Bargaining and Competition in Thin Markets

Francesc Dilmé[*]

Spring 2018

Abstract

This paper analyzes price dynamics and trade delay in an endogenously-evolving thin market. We obtain that, differently from large markets, thin markets may feature trade delay even when buyers and sellers are homogeneous, and we provide necessary and sufficient conditions for it to vanish as the bargaining frictions disappear. Under such conditions, different pricing mechanisms (like decentralized bargaining or centralized price-posting) generate very similar transaction prices: they are approximately proportional to the expected discounted future time where the market has an excess demand. This is shown to imply that prices drift towards the price in a balanced market, and that increments on the interest rate generate mean-preserving spreads of their ergodic distribution.

Keywords: Thin markets, Decentralized bargaining.
JEL Classifications: C73, C78, D53, G12.

[*]University of Bonn. `fdilme@uni-bonn.de`. Previous versions of the paper were titled "Dynamic Asset Trade a la Bertrand". I thank Syed Nageeb Ali, Stephan Lauermann, Peter Norman and Caroline Thomas for their useful comments, as well as the audiences in the EEA 2016 conference in Geneva and the Workshop Decentralized Markets with Informational Asymmetries 2016 (Turin), and the seminar workshops in UC Berkeley and UNC Chapel Hill.

1 Introduction

Many markets are thin: they have, at any given moment in time, a small but permanently changing set of traders. Examples include housing/rental markets in given locations, job markets for specific occupations, or some over-the-counter (OTC) financial markets. In such markets, the market composition evolves endogenously and stochastically over time. When traders trade they typically exit the market, and new traders enter the market when they can offer or need goods or services.

In a thin market, at each given instant in time, traders in the market have a limited set of current trading opportunities. Still, even though the possibility of waiting for arrivals of new traders is costly, it also serves them as an outside option. As a result, the market (or bargaining) power of the traders in the market and the realized transaction prices not only depend on the pricing mechanism or the current trade opportunities, but also on the endogenous expectation about the future ones.[1]

This paper analyzes how the pricing mechanism and arrival process affect the trade outcome (efficiency, timing and price of transactions,...) in thin markets, and compares its features to those of the trade outcomes of large markets. We obtain that the trade outcome of a thin market features two main departures from the usual "large markets" literature findings. First, owed to the endogeneity of the future market evolution, thin markets may feature inefficient trade delay even when buyers and sellers are homogeneous. Second, owed to the stochastic evolution of their composition, they feature price dispersion even in the limit where bargaining frictions disappear. We characterize the resulting price dynamics and show that they are similar across different pricing mechanisms.

Our model is a thin-market version of Gale (1987). It consists, at any given moment in time, of a finite number of sellers who own one unit of a homogeneous indivisible good, and a finite number of homogeneous buyers with a unit demand. Once in the market, each trader keeps meeting traders from the other side of the market. In our base model, within each meeting, one of the traders is randomly chosen to make a take-it-or-leave-it offer, and the other trader accepts or rejects it. We later show that our results apply to any general Nash-bargaining outcome or to a centralized price-posting mechanism. The arrival process of buyers and sellers, the matching rate and the probability of making offers are allowed

[1]Depending on the market, the price mechanism ranges from a decentralized one-to-one bargaining (job markets, OTC markets), decentralized one-to-many auctions (real state markets), or centralized price posting (online rental markets such as AirBnB.com).

to depend on the market composition, that is, the numbers of buyers and sellers in the market. We focus on symmetric Markov strategies using the market composition as state variable.

For some parameter values, the delay in trade is significantly large, even if the frequency with which traders in the market meet is very high. This happens because single transactions affect the future composition of a thin market and, as a result, some traders may benefit from letting other traders trade first. To see this, consider the following situation. Assume that, at some point in time, sellers benefit if another seller trades, as the transaction may make the market more visible and attract new buyers into it. Assume also that, while no transaction occurs, additional sellers may enter the market, so buyers in the market may benefit from waiting for the sellers' competition to increase. In this situation, the transaction price cannot be high, because waiting provides buyers with a reasonably high outside option. Also, if sellers would expect other sellers to trade fast, they would have the incentive to let others trade and sell at a high price afterwards. As a result, immediate trade is not possible. In this case, in equilibrium, the sellers in the market engage in a war of attrition, where each of them delays trade in the hope that other sellers will trade.

We provide conditions for the equilibrium trade delay to shrink to zero as the meeting frequency increases. They take the form of bounds on the effect that current transactions have on future arrivals. In this limit, for each composition of the market, the transaction prices converge to a unique *market price*, which is independent of who makes the offer. When the meeting frequency is high, traders on the long side of the market Bertrand-compete: the market price equals their (endogenous) outside option, which is given by letting the short-side of the market clear and waiting to trade until the market is balanced. When the market is balanced, the market price equals the seller's payoff in the stochastic Rubinstein bargaining game played by the last buyer and seller traders left before the market clears. This outcome is determined by their endogenous outside options of not trading and waiting for arrival of new traders.

We use our results to provide a simple characterization of the equilibrium transaction prices. If a transaction takes place at a given moment in time, the transaction price is proportional to the discounted future time the market exhibits excess demand, adjusted by a factor proportional to the discounted time the market is balanced. This has two implications for the price process. The first is that, even though the market composition may not drift towards being balanced, the market price always moves, in expectation, towards the price of the balanced market: on average, it increases when there is excess supply, and it decreases when there is excess demand. The second implication is that a decrease in the interest (or discount) rate reduces the price dispersion: for a given size of the excess supply (demand) in the market, a lower interest rate lowers the discounted time it takes for the market to clear, and therefore increases (depresses) the price.

We study the limit where traders become patient. This limit can be reinterpreted as the result of the unification of similar markets, which results on a proportional increase of the arrival rates of both sellers and buyers. Since waiting to trade with future arrivals is cheap, the effective market accessible to each trader increases, so the equilibrium outcome approaches that of a large market. In this limit, the distribution of transaction prices degenerates towards a "competitive price", which is proportional to the ergodic probability of the market exhibiting excess demand. Thus, in a thin market, the absence of both bargaining and delay frictions does not necessarily imply that the surplus from trade is fully captured by any side of the market, as the endogenous arrival of traders may prevent the market to become extremely unbalanced.

Our results are robust to several extensions of our model. We first show that they hold for arrival processes following a general multi-dimensional Markov chain, with some exogenously-evolving components (such as the economic cycle of the economy) and some endogenously evolving components (such as idiosyncratic demand or supply shocks). We show that the conditions for no trade delay, as well as the features of the price dynamics, have the same qualitative features as in our base model. Second, we extend our results to letting the outcome of each meeting arise from a general Nash bargaining, and also letting the pricing mechanism be centralized price-posting by the sellers. In both cases, we obtain that changing the pricing mechanism only affects the price through changing the adjusting factor when the market is balanced, and hence it is irrelevant if the market is rarely balanced.

The organization of the paper is as follows. After this introduction, we review the literature related to our paper. Section 2 introduces our model, and Section 3 characterizes Markov perfect equilibria that exhibit no trade delay when the bargaining frictions are small. In Section 4 we characterize the parametric conditions that guarantee that trade delay shrinks as the bargaining frictions disappear. Finally, Section 5 discusses general arrival processes and price mechanisms, and Section 6 concludes. The Appendix provides the proofs of the results in the previous sections.

1.1 Literature review

Our paper contributes to the literature on thin markets with stochastic arrival of traders. The paper closest to ours, Taylor (1995), analyzes a centralized market where buyers and sellers arrive over time at constant rates. In every period, agents on the short side of the market make offers, while each side makes an offer with probability $\frac{1}{2}$ when the market is balanced. Coles and Muthoo (1998) consider a similar model where buyers and sellers arrive in pairs, and they allow for heterogeneity in both buyers and goods. Similarly, Said (2011) studies dynamic market in which buyers compete in a sequence of private-value second-price auctions for differentiated goods. These papers analyze price dynamics under different price mechanisms in centralized markets with constant arrival rates. Our focus is, instead, on analyzing decentralized bargaining with an endogenous arrival process. We characterize the effect that the arrival process and bargaining asymmetries have on the price dynamics and trade delay. This allows us to compare our results with some of the literature on big markets (see below).

Our paper is also related to the extensive literature on bargaining and matching in large markets, reviewed in Osborne and Rubinstein (1990) and Gale (2000).[2] Models in this literature typically contain a continuum of agents and feature non-stochastic population dynamics, many times assumed to be in a stationary state. We focus, instead, how the arrival process of traders and the bargaining protocol affect trade delay and price dynamics in thin markets. We compare (in Section 3.3) the limit where traders become patient, which can be interpreted as the market growing by replication, to the results on convergence to the competitive outcome of this literature.

Finally, there has been some recent interest on thin markets in a network of agents. For example, Condorelli, Galeotti, and Renou (2016), Elliott and Nava (2016) and Talamàs (2016) look at bargaining in networks without arrival and with replacement, and allow for differences in the valuation of the good by sellers and buyers. Our analysis, instead, focusses on understanding how the dynamics of the population determines the price process in a homogeneous complete network (so, in the limit where bargaining frictions disappear, in each moment in time, there is a unique market price) and, in particular, the properties of the resulting long-run distribution of transaction prices.

[2]Important contributions in this literature are Rubinstein and Wolinsky (1985), Gale (1987), Burdett and Coles (1997), Shimer and Smith (2000), Atakan (2006), Satterthwaite and Shneyerov (2007), Manea (2011) and Lauermann (2012).

2 The model

In this section we introduce a model similar to Rubinstein and Wolinsky (1985) and Gale (1987). The main distinguishing feature of our model is that the market is assumed to be "small", that is, the number of traders in the market at each moment in time is a non-negative integer number (instead of a mass), which stochastically changes over time.

State of the market. Time is continuous with an infinite horizon, $t \in \mathbb{R}_+$. There is an infinite number of potential (male) buyers and (female) sellers. At a given moment in time t, there are $B_t \in \mathbb{Z}_+$ buyers and $S_t \in \mathbb{Z}_+$ sellers in the market. The state (of the market) at time t is defined to be (B_t, S_t).

Arrival process. Buyers arrive in the market at a Poisson rate $\gamma_b \equiv \gamma_b(B_t, S_t) \in \mathbb{R}_+$, and sellers arrive in the market at a Poisson rate $\gamma_s \equiv \gamma_s(B_t, S_t) \in \mathbb{R}_+$. The total rate at which the state exogenously changes is denoted $\gamma \equiv \gamma_b + \gamma_s$. Section 5.1 considers a more general arrival process.

Bargaining. In our base model we focus, for the sake of clarity, on a simplistic (yet canonical) bargaining protocol. As it is pointed in Section 5.2, our results can be straightforwardly generalized to allowing for general Nash bargaining, and also to alternative price mechanisms such as price posting.

If, at time t, $B_t, S_t > 0$, meetings occur at a Poisson arrival rate $\lambda \equiv \bar{\lambda}\, \ell(B_t, S_t)$, where $\bar{\lambda} \in \mathbb{R}_{++}$ is a constant, and $\ell(B_t, S_t) \in \mathbb{R}_{++}$. When a meeting occurs, nature selects one of the buyers and one of the sellers in the market uniformly randomly, and also chooses the trader who makes a price offer. The probability that the seller is chosen is $\xi(B_t, S_t) \in (0, 1)$.[3] The trading counter-party decides then whether to accept the offer or not. If the offer is accepted, transaction happens and the traders leave the market, while if it is rejected they continue in the market.

Payoffs. Both buyers and sellers discount the future at rate $\rho > 0$. If a buyer and a seller trade at time t at price p they obtain, respectively, $e^{-\rho t}(1-p)$ and $e^{-\rho t} p$. If they never trade they both obtain 0. Both buyers and sellers are risk-neutral and expected-utility maximizers. Even though the formal expressions for the payoffs (and the conditions for the optimality of a strategy profile) are obtained using standard recursive analysis, their length makes it convenient to leave them to Appendix A.1.

[3] In order to avoid Diamond's paradox (Diamond (1971)) we assume that both types of trader make an offer with a positive probability (see Remark 3.1).

Strategies. To simplify the model setting, we focus directly on Markov strategies, the state variable being the state of the market. Thus, the strategy of an agent (buyer or seller) is a maps each state $(B, S) \in \mathbb{N}^2$ into a price offer distribution in $\pi(B, S) \in \Delta(\mathbb{R}_+)$, and a probability of acceptance for each offer received, $\alpha(\cdot; B, S) : \mathbb{R}_+ \to [0, 1]$, interpreted to be his/her strategy in the bargaining stage when he/she is matched and the market state is (B, S).[4]

Equilibrium concept. We focus, for each $\varepsilon > 0$, on *symmetric perfect (Markov) ε-equilibria*, ε-MPE henceforth, that is, equilibria where all agents on each side of the market use the same strategy, and their strategy is optimal in each state up to ε (see Appendix A.1 for the formal definition). Focussing on this equilibrium set (which includes standard symmetric MPE) is technically convenient and, as we will see, will provide us with a unique prediction in the limit where the bargaining frictions disappear and ε tends to 0.

3 Equilibrium analysis

3.1 Equilibria without trade delay

We devote this section to studying equilibrium behavior when trade delay disappears as the bargaining frictions vanish. As we will see in Sections 3.2 and 3.3, under such conditions we can characterize the price dynamics and obtain comparative statics results. Section 4 establishes general conditions for no trade delay, and provides an example of a setting where trade delay remains (in all equilibria of the model) even when the bargaining frictions disappear.

