Learning Competitive Equilibrium

Sean Crockett
Social and Information Sciences Laboratory
California Institute of Technology

Stephen Spear
Tepper School of Business
Carnegie Mellon University

Shyam Sunder
School of Management
Yale University

July 2006

Abstract

We consider a pure exchange economy repeated for an indefinite number of periods from a fixed endowment and posit a learning rule which directs convergence to competitive equilibrium. In each period trade converges to an allocation in the contract set, where agents interpret the current (common) normalized utility gradient as a vector of prices to determine the implied wealth redistribution relative to their endowments. Agents who are less wealthy at the new allocation are designated *subsidizers*, and demand to provide smaller subsidies in subsequent periods of economic activity. Our model is a globally stable alternative to Walras' *tâtonnement*.

Keywords: Learning, general equilibrium, stability, Scarf's example

Journal of Economic Literature Codes: C63, C68, D44, D51, D58, D61.

1 Introduction

The neo-classical synthesis in economics known as the extended Arrow-Debreu-McKenzie (ADM) model occupies a central position in contemporary economics. Modern macroeconomics relies on general equilibrium theory even more than conventional microeconomics does. It is a workhorse in the analysis of financial markets, and plays an increasingly important role in the analysis of contracts, and in the understanding of economic forces in organizations.

While these applications of the ADM model have been welcomed by most economists, the model has not been immune to criticism. Simon [1978] argued that the ADM model and its descendents are overly mathematical and the underlying assumptions on agents' information processing abilities are unrealistic, and, as a result, the models lack predictive power. Kirman [1989] surmised that two fundamental problems have plagued general equilibrium theory from its earliest development: The theory seemingly imposes only the mildest empirical restrictions on aggregate excess demand, and there is no plausible theory of how economies attain competitive equilibrium and whether the equilibrium is stable in any reasonable sense. These blemishes have led some scholars to question the validity of the basic assumptions underlying the fundamental model. The computability of equilibrium, the assumption of maximizing behavior, and the possibility that out-of-equilibrium behavior may have important effects have been three important foci of such questions.

In this paper, we address the question of equilibrium attainment. The problems facing the theory here are old and well-known. Indeed, the problem was addressed explicitly by Léon Walras, who introduced the idea that attaining equilibrium would involve a process by which markets groped their way toward equilibrium (which he called the *tâtonnement* process), with the help of a fictitious auctioneer who announced prices, then collected orders (or demands) from consumers as to how much of each good they would wish to purchase at the announced prices. If the demand in a market exceeded the supply, the price was adjusted upward. If the supply exceeded demand, the price was adjusted downward. Walras reasoned that this procedure would cause the economy to eventually settle into equilibrium.

Until the 1960's, this intuitively plausible mechanism for attaining equilibrium was widely accepted (even though there were no formal proofs of its validity). Scarf [1960] famously demonstrated the existence of an open set of economies having a unique equilibrium which was unstable under the Walrasian tâtonnement. This surprising result spawned a substantial literature on stability which came to the conclusion that it was always possible to construct a tâtonnement procedure of the form $\dot{p} = H_\xi [z(p)]$ specific to a given economy ξ for which some competitive equilibrium would be stable under the generalized tâtonnement procedure (see, for example, Smale [1976]). Unfortunately, actually constructing the mechanism requires information not only about prices, but about higher-order derivatives of all agents' utility functions in order to coordinate the rates at which different market prices converge to their equilibrium values. Since such a mechanism requires the collection of information which is inherently private to agents, it is highly implausible. Hence, the results from this literature can only be interpreted as strengthening the negative implications of Scarf's example.

There is an alternative restriction that guarantees the stability of the tâtonnement process. When preferences and endowments are such that all goods are gross substitutes, the Walrasian tâtonnement procedure works. A sufficient condition for the gross substitutes property to hold is that the initial endowments of agents are sufficiently close to being Pareto optimal. However, there is no reason to believe that endowments must always reside in the so-called Pareto component. When endowments are not near Pareto optimal, there are many instances in which the gross substitutes property does not hold. This well-known result was examined by Hirota [1981] in the context of Scarf's example, where it was shown that the set of endowments for which the tâtonnement process is stable is large, in the sense that random draws from all possible distributions of endowment holdings leads to stable tâtonnement roughly 80% of the time. While these results are intriguing, we do not believe they should be interpreted as saying that instability due to strong income effects doesn't matter. It is easy to think of real situations – Texas or Saudi Arabia – where large communities earn the bulk of their income from the sale of a single commodity, and hence, face significant income effects when the price of this commodity fluctuates. If we back away from the abstraction of exchange economies and take account of the fact that most people depend on their labor skills/human capital investments for their income, and that these tend to be difficult or impossible to modify in response to changes in labor demand, we again have a situation in which there will be strong income effects associated with the sale of this particular endowment. Thus, in our view, the potential instability of tâtonnement remains a critical weakness of the ADM model.

