gamma_{(1)}^{B}(\theta|\ell)}\left(\frac{\gamma_{(1)}^{B}(\theta|h)}{\gamma_{(1)}^{B}(\theta|\ell)} - \frac{\Gamma_{(1)}^{B}(\theta|h)}{\Gamma_{(1)}^{B}(\theta|\ell)}\right).\end{aligned}$$

Moreover, $\frac{\gamma_{(1)}^{B}(\theta|h)}{\gamma_{(1)}^{B}(\theta|\ell)} = \frac{\frac{D(h)}{S(h)}\theta\Gamma_{(1)}^{B}(\theta|h)}{\frac{D(\ell)}{S(\ell)}(1-\theta)\Gamma_{(1)}^{B}(\theta|\ell)}$ is also increasing in θ.

Therefore, we only to need to establish that $\theta \in supp\Gamma^B$ implies $\frac{S(\ell)}{S(h)}\frac{\theta}{1-\theta} > 1$ for δ large enough. Let the per period trades have mass $t(w)$ for $w = \ell, h$. Then $S(w) = \frac{s-t(w)}{1-\delta}$ in a steady-state.[42] By feasibility $t(\ell) \leq d^\ell$, and by the fact that the limit is competitive in the high state we have that $\lim S_k(h)(1-\delta_k) = s - \lim t_k(h) = 0$. Therefore, we obtain that $\lim S_k(\ell)/S_k(h) = \infty$ and thus $\frac{\Gamma_{(1)}^{B}(\theta|h)}{\Gamma_{(1)}^{B}(\theta|\ell)}$ is increasing for all $\theta \geq \underline{\theta}^B$. *Q.E.D.*

Next, we establish Lemma 8. Recall that $\theta_+^B(\theta, b)$ is the posterior of a buyer who starts with belief θ, and learns that his bid b did not win (either there was a higher bidder or the seller had set a higher reserve price).

Lemma 8 *Assume $d^h > s > d^\ell$. There exists a $\underline{\delta} < 1$ such that for all $\delta > \underline{\delta}$ the following holds. Buyers update upward, that is*

$$\theta_+^B(\theta, b) > \theta$$

for all $b \in [\beta(\underline{\theta}^B), \beta(1))$.

Proof. Let $F_b^w(x) = \Gamma_{(1)}^B(\beta^{-1}(x)|w)$ denote the probability that the highest bid is less than or equal to x in state w. Similarly, let $F_r^w(x) = \Gamma^S(\rho^{-1}(x)|w)$ denote the probability that the reserve price set is less than or equal to x in state w. Given previous results, it holds that $F_b^w(x) = e^{-\frac{D(w)}{S(w)}(1-\Gamma^B(\beta^{-1}(x)|w))}$. In the main text we show that $F_b^w(x) = \Gamma_{(1)}^B\left(\beta^{-1}(x)|w\right)$ and $F_r^w(x) = \Gamma^S\left(\rho^{-1}(x)|w\right)$. By definition,

$$\frac{\theta_+^B(\theta, b)}{1-\theta_+^B(\theta, b)} = \frac{\theta}{1-\theta}\frac{1-F_b^h(b)F_r^h(b)}{1-F_b^\ell(b)F_r^\ell(b)}.$$

[42] The mass of sellers present in the next period is $1 + \delta(S(w) - t(w))$, which needs to be equal to $S(w)$ to reach a steady-state. Therefore, $1 + \delta(S(w) - t(w)) = S(w)$ or $S(w) = \frac{1-t(w)}{1-\delta}$.