Fix a strategy profile. The *trade delay* at state $(B, S) \in \mathbb{N}^2$, denoted $T(B, S) \in \overline{\mathbb{R}}_+$, is the expected delay for the next transaction to happen. For each $\varepsilon > 0$, we use $\bar{T}^\varepsilon(B, S) \in \overline{\mathbb{R}}_+$ to denote the (superior) limit of maximum trade delay at state $(B, S) \in \mathbb{N}^2$ among all ε-MPE as $\bar{\lambda} \to \infty$. As it will become apparent in Section 4.1, for some choices of the values of the primitives of our model, there is trade delay in all ε-MPEa for ε small enough even when the bargaining frictions disappear, that is, $\bar{T}^\varepsilon(B, S) > 0$ for some $(B, S) \in \mathbb{N}^2$. The following condition will allow us to focus, in this section, on equilibria without trade

[4] We implicitly assume traders observe the state of the market. Markov perfect equilibria (see the definition below) remain equilibria independently of the information structure as long as the current state of the market is known to the traders in the market.

delay:

Condition 1. $|\gamma_\theta(B,S)-\gamma_\theta(B-1,S-1)|<\rho/3$ for all $(B,S)\in\mathbb{N}^2$ and $\theta\in\{b,s\}$.

Condition 1, which is relaxed in Section 4.2 and is assumed to hold in the rest of Section 3, requires the arrival rates not to change too much when there is a transaction, that is, when a buyer and a seller leave the market. The following result establishes that, under Condition 1, a strategy profile where there is trade in every match (called a *full-trade strategy profile*) is an ε-MPE (called a *full-trade ε-MPE*), and that trade delay disappears as $\varepsilon \to 0$.

Theorem 3.1. *Assume Condition 1 holds. Then, for any $\varepsilon>0$, a full-trade ε-MPE exists $\bar{\lambda}$ big enough. Furthermore, $\lim_{\varepsilon\to 0}\bar{T}^\varepsilon(B,S)=0$ for all $(B,S)\in\mathbb{N}^2$.*

Theorem 3.1 may seem natural in our environment, where all sellers and buyers are homogeneous and do not have private information, and given our focus on symmetric equilibria. Indeed, Condition 1 holds trivially in Rubinstein and Wolinsky (1985) and Gale (1987), since the equilibrium arrival rate of traders (which, in their models, is a discrete-time flow) is independent of whether a given trader trades or not. Intuitively, delaying trade does not change the continuation value of the traders, but delays the realization of the gains from trade. This argument cannot be applied, in general, to a thin market: as each transaction affects the aggregate state of the market, traders may have the incentive to let other traders trade, and trade when his/her bargaining power is higher.

The proof of Theorem 3.1 is divided into two steps. The first consists on verifying that, under Condition 1, there is a full-trade strategy profile which is an ε-MPE for $\bar{\lambda}$ high enough. Notice that, for a full-trade strategy profile to be an equilibrium, it must be that each trader's offer is ε-close to the continuation value of his/her trading counter-party from rejecting it. This can be seen using the standard hold-up logic: if an equilibrium offer gave a positive rent to the trading counter-party, another offer giving a fraction of this rent would be accepted for sure, giving a profitable deviation. Thus, the proof focusses on verifying that offering the continuation value (from rejection) of the trading counter-party is indeed incentive compatible. Checking such an incentive is equivalent to verifying that the sum of the continuation values of a buyer and a seller is (weakly) below the surplus from trade, which is equal to 1. This is where Condition 1 is used, as it ensures that there cannot be much gain from making unacceptable offers: only arrivals of traders can significantly change the future arrival rates, and so the delay cost of waiting for the arrival of a trader does not compensate the potential gain. The second step of the proof consists on verifying that, as $\bar{\lambda} \to \infty$, the delay in any equilibrium shrinks to 0.

Section 4 shows that when Condition 1 does not hold there may be equilibria where there is a significant trade delay even when the bargaining frictions are low, that is, when $\bar{\lambda}$ is big. For example, in some specifications of the model, sellers may benefit from other sellers' trades (because, after a transaction takes place, there may be a higher arrival of buyers) and buyers may not be willing to trade at low prices (because, if there is no trade, the expected state of the market conditional on a trader arriving is attractive to them). Condition 1 (and relaxed versions of it, see Section 4) prevents such a possibility ensuring that a single trade (which lowers the number of buyers and the number of seller by one unit) does not change much the arrival process.

3.2 Frictionless-bargaining limit

We now consider the limit where the bargaining frictions disappear, that is, where buyers and sellers in the market meet at a very high rate (so $\bar{\lambda} \to \infty$). In this limit, we will be able to fully characterize how the market price depends on the arrival process and the bargaining protocol.

The first result of its section characterizes the limit of the continuation payoff functions (see Appendix A.1 for their expressions before the limit):

Lemma 3.1 (Bertrand competition)**.** *Two functions $V_b, V_s : \mathbb{N}^2 \to [0,1]$ exist satisfying that, for all $\delta > 0$, there are ε_δ and $\bar{\lambda}_\delta$ such that for all $\bar{\lambda} > \bar{\lambda}_\delta$ and ε-MPE with $\varepsilon \in (0, \varepsilon_\delta)$,*

$$\left\| \left(V_b^\varepsilon(B,S), V_s^\varepsilon(B,S)\right) - \left(V_b(B,S), V_s(B,S)\right) \right\| < \delta$$

for each state (B,S), where $(V_b^\varepsilon(B,S), V_s^\varepsilon(B,S))$ are the ε-MPE's continuation payoffs in state (B,S). Furthermore, they satisfy the following equations

$$\left(V_b(B,S), V_s(B,S)\right) = \begin{cases} \left(V_b(B-S,0), 1-V_b(B-S,0)\right) & \text{if } B>S, \\ \left(1-V_s(1,1), V_s(1,1)\right) & \text{if } B=S, \\ \left(1-V_s(0,S-B), V_s(0,S-B)\right) & \text{if } B<S. \end{cases}$$

Finally, the equilibrium prices at which transactions happen in state $(B,S) \in \mathbb{N}^2$ converge to $V_s(B,S)$.

We will refer to the *frictionless-bargaining limit* when talking about the limit equilibrium outcome of our game when $\bar{\lambda}$ is large and ε is small (see Lemma 3.1). By Theorem 3.1 we know that, under Condition 1 (notice that it does not involve $\bar{\lambda}$), such a limit outcome features no trade delay. Hence, Lemma 3.1 establishes that, in the frictionless-bargaining limit, the trade outcome at each imbalanced state of the market converges to that of a (static) competitive market: the agents on the long side of the market do not obtain any rents (on top of their value from not trading) or, analogously, the agents on the short side of the market obtain all surplus from trade net of their counter-parties (endogenous) reservation values. If, instead, the market is balanced, both types of traders are indifferent between trading or letting the other traders trade and obtaining the continuation value when the state is $(1,1)$. In any case, the sum of the continuation values of a seller and a buyer is 1: the agents on the short side of the market trade immediately, while the agents on the long side of the market obtain their continuation value from not trading independently of whether they trade or not.

Our model micro-founds, using a decentralized approach, the assumption in Taylor (1995) that, at any given time where the market is imbalanced, the transaction price is equal to the one corresponding to a static market with Bertrand-competition in the long side. Intuitively, in our model, an agent on the long side of the market (for example a buyer when $B > S$) has the option of not trading, letting the other side of the market to quickly clear and obtain his/her continuation value (equal to $V_b(B-S,0)$). Thus, when for example there is excess demand, sellers cannot sell their good at a price higher than $1 - V_b(B-S,0)$, and since under Condition 1 selling the good is preferable to keeping it for some time, the equilibrium price (and the sellers' payoff) is $1 - V_b(B-S,0)$.

Market price

Lemma 3.1 establishes that, in the frictionless-bargaining limit, if a transaction happens in a given state $(B,S) \in \mathbb{N}^2$, the transaction price (which is equal to $V_s(B,S)$) is only a function of the *net state* $N \equiv S - B \in \mathbb{Z}$. Hence, it is useful to define the *market price* $p(N)$ as the transaction price when the net state is N. It can be obtained as the unique solution of the following system of equations:

1. If $N > 0$, that is, if there are more sellers than buyers in the market, sellers are indifferent on trading now (and obtaining a payoff equal to $p(N)$), or refusing to trade and waiting for the state to change (and trade then). Thus, if a seller refuses

to trade, the buyers in the market trade with other sellers until the resulting state is $(0, N)$. As a result, the market price at the net state $N > 0$ satisfies the following equation

$$p(N) = \frac{\gamma_b(0,N)}{\gamma(0,N)+\rho}\, p(N-1) + \frac{\gamma_s(0,N)}{\gamma(0,N)+\rho}\, p(N+1)\ . \tag{3.1}$$

2. If $N < 0$, the situation is reversed: buyers are indifferent on trading now (obtaining a payoff equal to $1 - p(N)$) or letting other buyers trade and waiting for the state to change. Also, since by Theorem 3.1 the trade delay shrinks to 0 in the frictionless-bargaining limit, we have $V_b(B,S) + V_s(B,S) = 1$ for all $(B,S) \in \mathbb{N}^2$. Thus, rearranging some terms in the expression analogous to equation (3.1) for the buyers, we obtain that the price at net state N satisfies

$$p(N) = \frac{\rho}{\gamma(-N,0)+\rho} + \frac{\gamma_b(-N,0)}{\gamma(-N,0)+\rho}\, p(N-1) + \frac{\gamma_s(-N,0)}{\gamma(-N,0)+\rho}\, p(N+1)\ . \tag{3.2}$$

3. Finally, if $N = 0$, that is, when the market is balanced, the price is determined by computing the outcome of the frictionless-bargaining limit of a two-player bargaining game *á la* Rubinstein (1982) with randomly arrival outside options (given by the potential arrival of other traders). The resulting price is equal to

$$p(0) = \frac{\rho}{\gamma(1,1)+\rho}\, \xi(1,1) + \frac{\gamma_b(1,1)}{\gamma(1,1)+\rho}\, p(-1) + \frac{\gamma_s(1,1)}{\gamma(1,1)+\rho}\, p(1)\ . \tag{3.3}$$

   In the previous expression, the last two terms on the right hand side can be interpreted as the "outside option" for the seller, that is, her payoff from delaying trade by making unacceptable offers and rejecting the received offers. A similar expression can be obtained for the buyer. Thus, the "size of the pie" over which they bargain is not 1, but the trade surplus net of the sum of the outside options, which is $\frac{\rho}{\gamma(1,1)+\rho}$. Therefore, as in the standard Rubinstein bargaining game, the seller obtains, on top of her outside option, a fraction of the size of the pie equal to the probability with which she makes offers, $\xi(1,1)$.

Equations (3.1)-(3.3) fully characterize the market price in the frictionless-bargaining limit. Notice that $p(\cdot)$ is independent of the probability with which sellers make offers, ξ, except for its value when exactly one seller and one buyer are in the market, $\xi(1,1)$. The reason is that when the market is imbalanced, there is Bertrand competition between the agents on the long side of the market, independently of the probability with which they make offers.

When, instead, the market is balanced, each of the last two agents can delay the next trade so, in this case, the price coincides with that of the frictionless limit of a Rubinstein bargaining game.

The previous observations allow us to give the following characterization of the equilibrium price:

Proposition 3.1. *In the frictionless-bargaining limit, there exists some $x \in \mathbb{R}$ such that the market price at time t is given by*

$$p(S_t - B_t) = \mathbb{E}\left[\int_0^\infty e^{-\rho\,(s-t)} \left(\mathbf{1}_{B_s > S_s} + x\, \mathbf{1}_{B_s = S_s} \right) \rho \, \mathrm{d}s \,\middle|\, (B_t, S_t) \right], \tag{3.4}$$

where, for a given statement A, $\mathbf{1}_A = 1$ if A is true and $\mathbf{1}_A = 0$ if A is false. Furthermore, $p(\cdot)$ is a strictly decreasing function.

To see why Proposition 3.1 holds, consider a fictitious agent with the following flow payoff. Assume she receives a flow payoff equal to 0 when the market has excess supply (i.e., $N_t > 0$), a flow payoff equal to 1 when the market has excess demand (i.e., $N_t < 0$) and, when the market is balanced, a flow payoff $x \in \mathbb{R}$ to be determined. Visual inspection of equations (3.1) and (3.2) shows that, when $N_t \neq 0$, the market price $p(N_t)$ follows the same equation as the continuation value of the fictitious agent. When the market is balanced, instead, notice that the price depends on the arrival rates *if* the last two traders do not agree on trading, $\gamma_\theta(1,1)$ for $\theta = L, H$ (and the probability which each of them makes offers and counter offers in state $(1,1)$). Still, since they immediately agree in equilibrium, the rate at which the net state changes when $N_t = 0$ depends on the arrival rates when the market is empty, $\gamma_\theta(0,0)$ for $\theta = L, H$. Thus, in this case, the flow payoff that makes the continuation value of the fictitious agent equal to the price is, in general, different from $\xi(1,1)$. Under this interpretation, it is clear that the right hand side of equation (3.4) corresponds to the continuation value of the fictitious agent at time 0, which equal to $p(S_0 - B_0)$ to when x is chosen appropriately (the proof of Proposition 3.1 obtains the value of x that makes equation (3.4) hold).

Proposition 3.1 implies that only the evolution of the sign of the net amount of sellers (or buyers) in the market, which we call *balancedness of the market*, is relevant for determining the market price. This is because the intensity of the competition between agents on the long side of the market is irrelevant for determining the price when the market is unbalanced: the price equals their reservation value independently of their number. Thus, the price is not directly affected by the expected amount of future transactions, but by the expected evolution of the balancedness of the market. The net state influences the price because it affects the expected time until the market is balanced.