A second alternative is to model the economy as a large strategic market game of the kind first introduced by Shapley and Shubik [1977], in which prices are determined explicitly from the bids and offers that agents in the economy make on "trading posts" for each of the goods available for trade. In this model, equilibrium prices are determined as the Nash equilibrium of the underlying game in which agents take the bids and offers of other agents as given and choose their own bids and offers as best responses. Hence, in this model, the question of how the economy arrives at an equilibrium turns on the stability or instability of mechanisms for implementing the Nash equilibrium. While the literature on this subject is not as large as that for the ADM formulation, there have been several papers that examine this approach. Chatterji and Ghosal [2004] examine a market game with a continuum of agents (corresponding, therefore, to a perfectly competitive economy) in which agents may trade in two goods. They show that an out-of-equilibrium adjustment mechanism based on rationalizability of observed bids and offers has features that closely resemble those of the Walrasian tâtonnement in the sense that any competitive equilibrium which is stable under the Walrasian procedure will also be stable under their procedure. Of course, the restriction of this result to the case of two commodities limits its usefulness since it is well-known that in this setting, the Walrasian tâtonnement will always converge to some competitive equilibrium.

Kumar and Shubik [2004] demonstrate a mechanism (the Cournot-Shubik mechanism) which generates convergence to competitive equilibrium in Scarf's example adapted to the market game setting, but go on to make the observation via other examples that the convergence properties of a given mechanism in the market game setting depend on the underlying parameters of the economy, an observation which is consistent with the findings in the tâtonnement literature.

There are also a number of loosely related approaches that have been developed in response to the failure of the Walras' tâtonnement. One literature, the so-called non-tâtonnement approach, combines local tâtonnement with rationing and shows that these processes converge to Pareto optimal allocations, but will generally not converge to the competitive equilibrium from the original initial endowment (see Hurwicz, Radner and Reiter [1975] for references to this literature). A second literature on stability properties for processes based on coalition formation in large economies for implementing core allocations in large economies does provide mechanisms which will implement competitive equilibrium (see Bewley [1973] or Green [1974] for early references, or Serrano and Volij [2002] for a recent approach). While the results from this literature are positive in the sense of generating convergence to the competitive equilibrium, the complexity of search required as the number of agents gets large, given that the set of all coalitions grows exponentially in the number of agents, makes it difficult to see how such processes might actually be implemented in the large economy settings to which they are most applicable.

In this paper, we explore an alternative to the "thought experiment" approach of tâtonnement by postulating that economic equilibria, unlike those of physical systems, must be *actively* learned (or discovered) by the agents in the system. This idea has it's genesis in the empirical fact that in the standard double auction supply and demand experiments pioneered by Vernon Smith and Charles Plott (see Smith [1962]), agents tend to coordinate on competitive equilibrium prices but only after numerous repetitions of a sufficiently stationary environment.

The model we adopt in this paper stems from Gode and Sunder [1993], who asked whether the process of finding the equilibrium price and allocation in the experiments requires sophisticated learning, or whether it can be implemented with "zero intelligence" (ZI) search procedures. To analyze this question, they replicated the basic experimental setup using computerized robots. The robot traders in their model generated simple random bids (if they were buyers) or offers (if they were sellers) with the only restriction on behavior being that no bid or offer, if accepted, should make an agent worse off. In simulations of the model, Gode and Sunder found that while prices don't converge to the competitive equilibrium (CE) prices (as they do with human subjects), the infra-marginal prices (i.e. the prices of the last observed transactions) always occur at or near the CE price, while the efficiency of the market is in excess of 90% of the maximum (which occurs when the quantity of the good traded is the CE quantity). These results tell us that the double auction mechanism of the classic supply and demand experiment will implement the competitive equilibrium allocation under very mild conditions on agents' behavior. Cliff and Bruten [1997] have extended the Gode and Sunder results to allow robots to remember prices at which past transactions occurred, and to modify their bids and offers accordingly, and have shown that simulations with these robots replicates the observed results for human traders.

The zero intelligence trading result does not, however, answer the question of whether the competitive paradigm can be implemented easily in environments where many agents trade many goods. In follow-on work, Gode, Spear and Sunder [2004] showed that, at least in the context of a two agent, two good exchange economy, simple ZI search easily finds Pareto optimal equilibria. The random search process does not, however, find the competitive equilibrium. The reason for this is self-evident. The random search process generates a uniform set of random trajectories from the initial endowment to the contract curve. The ending allocations are, therefore, distributed on the contract curve about the average trajectory generated by the search procedure.

