Retrading in Market Games *

Sayantan Ghosal[⋆] and Massimo Morelli[⋆⋆]

[⋆] *University of Warwick*

[⋆⋆] *The Ohio State University*

First draft: July 2000; This draft: April 2001.

Abstract

We study retrading via reopening of trading posts in a market game where the static Nash equilibria yield Pareto suboptimal allocations whenever endowments are Pareto suboptimal. We show that there are allocations on the Pareto frontier that can be approximated arbitrarily closely along some Subgame Perfect Equilibrium path of retrading. The approximation result can also be proved when traders are myopic. However, no allocation on the Pareto frontier can be attained in finite time.

Keywords: Market Games, Retrading, Myopic versus Far-sighted Behavior.

*We wish to thankk M. Cripps, B. Dutta, R. Guesnerie, M. Machina, J. Peck, and the participants of various workshops, for their comments. All remaining mistakes are obviously our own.

1 Introduction

Traditional analysis of market games,[1] markets where the rules of exchange allow all traders to influence prices by sending quantity signals, often results in non-cooperative equilibria with inefficient outcomes. This result has been made precise by Dubey and Rogowski (1990), who have shown that under some regularity assumptions on preferences, the allocations obtained at the Nash equilibria of a market game are Pareto optimal if and only if the initial endowments are Pareto optimal as well.

This analysis of market games ignores the incentives traders have to reopen trading posts before they consume their final allocations. We address this issue by studying a model of trade where traders can reopen trading posts a finite or infinite number of times before they consume. At each round of trade, the rules of exchange are according to the Shapley-Shubik market game, where all commodities are exchanged at trading posts except for the numeraire commodity, in which bids for all other commoditiues have to be made. For each non-numeriare commodity, traders can submit bids for the commodity and make offers of a quantity of the commodity, at the relevant trading post. In any new round of trade, the endowments of individuals are their final allocations from the previous round of trade. Using these endowments, individuals now make bids and offers in the trading posts and obtain allocations determined by the same price formation rule and allocation rule. The cost of reopening trading posts in any new round of trade is measured by a common discount factor for all traders. We assume that individuals observe the bids and offers made by all traders in all previous rounds of trade. However, our results hold even when traders can only observe the price history. We will consider myopic retrading as well as far-sighted retrading. A path of myopic retrading only requires that each period allocation be a Nash equilibrium outcome given the final allocation of the previous period. The results with myopic retrading will be a useful benchmark.

[1]Shapley and Shubik (1977), Dubey and Shubik (1978), Sahi and Yao (1989), Peck and Shell (1991).

We show that, along any equilibrium path of finite retrading, with or without far-sighted behavior, no allocation on the Pareto frontier can be attained even when the cost of reopening trading posts is negligible. The intuition behind this result is simple. As trade always concludes after some finite length of time, this implies that at some finite stage in the game it must be the case that while the traders' inherited allocation from the previous period is Pareto suboptimal, the final allocation they obtain after reopening trading posts is both (a) a Pareto optimal allocation and (b) is a Nash equilibrium allocation for the one-shot market game with the traders' inherited allocation as the endowment. However, under our assumptions on preferences and endowments, no Nash equilibrium allocation of the one-shot market game can ever be Pareto optimal. This ensures that traders cannot attain a Pareto optimal allocation along any equilibrium play. This result continues to be true even when traders can retrade an infinite number of times: in fact, with a discount factor less than one, all traders stop trading after some finite length of time.

However, we show that there are allocations on the Pareto frontier that can be approximated arbitrarily closely along some equilibrium path of retrading, as the discount factor is close enough to perfect patience and the number of allowable retrading periods is large enough. Once again, this result holds under both of the two extreme assumptions made on the rationality of traders. The same sequence of allocations that approximates a Pareto optimal allocation under myopic retrading can be sustained by a Subgame Perfect Equilibrium profile under far-sighted retrading. The sequence of actions along the corresponding Subgame Perfect Equilibrium path is such that at each new round of trade all traders gain in utility relative to the allocation inherited from the previous round of trade. We are also able to establish, as a corollary, that the approximation result only requires individual traders to condition their choice of current bids and offers only on the prices observed at each trading post in the preceeding round of trade and on their own final allocation from the preceeding round of trade.

For the general case of far-sighted retrading, the set of allocations supported by Subgame Perfect Equilibrium profiles is shown to expand as the cost of reopening trading posts falls. This weak monotonicity result holds with finite as well as infinite horizon.

All the results just described will first be proved under the simplifying assumption that traders can consume commodities (all tradeable) only after having stopped trading. However, we will then show that all results extend to the more general class of games where traders can decide to consume part of their current endowment at any time, while remaining on the market with the rest. Moreover, we will display an interesting version (or interpretation) of the model where the tradeable goods are actually assets. In this case the issue of consumption becomes irrelevant, since the assets owned by each individual at any given time cannot be consumed, they can only be kept or traded. For example, think of traders as being endowed with apple trees and banana trees: trees can be kept or traded, they cannot be eaten. The goods that agents consume can be simply viewed as derived from the yields of the currently owned assets (current production of fruits of the owned trees).

Finally, although we focus on the possibility of eventually reaching an efficient allocation of resources (or assets) through retrading, we point out that a new type of market failure also arises in market games with retrading: there are “bad” Subgame Perfect Equilibria where traders delay trade only because the other traders do the same.

The rest of the paper is organised as follows. The next subsection compares our retrading model and our results with the related literature. The next section presents the economy and the basic models of non-cooperative trade that we study. Section 3 gives a simple example, in which the unique equilibrium path of retrading converges to the competitive equilibrium. Section 4 characterizes the equilibria of the benchmark retrading model with myopic players. Section 5 characterizes the Subgame Perfect Equilibria of the market game with far-sighted retrading. Section 6 contains the extensions to the case where each trader can always choose between consuming and trading any subset of her own commodities and to the case where the tradeable goods are assets.

1.1 Related literature

There is a body of related work that studies dynamic noncooperative games of exchange. Rubinstein and Wolinsky (1990) study repeated trade in a model of noncooperative bargaining with a finite set of traders where traders can choose the partners they trade with. They obtain a multiplicity result that allows them to implement both competitive equilibrium and non-competitive outcomes. Freidman (1984) also studies repeated trade, in a setting with double auctions, and concludes that traders are able to converge to the competitive equilibrium allocation. Both these results contrast with our result that no allocation on the Pareto frontier can be attained even when the cost of reopening trading posts is negligible. The main reason for this contrast lies in the crucial differences between retrading and repeated trade.

Gale (1986a, 1986b, 1987) looks at a model where traders are repeatedly pairwise matched and bargain over the trades that they make with each other. With a continuum of traders, complete information, and endogenous replacement, there is a stationary equilibrium which conveges to the competitive equilibrium as the discount factor converges to one. There are at least three crucial differences between what Gale does and what we do: (1) We have a finite number of traders; (2) Gale's traders make direct transfers to each other in pairs, which are independent of the transfers made within other matched pairs at each round of trade (there is no common price for each commodity); (3) Once a pair agree to trade, they exit and are replaced by identical copies. In this sense, the same set of traders *never really agree to retrade* with each other along the equilibrium path of play. In that framework, retrading refers to the fact that any type has a positive probability of being repeatedly matched with any given other type of trader. In addition, to get to competitive equilibrium, Gale has to restrict attention to stationary equilibria. In contrast, we can construct examples where the competitive equilibrium is approximated by retrading but the corresponding Subgame Perfect Equilibrium profile is not stationary (see the constructive proof of Proposition 4).

Dubey, Sahi and Shubik (1993) is closer to our paper, as they also study retrading in market games. However, they have a model with a continuum of agents who, in addition, do not discount future consumption. They show that if equilibria in the one-shot market game fail to coincide with competitive equilibria due to the endowment constraints in the numeraire commodity binding for non-negligible subsets of traders, competitive equilibria can nevertheless be approximated arbitrarily when traders are allowed to reopen trading posts before they consume their final allocations. In our model, with a finite number of agents, the Nash equilibria of the market game is Pareto inefficient even when endowment constraints in the numeraire commodity don't bind for any individual trader.

Peck and Shell (1990)[2] study a model of a market game where traders can make arbitrarily large short sales, so that net trades are small relative to gross trades. Using this model they show that, at equilibrium, no individual action has a big effect on market prices, and therefore equilibrium allocations approximate competitive equilibrium allocations. Introducing the possibility of arbitrarily large short sales requires traders in their model to satisfy a budget constraint. They postulate some form of outside enforcement of the budget constraint via a bankruptcy rule.[3]

Finally, Postlewaite and Schmeidler (1978) show that Nash equilibrium allocations of the one-shot market game can approximate competitive equilibrium allocations as the number of traders become large. The results we obtain here complement the results they obtain, as we are able to show that, keeping the number of traders fixed but allowing them to retrade, also allows traders to approximate efficient allocations.