An implication of Proposition 3.1 is that, when the market is unbalanced, the market price always drifts towards the price of a balanced market.

Corollary 3.1. *For all t, using p_t to denote $p(N_t)$, we have*

$$\lim_{\Delta \searrow 0} \frac{\mathbb{E}_t[p_{t+\Delta} - p_t]}{\Delta} = \begin{cases} \rho\, p_t & \text{if } N_t > 0, \text{ that is, if } p_t < p(0), \\ -\rho\,(1 - p_t) & \text{if } N_t < 0, \text{ that is, if } p_t > p(0). \end{cases}$$

Thus, in expectation, the price changes toward the price of a balanced market.

The rationale for Corollary 3.1 is the following. By Proposition 3.1, the price at time t is below $p(0)$ only if there is an excess supply in the market, that is, if $N_t > 0$. Equation (3.1) implies that, in this case, each seller is indifferent between trading now (and obtaining $p(N_t)$) or waiting for the next arrival, where the price will become $p(N_t - 1)$ with probability $\frac{\gamma_b(0,N_t)}{\gamma(0,N_t)}$ and $p(N_t + 1)$ with the complementary probability. Given that a seller is, at time t, indifferent between trading now or not trading and waiting for the next arrival (which generates a delay cost), the price must increase in expectation.

Remark 3.1 (No Diamond's paradox)*.* Equations (3.1)-(3.3) show that, in the frictionless-bargaining limit, the payoff of each trader in each state is strictly positive as long as there is a positive probability that his or her side of the market becomes the short side of the market in the future. This may be surprising since, in bargaining models with one-sided offers (which in our model would correspond to $\xi \equiv 0$ or $\xi \equiv 1$), the side of the market making the offers obtains all surplus from trade, independently of the degree of balancedness of the market (usually known as the Diamond's paradox). In our model, the order of limits matters: our claim implicitly takes the frictionless-bargaining limit first, and the limit of one-sided offers afterwards. This result would not hold if we first assumed that $\xi(\cdot,\cdot)$ is constant and equal to either 0 or 1, and then we would take the frictionless-bargaining limit: in this case, the type of traders making all offers would obtain all gains from trade. Still, as Section 5.2 shows, our results are recovered when only sellers post prices, *but* buyers observe all offered prices at the moment they make their purchasing decision.

Example 3.1 (No arrival of sellers)*.* To illustrate Proposition 3.1, consider the case where no sellers arrive and where, for each state of the market, the arrival rate of buyers is positive. This may correspond, for example, to re-sale markets with a limited supply, such as limited editions (books, mobile phones, cars...) or housing units in housing projects in exclusive areas. Assume also, for simplicity, that initially there is no buyer in the market, $B_0 = 0$, and the initial number of sellers is $S_0 > 0$.

When no sellers arrive, equation (3.4) can be written in terms of the (stochastic) time it takes for the market to clear. Indeed, let τ be the first (i.e., infimum) time where $S_\tau = 0$. By Proposition 3.1, this corresponds to the first time when there is only one seller in the market and a buyer arrives. In this case, we have that when there is only one buyer and one seller in the market, the payoff of the seller (arising from the frictionless limit of the Rubinstein bargaining), is given by

$$V_s(1,1) = \underbrace{\frac{\xi(1,1)\,\rho}{\gamma_b(1,1)+\rho}}_{\text{Rubinstein payoff}} + \underbrace{\frac{\gamma_b(1,1)}{\gamma_b(1,1)+\rho}}_{\text{value from waiting}} \quad .$$

As a result, we can write the payoff of a seller as $V_s(0, S_0) = \frac{\gamma_b(1,1)+\xi(1,1)\,\rho}{\gamma_b(1,1)+\rho}\,\mathbb{E}[e^{-\rho\,\tau}]$. The transaction price at time t if there are $S_t > 1$ sellers in the market and a buyer arrives equals $V_s(0, S_t - 1)$, so prices increase over time, highlighting that the outside option that sellers have to wait until they are a monopolist becomes increasingly attractive.

3.3 Comparative statics

Changes on the discount rate

In the frictionless-bargaining limit of our model the net state N_t evolves following a Markov chain. We now assume that such a Markov chain has an ergodic (or stationary) distribution. This is guaranteed if, for example, when the market becomes heavily unbalanced, the arrival rate of agents on the short side of the market is bounded away from the arrival rate of agents on the long side of the market. Formally, a sufficient condition for N to have an ergodic distribution is the existence of some $\underline{\gamma} > 0$ such that

$$\liminf_{N\to\infty} \big(\gamma_b(0,N) - \gamma_s(0,N)\big) \geq \underline{\gamma} \ \text{ and } \ \liminf_{N\to-\infty} \big(\gamma_s(-N,0) - \gamma_b(-N,0)\big) \geq \underline{\gamma}\ .$$

This is plausible: the payoff of a buyer from entering the market is high when the market price is low due to a large excess supply (i.e., when N_t is high), while the payoff of the sellers is low in this case. As a result, the low price when the market features a large excess supply may attract more buyers than sellers into the market.

Corollary 3.2. *Assume that there is an ergodic distribution for the net state N. Then, if $\gamma_b(0,0) = \gamma_b(1,1)$ and $\gamma_s(0,0) = \gamma_s(1,1)$, an increase in ρ generates a mean-preserving spread of the ergodic distribution of market prices.*

The fact that the average long run price (i.e., the expected price under the ergodic distribution) is independent of the discount rate follows from equation (3.4). Indeed, in the long run, the expected flow payoff of the "fictitious agent" introduced after Proposition 3.1 is equal to the ergodic probability that the market has excess demand plus an adjustment term x multiplying the ergodic probability of the market being balanced. Furthermore, for any initial state $(B_0, S_0) \in \mathbb{Z}_+^2$ and time $t \geq 0$, the law of iterated expectations allows us to rewrite equation (3.4) as

$$\mathbb{E}[p_t] = \mathbb{E}\left[\int_t^\infty e^{-\rho\,(t'-t)} \, (\mathbb{I}_{B_{t'} > S_{t'}} + x\, \mathbb{I}_{B_{t'} = S_{t'}})\, \rho \, \mathrm{d}t' \right] .$$

It is then clear from the previous expression that the expected price under the ergodic distribution is equal to the long-run average payoff of the fictitious agent. Hence, the only dependence of the long run expected price on ρ comes through the value of x, weighted by the ergodic probability that the market is balanced. As we argued after Proposition 3.1, when $\gamma_\theta(0,0) = \gamma_\theta(1,1)$ for both $\theta = b, s$, we have that $x = \xi(1,1)$, and as a result the long run expected price is independent of ρ.[5]

Corollary 3.2 illustrates an important effect that an increase on the discount rate (or the interest rate) has on the trade outcome: it increases the ergodic price dispersion. This can be seen as a consequence of Corollary 3.1, which establishes that an increase in ρ increases the drift of the market price towards the price of a balanced market. This is necessary to make traders on the long side of the market willing to delay trade instead of trading now. Intuitively, for a fixed $N > 0$, an increase in the discount rate ρ lowers the discount factor of the time it takes the market to become balanced. Thus, in the limit where $\rho \to \infty$, we have $p(N) \to 0$ for all $N > 0$, and $p(N) \to 1$ for all $N < 0$.

Patient traders

One of the salient questions in the literature on decentralized bargaining in large markets is whether lowering the frictions in the market leads to a competitive outcome. This

[5] In general, even when $\gamma_\theta(0,0) \neq \gamma_\theta(1,1)$ for some $\theta \in \{b, s\}$, the effect of an increase in ρ is likely to generate price dispersion (but not necessarily preserve the mean), specially when the market is unlikely to be balanced in the ergodic distribution.

exercise allows analyzing whether and how frictions may be magnified or mitigated by the equilibrium behavior of the agents in the market, and shed light as a result on how robust the predictions of the models with markets without frictions are.[6] This section asks a similar question for a thin market, and does that by lowering the friction that remains in the frictionless-bargaining limit: the delay cost that an agent incurs when he or she is on the long side of the market.

In a thin market there is no natural analogous of a "competitive outcome", since the number of traders on each side of the market is, at each given moment in time, finite. Still, when traders are patient, it is cheap for them wait to trade (and compete) with future traders, enlarging the effective market that each trader faces. As we will see, this implies that, when traders are patient, the outcome of a thin market shares many features with that of a competitive market: the price is (approximately) constant, and depends only the (expected) of balancedness of the market. Nevertheless, the endogenous arrival process in a thin market implies that, in general, no side of the market obtains the full surplus from trade.

Assume first that the net state of the market has an ergodic distribution F, that is, for any value of $N \in \mathbb{Z}$ we have that $\lim_{t\to\infty} \Pr(N_t = N) = F(\{N\})$ independently of N_0. As traders become more patient, the current state of the market becomes progressively less relevant to determine the price, since each trader in the market can wait for the state of the market to change without incurring a big cost. In particular, the delay cost from not trading and waiting until the net state reaches some given state in the support of F tends to 0 as $\rho \to 0$. As a result, as the discount rate ρ shrinks, it may seem that, for each given state, the payoff from not trading becomes increasingly attractive for all traders in the market. Nevertheless, this is not possible in the frictionless-bargaining limit: the sum of the continuation values of a seller and a buyer in the market is always equal to 1, independently of their discount rate. Even though waiting is increasingly cheap, it also becomes increasingly invaluable, since also the price variation across states becomes

[6] For example, Gale (1987) characterizes the trade outcome in the large-market version of our model in the limit where the discount rate converges to 0, and obtains that it converges to that of a competitive market. In this limit, the price is either 0 (if there are more buyers than sellers) or equal to 1 (if there are more buyers than sellers). Other papers have identified some reasons for the failure of convergence. It may be caused, for example, by asymmetric information between traders (Satterthwaite and Shneyerov, 2007; Lauermann and Wolinsky, 2016), the heterogeneity in each side of the market (Lauermann, 2012), or lack of knowledge about the state of the market (Lauermann, Merzyn, and Virág, 2017). See also Lauermann (2013) for an analysis of other causes of delay.

increasingly small. The following result characterizes how the waiting options of buyers and sellers affect the market's outcome when the discount rate decreases.[7]

Proposition 3.2. *Assume that net state has an ergodic distribution F. Then, as $\rho \searrow 0$, the ergodic distribution of transaction prices converges to a degenerated distribution.*

Proposition 3.2 establishes that, as traders become more patient, the distribution of transaction prices converges to a distribution degenerated at some "competitive" price p^*. This is intuitive: since waiting for the state to change (instead of trading now) is increasingly cheap as $\rho \to 0$, the market price in all states of the market converges to a single price. Such a price can be obtained using equation (3.4), from which it is clear that when ρ is small the market price is close to the ergodic probability of the market having an excess demand (adjusted with term multiplying the probability of the market being balanced). It is then immediate to see that the distribution of transaction prices also becomes degenerated, as $\rho \to 0$, to the competitive price p^*.

To provide further intuition on the previous results, let $\bar{B}_t$ and $\bar{S}_t$ denote, respectively, the number of buyers and sellers into the market from 0 to t, including the ones "arrived" (or present) at time 0. Then, trivially, $N_t = S_t - B_t = \bar{S}_t - \bar{B}_t$, so equation (3.4) is valid replacing B_t and S_t by $\bar{B}_t$ and $\bar{S}_t$, respectively. As traders become more patient, the price (at time 0) approximates the (ergodic) probability that more sellers than buyers arrive in the future. Hence, as ρ decreases, the effective market that a trader (at time 0) faces grows intertemporally. The endogenity of the price process implies that, in general (and differently from the large market case), the price is not degenerated at 0 (when there is excess supply) or 1 (when there is excess demand). Instead, in a thin market, the competitive price is a convex combination of the two extremes, each of them weighted according to the probability that the market features excess supply and demand.

Remark 3.2. In our model, making traders more patient (through decreasing ρ by a factor $k < 1$) can be reinterpreted as enlarging the market by replication, that is, increasing the arrival rates (of buyers and sellers) by a factor $1/k > 1$. Indeed, increasing the arrival rates accelerates the pace of our model: the distribution of (B_t, S_t) under arrival rates (γ_b, γ_s) is the same as the distribution of $(B_{k\,t}, S_{k\,t})$ under arrival rates $(\gamma_b/k, \gamma_s/k)$. Using equation (3.4), it is not difficult to see that the value of $p(S_0, B_0)$ is the same under $(\rho, \gamma_b, \gamma_s)$ and

[7] Notice that, letting ρ decrease (while fixing rest of the primitives of the model) makes Condition 1 hold only if γ_b and γ_s are functions of the net state. Still, it is easy to see that this is not necessary, in general, to ensure that there is no trade delay in the frictionless bargaining limit (see Condition 2 below).

under $(k\,\rho, \gamma_b/k, \gamma_s/k)$. So, replicating the market $1/k$ times is equivalent (in terms of the ergodic distribution of prices) to making it $1/k$ times faster or, equivalently, to make its traders' impatience to decrease by a factor k.

Our interpretation of a market replication may correspond, in practice, to the unification of similar markets into bigger ones. This may result, for example, from the introduction of websites providing information on local rental or housing prices in close locations, or used durable goods. The introduction of such webpages may make it easier for buyers to compare prices across markets, which may *de facto* transform them into a single market. Our result implies that even though the unification of markets may not change the ergodic distribution of the market composition much, it makes prices fluctuate faster (in the sense that changes in the market price happen more frequently), and makes their ergodic distribution gets more concentrated around a given value.