Note, that $F_b^w(r) = \Gamma_{(1)}^B\left(\beta^{-1}(r)|w\right)$ and thus the fact that $\frac{\Gamma_{(1)}^B(\theta|h)}{\Gamma_{(1)}^B(\theta|\ell)}$ is strictly increasing in θ for all $\theta \geq \underline{\theta}^B$ by Lemma 7 implies that $\frac{F_b^h(r)}{F_b^\ell(r)}$ is increasing in r for all $r \geq \beta(\underline{\theta}^B)$. By construction, $\frac{F_b^h(\beta(1))}{F_b^\ell(\beta(1))} = 1$ holds. These two observations imply that for all $b \in [\beta(\underline{\theta}^B), \beta(1))$, it holds that $\frac{F_b^h(b)}{F_b^\ell(b)} < 1$. Therefore, it is sufficient to show that $\frac{F_r^h(b)}{F_r^\ell(b)} \leq 1$. We show that $\frac{F_r^h(b)}{F_r^\ell(b)}$ is weakly increasing in b and we know that by construction $\frac{F_r^h(\beta(1))}{F_r^\ell(\beta(1))} = 1$, which concludes our proof. The requirement that $\frac{F_r^h(b)}{F_r^\ell(b)}$ is weakly increasing in b for all $b \geq \beta(\underline{\theta}^B)$ is equivalent to $\frac{\Gamma^S(\theta|h)}{\Gamma^S(\theta|\ell)}$ is increasing in θ for all $\theta \geq \underline{\theta}^S$. The no-introspection condition for the sellers states that $\frac{\gamma^S(\theta|h)}{\gamma^S(\theta|\ell)} = \frac{S(\ell)}{S(h)}\frac{\theta}{1-\theta}$. Therefore, $\frac{\gamma^S(\theta|h)}{\gamma^S(\theta|\ell)}$ is strictly increasing in θ, which implies that $\frac{\Gamma^S(\theta|h)}{\Gamma^S(\theta|\ell)}$ is increasing in θ as well.[43] *Q.E.D.*

Another consequence of Lemma 7 is that updating is monotone for the sellers as well. Formally, let $\theta_+^S(\theta, r)$ be the posterior of a seller who starts with belief θ, and learns that the highest bid is less than r (including the event that there is no bidder present at all).

Lemma 9 *Assume $d^h > s > d^\ell$. There exists $\underline{\delta} < 1$ such that for all $\delta > \underline{\delta}$ the following holds. In every equilibrium, sellers update downward, that is*

$$\theta_+^S(\theta, r) < \theta$$

for all $r \in [c, v]$ and $\theta \in (0, 1)$. Moreover, a lower reserve price yields stronger updating, that is, $\theta_+^S(\theta, r)$ is weakly increasing in r for all $r \geq \beta(\underline{\theta}^B)$, and if $\Gamma_{(1)}^B(r'|w) > \Gamma_{(1)}^B(r|w)$, then $\theta_+^S(\theta, r') > \theta_+^S(\theta, r)$.

Proof. Lemma 8 implies that since buyers update upwards for any $\theta \geq \underline{\theta}^B$, and strategies are monotone thus the buyers never place any bid lower than $\beta(\underline{\theta}^B)$ in equilibrium. Therefore, for all $r < \beta(\underline{\theta}^B)$ it holds that $\theta_+^S(\theta, r) = \theta_+^S(\theta, \beta(\underline{\theta}^B))$ and thus it is sufficient to prove $\theta_+^S(\theta, r) < \theta$ for all $r \geq \beta(\underline{\theta}^B)$. By definition,

$$\frac{\theta_+^S(\theta, r)}{1 - \theta_+^S(\theta, r)} = \frac{\theta}{1 - \theta}\frac{F_b^h(r)}{F_b^\ell(r)}. \tag{25}$$

In the proof of Lemma 8, we showed that $\frac{F_b^h(r)}{F_b^\ell(r)} < 1$ for all $r \in [\beta(\underline{\theta}^B), \beta(1))$, which then implies the first claim via (25). The second result is a direct consequence of the fact that $\frac{F_b^h(r)}{F_b^\ell(r)}$ is increasing in r by 7. *Q.E.D.*

[43] Note, that

$$\lim_{\theta \to 0}\frac{\gamma^S(\theta|h)}{\gamma^S(\theta|\ell)} = \lim_{\theta \to 0}\frac{\Gamma^S(\theta|h)}{\Gamma^S(\theta|\ell)} = 0,$$

and thus the fact that $\frac{\gamma^S(\theta|h)}{\gamma^S(\theta|\ell)}$ is increasing in θ implies that $\frac{\Gamma^S(\theta|h)}{\Gamma^S(\theta|\ell)}$ is also increasing in θ.

