Learning and Price Discovery in a Search Market*

Stephan Lauermann† Wolfram Merzyn‡ Gábor Virág§

January 23, 2017

Abstract

We study a dynamic, decentralized exchange economy with aggregate uncertainty about the relative scarcity of a commodity. We characterize price discovery and show how traders gradually learn about the state of the market through equilibrium actions. Such learning leads to equilibrium outcomes that are approximately competitive when the frictions are small. We derive equilibrium price and trading patterns related to learning, experimentation, and regret.

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 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 the tradition of Mortensen (1982). These models assume that the agents know the market

*We have benefited from comments by Georg Nöldeke, Tilman Börgers, Jan Eeckhout, Leonardo Felli, Chris House, Daniel Krähmer, Philipp Kircher, Benny Moldovanu, Tymofiy Mylovanov, Marco Ottaviani, Art Shneyerov, and the anonymous referees, as well as seminar audiences at Boston College, Cambridge, Concordia University, Essex, Michigan State, NYU Stern, Queen Mary University, Penn State University, the University of Toronto, the Canadian Economic Theory Conference, the MaCII 2016 Workshop, the Decentralization Conference at Washington University, the 2nd MEDS Theory workshop in Graz, and the Midwest Theory Conference. We thank Deniz Kattwinkel, 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 Netz AG, wolfram.merzyn@gmx.de.

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

conditions, while we assume that the agents do not. 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.

By studying those questions, we contribute 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. The identified equilibrium allows for a detailed analysis of trading strategies under market uncertainty in our model. 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 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 of whether markets aggregate information has implications for policies that are intended to increase the transparency of markets, to remove insider information, or to centralize markets to decrease informational frictions; see, for example, the discussion of benchmarks for increasing transparency in search markets in Duffie, Dworczak, and Zhu (2016). The main policy implication of our work, however, is that decentralized markets may work reasonably well in achieving market clearance even in the presence of aggregate uncertainty, at least for small frictions. This was not immediately evident from existing work on decentralized markets, especially given negative results for related problems in the literature; see below.

As an illustration of the setting we have in mind, consider a bidder on eBay who seeks to purchase a consumer good. eBay is a decentralized market but it still 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 the 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 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 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 observations are hard to reconcile with models of conventional search theory where buyers fully know market conditions from the outset.[1] However, they are consistent with an environment with aggregate uncertainty where buyers need to learn the market conditions.

We develop a model to study markets with aggregate uncertainty. 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.[2] 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 exit rate acts similar to a discount factor: it makes waiting costly and is interpreted as the "friction" of trade.

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 trader receives a noisy signal about the state of the world upon arrival. Moreover, at the end of the period, those who did not trade obtain additional information regarding the state because they draw an inference from the fact that their respective bids or reserve price lost. Feedback is otherwise minimal: traders do not observe the reserve price or bids after the auction. We concentrate on steady-state equilibria that are monotone in the beliefs of the agents: a buyer who attaches a higher probability to the high state bids more, and a seller who attaches a

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

[2] The auction protocol captures that there often is some degree of direct competition and helps with the analysis; see Section 4.4.

higher probability to the high state sets a higher reserve 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.[3] Therefore, the winner's curse implies a lower value for winning the good than was 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.[4]

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.

We are particularly interested in the equilibrium outcomes when the exogenous exit rate is small, which is interpreted as the frictionless limit of the decentralized market. We show 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.

[3] 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, as sellers set low reserve prices when the state is likely to be low.

[4] 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*.

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. In Section 4, we introduce a class of full-trade equilibria. We show that such equilibria exist and use these equilibria to discuss and illustrate our main points. In Section 5, we characterize all equilibria for small frictions. In Section 6, we introduce the refinement of monotone beliefs and show that, with this refinement, all equilibria become competitive for small frictions. In Section 7, we provide a discussion of our other assumptions, applications, and policy implications. The Appendix contains the proofs of the characterization results from Sections 5 and 6. A supplementary online Appendix contains the proofs of all results about full-trade equilibria from Section 4.

2 Contribution to the Literature

We contribute to 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).[5] 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.

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. The absence of aggregate uncertainty is a substantial restriction for at least two reasons. First, assuming that all participants know the aggregate market conditions is unrealistic in many markets. Second, price discovery has been emphasized as an integral function of markets. Our model allows us to investigate whether and under which conditions decentralized markets can serve this function.

There is a related strand of literature on consumer search, some of which considers uncertainty on the buyers' side; see, for example, Benabou and Gertner (1993), Dana (1994), and Janssen and Shelegia (2015). Unlike in our model, in that literature, sellers know the state, and these contributions typically focus on "reserve-price equilibria."[6] In our model, no trader knows the state, we consider the full set of equilibria, and we study information aggregation with small frictions.

[5] For recent contributions, see, for example, Satterthwaite and Shneyerov (2007, 2008), Shneyerov and Wong (2010), Kunimoto and Serrano (2004), Lauermann (2012), and Lauermann (2013), and the references therein. Most of these contributions study settings with private information.

[6] An exception is Janssen, Parakhonyak, and Parakhonyak (2014).

Our work is also related to work on matching and bargaining with exogenously assumed *common values*.[7] Particularly prominent contributions are Wolinsky (1990) and Blouin and Serrano (2001).[8] 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 implies that traders increasingly insist on favorable terms, turning the search market into "a vast war of attrition" (Blouin and Serrano (2001, p. 324)). As a result, non-competitive prices can be sustained in Wolinsky (1990) and Blouin and Serrano (2001) even when search frictions are small.[9] 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. Nevertheless, this "insistence problem" is overcome by the fact that agents who did not trade become more accommodating in subsequent meetings.

Golosov, Lorenzoni, and Tsyvinski (2014) consider a related search model with common values in which the traded good is divisible and study the long-run outcome for fixed frictions. They do not study whether outcomes become competitive in the "frictionless" limit, and learning is qualitatively different with a divisible good.

There is a large body of related work on the foundation for rational expectation equilibrium in centralized institutions.[10] 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, learning from other traders is endogenous and depends on the chosen action (bid).

[7] Majumdar, Shneyerov, and Xie (2015) is the only other paper that considers a dynamic matching and bargaining game with aggregate uncertainty. However, they assume that traders are subjectively certain about the market conditions.

[8] Serrano and Yosha (1993) consider a related problem with one-sided private information, and Gottardi and Serrano (2005) consider a "hybrid" model of decentralized and centralized trading. Lauermann and Wolinsky (2016) study information aggregation if a single, privately informed buyer searches among many sellers.

[9] To make their model tractable, Wolinsky (1990) and Blouin and Serrano (2001) assume that traders choose between only two price offers (bargaining postures). It is an open question whether the trading outcomes in these models are competitive with small frictions if prices can be chosen freely.

[10] Examples are the work on large double auctions by Reny and Perry (2006), Cripps and Swinkels (2006), and Pesendorfer and Swinkels (1997, 2000), as well as the work on information aggregation in Cournot models, summarized in Vives (2010), and in financial markets (e.g., Kyle (1989), Ostrovsky (2012), and Rostek and Weretka (2012)).

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, with $v > c \geq 0$. 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. Buyers and sellers maximize the expected payoffs.

There are two states of nature, a high state and a low state $\omega \in \{h, \ell\}$. Both states are equally likely. The realized state of nature is fixed throughout and is unknown to the traders, similar to Wolinsky (1990). For each realization of the state of nature, we consider the corresponding steady-state outcome, indexed by ω. The state of nature determines the constant and exogenous number of new traders who enter the market (the *flow*), and indirectly, the constant and endogenous number of traders in the market (the *stock*). The mass of buyers entering each period is d^ℓ in the low state and d^h in the high state, with $d^h > d^\ell > 0$. The mass of sellers who enter each period is the same in both states and is equal to s. There is uncertainty about who is on the long side of the market whenever $d^h > s > d^\ell$.

The buyers and the sellers are characterized by their beliefs $\theta \in [0, 1]$, the probability that they assign to the high state. (Sometimes θ is called a trader's *type*.) Each buyer who enters the market privately observes a noisy signal x. The signal is distributed with support $[0, 1]$ and with a cumulative distribution function (c.d.f.) $G_s^B(x|\omega)$ that admits a continuous density function, $g_s^B(x|\omega)$. The likelihood ratio $\frac{g_s^B(x|h)}{g_s^B(x|\ell)}$ is strictly increasing. Therefore, the Bayesian posterior $\theta(x) = \frac{(0.5)d^h g_s^B(x|h)}{(0.5)d^h g_s^B(x|h)+(0.5)d^\ell g_s^B(x|\ell)}$ is strictly increasing in x and, hence, invertible.

We will work directly with the distribution of the induced posteriors to simplify the exposition later. The support of this distribution is $[\underline{\theta}^B, \bar{\theta}^B] = [\theta(0), \theta(1)]$. Its c.d.f. is $G^B(\theta|\omega) = G_s^B(\theta^{-1}(\theta)|\omega)$, and its density is denoted $g^B(\theta|\omega)$.[11]

For a buyer, the mere fact of entering the market contains news because the inflow is larger in the high state. This is expressed by the likelihood ratio $d^h/d^\ell > 1$.[12] To avoid

[11] One can easily verify that the density of posteriors must satisfy $\theta = \frac{(0.5)d^h g^B(\theta|h)}{(0.5)d^h g^B(\theta|h)+(0.5)d^\ell g^B(\theta|\ell)}$, and, hence, $\frac{\theta}{1-\theta} = \frac{d^h}{d^\ell}\frac{g^B(\theta|h)}{g^B(\theta|\ell)}$. This formula expresses the Bayesian consistency of the distribution of posterior beliefs. Smith and Sorensen (2013) call it the no-introspection property: "If the individual further updates his private belief θ by asking of its likelihood in the two states of the world, he must learn nothing more." An implication is that $G^B(\cdot|h)$ dominates $G^B(\cdot|\ell)$ in the likelihood ratio ordering. We will later use an analogous property for the distribution of beliefs in the entire population, and not just for beliefs of the entering cohort.

[12] 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 ω, a mass $d(\omega)$ of the potential buyers actually enters the market. Alternatively,

technical difficulties, and to simplify exposition, we assume that[13]

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

The sellers' side is analogous. In state ω, the induced posteriors of the entering sellers are distributed on the support $[\underline{\theta}^S, \bar{\theta}^S]$ with $0 < \underline{\theta}^S < \bar{\theta}^S < 1$, with a c.d.f. $G^S(\theta|\omega)$ and density $g^S(\theta|\omega)$.

Each period unfolds as follows:

1. Entry occurs (the *"inflow"*): A mass s of sellers and a mass d^ω 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[14] and is equal to $e^{-\mu}\mu^n/n!$, where $\mu(\omega) = D(\omega)/S(\omega)$ 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 $\mu(\omega)$, of course.

3. Each seller runs a first-price sealed-bid auction with a secret reserve price r. Feedback is minimal: 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 a buyer learns only whether she has won, and a seller observes only whether the highest bid is above r.

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 her 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.

5. Upon losing, the remaining traders who did not exit update their beliefs based on the information gained from losing with their submitted bids or not trading with their chosen reserve price r. Together with the inflow, these traders make up the stock for the next period.

---

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).

[13]Because the belief conditional on entering (but without conditioning on the signal) is equal to $d^h/(d^h + d^\ell) > 1/2$, the assumption holds if $\frac{g_s^B(x|h)}{g_s^B(x|\ell)}$ is sufficiently close to one for all $x \in [0, 1]$ (the initial signal is not too precise).

[14]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 time. An immediate consequence is that in any period, the set of optimal actions (bids and 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 essentially a special case of theirs if the state is known, with the main difference being that they allow for heterogeneous values and costs.[15] We discuss our setup with a known state as a benchmark now, before returning to the detailed description of equilibrium.