4 Conditions for no trade delay

In this section we relax Condition 1. We first show an example of a setting where Condition 1 does not hold and exhibits trade delay even when in the frictionless bargaining limit. Next, we provide weaker conditions that ensure that, as the bargaining frictions disappear, trade becomes efficient.

4.1 An example with delay

In this section we illustrate how equilibria with delay may arise whenever Condition 1 does not hold. In order to keep the example simple, we focus on a given state of the market, and we exogenously fix the continuation payoffs when such a state changes without explicitly modeling the continuation play. Considering this "reduced version" of our model simplifies the computations, and it is easy to verify that there exist full specifications of our model with the same equilibrium features.

Consider the following reduced version of our model. Initially, there is one buyer and two sellers in the market. We assume that $\gamma_s(1,2) > \gamma_b(1,2) = 0$, and denote $\gamma \equiv \gamma_s(1,2)$ and $\lambda(1,2) = \bar{\lambda}$. If there is one trade before the arrival of a seller, the market becomes visible to other buyers, so the remaining seller obtains a high continuation payoff, which for simplicity is assumed to be equal to 1.[8] If, instead, a seller arrives, the strong competition between sellers gives the buyer a high continuation payoff, which is again assumed to be 1 (see footnote 8), and the sellers obtain 0.

We first compute the continuation values of the initial buyer and sellers in the market under a full-trade strategy profile. They solve the following system of equations:

$$V_b(1,2) = \frac{\bar\lambda}{\bar\lambda+\gamma+\rho}\left(\xi\, V_b(1,2)+(1-\xi)\,(1-V_s(1,2))\right) + \frac{\gamma}{\bar\lambda+\gamma+\rho}\,1\ ,$$
$$V_s(1,2) = \frac{\bar\lambda/2}{\bar\lambda+\gamma+\rho}\left(\xi\,(1-V_b(1,2)) + (1-\xi)\,V_s(1,2)\right) + \frac{\bar\lambda/2}{\bar\lambda+\gamma+\rho}\,1\ .$$

Solving the previous system of equations, and using simple algebra, it is easy to show that

$$\lim_{\rho\to 0}\left(V_b(1,2)+V_s(1,2)\right) = 1 + \frac{\gamma\,\bar\lambda}{(\gamma+\bar\lambda)\,(2\,\gamma+\xi\,\bar\lambda)} > 1\ .$$

Thus, if ρ is low, a full-trade strategy is not an ε-MPE for ε small enough, and in this case any ε-MPE of this reduced version of our model involves randomization in the acceptance of offers. Using β to denote the probability of agreement in a meeting in state $(1,2)$, in any ε-MPE of the (reduced) game, we have

$$\beta = \min\left\{1, \frac{2\,\rho\,(\gamma+\rho)}{\gamma\,\bar\lambda}\right\} + O(\varepsilon)\ .$$

Notice that the rate at which an agreement occurs in state $(1,2)$ (which equals $\beta\,\bar\lambda$) converges to $\frac{2\,\rho\,(\gamma+\rho)}{\gamma} + O(\varepsilon)$ as $\bar\lambda$ becomes big, that is, a significant trade delay remains even in the limit where bargaining frictions disappear.

Our example shows that, in some specifications, agents on one side of the market benefit from other traders' transaction, while agents on the other side of the market benefit from the arrival of new agents. In our example, sellers obtain a high continuation payoff if a transaction occurs, and the buyer gets a high payoff if a trader arrives. The buyer is then not willing to accept a high price, given that he has the option of waiting for the arrival of another seller and then obtain a high payoff. As a result, immediate agreement is not possible: otherwise, each seller would have the incentive to let the other seller trade at

[8] All numbers can be perturbed while keeping the features of the equilibrium the same. A continuation payoff for the seller arbitrarily close to 1 when the state is $(0,1)$ can be supported assuming that $\gamma_b(0,1) \gg \gamma_s(0,1)$, which requires Condition 1 to fail, that $\gamma_b(1,1) \gg \gamma_s(1,1)$ and that the arrival sellers is very low afterwards. Analogously, a high continuation value for the buyer in state $(1,3)$ can be supported if, for example, no more buyers arrive afterwards.

a low price, and obtain a high continuation payoff afterwards. The equilibrium behavior of the sellers in the market resembles then a war of attrition: each of them trades at the rate that makes the other seller indifferent between trading or not. Such delay lowers the value of making unacceptable offers from the seller's perspective, since doing so comes at the risk of another seller arriving. As time passes, either one of the sellers trade (and the remaining seller obtains a high payoff), or another seller arrives (and both of them obtain a low continuation payoff).

Remark 4.1. Inefficient delay can also be found in other bargaining models with complete information. For example, Cai (2000) analyzes a model with one-to-many bargaining between farmers and a railroad company, where the gains from trade are realized only if all farmers agree. Similar to us, farmers want other farmers to trade, to gain monopsony power. Also, in models of bargaining in networks such as Elliott and Nava (2016), delay may happen because traders are heterogeneous. Our example illustrates that trade delay may appear even when bargaining is decentralized and traders are homogeneous, the reason being that some traders may benefit from other traders' trades, while others benefit from arrivals.

4.2 Conditions for no delay

Condition 1 is attractive because of its simplicity and easiness to interpret, as it establishes a bound on the variation of arrival rates that a single trade can generate. Theorem 3.1 establishes that such a condition is enough to ensure that, in the frictionless-bargaining limit, the maximum trade delay in any equilibrium shrinks to 0.

This section provides two alternative conditions that relax Condition 1 while ensuring that, in the frictionless-bargaining limit, there is no trade delay. The first is necessary and sufficient, but involves solving the price of a full-trade strategy profile in the frictionless-bargaining limit:

Condition 2. Let p be the solution of the system of equations (3.1)-(3.3). Then,

$$p(N) > \tfrac{\gamma_b(1-N,1)}{\gamma(1-N,1)+\rho}\, p(N-1) + \tfrac{\gamma_s(1-N,1)}{\gamma(1-N,1)+\rho}\, p(N+1) \qquad \text{if } N<0, \tag{4.1}$$

$$1-p(N) > \tfrac{\gamma_b(1,N+1)}{\gamma(1,N+1)+\rho}\,(1-p(N-1)) + \tfrac{\gamma_s(1,N+1)}{\gamma(1,N+1)+\rho}\,(1-p(N+1)) \qquad \text{if } N>0. \tag{4.2}$$

Equations (4.1) and (4.2) ensure that a full-trade strategy profile is an equilibrium if $\bar{\lambda}$ is high enough. Indeed, in the frictionless-bargaining limit, equation (4.1) imposes that, when sellers are on the short side of the market, they do not have the incentive to make unacceptable offers. Its left hand side is the equilibrium payoff they obtain in state (B,S) with $N=S-B$, while the right hand side is the value of making unacceptable offers and waiting for the state to change. Thus, if the left hand side is higher than the right hand side, deviating from a full-trade strategy profile is not optimal. Equation (4.2) imposes the same condition for the buyers (recall that $V_b(B,S)=1-p(S-B)$).

The following result establishes that Condition 2 is a necessary and sufficient condition for the absence of delay in equilibrium behavior. In it, for each $\varepsilon>0$, we use $\underline{T}^{\varepsilon}(B,S)\in\overline{\mathbb{R}}_+$ to denote the (inferior) limit of minimum trade delay at state $(B,S)\in\mathbb{N}^2$ among all ε-MPE as $\bar{\lambda}\to\infty$.

Theorem 4.1. *Theorem 3.1 holds under Condition 2. Conversely, if one of the conditions in Condition 2 fails with strict inequality for some $N\neq 0$, there is some state $(B,S)\in\mathbb{N}^2$ such that* $\lim_{\varepsilon\to 0}\underline{T}^{\varepsilon}(B,S)>0$.

Condition 2 is (generically) necessary and sufficient for trade delay to be part of equilibrium behavior. Nevertheless, such a condition may be difficult to verify in some situations, as it requires solving a (infinite) system of equations (to find $p(N)$ for each $N\in\mathbb{Z}$). The following is condition can, instead, be verified state by state and it is easy to prove (Lemma A.1 in the Appendix) that relaxes Condition 1:

Condition 3. For all $B,S>1$, we have

$$\frac{\gamma_b(B,1)}{\gamma(B,1)+\rho}-\frac{\gamma_b(B-1,0)}{\gamma(B-1,0)+\rho}<\frac{\rho}{\gamma(B-1,0)+\rho}\quad\text{and}\tag{4.3}$$

$$\frac{\gamma_s(1,S)}{\gamma(1,S)+\rho}-\frac{\gamma_s(0,S-1)}{\gamma(0,S-1)+\rho}<\frac{\rho}{\gamma(0,S-1)+\rho}\,.\tag{4.4}$$

Condition 3, in the same spirit that Condition 1 requires that, when there is one trader in one of the sides of the market, rejecting offers is not too attractive.

Corollary 4.1. *Theorem 3.1 holds under Condition 3.*

5 Generalizations and extensions

5.1 State of the economy

In our base model we assume that the arrival rates of traders depend only on the current number of buyers and sellers in the market. In practice, the arrival of traders in markets may depend on other factors, like the state of the economy (economic booms or downturns), changes in similar markets, idiosyncratic demand/supply shocks in the market, etc. In this section, we analyze the robustness of our results to enriching the arrival process to depend on a multi-dimensional state.

We now assume that the state of the market at time t is (B_t, S_t, ω_t), where ω_t is the value of a stochastic process taking values in some $\Omega \subset \mathbb{R}^n$ for some $n \in \mathbb{Z}_+$, which we call the *market's cycle*. Hence, with some abuse of notation, we use $\gamma_b \equiv \gamma_b(B_t, S_t, \omega_t)$ and $\gamma_s \equiv \gamma_s(B_t, S_t, \omega_t)$ to denote, respectively, the arrival rates of buyers and sellers of buyers into the market at time t, and $\gamma = \gamma_b + \gamma_s$ as before. Also, we assume that the market's cycle ω_t changes at a Poisson rate $\eta(B_t, S_t, \omega_t)$, where $\eta : \mathbb{Z}_+ \times \mathbb{Z}_+ \times \Omega \to \mathbb{R}_+$. If there is an arrival of a change in ω at time t, the new state is determined by a random variable $\tilde{\omega}$ which only depends on (B_t, S_t, ω_t). Some components of the market's cycle can be assumed to evolve exogenously (like, for example, the state of the economy) and some may depend on the endogenous variables of the market (such as the number of traders in the market or the regional economic variables if the market is geographically located).

We focus again on the frictionless-bargaining limit without trade delay. Using similar arguments as in Section 3 it is easy to see that, if a full-trade strategy is an ε-MPE for ε and $\bar{\lambda}^{-1}$ small enough, the frictionless-bargaining limit of the market price in a given state (B, S, ω) only depends on the net number of traders, N, and the market's cycle, ω, so we use $p(N, \omega)$ to denote it. Furthermore, the expected new market price conditionally on the state changing either through an arrival or a change in the market's cycle (but not a a trade), denoted $\tilde{p}(B, S, \omega)$, is given by

$$\tilde{p}(B,S,\omega) \equiv \frac{\eta(B,S,\omega)\ \mathbb{E}_{\tilde{\omega}}[p(B,S,\tilde{\omega})|(B,S,\omega)]}{\eta(B,S,\omega)+\gamma(B,S,\omega)} + \frac{\gamma_b(B,S,\omega)\, p(B{+}1,S,\omega)}{\eta(B,S,\omega){+}\gamma(B,S,\omega)} + \frac{\gamma_s(B,S,\omega)\, p(B,S{+}1,\omega)}{\eta(B,S,\omega){+}\gamma(B,S,\omega)}\ ,$$

where $\mathbb{E}_{\tilde{\omega}}[\,\cdot\,|(B,S,\omega)]$ denotes the expectation with respect to the random variable $\tilde{\omega}$. The first term on the right hand side corresponds to the change in the market price when the market's cycle changes. The other two terms express, as in our base model, the changes of the market price owed to the arrival of buyers and sellers into the market.

We can then rewrite the equations defining the market price, (3.1)-(3.3), which become

$$p(N,\omega) = \begin{cases} \frac{\eta(0,N,\omega)+\gamma(0,N,\omega)}{\eta(0,N,\omega)+\gamma(0,N,\omega)+\rho}\,\tilde{p}(0,N,\omega) & \text{if } N > 0, \\ \frac{\rho}{\eta(1,1,\omega)+\gamma(1,1,\omega)+\rho}\,\xi(1,1) + \frac{\eta(1,1,\omega)+\gamma(1,1,\omega)+\rho}{\eta(1,1,\omega)+\gamma(1,1,\omega)+\rho}\,\tilde{p}(1,1,\omega) & \text{if } N = 0, \\ \frac{\rho}{\eta(-N,0,\omega)+\gamma(-N,0,\omega)+\rho} + \frac{\eta(-N,0,\omega)+\gamma(-N,0,\omega)}{\eta(-N,0,\omega)+\gamma(-N,0,\omega)+\rho}\,\tilde{p}(-N,0,\omega) & \text{if } N < 0. \end{cases} \tag{5.1}$$

Finally, Condition 2 now can be rewritten as:

Condition 4. Let p be the solution of the system of equations (5.1). Then,

$$p(N,\omega) > \frac{\eta(-N+1,1,\omega)+\gamma(-N+1,1,\omega)}{\eta(-N+1,1,\omega)+\gamma(-N+1,1,\omega)+\rho}\,\tilde{p}(-N+1,1,\omega) \qquad \text{if } N<0, \tag{5.2}$$

$$1-p(N,\omega) > \frac{\eta(1,N+1,\omega)+\gamma(1,N+1,\omega)}{\eta(1,N+1,\omega)+\gamma(1,N+1,\omega)+\rho}\,(1-\tilde{p}(1,N+1,\omega)) \qquad \text{if } N>0. \tag{5.3}$$

We can now use Condition 4 to write a result analogous to Theorem 4.1:

Proposition 5.1. *Under the general arrival process, Theorem 4.1 holds replacing Condition 4 by Condition 2.*

The logic behind Proposition 3.1 and Corollary 3.1 remain the same: at any given moment in time, the traders on the long side of the market are indifferent on trading or not. Indeed, it is easy to see that Proposition 5.1 implies that equation (3.4) is valid, now for function $x : \Omega \to \mathbb{R}$ (instead of a constant), that is, the price at a given time remains the expected discounted time that the market has excess demand, adjusted by the times where the market is balanced. Similarly, Proposition 3.2 also applies, where the "competitive price" still can be written as the ergodic probability of the market exhibiting excess demand plus a factor multiplying the probability that the market is balanced.

