Efficiency in Decentralized Markets with Aggregate Uncertainty

Braz Camargo* Dino Gerardi† Lucas Maestri‡

June 2016

Abstract

We study efficiency in decentralized markets with aggregate uncertainty and one-sided private information. There is a continuum of mass one of uninformed buyers and a continuum of mass one of informed sellers. 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.

*Corresponding author. Sao Paulo School of Economics - FGV. E-mail: braz.camargo@fgv.br. The author gratefully acknowledges financial support from CNPq.

†Collegio Carlo Alberto, Università di Torino. E-mail: dino.gerardi@carloalberto.org. The author gratefully acknowledges financial support from the European Research Council (Consolidator Grant INFOASYMMARKETS).

‡FGV/EPGE - Escola Brasileira de Economia e Finanças. E-mail: maestri.lucas@gmail.com. The author gratefully acknowledges financial support from CNPq.

1 Introduction

Market efficiency is a central concern in economics. In idealized centralized markets, where information is perfect, the First Welfare Theorem shows that market outcomes are efficient. However, several relevant markets, including markets for many financial securities, function differently. First, trade is decentralized: rather than trade taking place at a single price that clears the market, buyers and sellers privately negotiate the terms of trade. Second, information is asymmetric: agents have access to different information about important features of the environment. In this paper, we study a large decentralized market with aggregate uncertainty in which one side of the market knows the aggregate state but the other side does not, and ask whether the presence of one-sided private information about the returns from trade is, by itself, an impediment to market efficiency.

The environment we consider is as follows. There is a mass one of buyers and a mass one of sellers. The payoffs from trade depend on an aggregate state, which only the sellers know. There is a finite number of such states and gains from trade are nonnegative in all states. Time is discrete and in every period the buyers and sellers in the market are randomly and anonymously matched in pairs. In each 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 the agents exit the market. Otherwise, the agents remain in the market.

In our environment there are two frictions that can prevent agents from realizing gains from trade. First, there is the usual trading friction in dynamic matching and bargaining environments, namely, a real time between two consecutive trading opportunities. Second, and foremost, sellers have private information about the aggregate state, which they can use to extract more favorable terms of trade from buyers.

Our main result is that one-sided private information alone is not enough to prevent market efficiency. We show that as trading frictions vanish, i.e., as the real time between two consecutive trading opportunities converges to zero, welfare in *any* equilibrium approaches the first best welfare, which is the welfare one obtains when buyers also know the aggregate state.

There are a number of papers in the dynamic matching and bargaining literature that study how asymmetric information affects market efficiency. Most consider environments with private values (see, e.g., Wolinsky 1988, De Fraja and Sákovics 2001, Serrano 2002, Satterthwaite and Shneyerov 2007, and Lauermann 2013). Moreover, of the few papers on dynamic matching and bargaining that study environments with common values, see, e.g., Blouin and Serrano (2001), Blouin (2003), Camargo and Lester (2014), and Moreno and Wooders (2016), only Blouin and Serrano (2001) considers aggregate uncertainty.

The environment we consider is similar to the environment in Blouin and Serrano (2001), but with a few important differences. First, we mainly focus on the case of one-sided private information. We show that this is the only case in which (limit) efficiency is guaranteed. Second, we allow for any finite number of aggregate states, instead of just two, and place no restrictions on the payoffs from trade except that gains from trade are nonnegative in every state. Finally, and crucially, we depart from Blouin and Serrano (2001) in the bargaining protocol. These authors consider a stylized bargaining game which amounts to restricting the set of prices at which trade can take place, and show that both in the one-sided and the two-sided private information case the market outcomes remain inefficient even when discounting vanishes. Our analysis thus shows that restricting prices can have a critical impact on market efficiency.

Our limit efficiency result contrasts strongly with (limit) inefficiency results in dynamic decentralized markets with adverse selection but no aggregate uncertainty (Blouin 2003, Moreno and Wooders 2016), and also with inefficiency results in the literature on bargaining with common value uncertainty, see, e.g., Deneckere and Liang (2006), Hörner and Vieille (2009), and Gerardi, Hörner, and Maestri (2014). We discuss the reasons for this difference, as well as other aspects of our analysis, in our concluding remarks.

2 The Environment

Time is discrete and indexed by $t \in \{0, 1, \ldots\}$. There is an equal mass of buyers and sellers, all with the same discount factor $\delta \in (0, 1)$. 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 $\Theta = \{1, \ldots, K\}$ and the probability that the state is $k \in \Theta$ is $\pi_k > 0$. The sellers know the state, but the 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* $k \in \Theta$.

Assumption 1 is fairly weak. In particular, single-crossing of preferences, i.e., $v_k - c_k$ increasing in k, is not necessary for our results. Moreover, as we show in the next section, this assumption cannot be relaxed. Assumption 1 implies that the first best welfare is

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

Trade takes place as follows. In each period $t \geq 0$, the buyers and sellers in the market are randomly and anonymously matched in pairs, and the buyer makes a take-it-or-leave-it offer to the seller. If the seller accepts the offer, then trade occurs and both agents exit the market. Otherwise, the match is dissolved and the agents remain in the market. In order to ensure existence of equilibria, we assume that buyers are restricted to make offers in a set $P = \{p_0, p_1, \ldots, p_M\}$, where $p_i < p_{i+1}$ for all $i \in \{1, \ldots, M\}$, $p_0 < \min_k c_k$, and $p_M > \max_k c_k$.[1] In what follows, we let $\mathcal{C}(P) = \max_{0 \in \{1,\ldots,M-1\}} |p_{i+1} - p_i|$ be the coarseness of P. Our efficiency result is obtained in the limit when $\mathcal{C}(P)$ converges to zero, so that the grid of price offers becomes arbitrarily fine.[2]