2 The Economy

We study trade in pure exchange economies with a finite set of commodities L (indexed by l), a finite set of individuals I (indexed by i). Each individual's consump-

[2]For a related liquidity based approximation result see also Okuno and Schmeidler (1986).

[3]Peck and Shell (1990) refer, in footnote 6, to the possibility of substituting short sales with a dynamic model. However, their suggested dynamic model still requires some form of bankruptcy rule, which we do not need.

tion set is $\Re^L_+$, and his endowment is denoted by $w^i \in \Re^L_{++}$. The utility function is $u^i : \Re^L_+ \to \Re$. A pure exchange economy is $E = \{L, (u^i, w^i) : i \in I\}$. An allocation $x = (x^1, ..., x^I)$ such that $x^i \in \Re \times \Re^{L-1}_+$ for all $i \in I$ is feasible if, in addition, $\sum_{i \in I} x^i = \sum_{i \in I} w^i$. A feasible allocation x is Pareto optimal if there is no other feasible allocation y such that $u^i(y^i) \geq u^i(x^i)$ for all $i \in I$ with $u^i(y^i) > u^i(x^i)$ for some $i \in I$. Let P denote the set of Pareto optimal allocations in E. Let IR denote the set of individually rational allocations x such that $u^i(x) \geq u^i(w)$ for all $i \in I$. For any commodity l, $l = 1, ..., L$, let $w_l > 0$ denote the total endowment of commodity l. Let $F(w)$ denote the set of feasible allocations, i.e., $F(w) \equiv \{x \in \Re^{LI}_+ : \sum_{i \in I} x^i_l = w_l, l = 1, ..., L\}$. As before, P denotes the associated set of Pareto optimal allocations.

2.1 The one-shot market game

In this section we describe the Shapley-Shubik (Shapley and Shubik (1977)) market game of non-cooperative exchange[4] Each trader makes bids and offers of commodities at trading posts where commodities are exchanged; moreover, all bids are denoted in some numeraire commodity, which we set to be commodity 1. Traders are allowed to make offers in all the other commodities $2, ..., L$. There are $L-1$ trading posts for commodities $2, ..., L$. A strategic action for a trader i is a vector $s^i = (b^i_2, ..., b^i_L, q^i_2, ..., q^i_L)$ where b^i_l denotes the bid for commodity l while q^i_l denotes the offer of commodity l, $l = 2, ..., L$. The corresponding set of strategic actions for each trader i is $S^i(w^i) = \{(b^i_2, ..., b^i_L, q^i_2, ..., q^i_L)$ such that $b^i_l \geq 0, \sum_{i \in I} b^i_l \leq w^i_1$, $0 \leq q^i_l \leq w^i_l,\ l = 2, ..., L\}$. All bids and offers have to be non-negative and the offer of a commodity made by a trader cannot *exceed* his endowment of that commodity. For each profile of strategic actions $s = (s^1, ..., s^I)$, at the trading post for commodity l, the aggregate bid is $B_l = \sum_{i \in I} b^i_l$ while the aggregate offer is $Q_l = \sum_{i \in I} q^i_l$. Define the price at the trading post l to be $\pi_l(s) = \frac{B_l}{Q_l}$ if $B_l > 0$, $Q_l > 0$, with $\pi_l(s) = 0$ otherwise. For each trader i, the allocation rule determines commodity holdings as

[4] The reason for choosing this particular type of market game for our analysis is that, as will be clarified below, such a game always gives traders incentives to *retrade*. Hence, it makes sense to study the effects of retrading in this context.

follows: If $\pi_l(s) \neq 0$, $x_1^i(s) = w_1^i - \sum_{l=2}^{L} b_l^i + \sum_{l=2}^{L} q_l^i \pi_l(s)$ and $x_l^i(s) = w_l^i - q_l^i + \frac{b_l^i}{\pi_l(s)}$, $l = 2, ..., L$. If $\pi_l = 0$, $x_l^i(s) = w_l^i$, for all $i \in I$. Let $v^i(s^i, s_{-i})$ be the payoff associated with s. A Nash equilibrium profile of strategic actions is s^* such that $v^i(s^{*i}, s_{-i}^*) \geq v^i(s^i, s_{-i}^*)$, for all $s^i \in S^i(w^i)$ and $i \in I$. Let $N(w)$ denote the set of Nash equilibrium allocations of the market game.

In the one-shot market game with varaible offers observe that the trivial Nash equilibrium where $b_l^{*i} = q_l^{*i} = 0$ for all l and $i \in I$, always exists and yields the initial endowments as the final allocation.[5] What happens when $w \notin P$? Consider the following three properties, which turn out to characterize the set of Nash equilibrium allocations of the one-shot market game:

- **(P1)** *(Static inefficiency)* If $w \notin P$, then $N(w) \cap P = \phi$.
- **(P2)** *(Weak gains from trade)* If $w \notin P$, there exists $x \in N(w)$ such that $u^i(x^i) \geq u^i(w^i)$ for all $i \in I$, with $u^i(x^i) > u^i(w^i)$ for some $i \in I$.
- **(P3)** *(Strong gains from trade)* If $w \notin P$, there exists $x \in N(w)$ such that $u^i(x^i) > u^i(w^i)$ for all $i \in I$.

(P1) requires that whenever the endowments in an exchange economy are Pareto suboptimal, there is no Nash equilibrium allocation that is also Pareto optimal. **(P2)** requires that whenever the endowments in an exchange economy are Pareto suboptimal, there is nevertheless some Nash equilibrium allocation that makes at east one trader better-off relative to his endowments. **(P3)** requires that whenever the endowments in an exchange economy are Pareto suboptimal, there is nevertheless some Nash equilibrium allocation that makes all traders better-off relative to their endowments.

Suppose preferences and endowments satisfy the following regularity assumption:

Assumption 1 *For each $i \in I$, u^i is strictly monotone, strictly-concave, twice continuously differentiable, and the closure of the indifference curves through w^i are*

[5]When $w \in P$ the existence of the trivial Nash equilibrium also guarantees that $N(w) \cap P \neq \phi$, since w would be an element of such an intersection.

contained in $\Re^L_{++}$ and remain bounded away from the boundary of the consumption set.

An important class of economies for which **(P1)**, (**P2)**, **(P3)** characterize the set of static Nash equilibria whenever $w \notin P$, is identified by the following result, due to Dubey and Rogowski (1990) (see also Peck, Shell and Spear (1992) for similar results in a related market game).

Lemma 1 *(Dubey and Rogowski (1990)) Suppose $w \notin P$. If Assumption 1 is satisfied, then, $N(w)$ satisfies (***P1)***, (***P2), (P3)***.*

Proof. As endowments and utility functions satisfy Assumption 1, the existence of an equilibrium point, as defined and used in Theorem 1 in Dubey and Shubik (1978), also implies the existence of a Nash equilibrium where $x_l^{i*} > 0$ for all $l = 1, ..., L$ and $i \in I$. But then the assumptions used in Proposition 1, Remarks 1-5, Section 4 and Section 5.1, in Dubey and Rogowski (1990) are satisfied, and the conclusion immediately follows. **QED.**

In later sections, when we state our results, we will often directly assume that one or all of **(P1),(P2),(P3)** characterize $N(w)$ (respectively, $N_f(w)$) whenever $w \notin P$, without invoking Assumption 1.

2.2 The market game with retrading

Given Lemma 1, there are gains from trade that are not exhausted after a one-period exchange. Therefore there are always *incentives to retrade*. In this section we describe an exchange mechanism that takes into account these incentives. Now, trading posts can reopen over a sequence of finite or infinite periods, $t = 0, 1, ..., T$. At each t an action for trader i is a vector s_t^i. The corresponding set of strategic actions at t for each trader i is $S_t^i(x_{t-1}^i)$, starting from $s_{-1}^i = (0, ..., 0)$ for all $i \in I$ and $x_{l,-1}^i = w_l^i$, for all $l = 1, ..., L$ and for all $i \in I$. For each strategic action profile s_t, in the trading post for commodity l, the aggregate bid is $B_{l,t}$ while the aggregate offer is $Q_{l,t}$, with the corresponding price $\pi_{l,t}(s_t)$, defined as in the static game. For each trader i, the allocations $x_t^i(s_t)$ are also defined as before. Along a sequence of action profiles $s = \{s_0, ..., s_t, ...\}$, we say that player i *stops trading* after period $\tilde{T}^i$ iff $b^i_{l,t'} = q^i_{l,t'} = 0$ for all $t' \geq \tilde{T}^i$, $l = 2, ..., L$. Even though traders can stop trading at different times, it is convenient not to complicate notation by explicitly keeping track of traders who drop out. We can do so without loss of generality as the bids and offers of a trader can be zero at any round of trade and hence a trader i who stops trading at some period $\tilde{T}^i$ can be counted as a market player who makes zero bids and offers in all periods including and subsequent to $\tilde{T}^i$.

In what follows we shall consider two models of retrading, labelled as **myopic** and **far-sighted**.

