

Efficiency in Decentralized Markets with Aggregate Uncertainty

Braz Camargo*

Dino Gerardi[†]

Lucas Maestri[‡]

December 2015

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.

[‡]FGV/EPGE - Escola Brasileira de Economia e Finanças. E-mail: maestri.lucas@gmail.com.

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 the sellers know but the buyers do not know. There is a finite number of such states and the only assumption we make is that gains from trade are nonnegative in all states. Time is discrete and in every period the buyers and sellers still 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. Buyers are restricted to make offers in a finite set. This assumption ensures the existence of equilibria.

In our environment there are three 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, there is the additional trading friction that buyers are restricted in the prices that they can offer to sellers. Third, 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 and the set of possible price offers by buyers becomes

arbitrarily fine, 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 Sákovicš 2001, Serrano 2002, Satterthwaite and Shneyerov 2007, and Lauermann 2013). Moreover, of the few papers that study environments with common values, see, e.g., Blouin and Serrano (2001), Blouin (2003), and Moreno and Wooders (2010), 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 restrict attention to one-sided private information. This assumption constitutes a natural starting point for an analysis of market efficiency, as it allows us to avoid signalling problems. 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 the market outcome remains inefficient even when discounting vanishes. Our analysis shows that restricting prices has a critical impact on market efficiency.

2 The Environment

Time is discrete and indexed by $t \in \{0, 1, \dots\}$. 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 $\Theta = \{1, \dots, K\}$ 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, Assumption 1 cannot be relaxed. Notice that 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. Buyers are restricted to make offers in a set $P = \{p_0, p_1, \dots, p_M\}$, where $p_i < p_{i+1}$ for all $i \in \{1, \dots, M\}$, $p_0 < \min_k c_k$, and $p_M > \max_k c_k$.¹ It is possible to show that perfect Bayesian equilibria may fail to exist when the set of offers is the real line. In what follows, we let $\mathcal{C}(P) = \max_{0 \in \{1, \dots, M-1\}} |p_{i+1} - p_i|$ be the *coarseness* of P .

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 .² A pure strategy for a buyer is a sequence $s^B = \{s_t^B\}$, where $s_t^B(h^t)$ is the price the buyer offers in period t if his private history is h^t . A pure strategy for a seller is a sequence $s^S = \{s_t^S\}$, where $s_t^S(k, h^t, p)$ is the seller's acceptance decision in period t —accept or reject—when the state is k if his private history is h^t and the offer he receives is p . A belief system for the buyers is a sequence $\mu = \{\mu_t\}$ such that $\mu_t(h^t)$ is a buyer's (posterior) belief about the state if he is in the market in period t with private history h^t . Given the continuum population assumption, it is without loss of generality to assume that all buyers follow the same (mixed) strategy σ^B and all sellers follow the same (mixed) strategy σ^S .

Let σ be a strategy profile. 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

¹The 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 $\max_k c_k$.

²A private history for an agent in the market in period t is a sequence $h^t = (p_1, \dots, p_{t-1})$ of price offers. If the agent is a buyer, h^t is the sequence of price offers that the agent made and were rejected. On the other hand, if the agent is seller, then h^t is the sequence of price offers that the agent received and rejected.

payoffs to agents are well-defined even if there is a zero mass of agents in the market.³ 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. The assumption of a finite set of possible price offers ensures the existence of a sequential equilibrium.⁴

3 Market Efficiency

Besides a real time between two consecutive trading opportunities, our environment features the additional trading friction that buyers are restricted in the prices that they can offer. Trading frictions disappear when both δ converges to one, so that matching becomes frictionless, and $\mathcal{C}(P)$ converges to zero, so that the set of possible price offers by the buyers becomes arbitrarily fine.

In this section, we first show that restricting the set of possible prices that a buyer can offer to be finite can lead to equilibria which remain inefficient as δ converges to one even when there is no aggregate uncertainty. We then show our main result, namely, that in any equilibrium, welfare approaches the first best welfare as trading frictions disappear. We finish with a number of remarks.

3.1 Inefficiency with Restricted Price Offers

Clearly, equilibria can be inefficient even as discounting disappears if buyers can offer prices only in a finite set. For instance, if buyers can offer only a single positive price, p , to the sellers, then all equilibria are inefficient even without matching frictions when $p < \max_k c_k$. Here, we show that no matter how fine the set P is, i.e., no matter how small $\mathcal{C}(P)$ is, there can exist equilibria which remain inefficient as matching frictions disappear even when there is no aggregate uncertainty, i.e., the set Θ is a singleton. We also show that our equilibrium construction is robust to the introduction of aggregate uncertainty.