5.2 Other price mechanisms

In this section we extend our results to different price mechanisms. We first model the outcome of each meeting between a buyer and a seller as a general Nash-bargaining outcome. This generalizes the take-it-or-leave-it offer of our base model to the outcome of a general bargaining protocol. The second extension studies a centralized mechanism where sellers post prices in a price aggregator, which is commonly used in some online markets. As we will see, both cases deliver results which are similar to those obtained in our base model.

Nash bargaining

Our bargaining protocol within a meeting between a buyer and a seller can be easily generalized to be the outcome of a general Nash bargaining problem. Notice that, in this case, in a given equilibrium and for each state of the market (B, S), the "size of the cake" over which a buyer and a seller bargain is $1-V_b(S,B)-V_s(S,B)$ (which is equal to the gains from trade minus the reservation value of the traders). Indeed, since both traders retain the option of not trading, waiting and obtaining their continuation value, the outcome of the bargain is a price $p(S,B)$ satisfying

$$p(S,B) = y(S,B)\,\big(\underbrace{1 - V_b(S,B) + V_s(S,B)}_{\text{size of the cake}}\big) + \underbrace{V_s(S,B)}_{s\text{'s res. pric.}}\ .$$

for some potentially stochastic $y(S,B)\in[0,1]$. When the matching rate increases, the size of the cake shrinks to 0 when the market is imbalanced: since re-matching happens fast, the sum of the continuation payoffs if there is no agreement is close to 1. Furthermore, as long as y is not degenerated at 0 or 1,[9] Lemma 3.1 holds, that is, the outcome of an imbalanced market is equivalent to that of where there is Bertrand-competition in one of the sides of the market (see Remark 3.1).

Price posting

In some thin markets prices are posted by sellers and collected by centralized entities (or market makers), so buyers can easily access the different offers in the market. Examples are online price aggregators, such as *AirBnB.com* or *SabbaticalHomes.com* for apartment rentals or *CraigsList.com* for, among others, durable goods. In this section we analyze a market with centralized price posting, and we verify that the results we obtained in our base model still hold.

We consider a model similar to our base model in Section 2 with the following differences. Now we assume that buyers in the market are imperfectly attentive. More precisely, we assume that they draw attention times at rate $\bar{\lambda} > 0$. When a buyer draws an attention time, each seller in the market posts a price, and the buyer decides whether to accept one of them or not.[10] The assumption that buyers face inattention frictions is technically

[9] Notice that our base model corresponds to the case where $y(S,B) = 1$ with probability $\xi(S,B)$ and $y(S,B) = 0$ with probability $1 - \xi(S,B)$.

[10] Only allowing sellers to post prices when buyers draw attention times is technically convenient as it

convenient since, as we mentioned before, there is no coordination motives between the buyers in the market to which offer to accept. Indeed, our assumption guarantees that, in a given instant, at most one buyer has the opportunity to accept an offer.

The following result characterizes the conditions under which the price-posting market does not exhibit trade delay, and the corresponding equilibrium market price.

Proposition 5.2. *Theorem 4.1 holds under the price posting, where the market price is computed using* $\xi(1,1) = 1$.

Proposition 5.2 establishes that all results in Section 3 for our base model also apply in a market with centralized price posting by the sellers. Now, when there is one seller in the market, she enjoys monopolistic power. As a result, when $\bar{\lambda} \to \infty$, equations (3.2) and (3.3) (with $\xi(1,1) = 1$, since, when there is one seller in the market, she makes all offers) hold: each seller extracts the trade surplus net of outside option of the buyer she trades. When, instead, there are more sellers than buyers in the market, they compete *á la* Bertrand so, as in our base model, each seller is indifferent between trading now, or waiting and trading when she is a monopolist (that is, until $N_t = 1$ and a buyer arrives). Thus, if τ indicates the first time where $N_t = 0$, we have that the following equation applies:

$$p(N) = \begin{cases} \mathbb{E}\left[\, e^{-\rho\tau} \,\middle|\, N_0 = N \right] p(0) & \text{if } N > 0, \\ 1 - \mathbb{E}\left[\, e^{-\rho\tau} \,\middle|\, N_0 = N \right] (1 - p(0)) & \text{if } N < 0. \end{cases} \tag{5.4}$$

Notice that equation (5.4) also applies in the frictionless-bargaining limit of our base model (independently of the value of $\xi(1,1)$).

Remark 5.1. The assumption that the market features price posting brings our model closer to Taylor (1995). Still, some crucial differences between the models remain. First, we assume that, consistently with the price mechanism in online markets, only the sellers set prices, instead of the current short side of the market (and each side with probability $\frac{1}{2}$ when the market is balanced). Second, we allow the arrival process to have a general form (which can be further generalized proceeding similar to Section 5.1), which allows us to analyze its influence on the price dynamics. Finally, we assume that buyers face

avoids technicalities required to handle strategies in continuous-time. In a Markov strategy, allowing sellers to post a price at every instant (without knowing whether a buyer draws an attention time) or only at instants where a buyer draws an attention time makes no difference, given that prices posted at a given instant are relevant only conditional on a buyer drawing an attention at that instant.

some attention frictions, which eliminates coordination considerations among buyers that, otherwise, have to be exogenously assumed.

Remark 5.2. The bargaining protocol of our model can be generalized to allowing matches of not only two agents, but more. In this case, for each state (B, S), the meeting rate λ can be generalized to be a function from $\{1, ..., B\} \times \{1, ..., S\}$ to $\mathbb{R}_{++}$, where $\lambda_{B,S}(B', S')$ now indicates the rate at which B' buyers and S' sellers meet. Analogously, $\xi_{B,S}$ can be generalized to a function from $\{1, ..., B\} \times \{1, ..., S\}$ to $[0, 1]$, where $\xi_{B,S}(B', S')$ indicates the probability that the sellers make (simultaneous) offers when B' buyers and S' sellers meet. In this case, for example, if only sellers could make offers ($\xi \equiv 1$), they would still Bertrand-compete when two or more of them would meet a buyer, ending up offering $V_s(B-1, S-1)$ when $S > B$. Thus, irrespective of the value of ξ, the market price would have the similar dynamics as in our model.

6 Conclusions

We have studied decentralized bargaining in a thin market. In such a market, each agent enjoys some market power, so the outcome of bilateral bargaining not only depends on the current market conditions, but also on the expectations about them. We have showed how the trade outcome is affected by the price mechanism and the arrival process.

Our results highlight that trade outcomes in thin markets differ from those of large markets in several important dimensions. Most saliently, we show that when buyers and sellers are homogeneous, trade delay may be part of equilibrium behavior. Delay occurs when individual trades influence the future market evolution.

In the absence of trade delay, the short side of the market has all bargaining power. Nevertheless, each agent has the possibility of waiting to trade in the future, which gives him/her an endogenous outside option. This implies that, in thin markets, the stochasticity of the market composition generates rich price dynamics, which have three main features. First, the competition between the traders on the long side of the market makes prices drift towards the price of a balanced market. Second, as increases on the interest rate make waiting more costly, they generate mean preserving spreads. Finally, in the large market limit (that is, when the agents become increasingly patient), the distribution of transaction prices degenerates towards a competitive price. Such a competitive is, in general, not degenerated: no side of the market obtains all surplus from trade, and the price is proportional to the ergodic probability of the market having more sellers than buyers.

Our model can be generalized in multiple directions. One that may be particularly interesting is allowing buyers and sellers to be heterogeneous, both in terms of the quality of their goods and their valuation for them. This would make the analysis much more involved, as it would enlarge the dimensionality of the state of the market. For example, Elliott and Nava (2016) show that, in a model of bargaining in networks, the outcome of the bargaining is stochastic even in the limit when bargaining frictions vanish, as sometimes transactions with low gains from trade are realized in the presence of more beneficial trade opportunities. The analysis of this and other extensions is left to future research.

A Omitted expressions and proofs of the results

A.1 Payoffs and ε-equilibria

We fix a strategy for the sellers, (π_s, α_s), and for the buyers, (π_b, α_b), where, for each type $\theta \in \{b, s\}$ and state $(B, S) \in \mathbb{N}^2$, $\pi_\theta(B, S) \in \Delta(\mathbb{R})$ is the distribution of price offers that type-θ traders make, while $\alpha_\theta(\cdot; B, S) : \mathbb{R} \to [0, 1]$ maps each price offer to a probability of acceptance. Given that we are going to consider models with different values of $\bar\lambda$, it is convenient to make $\bar\lambda$ explicit in the continuation payoffs of buyers and sellers, denoted $V_b^{\bar\lambda}$ and $V_s^{\bar\lambda}$, respectively. Hence, the continuation values that such a strategy profile generates are given by

$$V_b^{\bar\lambda}(B,S) = \frac{\frac{1}{B}\,\bar\lambda\,\ell\,\hat{V}_b^{\bar\lambda}(B,S)}{\bar\lambda\,\ell+\gamma+\rho} + \frac{\frac{B-1}{B}\,\bar\lambda\,\ell\,\check{V}_b^{\bar\lambda}(B-1,S-1)}{\bar\lambda\,\ell+\gamma+\rho} + \frac{\gamma\,\tilde{V}_b^{\bar\lambda}(B,S)}{\bar\lambda\,\ell+\gamma+\rho}\ , \tag{A.1}$$

$$V_s^{\bar\lambda}(B,S) = \frac{\frac{1}{S}\,\bar\lambda\,\ell\,\hat{V}_s^{\bar\lambda}(B,S)}{\bar\lambda\,\ell+\gamma+\rho} + \frac{\frac{S-1}{S}\,\bar\lambda\,\ell\,\check{V}_s^{\bar\lambda}(B-1,S-1)}{\bar\lambda\,\ell+\gamma+\rho} + \frac{\gamma\,\tilde{V}_s^{\bar\lambda}(B,S)}{\bar\lambda\,\ell+\gamma+\rho}\ ; \tag{A.2}$$

where, to keep notation short, γ, ℓ and ξ should be interpreted as evaluated at (B, S), where

$$\tilde{V}_\theta^{\bar\lambda}(B,S) \equiv \frac{\gamma_b(B,S)}{\gamma(B,S)}\,V_\theta^{\bar\lambda}(B+1,S) + \frac{\gamma_s(B,S)}{\gamma(B,S)}\,V_\theta^{\bar\lambda}(B,S+1) \tag{A.3}$$

is the expected continuation payoff of a trader of type θ conditional on a trader arriving in the market, where

$$\begin{aligned}\hat{V}_b^{\bar\lambda}(B,S) \equiv{}& \xi\,\mathbb{E}_{\tilde p}\big[\,\alpha_b(\tilde p;B,S)\,(1-\tilde p) + (1-\alpha_b(\tilde p;B,S))\,V_b^{\bar\lambda}(B,S)\,\big|\,\pi_s(B,S)\,\big] \\ &+ (1-\xi)\,\mathbb{E}_{\tilde p}\big[\,\alpha_s(\tilde p;B,S)\,(1-\tilde p) + (1-\alpha_s(\tilde p;B,S))\,V_b^{\bar\lambda}(B,S)\,\big|\,\pi_b(B,S)\,\big]\ ,\end{aligned} \tag{A.4}$$

$$\begin{aligned}\hat{V}_s^{\bar\lambda}(B,S) \equiv{}& \xi\,\mathbb{E}_{\tilde p}\big[\,\alpha_b(\tilde p;B,S)\,\tilde p + (1-\alpha_b(\tilde p;B,S))\,V_s^{\bar\lambda}(B,S)\,\big|\,\pi_s(B,S)\,\big] \\ &+ (1-\xi)\,\mathbb{E}_{\tilde p}\big[\,\alpha_s(\tilde p;B,S)\,\tilde p + (1-\alpha_s(\tilde p;B,S))\,V_s^{\bar\lambda}(B,S)\,\big|\,\pi_b(B,S)\,\big]\end{aligned} \tag{A.5}$$

are the expected continuation values conditional on being selected in the match, and where

$$\begin{aligned}\check{V}_b^{\bar\lambda}(B,S) \equiv{}& \xi\,\mathbb{E}_{\tilde p}\big[\,\alpha_b(\tilde p;B,S)\,V_b^{\bar\lambda}(B-1,S-1) + (1-\alpha_b(\tilde p;B,S))\,V_b^{\bar\lambda}(B,S)\,\big|\,\pi_s(B,S)\,\big] \\ &+ (1-\xi)\,\mathbb{E}_{\tilde p}\big[\,\alpha_s(\tilde p;B,S)\,V_b^{\bar\lambda}(B-1,S-1) + (1-\alpha_s(\tilde p;B,S))\,V_b^{\bar\lambda}(B,S)\,\big|\,\pi_b(B,S)\,\big]\ , \\ \check{V}_s^{\bar\lambda}(B,S) \equiv{}& \xi\,\mathbb{E}_{\tilde p}\big[\,\alpha_b(\tilde p;B,S)\,V_s^{\bar\lambda}(B-1,S-1) + (1-\alpha_b(\tilde p;B,S))\,V_s^{\bar\lambda}(B,S)\,\big|\,\pi_s(B,S)\,\big] \\ &+ (1-\xi)\,\mathbb{E}_{\tilde p}\big[\,\alpha_s(\tilde p;B,S)\,V_s^{\bar\lambda}(B-1,S-1) + (1-\alpha_s(\tilde p;B,S))\,V_s^{\bar\lambda}(B,S)\,\big|\,\pi_b(B,S)\,\big]\end{aligned}$$

are the continuation values conditional on some other traders being selected in the match.