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

$$r - c = \delta V^S. \tag{2}$$

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' equilibrium 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.[16] 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 $e^{-\mu}$. 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.

[15] Another difference is that we do not have a separate entry stage.

[16] Because of the Poisson distribution, $e^{-\mu}$ is both, the probability that a seller is matched with 0 buyers and the probability that a buyer has no competitor.

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 of prices. It is a (very) special case of Satterthwaite and Shneyerov (2008), as mentioned before. 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 model with $d(h) > d(\ell)$. 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.

Strategies and Stocks.

Formally, the masses of buyers and sellers in the stock are $D(\omega)$ and $S(\omega)$. The distributions of beliefs are given by c.d.f.s $\Gamma^j(\cdot|\omega)$ for $j = B, S$ and $\omega = \ell, h$. We assume that each function Γ^j is absolutely continuous, with a density $\gamma^{\omega,j}$ that is right continuous on $[0, 1)$.

The bidding strategy β is a weakly increasing function and maps beliefs from $[0, 1]$ to $[c, v]$. Moreover, we assume that β is strictly increasing on 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 $[c, v]$.

Trading Probabilities.

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 $\Gamma^B_{(1)}(x|\omega)$ denote the probability that the highest belief in the auction is below x. The event in which all the buyers have a belief below x includes the event in which there are no buyers present at all. The probability of having no buyer present is $\Gamma^B_{(1)}(0|\omega)$ by our assumption that there is no atom in the distribution of beliefs at zero. The Poisson distribution implies $\Gamma^B_{(1)}(0|\omega) = e^{-\mu(\omega)}$, where $\mu(\omega) = D(\omega)/S(\omega)$ as defined before. In general, the first order statistic of the distribution of beliefs is given by

$$\Gamma^B_{(1)}(x|\omega) = e^{-\mu(\omega)(1-\Gamma^B(x|\omega))}. \tag{3}$$

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

Given the assumption that the bidding strategies are strictly increasing on the support of Γ^B, a bid b wins if (i) there is no bidder in the match with a belief above $x = \beta^{-1}(b)$ and

(ii) the seller sets a reserve price of at most b.[17] So, the probability that a buyer wins with bid b in state ω is

$$q^B(b|\omega) := \Gamma^B_{(1)}\left(\beta^{-1}(b)\,|\omega\right)\Gamma^S(\rho^{-1}(b)|\omega).$$

We also refer to q^B as the per-period trading probability. Similarly, let $q^S(r|\omega) = 1 - \Gamma^B_{(1)}\left(\beta^{-1}(r)\,|\omega\right)$ denote the probability that a seller trades with reserve price r in state ω.

Updating.

We first derive the posterior of a buyer upon not trading with bid b. Bayes' rule requires that

$$\theta^B_+(\theta, b) = \frac{\theta(1 - q^B(b|h))}{1 - (\theta q^B(b|h) + (1-\theta)q^B(b|\ell))}, \tag{4}$$

if the denominator is strictly positive; otherwise, θ^B_+ is arbitrary.

The posterior of a seller who did not trade with reserve price r is

$$\theta^S_+(\theta, r) = \frac{\theta(1 - q^S(r|h))}{1 - (\theta q^S(r|h) + (1-\theta)q^S(r|\ell))}; \tag{5}$$

the denominator is always strictly positive.

To discuss marginal incentives, it is useful to define the "tieing posterior". Let $b_{(1)}$ be the highest bid (with $b_{(1)} = c$ if there is no bid) and let $\theta^S_0(\theta, A) = \Pr\left(h|b_{(1)} \in A, \theta\right)$, be the posterior probability of h conditional on the highest bid being in a (measurable) set A.[18]

Taking the conditioning event A to be a single bid b, we obtain $\theta^S_0(\theta, b) = \Pr\left(h|b_{(1)} = b, \theta\right)$. Likewise, $\theta^B_0(\theta, A)$ is a buyer's posterior probability of h conditional on the highest *other* bid being in a set A. Because of the Poisson distribution, $\theta^B_0(\theta, A) = \theta^S_0(\theta, A)$.

Optimality Conditions.

Let $V^B(\theta)$ denote the buyers' value function, which satisfies

$$\max_b\; q^B(\theta, b)(v - b) + \delta(1 - q^B(\theta, b))V^B\left(\theta^B_+(\theta, b)\right), \tag{6}$$

with $q^B(\theta, b) = \theta q^B(b|h) + (1-\theta)q^B(b|\ell)$. A bidding strategy β is optimal if $b = \beta(\theta)$ solves problem (6) for every θ.

Let $V^S(\theta)$ denote the sellers' value function, which satisfies

$$\max_r\; q^S(\theta, r)(r - c) + \delta(1 - q^S(\theta, r))V^S\left(\theta^S_+(\theta, r)\right), \tag{7}$$

with $q^S(\theta, r) = \theta q^S(r|h) + (1-\theta)q^S(r|\ell)$. A reserve-price strategy ρ is optimal if $r = \rho(\theta)$ solves the problem (7) for every θ.

[17] Here and in the following, we use the generalized inverse of β, given by $\beta^{-1}(b) = \inf\{\theta \in [0,1]\,|\beta(\theta) \geq b\}$, where $\beta^{-1}(b) = 1$ if $\beta(\theta) < b$ for all θ.

[18] So, $\theta^S_0(\theta, A) = \frac{\theta \Pr(b_{(1)} \in A|h)}{\theta \Pr(b_{(1)} \in A|h) + (1-\theta)\Pr(b_{(1)} \in A|\ell)}$, with $\Pr(b_{(1)} \in A|\omega) = \Gamma^B_{(1)}(A|\omega)$.

We impose a notion of perfectness on the sellers' equilibrium strategies, and require

$$\rho(1) - c = \delta V^S(1) \qquad \text{and} \qquad \rho(0) - c = \delta V^S(0). \tag{8}$$

Thus, a seller who knows the state accepts prices if and only if accepting the price yields larger payoffs than continued searching. If ρ satisfies (8), it is said to be *undominated*. The condition is similar to the perfection requirement (2) in the full-information benchmark.

The main bite of (8) comes in combination with the monotonicity of ρ, namely, for all $\theta \in (0,1)$,

$$\delta V^S(0) \leq \rho(\theta) - c \leq \delta V^S(1). \tag{9}$$

In particular, (9) rules out "no-trade" equilibria with $\beta(1) = c$ and $\rho(0) = v$ because $V^S(1) = 0$ requires $\rho(\theta) = c$ for all θ.

Steady-State Conditions.

The steady-state stock of sellers needs to satisfy

$$S(\omega)\Gamma^S(\theta|\omega) = sG^S(\theta|\omega) + \delta S(\omega) \int_{\left\{\tau:\theta_+^S(\tau,\rho(\tau))\leq\theta\right\}} \Gamma^B_{(1)}\left(\beta^{-1}(\rho(\tau))|\omega\right) d\Gamma^S(\tau|\omega). \tag{10}$$

To see why, note that the left side is equal to the mass of sellers in the stock at the *beginning* of a period who have a type below θ. The right side is equal to the mass of such sellers at the *end* of the period, which consists of all those sellers in the inflow with type less than θ (the first term on the right side) plus the mass of sellers who lose, survive, and update to some type less than θ (the second term). In steady state, the mass at the end of a period must be identical to the mass at the beginning.[19]

The analogous steady-state condition for buyers is

$$\begin{aligned} D(\omega)\Gamma^B(\theta|\omega) \;&=\; d^\omega G^B(\theta|\omega) \\ &+\delta D(\omega) \int_{\left\{\tau:\theta_+^B(\tau,\beta(\tau))\leq\theta\right\}} \left(1 - \Gamma^B_{(1)}(\tau|\omega)\,\Gamma^S(\rho^{-1}(\beta(\tau))|\omega)\right) d\Gamma^B(\tau|\omega). \end{aligned} \tag{11}$$

Steady-State Equilibrium.

A steady-state equilibrium in undominated, symmetric, monotone strategies with an atomless distribution of types (abbreviated to *steady-state equilibrium* or just *equilibrium*) consists of (i) masses of buyers and sellers, $S(\omega), D(\omega)$, and distribution functions $\Gamma^B(\cdot|\omega)$, $\Gamma^S(\cdot|\omega)$ for $\omega \in \{\ell, h\}$, such that the steady-state conditions (10) and (11) hold for all θ; (ii) updating functions θ_+^B, θ_+^S that are consistent with Bayes' rule (4), (5); (iii) weakly increasing

[19] For the purpose of this paper, the steady-state model is *defined* by (10) and (11). 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).

functions β and ρ that are optimal (solve (6) and (7), respectively), with β strictly increasing on the support of Γ^B, and (iv) ρ is undominated, satisfying (8).

Every equilibrium determines (lifetime) trading probabilities $Q^j(\theta|\omega)$, expected trading prices $P^j(\theta|\omega)$, and expected payoffs $EU^j(\theta|\omega)$, for $j \in \{B, S\}$. With this notation, $EU^B(\theta|\omega) = Q^B(\theta|\omega)\left(v - P^B(\theta|\omega)\right)$ and $EU^S(\theta|\omega) = Q^S(\theta|\omega)\left(P^S(\theta|\omega) - v\right)$. Similarly, $V^j(\theta) = \theta EU^j(\theta|h) + (1-\theta)EU^j(\theta|\ell)$. We will be interested in how equilibrium prices and trading probabilities compare with those predicted by market clearing.

Roadmap.

Section 4 studies a class of equilibria in which sellers accept all equilibrium bids ("full-trade equilibria"). We show that such equilibria exist, and that, except for the trading price in the low state, the frictionless limit is competitive. Using the explicitly constructed full-trade equilibria, we discuss how learning unfolds. Section 5 then shows that the limit properties of the full-trade equilibria are shared by all steady-state equilibria. Section 6 shows that under a natural refinement, the limit outcome is fully competitive.

4 Full-Trade Equilibria

In this Section, we prove the existence of "full-trade" equilibria for large δ. We start by assuming that sellers are offering a common, exogenously fixed reserve price r_0 and show that there exists a unique "bidding equilibrium" given r_0 for all δ. Then, we characterize the limit properties of these bidding equilibria for $\delta \to 1$. Finally, we use these limit properties to argue that for all sufficiently large δ, the constraint on the sellers' behavior does not bind and the bidding equilibria constitute full equilibria of the original game. All proofs related to the full-trade equilibria are in the online Appendix.

4.1 Bidding Equilibrium: Construction and Uniqueness

Here, we fix an arbitrary $r_0 \in [c, v)$ and assume that sellers set reserve price r_0. A *steady-state bidding equilibrium* given r_0 is a combination of stocks and strategies such that (i) the stocks satisfy the steady-state conditions given the strategies, (ii) the buyers' strategy is optimal, and (iii) the sellers' strategy is $\rho(\theta) \equiv r_0$; equivalently, we replace the sellers' optimality condition by $\rho(\theta) \equiv r_0$ in the original steady-state equilibrium definition (and drop undominatedness) .

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

The proof is in the online Appendix. Below is a sketch of the arguments.

Decoupling and the Construction of Steady-State Stocks. The buyers bid at least r_0, and thus, all bids are accepted. Moreover, by assumption, β is strictly increasing on the support of Γ^B. Therefore, given any stock of buyers and sellers, the winning probabilities of any buyer type θ is independent of the exact form of β. This implies also that the posterior belief conditional on losing with an equilibrium bid is independent of β. Together, this allows us to "decouple" the construction of the steady-state stock and β.

The construction of the steady-state stock in the online Appendix uses a recursive algorithm on the cohorts of buyers. This algorithm starts with the exogenously given entering cohort and then recursively determines the cohort that remains in the next period. The difficulty of the construction is to ensure that updating is well-behaved during the recursion. Finally, the constructed steady-state stock is shown to be the unique stock that satisfies the steady-state equations.