Case 1 (**Myopic retrading**): When retrading is *myopic*, at each new round of retrading traders behave in a very simple way: at each new round of retrading, they choose a vector of bids and offers that constitutes a static Nash equilibrium to the final allocation obtained from the previous round of trade. In the notation developed before, at each t, the strategy profile chosen, s_t, satisfies the condition that $x_t(s_t) \in N(x_{t-1})$. Traders consume when they stop trading. As the utility function of each trader is continuous and the set of feasible allocations compact, we remark that even when $\tilde{T}^i = \infty$, the payoff to any player i remains well-defined. We study the properties of the stationary allocations of this retrading process. Myopic traders can be seen as traders who do not expect that trading posts can be reopened, so they play their best responses as if the current trading round were the last. Consistent with this, we will study myopic retrading without discounting, even though the results extend to the case where discounting occurs.

Case 2 (**Far-sighted retrading**): When retrading is *far-sighted*, all traders know that future play will, in general, be conditioned on the outcomes of the current round of trade. Here, as before, we assume that an individual trader consumes only when she has stopped trading. However, now we endow each trader i with a common discount factor δ. When T is finite, δ lies in $[0, 1]$. When $T = \infty$, δ lies in $[0, 1)$.[6] Trader i's payoff, once she has stopped trading in period $\tilde{T}^i$, is $\delta^{\tilde{T}^i} u^i(x^i_{\tilde{T}^i})$. A history of play at period t is $h_t = \{s_0, ..., s_{t-1}\}$. The corresponding set of histories

[6] We interpret δ as the cost of reopening trading posts in any new round of trade.

is denoted by H_t. A pure strategy for trader i is a sequence $\sigma^i = \{\sigma_0^i, ..., \sigma_t^i, ...\}$ with $\sigma_t^i : H_t \to S_t^i$ for all t. Denote by $\sigma^i|h_t$ the restriction of σ^i to the subgame from period t after history h_t. A pure strategy profile $\sigma = (\sigma^1, ..., \sigma^I)$ is a Subgame Perfect Equilibrium (SPE henceforth) if, for every h_t, the restriction $\sigma^i|h_t$ for all traders $i \in I$ is a Nash equilibrium in the subgame from period t. Let $\tilde{X}(\delta, w, T)$ denote the set of SPE allocations of the market game with far-sighted retrading, where trading posts are allowed to be reopened up to $T + 1$ times.

3 An example

In this section, we analyze retrading in an example. There are two individuals and two commodities. Both individuals have quasi-linear utility functions with $u^i(x) = x_1^i + f^i(x_2^i)$ (and similarly for j). We assume that $f^k(.)$ $(k = i, j)$ is strictly monotone, strictly concave, twice-continuously differentiable and satisfies the boundary condition that $\lim_{x_2 \to 0} \partial f^k(x_2) = \infty$ for $k = i, j$. Further, for simplicity, we choose the units in which commodities are measured so that $\sum_k w_2^k = 1$. We focus on retrading in the "sell-all" market game. The "sell-all" version of the Shapley-Shubik market game is obviously simpler than the variable offers version: At each time t where trader i is still active, his offer is assumed to equal $x_{2,t-1}^i$, which is the endowment of commodity 2 inherited from the trades of the previous period. Other than for this simplification, the strategies, aggregate variables, and the allocation rules are identical to the more general variable offers model described before.[7] In this case there is a unique Nash equilibrium with trade in the one-shot market game (the no trade equlibrium does not exist).[8] This means that finitely repeated trade would not add anything, while we now show that finite retrading leads the traders

[7] One additional difference would be in the precise definition of what it means to stop trading in the sell-all market game. We avoid the formal definitions since they are not relevant for this example, but the intuitive feature of any such definition is that traders must bid the exact amounts that give them back the endowments obtained with their last real trade.

[8] Endowments and utility functions satisfy Assumption 1, hence the existence of an equilibrium with trade follows from Dubey (1980), Remark 2. Further, using Remark 5 in Dubey and Rogowski (1990), it also follows that if $w \notin P$ then $N_f(w)$ satisfies **(P1)**, **(P2)**, **(P3)**.

towards the competitive allocation *even* if they are myopic.

It will be convenient to refer to $w_2^i = \alpha_0^i \in (0,1)$ as individual i's initial share of commodity 2 and α_t^i as individual i's share at the end of round $t-1$ of retrading. In what follows, we will make an assumption that each individual is endowed with enough of commodity 1 to ensure that the endowment constraint of the numeraire commodity does not bind. For the moment, we simply assume that at any round of trade, all traders have enough of the numeraire commodity to ensure existence of an interior one-shot Nash equilibrium in any one round of trade[9]. Using the allocation rule, we obtain that at any round of retrading t, $t = 0, 1, \ldots$, if the current profile of actions is $s_t = (b_{2,t}^j, b_{2,t}^i)$, player $i's$ objective function at time t is

$$x_{1,t}^i - b_t^i + \alpha_t^i B_t + f^i\left(\frac{b_t^i}{B_t}\right)$$

where $B_t = b_{2,t}^i + b_{2,t}^j$. Using the fact that the ratio $\frac{b_t^i}{B_t} = \alpha_{t+1}^i$, if the current profile of actions s_t is an interior Nash equilibrium, we can rewrite the first-order conditions of traders to obtain the dynamical system that characterizes the evolution of the sequence of allocations generated by myopic retrading:

$$\frac{\partial f^i\left(\alpha_{t+1}^i\right)\left(1-\alpha_{t+1}^i\right)}{\partial f^j\left(1-\alpha_{t+1}^i\right)\alpha_{t+1}^i} = \frac{\left(1-\alpha_t^i\right)}{\alpha_t^i}$$

Evidently, a stationary point of the preceeding map is an interior allocation on the Pareto frontier. Moreover, as both individuals have quasi-linear utility functions, the allocations of commodity 2 is uniquely determined at an interior Pareto optimum. Let $\bar{\alpha}^i$ denote individual $i's$ share of commodity 2 at the interior Pareto optimum. Suppose $\alpha_0^i < \bar{\alpha}^i$. Then, as $f^k(.)$ is strictly concave, we must have that $\frac{\partial f^i\left(\alpha_0^i\right)}{\partial f^j\left(1-\alpha_0^i\right)} > 1$. Moreover, $\frac{\left(1-\alpha_0^i\right)}{\alpha_0^i} > \frac{\left(1-\bar{\alpha}^i\right)}{\bar{\alpha}^i}$. For all $t > 0$ such that $\frac{\partial f^i\left(\alpha_t^i\right)}{\partial f^j\left(1-\alpha_t^i\right)} > 1$, $\frac{\left(1-\bar{\alpha}^i\right)}{\bar{\alpha}^i} < \frac{\left(1-\alpha_t^i\right)}{\alpha_t^i} < \frac{\left(1-\alpha_{t-1}^i\right)}{\alpha_{t-1}^i}$. If there exists $\hat{t}$ such that $\frac{\partial f^i\left(\alpha_{\hat{t}}^i\right)}{\partial f^j\left(1-\alpha_{\hat{t}}^i\right)} < 1$, $\frac{\left(1-\bar{\alpha}^i\right)}{\bar{\alpha}^i} > \frac{\left(1-\alpha_{\hat{t}}^i\right)}{\alpha_{\hat{t}}^i}$ and as

[9] Formally, if retrading can take place over T periods, at any T' consider the sum $\sum_{t=0}^{T'-1}\left(\frac{1-\alpha_{t+1}^i\left(1-\alpha_{t+1}^i\right)}{\alpha_t^i}\right)\frac{\alpha_t^i}{1-\alpha_t^i}\partial f^i\left(\alpha_{t+1}^i\right)$. As $\lim_{x_2\to 0}\partial f^i(x_2) = \infty$ and $w_2^i > 0$ for $i \in I$, without loss of generality we can assume that α_t is an element of some compact set bounded away from zero and one, independent of t. Let $K^{i,T'}$ denote the maximum of the above sum taken over α_t in this compact set. We assume that $w_1^k - \max_{T\geq T'\geq 1} K^{k,T'} > 0$.

long as for all $t > \hat{t}$, $\frac{\partial f^i(\alpha^i_{\hat{t}})}{\partial f^j(1-\alpha^i_{\hat{t}})} < 1$, we must have that $\frac{(1-\bar{\alpha}^i)}{\bar{\alpha}^i} > \frac{(1-\alpha^i_t)}{\alpha^i_t} > \frac{(1-\alpha^i_{t-1})}{\alpha^i_{t-1}}$. Suppose there exists $\tilde{t} > \hat{t}$ such that $\frac{\partial f^i(\alpha^i_{\tilde{t}})}{\partial f^j(1-\alpha^i_{\tilde{t}})} > 1$. Consider the ratio

$$\frac{\partial f^i\left(\alpha^i_{\tilde{t}}\right)...\partial f^i\left(\alpha^i_{\hat{t}}\right)}{\partial f^j\left(1-\alpha^i_{\tilde{t}}\right)...\partial f^j\left(1-\alpha^i_{\hat{t}}\right)}.$$

