

Efficiency in Decentralized Markets with Aggregate Uncertainty*

Braz Camargo[†]

Dino Gerardi[‡]

Lucas Maestri[§]

June 2019

Abstract

We study efficiency in non-stationary decentralized markets with common-value uncertainty and correlated asset values. There is an equal mass of buyers and sellers and payoffs from trade depend on an aggregate state, which only the sellers know. Buyers and sellers are randomly and anonymously matched in pairs over time, and buyers make the offers. We show that all equilibria become efficient as trading frictions vanish.

*We thank the Editor and two referees for their comments and suggestions. We are also grateful to seminar participants at FGV EESP, the SED meetings in Toulouse, the SAET meetings in Rio, and the Workshop on Decentralized Markets with Informational Asymmetries at Collegio Carlo Alberto for their input. Braz Camargo and Lucas Maestri gratefully acknowledge financial support from CNPq. Dino Gerardi gratefully acknowledges financial support from the European Research Council (Consolidator Grant INFOASYMMARKETS). This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001.

[†]Sao Paulo School of Economics - FGV. E-mail: braz.camargo@fgv.br.

[‡]Collegio Carlo Alberto, Università di Torino. E-mail: dino.gerardi@carloalberto.org.

[§]Corresponding author. FGV EPGE. E-mail: lucas.maestri@fgv.br

1 Introduction

Market efficiency is a central concern in economics. In idealized markets, trade is centralized and information is perfect. In this case, the first welfare theorem shows that market outcomes are efficient. However, trade is often decentralized: rather than trade taking place at a single price that clears the market, in many markets buyers and sellers negotiate the terms of trade bilaterally. Moreover, information is typically asymmetric: rather than buyers and sellers being perfectly, and thus equally, informed, in many markets sellers often have better information about underlying features of their assets than buyers.

It is well-known that both decentralized trade and asymmetric information can, by themselves, hurt market efficiency. When trade is decentralized, it may take time for an agent to find a desirable trading partner. This delay in trade represents a loss of efficiency. On the other hand, when sellers are better informed about their assets than buyers, they can use their private information to extract rents from buyers, distorting the terms of trade. This distortion also leads to a loss of efficiency.

A question that has attracted substantial attention is how decentralization of trade and asymmetry of information interact to affect market efficiency. However, the literature on this topic has mainly focused on the case in which asset values are independent across *sellers*, i.e., the value of a seller's asset is independent of the value of any other seller's asset. While the assumption of independent asset values is reasonable in some markets, there are many relevant markets, both real and financial, in which asset values are correlated across sellers. Real estate markets and markets for asset-backed securities are prominent examples.¹

In this paper, we study market efficiency in non-stationary decentralized markets with common value uncertainty and correlated asset values. The environment we consider is as follows. There is an equal mass of buyers and sellers, who enter the market in some initial period. Payoffs from trade depend on an aggregate state, which only the sellers know. Thus, asset values are perfectly correlated across sellers and the asymmetry of information is one-sided. The number of aggregate

¹For instance, in the case of mortgage-backed securities, the loans in the underlying pool of loans backing different securities could have been issued to borrowers with similar characteristics. In the housing market, sellers are typically better informed about neighborhood characteristics than buyers, which affects house values; see Kurlat and Stroebel (2015) for evidence on this.

states is finite and gains from trade are nonnegative in all states. Time is discrete and in every period buyers and sellers in the market are randomly and anonymously matched in pairs with some probability. In a buyer-seller match, the buyer makes a take-it-or-leave-it offer to the seller, which the seller either accepts or rejects. If the seller accepts the offer, then trade takes place and both agents exit the market. Otherwise, the match is dissolved and both agents remain in the market.

Our main result is that, somewhat surprisingly, welfare in *any* equilibrium approaches welfare in the complete information case as trading frictions vanish, i.e., as the real time between two consecutive trading opportunities converges to zero. Thus, by itself, the asymmetry of information between buyers and sellers is not enough to prevent market efficiency when asset values are perfectly correlated across sellers.

The rest of the paper is organized as follows. In the remainder of this section we discuss the related literature. In Section 2 we introduce our environment. In Section 3 we explicitly construct an equilibrium which achieves the first best in the limit as trading frictions vanish. In Section 4 we generalize the result of Section 3 and show that every equilibrium achieves the first best in the limit as trading frictions vanish. In Section 5 we discuss robustness and extensions of our efficiency result. In particular, we show that our efficiency result survives when we relax the assumption of perfect correlation among asset values by introducing private values. We also show that the assumption of one-sided asymmetric information is important for our efficiency result. In Section 6 we conclude. The Appendix contains omitted details and proofs.

Related Literature The literature on market efficiency in decentralized markets with correlated asset values is scant. Similarly to us, Blouin and Serrano (2001) study this question in a market with aggregate uncertainty. There are important differences between our analysis and the analysis in Blouin and Serrano's paper, though. First, we focus on the case in which sellers know the aggregate state but buyers do not. Second, we allow for any finite number of aggregate states, instead of just two, and place no restrictions on payoffs from trade except that gains from trade are nonnegative in every state. Finally, we depart from Blouin and Serrano (2001) in the bargaining protocol. While they consider a stylized bargaining game which amounts to restricting the set of

prices at which trade can take place, we place no restrictions on transaction prices.² Unlike Blouin and Serrano (2001), we show that market outcomes become efficient as trading frictions vanish.³

Golosov, Lorenzoni, and Tsyvinski (2014) study information aggregation in decentralized markets with aggregate uncertainty and divisible goods. They provide conditions under which in the long-run information is fully aggregated and trading outcomes become efficient.⁴ In contrast, we show in an environment with indivisible goods that market efficiency is obtained as trading frictions vanish. As we discuss in Section 4, equilibria need not fully aggregate information, though.

Asriyan, Fuchs, and Green (2017a) study information spillovers in a dynamic market with imperfectly correlated asset values in which sellers are privately informed about the quality of their assets. They show that as long as asset values are sufficiently correlated, making transaction outcomes public, i.e., introducing transaction transparency, leads to multiple equilibria which are Pareto ranked. In our environment there are no information spillovers. The only way buyers can learn about the aggregate state is through their own experience in the market.

There are a number of papers that study decentralized trade with common-value uncertainty and independent asset values; see e.g., Blouin (2003), Camargo and Lester (2014), and Kim (2017). Our efficiency result contrasts strongly with inefficiency results in this literature. When asset values are independent, multiple types of assets co-exist in the market, allowing owners of lower-quality assets to extract informational rents from buyers. This is not possible in the presence of aggregate uncertainty. Our efficiency result also contrasts strongly with inefficiency results in the literature on bargaining with common-value uncertainty; see, e.g., Deneckere and Liang (2006) and Gerardi, Hörner, and Maestri (2014) for models of bargaining between two long-lived parties, and Hörner

²We also differ from Blouin and Serrano (2001) in that we allow for the probability that agents in the market are matched to a trading partner to be smaller than one.

³Serrano (2002) considers the same bargaining protocol of Blouin and Serrano (2001) in a private-values setting and shows that equilibria which become efficient as trading frictions vanish can exist. It is possible to show that the same holds in our setting when price offers are restricted. However, with restricted price offers, there will also exist equilibria which remain inefficient as trading frictions vanish even if there is no asymmetry of information between buyers and sellers (see Serrano and Yosha (1995) for a similar result in a stationary common-values environment).