The Bidding Strategies. Given the unique steady-state stock, we construct the bidding strategies. First, we characterize the bidding equilibrium in a static, reduced-form auction in which we take continuation values as given. Note that the continuation values depend not only on the state but also on the buyers' beliefs, since the beliefs determine the buyers' behavior. Specifically, we suppose that $W(\theta, y)$ is the expected continuation payoff of a buyer if the probability of the high state is y and the buyer acts according to belief θ. Then, we show that for any W that satisfies certain regularity conditions, there exists a unique bidding strategy that forms a mutual best response in the reduced-form auction. This bidding strategy satisfies a standard differential equation with boundary condition $\beta(0) = r_0$. (For more details and a statement of the differential equation determining β, see Section 4.5.)

The second step shows that the continuation values are unique. Specifically, we show that in every steady-state bidding equilibrium, continuation value functions W have to satisfy certain regularity properties.[20] Then, for any given continuation value function that satisfies these properties, the unique bidding strategy found in Step 1 defines a new continuation value if all buyers are bidding according to it. The implied mapping is then shown to be a contraction on the set of "regular" continuation value functions, which establishes both the existence and the uniqueness of the bidding equilibrium.

Note that equilibrium existence is shown through an explicit construction: First, the steady-state stocks are constructed using a recursive algorithm on the cohorts. Second, a contraction mapping determines the continuation values, and, third, the bidding strategy is a solution to a standard differential equation. Thus, in principle, the bidding equilibria can be numerically constructed and evaluated.

[20] In particular, continuation values $V^B(\theta)$ are decreasing and convex, and a regular $W(\theta, \cdot)$ is a family of hyperplanes supporting some decreasing and convex $V^B(\theta)$.

4.2 Limit Properties of Full-Trade Bidding Equilibria

This section characterizes the outcome of full-trade equilibria for $\delta \to 1$. Let $\{\delta_k\}_{k=1}^{\infty}$ be a sequence of such that the exit rate converges to zero, $\lim_{k\to\infty}(1-\delta_k)=0$. From Proposition 1, there is a unique bidding equilibrium for each δ_k. 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.

To simplify notation, throughout the paper, the term "limit" (and the operator lim) refers to a limit over a subsequence such that all the magnitudes of interest are converging.[21] We will not repeat this qualification each time, but it is always there. Also, since most limits are with respect to $k \to \infty$, we suppress it and write $\lim(1-\delta_k)=0$ etc.

Proposition 2 *Suppose $d^\ell < s < d^h$ and take any $r_0 \in [c, v)$. Then, for the unique steady-state bidding equilibria given r_0,*

$$\begin{aligned}\lim Q_k^S(\theta|h) &= 1, \text{ and } \lim P_k^S(\theta|h) = v \text{ for all } \theta \in [\underline{\theta}^S, \bar{\theta}^S],\\ \lim Q_k^B(\theta|\ell) &= 1, \text{ and } \lim P_k^B(\theta|\ell) = r_0 \text{ for all } \theta \in [\underline{\theta}^B, \bar{\theta}^B].\end{aligned}$$

Thus, the limit outcome is competitive if $r_0 = c$: The short side of the market trades with probability converging to one, and the price is v and c in the high and in the low state, respectively. If $r_0 > c$, then the limit outcome is almost competitive, except that the price is higher than the competitive price when $\omega = \ell$. In the following, we discuss why the proposition holds. Because this discussion is already an almost complete proof and because the proposition is a special case of subsequent results, we do not provide a separate proof.[22]

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

Consider $\omega = h$. Since $s < d(h)$, the steady-state conditions imply for any δ_k, it must be the case that $S_k(h) < D_k(h)$.[23] So, each seller is matched with positive probability with a buyer in every period since $\mu_k(h) = D_k(h)/S_k(h) > 1$ implies $\lim e^{-\mu_k(h)} < 1$.[24] Hence, sellers trade with probability converging to one as $\delta_k \to 1$.

Similarly for $\omega = \ell$. Here, $d(\ell) < s$ implies that $\mu_k(\ell) = D_k(\ell)/S_k(\ell) < 1$. So, $\lim e^{-\mu_k(\ell)} > 0$, that is, with positive and nonvanishing probability any buyer is the sole bidder. Therefore, as $\delta_k \to 1$, buyers must trade with probability converging to one.

[21] This will be sufficient for our purposes since for the sequences under consideration, any property that holds for all converging subsequences also holds for the original sequence itself. Without this convention, we would need to introduce multiple layers of sub(-sub)sequences throughout many of our proofs.

[22] To be precise, the following discussion gives a complete proof for why the trading probabilities are competitive. The argument for $\lim P_k^B(\theta|\ell) = r_0$ is also essentially complete. The argument for $\lim P_k^S(\theta|h) = v$ is the only place where the discussion is incomplete. The interested reader will find a formalization of the intuition from the text in the proof of Proposition 4 on Page 36 of the Appendix.

[23] This inequality is formally shown in Equation (29) in the online Appendix. It holds because buyers and sellers either exit in pairs or with the same exogenous exit probability $1-\delta$. This inequality is true for any strategy profile (β, ρ), not just the full-trade profile considered here.

[24] Recall our convention that lim refers to a subsequence for which the limit exists in the extended reals.

(ii) *In the limit, almost all trades take place at prices close to r_0 in the low state and close to v in the high state.*

Intuitively, bidding in the low state is very non-competitive: Any bid $r_0 + \varepsilon$ guarantees sure winning in the low state in the limit because with positive probability any buyer is the sole bidder, as observed before. Therefore, buyers trade at a price close to r_0 in the low state. Conversely, in the high state, it is not feasible for all buyers to trade with probability one since $s < d(h)$. Therefore, some buyers will be rationed and have to stay in the market for many periods as $\delta_k \to 1$. Thus, there are many buyers who are essentially certain that the state is high, and these buyers drive up the price to v. We return to this discussion later in Section 4.5, where we relate the dynamic bidding behavior to the winner's and loser's curse.

For later reference, note that the sellers' trading probability in the low state satisfies

$$\lim Q_k^S(\theta|\ell) = \frac{d(\ell)}{s} < 1,$$

which follows from $\lim Q_k^B(\theta|\ell) = 1$, feasibility,[25] and the sellers' common reserve price.

4.3 Full-Trade Equilibria of the Entire Game

We show that when δ is high, then for any interior 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.[26]

Proposition 3 *Suppose $d^\ell < s < d^h$. Take any $r_0 \in (c, v)$. If δ 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 uses the limit properties identified in Proposition 2. We fix some sequence $\delta_k \to 1$ and the corresponding sequence of bidding equilibria β_k.

Setting r_0 is a Best Response.

First, the proof shows that setting a reserve price above r_0 is not optimal for any seller with any non-degenerate belief, for δ_k close enough to 1. Specifically, consider any $r' \in (r_0, v)$. Then, it follows from Proposition 2 that the probability of the low state conditional on the highest bid being below r' converges to 1 as $\delta_k \to 1$. This is because sellers trade with probability converging to 0 at such prices in the high state but with positive probability in the low state. Now, from Proposition 2 and the subsequent remark, the sellers' continuation

[25] Feasibility requires that the mass of buyers and the mass of sellers who end up trading are equal; see (20) for a formal definition.

[26] For the cases $d^\ell < d^h < s$ and $s < d^\ell < d^h$, existence of a full-trade equilibrium is implied by Proposition 6. In these cases, however, r_0 depends on δ and cannot be chosen freely in (c, v).

payoff in state ℓ is roughly $\frac{d(\ell)}{s}(r_0 - c) < r_0 - c$. Thus, raising the reserve price from r_0 to any $r' \in (r_0, v)$ is not optimal for δ_k large enough. The proof in the appendix verifies that for sufficiently large k, there is no profitable deviation to any $r' \in (r_0, v]$. It is clearly not strictly profitable to set $r' < r_0$ given $\beta_k(0) \geq r_0$.

The Reserve Price r_0 is "Undominated"

Second, the proof shows that the equilibrium satisfies undominatedness. Since $\rho(\theta) = r_0$ for $\theta \in (0,1)$ this requires $\delta_k V_k^S(0) \leq r_0 - c \leq \delta_k V_k^S(1)$; see (9). Proposition 2 implies that for the constructed equilibria, these inequalities hold for all $r_0 \in (c, v)$ and δ_k large enough. In particular, we observed in the paragraph above that $\lim V_k^S(0) = \frac{d(\ell)}{s}(r_0 - c)$, and it is immediate from Proposition 2 that $\lim V_k^S(1) = v - c$.

Remark. Seller Observing Bids before Accepting

The proof of Proposition 3 also implies that if the seller *observes* the highest bid b before setting a reserve price, then the full-trade equilibrium constructed above would remain an equilibrium for small frictions. For example, we already argued that the seller would accept any bid between r_0 and $r' < v$, for δ_k sufficiently large. For bids below r_0, one can assign off-equilibrium beliefs in an appropriate way to rationalize rejection.

4.4 Full-Trade Equilibria and Equilibrium Existence

There are a number of technical challenges that any existence proof for a model of search with learning faces. Our model is set up to solve or circumvent these, and the full-trade equilibria are particularly useful in this regard.

First, as is well-known, proving the existence of equilibrium is a non-trivial problem in any search model because of the endogeneity of the distribution of population characteristics; see Smith (2011). This problem is magnified with learning because the relevant types (beliefs in our model) are evolving according to an endogenous transition rule. The assumption that feedback is minimal helps. Full-trade equilibria help further because we can decouple the determination of the stock and the strategies; see above.

Second, the effective 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, we are solving a bargaining problem with two-sided asymmetric information and interdependent values, for which little progress has been made; see, e.g., Ausubel, Cramton, and Deneckere (2002). The auction protocol helps because—at least, in the full-trade equilibria—one market side can be taken to be essentially non-strategic. Moreover, the protocol reduces the multiplicity problem that is common in two-sided asymmetric information bargaining problems (although the multiplicity problem is still there).

There are two other technical problems. First, as Reny and Perry (2006) show, the first order statistic of affiliated random variables that are mapped by two different monotone functions (reserve prices and bidding strategies) does not have to be affiliated in general, which makes it difficult to prove existence of monotone equilibria in general. Here, full-trade equilibria render this problem moot. Second, 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 generally lost. We solve this problem by placing a lower bound on the bidders' types in the inflow, which implies that belief updating is monotone upward in any full-trade equilibrium.

Let us note that full-trade equilibria have been studied in the related literature before to address existence problem; see Satterthwaite and Shneyerov (2007) and the consumer search literature discussed in our literature review for examples.

Finally, full-trade equilibria allow for explicit characterization results. For example, we are able to derive the distribution of beliefs explicitly so the model is amenable to numerical analysis. Also, as shown below, in our model most of the important characterization results extend to other types of equilibria; see Section 5.

4.5 Price Discovery: Winner's and Loser Curse

It follows from Propositions 2 and 3 that for δ close to 1, there exists a full-trade equilibrium in which sellers set a common reserve price close to c and the outcome is close to being competitive.[27] We use these equilibria to show how price discovery takes place through a "decentralized tâtonnement process," by which we mean the dynamic bidding behavior of the individual buyers and sellers.

In the full-trade equilibrium, the buyers and the sellers behave differently: The sellers set the same reserve price r_0 regardless of their time on 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. This ensures that prices are close to r_0 in the low state and close to v in the high state. Given this, the sellers can "trust the market" to ensure them the right price even when setting a low reserve price. (Indeed, one can observe sellers using "absolute auctions" with no reserve price in many online markets such as eBay; see Jehiel and Lamy (2015).) Note that the critical piece of information about the correct price comes through the behavior of buyers who increase their bids after having been unable to trade at the low price. This may be compared to the tâtonnement process of a hypothetical Walrasian auctioneer who increases prices after observing excess demand.