Remark that if the above ratio is equal to one we must be on the Pareto frontier. On the other hand the above ratio must be strictly greater than one as otherwise $\frac{1-\alpha^i_{\tilde{t}}}{\alpha^i_{\tilde{t}}} > \frac{(1-\bar{\alpha}^i)}{\bar{\alpha}^i}$ a contradiction. Therefore, $\frac{\partial f^i(\alpha^i_{\tilde{t}})...\partial f^i(\alpha^i_{\hat{t}})}{\partial f^j(1-\alpha^i_{\tilde{t}})...\partial f^j(1-\alpha^i_{\hat{t}})} > 1$, which implies that $\frac{(1-\bar{\alpha}^i)}{\bar{\alpha}^i} < \frac{1-\alpha^i_{\tilde{t}}}{\alpha^i_{\tilde{t}}} < \frac{\left(1-\alpha^i_{\tilde{t}-1}\right)}{\alpha^i_{\tilde{t}-1}}$. By repeating the above argument from $\tilde{t}$, it follows that that ratio $\frac{(1-\alpha^i_t)}{\alpha^i_t}$ converges to $\frac{(1-\bar{\alpha}^i)}{\bar{\alpha}^i}$ and therefore, α^i_t to $\bar{\alpha}^i$. A symmetric argument establishes convergence when $\alpha^i_0 > \bar{\alpha}^i$. An immediate consequence is that the sequence of allocations generated by myopic retrading must converge to the Pareto frontier. Moreover, note that from the equations that determine final allocations, we also obtain that individuals consumption of commodity 1 is identical to that at the competitive equilibrium.

What about far-sighted retrading? We show that the sequence of allocations generated by myopic retrading can be supported as SPE outcomes when T is very large but finite. Consider the sequence of allocations $y_1, ..., y_t, ...$, with $y_0 = w$, associated with myopic retrading. Remark that $y_t = N_f(y_{t-1})$,[10] $t = 1, ...$, with the associated sequence of payoffs $u(y_1), ..., u(y_t), ...$ in utility space $\Re^2$. Consider the following strategy profile $\tilde{\sigma}$. For $t \leq \underline{T} + 1$, play $\tilde{s}_t$ such that $y^i_t = x^i(\tilde{s}_t)$ (and $\tilde{u}^i_t = u^i(x^i(\tilde{s}_t))$) as long as $\tilde{h}_t = \{\tilde{s}_0, ..., \tilde{s}_{t-1}\}$; otherwise, if there has been a deviation, play s'_t such that $x(s'_t) = N_f(x_{t-1})$, for all $t \leq T$. We need to show that $\tilde{\sigma}$ is a SPE. By construction, after any deviation, both players continue to choose bids according to one-shot Nash equilibria. As all sequences of allocations generated by one-shot Nash equilibria converge to the same allocation for both commodities, no player has an incentive to deviate when T is large.

[10]The subscript f refers to "fixed" offers, since in this example offers are not strategic.

Finally, it is worth noting that this example generalizes to more goods and individuals with quasi-linear utility functions.

4 Myopic retrading

We start the general analysis with myopic retrading. We show that, starting from an arbitrary configuration of initial endowments, traders are able to converge to some allocation in the Pareto set. Nevertheless, convergence cannot take place in a finite number of rounds of myopic retrading. We state the results only for the market game with variable offers described in Section 2.2, but we note that the same results also obtain for the "sell-all" market game (as one could guess from the previous section).

The following definition identifies the sequences of allocations that are consisitent with myopic retrading.

Definition of Myopic Retrading: A sequence of allocations $\{x_t\}$, $t = 1, ...$ is generated by myopic retrading if and only if it satisfies the inclusion $x_t \in N(x_{t-1})$, for all $t \geq 1$.[11]

Some notation is needed before proving the main result of this section. For any allocation y, let $u(y) = (u^1(y^1), ..., u^I(y^I))$. For any $K \subset \Re^{LI}$, let $u(K) = \{u(y) : y \in K\}$. Observe that $u(K) \subset \Re^I$, for all K. Let $\|.\|$ denote the Euclidian norm. Then, we define the distance between a vector y and a set K as $d(u(y), u(K)) \equiv \inf_{\hat{u} \in u(K)} \|u(y) - \hat{u}\|$.

Proposition 1 *Suppose $N(w)$ satisfies* **(P1)-(P2)** *whenever $w \notin P$. For any initial endowment vector $w = y_0 \in \Re^{LI}_{++}$, there exists a sequence of allocations $\{y_t\}$, $t = 0, 1, ...$, $y_t \in N(y_{t-1})$ for all $t \geq 1$,* such that*, for any $\varepsilon > 0$, there is a $T > 0$ such that for all $t > T$ $d(u(y_t), u(P \cap IR)) < \varepsilon$.*

Proof. If $w \in P$, then $w \in N(w)$ and we are done. Therefore assume that $w \notin P$. Consider the sequence of sets $N_1, ..., N_t, ...$, with $y_0 = w$, and $N_t = \{x :$

[11] x_0 is obviously the initial endowment w.

$x \in N(y),\ \textit{for some } y \in N_{t-1}\}$, $t = 1, ...$, with the associated sequence of sets $u(N_1), ..., u(N_t), ...$ in utility space $\Re^I$. By **(P2)**, we can extract a sequence $\tilde{u}_t$, $t = 0, 1, ...$ such that $\tilde{u}_t \in u(N_t)$ and $\tilde{u}_{t+1} > \tilde{u}_t$, at each t, with $y_0, ..., y_t, ...$ the associated sequence of allocations. Note that for each $i \in I$, the sequence $\tilde{u}^i_t$, $t = 0, 1, ...$ is bounded above, as the utility of each individual is continuous and the set of feasible allocations is compact. Let $\bar{u}^i$ denote the supremum of the sequence $\tilde{u}^i_t$, $t = 0, 1, ...$ As every increasing sequence converges to the supremum, it follows that the sequence $\tilde{u}_t$, $t = 0, 1, ...$, converges to $\bar{u} = (\bar{u}^1, ..., \bar{u}^I)$. Moreover, by passing to subsequence if necessary, without loss of generality, we may assume that the associated sequence of allocations $y_t, t = 0, 1, ...$ converges to some allocation $\bar{y}$ such that $u(\bar{y}) = \bar{u}$. We claim that $\bar{y} \in P$. Suppose, to the contrary, $\bar{y} \notin P$. Then, by **(P2)**, there exists an allocation $y' \in N(\bar{y})$ and $i \in I$ such that $u^i(y'^i) > u^i(\bar{y}^i)$. As the utility for each individual is continuous, by the upper semicontnuity of the Nash correspondence in the profile of initial endowments, there exists $\tilde{\varepsilon} > 0$ such that for all $\|\hat{y} - \bar{y}\| < \tilde{\varepsilon}$, if $\hat{y}' \in N(\hat{y})$, $u^i(\hat{y}'^i) > u^i(\bar{y}^i)$. As the sequence of allocations $y_t, t = 0, 1, ...$ converges, by construction, to $\bar{y}$, it follows that there exists $\tilde{T}$ such that for all $t > \tilde{T}$, $\|y_t - \bar{y}\| < \tilde{\varepsilon}$, a contradiction, as for each $i \in I$, $\bar{u}^i$ is the supremum of the sequence $\tilde{u}^i_t$, $t = 0, 1, ...$ But, then, it follows that for every $\varepsilon > 0$, there exists $\bar{T}$ such that for all $t > \bar{T}$, $d(\tilde{u}_t, u(P \cap IR)) < \varepsilon$. **QED.**

The above proposition demonstrates that myopic retrading eventually allows traders to approximate some allocation on the Pareto set. Each profile of strategic actions along the myopic retrading sequence constitutes a static Nash equilibrium to the allocation inherited from the preceeding round of trade. By **(P2),** for every configuration of Pareto suboptimal endowments, there is a static Nash equilibrium at which allocations are such that every trader is atleast as well-off and some trader(s) strictly better-off relative to their initial endowments. This implies that the sequence of utility profiles associated with the sequence of allocations generated by myopic retrading is an increasing sequence. But then, along each dimension, corresponding to a specific individual, this sequence of utilities must converge to its supremum, which in turn determines the limit of the sequence of utility profiles generated by myopic retrading. Using the continuity of utility functions, by passing to subsequence if necesary, we can conclude that the limit of the sequence of allocations generated by myopic retrading yields the utility vector that is the limit of the sequence of utilities generated by myopic retrading. It follows that the limit allocation is in the Pareto set, as otherwise there would be a contradiction with the fact that the limit of the sequence of utility profiles generated by myopic retrading is determined by the supremum, along each dimension corresponding to a specific individual. Therefore, some allocation on the Pareto set will be approximated by myopic retrading.