The present paper focuses on the question of how much additional "intelligence" or memory is required of agents in the zero-intelligence exchange environment in order to find a competitive equilibrium. The answer turns on the issue of whether agents can price the Pareto optimal allocation they find, in the sense of learning the (common) normalized utility gradient at the optimal allocation. We treat prices as endogenous objects that reflect agents' marginal valuations. While these valuations will be the same for all agents at a Pareto optimum, they need not agree on these prices out of equilibrium. A learning rule is developed by combining directed random search for Pareto optima with an application of the welfare theorems that makes use of the learned gradient.

Since the First Welfare Theorem implies that every competitive equilibrium is a Pareto optimum, the exchange process begins by searching for Pareto optima as candidate CE allocations. The Second Welfare Theorem implies that every Pareto optimum can be supported as a competitive equilibrium after some redistribution of endowments. Rather than redistributing endowments, we posit that agents use the common normalized utility gradient to price the Pareto optimal allocation and the original endowment, in order to determine the associated implied wealth redistribution. Obviously, if the agents have reached a Pareto optimum that is not a competitive equilibrium allocation, some agents (although better off) will have less implied wealth under the current allocation than with their endowments, i.e. these agents will be subsidizing the consumption of other agents at the observed prices. In such a situation, we find it plausible to posit that these agents would aspire to trade to a different allocation, given the opportunity, which required less subsidization on their part, and, presumably, a higher utility payoff associated with the reduced subsidization.[1]

We implement this idea over repeated periods of trading which can be viewed either as actual exchange in a stationary repeated environment, or as fictitious trade as is standard in the tâtonnement literature. In each period, agents move in small random Pareto-improving steps from the endowment allocation to something in the Pareto set. At this new allocation, agents will agree on relative prices and can determine whether or not they are subsidizing other agents. If no one is subsidizing anyone else, we are at a competitive equilibrium. If some agents are providing subsidies,

[1]Consider the fact that if the Pareto optimal "price" vector had governed the system from the outset, when everyone held their endowments, the optimal portfolio demanded by a price-taking subsidizer would have resulted in greater utility (had that demand been fulfilled) than the allocation he actually reached.

a new period of trading occurs from the original endowment, in which agents who have previously provided subsidies only accept trades in which they pay (weakly) smaller subsidies under the prices at which they had most recently subsidized.

The remainder of the paper outlines the theoretical basis for this research, provides a proof of the convergence of a mild strengthening of this learning process to the competitive equilibrium, and reports on computer simulations of a smooth version of Scarf's example.

2 The Model

A finite number of agents trade a finite number of goods and services in a pure exchange economy. We index agents as $i = 1, \ldots, M < \infty$ and goods as $j = 1, \ldots, \ell < \infty$. Preferences for agent i are represented by a utility function $u_i : \mathbb{R}_+^\ell \to \mathbb{R}$, which we assume is smooth, strictly monotonic, strictly quasi-concave, and satisfies the Inada condition. An allocation is defined as an m-tuple of commodity vectors $\mathbf{X} = [x_1, \ldots, x_M]$, where $x_i \in \mathbb{R}_+^\ell$ for $i = 1, \ldots, M$ and x_{ij} is agent i's quantity of good j. The initial endowment allocation is denoted $\Omega = [\omega_1, \ldots, \omega_M]$, where $\omega_i \in \mathbb{R}_{++}^\ell$ for each i. Feasible allocations satisfy

$$\sum_{i=1}^{M} x_i = \sum_{i=1}^{M} \omega_i = R,$$

where R is the vector of total available resources. We let Θ denote the set of feasible allocations; that is, $\Theta \equiv \{\mathbf{X} \in \mathbb{R}_+^{\ell M} : \sum_{i=1}^{M} x_i = R\}$. Let PO be the set of feasible Pareto optima. Let $IR(\mathbf{X}) \equiv \{\mathbf{Y} \in \Theta : u_i(y_i) \geq u_i(x_i) \, \forall i\}$ be the individually rational set at $\mathbf{X}$. Let $CS(\mathbf{X}) \equiv PO \cap IR(\mathbf{X})$ be the contract set at $\mathbf{X}$. The allocation $\mathbf{X}^*$ is a competitive equilibrium if $\mathbf{X}^* \in PO$ and $p(\mathbf{X}^*) \cdot x_i^* = p(\mathbf{X}^*) \cdot \omega_i$ for each i, where $p(\mathbf{X})$ is the (common) normalized utility gradient at $\mathbf{X}$.