⁴The seminal reference in the literature on information aggregation in decentralized markets is Wolinsky (1990). Serrano and Yosha (1993) shows that Wolinsky's negative result depends on the assumption of two-sided incomplete information. Blouin and Serrano (2001) extends the analysis in Wolinsky (1990) to non-stationary environments. Other recent papers in the literature on information aggregation in decentralized markets are Lauermaann and Wolinsky (2016) and Asriyan, Fuchs, and Green (2017b).

and Vieille (2009), Daley and Green (2012), Fuchs, Öry, and Skrzypacz (2016), and Kaya and Kim (2017) for models of bargaining between a long-lived seller and a sequence of short-lived buyers. Intuitively, unlike in a single-seller setting, an individual seller cannot affect aggregate behavior in our large market setting.

2 Environment

Time is discrete and indexed by $t \in \{0, 1, \dots\}$. There is an equal mass of buyers and sellers with common discount factor $\delta \in (0, 1)$; trading frictions vanish as δ converges to one.⁵ Each seller can produce one unit of an indivisible good and each buyer wants to consume one unit of the good. The set of (aggregate) states is $\mathcal{K} = \{1, \dots, K\}$ and the probability that the state is $k \in \mathcal{K}$ is $\pi_k > 0$. Sellers know the state, but buyers do not. Agents have quasi-linear preferences. The value to a buyer from consuming the good in state k is v_k , while the cost to a seller of producing the good in the same state is $c_k \geq 0$. We assume nonnegative gains from trade in every state.

Assumption 1. $v_k - c_k \geq 0$ for all $k \in \mathcal{K}$.

Assumption 1 is fairly weak. In particular, single-crossing preferences, i.e., $v_k - c_k$ strictly increasing in k , is not necessary for our results. Moreover, as we show in Section 5, Assumption 1 cannot be relaxed. It implies that the (ex-ante) first-best welfare is

$$W^* = \sum_{k=1}^K \pi_k (v_k - c_k).$$

Trade takes place as follows. In every period $t \geq 0$, a buyer in the market is randomly and anonymously matched to a seller with probability $\lambda \in (0, 1)$ and vice-versa. In each buyer-seller pair, the buyer makes a take-it-or-leave-it offer $p \in \mathbb{R}_+$ to the seller. If the seller accepts the offer, then trade occurs and both agents exit the market. Otherwise, the match is dissolved and both agents remain in the market. The assumption that $\lambda < 1$ ensures that there is a positive mass of agents in the market in every period.⁶

⁵More formally, one can think that agents discount the future at a common rate $\rho > 0$ and $\delta = e^{-\rho\Delta}$, where Δ is the time interval between two consecutive periods. Trading frictions vanish as Δ converges to zero.

⁶In Section 5 we discuss how to extend our analysis to the case in which $\lambda = 1$.

We now define strategies and equilibria. Let \mathcal{H}_t , with typical element h^t , be the set of private histories for an agent in the market in period t .⁷ A behavior strategy for a buyer is a sequence $\sigma^B = \{\sigma_t^B\}$, where $\sigma_t^B : \mathcal{H}_t \rightarrow \Delta(\mathbb{R}_+)$ and $\sigma_t^B(h^t)$ is the (random) price offer that the buyer makes in period t if he is matched to a seller when his private history is h^t . A behavior strategy for a seller is a sequence $\sigma^S = \{\sigma_t^S\}$, where $\sigma_t^S : \mathcal{H}_t \times \mathcal{K} \times \mathbb{R}_+ \rightarrow [0, 1]$ and $\sigma_t^S(h^t, k, p)$ is the probability that the seller accepts an offer of p in period t when the state is k and his private history is h^t . A belief system for a buyer is a sequence $\mu = \{\mu_t\}$, where $\mu_t : \mathcal{H}_t \rightarrow \Delta(\mathcal{K})$ and $\mu_t(h^t)$ is the buyer's (posterior) belief about the state in period t when his private history is h^t . We let Σ and θ denote, respectively, a strategy profile and a profile of belief systems. We consider pairs (Σ, θ) which constitute a perfect Bayesian equilibrium (PBE).

We now compute ex-ante welfare for any strategy profile.⁸ An outcome for a seller is a triple $(k, \mathbf{T}, \mathbf{p})$, where $k \in \mathcal{K}$ is the aggregate state, $\mathbf{T} \in \mathbb{Z}_+ \cup \{\infty\}$ is the time at which the seller trades, and $\mathbf{p} \in P$ is the price at which the seller trades; $\mathbf{T} = \infty$ corresponds to the event in which the seller does not trade.⁹ Together with the prior over the set of aggregate states, a strategy profile Σ uniquely determines a probability distribution over the set of outcomes for each seller. Let \mathbb{E} denote the expectation with respect to this probability distribution. Welfare under Σ is

$$W(\Sigma) = \sum_{k=1}^K \pi_k \mathbb{E} [\delta^{\mathbf{T}} | k] (v_k - c_k);$$

the term $\mathbb{E} [\delta^{\mathbf{T}} | k]$ is the discounted probability of trade in aggregate state k .

Our centralized benchmark is the centralized mechanism in which all sellers announce the aggregate state to the planner and trade takes place at price $p(k) \in [c_k, v_k]$ if all sellers announce the same aggregate state k and trade does not take place otherwise. Truth-telling is satisfied, as no seller has an incentive to lie about the aggregate state when all other sellers report the aggregate state truthfully. Therefore, the first-best is incentive feasible in our environment.¹⁰

⁷A private history for a buyer in the market in period t is a sequence $h^t = (\tilde{p}_0, \dots, \tilde{p}_{t-1})$, where $\tilde{p}_s \in \mathbb{R}_+$ is the buyer's (rejected) price offer in period s if he was matched in this period and $\tilde{p}_s = \emptyset$ if the buyer was not matched in period s . Private histories for sellers are defined similarly.

⁸Since there are finitely many states, the assumption that all states have positive probability ensures that we have ex-ante efficiency if, and only if, we have interim efficiency for each state.

⁹When $\mathbf{T} = \infty$, the seller's transaction price is undetermined. We adopt the convention that $\mathbf{p} = 0$ in this case.

¹⁰When the prices $p(k)$ are distinct, the centralized mechanism described above implements a Fully Revealing Rational Expectations Equilibrium.

3 A Limit-Efficient Equilibrium

We know from the previous section that the first best is incentive feasible. The question of interest is whether the outcome of decentralized trade always approaches the first best as trading frictions vanish. In this section we show that there exists a PBE whose welfare approaches the first-best welfare as δ converges to one. For ease of exposition, we consider the case of two aggregate states and discuss how to extend our equilibrium construction to the case of more than two aggregate states in the Appendix.

Assume, without any loss, that $c_1 < c_2$ and refer to a seller when the aggregate state is k as a type- k seller.¹¹ Consider the following symmetric assessment (Σ, θ) : (i) a type- k seller accepts an offer of p if, and only if, $p \geq c_k$; (ii) a buyer offers $p = c_1$ the first time he is matched to a seller and offers $p = c_2$ afterwards; and (iii) a buyer in the market assigns probability π_1 to $k = 1$ if either he has not made any offer or the highest offer he has made is less than c_1 , otherwise he assigns probability 0 to $k = 1$. Under Σ , all buyers trade after at most two offers, and so welfare approaches the first-best welfare as trading frictions vanish.