Remark 1 Evidently, the preceeding proposition goes through with the stronger requirement that $N(w)$ satisfies **(P3)** whenever $w \notin P$. In the formal results of this section we keep the Assumption $\delta = 1$ because this is the case where the other assumptions of myopic retrading make the most sense intuitively: in fact, a myopic player who does not discount the future can be assumed to believe he will consume right away. So myopia here means that traders can't predict that after trading they will change their mind, trading again instead of consuming.[12]

Remark 2 Remark that at each stage of myopic retrading, the final allocation from the preceeding round of trade defines the distribution of endowments for a "new" economy. As the sequence of allocations converge to some allocation on the Pareto frontier, in limit, we obtain an economy with Pareto optimal endowments. As no trade is the only outcome at the competitive equilibrium of an economy with Pareto optimal endowments, in this sense the sequence of allocations associted with myopic retrading converges to competitive equilibria of the limit economy as well.

Although Proposition 1 demonstrates that traders will obtain allocations in the vicinity of the Pareto set, it still leaves open the question of whether traders are able to converge to an allocation *on* the Pareto frontier after a *finite* number of rounds of myopic retrading.

[12]The results of this section can easily be extended to $\delta < 1$.

Proposition 2 *If $w \notin P$, there is no $T < \infty$, and no sequence of allocations $\{y_t\}$, $t = 1, ..., y_t \in N(y_{t-1})$, with $y_0 = w$ and $t = 0, ..., T$, such that $y_T \in P$.*

Proof. Given that $y_T \in P$, $y_T = N(y_T)$. Moreover, as $w \notin P$, there must be some $T' < T$ such that the allocation obtained at $T' - 1$, $y_{T'-1}$, is not in P, while for all $t \geq T'$, $y_t \in P$. Then we must have that $y_{T'} \in N(y_{T'-1}) \cap P$: a contradiction with Lemma 1. **QED.**

The intuition behind this result is simple. If trade concludes after some finite length of time, at some finite stage in the game it must be the case that while the traders' inherited allocation from the previous period is Pareto suboptimal, the final allocation they obtain after reopening trading posts is both (a) a Pareto optimal allocation, and (b) satisfies the inequalities for a Nash equilibrium allocation for the one-shot market game with the traders' inherited allocation as the endowment. But by **(P1),** with Pareto suboptimal endowments, no Nash equilibrium allocation of the one-shot market game can ever be Pareto optimal. This guarantees that no allocation *on* the Pareto set will be attained by traders after a *finite* number of rounds of myopic retrading. Without discounting, this implies that trading posts will *always* be reopened. This makes the assumption of myopic traders hard to swollow, but the next section shows that not only the results above extend to far-sighted behavior, but also that far-sighted behavior becomes indistinguishible from myopic behavior over the process of retrading.

5 Far-sighted retrading

In this section, we extend the analysis of myopic retrading to the study of the Subgame Perfect Equilibria (SPE) with far-sighted retrading. The results of the previous section are confirmed, and we obtain some additional characterization results on equilibrium behavior. We also point out that a new kind of market failure emerges with far-sighted retrading: traders may delay trade along a SPE path merely because they expect other traders to do the same. Finally, we show that as traders become more patient, the set of SPE allocations expands.

Observe that allowing for far-sighted retrading doesn't get rid of the no trade outcome, which is also a feature of the one-shot market game: there is always a SPE where $b^i_{l,t} = q^i_{l,t} = 0$ for all $t = 0, 1, ..., l = 2, ..., L$, and $i \in I$.[13] In this section the discount factor can take any value, and we keep the assumption that individual consumption occurs only after a player has stopped trading. This assumption is only made for simplicity, and it is not crucial for the results, as discussed in Section 6.

The next proposition and its corollary, provide a negative result, which strengthen the result of Proposition 2.

Proposition 3 *If $w \notin P$, $\tilde{X}(\delta, w, T) \cap P = \phi$, for all $\delta \in [0, 1]$, and all $T < \infty$.*

Proof. Let $\bar{T} \leq T$ be the first period at which an allocation $x_{\bar{T}} \in P$ is obtained along some SPE path. Given that trade cannot take place after reaching the Pareto set, it must be the case that traders stop trading at some $\bar{T} \leq T$, i.e., $b^i_{l,t} = q^i_{l,t} = 0 \ \forall i, \forall l, \forall t > \bar{T}$. Moreover, since $\bar{T}$ is the first period where p is reached, $x_{\bar{T}-1} \notin P$. As $x_{\bar{T}-1} \notin P$, by **(P1),** $N(x_{\bar{T}-1}) \cap P = \phi$. This is a contradiction, since, at the last round of trade, any SPE profile requires the final allocation to be in the set of Nash equilibrium allocations with respect to the inherited allocation. **QED.**

Corollary 1 *If $w \notin P$, $\tilde{X}(\delta, w, \infty) \cap P = \phi$, for all $\delta \in [0, 1)$.*

Proof. When $\delta \in [0, 1)$, any trader gets a payoff of zero if he trades indefinitely. Therefore, along any SPE path, all traders will stop trading after some finite length of time, implying that there exists a $\bar{T} < \infty$ such that $b^i_{l,t} = q^i_{l,t} = 0$ for all $t \geq \bar{T}$, $l = 2, ..., L$. Trade stops before $\bar{T}' = \inf_T \{T : b^i_{l,t} = q^i_{l,t} = 0 \ for \ all \ t \geq T, \ l = 2, ..., L, \ i \in I\}$. Then, the proof immediately follows from Proposition 3. **QED.**

The above proposition and its corollary show that, with varaible offers, traders cannot obtain allocations on the Pareto set. As trade always concludes after some finite length of time, at some finite stage in the game, it must be the case that

[13] In the sell-all version analyzed in the example the unique equilibrium has trade, but with variable offers the no trade equilibrium is always a possibility.

both the traders' inherited allocation from the previous period and the final one are Pareto suboptimal, otherwise there would be a contradiction with **(P1)**.

The above results leave open the possibility that traders may nevertheless obtain a final allocation that is arbitrarily close to the Pareto frontier with far sighted retrading. To this end, we find it convenient to prove the following lemma, which provides a useful characterization of the set of SPE strategy profiles. We use the following notation. For any SPE strategy σ, let $s(\sigma) = (s_1(\sigma), ..., s_t(\sigma), ...)$ denote the equilibrium path of play associated with σ. Let $y' \in \tilde{X}(\delta', w, T)$ when $\delta = \delta'$. By assumption, there exists a SPE strategy, denoted by $\sigma(\delta', y')$, which yields $y' \in \tilde{X}(\delta', w, T)$. Further, as at any SPE all traders stop trading after some finite period of time, there exits a final time period $T(\sigma(\delta', y'))$, after which no trader trades. Consider the strategy profile $\hat{\sigma}(\delta', y')$ which is identical to $\sigma(\delta', y')$ on the equilibrium path but differs off the equilibrium path in that, after any deviation from the equilibrium path of play at some time $t < T(\sigma(\delta', y'))$, $b^i_{t'} = q^i_{t'} = 0$, $i \in I$, for all $t' > t$. Let $\hat{\Sigma}$ denote the corresponding set of strategies.

Lemma 2 *For any $T < \infty$, for all $\delta \in [0, 1]$, $y' \in \tilde{X}(\delta', w, T)$ if and only if there is a $\hat{\sigma}(\delta', y') \in \hat{\Sigma}$ that supports y'.*

Proof. When $\delta = 0$, all traders stop trading at $t = 0$, implying that $x_0 \in N(w)$. It follows that any SPE strategy profile must be an element of $\hat{\Sigma}$. Suppose $\delta \in (0, 1]$. If there is a $\hat{\sigma}(\delta', y') \in \hat{\Sigma}$ that supports y', by definition $y' \in \tilde{X}(\delta', w, T)$. Next, suppose that $\sigma(\delta', y')$ is a SPE strategy profile that yields $y' \in \tilde{X}(\delta', w, T)$. Then, $\hat{\sigma}(\delta', y')$ is also a SPE strategy that yields $y' \in \tilde{X}(\delta', w, T)$. By construction, observe that no player has an incentive to deviate after $T(\sigma(\delta', y')) + 1$ or after observing a deviation from the equilibrium path of play. Therefore, suppose player i deviates at t choosing some action s'^i_t. As $b^i_{\bar{t}} = q^i_{\bar{t}} = 0$, $i \in I$, for all $\bar{t} > t$, denote i's maximum payoff from such a deviation by $u^{d,i}_t(\delta') = (\delta')^{t+1} u^i(x^i(s'^i_t, s_{-i,t}(\sigma(\delta', y')))$ where $x^i(s'^i_t, s_{-i,t}(\sigma(\delta', y')))$ be the resulting allocation for i when he chooses s'^i_t while all other players choose according to $\hat{\sigma}(\delta', y')$. Observe that as $\sigma(\delta', y')$ is itself a SPE, it must be the case that i's maximum payoff from deviating from the equilibrium path of play under the strategy profile $\sigma(\delta', y')$ cannot be less than

$u_t^{d,i}(\delta')$. Therefore, if player i has no incentive to deviate from the equilibrium path of play under $\sigma(\delta', y')$, she cannot have an incentive to deviate from the equilibrium path of play under $\hat{\sigma}(\delta', y')$. **QED.**

Lemma 2 shows that for the purposes of characterizing SPE strategy profiles, we can, without loss of generality, restrict attention to only those strategy profiles where all traders stop trading after observing a deviation from the equilibrium path of play. Intuitively, the reason for this is that the worst punishment that can be imposed on a deviating trader is for all traders not to trade in all subsequent periods. It may be that the sequence of allocations along a SPE path of play can be sustained by a strategy profile where punishments are weaker in that non-deviating traders continue to trade even after they observe that a trader has deviated in a previous round. As shown in the previous sections, punishments are certainly not necessary for the convergence result. But the same sequence of allocations can be sustained by a strategy profile which is identical to the preceeding strategy profile along the equilibrium path of play but punishes any current deviation by shutting down trading posts in all subsequent periods.