In each period $t = 1, 2, 3, \ldots$, we assume that through some process of exchange the economy moves from Ω to $\widehat{\mathbf{X}}^t \in CS(\Omega)$. We impose a learning rule that further restricts the reallocation to occur within a subset of the contract set refined during the previous periods, which we will call $CS^{t-1}(\Omega)$. For any subset of $CS^{t-1}(\Omega)$ with positive Lesbesgue measure, we assume that the probability $\widehat{\mathbf{X}}^t$ is an allocation in this subset is greater than zero. We prove in the next section that $\left\langle \widehat{\mathbf{X}}^t \right\rangle$ converges to a competitive equilibrium. We do not place explicit structure on intermediate allocations in period t, although in Section 4 we develop some examples.

We next introduce variables that will serve to increasingly constrain exchange over time. For $t \geq 1$, let $\Lambda^{t-1} = \left[\lambda_1^{t-1}, \lambda_2^{t-1}, \ldots, \lambda_M^{t-1}\right]$, $P^{t-1} = \left[\rho_1^{t-1}, \rho_2^{t-1}, \ldots, \rho_M^{t-1}\right]$, and $\Psi^{t-1} = \left[\chi_1^{t-1}, \chi_2^{t-1}, \ldots, \chi_M^{t-1}\right]$, where $\lambda_i^{t-1} \in \mathbb{R}_-$, $\rho_i^{t-1} \in \mathbb{R}_{++}^\ell$ and $\chi_i^{t-1} \in \mathbb{R}_+^\ell$ for $i = 1, 2, \ldots, M$. We will refer to the scalar λ_i^{t-1} as the *subsidization constraint* for agent i during period t. The vectors ρ_i^{t-1} and χ_i^{t-1} will be a reference price vector and consumption bundle, respectively, corresponding to the subsidization constraint for i at time t. The learning rule requires that $\widehat{\mathbf{X}}^t$ can be the final allocation adopted in period t only if

$$\rho_i^{t-1} \cdot \widehat{x}_i^t > \lambda_i^{t-1} \; \forall i. \tag{1}$$

To initialize the process, we set $\Psi^0 = \Omega$, $P^0 \gg 0$, and $\lambda_i^0 = -\infty$ for all i, so that $CS^0(\Omega) = CS(\Omega)$.

For $\widehat{\mathbf{X}}^t \in CS^{t-1}(\Omega)$, define agent i's *gain* as

$$g_i\left(\widehat{\mathbf{X}}^t\right) = p\left(\widehat{\mathbf{X}}^t\right) \cdot \left(\widehat{x}_i^t - \omega_i\right),$$

recalling that $p(\mathbf{X})$ is a normalization of the utility gradient for all agents at $\mathbf{X} \in PO$. If $g_i\left(\widehat{\mathbf{X}}^t\right) < 0$, then we say that i has *subsidized* other agents in period t. If no agent is providing subsidies, then the economy must be at a competitive equilibrium, since $g_i\left(\widehat{\mathbf{X}}^t\right) \geq 0$ for all i implies $p\left(\widehat{\mathbf{X}}^t\right) \cdot \sum_{i=1}^M \left(\widehat{x}_i^t - \omega_i\right) \geq 0$. From strict monotonicity and feasibility, respectively, we know that $p\left(\widehat{\mathbf{X}}^t\right) \gg 0$ and $\sum_{i=1}^M \left(\widehat{x}_i^t - \omega_i\right) = 0$, so we can infer that $p\left(\widehat{\mathbf{X}}^t\right) \cdot \left(\widehat{x}_i^t - \omega_i\right) = 0$ for all i. Since $\widehat{\mathbf{X}}^t \in PO$, it must be CE.

If $\widehat{\mathbf{X}}^t$ is not a competitive equilibrium, the constraints Λ^t, P^t, and Ψ^t are updated as follows. For all i such that $g_i^t < 0$, if $g_i^t > \lambda_i^{t-1}$ (that is, if the current loss for i is smaller in absolute value than the loss represented by her constraint), then $\lambda_i^t = g_i^t$, $\rho_i^t = p\left(\widehat{\mathbf{X}}^t\right)$, and $\chi_i^t = \widehat{x}_i^t$. For all other $j \neq i$, $\lambda_j^t = \lambda_j^{t-1}$, $\rho_j^t = \rho_j^{t-1}$, and $\chi_i^t = \chi_i^{t-1}$. The reason we require a loss to be smaller in absolute value than the current subsidization constraint in order for constraints to be updated will be explained in the *No cycling* subsection of the paper. Therefore, at the beginning of period $t+1$, $CS^t(\Omega)$ is the region of the contract set that survives all subsidization constraints.