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.[3] A behavior strategy for a buyer is a sequence $\sigma^B = \{\sigma_t^B\}$, where $\sigma_t^B : \mathcal{H}_t \to \Delta(P)$ and $\sigma_t^B(h^t)$ is the (random) price offer that the buyer makes in period t when his private history is h^t. A behavior strategy for a seller is a sequence $\sigma^S = \{\sigma_t^S\}$, where $\sigma_t^S : \mathcal{H}_t \times \Theta \times P \to [0,1]$ and $\sigma_t^S(h^t, k, p)$ is the probability that the seller accepts an offer

[1] The assumption that $p_0 < \min_k c_k$ is natural, as it implies that buyers can make offers that are 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. Notice that a version of Diamond's Paradox holds in our setting, so in equilibrium sellers accept any offer greater than $\max_k c_k$.

[2] In the working paper version of our paper we show that if the set P is finite, then even when there is no asymmetry of information between buyers and sellers there can exist equilibria which remain inefficient as δ converges to one. Serrano and Yosha (1995) obtain the same result in a stationary version of our environment.

[3] A private history for an agent in the market in period t is a sequence $h^t = (p_1, \ldots, p_{t-1})$ of price offers. If the agent is a buyer, then h^t is the sequence of price offers that the agent made and were rejected. If the agent is a seller, then h^t is the sequence of price offers that the agent received and rejected.

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 $\mu^B = \{\mu_t^B\}$, where $\mu_t^B : \mathcal{H}_t \to \Delta(\Theta)$ and $\mu_t^B(h^t)$ is the buyer's (posterior) belief about the state in period t when his private history is h^t. We let σ and μ denote, respectively, a strategy profile and a profile of belief systems.

We consider pairs (σ, μ) which constitute a sequential equilibrium. This is the natural equilibrium concept in our environment, as in any sequential equilibrium the payoffs to agents are well-defined even if there is a zero mass of agents in the market. In particular, the payoff to an agent is well-defined if aggregate behavior is such that the market clears but the agent behaves in a such a way that he does not trade.[4] The existence of a sequential equilibrium (with a finite set of possible price offers) follows from a standard argument.[5]

3 Market Efficiency

In this section we show our main result, namely, that in any equilibrium welfare approaches the first best welfare as δ converges to one and $\mathcal{C}(P)$ converges to zero.

We begin with some definitions. An outcome for a given seller is a triple $(k, \mathbf{T}, \mathbf{p})$, where $k \in \Theta$ is the state, $\mathbf{T} \in \mathbb{Z}_+ \cup \{\infty\}$ is the time at which trade takes place, and $\mathbf{p} \in P$ is the price at which trade takes place; $\mathbf{T} = \infty$ corresponds to the event in which trade does not take place. Denote the set of all possible outcomes by $\mathcal{O}$. For any equilibrium (σ, μ), the probability distribution over the set of states and the strategy profile σ uniquely determine a probability distribution ξ over $\mathcal{O}$.[6] Let $\mathbb{E}_\xi$ denote the expectation with respect to ξ. Welfare in the equilibrium (σ, μ) is

$$W(\sigma, \mu) = \sum_{k=1}^{K} \pi_k \mathbb{E}_\xi \left[\delta^{\mathbf{T}}(v_k - c_k)|k\right].$$

Clearly, $W(\sigma, \mu)$ is bounded above by W^*.

We can now state and prove our main result. For this, let $\mathcal{P}$ be the set of all non-empty finite set P of prices such that $\min_{p \in P} p < \min_k c_k$ and $\max_{p \in P} p > \max_k c_k$.

[4] Indeed, first notice that if the pair (σ, μ) 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, μ) is such that σ has full support.

[5] For completeness, we provide a sketch of the argument in a Supplementary Appendix.

[6] When $\mathbf{T} = \infty$, the transaction price is undetermined. We adopt the convention that $\mathbf{p} = p_0$ in this case.

Theorem 1. *Let $\{\delta_n\}$ be a sequence of discount factors such that $\lim_n \delta_n = 1$ and $\{P_n\}$ be a sequence in $\mathcal{P}$ such that $\lim_n \mathcal{C}(P_n) = 0$. For any sequence $\{(\sigma_n, \mu_n)\}$ of equilibria such that (σ_n, μ_n) is an equilibrium when $\delta = \delta_n$ and $P = P_n$, the sequence $\{W(\sigma_n, \mu_n)\}$ converges to W^*.*

Proof. Fix the grid $P = \{p_0, p_1, \ldots, p_M\}$ of price offers with $p_0 < \cdots < p_M$ and the discount factor δ, and let (σ, μ) be a sequential equilibrium. First notice that even though (σ, μ) need not be symmetric, all the agents on a given side of the market obtain the same equilibrium payoff. In fact, since there is a continuum of agents, buyer (seller) i can obtain the same payoff as buyer (seller) j by mimicking j's behavior. We write V^B for the buyers' ex-ante (equilibrium) payoff and V^k for the sellers' payoff in state k. When the state is k, we refer to a seller as a type-k seller and write V_t^k for his payoff in period t. Since there is a continuum of agents and sellers know the state, V_t^k does not depend on the private history of a seller, only on the period t.

The following result establishes some useful properties of the payoffs V_t^k. Let $z = \max_k v_k$.