As a corollary of this lemma, it is also possible to prove the same property for the set of SPE when traders can retrade indefinitely.

Corollary 2 *For all $\delta \in [0, 1)$, $y' \in \tilde{X}(\delta', w, \infty)$ if and only if there is a $\hat{\sigma}(\delta', y') \in \hat{\Sigma}$ that supports y'.*

Proof. When $\delta \in [0, 1)$, any trader gets a payoff of zero if he trades indefinitely. Therefore, along any SPE path, all traders will stop trading after some finite length of time, implying that $T(\sigma(\delta', y')) < \infty$. Thus, the proof immediately follows from Lemma 2. **QED.**

We are now in a position to extend the approximation result obtained under myopic retrading to this world of far-sighted, and potentially impatient, players.

Proposition 4 *Suppose $N(w)$ satisfies* **(P1)-(P3)** *whenever $w \notin P$. For every $\varepsilon > 0$, there is a $\underline{T}$ and $\underline{\delta}$ and $y \in \tilde{X}(\delta, w, T)$ such that $d(u(y), u(P)) < \varepsilon$ for all $\delta \in [\underline{\delta}, 1]$, $T \geq \underline{T}$.*

Proof. Consider the sequence of sets $N_1, ..., N_t, ...$, with $y_0 = w$, and $N_t = \{x : x \in N(y),\ for\ some\ y \in N_{t-1}\}$, $t = 1, ...$, with the associated sequence of sets $u(N_1), ..., u(N_t), ...$ in utility space $\Re^I$. By **(P3)** it is possible to extract a sequence $\tilde{u}_t$, $t = 0, 1, ...$ such that $\tilde{u}_t \in u(N_t)$ and $\tilde{u}_{t+1} > \tilde{u}_t$ for each t, with $y_0, ..., y_t, ...$ the associated sequence of allocations. Using Proposition 1, we obtain that for every $\varepsilon > 0$, there exists $\underline{T}$ such that for all $t > \underline{T}$, $d(\tilde{u}_t, u(P \cap IR)) < \varepsilon$. Now construct the following strategy profile $\tilde{\sigma}$. For $t \leq \underline{T} + 1$, play $\tilde{s}_t$ such that $y_t^i = x^i(\tilde{s}_t)$ (and $\tilde{u}_t^i = u^i(x^i(\tilde{s}_t)))$ as long as $\tilde{h}_t = \{\tilde{s}_0, ..., \tilde{s}_{t-1}\}$; otherwise, if there has been a deviation, play $b_{\bar{t}}^i = q_{\bar{t}}^i = 0$, $i \in I$, for all $\bar{t} > t$. Finally, when $t > \underline{T} + 1$, play $b_t^i = q_t^i = 0$. To complete the proof, we need to show that $\tilde{\sigma}$ is a SPE. By construction, observe that no player has an incentive to deviate after $\underline{T}+1$ or in any subgame following a deviation from the SPE path. It remains to check that no player has an incentive to deviate at any $t \leq \underline{T} + 1$. Indeed, consider player i who deviates at t choosing some action $s_t^{\prime i}$. As $b_{t'}^i = q_{t'}^i = 0$, $i \in I$, for all $t' > t$, denote i's maximum payoff from such a deviation by $\delta^{t+1} v^i(s_t^{\prime i}, \tilde{s}_{-i,t})$, where $x^i(s_t^{\prime i}, \tilde{s}_{-i,t})$ is the resulting allocation for i when i chooses $s_t^{\prime i}$ while all other players choose according to $\tilde{\sigma}$. On the other hand, his payoff from continuing to choose according to $\tilde{\sigma}$ is $\delta^{T^i(\tilde{\sigma})} u^i(y)$. As $y_t \in N(y_{t-1})$, we must have

$$v^i(s_t^{\prime i}, \tilde{s}_{-i,t}) \leq u^i(y_t^i) < u^i(y_{T^i(\tilde{\sigma})}^i)$$

Consider $\delta^{T^i(\tilde{\sigma})} u^i(y_{T^i(\tilde{\sigma})}^i) - \delta^{t+1} v^i(s_t^{\prime i}, \tilde{s}_{-i,t}) = \delta^{t+1}[\delta^{T^i(\tilde{\sigma})-t-1} u^i(y_{T^i(\tilde{\sigma})}^i) - v^i(s_t^{\prime i}, \tilde{s}_{-i,t})]$. Let δ_i^{t+1} be such that $[\delta^{T^i(\tilde{\sigma})-t-1} u^i(y_{T^i(\tilde{\sigma})}^i) - v^i(s_t^{\prime i}, \tilde{s}_{-i,t})] = 0$. Set $\underline{\delta} = \inf_{i, 0 \leq t \leq T^i(\tilde{\sigma})-1} \delta_i^{t+1}$. It follows that for all $\delta \in [\underline{\delta}, 1)$, $\tilde{\sigma}$ is a SPE strategy profile. **QED.**

The proof of Proposition 4 selects a sequence of allocations generated by myopic retrading that converges to the Pareto set and shows that such a sequence of allocations can be supported as SPE outcome with far-sighted retrading when traders are sufficiently patient.[14] In order to explain the main idea of the proof, it is convenient to initially restrict attention to the case where $\delta = 1$. The strategy profile is constructed so that traders continue to choose the bids and offers that implement the sequence of allocations generated by myopic trade. If a trader deviates at some

[14]Note that the corresponding SPE profile does not need to be stationary.

round of trade, in all subsequent rounds of trade all traders make null bids and offers at the trading post for each commodity, thus ensuring that no trade is the outcome. In the no trade phase, no individual trader has an incentive to deviate. As no other trader is making positive bids and offers on any trading post, whatever an individual trader does makes no difference to her final allocation. Given this, no individual trader has an incentive to deviate from the sequence of bids and offers that implement the sequence of allocations generated by myopic retrading. This is because the bids and offers at any round of trade constitute a static Nash equilibrium to the final allocation from the previous round of trade, which implies no individual trader can gain by deviating, as a deviation will be followed by no trade in all subsequent rounds. Under **(P3)**, all traders *strictly* gain in utility along some sequence of allocations generated by myopic retrading. This implies that if traders are sufficiently patient, they will prefer to retrade over consuming their current allocation.

Remark that in order to obtain the above approximation result, strategies *need not* to be conditioned on the observability of individual deviations. As stated, Proposition 4 implicitly assumes that each trader chooses strategies in the market game with retrading that are conditioned on the entire history of play. However, the following corollary shows that Proposition 4 goes through even when strategies are conditioned only on a subset of the entire history of play. Denote by σ_M a strategy profile where each player i conditions his choice of bids and offers in period t, (b_t^i, q_t^i), only on the preceeding period's *aggregate* bids and offers, $B_{l,t-1}$ and $Q_{l,t-1}$, $l = 2, ..., L$ (and therefore on the preceeding period's market price vector $\pi_{t-1}(s_{t-1})$), and on her own individual allocation $x_{t-1}^i(s_{t-1})$. Let $\tilde{X}_M(\delta, w, T)$ the set of SPE allocations for strategy profiles in Σ_M.

Corollary 3 *For every $\varepsilon > 0$, there is a $\underline{T}$ and $\underline{\delta}$ and $y \in \tilde{X}_M(\delta, w, T)$ such that $d(u(y), u(P \cap IR)) < \varepsilon$ for all $\delta \in [\underline{\delta}, 1]$, $T \geq \underline{T}$.*

Proof. It is sufficient to observe that the sequence of allocations along the SPE path, $y_0, ..., y_t, ...$, used in the proof of Proposition 4 can also be supported by a strategy profile $\tilde{\sigma}_M$ specified as follows. For $t \leq \underline{T}$, play $\tilde{s}_t$ as long as $B_{t-1} = \tilde{B}_{t-1}$ and $Q_{t-1} = \tilde{Q}_{t-1}$; otherwise, if there has been a deviation, play $b_{\tilde{t}}^i = q_{\tilde{t}}^i = 0$, $i \in I$, for all $\bar{t} > t$. Finally, when $t > \underline{T}$, play $b_t^i = q_t^i = 0$. It is immediate that $\sigma_M \in \Sigma_M$ is also a SPE strategy profile. **QED.**

The fact that the approximation result holds using strategy profiles in Σ_M where players can condition current bids and offers only on aggregate variables, not only shows that conditioning on a specific individual's deviations *is never needed*, but it also suggests that proofs based on keeping track of individual deviations *may not be viable*. In particular, if it is the case that whenever trading posts are open there are always at least three traders on each market, it is possible to show that no deviation from an equilibrium profile can ever be attributed to a specific individual, if the only observable history is that of past aggregate variables. This highlights the minimal nature of the information used in obtaining the approximation result in Proposition 4.