The system of equations has a unique solution by the standard fixed-point argument. Indeed, one can replace V_b by $W_s \equiv 1 - V_b$ and verify that the previous equations can be understood as an operator which maps any pair of functions $(V_s, W_s) : \mathbb{Z}^2 \to \mathbb{R}^2$ to another pair of similar functions, and that such operator satisfies the sufficient Blackwell conditions for a contraction (where we use the boundedness of λ, γ_b and γ_s).

Then, using the principle of optimality, we define $\{(\pi_\theta, \alpha_\theta)\}_{\theta\in\{b,s\}}$ to be an ε-MPE if for each state (B, S) and $\theta \in \{b, s\}$, fixing the continuation values in all other states (solving the previous system of equations), as well as the strategy and continuation value of type $\bar{\theta} \neq \theta$ at (B, S), we have that $(\pi_\theta(B, S), \alpha_\theta(\cdot; B, S))$ maximizes the value of $V_\theta(B, S)$.

A.2 Proofs of the results

Proof of Theorem 3.1

Proof. Theorem 3.1 is an immediate consequence of Corollary 4.1 (proven below) and the following result:

Lemma A.1. *Condition 1 implies Condition 3.*

Proof. Assume that Condition 1 holds. Then, using simple algebra, we have

$$\frac{\gamma_b(B,1)}{\gamma(B,1)+\rho}-\frac{\gamma_b(B-1,0)}{\gamma(B-1,0)+\rho}=\underbrace{\frac{\gamma_b(B,1)-\gamma_b(B-1,0)}{\gamma(B-1,0)+\rho}}_{\leq\frac{1}{3}\rho/(\gamma(B-1,0)+\rho)}+\underbrace{\frac{\gamma_b(B,1)}{\gamma(B,1)+\rho}}_{\leq 1}\underbrace{\frac{\gamma(B-1,0)-\gamma(B,1)}{\gamma(B-1,0)+\rho}}_{\leq\frac{2}{3}\rho/(\gamma(B-1,0)+\rho)}$$
$$\leq \frac{\rho}{\gamma(B-1,0)+\rho}\ ,$$

where we used the fact that, as an immediate consequence of Condition 1, we have $\gamma(B-1,0)-\gamma(B,1) \leq \frac{2}{3}\rho$. □

□

Proof of Lemma 3.1

Proof. See the proof of Theorem 4.1 and Corollary 4.1 (in particular, equations (A.12) and (A.13)). □

Proof of Proposition 3.1

Proof. Equation (3.4) follows immediately from equations (3.1)-(3.3), where

$$x \equiv \frac{\frac{\rho}{\gamma(1,1)+\rho}\,\xi(1,1) + \frac{\gamma_b(1,1)}{\gamma(1,1)+\rho}\,V_s(-1) + \frac{\gamma_s(1,1)}{\gamma(1,1)+\rho}\,V_s(1) - \frac{\gamma_b(0,0)}{\gamma(0,0)+\rho}\,V_s(-1) - \frac{\gamma_s(0,0)}{\gamma(0,0)+\rho}\,V_s(1)}{\frac{\rho}{\gamma(0,0)+\rho}} \ . \tag{A.6}$$

We now prove that p is decreasing. We prove the case $N \equiv S - B > 0$ (the other cases are proven analogously). We can use equation (3.1) to obtain

$$\begin{aligned} p(N) - p(N-1) &= -\frac{\rho}{\gamma_b+\gamma_s+\rho}\,p(N-1) + \frac{\gamma_s}{\gamma_b+\gamma_s+\rho}\,(p(N+1) - p(N-1)) \\ &= -\frac{\rho}{\gamma_b+\rho}\,p(N-1) + \frac{\gamma_s}{\gamma_b+\rho}\,(p(N+1) - p(N)) \ , \end{aligned}$$

where γ_b and γ_s are evaluated at state $(0, N)$, and the second expression is obtained rearranging the terms of each side of the first equality. Clearly, if $p(N+1) - p(N) \leq 0$ then $p(N) - p(N-1) \leq 0$. Since, by equation (3.1), we have that $\lim_{N\to\infty} p(N) = 0$, we also have that for each $N \in \mathbb{Z}$ there exists $N' > N$ such that $p(N'+1) - p(N') \leq 0$. Therefore, using a standard induction argument, we have that $p(\cdot)$ is decreasing function. □

Proof of Corollary 3.1

Proof. Assume $N_t > 0$ (the other case is analogous). Notice that, as $\Delta \to 0$, we can write

$$\begin{aligned} \mathbb{E}_t[p_{t+\Delta}] = &\,(1 - \gamma_b(0, N_t)\,\Delta - \gamma_s(0, N_t)\,\Delta)\,p(N_t) \\ &+ \gamma_b(0, N_t)\,\Delta\,p(N_t-1) + \gamma_s(0, N_t)\,\Delta\,p(N_t+1) + O(\Delta^2) \ . \end{aligned}$$

As a result, we have that

$$\lim_{\Delta\to 0} \frac{\mathbb{E}_t[p_{t+\Delta} - p_t]}{\Delta} = \gamma_b(0, N_t)\,\big(p(N_t-1) - p(N_t)\big) + \gamma_s(0, N_t)\,\big(p(N_t+1) - p(N_t)\big) \ .$$

Using equation (3.1) the result holds. □

Proof of Corollary 3.2

Proof. Notice that when $\gamma_\theta(0, 0) = \gamma_\theta(1, 1)$ then x defined in the proof of Proposition 3.1 (see equation (A.6)) is equal to $\xi(1, 1)$. Let F the ergodic distribution of N, so the expected price under such a distribution is

$$\mathbb{E}[p(\tilde{N})|F] = \sum_{N\in\mathbb{Z}} F(\{N\})\,p(N) \ .$$

It is also the case that

$$\lim_{t\to\infty} \mathbb{E}[p_t] = \mathbb{E}[p(\tilde{N})|F] = \lim_{t\to\infty} \mathbb{E}\left[\int_t^\infty e^{-\rho\,(s-t)}\,(\mathbb{I}_{N_s<0}+x\,\mathbb{I}_{N_s=0})\,\rho\,\mathrm{d}s\right]$$
$$= \mathbb{E}\big[I_{\tilde{N}<0}+x\,\mathbb{I}_{\tilde{N}=0}\big|F\big] = F(\mathbb{Z}_{--}) + \xi(1,1)\,F(\{0\})\ .$$

This proves that the ergodic mean of the market price is independent of ρ.

Note now that equation (5.4) (which follows from equations (3.1)-(3.3)) holds in our base model (not only the case of centralized price-posting). Consider an increase on discount rate from ρ_1 to ρ_2, with $\rho_1 < \rho_2$, and let $p^{\rho_i}(\cdot)$ denote the market price function for each ρ_i, $i = 1, 2$. Assume that $p^{\rho_1}(0) \geq p^{\rho_2}(0)$ (the reverse is analogous). In this case, for all $N > 0$ the price $p^{\rho_1}(N) > p^{\rho_2}(N)$: we have that, using τ to denote the (stochastic) time it takes for the market to become balanced (which is independent of ρ), we can write (using equation (5.4))

$$p^{\rho_1}(N) = \mathbb{E}\big[\,e^{-\rho_1\,\tau}\,\big|\,N_0\!=\!N\big]\,p^{\rho_1}(0) > \mathbb{E}\big[\,e^{-\rho_2\,\tau}\,\big|\,N_0\!=\!N\big]\,p^{\rho_2}(0) = p^{\rho_2}(N)\ . \tag{A.7}$$

Since the mean of the market price under the ergodic distribution is independent of the discount rate there must be some $N < 0$ such that $p^{\rho_1}(N) < p^{\rho_2}(N)$. Let $\bar{N}$ be the maximum satisfying this property. Notice that equation (5.4) can be rewritten, for any $N \leq \bar{N} < 0$ and $i \in \{1, 2\}$, as

$$p^{\rho_i}(N) = 1 - \mathbb{E}\big[\,e^{-\rho_i\,\bar{\tau}}\,\big|\,N_0\!=\!N\big]\,(1-p^{\rho_i}(\bar{N}))$$

where $\bar{\tau}$ is the first time where $N_t = \bar{N}$. It is then clear, using equation (A.7) and $p^{\rho_1}(N) < p^{\rho_2}(N)$, that for all $N \leq \bar{N}$ we have $p^{\rho_1}(N) < p^{\rho_2}(N)$. Thus, in fact, $\bar{N}$ is such that

$$p^{\rho_1}(N) \geq p^{\rho_2}(N) \text{ for all } N > \bar{N} \text{ and } p^{\rho_1}(N) < p^{\rho_2}(N) \text{ for all } N \leq \bar{N}\ .$$

This property (and the fact that the ergodic distribution of N is independent of the discount rate) ensures that the distribution of $p^{\rho_2}(N)$ is a mean-preserving spread of $p^{\rho_1}(N)$. □

Proof of Proposition 3.2

Proof. If the ergodic distribution F does not have 0 in its support, then the support is either entirely contained in $\mathbb{N}$ or in $-\mathbb{N}$. In this case, it is clear from equation (3.4) that either either $\lim_{\rho\to 0} p(N) = 0$ for all $N \in \mathbb{Z}$ (if the support of F is a subset of $\mathbb{N}$) or

$\lim_{\rho\to 0} p(N) = 1$ for all $N \in \mathbb{Z}$ (if the support of F is a subset of $-\mathbb{N}$). Assume then that the support of F contains 0.

Let us define the random variable $\tilde{N}_t$ which increases by one unit at a Poisson rate $\tilde{\gamma_s}(\tilde{N}_t)$ and decreases by one unit at Poisson rate $\tilde{\gamma_b}(\tilde{N}_t)$, where

$$(\tilde{\gamma_b}(\tilde{N}_t), \tilde{\gamma_s}(\tilde{N}_t)) = \begin{cases} (\gamma_b(-\tilde{N}_t, 0), \gamma_s(-\tilde{N}_t, 0)) & \text{if } \tilde{N}_t < 0, \\ (\gamma_b(1,1), \gamma_s(1,1)) & \text{if } \tilde{N}_t = 0, \\ (\gamma_b(0, \tilde{N}_t), \gamma_s(0, \tilde{N}_t)) & \text{if } \tilde{N}_t > 0. \end{cases}$$

Notice that $\tilde{N}_t$ evolves similarly to the equilibrium value of the net state N_t. The only difference between two dynamics arises from the fact that when the state is balanced: $\tilde{N}_t$ changes according to the arrival rates $(\gamma_b(1,1), \gamma_s(1,1))$, while the corresponding arrival rates for N_t are $(\gamma_b(0,0), \gamma_s(0,0))$. It is then clear that, using equations (3.1)-(3.3), we can replace equation (3.4) by

$$p(\tilde{N}_0) = \mathbb{E}\left[\int_0^\infty e^{-\rho\, t} \, (\mathbb{I}_{\tilde{N}_t > 0} + \xi(1,1) \, \mathbb{I}_{\tilde{N}_t = 0}) \, \rho \, \mathrm{d}t \right],$$

where now x is replaced by $\xi(1,1)$.

Given that N_t has an ergodic distribution with 0 in its support, $\tilde{N}_t$ also has an ergodic distribution with 0 in its support. Denote it by $\tilde{F}$. Hence, we have that, for any $\tilde{N}_0$,

$$\lim_{\rho\to 0} p(\tilde{N}_0) = \lim_{t\to\infty} p(\tilde{N}_t) = \Pr(\tilde{N} < 0 | \tilde{F}) + \xi(1,1) \, \Pr(\tilde{N} = 0 | \tilde{F}) \; .$$

As we see, the distribution of prices degenerates, as $\rho \to 0$, to the term on the right hand side of the previous expression, which proves our result. □

Proof of Theorem 4.1 and Corollary 4.1

Proof. We divide the proof in 5 steps:

Step 1. Continuation values for a full-trade strategy profile. We begin proving that, under Condition 2 and for each $\varepsilon > 0$, a full-trade strategy strategy profile (that is, such that on-path offers are accepted with probability one) is an ε-MPE whenever $\bar{\lambda}$ is high enough. We fix throughout the proof the following full-trade strategy profile, and we denote the corresponding continuation values using $V_b^{\bar{\lambda}}$ and $V_s^{\bar{\lambda}}$:

1. the price offer is $1 - V_b^{\bar{\lambda}}(B, S)$ if the seller makes the offer, and $V_s^{\bar{\lambda}}(B, S)$ if the buyer makes the offer, and

2. buyers accept all offers equal or lower than $1 - V_b^{\bar{\lambda}}(B,S)$ and reject all other offers, and sellers accept all offers equal or higher than $V_s^{\bar{\lambda}}(B,S)$, and accept all other offers.