Lemma 1. *The payoffs V_t^k have the following properties:*
(i) For every $k \in \Theta$ and $s > t \geq 0$, we have $V_t^k \geq \delta^{s-t} V_s^k$;
(ii) For every $k \in \Theta$ and $t \geq 0$, we have $V_t^k \leq \min\{p_M - c_k, z\}$.

Proof of Lemma 1. Let $s > t \geq 0$. The first part follows immediately from the fact that a strategy for a seller in period t is to reject all offers between periods t and $s - 1$, and then follow the equilibrium behavior starting from period s.

We now prove the second part. Since $p_M > \max_k c_k$ and $\min_k c_k \geq 0$, a type-k seller accepts the offer p_M with probability one in equilibrium. Hence, $V_t^k \leq p_M - c_k$. Now observe that since P is finite, there exist $p^* \in P$ and a history h^t such that a buyer offers p^* with positive probability after h^t and no buyer makes an offer greater than p^* after any other history. Clearly, the seller who receives the offer p^* accepts it with probability one. Then $p^* \leq z$, otherwise a buyer who offers p^* obtains a negative payoff, in which case he would have a profitable deviation. Thus, $V_t^k \leq z$. $\square$

We divide the argument in several steps.

Step 1. First, for each $k \in \Theta$ we construct an offer $\widehat{p}_k$ that is feasible, i.e., belongs to P, and is accepted with probability one by a type-k seller in every period $t \in \{0, \ldots, K - 1\}$. We start with the following auxiliary result.

Lemma 2. *Consider the equilibrium* (σ, μ)*. For every* $k \in \Theta$ *and* $t \in \{0, \ldots, K-1\}$*, let*

$$\underline{p}_{k,t} := \min\left\{c_k + \frac{V_1^k}{\delta^{t-1}}, (1-\delta)c_k + \delta p_M\right\}.$$

If a type-k *seller receives an offer* $p > \underline{p}_{k,t}$ *in period* $t \in \{0, \ldots, K-1\}$*, then he accepts it with probability one.*

Proof of Lemma 2. Consider first the case $(1-\delta)c_k + \delta p_M \leq c_k + V_1^k/\delta^{t-1}$. If in period t a type-k seller accepts an offer $p > (1-\delta)c_k + \delta p_M$, then his payoff is strictly larger than $\delta(p_M - c_k)$. If instead he rejects it, then his payoff is equal to δV_{t+1}^k. The result follows from Lemma 1(*ii*).

Consider now the case $(1-\delta)c_k + \delta p_M > c_k + V_1^k/\delta^{t-1}$ and let $t \in \{0, \ldots, K-1\}$. It is strictly more profitable for a type-k seller to accept an offer p than to reject it if $p - c_k > \delta V_{t+1}^k$. From Lemma 1(*i*) we have $V_1^k/\delta^t \geq V_{t+1}^k$; notice that $V_1^k/\delta^t = V_{t+1}^k$ when $t = 0$. So, $p > \underline{p}_{k,t}$ implies

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

which concludes the proof of the lemma. □

We now use Lemma 2 to construct the desired prices $\widehat{p}_k$. For each $k \in \Theta$, let $\widehat{p}_k$ be the smallest element of P that is strictly greater than $c_k + V_1^k/\delta^{K-2}$ if $p_M > c_k + V_1^k/\delta^{K-2}$ and let $\widehat{p}_k = p_M$ otherwise. Notice that if $p_M > c_k + V_1^k/\delta^{K-2}$, then

$$\widehat{p}_k > c_k + \frac{V_1^k}{\delta^{K-2}} \geq c_k + \frac{V_1^k}{\delta^{t-1}} \tag{1}$$

for every $t \in \{0, \ldots, K-1\}$. On the other hand, if $p_M \leq c_k + V_1^k/\delta^{K-2}$, then

$$\widehat{p}_k = p_M > (1-\delta)c_k + \delta p_M \tag{2}$$

since $p_M > \max_k c_k$. Lemma 2 and inequalities (1) and (2) imply the following corollary.

Corollary 1. *Consider the equilibrium* (σ, μ)*. If a type-*k *seller receives the feasible offer* $p = \widehat{p}_k$ *in period* $t \in \{0, \ldots, K-1\}$*, then he accepts it with probability one.*

To finish the first step, we observe that

$$\widehat{p}_k \leq \left(c_k + \frac{V_1^k}{\delta^{K-2}}\right) + \mathcal{C}(P). \tag{3}$$

for every $k \in \Theta$; recall that $\mathcal{C}(P) = \max_{i\in\{0,\ldots,M-1\}} |p_{i+1} - p_i|$. In fact, inequality (3) is trivially satisfied if $p_M \leq c_k + V_1^k/\delta^{K-2}$ since, in this case, $\widehat{p}_k = p_M$. Suppose then that $p_M > c_k + V_1^k/\delta^{K-2}$ and let $\widetilde{p}_k$ be the greatest element in P that is strictly smaller than $\widehat{p}_k$; notice that $\widetilde{p}_k$ is well defined since $p_0 < \min_k c_k$. It follows from the definition of $\widehat{p}_k$ that $\widetilde{p}_k \leq c_k + V_1^k/\delta^{K-2}$. Thus,

$$\mathcal{C}(P) \geq \widehat{p}_k - \widetilde{p}_k \geq \widehat{p}_k - c_k - \frac{V_1^k}{\delta^{K-2}},$$

which is the desired result.