Some geometric intuition is helpful to motivate this learning process. Suppose i updated her constraint in period t. Then equation (1) requires future allocations to lie in the halfspace above the hyperplane passing through $\widehat{\mathbf{X}}^t$ and orthogonal to the (common) utility gradient at that allocation. Interpreting a normalization of this gradient as a vector of prices, this constraint hyperplane can be viewed as a budget set through $\widehat{\mathbf{X}}^t$ at which the endowment is not affordable. Viewed another way, if the prices on which all agents implicitly agree at the end of period t had been in effect at the beginning of the period, agent i could have afforded to reach a higher indifference surface by waiting until the end of the period to trade (if her relative influence in the market is small). The agent interprets subsidization as a signal that she might have secured better terms of trade (and presumably greater utility) than those she actually realized, and equation (1) reflects a simple heuristic designed to reduce such missed opportunities in the future. For example, the agent might reject a utility-improving trade in a subsequent period that would cross her constraint hyperplane, recalling the precedent in period t and hopeful that terms of trade will again move in a more favorable direction during the period.

We illustrate the geometry of the learning process in an Edgeworth box for a 2×2 economy in Figure 1. Here, at some period-ending allocation $\mathbf{X}$, agent 1 was a subsidizer. We know this because the initial endowment is drawn in the half-space above her constraint hyperplane. The final allocation in the following period must lie in the region of $CS(\Omega)$ above this hyperplane. Now suppose in the following period

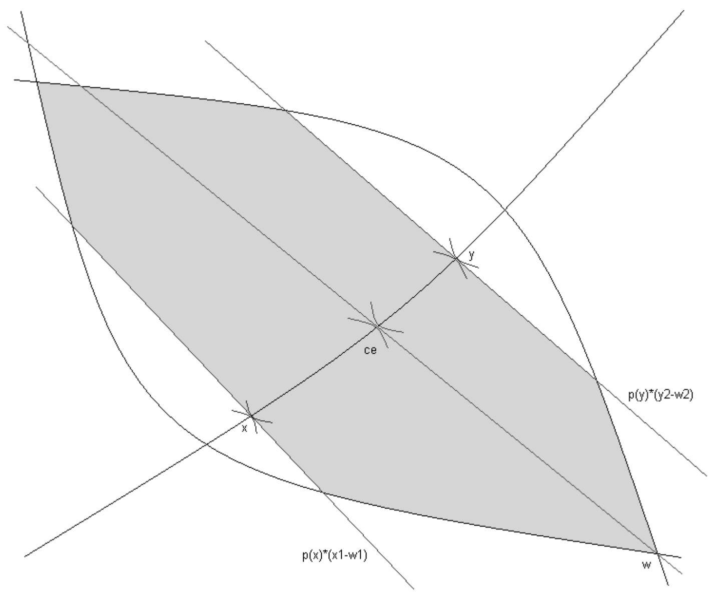

Figure 1: Across-Period Learning in a 2×2 Economy

the allocation $\mathbf{Y}$ is implemented (note that $p(\mathbf{X}) \cdot (y_1 - x_1) > 0$, as required by equation (1)). At this allocation agent 2 incurs a loss, so subsequent allocations must lie above (below relative to agent 1's coordinate axis) her constraint hyperplane. Since agent 1 is a subsidy receiver in the current period, her subsidization constraint remains unchanged. Thus potential reallocations in the following period are restricted to the subset of $CS(\Omega)$ contained in the shaded region of the figure. And so forth. It turns out that in the 2×2 case a large class of decentralized within-period Pareto-improving processes, coupled with the restriction provided in equation (1), will converge to the competitive equilibrium, as demonstrated in a later section of the paper. A common theme of these processes is that all trade must be individually rational and remain above the constraint hyperplanes; with respect to the figure, trade from the endowment must remain in the shaded region.

As drawn, the figure basically assumes that a competitive equilibrium always lies in the region defined by the intersection of the contract set and the subsidization constraints (that is, for all $t \geq 0$, there exists $CE \in CS^t(\Omega)$), so next we turn to the question of convergence.

3 Convergence

Under our standing assumptions on preferences, it is well-known that generically, PO is an $M-1$ dimensional manifold in the space of feasible allocations (see, for example, Smale [1974] or Balasko [1988]). Since this set is bounded while any manifold is closed, it follows that PO is compact. From our specification of restrictions on exchange, one of three things must happen in any period $t \geq 1$: (a) The reallocation in t hits a competitive equilibrium, in which case convergence to

CE has occurred; (b) $CS^{t-1}(\Omega)$ is empty, in which case the learning process has failed to converge to CE; or (c) An allocation in $CS^{t-1}(\Omega)$ is adopted, constraints are updated if necessary, and a new period begins. Obviously, if we attain a competitive equilibrium, we are done.