We claim that there exists $\delta^* \in (0, 1)$ such that (Σ, θ) is a PBE if $\delta \geq \delta^*$. First notice that buyer beliefs are consistent with Bayes' rule on the path of play. Moreover, sellers' behavior is sequentially rational. Clearly, a type-2 seller has no incentive to deviate. A type-1 seller also has no incentive to deviate by rejecting an offer $p \geq c_1$. Indeed, a type-1 seller who rejects an offer knows that in the future he is matched with probability 1 to a buyer who has not had the chance to make an offer (and so will offer $p = c_1$). Finally, notice that the only possibly profitable deviation for a buyer is to offer a price $p \geq c_2$ and trade immediately if the highest price he has offered so far is smaller than c_1 . The expected payoff from this deviation is bounded above by $\bar{u} = \pi_1 v_1 + \pi_2 v_2 - c_2$. On the other hand, if the highest price offer a buyer has made so far is smaller than c_1 , then the buyer's expected payoff from following the equilibrium strategy is

$$\pi_1(v_1 - c_1) + \pi_2 \delta \sum_{s=0}^{\infty} \delta^s (1 - \lambda)^s \lambda (v_2 - c_2) = \pi_1(v_1 - c_1) + \pi_2 \frac{\delta \lambda}{1 - \delta(1 - \lambda)} (v_2 - c_2),$$

¹¹By re-ordering the states if necessary, we have that $c_1 \leq c_2$. When $c_1 = c_2$, it is immediate to see that there exists a symmetric PBE where buyers offer $p = c_1$ when matched to a seller and sellers accept and offer of p if, and only if, $p \geq c_1$. Welfare in this equilibrium approaches the first-best welfare as trading frictions vanish.

which is greater than \bar{u} as long as δ is sufficiently close to one. This establishes the desired result.

4 Market Efficiency

In Section 3 we constructed a PBE that approaches the first best as trading frictions vanish. This, of course, does not rule out the possibility that there are PBE whose welfare is bounded away from W^* no matter how small trading frictions are. In this section we show that this is not the case.

Theorem 1. *Let $\{\delta_n\}$ be a sequence of discount factors such that $\lim_n \delta_n = 1$. For any sequence $\{(\Sigma_n, \theta_n)\}$ such that (Σ_n, θ_n) is a PBE when $\delta = \delta_n$, the sequence $\{W(\Sigma_n)\}$ converges to W^* .*

In what follows we present a sketch of the proof of Theorem 1. For simplicity, we consider the limiting case in which λ is close enough to one that the probability that a buyer who stays in the market is matched to a seller is essentially equal to one.¹² The proof of the general case—when λ assumes any value in $(0, 1)$ —is in the Appendix; we discuss how to extend the argument that follows to the general case at the end of the section.

Let (Σ, θ) be a PBE when the discount factor is $\delta \in (0, 1)$. Notice that even though (Σ, θ) need not be symmetric, all agents on a given side of the market obtain the same payoff; this property of equilibria holds regardless of the value of λ . In fact, since there is a continuum of buyers and matching is random and anonymous, a buyer can obtain the same payoff as any other buyer by mimicking the other buyer's behavior. The same applies for sellers. Let V^B be the buyers' ex-ante equilibrium payoff and V^k be the type- k seller's equilibrium payoff. Moreover, let V_t^k be the type- k sellers' payoff in period t ; by construction, $V_0^t = V^k$. Since sellers know the aggregate state, V_t^k does not depend on the private history of a seller, only on the period t .

Now observe that for every $k \in \mathcal{K}$ and $s > t \geq 0$, we have that $V_t^k \geq \delta^{s-t} V_s^k$. This follows immediately from the fact that a possibility for a seller in period t is to reject all offers between periods t and $s - 1$, and then follow the equilibrium strategy from period s on. Moreover, for every $k \in \mathcal{K}$ and $t \geq 0$, we have that $V_t^k \leq z = \max_k v_k$. This follows from the fact buyers do not find

¹²Formally, we show that if $\{\delta_n\}$ and $\{\lambda_n\}$ are, respectively, a sequence of discount factors and a sequence of matching probabilities such that $\lim_n \delta_n = \lim_n \lambda_n = 1$, then $\lim_n W(\Sigma_n) = W^*$ for any sequence $\{(\Sigma_n, \theta_n)\}$ such that (Σ_n, θ_n) is a PBE when $\delta = \delta_n$ and $\lambda = \lambda_n$.

it optimal to offer $p > z$ on the path of play if such an offer is accepted with positive probability. Both these facts about the payoffs V_t^k hold regardless of λ .

We proceed in three steps. First, we construct prices \hat{p}_1 to \hat{p}_K such that for all $k \in \mathcal{K}$ a type- k seller accepts an offer of \hat{p}_k in the first K periods. Then, we use these "reservation" prices to construct a lower bound to V^B . Finally, we use the lower bound to V^B to construct a lower bound to $W(\Sigma)$ that is independent of Σ and show that the lower bound converges to the first-best welfare as δ converges to one.

Step 1. Reservation Prices for Sellers

Fix $\varepsilon > 0$ and let $T \geq K$. For the argument that follows it suffices to set $T = K$. Letting T be arbitrary helps us when we discuss the extension of our argument to the case in which λ can assume any value in $(0, 1)$. Now, for each $k \in \mathcal{K}$, let

$$\hat{p}_k = c_k + \frac{V_1^k}{\delta^{T-2}} + \frac{\varepsilon}{4}.$$

We claim that a type- k seller accepts an offer of \hat{p}_k in the first $T - 1$ periods. Indeed, it is strictly optimal for a type- k seller to accept an offer of p in period t if $p - c_k > \delta V_{t+1}^k$. Now observe that if $t \in \{0, \dots, T - 1\}$, then $V_1^k \geq \delta^t V_{t+1}^k \geq \delta^{T-1} V_{t+1}^k$. Hence, $t \in \{0, \dots, T - 1\}$ implies that $\hat{p}_k - c_k \geq \delta V_{t+1}^k + \varepsilon/4 > \delta V_{t+1}^k$, and so a type- k seller accepts an offer of \hat{p}_k in period t .