Reordering the States. Now reorder the states so that $\widehat{p}_k$ is (weakly) increasing in k. Since we have not imposed any order on the set Θ of states, this is without loss of generality.

Step 2. As our second step, we use Corollary 1 to derive a lower bound to the buyer's equilibrium payoff V^B. To do this, we propose a strategy $\widehat{\sigma}^B$ for the buyer and compute a lower bound to his payoff when he plays $\widehat{\sigma}^B$ and all other players follow the equilibrium strategy.

The strategy $\widehat{\sigma}^B$ prescribes to offer $\widehat{p}_{t+1}$ in period $t \in \{0, \ldots, K-1\}$ and offer $\widehat{p}_K$ in every period $t \geq K$. Let $u\left(\widehat{\sigma}^B; (\sigma, \mu)\right)$ denote the buyer's ex-ante payoff when he follows the strategy $\widehat{\sigma}^B$ and all other players follow the equilibrium strategy. It follows from Corollary 1 that if a buyer plays $\widehat{\sigma}^B$ and the state is k, then he purchases the good with probability one in period $t \leq k-1$. Thus, given that $\widehat{p}_k$ is increasing in k and $v_k \geq 0$ for all k, the payoff to a buyer from playing $\widehat{\sigma}^B$ is at least $\delta^{k-1}v_k - \widehat{p}_k$ in state k, so that

$$u\left(\widehat{\sigma}^B; (\sigma, \mu)\right) \geq \sum_{k=1}^{K} \pi_k \left(\delta^{k-1} v_k - \widehat{p}_k\right). \tag{4}$$

In equilibrium, $\widehat{\sigma}^B$ cannot be a profitable deviation for a buyer, and so $V^B \geq u\left(\widehat{\sigma}^B; (\sigma, \mu)\right)$. From inequalities (3) and (4), we then obtain

$$\begin{aligned} V^B &\geq \sum_{k=1}^{K} \pi_k \left(\delta^{k-1} v_k - c_k - \frac{V_1^k}{\delta^{K-2}} - \mathcal{C}(P)\right) \\ &= \sum_{k=1}^{K} \pi_k \left[(v_k - c_k) - \delta V_1^k - \left(1 - \delta^{k-1}\right) v_k - \frac{V_1^k\left(1 - \delta^{K-1}\right)}{\delta^{K-2}} - \mathcal{C}(P)\right]. \end{aligned}$$

From the last inequality and Lemma 1(*ii*) we can then conclude that

$$V^B \geq \sum_{k=1}^{K} \pi_k \left[(v_k - c_k) - \delta V_1^k - \left(1 - \delta^{k-1}\right) v_k - \frac{z\left(1 - \delta^{K-1}\right)}{\delta^{K-2}} - \mathcal{C}(P)\right]. \tag{5}$$

Step 3. We can now conclude the proof. Since preferences are quasi-linear, welfare is the sum of the buyers' and sellers' ex-ante equilibrium payoffs:

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

From Lemma 1(i), we have $V^k = V_0^k \geq \delta V_1^k$ for every $k \in \Theta$. This and inequality (5) imply

$$\begin{aligned}
W(\sigma,\mu) &\geq V^B + \sum_{k=1}^{K} \pi_k \delta V_1^k \\
&\geq \sum_{k=1}^{K} \pi_k (v_k - c_k) - \sum_{k=1}^{K} \pi_k \left[\left(1-\delta^{k-1}\right) v_k + \frac{z\left(1-\delta^{K-1}\right)}{\delta^{K-2}} + \mathcal{C}(P)\right] \\
&= W^* - \sum_{k=1}^{K} \pi_k \left[\left(1-\delta^{k-1}\right) v_k + \frac{z\left(1-\delta^{K-1}\right)}{\delta^{K-2}} + \mathcal{C}(P)\right].
\end{aligned}$$

Consequently $W(\sigma,\mu)$ converges to W^* as $\delta \to 1$ and $\mathcal{C}(P) \to 0$, which is the desired result. $\square$

4 Final Remarks

We conclude our analysis with several remarks. First, we show that the assumption of nonnegative gains from trade cannot be relaxed. Then, we discuss the robustness of our efficiency result to alternative bargaining protocols and to the presence of uninformed sellers. After that, we discuss our assumption of restricted price offers and how we can relax it. We also discuss the role of random matching and aggregate uncertainty in our efficiency result. Finally, we briefly discuss information aggregation.

Gains From Trade. Theorem 1 shows that the assumption of nonnegative gains from trade in every state is sufficient for all equilibria to become efficient 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 in $k = 1$. In this case, the first best welfare is $W^* = \pi_2(v_2 - c_2)$. Take a sequence $\{\delta_n\}$ of discount factors that converges to one and, for each $n \in \mathbb{N}$, let (σ_n, μ_n) be a sequential equilibrium when the agents' discount factor is δ_n. Let W_n be welfare in (σ_n, μ_n) and assume towards a contradiction that W_n

converges to W^*. By assumption, $\lim_n \mathbb{E}_{\xi_n}\left[\delta_n^{\mathbf{T}}|k=1\right]=0$ and $\lim_n \mathbb{E}_{\xi_n}\left[\delta_n^{\mathbf{T}}|k=2\right]=1$, where ξ_n is the probability distribution over the set of outcomes induced by σ_n (and the distribution over the set of states) and $\mathbf{T}$ is the (random) time at which trade occurs. In particular, the expected payoff of the sellers when $k=1$ converges to zero. Now let $\mathbf{Q}$ be the first (random) period in which a buyer makes an offer at least as large as c_2. Then $\lim_n \mathbb{E}_{\xi_n}\left[\delta_n^{\mathbf{Q}}|k=2\right]=1$. However, it is easy to show that $\lim_n \mathbb{E}_{\xi_n}\left[\delta_n^{\mathbf{T}}|k=1\right]=0$ and $\lim \mathbb{E}_{\xi_n}\left[\delta_n^{\mathbf{Q}}|k=2\right]=1$ together imply that $\lim_n \mathbb{E}_{\xi_n}\left[\delta_n^{\mathbf{Q}}|k=1\right]=1$. 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 if, and only if, the offer is c_2 or more, a contradiction.