Finally, we prove Lemma 2.

Restatement of Lemma 2. *Suppose that δ is large enough. For all $b' > b \geq 0$ such that $b, b' \in supp(\beta)$, and $\theta \in (0,1)$,*

$$\theta_0^S(\theta, b') > \theta_0^S(\theta, b).$$

Moreover, let $\theta_0^S(\theta, \varnothing)$ be the posterior if no bid is received. Then for all $b \geq 0$ such that $b \in supp(\beta)$,

$$\theta_0^S(\theta, b) \geq \theta_0^S(\theta, \varnothing).$$

Proof. Take $b = \beta(\theta^B), b' = \beta(\widetilde{\theta}^B) \in supp(\beta)$ with $b' > b$, and notice that $\widetilde{\theta}^B > \theta^B \geq \underline{\theta}^B$ by the bids being in the support of equilibrium bids as discussed after the proof of the previous Lemma. Bayes rule implies that

$$\frac{\theta_0^S(\theta, b)}{1 - \theta_0^S(\theta, b)} = \frac{\theta}{1-\theta} \frac{\gamma_{(1)}^B \left(\theta^B | h\right)}{\gamma_{(1)}^B \left(\theta^B | \ell\right)} < \frac{\theta}{1-\theta} \frac{\gamma_{(1)}^B \left(\widetilde{\theta}^B | h\right)}{\gamma_{(1)}^B \left(\widetilde{\theta}^B | \ell\right)},$$

which establishes the first result upon using Lemma 7.

To establish the second result, it is sufficient to prove $\theta_0^S(\theta, \beta(\underline{\theta}^B)) \geq \theta_0^S(\theta, \varnothing)$, which boils down to

$$\frac{\gamma_{(1)}^B \left(\underline{\theta}^B | h\right)}{\gamma_{(1)}^B \left(\underline{\theta}^B | \ell\right)} \geq \frac{\Gamma_{(1)}^B \left(\underline{\theta}^B | h\right)}{\Gamma_{(1)}^B \left(\underline{\theta}^B | \ell\right)} \tag{26}$$

upon noting that $\theta_0^S(\theta, \varnothing) = \frac{\theta}{1-\theta} \frac{\Gamma_{(1)}^B (\underline{\theta}^B | h)}{\Gamma_{(1)}^B (\underline{\theta}^B | \ell)}$ because the probability of not receiving any bid in state w is $\Gamma_{(1)}^B \left(\underline{\theta}^B | w\right)$. Using (24), (26) can be rewritten as

$$\frac{S(\ell)}{S(h)} \frac{\underline{\theta}^B}{1 - \underline{\theta}^B} \geq 1,$$

which holds for a large enough δ as we argued in the proof of Lemma 7. *Q.E.D.*

9.3 Proof of Proposition 4

Restatement of Proposition 4: *Every equilibrium that satisfies the refinement of monotone beliefs converges to the competitive limit.*

Restatement of Lemma 1: *Take any sequence of equilibria that satisfies the refinement of monotone beliefs, and suppose that $\lim \theta_0^S(\theta_k, b_k) = 0$ for some sequence $b_k \in supp\beta_k$,* and *$b_k = \rho_k(\theta_k) \to b \in [c, v]$. Then*

$$\lim V_k^S(0) + c \geq b.$$

Proof. Suppose Lemma 1 does not hold, and take b such that $b > \lim V_k^S(0) + c$, and $b = \lim b_k$ with $b_k = \rho_k(\theta_k)$. Take any $r \in (\lim V_k^S(0) + c, b)$, and we show that the condition for the refinement of monotone beliefs is violated for high enough k. By point i) of the refinement for every $\varepsilon > 0$, $\lim \theta_0^S(\theta_k, [b-\varepsilon, b_k]) = 0$. By point ii) and the continuity of V_k^S, then $\lim V_k^S(\theta_0^S(\theta_k, [b-\varepsilon, b_k])) = \lim V_k^S(0) > b - \varepsilon - c$ for all $\varepsilon > 0$, which implies our claim. *Q.E.D.*

Proof of Proposition 4.