To explain the dynamics of the buyers' bidding strategies and the difference to the sellers'

[27] For every $\varepsilon > 0$, there is some $\bar{\delta}(\varepsilon)$ such that for $\delta \geq \bar{\delta}(\varepsilon)$ there is a full-trade equilibrium with $r_0 = c + \varepsilon$.

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 she is almost certain that the state favors the sellers. Second, the loser's curse means that if a buyer has not traded for sufficiently many periods, then she eventually increases her bid.

To describe this more formally, we derive the optimality condition that determines the equilibrium bidding strategy β. (The following can be skipped by the reader with no loss of continuity.) Recall that $EU^B(\theta|\omega)$ is the expected payoff of a buyer with belief θ in state ω. If the probability of state h is y, the expected payoff of a buyer who holds belief θ is

$$W(\theta, y) = yEU^B(\theta|h) + (1-y)\, EU^B(\theta|\ell).$$

The expected distribution of the highest type among the other bidders is $\Gamma^y_{(1)}(\theta) = y\Gamma^B_{(1)}(\theta|h) + (1-y)\,\Gamma^B_{(1)}(\theta|\ell)$, where $\Gamma^B_{(1)}(\theta|\omega)$ is the probability that the highest (other) type is below θ in state ω. Then, the payoff of a buyer who has type θ but mimics θ' by bidding $\beta\left(\theta'\right)$ is

$$\Gamma^\theta_{(1)}(\theta')\left(v - \beta\left(\theta'\right)\right) + \delta(1 - \Gamma^\theta_{(1)}(\theta'))W\left(\theta^B_+, \theta^B_+\right),$$

where $\theta^B_+ = \theta^B_+(\theta, \beta\left(\theta'\right))$ is the posterior of a buyer who starts with belief θ and loses with a bid $\beta\left(\theta'\right)$. Using that in equilibrium choosing $\theta' = \theta$ is optimal, and solving the necessary first-order condition for β to be a best response for given continuation payoffs W yields the differential equation[28]

$$\beta'(\theta) = \frac{\gamma^\theta_{(1)}(\theta)}{\Gamma^\theta_{(1)}(\theta)}\left(v - \beta(\theta) - \delta W\left(\theta^B_+, \theta^B_0\right)\right), \tag{12}$$

with initial condition $\beta(0) = r_0$ and where $\theta^B_0 = \theta^B_0(\theta, \beta(\theta))$ is the posterior of a buyer who starts with belief θ conditional on being tied at bid $\beta(\theta)$. The main point is that the buyer's "effective valuation" is $v - \delta W\left(\theta^B_+, \theta^B_0\right)$, which takes into account the continuation payoff. Importantly, the continuation payoff is evaluated conditional on being tied because this is the pivotal event for a marginal change in the bid.[29]

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 $\theta^B_0(\theta, \beta(\theta))$ 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

[28] See Lemma 17 and its proof in the online Appendix for the derivation of (12).

[29] Here, $W\left(\theta^B_+, \theta^B_0\right)$ is the expected payoff of a buyer who behaves according to belief θ^B_+ (the losing posterior), evaluated given θ^B_0 (the tieing posterior).

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.

We also describe the loser's curse in full-trade equilibria. In particular, buyers who have not traded for sufficiently many periods know that the state likely favors the sellers—even conditional on being tied at the top—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 $\theta_{+}^{B}(\theta, \beta(\theta)) > \theta$.

The winner's curse works differently for the sellers. The pivotal event for a marginal increase in r 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, indicating a low opportunity cost of selling. Therefore, the winner's curse does not arise for the sellers, and the sellers do not have an incentive to increase their reserve prices. Given that the sellers accept all bids anyway, 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.

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 ($d^{\ell} < s < d^{h}$). Here, we show that this characterization extends to *all* steady-state equilibria.[30]

Proposition 4 *Take any sequence $\delta_k \to 1$ and consider any corresponding sequence of steady-state equilibria.*

- *If $d^{\ell} < s < d^{h}$, then there exists some $p_0 \in [c, v]$ such that*

$$\begin{aligned} \lim Q_k^S(\theta|h) &= 1, \text{ and } \lim P_k^S(\theta|h) = v \text{ for all } \theta \in [\underline{\theta}^S, \bar{\theta}^S], \\ \lim Q_k^B(\theta|\ell) &= 1, \text{ and } \lim P_k^B(\theta|\ell) = p_0 \text{ for all } \theta \in [\underline{\theta}^B, \bar{\theta}^B]. \end{aligned}$$

- *If $d^{\ell} < d^{h} < s$, then*

$$\lim Q_k^B(\theta|\omega) = 1, \text{ and } \lim P_k^B(\theta|\omega) = c \text{ for all } \theta \in [\underline{\theta}^B, \bar{\theta}^B] \text{ and } \omega = \ell, h.$$

- *If $s < d^{\ell} < d^{h}$, then*

$$\lim Q_k^S(\theta|h) = 1, \text{ and } \lim P_k^S(\theta|\omega) = v \text{ for all } \theta \in [\underline{\theta}^S, \bar{\theta}^S] \text{ and } \omega = \ell, h.$$

[30] The proof in the Appendix shows a stronger result, proving also the law-of-one-price: The realized transaction prices converge in distribution to a point in each state (to v and p_0).

The main difficulty is showing that trading probabilities become competitive in all steady-state equilibria in the limit. In the full-trade equilibria considered before, this was easily shown to be true. Here, however, we need to rule out "war-of-attrition" type structures. This problem is self-enforcing: Once buyers are insisting on low prices for sufficiently many periods, it is no longer optimal for sellers to "trust the market" and set very low reserve prices. However, once sellers are insisting on high prices themselves, not trading is less informative for buyers, and so it becomes optimal for buyers to indeed bid low for longer.

5.1 Sketch of the Proof

The proposition follows from two lemmas. The lemmas highlight the role of "rationing" in price discovery. First, if a seller who knows that the state is high trades with probability less than one (is rationed), then prices unravel all the way down to c:

Lemma 1 *If* $\lim Q_k^S(1|h) < 1$*, then* $\lim \rho_k(1) = \lim \beta_k(1) = c$*.*

The lemma is proven as follows. First, $\lim Q_k^S(1|h) < 1$ implies that the probability of a bid above $\rho_k(1)$ vanishes to zero. Hence, a buyer who bids $\rho_k(1) + \varepsilon$ wins almost surely (for any $\varepsilon > 0$ and k large enough), so that the optimal bid satisfies $\lim \beta_k(1) \leq \lim \rho_k(1)$. It follows that $\lim P_k^S(1|h) = \lim \rho_k(1)$. So, $\lim V_k^S(1) = \lim Q_k^S(1|h)(\rho_k(1) - c)$. Undominatedness requires that $\rho_k(1) - c = \delta_k V_k^S(1)$. Combining these,

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

Now, $\lim Q_k^S(1|h) < 1$ implies $\lim \rho_k(1) - c = 0$.

In words, suppose a seller with type $\theta = 1$ who is sure of $\omega = h$ trades with probability less than one but insists on a relatively high reserve price. Then, the seller's continuation payoff is below the payoff from trading at the reserve price, leading the seller to accept lower prices as well by undominatedness—which eventually unravels all the way to c.

Lemma 2 *If* $\lim D_k(\ell)/S_k(\ell) > 0$*, then either* $\lim Q_k^B(0|\ell) = 1$ *or* $\lim \beta_k(0) = v$ *or both.*

The idea of the proof is as follows. First, $\lim D_k(\ell)/S_k(\ell) > 0$ means that sellers are matched with at least one buyer with a positive, non-vanishing probability every period. Now, optimal bidding requires that $\rho_k(0) \leq \beta_k(0)$, since, otherwise, no seller would accept $\beta_k(0)$. But for sellers to accept $\beta_k(0)$, it must be unlikely that any buyer bids much higher, meaning, almost all buyers are bidding between $\beta_k(0)$ and $\beta_k(0) + \varepsilon$, for any $\varepsilon > 0$ and k large enough.[31] So, as $\delta_k \to 1$, there is an atom in the bid distribution at $\beta_k(0)$. Finally, if

[31] Otherwise, $\lim D_k(\ell)/S_k(\ell) > 0$ would imply that the seller would surely trade in state ℓ even when he sets a reservation price of $\beta_k(0) + \varepsilon$ for some $\varepsilon > 0$. Hence, setting a reservation price below $\beta_k(0)$ would not be optimal so that $\rho_k(0) > \beta_k(0)$, a contradiction.

both, $\lim Q_k^B(0|\ell) < 1$ and $\lim \beta_k(0) < v$, a buyer having type $\theta = 0$ has a strict incentive to overbid this atom in order to increase her trading probability. So, if $\lim D_k(\ell)/S_k(\ell) > 0$, either $\lim Q_k^B(0|\ell) = 1$ or $\lim \beta_k(0) = v$—as claimed.

Using Lemmas 1 and 2 to prove Proposition 4.

Suppose $d^\ell < d^h < s$. Since more sellers than buyers enter the market, feasibility dictates that not all sellers can trade with probability one in the limit. In particular, by monotonicity of ρ_k, $\lim Q_k^S(1|h) < 1$. Now, Lemma 1 implies $\lim \rho_k(1) = \lim \beta_k(1) = c$. Therefore, a buyer who bids $c + \varepsilon$ is certain to trade as $\delta_k \to 1$. So, the buyers' payoffs are converging to $v - c$—which means they must trade with probability converging to one at a price converging to c, implying the claim.

Suppose $s < d^\ell < d^h$. As noted before on Page 15, $s < d^\ell$ implies that $S_k(\ell) < D_k(\ell)$ for all k. Hence, by Lemma 2, either $\lim Q_k^B(0|\ell) = 1$ or $\lim \beta_k(0) = v$. However, since more buyers than sellers enter the market, feasibility and monotonicity of β_k require $\lim Q_k^B(0|\ell) < 1$. So, $\lim \beta_k(0) = v$. But this means all buyers are bidding close to v in the limit and $S_k(\ell) < D_k(\ell)$ for all k means that sellers are certain to be matched with some buyer eventually. So, the sellers' payoffs are converging to $v - c$—which implies the claim.

Finally, suppose $d^\ell < s < d^h$. We start by verifying the trading probabilities and consider $\omega = h$ first. By Lemma 1, if $\lim Q_k^S(1|h) < 1$, then $\lim \rho_k(1) = \lim \beta_k(1) = c$. This would mean that all buyers could ensure themselves a payoff close to $v - c$ by bidding just slightly above c. However, since $d^h > s$, it is not feasible that all buyers have payoffs $v - c$, since this would require a trading probability of one. Therefore, $\lim Q_k^S(1|h) < 1$ cannot hold.

Now, consider $\omega = \ell$. If not all buyers are able to trade with probability converging to one, then a positive mass of them accumulates, implying that $\lim D_k(\ell)/S_k(\ell) > 0$. Also, by the monotonicity of β_k, $\lim Q_k^B(0|\ell) < 1$. So, by Lemma 2, $\lim \beta_k(0) = v$. But this and $\lim D_k(\ell)/S_k(\ell) > 0$ would allow all sellers to trade with probability converging to one at a price v, so that all sellers could ensure a payoff $v - c$ in the low state. But this contradicts feasibility since $d^\ell < s$ means that not all sellers can end up trading. Thus, it must be that almost all buyers can trade, as claimed.

Given that the trading probabilities of the shorter side of the market are one, all traders from the shorter side will trade at the same price (to see why, suppose two types of sellers both trade with probability one in the high state but at different prices; then, the type that trades at the lower price has an incentive to mimic the other). Thus, almost all trade must take place at the same price (the law-of-one-price). Now, to see that this price must be v for $\omega = h$, note that the trading probability of the buyers must be strictly smaller than one since $s < d^h$. But then each buyer would have an incentive to bid higher than the common trading price in order to increase its trading probability—which drives prices up to v.