Bargaining Protocol and Uninformed Sellers. It is possible to extend Theorem 1 to the case in which in every buyer-seller meeting the buyer makes a take-it-or-leave-it offer to the seller with positive probability and the seller makes a take-it-or-leave-it offer to the buyer with the remaining probability. A sketch of the proof—which is similar to the proof of Theorem 1—is as follows. A lower bound to a buyer's payoff in any equilibrium is obtained when the buyer: (*i*) rejects any offer that he receives from a seller; and (*ii*) offers the lowest price in P that is greater than the seller's reservation price when the state is k in the kth period in which the buyer gets to make an offer. As trading frictions vanish, this lower bound converges to the first best welfare net of the sellers' ex-ante payoff, which establishes the desired result.

Theorem 1 is not true when sellers have all the bargaining power, though. Not surprisingly, signalling opens up the possibility of equilibria which remain inefficient even as trading frictions vanish. Likewise, Theorem 1 is also not true when some sellers are uninformed about the aggregate state. Again, it is possible to construct inefficient equilibria in which signalling sustains inefficient outcomes: a buyer who deviates by making an offer to attract an uninformed seller changes the uninformed seller's belief in a way that precludes trade.

Limited Price Offers. We assume that buyers are restricted to make offers in a finite set P and consider the limiting case in which P becomes arbitrarily fine. The existence of perfect Bayesian equilibria cannot be guaranteed in our environment when there is a continuum of possible price offers for buyers. We can work with unrestricted price offers by changing our environment so that now in every period the agents in the market are matched with probability $\lambda \in (0, 1)$. In the Appendix we show *constructively* that in this alternative environment a perfect Bayesian equilibrium with unrestricted price offers exists as long as agents are patient enough. We also show that welfare in any perfect Bayesian equilibrium with unrestricted price offers approaches the first-best welfare as trading frictions vanish.

The Role of Aggregate Uncertainty and Random Matching. We depart from most of the literature that studies trading in dynamic decentralized markets with common value uncertainty by assuming that there is aggregate uncertainty. In trading environments with common value uncertainty but no aggregate uncertainty, such as Camargo and Lester (2014) and Moreno and Wooders (2016), 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.[7] In our environment, a single type of seller is present in the market at any point in time. As such, when the price grid is arbitrarily fine, 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. In the limit as δ converges to one, the inefficiency resulting from this strategy converges to zero. Since the buyer's payoff in any equilibrium is bounded below by the payoff they obtain using the strategy described above, all equilibria become efficient as trading frictions vanish.

Some of the driving-forces present in our model with a continuum of agents are also not present in models in which a single seller dynamically meets a with a sequence of short-run buyers (as in Hörner and Vieille 2009), or the same buyer (as in Deneckere and Liang 2006 and Gerardi, Hörner, and Maestri 2014). The main difference between our model and the aforementioned ones is that in our environment a single trader cannot influence the aggregate dynamics of the economy. 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.

[7] A reduction in delay costs reduces the cost for the former type of seller to imitate the latter type of seller.

Information Aggregation. A question that has attracted a lot of attention in the literature is whether markets fully aggregate the information dispersed among agents.[8] In our context, information aggregation is not necessarily achieved. In particular, it is easy to construct examples of (pooling) equilibria which are efficient but fail to aggregate information perfectly.

Appendix

We make two modifications in the benchmark model in this appendix. First, we assume that in every period a buyer and a seller are matched with probability $\lambda \in (0, 1)$. Second, we assume that the set of possible price offers is unrestricted, i.e., equal to $\mathbb{R}$.

We consider pairs (σ, μ), where σ is a strategy profile and μ is a profile of belief systems for buyers, that constitute a perfect Bayesian equilibrium. We first show that a perfect Bayesian equilibrium exists if agents are patient enough. We then show that all perfect Bayesian equilibria become efficient as discounting vanishes. As before, we refer to a seller in state k as a type-k seller.

Theorem 2. *There exits $\delta^* \in (0, 1)$ such that for all $\delta > \delta^*$ a perfect Bayesian equilibrium exists.*

Proof. Without loss of generality, order the set of states in such a way that c_k is increasing in k and let $\mathcal{C} = \{c \in \mathbb{R}_+ : c = c_k \text{ for some } k \in \Theta\}$. Then $\mathcal{C} = \{\bar{c}_1, \dots, \bar{c}_L\}$ with $\bar{c}_1 < \cdots < \bar{c}_L$ and $L \leq K$; notice that $L = K$ when the costs c_1 to c_K are distinct. Now let $f : \mathcal{C} \rightrightarrows \Theta$ be the correspondence such that $f(\bar{c}_\ell) = \{k \in \Theta : c_k = \bar{c}_\ell\}$.