Step 2. Lower Bound to Buyers' Payoff

Re-label the aggregate states so that \hat{p}_k is (weakly) increasing in k . Consider now the following alternative strategy $\hat{\sigma}^B$ for a buyer: offer \hat{p}_{t+1} if matched to a seller in period $t \in \{0, \dots, K - 1\}$ and offer \hat{p}_K if matched to a seller in period $t \geq K$. Denote by $u(\hat{\sigma}^B; (\Sigma, \theta))$ the payoff to a buyer who follows $\hat{\sigma}^B$ when all other agents behave according to Σ and the belief system is θ . Since $\lambda \approx 1$, the buyer transacts with probability essentially equal to one in at most $K \leq T$ periods and pays at most \hat{p}_k for the good when the aggregate state is k . Thus, a lower bound to the buyer's payoff in aggregate state k when he follows $\hat{\sigma}^B$ is $\delta^{T-1} v_k - \hat{p}_k - \varepsilon/4$, and so

$$\begin{aligned} u(\hat{\sigma}^B; (\Sigma, \theta)) &\geq \sum_{k=1}^K \pi_k (\delta^{T-1} v_k - \hat{p}_k) - \frac{\varepsilon}{4} \\ &= \sum_{k=1}^K \pi_k (\delta^{T-1} v_k - c_k) - \frac{1}{\delta^{T-2}} \sum_{k=1}^K \pi_k V_1^k - \frac{\varepsilon}{2}. \end{aligned}$$

Given that (Σ, θ) is an equilibrium, $V^B \geq u(\hat{\sigma}^B; (\Sigma, \theta))$, otherwise buyers would have a profitable deviation. Consequently,

$$\begin{aligned} V^B &\geq \sum_{k=1}^K \pi_k (v_k - c_k) - (1 - \delta^{T-1}) \sum_{k=1}^K \pi_k v_k - \frac{1}{\delta^{T-2}} \sum_{k=1}^K \pi_k V_1^k - \frac{\varepsilon}{2} \\ &\geq W^* - (1 - \delta^{T-1})z - \frac{1}{\delta^{T-1}} \sum_{k=1}^K \pi_k V^k - \frac{\varepsilon}{2}. \end{aligned}$$

where the second inequality follows from the fact that $V^k \leq z = \max_k v_k$ and $V^k = V_0^k \geq \delta V_1^k$ for every aggregate state k .

Step 3. Lower Bound to Welfare

Since preferences are quasi-linear,

$$W(\Sigma) = V^B + \sum_{k=1}^K \pi_k V^k.$$

Thus, from Step 2 and using again the fact that $V^k \leq z$ for all $k \in \mathcal{K}$, we have that

$$W(\Sigma) \geq W^* - (1 - \delta^{T-1}) \left(z + \frac{z}{\delta^{T-1}} \right) - \frac{\varepsilon}{2}.$$

Consequently, there exists $\bar{\delta} \in (0, 1)$ such that $W(\Sigma) \geq W^* - \varepsilon$ if $\delta \geq \bar{\delta}$. The desired result follows from the fact that ε is arbitrary.

Extending the above argument to the case in which λ is bounded away from one requires changing the strategy $\hat{\sigma}^B$ used to compute the lower bound to the buyers' equilibrium payoff to account for the fact that a buyer might not be matched to a seller in every period. Loosing speaking, when λ is bounded away from one, the strategy $\hat{\sigma}^B$ must be such that a buyer first attempts to trade with the type of seller with the lowest reservation price for sufficiently many periods, then attempts to trade with the type of seller with the second lowest reservation price for sufficiently many periods, and so on. This requires taking T in the above definition of reservation prices to be appropriately large. In the limit as δ converges to one, the delay in trading implied by the modified strategy $\hat{\sigma}^B$ converges to zero and one still obtains a lower bound to the equilibrium welfare that is independent of the equilibrium under play and converges to the first-best welfare. The details are in the Appendix. We conclude this section with some remarks about our results.

Role of Unrestricted Price Offers An important difference between our model and earlier models of dynamic trading with aggregate uncertainty is the assumption that buyers are not restricted in the prices they can offer. This assumption plays an important role in the proof of Theorem 1. In the proof, we make use of the fact that an option for a buyer is to offer the reservation prices of the different types of sellers in ascending order, thus extracting the residual surplus from sellers in a finite number of periods regardless of the aggregate state. With restricted price offers, buyers cannot fine-tune their offers to match the sellers' reservation values. This opens up the possibility of delay in trade for the following reason. A finite price grid allows agents to coordinate in an equilibrium with delay in trade by eliminating the possibility of deviations that could break the impasse between buyers and sellers.¹³

Role of Aggregate Uncertainty and Random Matching. Our efficiency result contrasts strongly with inefficiency results in dynamic decentralized markets with common-value uncertainty but uncorrelated asset values. In such environments, multiple types of seller co-exist in the market. As is well-known, the incentive that sellers with low valuation have to mimic the behavior of sellers with high valuation ensures that equilibria remain inefficient even as trading frictions vanish—a reduction in delay costs reduces the cost for the former type of seller to imitate the latter type of seller. In our environment, a single type of seller is present in the market at any point in time, ruling out the possibility just described.

Some of the driving forces present in our random-matching model with a continuum of agents are also not present in bargaining models in which a single seller dynamically meets a with a sequence of short-run buyers or the same buyer. In our setting, since rejected offers are not observable and there is a continuum of agents, a seller cannot influence the future behavior of buyers by rejecting an offer in an attempt to 'signal' that he is of a higher type.¹⁴ On the other hand, when there is a single seller in the market, his behavior can affect the future behavior of buyers. In this case, the only way to provide incentives for a low-valuation seller to trade at a low price is to have delay in trade with a high-valuation seller.

¹³This holds even without aggregate uncertainty. An example of an equilibrium with restricted price offers and no aggregate uncertainty which remains inefficient as trading frictions vanish is available from the authors upon request.

¹⁴Notice, however, that time on the market is observable. So, an option for sellers is to reject offers and stay longer in the market in the hopes of securing higher future offers.

Information Aggregation. A question that has attracted substantial interest is whether markets fully aggregate the information dispersed among agents. While the equilibrium of Section 3 aggregates information perfectly, information aggregation does not hold in general in our setting, as the following example of a pooling equilibrium shows.

Suppose that $K = 2$, $c_1 = 0$, $v_1 = 1$, $c_2 = 2$, and $v_2 = 3$, and consider the following symmetric assessment: (i) a buyer always offers c_2 ; (ii) a type-1 seller accepts an offer p if, and only if, $p \geq v_B = 2\delta\lambda/(1 - (1 - \lambda)\delta)$, while a type-2 seller accepts an offer p if, and only if, $p \geq 2$; and (iii) a buyer who has never made an offer or only made offers less than v_B assigns probability π_2 , the prior, to $k = 2$, while a buyer who has made an offer $p \geq v_B$ assigns probability one to $k = 2$. By construction, v_B is the continuation payoff to a type-1 seller who rejects an offer and stays in the market. In order to check that the assessment under consideration is a PBE, we need to show that offering c_2 is better than offering v_B , which sustains pooling. The payoff to a buyer from offering c_2 is $\pi_2 - (1 - \pi_2) = 2\pi_2 - 1$, while the payoff to a buyer from offering v_B is

$$(1 - \pi_2) \left(1 - \frac{2\delta\lambda}{1 - (1 - \lambda)\delta} \right) + \pi_2 \frac{\delta\lambda}{1 - (1 - \lambda)\delta}.$$

The first payoff is greater than the second if, and only if, $\pi_2 \geq 2/3$.

Finding necessary and sufficient conditions on primitives for every equilibrium to fully aggregate information is an interesting avenue for future research.

5 Robustness and Extensions

In this section we discuss robustness and extensions of our efficiency result.

Gains From Trade. Theorem 1 shows that the assumption of nonnegative gains from trade in every state is sufficient for welfare in all PBE to approach the first-best welfare as trading frictions disappear. The example below shows that this assumption is also necessary for limit efficiency.