5.2 Relation to Full-Trade Equilibria

The full-trade equilibria also capture another property common to steady-state equilibria: Suppose that $d^\ell < s < d^h$ and consider any sequence of steady-state equilibria that does not converge to the competitive allocation, that is, $p^\ell > c$. In the online Appendix, Proposition 9 shows the following: There exists a sequence of periods $\{t_k\}_{k=1}^\infty$ with $\lim(\delta_k)^{t_k} = 1$, such that all sellers who entered at least t_k periods ago set a reserve price r that is almost surely acceptable to any buyer he is matched with. Since $\lim(\delta_k)^{t_k} = 1$, the probability of exogenous exit before t_k vanishes to zero. Thus, similar to the full-trade equilibria, sellers are conceding quickly to buyers in any non-competitive equilibrium.[32]

6 Competitive Limit with Monotone Beliefs

6.1 Off-Equilibrium Beliefs and Information Aggregation

Proposition 4 shows that steady-state equilibrium outcomes are almost competitive for small frictions, except that prices may be too high in the low state if $d^\ell < s < d^h$—meaning, equilibria are non-competitive exactly when there is aggregate uncertainty about the market clearing price but not otherwise. Proposition 3 shows that non-competitive equilibria indeed exist.

What drives the non-competitive outcome is a combination of the trading protocol and aggregate uncertainty. Most fundamentally, for any $r_0 \in (c, v)$, it is trivially a mutual best response for sellers to set reserve prices not smaller than r_0 and for buyers to bid at least r_0—since no buyer bids below r_0, sellers have no incentives to set reserve prices below r_0 and, given that, buyers have no incentives to bid below r_0.

However, in the absence of aggregate uncertainty, this problem can be solved easily. To see this, recall from Proposition 4 that when $s > d^h > d^\ell$—sellers are known to be on the long side—prices are close to c in every undominated equilibrium. In particular, there is no full-trade equilibrium with sellers setting $r_0 > c$. This is because, given the corresponding bidding equilibrium from Proposition 1, even the highest continuation payoff $\delta V^S(1)$ is bounded strictly below $r_0 - c$ for $\delta \to 1$. Therefore, the reserve price r_0 would be weakly dominated by any reserve price $r' < r_0$ for which $r' > \delta V^S(1) + c$: Whatever a seller believes about the state, if the winning bid is in $[r', r_0]$, accepting the bid yields strictly higher payoffs than continuation, and otherwise, if the winning bid is not in $[r', r_0]$, then it does not matter

[32] When the limit is competitive, $p^\ell = c$, there may be steady-state equilibria in which the sellers do not concede quickly (relative to δ_k) in the limit. There is a possibility for an equilibrium in which the sellers set high reserve prices for a very long 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.

whether the reserve price is r' or r_0. Thus, requiring equilibrium to be undominated solves the problem.[33]

Such an argument no longer works if $d^\ell < s < d^h$. In this case $V^S(1) \approx v - c$, so no reserve price $r_0 \in (c, v)$ is weakly dominated: The seller may believe that conditional on the winning bid being in $[r', r_0]$, the state is very likely to be high so that it is optimal to reject it. Intuitively, non-competitive full-trade equilibria with $r_0 > c$ for $\delta \to 1$ are supported by off-equilibrium beliefs that put high probability on $\omega = h$ if the winning bid is below r_0.

However, such off-equilibrium beliefs seem unnatural. First, recall that we already argued that in any full-trade equilibrium, for any r_0 and δ sufficiently large, the probability of $\omega = h$ is close to zero conditional on the winning bid being in $[r_0, b']$, for any $b' < v$.[34] Thus, the equilibrium beliefs conditional on the winning bid have to be non-monotone in any non-competitive full-trade equilibrium. Second, for bids that are on the equilibrium path, a later result shows that beliefs are monotone in all steady-state equilibria; see Lemma 3.

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) steady-state 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 (measurable) set A.

Refinement of Monotone Beliefs. *An equilibrium satisfies the refinement of monotone beliefs if there exist beliefs $\theta_0^S(\theta, \cdot)$ 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*

$$\delta V^S(\theta_0^S(\theta, [b_1, b_2])) \geq b_1 - c. \tag{13}$$

Condition (i) states that beliefs need to satisfy monotonicity, including off-the-equilibrium path. Condition (ii) states that if a seller sets a reserve price of b_2, he can rationalize not decreasing his reserve price to $b_1 < b_2$ by his beliefs $\theta_0^S(\theta, \cdot)$: To see this, suppose that the seller switches from reserve price b_2 to reserve price b_1. If the highest bid is not in the interval

[33] Similarly, in the model by Satterthwaite and Shneyerov (2008) without aggregate uncertainty, there are strategy profiles that are mutual best-responses in which prices are too high because of the same problem. In that paper, the problem is solved by having the seller observe the bids before accepting and requiring sequentially rational acceptance decisions, which has the same implication as undominatedness here, namely, that sellers follow a reservation price strategy with $r = \delta V^S + c$.

[34] See Page 16, when we discuss the sellers' optimality of setting r_0.

$[b_1, b_2]$, then there is no change in the seller's payoff. Therefore, when changing his strategy, the seller can assume that the highest bid is in the interval $[b_1, b_2]$, and the seller's belief conditional on this event is $\theta_0^S(\theta, [b_1, b_2])$. In this event, if he does not accept the winning bid, then his continuation utility is at most $\delta V^S(\theta_0^S(\theta, [b_1, b_2]))$. If he does accept the winning bid, then his profit is at least $b_1 - c$.[35] Given this discussion, (13) means that the seller can rationalize setting b_2 rather than b_1 by the belief that, conditional on reducing the reserve price from b_2 to b_1 and the event $b_{(1)} \in [b_1, b_2]$, the expected payoff from trading at this bid is lower than the continuation payoff.

The non-competitive equilibria featured in Proposition 3 do not satisfy the refinement. The reason is that, as we discussed, $\theta_0^S(\theta, [r_0, r_0 + \varepsilon]) \to 0$ for any $\varepsilon > 0$. Hence, monotonicity requires $\theta_0^S(\theta, [r_0 - \varepsilon, r_0]) \to 0$. In addition, trade occurs with probability converging to $\frac{d^\ell}{s} < 1$ in the low state at a price close to r_0. Together, for any $\theta \in (0, 1)$, continuation payoffs conditional on $b_{(1)} \in [r_0 - \varepsilon, r_0]$ satisfy $\delta V^S(\theta_0^S(\theta, [r_0 - \varepsilon, r_0])) \cong \delta V^S(0) \cong \frac{d^\ell}{s}(r_0 - c)$. Thus, the refinement fails for $b_1 = r_0 - \varepsilon$ and $b_2 = r_0$, for ε small enough and δ large enough: Then, the left-hand side of (13) is roughly $\frac{d^\ell}{s}(r_0 - c)$, while the right-hand side is roughly $r_0 - c$.

The next result shows that this holds more generally. The refinement of monotone beliefs rules out *all* steady-state equilibria that are not competitive in the limit. Non-competitive equilibria can only be supported by non-monotone beliefs.

Proposition 5 *The allocation provided by any sequence of steady-state equilibria that satisfy the refinement of monotone beliefs converges to the competitive limit.*

The proposition is proven in the appendix. Given Proposition 4, we only need to show that, in the case $d(\ell) < s < d(h)$, the trading price in the low state is competitive, meaning, $p^\ell = c$. The key step of the proof (Step 1) is to show that, in the limit, when $p^\ell > c$, any on-path winning bid sufficiently close to p^ℓ makes the sellers believe that the state is almost surely low. Then, part (i) of the refinement extends this to bids below p^ℓ. Therefore, we can argue that it must be the case that sellers would accept a range of bids below p^ℓ, which upsets any equilibrium in which $p^\ell > c$.

In the next two sections, we show that equilibrium beliefs are monotone, and we show that there do exist full-trade equilibria that satisfy the refinement.

[35] The expected price conditional on $b_{(1)} \in [b_1, b_2]$ can be believed to be higher than b_1 but not lower. Similarly, the continuation payoff conditional on $b_{(1)} \in [b_1, b_2]$ may be believed to be lower than $V^S(\theta_0^S(\theta, [b_1, b_2]))$ but not higher.

6.3 Monotone Equilibrium Updating

To motivate the refinement of monotone beliefs, we show that $\theta_0^S(\theta, b) = \lim_{\varepsilon \to 0} \theta_0^S(\theta, [b, b+\varepsilon])$ is increasing in b on the support of β for every steady-state equilibrium for sufficiently large δ. This result implies that a higher winning bid in the support of equilibrium bids makes the high state more likely.[36] Thus, the refinement restricts the set of equilibria only through its implications for bids that are outside the support of β.

Lemma 3 *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 steady-state equilibria: For all $b' > b \geq c$ such that b and b' are in the support of the equilibrium bid distribution,*

$$\theta_0^S(\theta, b') > \theta_0^S(\theta, b). \tag{14}$$

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

$$\theta_0^S(\theta, b) \geq \theta_0^S(\theta, \varnothing). \tag{15}$$

The lemma uses the characterization from Proposition 4 given $d^\ell < s < d^h$. It follows from the Proposition that the trading probabilities are competitive. Hence, it must be that $S(\ell) > S(h)$ and $D(\ell) < D(h)$ for low enough frictions. But this implies that the winning bid tends to be higher in the high state: First, the expected number of buyers per seller is higher since $D(\ell)/S(\ell) < D(h)/S(h)$, and, second, the distribution of beliefs is higher in the high state (in likelihood ratio order) because, on average, beliefs must be correct.[37]

Note that we need the characterization from Proposition 4 to argue that $S(\ell) > S(h)$ and $D(\ell) < D(h)$. In general, we cannot prove that this is true. In particular, we have not been able to verify that θ_0^S is monotone for all δ or for the cases in which $d^\ell < d^h < s$ or $s < d^\ell < d^h$.

6.4 Equilibrium Existence

In the online Appendix, we prove that an equilibrium exists that satisfies the refinement.

Proposition 6 *There exists some $\underline{\delta} < 1$ such that for all $\delta > \underline{\delta}$ there exists an equilibrium that satisfies the refinement of monotone beliefs and that has full-trade: For some r_0 and for all θ^S and θ^B in the support of beliefs in the stocks, $\rho(\theta^S) \leq r_0 \leq \beta(\theta^B)$.*

[36] It also implies that $\theta_0^S(\theta, [a, b])$ is increasing in a and b, whenever $\theta_{(1)} \in [a, b]$ that has a positive probability.

[37] This follows from the Bayesian consistency of the distribution of posterior beliefs in the stock; see also Footnote 11.

For Proposition 6, we modify our previous construction of a full-trade equilibrium to make sure that for some off-equilibrium beliefs that satisfy the refinement of monotone beliefs, the sellers do not have an incentive to *decrease* the reserve price below the equilibrium reserve price. To do so, 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 $\beta_k(\underline{\theta}^B)$. The key idea is to have type $\bar{\theta}^S$ (the highest belief in the support) set a reserve price at which he is just indifferent to trade conditional on the winning bid being equal to $\beta_k(\underline{\theta}^B)$. Sellers with beliefs below $\bar{\theta}^S$ strictly prefer accepting $\beta_k(\underline{\theta}^B)$, and the refinement can be satisfied by assigning reserve prices strictly below $\beta_k(\underline{\theta}^B)$. The resulting equilibria has "full trade" by monotonicity of β_k.