Consider the pair (σ, μ) in which sellers follow the same strategy σ^S and buyers follow the same strategy σ^B and have the same belief system μ^B, where σ^S, σ^B, and μ^B are defined as follows:

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

ii) 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 made so far,

[8] The 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. More recent papers in this literature are Golosov, Lorenzoni, and Tsyvinski (2014) and Lauermann and Wolinsky (2016). None of these papers are concerned with market efficiency.

with the convention that $\overline{p}(h^t) = -\infty$ if $m(h^t) = 0$; and (*iii*) $\mathcal{C}(\overline{p}(h^t)) = \{c \in \mathcal{C} : c > \overline{p}(h^t)\}$.[9] Notice that $\mathcal{C}(\overline{p}(h^t)) = \mathcal{C}$ if $m(h^t) = 0$. A buyer's behavior after a history h^t is as follows. If $\mathcal{C}(\overline{p}(h^t)) \neq \emptyset$, then the buyer offers the smallest element of $\mathcal{C}(\overline{p}(h^t))$ if he is matched in the market. If, instead, $\mathcal{C}(\overline{p}(h^t)) = \emptyset$, then the buyer offers $\overline{c}_L$ if he is matched in the market. The strategy σ^B is such that on the path of play, for each $k \in \{1, \ldots, L\}$, a buyer offers $\overline{c}_k$ the kth time that he is matched in the market.

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

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

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

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

to every other state.

It is straightforward to check that σ^S is sequentially rational and that μ is consistent with Bayes' rule on the path of play. Therefore, it suffices to show that σ^B is sequentially rational to conclude the proof. Consider a buyer with history h^t and notice that there exists $j \in \{1, \ldots, L\}$ such that $\mu(h^t)$ assigns probability

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

to every state $k \in \bigcup_{k=j}^{L} f(\overline{c}_k)$ and assigns probability 0 to every other state. If $j = L$, then the buyer does not have a profitable deviation as he believes that $k \in f(\overline{c}_L)$ and expects that any seller accepts an offer of $\overline{c}_L$.

Suppose then that $j < L$ and let $\tau \geq L - 1$ be the random time at which the buyer is matched for the Lth time in the market. If the buyer follows the equilibrium strategy, then his expected

[9] Now a history for a buyer in period $t \geq 1$ is a sequence $(\widetilde{p}_1, \ldots, \widetilde{p}_{t-1})$, where $\widetilde{p}_s = \emptyset$ means that the buyer was not matched in period $s \in \{1, \ldots, t-1\}$ and $\widetilde{p}_s \in \mathbb{R}$ means that the buyer was matched in period s and his offer $\widetilde{p}_s$ was rejected. A history for a seller in period t is defined similarly.

payoff is bounded below by

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

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

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

where we used the fact that with probability at least $\sum_{\ell \in f(\overline{c}_j)} \widetilde{\pi}_\ell$ there is a state in which the buyer purchases the good at a price at least $(\overline{c}_{j+1} - \overline{c}_j) > 0$ greater than $\overline{c}_j$. Since $\lim_{\delta \to 1} \mathbb{E}[\delta^{\tau-1}] = 1$ by the dominated convergence theorem, there exists $\delta_j \in (0,1)$ such that $\delta > \delta_j$ implies that $u_j > u'_j$.

Letting $\delta^* = \max\{\delta_1, \ldots, \delta_{L-1}\}$, we can conclude that σ^B is sequentially rational whenever $\delta > \delta^*$. It then follows that for all $\delta > \delta^*$, the pair (σ, μ) is a perfect Bayesian equilibrium. □

Theorem 3. *Let $\{\delta_n\}$ be a sequence of discount factors converging to one. For any sequence $\{(\sigma_n, \mu_n)\}$ of equilibria such that (σ_n, μ_n) is an equilibrium when $\delta = \delta_n$ the sequence $\{W(\sigma_n, \mu_n)\}$ converges to the first-best welfare W^*.*

Proof. We show that for all $\varepsilon \in (0, z)$ there exists $\overline{\delta} \in (0,1)$ such that if $\delta > \overline{\delta}$, then $W(\sigma, \mu) > W^* - \varepsilon$ for every pair (σ, μ) that is a perfect Bayesian equilibrium when the agents' discount factor is δ; recall that $z = \max_k v_k$.

Fix $\varepsilon \in (0, z)$ and let (σ, μ) be a perfect Bayesian equilibrium for some discount factor δ. As in the main text, in equilibrium the payoffs to buyers are the same and the payoffs to sellers in each state are the same. Let V^B denote the buyers' ex-ante (equilibrium) payoff and V^k denote the type-k sellers' payoff. Also, let V_t^k denote a type-k seller's payoff in period t, which does not depend on a seller's private history. Lemma 1 is still valid and its proof is the same as before.

Step 1. We proceed as in the proof of Theorem 1 and first identify for each $k \in \Theta$ 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 3. *Consider the equilibrium (σ, μ). For every $k \in \Theta$ and $t \in \{0, \ldots, T(\kappa)K - 1\}$, 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 $t \in \{0, \ldots, T(\kappa)K-1\}$, then he accepts it with probability one.

The proof of Lemma 3 is identical to the proof of Lemma 2 and is, therefore, omitted. It follows from Lemma 3 that if a buyer offers

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

in any period $t \in \{0, \ldots, T(\kappa)K-1\}$, then a type-$k$ seller accepts this offer with probability one.

Reordering the States. Next, we reorder the states so that $\widehat{p}_k$ is (weakly) increasing in k. Since we have not imposed any order on the states, this is without loss of generality.