Hence, under such a full-trade strategy profile, equations (A.1) and (A.2) for the continuation values of the sellers (analogous equations can be obtained for the buyers) can be written as

$$V_s^{\bar{\lambda}}(B,S) = \frac{\frac{1}{S}\,\bar{\lambda}\,\ell\,\big(\xi\,(1-V_b^{\bar{\lambda}}(B,S)) + (1-\xi)\,V_s^{\bar{\lambda}}(B,S)\big)}{\bar{\lambda}\,\ell+\gamma+\rho} + \frac{\frac{S-1}{S}\,\bar{\lambda}\,\ell\,V_s^{\bar{\lambda}}(B-1,S-1)}{\bar{\lambda}\,\ell+\gamma+\rho} + \frac{\gamma\,\tilde{V}_s^{\bar{\lambda}}(B,S)}{\bar{\lambda}\,\ell+\gamma+\rho}\ ,$$

where $\tilde{V}_s^{\bar{\lambda}}(B,S)$ is defined in equation (A.3) as the expected continuation payoff of a seller conditional on a trader arriving in the market, and where throughout the proof γ, ℓ and ξ are evaluated at (B,S).

Notice that $V_b^{\bar{\lambda}}$ and $V_s^{\bar{\lambda}}$ have well defined limits when $\bar{\lambda} \to \infty$, denoted respectively V_b and V_s. Indeed, we have that V_b satisfies (an analogous equation can be obtained for V_s):

$$V_s(B,S) = \begin{cases} \frac{1}{S}\big(\xi\,(1-V_b(B,S)) + (1-\xi)\,V_s(B,S)\big) + \frac{S-1}{S}\,V_s(B-1,S-1) & \text{if } B,S \geq 1, \\ \frac{\gamma\,\tilde{V}_s(B,S)}{\gamma+\rho} & \text{if } S > B = 0, \end{cases} \tag{A.8}$$

where $\tilde{V}_s$ is defined as the limit of $\tilde{V}_s^{\bar{\lambda}}$ (defined in equation (A.3)) as $\bar{\lambda} \to \infty$, so for each $\theta \in \{b,s\}$:

$$\tilde{V}_\theta(B,S) \equiv \frac{\gamma_b}{\gamma}\,V_\theta(B+1,S) + \frac{\gamma_s}{\gamma}\,V_\theta(B,S+1)\ . \tag{A.9}$$

When $\bar{\lambda}$ is large, the continuation value of a seller in a "full trade strategy profile" when there are buyers in the market depends on whether she is he next seller to trade or not, while when there are no buyers in the market, it is affected by the delay and the type of the next trader who arrives.

Step 2. Proof that $V_b^{\bar{\lambda}}(B,S) + V_s^{\bar{\lambda}}(B,S) \to 1$ as $\bar{\lambda} \to \infty$: Consider first a state (B,S) such that $S > B \geq 1$. In this case we have that

$$V_s^{\bar{\lambda}}(B,S) \;\to\; V_s(0,S-B) = \frac{\gamma(0,S-B)\,\tilde{V}_s(0,S-B)}{\gamma(0,S-B)+\rho}\ , \tag{A.10}$$

$$V_b^{\bar{\lambda}}(B,S) \;\to\; 1 - V_s(0,S-B) = 1 - \frac{\gamma(0,S-B)\,\tilde{V}_s(0,S-B)}{\gamma(0,S-B)+\rho}\ , \tag{A.11}$$

Indeed, notice that when $B=1$, we can use the equation analogous to equation (A.8) for the buyers to obtain

$$V_b(1,S) = \xi(1,S)\, V_b(1,S) + (1-\xi(1,S))\,(1-V_s(1,S)) \quad \Rightarrow \quad V_b(1,S) = 1-V_s(1,S)\ .$$

Then, proceeding inductively from $B=1$ to any value of B, and using in each step that $V_b(B-1,S-1)+V_s(B-1,S-1)=1$, one can easily prove that equations (A.10) and (A.11) hold for any (B,S) with $S>B\geq 1$. This is intuitive: when the market has excess supply, agents on the long side of the market (in this case, the sellers) end up Bertrand competing.

The case $B>S\geq 1$ one can proceed analogously, and obtain that, in general, when $B\neq S$ we have

$$V_s(B,S) = 1 - V_b(B,S) = \begin{cases} \frac{\gamma(0,S-B)}{\gamma(0,S-B)+\rho}\, \tilde{V}_s(0,S-B) & \text{if } B<S, \\ 1 - \frac{\gamma(B-S,0)}{\gamma(B-S,0)+\rho}\, \tilde{V}_b(B-S,0) & \text{if } B>S. \end{cases} \tag{A.12}$$

Consider now the case $S=B\geq 1$. If $S=B=1$ then the payoffs are computed similarly to a standard Rubinstein bargaining game with two players, where the seller makes offers at rate $\xi(1,1)\,\lambda(1,1)$ and the buyer at rate $(1-\xi(1,1))\,\lambda(1,1)$, and where at an exogenous rate $\gamma(1,1)$ the state changes. Indeed, in this case $V_b^{\bar\lambda}(B,S)$ and $V_s^{\bar\lambda}(B,S)$ solve

$$V_s^{\bar\lambda}(1,1) = \frac{\bar\lambda\,\ell}{\bar\lambda\,\ell+\gamma+\rho}\left(\xi\,(1-V_b^{\bar\lambda}(1,1)) + (1-\xi)\,V_s^{\bar\lambda}(1,1)\right) + \frac{\gamma}{\bar\lambda\,\ell+\gamma+\rho}\,\tilde{V}_s^{\bar\lambda}(1,1)\ ,$$
$$V_b^{\bar\lambda}(1,S) = \frac{\bar\lambda\,\ell}{\bar\lambda\,\ell+\gamma+\rho}\left(\xi\,V_b^{\bar\lambda}(1,1) + (1-\xi)\,(1-V_s^{\bar\lambda}(1,1))\right) + \frac{\gamma}{\bar\lambda\,\ell+\gamma+\rho}\,\tilde{V}_b^{\bar\lambda}(1,1)\ .$$

So, the price (and the continuation value of the seller) as $\bar\lambda\to\infty$ is given by

$$V_s(1,1) = 1 - V_b(1,1) = \frac{\rho}{\gamma(1,1)+\rho}\,\xi(1,1) + \frac{\gamma(1,1)}{\gamma(1,1)+\rho}\,\tilde{V}_s(1,1)\ ,$$

where we used that $\tilde{V}_b(1,1)+\tilde{V}_s(1,1)=1$ from our previous result (since, when a trader arrives, the market becomes unbalanced, and we show above that $\tilde{V}_b(B,S)+\tilde{V}_s(B,S)=1$ in an unbalanced market, see equation (A.9)). Thus, similarly as before, one can proceed recursively from state $(1,1)$ to any state (B,S) with $B=S$, and show that the payoffs of a seller and a buyer can be written as

$$V_s(B,S) = 1 - V_b(B,S) = \xi(1,1)\left(1-\frac{\gamma(1,1)}{\gamma(1,1)+\rho}\right) + \frac{\gamma(1,1)}{\gamma(1,1)+\rho}\,\tilde{V}_s(1,1) \quad \text{if } B=S. \tag{A.13}$$

It is then clear that Lemma 3.1 holds and that $V_s(B,S) = p(S-B)$, where p solves equations (3.1)-(3.3).

Step 3. Proof that Condition 2 is necessary and sufficient: In this part of the proof we show that Condition 2 is necessary and sufficient for the following property to hold: "for all $\varepsilon > 0$ there exists some $\bar{\lambda}(\varepsilon)$ such that if $\bar{\lambda} > \bar{\lambda}(\varepsilon)$ then the full-trade strategy described above to be a ε-MPE." We do that by verifying that the sellers do not gain from deviating, and analogous arguments can be made for buyers.

Given that, for each state (B,S), the acceptance rule is optimal for both types of traders in the full-trade strategy profile specified above (given that they accept an offer only if it is equal or above their continuation value conditional on rejecting the offer). It is then only left to prove that it is incentive compatible for each seller to make the prescribed offer. Hence, the full-trade strategy profile specified above is an ε-MPE if only if, for all $(B,S) \in \mathbb{N}^2$, the payoff that a seller obtains from making the specified offer $1 - V_b^{\bar{\lambda}}(B,S)$, is not strictly higher than the value from making an acceptable offer plus ε, that is, $V_s^{\bar{\lambda}}(B,S) + \varepsilon$, and the analogous condition holds for the buyers.

Fix some $\bar{\lambda} > 0$. Consider first the state (B,S) with $B > S = 1$. Consider the deviation of a seller in such a state given consisting on rejecting all offers received and making unacceptable offers until the state changes (other deviations can be evaluated similarly). The payoff of the seller from following such a strategy is

$$\frac{\gamma}{\gamma+\rho}\,\tilde{V}_s^{\bar{\lambda}}(B,S)\ .$$

Condition 2 ensures that if $\bar{\lambda}$ is high enough the previous expression is strictly lower than $V_s^{\bar{\lambda}}(B,S)$, so the deviation is suboptimal when $B > S = 1$.[11] Conversely, if equation (4.1) holds with the reverse (strict) inequality, the gain for a seller from deviating is higher than any given $\varepsilon > 0$ if $\bar{\lambda}$ is high enough.

Fix now some (B,S) such that $B > S > 1$. Note that the value from deviating by rejecting all offers received and making unacceptable offers until the state changes is given by

$$\frac{\frac{S-1}{S}\,\bar{\lambda}\,\ell}{\frac{S-1}{S}\,\bar{\lambda}\,\ell+\gamma+\rho}\,V_s^{\bar{\lambda}}(B-1,S-1) + \frac{\gamma}{\frac{S-1}{S}\,\bar{\lambda}\,\ell+\gamma+\rho}\,\tilde{V}_s^{\bar{\lambda}}(B,S)\ . \tag{A.14}$$

[11] Notice that, since $V_s^{\bar{\lambda}}(B,S) = p(S-B)$, the right hand side of equation (4.1) (which corresponds to $B > S$) is equal to $\frac{\gamma}{\gamma+\rho}\tilde{V}_s^{\bar{\lambda}}(B,S)$.

Using the results in Step 1, the previous expression converges to $V_s(B-1, S-1)$ as $\bar{\lambda} \to \infty$. We have then that, by Step 2, both the value from deviating and the value from following the strategy (equal to $V_s^{\bar{\lambda}}(B-1, S-1)$) converge to $1 - V_b(B-S, 0)$ as $\bar{\lambda} \to \infty$. Hence, for each $\varepsilon > 0$ there exists some $\bar{\lambda}(\varepsilon)$ such that deviating is unprofitable whenever $\bar{\lambda} > \bar{\lambda}(\varepsilon)$.

Consider now the case $S > B \geq 1$. The seller's payoff from making an unacceptable offer is again equal to expression (A.14). Thus, we have that, as before, both the value from deviating and the value from following the strategy converge to $V_s(0, S-B)$ as $\bar{\lambda} \to \infty$. Consequently, for all $\varepsilon > 0$ there is some $\bar{\lambda}(\varepsilon)$ such that the gain from deviating is lower than any ε if $\bar{\lambda} > \bar{\lambda}(\varepsilon)$.

Finally, consider the case $B = S \geq 1$. If $B = S = 1$ we have that the seller's continuation value from rejecting all offers received and making unacceptable offers until the state changes is

$$\frac{\gamma}{\bar{\lambda}\ell + \gamma + \rho} \tilde{V}_s^{\bar{\lambda}}(B, S) \; .$$

The limit of the right hand side of the previous expression as $\bar{\lambda} \to 0$ is clearly lower than the equilibrium continuation value, given in equation (A.13). If, instead, $S = B > 1$ we have that, again, both the value from deviating (given in (A.14)) and the value from following the strategy converge to the same value as $\bar{\lambda} \to \infty$, now equal to $V_s(1, 1)$. Thus, for all $\varepsilon > 0$ there is some $\bar{\lambda}(\varepsilon)$ such that the gain from deviating is lower than any ε if $\bar{\lambda} > \bar{\lambda}(\varepsilon)$.

Notice that a sufficient condition for following the full-trade strategy defined above to be incentive compatible for $\bar{\lambda}$ sufficiently large when $B > S = 1$ is that

$$\frac{\gamma(B,1)}{\gamma(B,1)+\rho} \tilde{V}_s(B,1) < V_s(B,1) = 1 - V_s(B-1,0) = 1 - \frac{\gamma(B,1)}{\gamma(B-1,0)+\rho} \tilde{V}_s(B-1,0) \; .$$

Using the definition of $\tilde{V}_s$ (in expression (A.9)) and the fact that, by Proposition 3.1, $p(B-S) = V_s(B-S, 0)$ is a decreasing function of $S-B$, the previous inequality holds whenever[12]

$$\frac{\gamma_b(B-S+1,1)}{\gamma(B-S+1,1)+\rho} - \frac{\gamma_b(B-S,0)}{\gamma(B-S,0)+\rho} < \frac{\rho}{\gamma(B-S,0)+\rho} \; .$$

This is exactly the same condition as equation (4.3) in Condition 3. Equation (4.4) can be obtained analogously in the case $S > B = 1$, and checking the deviating incentives by

[12]The inequality can is obtained when, using the expressions for $\tilde{V}_s(B,1)$ and $\tilde{V}_s(B-1,0)$ (see equation (A.9)), we use the fact that $V_s(B-S+2,1) \leq 1$ and $V_s(B-S+1,2) \geq 0$.

buyers. Notice that the incentive compatibility holds trivially when $S = B$, so Condition 3 is a sufficient condition.