We need to show that $p^\ell = c$, where p^ℓ is as defined in Proposition 3 for the case where $d^h > s > d^\ell$.

Case 1: $p^\ell = \lim V_k^S(0) + c$.

From Proposition 3 and Lemma 3, for all $z \in (0,1)$, $\lim W_k^S(z \mid \ell) = d^\ell(p^\ell - c)$. Also from Lemma 3, $\lim V_k^S(0) = d^\ell(p^\ell - c)$. Since $d^\ell < 1$, the hypothesis that $p^\ell = \lim V_k^S(0) + c$ implies $p^\ell = c$.

Case 2: $p^\ell > \lim V_k^S(0) + c$.

Take any $z \in \left(\underline{\theta}^B, \bar{\theta}^B\right)$. From Proposition 3, the monotonicity of β_k, and Lemma 8 (monotone updating), $\lim Q_k^B(z \mid \ell) = 1$ and $\lim Q_k^B(z \mid h) < 1$. Take $\alpha \in (0,1)$ and let t_k be the smallest number such that type z wins with a probability of at least α by period t_k in state ℓ. With θ_k being the posterior after $t_k - 1$ periods, type z wins with a probability of at least α at a bid of at most $b_k = \beta_k(\theta_k)$ (by monotonicity of θ_+).[44] Such t_k exists for k large enough. From Proposition 3, $\lim b_k = p^\ell$. Also, $\left(\delta^k\right)^{t_k} \to 1$. Therefore, the probability that z wins with a bid b_k or lower in state h is vanishing to zero. Since z wins with a probability not more than α before t_k in state ℓ and with probability converging to zero in state h, we have $\theta_k < 1$. This follows from $\lim \frac{\theta_k}{1-\theta_k} = \lim \frac{z}{1-z} \frac{\Pr(\text{no win before } t_k - 1|h)}{\Pr(\text{no win before } t_k - 1|\ell)} < \frac{\bar{\theta}^B}{1-\bar{\theta}^B} \frac{1}{1-\alpha}$.

Step 1: For all $x_k \le \bar{\theta}^S$ and $b_k = \beta_k(\theta_k)$ (as defined above),

$$\lim \theta_0^S(x_k, b_k) = 0. \tag{27}$$

From $\gamma_{k(1)} = \gamma_k^B(\theta_k|w) \frac{D_k(w)}{S_k(w)} e^{-\frac{D_k(w)}{S_k(w)}(1-\Gamma_k^B(\theta_k|w))}$, we have

$$\frac{\theta_0^S(x_k, b_k)}{1 - \theta_0^S(x_k, b_k)} = \frac{x_k}{1 - x_k} \frac{\frac{D_k(h)}{S_k(h)}}{\frac{D_k(\ell)}{S_k(\ell)}} \frac{\gamma_k^B(\theta_k|h)}{\gamma_k^B(\theta_k|\ell)} \frac{e^{-\frac{D_k(h)}{S_k(h)}(1-\Gamma_k^B(\theta_k|h))}}{e^{-\frac{D_k(\ell)}{S_k(\ell)}(1-\Gamma_k^B(\theta_k|\ell))}}.$$

By the no-introspection condition (23), it holds that

$$\frac{D_k(h)}{D_k(\ell)} \frac{\gamma_k^B(\theta_k|h)}{\gamma_k^B(\theta_k|\ell)} = \frac{\theta_k}{1 - \theta_k}.$$

[44] For example, with $t_k = 2$,

$$q_k(\beta_k(z)|\ell) < \alpha \le$$
$$q_k(\beta_k(z)|\ell) + (1 - q_k(\beta_k(z)|\ell))(1 - \delta_k) q_k\left(\beta_k\left(\theta_+^k(z)\right)|\ell\right).$$

Also, by construction $S_k(\ell) \leq \frac{s}{1-\delta_k}$ and $S_k(h) \geq s$ hold. Therefore, upon substitution, we obtain

$$\frac{\theta_0^S(x_k, b_k)}{1-\theta_0^S(x_k, b_k)} \leq \frac{x_k}{1-x_k}\frac{\theta_k}{1-\theta_k}\frac{1}{1-\delta_k}\frac{e^{-\frac{D_k(h)}{S_k(h)}(1-\Gamma_k^B(\theta_k|h))}}{e^{-\frac{D_k(\ell)}{S_k(\ell)}(1-\Gamma_k^B(\theta_k|\ell))}}.$$