Step 2. We now use Lemma 3 to derive a lower bound to the buyers' equilibrium payoff. Consider a buyer who follows the strategy $\widehat{\sigma}^B$ that prescribes him to offer $\widehat{p}_k$ if he is matched in periods $t \in \{T(\kappa)(k-1), \ldots, T(\kappa)k-1\}$, with $k \in \Theta$, and offer $\widehat{p}_K$ if he his matched in period $t \geq T(\kappa)K$. Denote by $u\left(\widehat{\sigma}^B; (\sigma, \mu)\right)$ the payoff the buyer obtains when all other agents follow the equilibrium strategy. Notice that $V^B \geq u\left(\widehat{\sigma}^B; (\sigma, \mu)\right)$, otherwise the buyer would have a profitable deviation.

We obtain a lower bound for $u\left(\widehat{\sigma}^B; (\sigma, \mu)\right)$, and thus V^B, as follows. Suppose that the 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, which is only possible if $k > 1$, the buyer's expected payoff is at least $\delta^{T(\kappa)k-1}v_k - \widehat{p}_k$; this is because $\widehat{p}_k$ 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 $t \in \{T(\kappa)(k-1), \ldots, T(\kappa)k-1\}$ and obtains a payoff of at least $\delta^{T(\kappa)k-1}v_k - \widehat{p}_k$, or the buyer is not matched with a seller in any of these periods and obtains a payoff of a least $-\widehat{p}_K$. Since $(1-\lambda)^{T(\kappa)} < \kappa$, the probability that the buyer does not meet a seller in some period $t \in \{T(\kappa)(k-1), \ldots, T(\kappa)k-1\}$ is at most κ. So, the buyer's expected payoff in the second event is at least $(1-\kappa)\left(\delta^{T(\kappa)k-1}v_k - \widehat{p}_k\right) - \kappa\widehat{p}_K$. Given that the lower bound to the buyers' expected payoff is lower in the second event, we then have that

$$\begin{aligned} u\left(\widehat{\sigma}^B; (\sigma, \mu)\right) \;&\geq\; (1-\kappa)\sum_{k=1}^{K} \pi_k \left(\delta^{T(\kappa)k-1}v_k - \widehat{p}_k\right) - \kappa\widehat{p}_K \\ &\geq\; (1-\kappa)\sum_{k=1}^{K} \pi_k \left(\delta^{T(\kappa)K-1}v_k - \widehat{p}_k\right) - \kappa\widehat{p}_K. \end{aligned}$$

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

$$\kappa \widehat{p}_K \le \kappa \left(c_k + V_1^K + \varepsilon + \frac{\left(1 - \delta^{T(\kappa)K-2}\right) V_1^K}{\delta^{T(\kappa)K-2}} \right) \le 3\kappa z + z \left(\frac{1 - \delta^{T(\kappa)K-2}}{\delta^{T(\kappa)K-2}} \right). \tag{6}$$

Moreover, we also have that

$$\begin{aligned}
(1-\kappa) \sum_{k=1}^{K} \pi_k \left(\delta^{T(\kappa)K-1} v_k - \widehat{p}_k\right) &= (1-\kappa) \sum_{k=1}^{K} \pi_k \left(\delta^{T(\kappa)K-1} v_k - c_k - \frac{\delta V_1^k}{\delta^{T(\kappa)K-1}} - \frac{\varepsilon}{4} \right) \\
&= (1-\kappa) \left[\sum_{k=1}^{K} \pi_k (v_k - c_k) - \left(1 - \delta^{T(\kappa)K-1}\right) \sum_{k=1}^{K} \pi_k v_k - \frac{\varepsilon}{4} - \sum_{k=1}^{K} \pi_k \frac{\delta V_1^k}{\delta^{T(\kappa)K-1}} \right] \\
&\ge \sum_{k=1}^{K} \pi_k (v_k - c_k) - \kappa \sum_{k=1}^{K} \pi_k v_k - \left(1 - \delta^{T(\kappa)K-1}\right) \sum_{k=1}^{K} \pi_k v_k - \frac{\varepsilon}{4} - \sum_{k=1}^{K} \pi_k \frac{\delta V_1^k}{\delta^{T(\kappa)K-1}} \\
&\ge \sum_{k=1}^{K} \pi_k (v_k - c_k) - \kappa z - \left(1 - \delta^{T(\kappa)K-1}\right) z - \frac{\varepsilon}{4} - z \left(\frac{1 - \delta^{T(\kappa)K-1}}{\delta^{T(\kappa)K-1}} \right) - \sum_{k=1}^{K} \pi_k \delta V_1^k,
\end{aligned} \tag{7}$$

where the last inequality follows from the fact that $\sum_{k=1}^{K} \pi_k v_k \le z$ and $\sum_{k=1}^{K} \pi_k \delta V_1^k \le z$.

Using inequalities (6) and (7) and the facts that $\kappa z = \varepsilon/16$ and $(1 - \delta^t)/\delta^t$ is increasing in t, we then have that

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

Step 3. We conclude by using the above lower bound on V^B to obtain a lower bound on welfare. Since $W(\sigma, \mu) = V^B + \sum_{k=1}^{K} \pi_k V^k$ and, for all $k \in \Theta$, $V^k = V_0^k \ge \delta V_1^k$, we have

$$W(\sigma, \mu) \ge \sum_{k=1}^{K} \pi_k (v_k - c_k) - \left(1 - \delta^{T(\kappa)K-1}\right) z - \frac{\varepsilon}{2} - 2z \left(\frac{1 - \delta^{T(\kappa)K-1}}{\delta^{T(\kappa)K-1}} \right).$$

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

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

we can then conclude that $W(\sigma, \mu) > W^* - \varepsilon$ whenever $\delta > \overline{\delta}$, which is the desired result. $\square$

References

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.