To show that $CS^{t-1}(\Omega)$ is not empty for $t \geq 2$, assume trade in period $t-1$ resulted in Pareto optimal allocation $\widehat{\mathbf{X}}^{t-1}$, and for $\varepsilon > 0$, consider the ε-ball centered at $\widehat{\mathbf{X}}^{t-1}$. Suppose for the moment that for every $\mathbf{X}$ in this ball, the competitive equilibrium associated with $\mathbf{X}$ is unique and has some agent who is unconstrained at $\widehat{\mathbf{X}}^{t-1}$ but constrained at $\mathbf{X}$. Then for every sequence in the ε-ball converging to $\widehat{\mathbf{X}}^{t-1}$, some unconstrained agent is constrained at the competitive equilibrium associated with each point in the sequence. Thus we can find a subsequence at which there is one agent who is constrained at the associated competitive equilibrium for every point in the subsequence, but is not constrained in the limit. The existence of such a subsequence contradicts the fact that the agent's demand function depends continuously on the endowment and prices, while the equilibrium prices depend continuously on the endowments. The last assertion necessarily obtains for ε sufficiently small, because for any allocation in the Pareto component the competitive equilibrium price vector is unique, and regular in the sense that it depends continuously on the allocation under consideration.

Therefore it must be the case that for sufficiently small ε, the competitive equilibrium associated with every allocation in this ε-ball has all agents who are unconstrained at $\widehat{\mathbf{X}}^{t-1}$ similarly unconstrained at the associated competitive equilibrium. Now consider the allocation in such a ball that gives ε/K of good 1 to each of the K subsidizers at $\widehat{\mathbf{X}}^{t-1}$, and takes $\varepsilon/(M-K)$ of good 1 away from each of the subsidy receivers at $\widehat{\mathbf{X}}^{t-1}$. Then all agents are unconstrained at the competitive equilibrium associated with this allocation, and thus proving the non-emptiness of $CS^{t-1}(\Omega)$.

3.1 No cycling

So, we need to deal with the case where the learning process continues to the next period, and ask if the sequence of allocations generated will have the competitive equilibrium as its limit. We first assert that subsidization constraints cannot cycle, due to the condition that a loss is not adopted as a new subsidization constraint unless it is smaller than the current constraint in absolute value. Suppose this condition were not applied, so that if $g_i^t < 0$, then $\lambda_i^t = g_i^t$ (that is, the most recent period of subsidization would always provide the constraint). While it should be clear from the geometry of the 2×2 economy illustrated in Figure 1 that allocations will not cycle, in higher dimensional settings they may. To see this, consider a 2-good, 3-agent economy and some Pareto optimal allocation in which agent 1 consumes the bundle x^1 (the agent subscript has been omitted for convenience). Given agent 1's consumption, the remaining resources define an Edgeworth box for the other two agents. Since the allocation is in PO, agents 2 and 3 must be consuming bundles at which their marginal rates of substitution are the same as that of agent 1 at x^1.