Suppose that $K = 1$, with $v_1 > c_1$. Suppose also that the set P is fine enough that it contains two consecutive prices, p_ℓ and p_h , with $c_1 < p_\ell < p_h < v_1$. Consider then the following stationary

³Indeed, 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 (σ, μ) is such that σ has full support.

⁴The existence of a sequential equilibrium follows from a standard argument. For completeness, we provide a sketch of the argument in the Appendix.

strategy profile. In each period, a buyer randomizes between offering p_ℓ and p_h , while a seller accepts any $p \in P$ with $p \geq p_h$, rejects any $p \in P$ with $p < p_\ell$, and randomizes between accepting and rejecting p_ℓ . Let α be the probability that a seller accepts p_ℓ and β be the probability that a buyer offers p_h . The following conditions are necessary for an equilibrium:

$$v_1 - p_h = \alpha(v_1 - p_\ell) + (1 - \alpha)\delta(v_1 - p_h); \quad (1)$$

$$p_\ell - c_1 = \delta[(1 - \beta)p_\ell + \beta p_h - c_1]. \quad (2)$$

Condition (1) states that a buyer is indifferent between offering p_ℓ and p_h , while condition (2) states that a seller is indifferent between accepting and rejecting p_ℓ . Clearly, given the buyers' behavior, a seller accepts any offer greater than or equal to p_h . Moreover, since the set P is fixed, a seller rejects any offer less than p_ℓ if he is patient enough. Finally, given the sellers' behavior, a buyer has no incentive to offer more than p_h or less than p_ℓ . Thus, (1) and (2) are also sufficient for an equilibrium as long as δ is close enough to one.

Solving for α and β we obtain that:

$$\alpha = \frac{(1 - \delta)(v_1 - p_h)}{(1 - \delta)(v_1 - p_\ell) + \delta(p_h - p_\ell)};$$

$$\beta = \frac{(1 - \delta)(p_\ell - c_1)}{\delta(p_h - p_\ell)}.$$

Let $\gamma = 1 - (1 - \beta)(1 - \alpha) = \beta + \alpha - \beta\alpha$ be the probability that an offer from a buyer is accepted by a seller. The discounted probability of trade for the good is

$$\rho = \frac{\gamma}{1 - (1 - \gamma)\delta}.$$

We claim that $\rho = \rho(\gamma)$ does not converge to one as δ increases to one, so that the equilibrium under consideration remains inefficient even as matching frictions disappear. Since ρ is strictly increasing in γ , it suffices to show that $\lim_{\delta \rightarrow 1} \rho(\bar{\gamma}) < 1$, where $\bar{\gamma} = \alpha + \beta$ is an upper bound for γ . Now observe that

$$\bar{\gamma} = \frac{(1 - \delta)(v_1 - p_h)}{(1 - \delta)(v_1 - p_\ell) + \delta(p_h - p_\ell)} + \frac{(1 - \delta)(p_\ell - c_1)}{\delta(p_h - p_\ell)} \leq (1 - \delta) \underbrace{\left[\frac{v_1 - p_h}{v_1 - p_\ell} + \frac{2(p_\ell - c_1)}{p_h - p_\ell} \right]}_A,$$

where the second inequality holds as long as $\delta \geq 1/2$. Thus,

$$\rho \leq \frac{A(1-\delta)}{1 - [1 - A(1-\delta)]\delta} = \frac{A}{1 + A\delta}$$

for all δ sufficiently close to one. This establishes the desired result.

We can extend the above equilibrium construction to obtain an equilibrium which is inefficient even in the limit as δ converges to one. For example, let $K = 2$ and assume that $c_1 < c_2 < v_1 < v_2$. Suppose also that the set P is fine enough that it contains two consecutive prices, p_ℓ and p_h , with $c_2 < p_\ell < p_h < v_1$. Under these assumptions, when δ is close enough to one, it is straightforward to construct a sequential equilibrium with the following properties: (i) buyers offer p_ℓ in the first period, and sellers accept this offer if, and only if, $k = 1$; and (ii) behavior from the second period on is the same as in the stationary equilibrium constructed for the case when $K = 1$, with c_2 and v_2 replacing c_1 and v_1 , respectively, in conditions (1) and (2).