Deneckere, R., and M.-Y. Liang. 2006. Bargaining with Interdependent Values. *Econometrica*, 74(5), 1309–1364.

De Fraja, G., and J. Sákovics. 2001. Walras Retrouvé: Decentralized Trading Mechanisms and the Competitive Price. *Journal of Political Economy*, 109, 842–863.

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.

Lauermann, S. 2013. Dynamic Matching and Bargainin Games: A General Approach. *American Economic Review*, 103, 663–689.

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

Moreno, D., and J. Wooders. 2016. Dynamic markets for lemons: performance, liquidity, and policy intervention. *Theoretical Economics*, 11(2), 601–639.

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

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. 1988. Dynamic Markets with Competitive Bidding. *Review of Economic Studies*, 58, 351–374.

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

Supplementary Material (Not for Publication):

Here, we sketch the proof that a sequential equilibrum exists in our environment when sellers are restricted to make offers in a finite set P.

For each $n \geq 1$, let $\mathcal{G}_n$ be the game in which buyers are restricted to play behavior strategies assigning probability at least $1/n(M+1)$ to each element of P and sellers are restricted to play behavior strategies assigning probability at least $1/2n$ to each acceptance decision. Fix $n \geq 1$. Since the action sets of buyers and sellers are finite, a standard argument shows that $\mathcal{G}_n$ has a Nash equilibrium σ_n. Moreover, since under σ_n every history in $\mathcal{G}_n$ is reached with positive probability, there exists a belief system μ_n for buyers such that (σ_n, μ_n) is a sequential equilibrium of $\mathcal{G}_n$.

Consider now the sequence $\{(\sigma_n, \mu_n)\}$ of sequential equilibria. Since the set $\mathcal{H} = \bigcup_{t=1}^{\infty} \mathcal{H}_t$ is countable, a standard argument shows that $\{(\sigma_n, \mu_n)\}$ admits a subsequence $\{(\sigma_{n_k}, \mu_{n_k})\}$ such that the numerical sequence $\{(\sigma_{n_k}(h), \mu_{n_k}(h))\}$ is convergent for all $h \in \mathcal{H}$. Assume, without loss, that $\{(\sigma_n, \mu_n)\}$ itself has this property, and let $(\sigma_\infty, \mu_\infty)$ be its pointwise limit. We claim that σ_∞ is sequentially rational given μ_∞, so that $(\sigma_\infty, \mu_\infty)$ is a sequential equilibrium. We only consider buyers, as the proof for sellers is similar. In what follows, let Σ_n^B be the set of strategies for the buyers in $\mathcal{G}_n$ and notice that $\Sigma_{n_1}^B \supseteq \Sigma_{n_2}^B$ for all $n_1 > n_2$.

Fix $h \in \mathcal{H}$ and consider the continuation game after the history h. Let $V^B(\sigma^B|\sigma, \mu)$ be the expected payoff to a buyer who follows the strategy σ^B after h when aggregate behavior is given by the strategy profile σ and the belief system for the buyers is μ; we omit the dependence of the payoff $V^B(\sigma^B|\sigma, \mu)$ on h for ease of exposition. Suppose, by contradiction, that there exists a strategy $\widetilde{\sigma}^B$ for buyers such that

$$V^B(\widetilde{\sigma}^B|\sigma_\infty, \mu_\infty) \geq V^B(\sigma_\infty^B|\sigma_\infty, \mu_\infty) + \varepsilon \tag{8}$$

for some $\varepsilon > 0$. Because of discounting, there exists $n_1 \in \mathbb{N}$ and $\widehat{\sigma}^B \in \Sigma_n^B$ for all $n \geq n_1$ such that

$$V^B(\widehat{\sigma}^B|\sigma_\infty, \mu_\infty) \geq V^B(\widetilde{\sigma}^B|\sigma_\infty, \mu_\infty) - \frac{\varepsilon}{4}. \tag{9}$$

Moreover, by the construction of $(\sigma_\infty, \mu_\infty)$, there exists $n_2 \in \mathbb{N}$ such that if $n \geq n_2$, then

$$V^B(\widehat{\sigma}^B|\sigma_n, \mu_n) \geq V^B(\widehat{\sigma}^B|\sigma_\infty, \mu_\infty) - \frac{\varepsilon}{4} \tag{10}$$

and

$$V^B(\sigma_n^B|\sigma_n,\mu_n) \leq V^B(\sigma_\infty^B|\sigma_\infty,\mu_\infty) + \frac{\varepsilon}{4}. \tag{11}$$

Hence, $n \geq \max\{n_1, n_2\}$ implies that

$$V^B(\widehat{\sigma}^B|\sigma_n,\mu_n) \geq V^B(\widetilde{\sigma}^B|\sigma_\infty,\mu_\infty) - \frac{\varepsilon}{2} \geq V^B(\sigma_\infty^B|\sigma_\infty,\mu_\infty) + \frac{\varepsilon}{2} > V^B(\sigma_n^B|\sigma_n,\mu_n),$$

where the first inequality follows from (9) and (10), the second inequality follows from (8), and the third inequality follows from (11). Given that $\widehat{\sigma}^B \in \Sigma_n^B$, we can then conclude that σ_n^B is not sequentially rational for the buyers in $\mathcal{G}_n$, a contradiction.