Are there other SPE strategy profiles, which donot require players to implement allocations generated by myopic retrading but which, nevertheless, supports a sequence of allocations that approximate the Pareto frontier? While the answer is generally yes, the next proposition shows that any SPE strategy profile must, after some length of time, begin to look like a strategy profile that implements allocations genertaed by myopic retrading. Formally, for any SPE strategy profile σ, let $y_1(\sigma), ..., y_t(\sigma), ...$ where $y_t = x(s_t(\sigma))$ denote the allocations generated along the equilibrium path of play associated with σ. We say that a SPE strategy profile σ approximates the Pareto frontier if for $\varepsilon > 0$ there exits $\underline{T}$, $\underline{\delta}$ such that $y_T(\sigma) \in \tilde{X}(\delta, w, T)$, $d(u(y_T(\sigma)), u(P)) < \varepsilon$ for all $\delta \in [\underline{\delta}, 1]$, $T \geq \underline{T}$. Moreover, for any $\epsilon > 0$, and $w \in R_{++}^{LI}$, let $N\epsilon(w)$ denote the set of non-trivial $\epsilon-$Nash equilibrium allocations[15].

Proposition 5 *For any SPE strategy profile σ which approximates the Pareto frontier, for every ϵ, there is a $\tilde{T}$ such that for each $t > \tilde{T}$, $y_{t+1}(\sigma) \in N_{\epsilon_t}(y_t(\sigma))$.*

Proof. Consider the sequence of strategies along the equilibrium path of play of σ, $s_1(\sigma), ..., s_t(\sigma), ...,$. At any t such that $y_t(\sigma) \notin N\epsilon(y_{t-1}(\sigma))$, there is some

[15] A non-trivial $\epsilon-$Nash equilibtium allocation x satisfies the condition that there is a profile of strategies s' with $x^i = x^i(s')$ such that for all $i \in I$, $u^i(x^i(s')) \geq u^i(x^i(s^i, s'_{-i})) - \epsilon$ for all $s^i \in S^i(w)$.

player i whose maximum payoff from a deviation, denoted by $v^i(s^i_t, s_{-i,t}(\sigma))$, where $x^i(s^i_t, \tilde{s}_{-i,t}(\sigma))$ is the resulting allocation for i when she chooses s^i_t while all other players choose according to σ, is such that $v^i(s^i_t, s_{-i,t}(\sigma)) - u^i(y^i_t(\sigma)) > 0$. By choosing $b^i_{t'} = q^i_{t'} = 0$, for all $t' > t$, player i can obtain a payoff $\delta^{t+1} v^i(s^i_t, s_{-i,t}(\sigma))$. As σ is SPE, it follows that $\delta^{t+1} v^i(s^i_t, s_{-i,t}(\sigma)) \leq \delta^{\tilde{T}+1} u^i(y^i_{t'}(\sigma))$ for all $\delta \in [\underline{\delta}, 1]$ and $t' > t$ and therefore, $u^i(y^i_{t'}(\sigma)) > v^i(s^i_t, s_{-i,t}(\sigma)) > u^i(y^i_t(\sigma))$. As σ approximates the Pareto frontier, for every $\epsilon > 0$, there exists $\tilde{T}$ such that if $t > \tilde{T}$ and $t' > t$, $u^i(y^i_{t'}(\sigma)) - u^i(y^i_t(\sigma)) < \epsilon$ and therefore, $v^i(s^i_t, s_{-i,t}(\sigma)) - u^i(y^i_t(\sigma)) < \epsilon$ which implies that $y_{t+1}(\sigma) \in N_{\epsilon_t}(y_t(\sigma))$. **QED.**

When at some t players donot choose bids and offers according to myopic retrading, along they obtain an allocation $y_t \notin N(y_{t-1})$ where y_{t-1} is the allocation obtained from $t-1$. This implies that there will be a some individual i who will gain in utitity by deviating from the SPE strategy profile at t and then choosing $b^i_{t'} = q^i_{t'} = 0$ for all $t' > t$. Therefore, if $\{y_t : t \geq 0\}$ is generated along some SPE path of play, it must be the case that the gain in utility for i in the continuation game along the SPE path of play from $t+1$ outweighs the gain in utility from deviating at t. As individuals approach the Pareto frontier along a SPE, remark that gain in utility, for any individual player, in the continuation game along the SPE path of play becomes smaller so must the gain in utility by deviating from the equilibrium path of play.

Does far-sighted retrading always lead to gains in efficiency? The following remark shows that this is not necessarily the case. Proposition 4 shows the possibility of efficiency gains through retrading, but the Pareto set can be approximated only in terms of final allocation, whereas discounting makes the convergence process itself "inefficiently long" in terms of utility. Moreover, a *new* type of market failure can also arise: there are SPE where traders will delay trade only because all other traders do the same.

Remark 3 By Lemma 1, we know that there always exists a static Nash equilibrium in the one-shot variable-offers market game where all traders gain relative to to the no-trade equilibrium. Denote the bid-offer profiles that constitutes a Nash equilibrium with trade $s^* = (b^*, q^*)$. Now suppose that traders are allowed to retrade in an extra round of trade. Consider the following strategy profile $\tilde{\sigma}$: (1) for all $i \in I$, play $s^i_{0,l} = (b^i_{0,l}, q^i_{0,l}) = (0,0)$ for all $l = 2, ..., L$; (2) if $s^i_{0,l} = (b^i_{0,l}, q^i_{0,l}) = (0,0)$ for all $l = 2, ..., L$, and all $i \in I$, play $s = (b^*, q^*)$ next round; otherwise, play $b^i_1 = q^i_1 = 0$, for all $i \in I$. Then, for each $\delta \in (\frac{u^{\bar{i}}(w^{\bar{i}})}{u^{\bar{i}}(x^{\bar{i}}(s^*))}, 1]$, where $\bar{i} = \arg\max_{i \in I} \left\{ \frac{u^i(w^i)}{u^i(x^i(s^*))} \right\}$, $\tilde{\sigma}$ is a SPE. However, observe that for $\delta \in (\frac{u^{\bar{i}}(w^{\bar{i}})}{u^{\bar{i}}(x^{\bar{i}}(s^*))}, 1)$, at $\tilde{\sigma}$, all traders obtain payoffs which are Pareto dominated by their payoffs corresponding to the static Nash equilibrium.

The next result shows that the set of SPE allocations with far-sighted retrading expands as δ becomes larger.

Proposition 6 *Consider $\delta', \delta'' \in [0,1]$ such that $\delta' \leq \delta''$. For each $T < \infty$, then, $\tilde{X}(\delta', w, T) \subseteq \tilde{X}(\delta'', w, T)$.*

Proof. We show that if $y' \in \tilde{X}(\delta', w, T)$, then $y' \in \tilde{X}(\delta'', w, T)$. By Lemma 2, without loss of generality, we assume that any $y' \in \tilde{X}(\delta', w, T)$ is supported by a SPE strategy profile $\hat{\sigma}(\delta', y')$, which has the property that after any deviation from the equilibrium path of play at some period t, all traders stop trading. We need to show that $\tilde{\sigma}(\delta', y')$ remains a SPE strategy profile when $\delta = \delta''$. For each i, let $\tilde{T}_i$ denote the final period when $s^i_t(\tilde{\sigma}(\delta', y')) \neq 0$. Then we must have, at each $t \leq \tilde{T}_i$, $u^i(x^i(s_t(\tilde{\sigma}(\delta', y'))) \leq (\delta')^{\tilde{T}_i - t} u^i(x^i(s_{\tilde{T}_i}(\tilde{\sigma}(\delta', y')))$ and for all $t' > t$, $t' \leq \tilde{T}_i$, $u^i(x^i(s'_{i,t}, s_{-i,t}(\tilde{\sigma}(\delta', y'))) \leq (\delta')^{t'-t} u^i(x^i(s_{t'}(\tilde{\sigma}(\delta', y')))$. Finally, note that as $\delta'' > \delta'$, the above inequalities continue to hold when δ' is replaced by δ''. **QED.**

Proposition 6 shows that as δ becomes larger, the set of allocations supported by SPE expands as the cost of reopening trading posts in a new round of trade falls. By Lemma 2 we can restrict attention to strategy profiles where all trading posts shut down in all subsequent periods after a deviation from the equilibrium path of play is observed in the current period. Thus, when all traders use these strategies, any allocation that satisfies the inequlities that characterize the sequnece of allocations along the SPE path for a specific δ must continue to do so as δ becomes larger.