By assumption, $x_k \leq \bar{\theta}^S < 1$, and we already argued $\lim \theta_k < 1$. Also, $\lim D_k(\ell)/S_k(\ell) = 0$ by Lemma 5, and thus $\lim e^{-\frac{D_k(\ell)}{S_k(\ell)}(1-\Gamma_k^B(\theta_k|\ell))} = 1$. Therefore, it is sufficient to prove that

$$\lim \frac{e^{-\frac{D_k(h)}{S_k(h)}(1-\Gamma_k^B(\theta_k|h))}}{1-\delta_k} = 0.$$

From $\left(\delta^k\right)^{t_k} \to 1$ and $\lim Q_k^B(\theta \mid h) < 1$ for all $\theta \leq \bar{\theta}$, we have $\lim(1-\Gamma_k^B(\theta_k|h)) = 0$. Finally, by Lemma 20 it holds that $\liminf \frac{D_k(h)(1-\delta_k)}{S_k(h)} = y > 0$. Therefore, $\lim \frac{e^{-\frac{D_k(h)}{S_k(h)}(1-\Gamma_k^B(\theta_k|h))}}{1-\delta_k} = \lim \frac{e^{-\frac{y}{1-\delta_k}}}{1-\delta_k} = 0$, and so (27) holds.

Step 2. For any ε and k large enough, $\rho_k(\theta) \notin [V_k^S(0) + c + \varepsilon, b_k)$ for all $\theta \leq \bar{\theta}^S$ (and recall $b_k \to p^\ell > \lim V_k^S(0) + c$).

First, we show that for all $b \in [\lim V_k^S(0) + c + \varepsilon, p^\ell]$, $b \in supp\beta_k$ for a large enough k. Otherwise, there exists an interval $(\underline{b}, \bar{b}) \subset [\lim V_k^S(0) + c + \varepsilon, p^\ell)$ that is not in the support of the limiting bid distribution but $\bar{b}$ is. Then for it to be optimal to bid close to $\bar{b}$ in the limit, it must hold that $\bar{b}$ is in the support of the limiting reserve price distribution. But then Lemma 1 applies to imply that $\bar{b} \leq \lim V_k^S(0) + c$, a contradiction. Second, suppose that for some $r \in [\lim V_k^S(0) + c + \varepsilon, b_k]$ it holds that r is in the limiting reserve price distribution. Then again Lemma 1 yields a contradiction together with the previous observation that r is in the limiting bid distribution.

Step 3. Given any ε small enough and k large enough, bidding $b'_k = V_k^S(0) + c + \varepsilon$ is strictly more profitable for θ_k than bidding b_k.

By the hypothesis of the case, $\lim b'_k < \lim b_k = p^\ell$. In addition, the probability that there is no other buyer converges to one when $w = \ell$. Hence, given Step 2, conditional on $w = \ell$, bidding b'_k strictly increases payoffs for θ_k since b'_k wins with the same probability as b_k in the low state in the limit: $\lim \frac{q_k^B(b'_k|\ell)}{q_k^B(b_k|\ell)} = 1$.

Next, observe that by Step 1 and the refinement of monotone beliefs $\lim \theta_0^S(\theta_k, [b'_k, b_k]) = 0$. Therefore, conditional on winning against bids on $[b'_k, b_k]$, the state is almost sure to be low in the limit. But in the low state, it is more profitable to place bid b'_k as we argued above, and Step 3 is complete. Thus, we have reached a contradiction and so Case 2 cannot occur. Since the claim holds in Case 1, this completes the proof of Proposition **4**. *Q.E.D.*

10 Appendix B: Proof of Proposition 6

Restatement of Proposition 6. *For any sequence of undominated equilibria with $p^\ell > c$, there exists a period t_k with $\lim t_k(1-\delta_k)^\alpha = 0$ for all $\alpha > 0$, such that all sellers who entered at least t_k periods ago set a reserve price r that satisfies*

$$\lim \Pr(b_{(1),k} < r \mid b_{(1),k} > 0, \ell) = 0.$$

Proof.