Suppose that $K = 2$ and $v_1 < c_1 < c_2 < v_2$, so that gains from trade are negative when $k = 1$. In this case, $W^* = \pi_2(v_2 - c_2)$. Take a sequence $\{\delta_n\}$ of discount factors that converges to one and, for each $n \in \mathbb{N}$, let (Σ_n, θ_n) be a PBE when $\delta = \delta_n$. Assume towards a contradiction that

$W(\Sigma_n)$ converges to W^* . Then $\lim_n \mathbb{E}_{\xi_n} [\delta_n^T | k = 1] = 0$ and $\lim_n \mathbb{E}_{\xi_n} [\delta_n^T | k = 2] = 1$, where ξ_n is the probability distribution over the set of outcomes induced by Σ_n . Now let \mathbf{Q} be the first (random) period in which a buyer makes an offer $p \geq c_2$. Then $\lim_n \mathbb{E}_{\xi_n} [\delta_n^Q | k = 2] = 1$. It is possible to show that $\lim_n \mathbb{E}_{\xi_n} [\delta_n^T | k = 1] = 0$ and $\lim_n \mathbb{E}_{\xi_n} [\delta_n^Q | k = 2] = 1$ together imply that $\lim_n \mathbb{E}_{\xi_n} [\delta_n^Q | k = 1] = 1$. So, a seller in state 1 can secure a limit payoff of at least $c_2 - c_1 > 0$ by following the strategy in which he accepts an offer p if, and only if, $p \geq c_2$, a contradiction.

Bargaining Protocol. It is possible to extend Theorem 1 to the case in which in every buyer-seller pair the buyer makes a take-it-or-leave-it offer to the seller with positive probability. A sketch of the proof, which is similar to the proof of Theorem 1, is as follows. We again consider the limiting case in which λ is close enough to one that the probability that a buyer in the market is matched to a seller is essentially equal to one; the extension to the case in which λ is bounded away from one proceeds along the same lines as in the previous section. A lower bound to a buyer's payoff in any equilibrium is obtained when the buyer: (i) rejects any offer he receives from a seller; and (ii) offers the type- k seller's reservation price in the k th period in which he gets to make an offer. This strategy ensures that the buyer trades after making at most K offers. As trading frictions vanish, this lower bound on the buyers's equilibrium payoff converges to the first-best welfare net of the sellers' ex-ante equilibrium payoff, which establishes the desired result.¹⁵

Theorem 1 is not true when sellers have all the bargaining power, though. When only the informed agents can make offers, the signalling of private information through offers opens up the possibility of equilibria which remain inefficient even as trading frictions vanish.

Uninformed Sellers. We assume that all sellers are informed about the aggregate state. As it turns out, we cannot relax this assumption. When some sellers are uninformed about the aggregate state, offers to such sellers provide information about the aggregate state. In this case, it is possible to construct equilibria sustaining inefficient outcomes even as trading frictions vanish.¹⁶ The logic behind such inefficient equilibria is as follows. In equilibrium, uninformed sellers can interpret

¹⁵We can also extend our analysis to the case in which matched sellers meet with more than one buyer and buyers simultaneously make take-it-or-leave-it offers. Thus, imperfect competition is not crucial for our results.

¹⁶An example of such an equilibrium is available from the authors upon request.

intermediate offers as a signal of states in which the opportunity cost of trading is high, and so only be willing to accept high offers. Trade stops if buyers are not willing to make high offers.

Market Opacity. In our environment, the only information buyers have about the aggregate state is their history of trade. In most over-the-counter markets buyers have access to more information than their private histories, though. A natural concern is whether our efficiency result survives when we relax our assumption of market opacity. Indeed, the example of Section 3 is sensitive to this assumption, as buyers would have an incentive to deviate from their prescribed strategy if they had access to more information than their private histories.

As it turns out, our efficiency result extends to the case in which buyers observe private signals about the aggregate state. The signals can even be correlated, as when in every period all buyers in the market observe a public signal about the aggregate state. The key observation is that buyers can still compute the sellers' equilibrium payoffs in each period as a function of the aggregate state and use the same strategy described in the proof of Theorem 1. What is crucial is the continuum population assumption (or, more broadly, the anonymity assumption): since a buyer-seller pair will never meet again, the fact that the buyer has private information in a match is irrelevant for the seller's continuation payoff.

Matching Probabilities We can extend our analysis to the case in which $\lambda = 1$, and so in every period all buyers and sellers in the market are matched in pairs with probability one. In this case, the natural equilibrium concept to consider is sequential equilibrium, as it pins down beliefs, and thus payoffs, even when the mass of agents in the market is zero.¹⁷

Since the existence of a sequential equilibrium cannot be guaranteed when action sets are infinite, we make the additional assumption that buyers are restricted to make offers in a finite grid $P = \{p_0, p_1, \dots, p_M\}$ of prices, where p_i is strictly increasing in i , $p_0 < \min_k c_k$, and $p_M > \max_k c_k$.¹⁸ The assumption that $p_0 < \min_k c_k$ is natural, as it implies that buyers can make offers

¹⁷Indeed, first notice that if the pair (Σ, θ) is such that Σ has full support, then there is a positive mass of agents in the market in every period, in which case payoffs are well-defined after any history. Now observe that payoffs in a sequential equilibrium are the limits of payoffs when the pair (Σ, θ) is such that Σ has full support.

¹⁸A straightforward fixed-point argument shows that a sequential equilibrium exists when buyers are restricted to make offers in a finite price grid.

that are always rejected. The assumption that buyers can make offers that are greater than the highest cost of production is also natural; otherwise, it is trivial to generate inefficient equilibria. For any price grid P , let $\mathcal{E}(P) = \max_{0 \in \{1, \dots, M-1\}} |p_{i+1} - p_i|$ be the coarseness of P . Our efficiency result is obtained in the limit as $\mathcal{E}(P)$ converges to zero, and so the grid of possible price offers becomes arbitrarily fine. This limiting case approaches the case in which buyers are not restricted in the offers they can make to sellers.

Theorem 2. *Let $\{\delta_n\}$ be a sequence of discount factors such that $\lim_n \delta_n = 1$ and $\{P_n\}$ be any sequence of finite price grids such that $\min P_n < \min_k c_k$ and $\max P_n > \max_k v_k$ for all $n \in \mathbb{N}$ and $\lim_n \mathcal{E}(P_n) = 0$. For any sequence $\{(\Sigma_n, \theta_n)\}$ such that (Σ_n, θ_n) is a sequential equilibrium when $\delta = \delta_n$ and $P = P_n$, the sequence $\{W(\Sigma_n)\}$ converges to W^* .*

Notice that the sequence of price grids in the statement of Theorem 2 is independent of the sequence of discount factors. The proof of Theorem 2 is very similar to the proof of Theorem 1 and thus is omitted. Loosely speaking, when buyers are restricted to make offers in a finite set P , the reservation prices used to derive a lower bound to the buyers' equilibrium payoff need to be adjusted so that they are elements of P . As $\mathcal{E}(P)$ converges to zero, this adjustment becomes arbitrarily small and one obtains a lower bound to the buyers' equilibrium payoff that approaches the lower bound we obtain when buyers are not restricted in the offers they can make.

Matching Process. While standard in the literature on dynamic decentralized trading in the presence of common-value uncertainty, the assumptions that the initial mass of buyers and sellers are the same and matching is one-to-one simplify the analysis by implying that the outcome of matching is independent of the aggregate state. Two issues show up when this is not the case. First, the outcome of matching now provides information about the aggregate state. Second, matching probabilities for buyers change over time, affecting the payoff to buyers when they follow the alternative strategy described in the proof of Theorem 1.