7 Discussion

7.1 Alternative Matching and Bargaining Protocols

The analysis in Section 6 implies that under full information about market conditions, any steady-state equilibrium has a competitive limit. Therefore, in Proposition 4, the possibility of non-competitive limits arises due to the combination of incomplete information about the 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 yield competitive limits. We start the discussion by observing that, in the present model, the buyers compete directly by bidding, whereas there is no such competition between the sellers. This helps explain why the price is competitive when frictions are low in the case in which there are more buyers than sellers but the price may stay above the competitive level when there are more sellers than buyers. In particular, sellers do not compete away their profits, because they face no direct competition.[38]

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 $\varepsilon > 0$, each buyer is matched with two sellers (and with probability $1 - \varepsilon$ with one seller). Then the sellers compete, 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 its potential rival to make sure that the buyer chooses his offer. This is sufficient to force the reserve prices down to the cost of the sellers

[38] To understand the distribution of market power, note the following. The winning bid is essentially a take-it-or-leave-it offer to the seller. Without aggregate uncertainty, the seller accepts any offer that makes him better off than waiting for the next period. This limits the market power of sellers. With aggregate uncertainty, however, "belief-threats" allow the seller to rationalize rejecting low offers and commit to a high reserve price. This essentially flips the order of moves and makes the seller the proposer.

when there are more sellers than buyers in the economy.[39]

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.[40] 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.

7.2 Applications and Testable Implications

Online auction markets, especially eBay, are a natural application of our model. In particular, our model predicts that buyers increase their bids over time. Consequently, buyers may later regret not having won at a lower price earlier. Finally, buyers who have already spent more time on the market will win with higher probability. The empirical observations by Juda and Parkes (2006), discussed in the introduction, are broadly consistent with these predictions. It would be interesting to test systematically for the presence of learning on eBay.

More generally, the type of two-sided search model that we study has also been used in the literature to understand decentralized housing and labor markets; see Rogerson, Shimer, and Wright (2005). The housing market may fit well, since buyers and sellers are often small households (and, hence, have little experience with the market conditions) and some versions of an auction are indeed used in many countries.

Another application of our type of model is over-the-counter asset markets in which individual traders contact each other to bargain over terms or trade; see Duffie, Garleanu, and Pedersen (2005) and Duffie (2012).[41] Indirect evidence for the presence of aggregate uncertainty is the impact of a reform that increased the post-trade transparency of the over-the-counter U.S. corporate bond market via posting of past transaction terms (TRACE); see, for example, Bessembinder and Maxwell (2008) and the literature discussed therein. The fact that post-trade transparency affected the market outcome is consistent with the presence of uncertainty about the market condition, at least by some traders; see Duffie, Dworczak, and Zhu (2016) and our discussion of policy implications below.

[39] This discussion of direct competition is somewhat reminiscient of the role of competition for partners in Felli and Roberts (2016). A similar ε-competition model is also studied as a special case in Lauermann (2013).

[40] For example, the central arguments behind Proposition 4, namely, Lemmas 1 and 2, would extend immediately to this setting.

[41] Search models for trade of an indivisible good—like ours—are useful for studying asset markets, but the assumed indivisibility is an important restriction that should be kept in mind. The advantage of an indivisible good model is its tractability.

7.3 Policy Implications

Despite its stylized nature, our model has some—at least, suggestive—policy implications. Perhaps most significant is that our model shows that decentralized markets can, in principle, achieve the market clearing outcome, even if no one individually knows what it is. Thus, aggregate uncertainty is not a reason per se to intervene in a market. This may be seen as lending support to Hayek's conjecture. At least in theory, this was not immediately evident from existing work on decentralized markets—especially given prior negative results in related settings.[42]

Nevertheless, for large frictions, the outcome is not necessarily (constrained) efficient or market clearing. In this case, "soft" interventions in the form of information provision (transparency) may already be useful. Future extensions of our model could further allow for heterogeneity in values and costs and for costly participation decisions. In these extensions, aggregate uncertainty would imply inefficient trades and inefficient entry choices, allowing the study of interventions that reduce uncertainty. The recent contribution by Duffie, Dworczak, and Zhu (2016) studies such interventions in a model of over-the-counter asset markets. In their work, one market side knows the market conditions. Extensions of our model would allow studying such interventions in settings with more symmetric uncertainty.[43]

7.4 Discussion of Assumptions

Initial Signals.—Newly arriving traders receive an initial signal. This initial signal serves only a technical purpose to avoid the notation related to mixed bidding strategies that would otherwise be necessary. All our results continue to hold if the initial signal is uninformative.[44]

The fact that it is possible to have no external information in our setting may also clarify the relation to work on the foundation for rational expectations equilibrium with common values, as in Wolinsky (1990) and Golosov, Lorenzoni, and Tsyvinski (2014). Roughly speaking, in that literature, the question is whether information about the underlying value spreads from the initially "informed" to the "uninformed" traders through trade.

Minimal Feedback.—Information feedback in the trading protocol is minimal, so agents learn only through the failure to trade. Despite the minimal feedback, information aggregation is possible in equilibrium. In principle, it should be possible to extend the characterization results to trading protocols that reveal more information. For example, consider the

[42] Indeed, we show more: even for positive frictions, as long as they are not too large, there exists a natural class of full-trade equilibria that achieves a constrained efficient outcome, meaning that aggregate uncertainty and the war-of-attrition problem do not necessarily imply inefficiencies beyond those imposed by the decentralized meeting technology.

[43] Indeed, in the previous section, we mentioned an empirical work of Bessembinder and Maxwell (2008) documenting the effect of the introduction of post-trade transparency.

[44] A supplement with details is available on request.

central arguments behind Proposition 4, namely, Lemmas 1 and 2. These lemmas use very basic arguments that do not rely on details of the protocol such as minimal feedback.

Extending the existence results, however, would be more challenging because (i) minimal feedback enabled us to construct the steady-state stock and the strategies separately ("decoupling") and (ii) more information implies more signaling possibilities, requiring more demanding refinements for which equilibrium existence would be harder to establish.

We note that the full-trade equilibria remain equilibria if sellers observe the winning bid before making an acceptance decision (see Page 12). In addition, we establish in a companion paper that sellers have strict incentives to choose a protocol that limits information flows to buyers because the value of continued search is smaller when buyers are uncertain about the state, reducing the value of their outside option; see Lauermann and Virág (2012).

Steady-State and Learning with Overlapping Generations.—The overall stock of the market is in a steady state, and the underlying market conditions are assumed constant over time. This captures a setting in which traders believe market conditions to be stable relative to the duration of their own search. To us, this belief seems to be relevant in many search environments where there is entry by new traders over time, so that trade and learning continually take place. A model with entry in which different cohorts are continuously mixed to generate a steady distribution of beliefs may, therefore, be a realistic depiction.[45]

Moreover, despite the stationarity on the aggregate level, the history-dependent, dynamic behavior of the individuals and cohorts is at the heart of our model and analysis. In particular, potential buyers and sellers learn about market conditions through experimentation over time and change their behavior as their beliefs change. Finally, although there is no market-wide learning, individual buyers and sellers will almost surely learn the state as frictions become small—and they will do so even if there are essentially no initial signals (meaning that this is true even if all information is endogenously generated by the market).

Simple Economy with Unit Demand/Supply.— The underlying economy is a simple market with unit demand and supply and homogeneous values and costs. There are two reasons for concentrating on such a simple economy. First, the assumption that the good is indivisible captures markets for goods like houses or labor, in which traders engage mostly in one-time transactions. The indivisibility is critical to how learning and experimentation take place. Learning is different if the good is easily divisible because traders can engage in frequent transactions of incremental units; see Golosov et al. (2014). Thus, models with indivisible goods capture different economic scenarios from those with divisible goods.

There is also a technical reason for why we concentrate on the case of unit demand and

[45] A similar steady-state learning model with overlapping generations is Fudenberg and Levine (1993), albeit in a different literature.

supply. The advantage of our setup is that types are one-dimensional (beliefs about the state). This makes the analysis transparent, and also allows us to utilize available techniques for analyzing auctions and establish the existence of equilibrium. With divisible goods, the effective type becomes two-dimensional—the beliefs and the current endowment of an agent—and so one typically cannot establish equilibrium existence. It may be possible to push the characterization result further. For example, in a previous version (with a slightly different protocol and refinement), we considered a setting with heterogeneous valuations—giving rise to a downward sloping market demand function—and showed market clearing in the limit. However, our methods would not allow us to prove existence with two-dimensional types.

8 Conclusion

We study a search market in which the participants do not know the market conditions. We emphasize three contributions to the literature. First, we study the combined effects of search and learning in a novel, two-sided equilibrium search model with uncertainty about market conditions. Second, we identify tractable "full-trade equilibria" that provide insights into the economics of decentralized markets and demonstrate how bidding and learning over time takes place. Third, we test the hypothesis that even if decentralized, markets can nevertheless aggregate information that is dispersed among their participants. When frictions are small, trade takes place at the correct market clearing prices "as if" the participants actually knew demand and supply conditions. We relate the behavior that leads to this outcome to a "decentralized tâtonnement process."

For small frictions, our analysis suggests that decentralized markets may work well even with uncertainty about the market conditions, despite earlier negative results in related settings. Natural applications of our model include labor markets and housing markets, but models of this kind have also proven useful for over-the-counter asset markets.

A number of open questions remain. In particular, to understand the relation to the negative results in Wolinsky (1990) and Blouin and Serrano (2001) for rational expectations equilibrium with common values, it may be useful to study a more general model that nests both of the extreme settings; that is, their pure common value environment and our private value setting. Furthermore, future work may use our algorithm for equilibrium construction to study the market outcome for large frictions: How large is equilibrium price dispersion? How long do traders wait in the market? It would be interesting to quantify the impact of aggregate uncertainty: Can aggregate uncertainty magnify price/wage dispersion and waiting times (unemployment) relative to a market with known demand and supply conditions?

9 Appendix

Content. This Appendix contains the proofs for the general characterization results from the paper, namely, Propositions 4 and 5 and Lemma 3. All results concerning full-trade equilibria are relegated to a separate online Appendix.

Notation. We abuse notation and write $Q^B\left((\beta,\theta)\,|\omega\right)$, $Q^S\left((\rho,\theta)\,|\omega\right)$, $P^B\left((\beta,\theta)\,|\omega\right)$, etc., for the trading probabilities and the (expected) transaction prices for the allocation of type θ who follows strategy β or ρ. Let $Q^B\left(b|\omega\right)$, $Q^S\left(r|\omega\right)$, $P^B\left(b|\omega\right), Q^S\left(r|\omega\right)$ denote the trading probability and price variables for the constant bidding strategy $\beta \equiv b$ and constant reserve price strategy $\rho \equiv r$, respectively.

We continue to use the convention that lim refers to a subsequence for which the limit exists, possibly in the extended reals.

9.1 Proof of Proposition 4

We start the proof of Proposition 4 with a number of intermediate results. The first Lemma captures the implications of experimentation becoming cheap as frictions vanish.

Lemma 4 *Consider any sequence of steady-state equilibria for $\delta_k \to 1$.*

Sellers. If for some $\theta > 0$ it holds that $\lim Q_k^S(\theta|h) = 1$, 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 $\lim Q_k^S(\theta'|h) = \lim Q_k^S(\theta''|h) = 1$ for some $1 \geq \theta', \theta'' > 0$, then $p^h(\theta') = p^h(\theta'') = p^h$. Moreover, if $\lim Q_k^S(\theta|h) = 1$ for some $\theta < 1$, then $\lim W_k^S(\theta|\ell) = \lim V_k^S(0)$.

Buyers. If for some $\theta < 1$ it holds that $\lim Q_k^B(\theta|\ell) = 1$, then the realized transaction price of a buyer with belief θ converges in distribution to a single price $p^\ell(\theta) \in [c, v]$ in the low state. If $\lim Q_k^B(\theta'|\ell) = \lim Q_k^B(\theta''|\ell) = 1$ for some $0 \leq \theta', \theta'' < 1$, then $p^\ell(\theta') = p^\ell(\theta'') = p^\ell$. Moreover, if $\lim Q_k^B(\theta|\ell) = 1$ for some $\theta > 0$, then $\lim W_k^B(\theta|h) = \lim V_k^B(1)$.