We start the proof by a useful observation. By L'Hospital's rule, an equivalent way to express $\lim \delta_k^{\tau_k} = 1$ is to assume that $\lim(1-\delta_k)\tau_k = 0$. The same argument implies that Δ_k and $1-\delta_k$ converge to zero at the same rate. Therefore, an event occurs quickly in calendar time if and only if the time of the event τ_k satisfies $\lim(1-\delta_k)\tau_k = 0$.

Step 1. For any sequence of undominated equilibria we show that for any $\varepsilon > 0$, it holds that

$$\lim \frac{\Pr(b_{(1),k} \in [p^\ell - \varepsilon, \overline{\rho}_k] \mid h)(1-\delta_k)}{\Pr(b_{(1),k} \in [p^\ell - \varepsilon, \overline{\rho}_k] \mid \ell)} < \infty. \tag{28}$$

First, note that Proposition 3 implies that for all $\varepsilon > 0$,

$$\lim \frac{\Pr(b_{(1),k} \in [p^\ell - \varepsilon, \overline{\rho}_k] \mid \ell)}{1-\delta_k} > 0. \tag{29}$$

To see this, note that Proposition 3 implies that there are d^ℓ trades in the low state in the limit per period, almost all at the price close to p^ℓ. Thus the mass of buyers who bid close to p^ℓ and win in the limit divided by the mass of sellers is equal to $d^\ell/S(\ell)$. Then (29) follows upon observing that $\lim d^\ell/(S_k(\ell)(1-\delta_k)) > 0$ by fundamental balance conditions.[45]

Step 2. Take any sequence z_k with $\lim \frac{z_k}{(1-\delta_k)^\alpha} > 0$ for some $\alpha > 0$ and $\lim \frac{z_k}{1-\delta_k} = 0$, and denote the posterior of the seller who has not traded for z_k periods as $\widehat{\theta}_k$. We establish that this posterior is very low, that is

$$\lim \frac{\widehat{\theta}_k}{1-\delta_k} = 0. \tag{30}$$

Proposition 3 implies that for almost all types of entering sellers $\theta \in [\underline{\theta}^S, \bar{\theta}^S]$, the probability of trade converges to zero in a single periods in the low state. Formally, suppose that the time sequence of equilibrium reserve prices for type $\theta \in [\underline{\theta}^S, \bar{\theta}^S]$ is $(r_0, r_1, r_2, ...r_{z_k})$. Then Proposition 3 and the assumption that $\lim \frac{z_k}{1-\delta_k} = 0$ imply that for almost all $\theta \in [\underline{\theta}^S, \bar{\theta}^S]$ the

[45] Here we also used that if τ is the relevant ratio, then the sellers obtain a bid at a price close to p^ℓ with probability $1 - e^{-\tau}$. But $\lim \frac{1-e^{-\tau}}{\tau} = 1$ when $\tau \to 0$, and thus the sellers obtain such a bid with probability τ in the limit, which implies (29) then.

sellers trade with a probability that converges to 0 in the first z_k periods when the state is low.[46]

Each seller trades with a probability of at least $q_k^S(\overline{\rho}_k \mid h)$, and $q = \lim q_k^S(\overline{\rho}_k \mid h) > 0$ by the proof of Lemma 4. Therefore, the probability of not trading in the first z_k periods is less than $(1-q)^{z_k}$. Then straightforward calculations show that $\lim \frac{z_k}{(1-\delta_k)^\alpha} > 0$ for some $\alpha > 0$ implies that the probability of not trading in the high state in the first z_k periods converges to zero at a much faster rate than $1-\delta_k$ does:

$$\lim \frac{(1-q)^{z_k}}{1-\delta_k} = 0.$$

Given that the entering belief is always less than 1, the last displayed formula implies $\lim \frac{\widehat{\theta}_k}{1-\delta_k} = 0$ via the rules of Bayesian updating.