Going back to the example without aggregate uncertainty, notice that since $\gamma \geq \alpha$,

$$\rho \geq \frac{\alpha}{1 - (1 - \alpha)\delta} = \frac{v_1 - p_h}{v_1 - p_\ell}.$$

Hence, as $p_h - p_\ell$ converges to zero, ρ converges to one. Thus, even though the equilibrium we constructed remains inefficient as matching frictions disappear, this inefficiency disappears as P becomes arbitrarily fine. The same result holds in the example with aggregate uncertainty.

3.2 Limit Efficiency

The previous subsection showed by means of examples that restricting the prices that buyers can offer can lead to inefficient equilibria even when matching frictions disappear. However, the equilibria that we constructed are such that their inefficiency becomes negligible when the set of possible prices that buyers can offer becomes arbitrarily fine. We now show that this second result is not particular to the examples that we considered but, instead, it holds in general.

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 $\mathcal{O}$.⁵ Let $\mathbb{E}_\xi$ denote the expectation with respect to ξ . Welfare in the equilibrium (σ, μ) is

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

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 $\max_{p \in P} p > \max_k c_k$.

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 pair (δ, P) and let (σ, μ) be an equilibrium. Denote the buyer's ex-ante equilibrium payoff by V^B and the seller's equilibrium payoff in state k by V^k . Moreover, denote the present discounted payoff to a seller in the market in period t when the state is k by V_t^k ; given that a seller's history is his private information and there is a continuum of agents, V_t^k does not depend on a seller's history in period t . Notice that $V_0^k = V^k$ and that since an option for a seller is to reject any offers that he receives, $V_s^k \geq \delta^{t-s} V_t^k$ for all $k \in \Theta$ and all $t > s \geq 0$.

For each $k \in \Theta$, let $\underline{p}^k = c_k + \delta V_k^k$ be the seller's reservation price in period $k - 1$ when the state is k . Since $V_k^k \leq p_M - c_k$, we have that $\underline{p}^k \leq (1 - \delta)c_k + \delta p_M < p_M$. By reordering the set of states if necessary, we can assume that $\underline{p}^k$ is nondecreasing in k . Now, for each $k \in \Theta$, define $\hat{p}^k$ to be the smallest price in P such that $\hat{p}^k > c_k + V_1^k / \delta^{k-2}$ if such a price exists and $\hat{p}^k = p_M$ otherwise. Given that

$$c_k + \frac{V_1^k}{\delta^{k-2}} > c_k + \delta^{k-1} \frac{V_1^k}{\delta^{k-2}} \geq c_k + \delta V_k^k = \underline{p}^k,$$

we then have that

$$c_k + \frac{V_1^k}{\delta^{k-2}} + \mathcal{C}(P) \geq \hat{p}^k > \underline{p}^k.$$

Now observe that an option for a buyer is to behave according to the following strategy: offer $\hat{p}^{t+1}$ in period $t \in \{0, \dots, K - 1\}$, and then offer $\hat{p}^K$ from period K on. Thus, a lower bound to

⁵When $\mathbf{T} = \infty$, the transaction price is undetermined. We adopt the convention that $\mathbf{p} = p_0$ in this case.

the buyer's equilibrium payoff is

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

Consequently, since $W(\sigma, \mu) = V^B + \sum_{k=1}^K \pi_k V^k$ by quasi-linearity of preferences, we have that

$$W(\sigma, \mu) \geq \sum_{k=1}^K \pi_k \delta^{k-1} (v_k - c_k - \mathcal{C}(P)). \quad (3)$$

The desired result follows from the fact that the right side of (3) converges W^* as δ increases to one and $\mathcal{C}(P)$ decreases to zero. $\square$

3.3 Final Remarks

We conclude our analysis with a number of remarks. First, we show that the assumption of non-negative gains from trade cannot be relaxed. Then, we discuss the robustness of our main result to alternative bargaining protocols. 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 $v_1 < c_1 < c_2 < v_2$. In this case, the first best welfare is $W^* = \pi_2(v_2 - c_2)$. Suppose also that $\pi_2(v_2 - c_2) < \pi_1(c_2 - c_1)$. Fix an equilibrium (σ, μ) and let V^B be the buyers' (ex-ante) equilibrium payoff and V^k be the sellers's equilibrium payoff when the state is $k \in \{1, 2\}$. Moreover, let $\mathbf{Q} \leq \mathbf{T}$ be the first (random) period in which a buyer makes an offer of at least c_2 . Given that an option for a seller is to reject an offer less than c_2 and accept any other offer, we have that $V^k \geq \mathbb{E}_\xi [\delta^{\mathbf{Q}}(c_2 - c_k) | k]$ for each $k \in \{1, 2\}$. Likewise, since an option for a buyer is to always make an offer that is rejected, we also have that $V^B \geq 0$. On the other hand, $V^B + \pi_1 V^1 + \pi_2 V^2 \leq W^*$. Hence,

$$\mathbb{E}_\xi [\delta^{\mathbf{Q}} | k = 1] \leq \frac{\pi_2(v_2 - c_2)}{\pi_1(c_2 - c_1)} < 1.$$