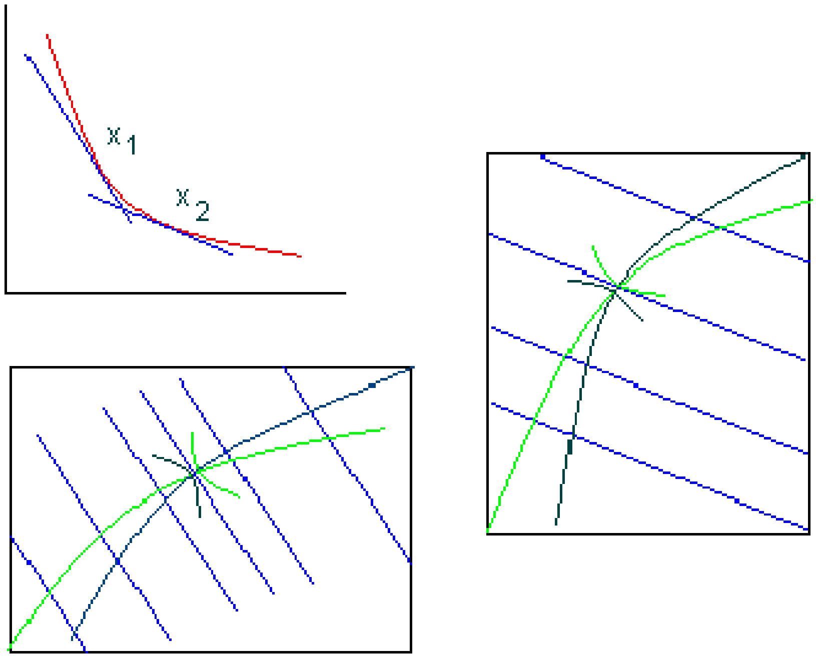

Figure 2: Potential for Cycling

Hence, the allocation for these two agents will lie at the intersection of the income expansion paths in the Edgeworth box for these two agents, as illustrated in Figure 2 (in the box below agent 1's consumption space). Now, consider a perturbation of the allocation which moves agent 1 from x^1 to x^2 in the diagram. Note that at the prices implied by agent 1's normalized utility gradient at x^2, agent 1 can't afford x^1 (and hence would accept x^1 as subsidy-improving from x^2). At x^1, however, the same thing is true: at the prices determined by the normalized gradient at x^1, agent 1 cannot afford x^2 and hence would accept it as subsidy-improving from x^1. Hence, if we can show that the perturbation from x^1 to x^2 can be made in such a way as to remain a Pareto optimum, then it follows that without the added condition on constraint updating, allocations can cycle. To see that this can be done, note that if the perturbation along agent 1's indifference curve is small enough, it will result in a small perturbation to the Edgeworth box in which the allocation of agents 2 and 3 lies, together with a small perturbation in the agents' income expansion paths. This is illustrated (in exaggerated form) in the right-hand Edgeworth box in Figure 2. Given that the agents' expansion paths intersect transversely, this will continue to be the case after a small perturbation, so that the resulting allocation will constitute a Pareto optimum.

Therefore, the extra condition on constraint updating after agent i realizes a loss (that is, $g_i^t > \lambda_i^{t-1}$) requires a comparison of subsidies not at the old prices, but rather at the prices determined by the new Pareto optimum, and clearly guarantees there can be no cycling of allocations. Of course, it is possible that in any given period trade can take place to some Pareto optimal allocation at which no subsidization constraints are updated, because of the monotonicity condition on updating constraints. However, since the Lesbesgue measure of the set of Pareto optimal allocations at which constraints will be updated is easily shown to be greater than zero, with probability one some such allocation will eventually be adopted.

3.2 Convergence to competitive equilibrium

We next prove that final allocations across periods converge to a competitive equilibrium. In looking at this result, it will also become apparent that how constraints are updated is closely related to the proof of existence of competitive equilibrium in Negishi [1960], so that the learning process considered here can be viewed as a behavioral implementation of the Negishi approach.

We need to show that the sequence of allocations adopted across periods converges to competitive equilibrium. Let

$$R_i(y_i) \subset CS(\Omega)$$

be the set of individually rational Pareto optimal allocations such that $p\left(\widehat{\mathbf{X}}\right) \cdot \widehat{z}_i = y_i$, for $\widehat{\mathbf{X}} \in CS(\Omega)$ and $\widehat{z}_i$ representing the vector of net trades for i at $\widehat{\mathbf{X}}$. We will call y_i for which $R_i(y_i)$ is non-empty *admissable* values of y_i. Now, let

$$\widehat{\Delta} = \left\{ p \in \Delta | p = p\left(\widehat{\mathbf{X}}\right) \text{ for all } \widehat{\mathbf{X}} \in CS(\Omega) \right\}.$$

We also let

$$\widehat{R}_i(y_i) \subset \mathbb{R}^{\ell}_{+}$$

be the set

$$\widehat{R}_i(y_i) = \left\{ \begin{array}{c} x \in \mathbb{R}^{\ell}_{+} | x = f_i(p, w_i), \\ w_i = p \cdot \omega_i + y_i, \ p \in \widehat{\Delta} \end{array} \right\}$$

where $f_i(p, w_i)$ is agent i's demand function at prices p and wealth w_i. $\widehat{R}_i(y_i)$ is i's offer surface when her wealth is $p \cdot \omega_i + y_i$. Next, let

$$\bar{R}_i(y_i) = \widehat{R}_i(y_i) \times_{i=1}^{M-1} \mathbb{R}^{\ell}_{+}.$$

Then,

$$R_i(y_i) = \bar{R}_i(y_i) \cap CS(\Omega)$$

Proposition 1: $\bar{R}_i(y_i) \pitchfork$ *PO for generic* y_i.