Proof. We prove the lemma for the sellers' side only. Fix $\theta, \theta' > 0$ such that $\lim Q_k^S(\theta|h) = \lim Q_k^S(\theta'|h) = 1$ and let $p^h(\theta) = \lim P_k^S(\theta|h)$.

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

Given $\lim Q_k^S(\theta|h) = 1$, 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.[46] 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 state h,

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

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 $\lim \delta_k^{z_k} = 1$.[47] 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 $q_k \geq \chi_k(p)$,[48] 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 $\lim V_k^S(0)$.[49]

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

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 $\lim Q_k^S(\theta'|h) = 1$. 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.[50] 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

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

[47] 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.

[48] 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.

[49] 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.

[50] 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.

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.*

Proof of Lemma 1. Suppose $\lim Q_k^S(1|h) < 1$. We want to show that $\lim \rho_k(1) = \lim \beta_k(1) = c$. From $\lim Q_k^S(1|h) < 1$, it follows that $\lim 1 - q_k^S(\rho_k(1)|h) = \lim q_k^B(\rho_k(1)|h) = 1$ (almost no buyer bids above $\rho_k(1)$). Hence, $\lim Q_k^B(\rho_k(1)|h) = 1$, so optimality of $\beta_k(1)$ implies $\lim \beta_k(1) \le \lim \rho_k(1)$. From undominatedness (8) and $V_k^S(1) \le \beta_k(1) - c$, we have $\rho_k(1) - c \le \delta_k(\beta_k(1) - c)$, and hence $\rho_k \le \beta_k(1)$ for all all k. Hence, $\lim \beta_k(1) = \lim \rho_k(1)$. It follows from the monotonicity of β_k that $\lim P_k^S(1|h) = \lim \rho_k(1)$. From the definition of V_k^S and (8), 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 $\lim \beta_k(1) = \lim \rho_k(1) = c$. *Q.E.D.*

Proof of Lemma 2. Suppose that $\lim D_k(\ell)/S_k(\ell) > 0$. We show that

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

which proves the lemma.

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

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

which implies that

$$\lim Q_k^S(\theta = 0|\ell) = 1. \tag{16}$$

Let $p^\ell := \lim P_k^S(0|\ell)$. From (8) and (16), $\lim \rho_k(0) = p^\ell$. Of course,

$$\lim Q_k^S(\rho_k(0) + \varepsilon|\ell) < 1, \tag{17}$$

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

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

where the first equality follows from definition of q_k^S and the second follows from (17). So, $\lim e^{-\mu_k^\ell} < 1$ and $\lim e^{-\mu_k^\ell(1-\Gamma_k^B(\beta^{-1}(b_k)|\ell)} = 1$. Hence,

$$\begin{aligned} \lim \frac{q_k^B(b_k + \varepsilon|\ell)}{q_k^B(b_k|\ell)} := \alpha &= \lim \frac{e^{-\mu_k^\ell(1-\Gamma_k^B(\beta^{-1}(b_k)|\ell)}\Gamma_k^S\left(\rho_k^{-1}(b_k + \varepsilon)|\ell\right)}{e^{-\mu_k^\ell}\Gamma_k^S\left(\rho_k^{-1}(b_k)|\ell\right)} \\ &\ge \lim \frac{e^{-\mu_k^\ell(1-\Gamma_k^B(\beta^{-1}(b_k)|\ell)}}{e^{-\mu_k^\ell}} > 1, \end{aligned} \tag{18}$$

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

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

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

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

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

$$\lim \frac{Q_k^B(b_k+\varepsilon|\ell)}{Q_k^B(b_k|\ell)}\frac{v-b_k-\varepsilon}{v-b_k} > 0. \tag{19}$$

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

We say that the law-of-one-price holds if the distribution of the realized transaction price converges in distribution to a point for almost all types in the inflow.

Lemma 5 *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*

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

Proof. Consider $\omega = h$.

Suppose $\lim \beta_k(1) > c$. Then, the monotonicity of ρ_k and Lemma 1 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 1, $\beta_k(1) \geq \rho_k(1)$, 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 $\theta \in \left[\underline{\theta}^B, \bar{\theta}^B\right]$—in contradiction to mass balance given $d^h > s$: Namely, feasibility requires that

$$d^\omega \int_{[0,1]} Q_k^B(\theta|\omega)g^B(\theta|\omega)\,d\theta = s\int_{[0,1]} Q_k^S(\theta|\omega)g^S(\theta|\omega)\,d\theta. \tag{20}$$

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 $\omega = h$.

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 4 (the distribution of trading prices collapses). Thus, the law of one price holds if $\omega = h$.

Consider $\omega = \ell$.

Case 1: $\lim D_k(\ell)/S_k(\ell) > 0$. Then, from Lemma 2, either $\lim Q_k^B(\beta_k(0)|\ell) = 1$ or $\lim \beta_k(0) = v$ (or both). If $\lim Q_k^B(\beta_k(0)|\ell) = 1$, then by monotonicity of β_k, $\lim Q_k^B(\theta|\ell) = 1$ for all θ, so the law-of-one-price holds. If $\lim \beta_k(0) = v$, then this and $\lim D_k(\ell)/S_k(\ell) > 0$ implies that sellers trade for sure when setting $r = v - \varepsilon$ for any $\varepsilon > 0$, $\lim Q_k^S(v - \varepsilon|\omega) = 1$ for all ε and $\omega = \ell, h$. 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 (20) given $d^\ell < s$. Thus, trading probabilities are competitive in case 1.

Case 2: $\lim D_k(\ell)/S_k(\ell) = 0$. Then, $\lim Q_k^B(\theta|\ell) = 1$ 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 $\lim Q_k^B(\theta^B|\ell) < 1$. Then there exists a sequence t_k with $\lim \delta_k^{t_k} < 1$ 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 $\lim \delta_k^{t_k} < 1$ is equivalent to $\lim t_k(1-\delta_k) > 0$, we obtain that $\lim D_k(\ell)(1-\delta_k) > 0$. But $S_k(\ell)(1-\delta_k) \le s$—Contradiction. Thus, $\lim Q_k^B(\theta|\ell) = 1$ for almost all θ in case 2. Therefore, trading probabilities are competitive in both possible cases.

Take some $\theta \in \left[\underline{\theta}^B, \bar{\theta}\right]$ 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 4. Thus, the law of one price holds if $\omega = \ell$.

Characterization of Payoffs. Consider $\omega = h$. From competitive trading probabilities and the law of one price, sellers' expected payoffs for almost all $\theta \in \left[\underline{\theta}^S, \bar{\theta}\right]$ are $\lim EU^S(\theta|h) = \left(p^h - c\right)$. Consider a buyer having type $\theta \in \left[\underline{\theta}^B, \bar{\theta}\right]$ such that $\lim Q_k^B(\theta|\ell) = 1$. From Lemma 4 and $\lim Q_k^B(\theta|\ell) = 1$ for almost all $\theta \in \left[\underline{\theta}^B, \bar{\theta}\right]$, we have $\lim EU^B(\theta|h) = \lim V_k^B(1)$ for almost all $\theta \in \left[\underline{\theta}^B, \bar{\theta}\right]$. 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, $\lim V_k^B(1) = \bar{Q}\left(v - p^\ell\right)$. Thus, for almost all θ, we have $\bar{Q} = \lim Q_k^B(\theta|h)$. Finally, from the mass balance requirement, $\lim d^h \int_{\underline{\theta}^B}^{\bar{\theta}} Q_k^B(\theta|h)\, d\theta = s$, and so we have $\lim Q_k^B(\theta|h) = \frac{d^h}{s}$ for almost all θ, proving $\lim EU^B(\theta|h) = \frac{d^h}{s}\left(v - p^h\right)$.

An analogous argument establishes the same for $\omega = \ell$. Together, this proves the characterization of payoffs. *Q.E.D.*

Proof of Proposition 4.

Proof of Proposition 4 for the case where $d^\ell < s < d^h$.

Lemma 5 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 5.

Suppose, $\lim \beta_k(1) = c$. As argued before in the proof of Lemma 5, if $\lim \beta_k(1) = c$ then $\lim V_k^B(\theta) = v - c$ for all θ. This requires $\lim Q_k^B(\theta|\omega) = 1$ for almost all $\theta \in [\underline{\theta}^B, \bar{\theta}]$ and $\omega = \ell, h$. This violates mass balance (20) since $s < d^h$. Thus, $\lim \beta_k(1) > c$. Then, Lemma 1 implies that $\lim Q_k^S(1|h) = 1$. So, Lemma 4 requires $P_k^S(1|h) = P_k^S(\theta|h)$ for all θ and from Lemma 5, $P_k^S(1|h) = p^h$. Thus, by Lemma 4, 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 (8), $P_k^S(1|h) \to p^h$ 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 5 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 4 for the case where $s < d^\ell < d^h$.

Consider $\omega = \ell$. From $d^\ell > s$, mass balance requires $\lim D_k(\ell)/S_k(\ell) > 0$. Thus, from Lemma 2, either $\lim Q_k^B(\beta_k(0)|\ell) = 1$ or $\lim \beta_k(0) = v$ (or both). Mass balance and $s < d^\ell$ prohibits $\lim Q_k^B(\beta_k(0)|\ell) = 1$ (since then $\lim Q_k^B(\theta|\ell) = 1$ for all θ). Thus, $\lim \beta_k(0) = v$. As argued before in the proof of Lemma 5: This and $\lim D_k(\ell)/S_k(\ell) > 0$ implies that sellers trade for sure when setting $r = v - \varepsilon$ for any $\varepsilon > 0$, that is, $\lim Q_k^S(v - \varepsilon|\omega) = 1$ for all ε and $\omega = \ell, h$. 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 4 for the case where $s > d^h > d^\ell$.

Suppose $\lim \beta_k(1) > c$. Then, the monotonicity of ρ_k and Lemma 1 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 5, when $\lim \beta_k(1) = c$ then $\lim V_k^B(\theta) = v - c$ for all θ. This requires $\lim Q_k^B(\theta|\omega) = 1$ for all $\theta \in [\underline{\theta}^B, \bar{\theta}]$ and $\omega = \ell, h$. Thus trading probabilities are competitive. Of course, if $\lim \beta_k(1) = c$, then $\lim P_k^B(\theta|\omega) = \lim P_k^S(\theta|\omega) = c$ (the law of one price holds). The characterization of payoffs follows from $p^h = p^\ell = c$.

This completes the proof of Proposition 4. *Q.E.D.*

9.2 Proof of Proposition 5

Remark: The proof of the Proposition uses Lemma 3 (monotone updating on-path) and some of the intermediate steps of its proof, stated in Section 9.3.

Lemma 6 *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* $\{b_k\}$ *with* $\beta_k^{-1}(b_k) \in$*supp* Γ_k^B *and* $b_k = \rho_k(\theta_k) \to b \in [c, v]$. *Then*

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

Proof. Suppose Lemma 6 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 5.

We need to show that $p^\ell = c$, where p^ℓ is as defined in Proposition 4 for the case where $d^h > s > d^\ell$.

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

By Proposition 4, and monotonicity of the bid functions and updating $\lim Q_k^B(x|\ell) = 1$ for all $x \geq \underline{\theta}^B$. Therefore, Lemma 4 implies that for $\omega = \ell$, almost all buyers trade in the limit at the price p^ℓ. Hence, so do the sellers, meaning, $\lim P_k^S(z|\ell) = p^\ell$ for almost all $z \in (0, 1)$. By Proposition 4, $\lim Q_k^S(z|h) = 1$, and thus Lemma 4 implies $\lim(W_k^S(z|\ell) = \lim V_k^S(0)$ for all $z \in (0, 1)$. So, the hypothesis implies $\lim(W_k^S(z|\ell) = p^\ell - c$ for almost all z. Since $\lim P_k^S(z|\ell) = p^\ell$, this implies that either $p^\ell = c$ or $\lim Q_k^S(z|\ell) = 1$ for almost all $z \in (0, 1)$— but $\lim Q_k^S(z|\ell) = 1$ for almost all $z \in (0, 1)$ cannot hold by feasibility if $s > d^\ell$. Hence, $p^\ell = c$, which implies the claim.