Step 3. By Bayesian updating, $\frac{\Pr(\ell | b_{(1),k} \in [p^\ell - \varepsilon, \overline{\rho}_k], \theta)}{1-\Pr(\ell | b_{(1),k} \in [p^\ell - \varepsilon, \overline{\rho}_k], \theta)} = \frac{(1-\theta)\Pr(b_{(1),k} \in [p^\ell - \varepsilon, \overline{\rho}_k] | \ell)}{\theta \Pr(b_{(1),k} \in [p^\ell - \varepsilon, \overline{\rho}_k] | h)}$. This posterior upon z_k periods satisfies

$$\lim \frac{(1-\widehat{\theta}_k)\Pr(b_{(1),k} \in [p^\ell - \varepsilon, \overline{\rho}_k] \mid \ell)}{\widehat{\theta}_k \Pr(b_{(1),k} \in [p^\ell - \varepsilon, \overline{\rho}_k] \mid h)} =$$

$$= \lim \frac{1-\delta_k}{\widehat{\theta}_k} \lim \frac{\Pr(b_{(1),k} \in [p^\ell - \varepsilon, \overline{\rho}_k] \mid \ell)}{\Pr(b_{(1),k} \in [p^\ell - \varepsilon, \overline{\rho}_k] \mid h)(1-\delta_k)} = \infty,$$

by Step 1 (see (28)) and Step 2 (see 30). By monotonicity of updating (see Lemma 2), it follows that for all $r_k \le \overline{\rho}_k$,

$$\lim \frac{\Pr(\ell \mid b_{(1),k} \in [p^\ell - \varepsilon, r_k], \widehat{\theta}_k)}{1-\Pr(\ell \mid b_{(1),k} \in [p^\ell - \varepsilon, r_k], \widehat{\theta}_k)} = \infty. \tag{31}$$

Step 4. Given (31), the rest of the proof is immediate: Straightforward algebra yields that setting a reserve price of $p^\ell - \varepsilon$ and thus accepting almost all highest bids in the low state in the limit dominates not accepting almost all highest bids in the low state if

$$\lim \Pr(\ell \mid b_{(1),k} \in [p^\ell - \varepsilon, r_k], \widehat{\theta}_k)(p^\ell - \varepsilon - (V_k^S(0) + c)) >$$

$$(1 - \lim \Pr(\ell \mid b_{(1),k} \in [p^\ell - \varepsilon, r_k], \widehat{\theta}_k))(v - (p^\ell - \varepsilon)),$$

which concludes the proof upon taking $\varepsilon > 0$ arbitrarily small, and noting that $\lim V_k^S(0) = (p^\ell - c)d^\ell < p^\ell - c$. *Q.E.D.*

[46] This follows because if the trading probability was positive in the first z_k periods with $\lim \frac{z_k}{1-\delta_k} = 0$, then with lower reserves in the latter periods the sellerswould end up trading with probability 1 in the low state in the limit, which violates Proposition 3.

References

Banerjee, A., and D. Fudenberg (2004): "Word-of-mouth learning," *Games and Economic Behavior*, 46(1), 1–22.

Benabou, R. (1993): "Search Market Equilibrium, Bilateral Heterogeneity, and Repeat Purchases," *Journal of Economic Theory*, 60(1), 140–158.

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

Burdett, K., and K. L. Judd (1983): "Equilibrium Wage Dispersion," *Econometrica*, 51(4), 955–969.

Dana, J. D. (1994): "Learning in an Equilibrium Search Model," *International Economic Review*, 35(3), 745–771.

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

Duffie, D., and G. Manso (2007): "Information Percolation in Large Markets," *American Economic Review*, 97(2), 203–209.

Duffie, D., and Y. Sun (2007): "Existence of independent random matching," *The Annals of Applied Probability*, 17(1), 386–419.

Fudenberg, D., and D. K. Levine (1993): "Steady State Learning and Nash Equilibrium," *Econometrica*, 61(3), 547–573.

Gale, D. (1987): "Limit Theorems for Markets with Sequential Bargaining," *Journal of Economic Theory*, 43(1), 20–54.

Golosov, M., G. Lorenzoni, and A. Tsyvinski (2011): "Decentralized Trading with Private Information," Mimeo.