Step 4. Difference in payoffs: In order to prove that under Condition 2 there is no trade delay in the frictionless-bargaining limit it is useful to first show a technical result on the speed of convergence of $V_b^{\bar\lambda}(B,S) - V_b^{\bar\lambda}(B-1,S-1)$ to 0 as $\bar\lambda \to \infty$. Using equation (A.8) and simple algebra it is easy to show that, for any $B > 1$, we have that, as $\bar\lambda \to \infty$,

$$
\begin{aligned}
&V_b^{\bar\lambda}(B,1) - V_b^{\bar\lambda}(B-1,0) \\
&\quad = \frac{B\,\gamma\,\xi\,\tilde V_b(B,1) - \gamma\,(1-\xi)\,\tilde V_s(B,1) - (\gamma+\rho)\big(((B-1)\,\xi+1)\,V_b(B-1,0)-1+\xi\big)}{(B-1)\,\xi\,\bar\lambda\,\ell} \\
&\qquad + O(\bar\lambda^{-2}) \ ,
\end{aligned}
$$

where $\gamma \equiv \gamma(B,1)$, $\xi \equiv \xi(B,1)$ and $\ell \equiv \ell(B,1)$. Then, proceeding recursively, it is easy to show that $V_b^{\bar\lambda}(B,S) - V_b^{\bar\lambda}(B-1,S-1) = O(\bar\lambda^{-1})$ as $\bar\lambda \to \infty$. The same holds for $V_s^{\bar\lambda}$.

Step 5. Proof of no delay in the limit: We now prove that under Condition 2, as $\bar\lambda \to 0$, transactions happen increasingly fast (conditional on both buyers and sellers being present in the market) in any ε-MPE for $\varepsilon > 0$ small. To prove this, fix a sequence $(\varepsilon_n)_n$ decreasing toward 0 and a sequence $(\bar\lambda_n)_n$ tending to $+\infty$ (so that the full-trade strategy described above is an equilibrium). We also fix, for each n, a ε_n-equilibrium. For each state $(B,S) \in \mathbb{N}^2$, let $\beta^{\bar\lambda_n}(B,S)$ be the probability that there is trade in a meeting between a buyer and a seller in state (B,S) (which is equal to 1 in a full-trade strategy profile). Assume there exists, for each n, an equilibrium where $\beta^{\bar\lambda_n}(B,S) < 1$ for some (B,S). Let $V_b^{\dagger\bar\lambda_n}$ and $V_s^{\dagger\bar\lambda_n}$ denote the respective continuation values of a buyer and a seller in this equilibrium (we keep $V_b^{\bar\lambda_n}$ and $V_s^{\bar\lambda_n}$ for the full-trade equilibrium). Let $\Delta_\theta^{\bar\lambda_n}(B,S) \equiv V_\theta^{\dagger\bar\lambda_n}(B,S) - V_\theta^{\bar\lambda_n}(B,S)$, for $\theta \in \{b,s\}$. Then, we have that

$$
\begin{aligned}
\Delta_s^{\bar\lambda_n}(B,S) = & \frac{\frac{1}{S}\,\bar\lambda_n\,\ell\,\big(-\xi\,\Delta_b^{\bar\lambda_n}(B,S) + (1-\xi)\,\Delta_s^{\bar\lambda_n}(B,S)\big)}{\bar\lambda_n\,\ell+\gamma+\rho} \\
& + \frac{\frac{S-1}{S}\,\bar\lambda_n\,\ell\,\beta^{\bar\lambda_n}\,\Delta_s^{\bar\lambda_n}(B-1,S-1)}{\bar\lambda_n\,\ell+\gamma+\rho} \\
& + \frac{\frac{S-1}{S}\,\bar\lambda_n\,\ell\,(1-\beta^{\bar\lambda_n})\,(\Delta_s^{\bar\lambda_n}(B,S)+V_s^{\bar\lambda_n}(B,S)-V_s^{\bar\lambda_n}(B-1,S-1))}{\bar\lambda_n\,\ell+\gamma+\rho} \\
& + \frac{\gamma\,\mathbb{E}[\Delta_s^{\bar\lambda_n}(B',S')|(B,S)]}{\bar\lambda_n\,\ell+\gamma+\rho}
\end{aligned}
$$

where in this expression, consistently with the rest of the expressions in this proof, ℓ, γ, ξ and $\beta^{\bar\lambda_n}$ are evaluated at (B,S), and where the expectation in the last expression is taken over the new state (B', S') conditional on a trader arriving when the state is (B, S) (see equation (A.9)). A similar expression can be obtained for $\Delta_b^{\bar{\lambda}_n}(B,S)$. Let $\bar{\Delta}^{\bar{\lambda}_n} \equiv \sup_{(B,S)}\{|\Delta_b^{\bar{\lambda}_n}(B,S)|, |\Delta_s^{\bar{\lambda}_n}(B,S)|\}$. Assume that a state (S, B) is such that $|\Delta_s^{\bar{\lambda}_n}(B,S)| \geq \frac{(\bar{\lambda}_n\,\ell+\gamma+\rho/2)}{\bar{\lambda}_n\,\ell+\gamma+\rho}\,\bar{\Delta}^{\bar{\lambda}_n}$ (note that, by definition of $\bar{\Delta}^{\bar{\lambda}_n}$, and since γ is bounded, there exists a state with this property or with the analogous property with $|\Delta_b^{\bar{\lambda}_n}(B,S)|$). Then, from the previous expression, we obtain

$$\frac{(\bar{\lambda}_n\,\ell+\gamma+\rho/2)}{\bar{\lambda}_n\,\ell+\gamma+\rho}\,\bar{\Delta}^{\bar{\lambda}_n} \leq \frac{(\bar{\lambda}_n\,\ell+\gamma)\,\bar{\Delta}^{\bar{\lambda}_n}}{\bar{\lambda}_n\,\ell+\gamma+\rho} + \frac{\frac{S-1}{S}\,(1-\beta^{\bar{\lambda}_n})\,\bar{\lambda}_n\,\ell\left|V_s^{\bar{\lambda}_n}(B,S)-V_s^{\bar{\lambda}_n}(B-1,S-1)\right|}{\bar{\lambda}_n\,\ell+\gamma+\rho}\,,$$

that is,

$$\frac{\bar{\Delta}^{\bar{\lambda}_n}}{2}\,\rho \leq \frac{S-1}{S}\,(1-\beta^{\bar{\lambda}_n})\,\bar{\lambda}_n\,\ell\left|V_s^{\bar{\lambda}_n}(B,S)-V_s^{\bar{\lambda}_n}(B-1,S-1)\right|\,. \tag{A.15}$$

Now assume, for the sake of contradiction, that we have that $\lim_{n\to\infty}\bar{\Delta}^{\bar{\lambda}_n} > 0$. Let (B_n, S_n) be a state satisfying (A.15) when the matching rate is $\bar{\lambda}_n\,\ell$. From our previous analysis we have that when $n\to\infty$ we have $V_s^{\bar{\lambda}_n}(B_n,S_n)-V_s^{\bar{\lambda}_n}(B_n-1,S_n-1) = O(\bar{\lambda}_n^{-1})$. Therefore, necessarily, $\liminf_{n\to\infty}(1-\beta^{\bar{\lambda}_n}) > 0$. Nevertheless, this implies that the rate at which an offer is accepted in equilibrium increases towards infinity as n increases and, as a result, $V_s^{\bar{\lambda}_n}(B,S) \to V_s(B,S)$ for any (B,S).

Therefore, as ε decreases, the continuation values of any equilibrium become arbitrarily close to $(V_b(B,S), V_s(B,S))$ for any (B,S) if $\bar{\lambda}$ is large enough. This implies that the trade delay necessarily shrinks to 0 as $\varepsilon\to 0$: if there is trade delay, the realized trade surplus is lower than in a full-trade strategy, so necessarily the continuation payoffs of either buyers or sellers are strictly lower for some state. □

Proof of Proposition 5.1

Proof. The proof is analogous to the proof of Theorem 4.1 and Corollary 4.1. □

Proof of Proposition 5.2

Proof. Fix $\bar{\lambda} > 0$. If a full-trade strategy profile is an ε-MPE, the standard Bertrand-competition and Diamond paradox arguments imply that the price at a given state (B,S) satisfies the following equations hold up to a factor $O(\varepsilon)$:

- If $S > 1$ then $p(B,S) = V_s(B-1, S-1)$.
- If $S = 1$ then $p(B,S) = 1 - V_b(B,S)$.

We first verify that sellers do not have an incentive to deviate. It is clear a seller does not have an incentive to deviate when $S > 1$, since offering a lower price lowers her payoff (in this case she sells her good for sure, but she obtains a payoff lower than the continuation value of making an unacceptable offer), and if she increases it she obtains the same continuation payoff (since the equilibrium price is equal to the continuation value if the offer made rejected by the buyer). To verify that there is no incentive to deviate when $S = 1$, we proceed as in Step 3 of the proof of Theorem 4.1 and Corollary 4.1. In this case, it is not profitable to deviate compatible only if the payoff from offering the price $1 - V_b(B, 1)$ (which is accepted), is not lower than the payoff from making unacceptable offers and waiting for a trader to arrive, $\frac{\gamma(B,1)}{\gamma(B,1)+\rho} \tilde{V}_s(B, 1)$, where $\tilde{V}_\theta$ is defined in equation (A.9). Then, the same argument can be used to argue that, if Condition 2 holds, ε is small and $\bar{\lambda}$ is high enough, the seller is worse off when she deviates.

Finally, we verify that buyers do not have an incentive to deviate. When $S = 1$ it is clear that buyers have no incentive to deviate, as they are offered their continuation value from rejecting the offer. When $S > 1$, buyers have an incentive to deviate only if their payoff from accepting the specified price offer $V_s(B-1, S-1)$ is strictly lower than the payoff from rejecting it. As $\bar{\lambda}$ increases, the payoff from rejecting the offer $V_s(B-1, S-1)$ and waiting for a trader to arrive is

$$\begin{cases} \frac{\gamma(1,S-B+1)}{\gamma(1,S-B+1)+\rho} \tilde{V}_b(1, S-B+1) & \text{if } S \geq B-1 \ , \\ \frac{\gamma(B-S,0)}{\gamma(B-S,0)+\rho} \tilde{V}_b(B-S, 0) & \text{if } S < B-1 \ . \end{cases}$$

Again, arguing similarly to the proof of Theorem 4.1 and Corollary 4.1 we have that if Condition 2 is satisfied and $\bar{\lambda}$ is high enough, such a deviation is suboptimal for a buyer. The absence of trade delay as $\bar{\lambda} \to \infty$ is proven analogously to Theorem 4.1 and Corollary 4.1. □

References

Atakan, A. E., 2006, "Assortative matching with explicit search costs," *Econometrica*, 74(3), 667–680.

Burdett, K., and M. G. Coles, 1997, "Steady state price distributions in a noisy search equilibrium," *Journal of Economic Theory*, 72(1), 1–32.

Cai, H., 2000, "Delay in multilateral bargaining under complete information," *Journal of Economic Theory*, 93(2), 260–276.

Coles, M. G., and A. Muthoo, 1998, "Strategic Bargaining and Competitive Bidding in a Dynamic Market Equilibrium," *Review of Economic Studies*, 65(2), 235–60.

Condorelli, D., A. Galeotti, and L. Renou, 2016, "Bilateral trading in networks," *The Review of Economic Studies*, 84(1), 82–105.

Diamond, P. A., 1971, "A model of price adjustment," *Journal of Economic Theory*, 3(2), 156–168.

Elliott, M., and F. Nava, 2016, "Decentralized Bargaining: Efficiency and the Core," working paper.

Gale, D., 1987, "Limit theorems for markets with sequential bargaining," *Journal of Economic Theory*, 43(1), 20–54.

——— , 2000, *Strategic foundations of general equilibrium: dynamic matching and bargaining games*. Cambridge University Press.

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

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

Lauermann, S., W. Merzyn, and G. Virág, 2017, "Learning and Price Discovery in a Search Market," working paper.

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

Manea, M., 2011, "Bargaining in stationary networks," *The American Economic Review*, 101(5), 2042–2080.

Osborne, M. J., and A. Rubinstein, 1990, *Bargaining and markets*. Academic press.

Rubinstein, A., 1982, "Perfect equilibrium in a bargaining model," *Econometrica: Journal of the Econometric Society*, pp. 97–109.

Rubinstein, A., and A. Wolinsky, 1985, "Equilibrium in a market with sequential bargaining," *Econometrica: Journal of the Econometric Society*, pp. 1133–1150.

Said, M., 2011, "Sequential auctions with randomly arriving buyers," *Games and Economic Behavior*, 73(1), 236–243.

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.

Shimer, R., and L. Smith, 2000, "Assortative matching and search," *Econometrica*, 68(2), 343–369.

Talamàs, E., 2016, "Prices and Efficiency in Networked Markets," working paper.

Taylor, C. R., 1995, "The Long Side of the Market and the Short End of the Stick: Bargaining Power and Price Formation in Buyers', Sellers', and Balanced Markets," *The Quarterly Journal of Economics*, 110(3), 837–55.