The discussion about market opacity shows that the first issue discussed above is not an impediment to efficiency. While we do not have a general proof that our efficiency result extends to settings in which the matching probability for buyers changes over time, it is possible to show

that it extends to a number of different settings in which this is the case. One such setting is the generalization of our environment to the case in which the initial mass of sellers is different from the initial mass of buyers and there are two aggregate states, a fairly standard assumption in the literature. Another such setting is the case in which the initial mass of buyers is greater than the initial mass of sellers and matching is many-to-one, with the number of buyers a seller meets given by a Poisson random variable with expected value equal to the (state- and time-dependent) ratio of buyers to sellers in the market.¹⁹

Private Values We assume that asset values are perfectly correlated across buyers and sellers. It turns out that this assumption is not necessary for our results. As we now discuss, an efficiency result is possible when asset values have a common- and a private-value component.

The environment is the same as before, except that now buyers and sellers have idiosyncratic tastes. The set of possible buyer types is $\mathcal{D} = \{1, \dots, D\}$, with $D \geq 1$, while the set of possible seller types is $\mathcal{L} = \{1, \dots, L\}$, with $L \geq 1$. We denote a typical element of \mathcal{D} by d and a typical element of \mathcal{L} by ℓ . An agent's type is his private information and is independent of any other agent's type. The probability that a buyer is of type d is $\varphi_d \in [0, 1]$, while the probability that a seller is of type ℓ when the aggregate state is k is $\gamma_\ell^k \in [0, 1]$. The payoff to a type- d buyer from consuming the good in state k is $v_{k,d}$, while the cost to a type- ℓ seller of producing the good in state k is $c_{k,\ell}$. As in the setting of Theorem 1, we assume nonnegative gains from trade in every state, in which case the first-best welfare is given by

$$W^* = \sum_{k=1}^K \sum_{\ell=1}^L \sum_{d=1}^D \pi_k \gamma_\ell^k \varphi_d (v_{k,d} - c_{k,\ell}).$$

Assumption 2. $v_{k,d} - c_{k,\ell} \geq 0$ for all $(k, d, \ell) \in \mathcal{K} \times \mathcal{D} \times \mathcal{L}$.

It is possible to extend Theorem 1 to the case described here. While the idea behind this more general efficiency result is similar to the idea behind Theorem 1, the proof needs to be adapted to take into account the fact that when there is more than one type of seller in each aggregate state, these different types of seller can exit the market at different rates in equilibrium.

¹⁹Details are available upon request.

6 Concluding Remarks

Most of the literature that studies the impact of asymmetric information on market efficiency in decentralized markets with common-value uncertainty focuses on the case of independent asset values. In this paper, we show that allowing correlated asset values—a more realistic assumption in many relevant markets—can lead to starkly different results. While our assumption of perfectly correlated asset values is strong, it constitutes a useful first step in relaxing the assumption of independent asset values. The question of how asymmetric information affects market efficiency in decentralized markets with common-value uncertainty and *imperfectly* correlated asset values remains an important open question.

References

- Asriyan, V., Fuchs, W., and B. Green. 2017. Information Spillovers in Asset Markets with Correlated Values. *American Economic Review*, 107(7), 2007–2040.
- Asriyan, V., Fuchs, W., and B. Green. 2017. Information Aggregation in Dynamic Markets with Adverse Selection. Unpublished manuscript.
- Blouin, M. R. 2003. Equilibrium in a Decentralized Market with Adverse Selection. *Economic Theory*, 22, 245–262.
- Blouin, M. R., and R. Serrano. 2001. A Decentralized Market with Common Values Uncertainty: Non-Steady States. *Review of Economic Studies*, 68, 323–346.
- Camargo, B., and B. Lester. 2014. Trading Dynamics in Decentralized Markets with Adverse Selection. *Journal of Economic Theort*, 153, 534–568.
- Daley, B., and B. Green. 2012. Waiting for News in the Market for Lemons. *Econometrica*, 80(4), 1433–1504.
- Deneckere, R., and M.-Y. Liang. 2006. Bargaining with Interdependent Values. *Econometrica*, 74(5), 1309–1364.

- Fuchs, W., Öry, A., and A. Skrzypacz. 2016. Transparency and Distressed Sales Under Assymmetric Information. *Theoretical Economics*, 11(3), 1103–1144.
- Gerardi, D., Hörner, J., and L. Maestri. 2014. The Role of Commitment in Bilateral Trade. *Journal of Economic Theory*, 154, 578–603
- Golosov, M., Lorenzoni, G., and A. Tsyvinski. 2014. Decentralized Trading with Private Information. *Econometrica*, 82, 1055–1091.
- Hörner, J., and N. Vieille. 2009. Public vs. Private Offers in the Market for Lemons. *Econometrica*, 77(1), 29–69.
- Kaya, A., and K. Kim. 2015. Trading Dynamics with Private Buyer Signals in the Market for Lemons. Unpublished Manuscript.
- Kurlat, P., and J. Stroebe. 2015. Testing for Information Asymmetries in Real Estate Markets. *Review of Financial Studies*, 28(8), 2429–2461.
- Kim, K. 2017. Information About Seller’s Past Behavior in the Market for Lemons. *Journal of Economic Theory*, 169, 365–399.
- Lauerermann, S., and A. Wolinsky. 2016. Search with Adverse Selection. *Econometrica*, 84(1), 243–315.
- Serrano, R., and O. Yosha. 1993. Information Revelation in a Market with Pairwise Meetings: The One Sided Information Case. *Economic Theory*, 3(3), 481–499.
- Serrano, R., and O. Yosha. 1995. Decentralized Markets with Pairwise Meetings: Recent Developments. *Journal of International and Comparative Economics*, 4(4), 223–241.
- Serrano, R. 2002. Decentralized Information and the Walrasian Outcome: A Pairwise Meetings Market with Private Values. *Journal of Mathematical Economics*, 38, 65–89.
- Wolinsky, A. 1990. Information Revelation in a Market with Pairwise Meetings. *Econometrica*, 58, 1–23.

7 Appendix

In this Appendix we first extend the equilibrium construction of Section 3 to the case of any finite number of aggregate states. We then provide the proof of Theorem 1 in the general case in which λ can assume any value in $(0, 1)$.

Equilibrium of Section 3

Let $\mathcal{C} = \{\bar{c}_1, \dots, \bar{c}_L\}$, with $\bar{c}_1 < \dots < \bar{c}_L$ and $L \leq K$, be the set of possible production costs for sellers and define $f : \mathcal{C} \rightrightarrows \mathcal{K}$ to be the correspondence such that $f(\bar{c}_\ell) = \{k \in \mathcal{K} : c_k = \bar{c}_\ell\}$. Now consider the symmetric assessment (Σ, θ) in which the sellers' strategy, the buyers' strategy, and the buyers' belief system are as follows:

Sellers' Strategy σ^S : A type- k seller in a match accepts an offer p if, and only if, $p \geq c_k$.