Golub, B., and M. O. Jackson (2010): "Naï£¡ve Learning in Social Networks and the Wisdom of Crowds," *American Economic Journal: Microeconomics*, 2(1), 112–49.

Gottardi, P., and R. Serrano (2005): "Market Power And Information Revelation In Dynamic Trading," *Journal of the European Economic Association*, 3(6), 1279–1317.

Hayek, F. A. (1945): "The Use of Knowledge in Society," *The American Economic Review*, 35(4), pp. 519–530.

Janssen, M., A. Parakhonyak, and A. Parakhonyak (2014): "Non-reservation Price Equilibria and Consumer Search," .

Janssen, M., and S. Shelegia (2015): "Consumer Search and Double Marginalization," *American Economic Review*, 105(6), 1683–1710.

Juda, A. I., and D. C. Parkes (2006): "The Sequential Auction Problem on eBay: An Empirical Analysis and a Solution," *Proceedings of the 7th ACM Conference on Electronic Commerce, Ann Arbor, MI*, pp. 180–189.

Kunimoto, T., and R. Serrano (2004): "Bargaining and Competition Revisited," *Journal of Economic Theory*, 115(1), 78–88.

Kyle, A. S. (1989): "Informed Speculation with Imperfect Competition," *Review of Economic Studies*, 56(3), 317–55.

Lauermann, S., and G. Virag (2012): "Auctions in Markets: Common Outside Options and the Continuation Value Effect," *American Economic Journal: Microeconomics*, 101(5), 107–130.

Lauermann, S., and A. Wolinsky (2016): "Search with Adverse Selection," *Econometrica*, 84(1), 243–315.

Majumdar, D., A. Shneyerov, and H. Xie (2015): "An Optimistic Search Equilibrium," *Review of Economic Design*, pp. 1–26.

Milchtaich, I. (2004): "Random-player games," *Games and Economic Behavior*, 47(2), 353–388.

Milgrom, P., and I. Segal (2002): "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70(2), 583–601.

Myerson, R. B. (1998): "Extended Poisson Games and the Condorcet Jury Theorem," *Games and Economic Behavior*, 25(1), 111–131.

Ostrovsky, M. (2011): "Information Aggregation in Dynamic Markets with Strategic Traders," Forthcoming: Econometrica.

Pesendorfer, W., and J. Swinkels (1997): "The Loser's Curse and Information Aggregation in Common Value Auctions," *Econometrica*, 65(6), 1247–1282.

——— (2000): "Efficiency and Information Aggregation in Auctions," *American Economic Review*, 90(3), 499–525.

Reny, P. J., and M. Perry (2006): "Toward a Strategic Foundation for Rational Expectations Equilibrium," *Econometrica*, 74(5), 1231–1269.

Rostek, M., and M. Weretka (2012): "Price Inference in Small Markets," *Econometrica*, 80(2), 687–711.

Rubinstein, A., and A. Wolinsky (1985): "Equilibrium in a Market with Sequential Bargaining," *Econometrica*, 53(5), 1133–50.

Satterthwaite, M., and A. Shneyerov (2007): "Dynamic Matching, Two-Sided Incomplete Information, and Participation Costs: Existence and Convergence to Perfect Competition," *Econometrica*, 75(1), 155–200.

——— (2008): "Convergence to Perfect Competition of a Dynamic Matching and Bargaining Market with Two-sided Incomplete Information and Exogenous Exit Rate," *Games and Economic Behavior*, 63(2), 435–467.

Serrano, R., and O. Yosha (1993): "Information Revelation in a Market with Pairwise Meetings: The One Sided Information Case," *Economic Theory*, 3(3), 481–99.

Shneyerov, A., and C. L. Wong (2010): "The Rate of Convergence to Perfect Competition of Matching and Bargaining Mechanisms," *Journal of Economic Theory*, 145(3), 1164–1187.

Vives, X. (2010): *Information and Learning in Markets: The Impact of Market Microstructure*. Princeton University Press, Princeton.

Wolinsky, A. (1990): "Information Revelation in a Market with Pairwise Meetings," *Econometrica*, 58(1), 1–23.