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

Take any $z \in \left(\underline{\theta}^B, \bar{\theta}^B\right)$. From Proposition 4, the monotonicity of β_k, and Lemma 8 (monotone updating), $\lim Q_k^B(z|\ell) = 1$ and $\lim Q_k^B(z|h) < 1$. 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_+).[51] Such t_k exists for k large enough. From Proposition 4, $\lim b_k = p^\ell$. Also, $(\delta_k)^{t_k} \to 1$. 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

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

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

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 $x_k \leq \bar{\theta}^S$ and $b_k = \beta_k(\theta_k)$ (as defined above),

$$\lim \theta_0^S(x_k, b_k) = 0. \tag{21}$$

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

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

Note that the distribution of beliefs must satisfy

$$\frac{D_k(h)}{D_k(\ell)} \frac{\gamma_k^B(\theta_k|h)}{\gamma_k^B(\theta_k|\ell)} = \frac{\theta_k}{1 - \theta_k}. \tag{22}$$

This follows from the Bayesian consistency of the distribution of posterior beliefs.[52] 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 2, 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 $(\delta_k)^{t_k} \to 1$ and $\lim Q_k^B(\theta|h) < 1$ for all $\theta \leq \bar{\theta}$, we have $\lim(1 - \Gamma_k^B(\theta_k|h)) = 0$. Finally, by Lemma 5 it holds that $\lim \frac{D_k(h)(1-\delta_k)}{S_k(h)} = z > 0$. Therefore, $\lim \frac{e^{-\frac{D_k(h)}{S_k(h)}(1-\Gamma_k^B(\theta_k|h))}}{1-\delta_k} = \lim \frac{e^{-\frac{z}{1-\delta_k}}}{1-\delta_k} = 0$, and so (21) 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 $b \in [\lim V_k^S(0) + c + \varepsilon, p^\ell]$, $\beta_k^{-1}(b) \in \text{supp } \Gamma_k^B$ for a large enough k. Otherwise, there exists an interval $(\underline{b}, \overline{b}) \subset [\lim V_k^S(0) + c + \varepsilon, p^\ell]$ that is not in the limit of the supports of the bid distributions but $\overline{b}$ is. Then for it to be optimal to bid close to $\overline{b}$ in the limit, it must hold that $\overline{b}$ is in the limit of the supports of the reserve price distributions.

[52] We explicitly verify that steady-state distributions satisfy this condition in Lemma 11 in the online appendix for the case of full-trade equilibria.

But then Lemma 6 implies 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 support of the reserve price distribution. Then again Lemma 6 yields a contradiction.

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 $\omega = \ell$. Hence, given Step 2, conditional on $\omega = \ell$, 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 **5**. *Q.E.D.*

9.3 Monotone Updating and Proof of Lemma 3

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 $\frac{\Gamma_{(1)}^B(\theta|h)}{\Gamma_{(1)}^B(\theta|\ell)}$ and $\frac{\gamma_{(1)}^B(\theta|h)}{\gamma_{(1)}^B(\theta|\ell)}$ are strictly increasing in θ for all $\theta \geq \underline{\theta}^B$ on the support of $\Gamma_{(1)}^B$.*

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

$$\gamma_{(1)}^B(\theta|\omega) = \frac{D(\omega)}{S(\omega)}\gamma^B(\theta|\omega)\,\Gamma_{(1)}^B(\theta|\omega).$$

The likelihood ratio can be written as

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

From (22),

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

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 $\frac{\gamma^B_{(1)}(\theta|h)}{\gamma^B_{(1)}(\theta|\ell)} > \frac{\Gamma^B_{(1)}(\theta|h)}{\Gamma^B_{(1)}(\theta|\ell)}$, so $\frac{\Gamma^B_{(1)}(\theta|h)}{\Gamma^B_{(1)}(\theta|\ell)}$ is increasing in θ because

$$\left(\frac{\Gamma^B_{(1)}(\theta|h)}{\Gamma^B_{(1)}(\theta|\ell)}\right)' = \frac{\gamma^B_{(1)}(\theta|h)\,\Gamma^B_{(1)}(\theta|\ell) - \gamma^B_{(1)}(\theta|\ell)\,\Gamma^B_{(1)}(\theta|h)}{\left(\Gamma^B_{(1)}(\theta|\ell)\right)^2}$$

$$= \frac{\gamma^B_{(1)}(\theta|\ell)}{\Gamma^B_{(1)}(\theta|\ell)}\left(\frac{\gamma^B_{(1)}(\theta|h)}{\gamma^B_{(1)}(\theta|\ell)} - \frac{\Gamma^B_{(1)}(\theta|h)}{\Gamma^B_{(1)}(\theta|\ell)}\right).$$

Moreover, $\frac{\gamma^B_{(1)}(\theta|h)}{\gamma^B_{(1)}(\theta|\ell)} = \frac{\frac{D(h)}{S(h)}\theta\Gamma^B_{(1)}(\theta|h)}{\frac{D(\ell)}{S(\ell)}(1-\theta)\Gamma^B_{(1)}(\theta|\ell)}$ 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 $t(\omega)$ for $\omega = \ell, h$. Then $S(\omega) = \frac{s-t(\omega)}{1-\delta}$ in a steady-state.[53] By feasibility $t(\ell) \le d^\ell$, 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 $\frac{\Gamma^B_{(1)}(\theta|h)}{\Gamma^B_{(1)}(\theta|\ell)}$ is increasing for all $\theta \ge \underline{\theta}^B$. *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 the 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 $F^\omega_b(x) = \Gamma^B_{(1)}(\beta^{-1}(x)|\omega)$ denote the probability that the highest bid is less than or equal to x in state ω. Similarly, let $F^\omega_r(x) = \Gamma^S(\rho^{-1}(x)|\omega)$ denote the probability that the reserve price set is less than or equal to x in state ω. Given previous results, it holds that $F^\omega_b(x) = e^{-\frac{D(\omega)}{S(\omega)}(1-\Gamma^B(\beta^{-1}(x)|\omega))}$. In the main text we show that $F^\omega_b(x) = \Gamma^B_{(1)}\left(\beta^{-1}(x)|\omega\right)$ and $F^\omega_r(x) = \Gamma^S\left(\rho^{-1}(x)|\omega\right)$. By definition,

$$\frac{\theta^B_+(\theta,b)}{1-\theta^B_+(\theta,b)} = \frac{\theta}{1-\theta}\frac{1-F^h_b(b)F^h_r(b)}{1-F^\ell_b(b)F^\ell_r(b)}.$$

[53] 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 $F_b^\omega(r) = \Gamma_{(1)}^B\left(\beta^{-1}(r)|\omega\right)$ 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 analogue of (22) for the sellers implies 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.[54] *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 $\Gamma_{(1)}^B(r'|\omega) > \Gamma_{(1)}^B(r|\omega)$, then $\theta_+^S(\theta, r') > \theta_+^S(\theta, r)$.

Proof. Recall that $F_b^\omega(x) = \Gamma_{(1)}^B(\beta^{-1}(x)|\omega)$. 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.*

[54] 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 3.

Proof of Lemma 3.

Take any $b' > b \geq c$ such that $\beta^{-1}(b')$ and $\beta^{-1}(b)$ are in the support of Γ^B. We want to show that $\theta_0^S(\theta, b) > \theta_0^S(\theta, b')$. Let θ^B and $\widetilde{\theta}^B$ be such that $b = \beta(\theta^B)$ and $b' = \beta(\widetilde{\theta}^B)$. 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.

The posterior if no bid is received is $\theta_0^S(\theta, \varnothing)$. We want to show that for all $b \geq 0$ such that $\beta^{-1}(b) \in supp(\Gamma^B)$, $\theta_0^S(\theta, b) \geq \theta_0^S(\theta, \varnothing)$. To establish this, 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 ω is $\Gamma_{(1)}^B\left(\underline{\theta}^B|\omega\right)$. 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.*

References

Ausubel, L. M., P. Cramton, and R. J. Deneckere (2002). Bargaining with incomplete information. *Handbook of game theory with economic applications 3*, 1897–1945.

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

Benabou, R. and R. Gertner (1993). Search with learning from prices: does increased inflationary uncertainty lead to higher markups? *The Review of Economic Studies 60*(1), 69–93.

Bessembinder, H. and W. Maxwell (2008, Spring). Markets: Transparency and the corporate bond market. *Journal of Economic Perspectives 22*(2), 217–234.

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 price dispersion. *Econometrica 51*(4), 955–969.

Cripps, M. W. and J. M. Swinkels (2006). Efficiency of large double auctions. *Econometrica 74*(1), 47–92.

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. (2012). *Dark markets: Asset pricing and information transmission in over-the-counter markets*. Princeton University Press.

Duffie, D., P. Dworczak, and H. Zhu (2016). Benchmarks in search markets. *Journal of Finance Forthcoming*.

Duffie, D., N. Garleanu, and L. H. Pedersen (2005). Over-the-counter markets. *Econometrica 73*(6), 1815–1847.

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.

Felli, L. and K. Roberts (2016). Does competition solve the hold-up problem? *Economica 83*(329), 172–200.

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 (2014). Decentralized trading with private information. *Econometrica 82*(3), 1055–1091.

Golub, B. and M. O. Jackson (2010). Naive 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), 519–530.

Janssen, M., A. Parakhonyak, and A. Parakhonyak (2014). Non-reservation price equilibria and consumer search. Mimeo.

Janssen, M. and S. Shelegia (2015). Consumer search and double marginalization. *American Economic Review 105*(6), 1683–1710.

Jehiel, P. and L. Lamy (2015). On absolute auctions and secret reserve prices. *The RAND Journal of Economics 46*(2), 241–270.

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*, 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. (2012). Asymmetric information in bilateral trade and in markets: An inversion result. *Journal of Economic Theory 147*(5), 1969–1997.

Lauermann, S. (2013). Dynamic matching and bargaining games: A general approach. *The American Economic Review 103*(2), 663–689.

Lauermann, S. and G. Virág (2012). Auctions in markets: Common outside options and the continuation value effect. *American Economic Journal: Microeconomics 4*(4), 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 20*(2), 89–114.

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

Mortensen, D. T. (1982). The matching process as a noncooperative bargaining game. In *The Economics of Information and Uncertainty*, pp. 233–258. University of Chicago Press.

Myerson, R. B. (1998). Extended poisson games and the condorcet jury theorem. *Games and Economic Behavior 25*(1), 111–131.

Ostrovsky, M. (2012). Information aggregation in dynamic markets with strategic traders. *Econometrica 80*(6), 2595–2647.

Pesendorfer, W. and J. Swinkels (1997). The loser's curse and information aggregation in common value auctions. *Econometrica 65*(6), 1247–1282.

Pesendorfer, W. and J. Swinkels (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.

Rogerson, R., R. Shimer, and R. Wright (2005). Search-theoretic models of the labor market: A survey. *Journal of economic literature 43*(4), 959–988.

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.

Satterthwaite, M. and A. Shneyerov (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.

Smith, L. (2011). Frictional matching models. *Annual Review of Economics 3*(1), 319–38.

Smith, L. and P. N. Sorensen (2013). Rational social learning by random sampling. *Available at SSRN 1138095*.

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

Wolinsky, A. (1990, January). Information revelation in a market with pairwise meetings. *Econometrica 58*(1), 1–23.