Now observe that in equilibrium any offer less than c_2 must be rejected by a seller; a seller would accept such an offer only if $k = 1$ (and the offer is not smaller than c_1), in which case buyers are better off not making the offer. Hence, $\mathbb{E}_\xi [\delta^Q | k = 2] = \mathbb{E}_\xi [\delta^Q | k = 1]$, and so

$$W(\sigma, \mu) \leq \mathbb{E}_\xi [\delta^Q (v_2 - c_2) | k = 2]$$

is bounded away from W^* regardless of δ and $\mathcal{C}(P)$. Thus, *all* equilibria are inefficient regardless of how small trading frictions are.

Bargaining Protocol It is straightforward 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 k th 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, the multiplicity of equilibria afforded by signalling opens up the possibility of equilibria which remain inefficient even as trading frictions vanish.

Information Aggregation A question that has attracted much attention in the literature is whether markets fully aggregate the information dispersed among agents.⁶ 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.

⁶The seminal reference in the literature on information aggregation in decentralized markets is Wolinsky (1990). Blouin and Serrano (2001) extends the analysis in Wolinsky (1990) to non steady-state environments. More recent papers in this literature are Golosov, Lorenzoni, and Tsyvinski (2014) and Lauermaun and Wolinsky (2015).

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.
- De Fraja, G., and J. Sákovics. 2001. Walras Retrouvé: Decentralized Trading Mechanisms and the Competitive Price. *Journal of Political Economy*, 109, 842–863.
- Golosov, M., Lorenzoni, G., and A. Tsyvinski. 2014. Decentralized Trading with Private Information. *Econometrica*, 82, 1055–1091.
- Lauermann, S. 2013. Dynamic Matching and Bargaining Games: A General Approach. *American Economic Review*, 103, 663–689.
- Lauermann, S., and A. Wolinsky. 2015. Search with Adverse Selection. *Econometrica*, forthcoming.
- Moreno, D., and J. Wooders. 2010. Decentralized Trade Mitigates the Lemons Problem. *International Economic Review*, 51, 383–399.
- 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. 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.

Appendix (Not for Publication)

Existence of Sequential Equilibria (Sketch)

For each $n \geq 1$, let $\mathcal{G}_n$ be the auxiliary 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 both buyers and sellers are finite, a standard argument shows that $\mathcal{G}_n$ has a Nash equilibrium σ_n . Moreover, given that under σ_n every history in $\mathcal{G}_n$ is reached with positive probability, there exists a belief system μ_n for the 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 $\sigma_\infty = (\sigma_\infty^B, \sigma_\infty^1, \dots, \sigma_\infty^K)$ is sequentially rational given μ_∞ , so that $(\sigma_\infty, \mu_\infty)$ is a sequential equilibrium. We only consider the buyers, as the proof for the 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 $\tilde{\sigma}^B$ for the buyers such that

$$V^B(\tilde{\sigma}^B | \sigma_\infty, \mu_\infty) \geq V^B(\sigma_\infty^B | \sigma_\infty, \mu_\infty) + \varepsilon \quad (4)$$

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

$$V^B(\hat{\sigma}^B | \sigma_\infty, \mu_\infty) \geq V^B(\tilde{\sigma}^B | \sigma_\infty, \mu_\infty) - \frac{\varepsilon}{4}. \quad (5)$$

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

$$V^B(\hat{\sigma}^B | \sigma_n, \mu_n) \geq V^B(\hat{\sigma}^B | \sigma_\infty, \mu_\infty) - \frac{\varepsilon}{4} \quad (6)$$

and

$$V^B(\sigma_n^B|\sigma_n, \mu_n) \leq V^B(\sigma_\infty^B|\sigma_\infty, \mu_\infty) + \frac{\varepsilon}{4}. \quad (7)$$

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

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

where the first inequality follows from (5) and (6), the second inequality follows from (4), and the third inequality follows from (7). Given that $\hat{\sigma}^B \in \Sigma_n^B$, we can then conclude that σ_n^B is not sequentially rational for the buyers in $\mathcal{G}_n$, a contradiction.