Buyers' Strategy σ^B : For any history h^t for a buyer, let: (i) $m(h^t) \in \mathbb{N}_+$ be the number of times the buyer was matched in the market so far; (ii) $\bar{p}(h^t)$ be the largest offer the buyer has made so far, where $\bar{p}(h^t) = -\infty$ if $m(h^t) = 0$; and (iii) $\mathcal{C}(\bar{p}(h^t)) = \{c \in \mathcal{C} : c > \bar{p}(h^t)\}$. Notice that $\mathcal{C}(\bar{p}(h^t)) = \mathcal{C}$ if $m(h^t) = 0$. A buyer's behavior after a history h^t is as follows. If $\mathcal{C}(\bar{p}(h^t)) \neq \emptyset$, then the buyer offers the smallest element of $\mathcal{C}(\bar{p}(h^t))$ if he is matched in the market. If, instead, $\mathcal{C}(\bar{p}(h^t)) = \emptyset$, then the buyer offers \bar{c}_L if he is matched in the market.

Buyers' Belief System μ^B : If $m(h^t) = 0$, then $\mu(h^t)$ assigns probability π_k to $k \in \mathcal{K}$. Now suppose that $m(h^t) > 0$. If $\mathcal{C}(\bar{p}(h^t)) \neq \emptyset$, then $\mathcal{C}(\bar{p}(h^t)) = \{\bar{c}_j, \dots, \bar{c}_L\}$ for some $j \in \{1, \dots, L\}$. In this case, $\mu(h^t)$ assigns probability 0 to every $k \in \bigcup_{s=1}^{j-1} f(\bar{c}_s)$ and assigns probability

$$\frac{\pi_k}{\sum_{\ell \in \bigcup_{s=j}^L f(\bar{c}_s)} \pi_\ell}$$

to every other aggregate state. On the other hand, if $\mathcal{C}(\bar{p}(h^t)) = \emptyset$, then $\mu(h^t)$ assigns probability 0 to every $k \in \bigcup_{s=1}^{L-1} f(\bar{c}_s)$ and assigns probability

$$\frac{\pi_k}{\sum_{\ell \in f(\bar{c}_L)} \pi_\ell}$$

to the remaining aggregate states.

Notice that the strategy profile $\Sigma = (\sigma^B, \sigma^S)$ is such that on the path of play a buyer offers \bar{c}_k , $k \in \{1, \dots, L\}$, the k th time that he is matched to a seller and trades after making at most L offers.

So, when play is given by Σ , welfare approaches the first-best welfare as trading frictions vanish. The assessment (Σ, Θ) reduces to the assessment in the example of Section 3 when there are two aggregate states (and $c_1 \neq c_2$).

We claim that there exists $\delta^* \in (0, 1)$ such that (Σ, θ) is a PBE if $\delta > \delta^*$. It is easy to see that μ is consistent with Bayes' rule on the path of play. The same argument as in the case of two aggregate states shows that σ^S is sequentially rational regardless of δ : with probability one, a type- k seller who rejects an offer in a given period receives no offer greater than c_k in the future. In what follows we show that σ^B is sequentially rational δ when sufficiently close to one.

Consider a buyer with history h^t . From above, there exists $j \in \{1, \dots, L\}$ such that $\mu(h^t)$ assigns probability

$$\tilde{\pi}_k = \frac{\pi_k}{\sum_{\ell \in \bigcup_{s=j}^L f(\bar{c}_s)} \pi_\ell}$$

to every aggregate state $k \in \bigcup_{s=j}^L f(\bar{c}_s)$ and assigns probability 0 to every other aggregate state. If $j = L$, then the buyer does not have a profitable deviation as he believes that $k \in f(\bar{c}_L)$ and expects that any seller accepts an offer of \bar{c}_L . Suppose then that $j < L$ and let $\tau \geq L - 1$ be the random time at which the buyer is matched for the L th time in the market. If the buyer follows the equilibrium strategy, then his expected payoff is bounded below by

$$u_j = \sum_{\ell \in \bigcup_{k=j}^L f(\bar{c}_k)} \mathbb{E}[\delta^{\tau-1}] \tilde{\pi}_\ell (v_\ell - c_\ell).$$

Clearly, the most profitable deviation for the buyer consists in offering a price $\bar{c}_k \in \{\bar{c}_{j+1}, \dots, \bar{c}_L\}$ to trade earlier. The payoff from this deviation is bounded above by

$$u'_j = \underbrace{\sum_{\ell \in \bigcup_{k=j}^L f(\bar{c}_k)} \tilde{\pi}_\ell (v_\ell - c_\ell)}_{\bar{u}} - \underbrace{\mathbb{E}[\delta^{\tau-1}] (\bar{c}_{j+1} - \bar{c}_j) \sum_{\ell \in f(\bar{c}_j)} \tilde{\pi}_\ell}_{\varepsilon}.$$

The payoff \bar{u} is the agent's expected payoff from trading immediately at the correct price, while ε is a lower bound to the agent's expected loss from trading at a higher price; with probability at least $\sum_{\ell \in f(\bar{c}_j)} \tilde{\pi}_\ell$ there is an aggregate state in which the buyer purchases the good at a price at least $(\bar{c}_{j+1} - \bar{c}_j) > 0$ greater than \bar{c}_j . Since $\lim_{\delta \rightarrow 1} \mathbb{E}[\delta^{\tau-1}] = 1$ by the dominated convergence theorem, there exists $\delta_j \in (0, 1)$ such that $\delta > \delta_j$ implies that $u_j > u'_j$.

Letting $\delta^* = \max\{\delta_1, \dots, \delta_{L-1}\}$, we can then conclude that σ^B is sequentially rational whenever $\delta > \delta^*$, and so (Σ, θ) is a PBE if $\delta > \delta^*$.

Proof of Theorem 1

We show that for all $\varepsilon \in (0, z)$, there exists $\bar{\delta} \in (0, 1)$ such that if $\delta > \bar{\delta}$, then $W(\Sigma) > W^* - \varepsilon$ for every assessment (Σ, θ) which is a PBE when the discount factor is δ . Recall that $z = \max_k v_k$.

Fix $\varepsilon \in (0, z)$ and let (Σ, θ) be a PBE for some discount factor δ . We know from the main text that buyers obtain the same payoff V^B and each type k of seller obtains the same payoff V^k . Moreover, if we let V_t^k denote a type- k seller's payoff in period t , then this payoff does not depend on a seller's private history. Finally, the payoffs V_t^k are such that: (i) $V_t^k \geq \delta^{s-t} V_s^k$ for every $k \in \mathcal{K}$ and $s > t \geq 0$; and (ii) $V_t^k \leq z$ for every $k \in \mathcal{K}$ and $t \geq 0$.

Step 1. Reservation Prices for Sellers

We proceed as in the main text and first identify for each $k \in \mathcal{K}$ a set of offers that, in equilibrium, a type- k seller accepts with probability one. Let $\kappa = \varepsilon/16z > 0$ and define $T(\kappa)$ as the smallest positive integer such that $(1 - \lambda)^{T(\kappa)} < \kappa$.

Lemma 1. *Consider the equilibrium (Σ, θ) . For every $k \in \mathcal{K}$ and $t \geq 0$, let*

$$\underline{p}_{k,t} = c_k + \frac{V_1^k}{\delta^{t-1}}.$$

If a type- k seller receives an offer $p > \underline{p}_{k,t}$ in period t , then he accepts it with probability one.