The following corollary extends this characterization to the case where traders can retrade infinitely often, if they so wish.

Corollary 4 *Consider* $\delta', \delta'' \in [0,1)$ *such that* $\delta' \leq \delta''$. *Then,* $\tilde{X}(\delta', w, \infty) \subseteq \tilde{X}(\delta'', w, \infty)$.

Proof. When $\delta \in [0,1)$, any trader gets a payoff of zero if he trades indefinitely; therefore, along any SPE path, all traders will stop trading after some finite length of time, implying that $T(\sigma(\delta', y')) < \infty$. But, then, the proof is an immediate consequence of Proposition 6. **QED.**

6 Discussion on consumption and asset trading

Throughout the paper we have made the simplifying assumption that consumption by trader i may occur only after he has stopped trading. A natural question to ask is whether therefore our results hold when individual traders can decide otherwise, i.e., when they can opt to consume part of their current endowment instead of using it all for trading purposes. We will divide the analysis of this issue in two parts. First, we give a direct answer to this question, keeping the assumption that all tradeable goods are also consumable. The second part of the analysis makes an argument that in fact one of the best interpretations of our model ought to be the case where the tradeable goods on the trading posts are assets, which are long-lived, yield consumption indirectly, but are not directly consumable themselves. In the second case, of course, the consumption issue becomes irrelevant.

Let us start by keeping the assumption that all tradeables are consumables. Clearly, when the discount factor is equal to one it cannot make a difference, and the SPE profiles that approximate points on the Pareto set remain SPE profiles even when consumption is in principle allowed at any time. On the other hand, when the discount factor is in the open interval $(0,1)$, individuals will typically have an incentive to consume (part of) their endowments even before leaving the market. An important observation is that along a SPE path, the bids and offers typically do not exhaust the endowments at any round of trade. In other words, considering a sequence of actions $s_0, ..., s_t, ...$ that constitutes a SPE path of far-sighted retrading, it is typically the case that $q_t^i < x_{t-1}^i$ at all times.[16] Therefore, individuals can consume a small fraction of their current endowments at each new round of trade and not necessarily affect the SPE profile of retrading. Hence, it is obvious that allowing traders to consume whenever they want the final utilities must be higher. But what matters here is that if δ is high enough the same path $s_0, ..., s_t, ...$ can remain an equilibrium path of retrading. The equilibrium path used in the approximation results is such that everybody is made better off by each successive round of trade, and hence, for δ high enough, the utility difference can always compensate for the longer wait to consume. A deviation to consume current endowments that affects the feasibility of bids and offers at the current or subsequent rounds of trade will not be profitable.

Even though it should be clear from the above discussion that our assumption of "consumption at the end" is irrelevant for the main results, it is also worth noting that this consumption issue would not even be raised if the trading posts were just markets for assets. Let $x = (x_1, ..., x_L)$ be reinterpreted as an allocation of assets. For any x^i, let $y^i = (y_1^i, ..., y_M^i)$ be the associated allocation of commodities.[17] Let $v^i(y)$ represent trader $i's$ preferences over the commodity bundle y. Traders are endowed with assets but not commodities. A feasible allocation of assets generates a feasible allocation of commodities. An allocation of assets is Pareto optimal if and only if the associated allocation of commodities is Pareto optimal. Traders trade assets $2, ..., L$ using asset $l = 1$ as numeraire. For simplicity, we assume that traders cannot trade commodities directly. They can only trade commodities indirectly, by trading assets. The retrading process, both myopic and far-sighted, is as in the previous sections. The difference is that now at each round of trade, if x_t^i is trader $i's$ current allocation of assets, then y_t^i is trader $i's$ current commodity bundle, which

[16] Recall that we are looking at environments in which no trader has a shortage of tradeable goods.

[17] As a metaphore, think of the allocation of assets as being allocation of trees, and the vector y would be the corresponding allocation of fruits. People consume fruits, not trees, but trade trees only in this interpretation of the model.

he consumes to obtain a current utility of $v^i(y_t^i)$. Then, trader $i's$ total utility from retrading will be $\sum_{t=0}^{T} \delta^t v^i(y_t^i)$.

With this specification, all our previuous results apply by appropriately rephrasing the propositions and proofs. After all players have stopped trading, the final allocation of assets will keep giving the same consumption bundle to all traders thereafter every period. If we extended the model to allow for stochastic yields of assets, then asset trading could continue forever, since every shock on the productivity of assets may change the incentives (or needs) of traders to readjust their asset portfolio. Issues related to uncertainty and/or asymmetric information are however beyond the objective of this paper.

7 Conclusion

The main result of this paper has been to show that allowing retrading in markets where the one-shot allocations are inefficient allows traders to approximate allocations on the Pareto frontier arbitrarily closely.

This "approximation" result, however, needs to be qualified on the following grounds: (1) allocations on the Pareto frontier are never attained in finite time by retrading; (2) getting to an allocation close to the Pareto frontier may take several rounds of retrading and therefore, when traders discount future consumption heavily, in payoff space traders may stilll be far away from the Pareto frontier of utilities; (3) there is a huge multiplicity of equilibria with retrading, and therefore not all Subgame Perfect Equilibrium allocations with retrading are close to the Pareto frontier; (4) the possibility of retrading introduces new types of market failure, as now traders might delay or fragment trade along a Subgame Perfect Equilibrium path; (5) in other contexts (see for instance Jehiel and Moldovanu (1999)), where there are externalities in consumption and traders use trading mechanisms which allow some subset of traders to be excluded from the market, retrading may not approximate allocations on the Pareto frontier.

Beside the issue of efficiency, this paper has also demonstrated some interesting "behavioral" properties of retrading processes. In particular, we have shown that the set of equilibrium paths of retrading that converge to the Pareto frontier when agents are forward looking shrinks towards the converging path of myopic retrading. We have also shown by example that convergence holds even when there is a unique Nash equilibrium in the one-shot game, i.e., in a context where finitely repeated trade could not have efficient equilibrium outcomes. The properties of retrading that we have studied seem therefore to be quite general, and independent on the assumptions made on the rationality of traders.

References

[1] Benoit, J.P. and V. Krishna (1985), "Finitely repeated games", *Econometrica*, 53, pp.890-904.

[2] Dubey, P. and M. Shubik (1978),"The non-cooperative equilibria of a closed trading economy with market supply and bidding strategies", *Journal of Economic Theory*, 17, pp.1-20.

[3] Dubey, P. and J.D. Rogowski (1990), "Inefficiency of smooth market mechanisms", *Journal of Mathematical Economics*, 19, pp.285–304.

[4] Dubey, P., S.Sahi and M.Shubik (1993), "Repeated trade and the velocity of money", *Journal of Mathematical Economics*, 22, pp.125–137.

[5] Friedman, D. (1984): "On the efficiency of double auction markets", *American Economic Review*, 74, 60-72.

[6] Fudeberg, D. and E. Maskin (1986), "The folk theorem in repeated games with discounting or with incomplete information", *Econometrica*, 54, pp.533-556.

[7] Gale, D. (1986a), "Bargaining and Competition Part I: Characterization," *Econometrica* **54**, 785-806.

[8] Gale, D. (1986b) "Bargaining and Competition Part II: Existence," *Econometrica* **54**, 807-818.

[9] Gale, D. (1987), "Limit Theorems for markets with sequential bargaining," *Journal of Economic Theory* **43**, 20-54.

[10] Okuno, M., and D. Schmeidler, (1986), "An allocation rule using linear excess demand functions", in *Contributions to Mathematical Economics: In Honor of Gerard Debreu*, (W.Hildenbrand and A. Mas-Collel, Eds.), pp. 361-385, Amsterdam: North Holland.

[11] Peck, J. and K.Shell (1990), "Liquid markets and competition", *Games and Economic Behavior*, 2, pp.362–377.

[12] Peck, J. and K. Shell (1991), "Market uncertainty: correlated equilibria and sunspot equilibria", *Review of Economic Studies*, 58, pp.1011-1029.

[13] Peck,J., K.Shell, S.E. Spear (1992), "The market game: existence and structure of equilibrium", *Journal of Mathematical Economics*, 21, pp.271–299.

[14] Posltlewaite, A and D. Schmeidler (1978), "Approximate efficiency of non-Walrasian Nash equilibria", *Econometrica*, 46, pp.127–135.

[15] Rubinstein A. and A. Wolinsky (1990), "Decentralized trading, strategic behaviour and the Walrasian outcome", *Review of Economic Studies*, 57, pp.63-78.

[16] Sahi, S. and S. Yao (1989), "The non-cooperative equilibria of a trading economy with complete markets and consistent prices", *Journal of Mathematical Economics*, 18, pp.325-346.

[17] Shapley, L.S. and M. Shubik (1977), "Trade using one-commodity as a means of payment", *Journal of Political Economy*, 85, pp.937-968.