Proof. Consider a type- k seller in period $t \geq 0$. He accepts an offer of p if $p - c_k > \delta V_{t+1}^k$. The desired result follows from the fact that $\underline{p}_{k,t} - c_k = V_1^k / \delta^{t-1} \geq \delta V_{t+1}^k$, where the inequality follows from the fact that $V_t^k \geq \delta^t V_{t+1}^k$ for all $k \in \mathcal{K}$ and $t \geq 0$. \square

Now, for each $k \in \mathcal{K}$, define \hat{p}_k to be such that

$$\hat{p}_k = c_k + \frac{V_1^k}{\delta^{T(\kappa)K-2}} + \frac{\varepsilon}{4}.$$

Since, by construction, $\hat{p}_k > \underline{p}_{k,t}$ for all $t \in \{0, \dots, T(\kappa)K - 1\}$, Lemma 1 implies that a type- k seller accepts an offer of \hat{p}_k in any period $t \in \{0, \dots, T(\kappa)K - 1\}$.

Step 2. Lower Bound to Buyers' Payoff

As in the main text, we now use the prices \hat{p}_1 to \hat{p}_K to derive a lower bound to the buyers' equilibrium payoff. Re-label the aggregate states so that \hat{p}_k is (weakly) increasing in k . Consider

the following alternative strategy $\hat{\sigma}^B$ for a buyer: offer \hat{p}_k if matched to a seller in periods $t \in \{T(\kappa)(k-1), \dots, T(\kappa)k-1\}$ and offer \hat{p}_K if matched to a seller in any period $t \geq T(\kappa)K$. Denote by $u(\hat{\sigma}^B; (\Sigma, \theta))$ the payoff the buyer obtains when all other agents behave according to the strategy profile Σ and the belief system is θ . Notice that $V^B \geq u(\hat{\sigma}^B; (\Sigma, \theta))$, otherwise the buyer would have a profitable deviation.

We obtain a lower bound for $u(\hat{\sigma}^B; (\Sigma, \theta))$, and thus V^B , as follows. Suppose the aggregate state is k . There are two mutually exclusive and exhaustive events to consider: the buyer transacts in period $t < T(\kappa)(k-1)$ or the buyer is still in the market in period $T(\kappa)(k-1)$. In the first event, the buyer's expected payoff is bounded below by $\delta^{T(\kappa)K-1}v_k - \hat{p}_k$; this is because \hat{p}_k is increasing in k and $v_k \geq 0$. Consider now the second event. Either the buyer is matched with a seller in some period $t \in \{T(\kappa)(k-1), \dots, T(\kappa)k-1\}$ and obtains a payoff of at least $\delta^{T(\kappa)K-1}v_k - \hat{p}_k$, or the buyer is not matched with a seller in any of these periods and obtains a payoff of at least $-\hat{p}_K$. Given that $(1-\lambda)^{T(\kappa)} < \kappa$, the probability the buyer does not meet a seller in some period $t \in \{T(\kappa)(k-1), \dots, T(\kappa)k-1\}$ is at most κ . So, the buyer's expected payoff in the second event is bounded below by $(1-\kappa)(\delta^{T(\kappa)K-1}v_k - \hat{p}_k) - \kappa\hat{p}_K$. Given that the lower bound to the buyers' expected payoff is lower in the second event, we then have that

$$u(\hat{\sigma}^B; (\Sigma, \theta)) \geq (1-\kappa) \sum_{k=1}^K \pi_k (\delta^{T(\kappa)K-1}v_k - \hat{p}_k) - \kappa\hat{p}_K.$$

Now observe that $\max_k \{c_k, V_1^k\} \leq z$ and $\varepsilon \in (0, z)$ implies that

$$\kappa\hat{p}_K \leq \kappa \left(c_k + V_1^K + \varepsilon + \frac{(1-\delta^{T(\kappa)K-2})V_1^K}{\delta^{T(\kappa)K-2}} \right) \leq 3\kappa z + z \left(\frac{1-\delta^{T(\kappa)K-2}}{\delta^{T(\kappa)K-2}} \right). \quad (1)$$

Moreover, we also have that

$$\begin{aligned} (1-\kappa) \sum_{k=1}^K \pi_k (\delta^{T(\kappa)K-1}v_k - \hat{p}_k) &= (1-\kappa) \sum_{k=1}^K \pi_k \left(\delta^{T(\kappa)K-1}v_k - c_k - \frac{\delta V_1^k}{\delta^{T(\kappa)K-1}} - \frac{\varepsilon}{4} \right) \\ &= (1-\kappa) \left[\sum_{k=1}^K \pi_k (v_k - c_k) - (1-\delta^{T(\kappa)K-1}) \sum_{k=1}^K \pi_k v_k - \frac{\varepsilon}{4} - \sum_{k=1}^K \pi_k \frac{\delta V_1^k}{\delta^{T(\kappa)K-1}} \right] \\ &\geq \sum_{k=1}^K \pi_k (v_k - c_k) - \kappa \sum_{k=1}^K \pi_k v_k - (1-\delta^{T(\kappa)K-1}) \sum_{k=1}^K \pi_k v_k - \frac{\varepsilon}{4} - \sum_{k=1}^K \pi_k \frac{\delta V_1^k}{\delta^{T(\kappa)K-1}} \\ &\geq \sum_{k=1}^K \pi_k (v_k - c_k) - \kappa z - (1-\delta^{T(\kappa)K-1}) z - \frac{\varepsilon}{4} - z \left(\frac{1-\delta^{T(\kappa)K-1}}{\delta^{T(\kappa)K-1}} \right) - \sum_{k=1}^K \pi_k \delta V_1^k; \quad (2) \end{aligned}$$

the first inequality follows from the fact that c_k , v_k , and V_1^k are non-negative for all $k \in \mathcal{K}$, while the second inequality follows from the fact that $\sum_{k=1}^K \pi_k v_k \leq z$ and $\sum_{k=1}^K \pi_k \delta V_1^k \leq z$.

Using inequalities (1) and (2) and the facts that $4\kappa z = \varepsilon/4$ and $(1 - \delta^t)/\delta^t$ is increasing in t , we can then conclude that

$$V^B \geq W^* - \sum_{k=1}^K \pi_k \delta V_1^k - (1 - \delta^{T(\kappa)K-1}) z - \frac{\varepsilon}{2} - 2z \left(\frac{1 - \delta^{T(\kappa)K-1}}{\delta^{T(\kappa)K-1}} \right).$$

Step 3. Lower Bound to Welfare

We conclude by using the above lower bound to V^B to obtain a lower bound to welfare. Since $W(\Sigma) = V^B + \sum_{k=1}^K \pi_k V^k$ and, for all $k \in \mathcal{K}$, $V^k = V_0^k \geq \delta V_1^k$, we have

$$W(\Sigma) \geq W^* - (1 - \delta^{T(\kappa)K-1}) z - \frac{\varepsilon}{2} - 2z \left(\frac{1 - \delta^{T(\kappa)K-1}}{\delta^{T(\kappa)K-1}} \right).$$

Taking $\bar{\delta} \in (0, 1)$ such that $\delta > \bar{\delta}$ implies that

$$(1 - \delta^{T(\kappa)K-1}) z + 2z \left(\frac{1 - \delta^{T(\kappa)K-1}}{\delta^{T(\kappa)K-1}} \right) < \frac{\varepsilon}{2},$$

we can then conclude that $W(\sigma, \mu) > W^* - \varepsilon$ whenever $\delta > \bar{\delta}$. The desired result follows from the fact that ε was arbitrary.