Proof. Taking y_i variable, the tangent space to $\bar{R}_i(y_i)$ is spanned by the columns of

$$\begin{bmatrix} I & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & K_i - \mu_i z_i^T & \mu_i & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & 0 & \cdots & I \end{bmatrix}$$

where $K_i - \mu_i z_i^T$ is the derivative of agent i's demand function with respect to p (keeping in mind that i's income depends on p through the value of the endowment)

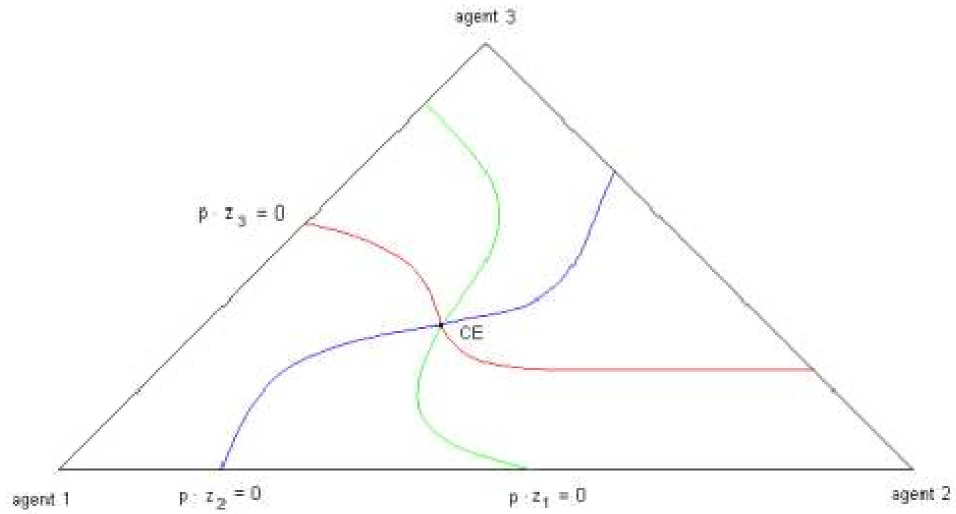

Figure 3: Subsidization Manifolds at Competitive Equilibrium

and μ_i is the derivative of i's demand with respect to y_i. In these expressions, K_i is the substitution matrix, while μ_i is the vector of income effects for i. Now, from Theorem 2.4.1 of Balasko [1988], individual demand functions when considered as functions on relative prices and wealth map $\Delta \times \mathbb{R}$ diffeomorphically (i.e. smoothly bijectively) onto $\mathbb{R}^{\ell}_{+}$. The mapping from $\Delta \to \Delta \times \mathbb{R}$ given by $(p, p \cdot \omega)$ for a fixed endowment vector ω is easily shown to be an immersion. Since the composition of an immersion with a diffeomorphism is again an immersion, it follows that the derivative of $f_i\,(p, p \cdot \omega_i)$ with respect to p has rank $\ell - 1$. Since we know that

$$\left[K_i - \mu_i z_i^T\right] p = 0$$

while $\mu_i \cdot p = 1$, it now follows that the matrix above has rank ℓM. Thus, $\bar{R}_i\,(y_i)$ is transverse to anything as long as y_i is variable. By the transversal density theorem, then, for generic y_i, we will have $\bar{R}_i\,(y_i) \pitchfork PO$ for y_i fixed. ■

Proposition 2: *Let $\langle \Psi^{\tau} \rangle$ be the maximal subsequence of $\langle \Psi^{t} \rangle$ such that $\Psi^{\tau+1} \neq \Psi^{\tau}$ for $\tau \geq 0$. For admissible y_i, let $\breve{R}_i\,(y_i) \equiv \{\boldsymbol{X} \in CS\,(\Omega) \mid R_i\,(x_i) \geq R_i\,(y_i)\}$. Then $\bigcap_{i=1}^{M} \breve{R}_i \left(\chi_i^{\tau+1}\right) \subset \bigcap_{i=1}^{M} \breve{R}_i\,(\chi_i^{\tau})$.*

Proof. We will consider what happens at a critical value of y_i below, but for now, we focus on regular values of y_i. Since the offer surface for any agent is an $\ell - 1$ dimensional (i.e. codimension 1) submanifold of $\mathbb{R}^{\ell}_{+}$ it follows that $\bar{R}_i\,(y_i)$ is codimension 1 in $\mathbb{R}^{\ell M}_{+}$, and hence, via the codimension formulas for transverse intersections, $R_i\,(y_i)$ will be a codimension 1 submanifold of $CS\,(\Omega)$ with boundary (making the usual identification of PO with the $M - 1$ dimensional simplex) for regular, admissible y_i. Hence, we have that for every non-critical y_i, $R_i\,(y_i)$ separates $CS\,(\Omega)$ locally into two half spaces. On one side, allocations are such that $p(\widehat{\mathbf{X}}) \cdot z_i < y_i$, while on the other $p(\widehat{\mathbf{X}}) \cdot z_i > y_i$. Figure 3 shows the $R_i\,(y_i)$ sets for $y_i = 0$ for a 3-agent economy having a unique competitive equilibrium. Figure 4 shows the y_i

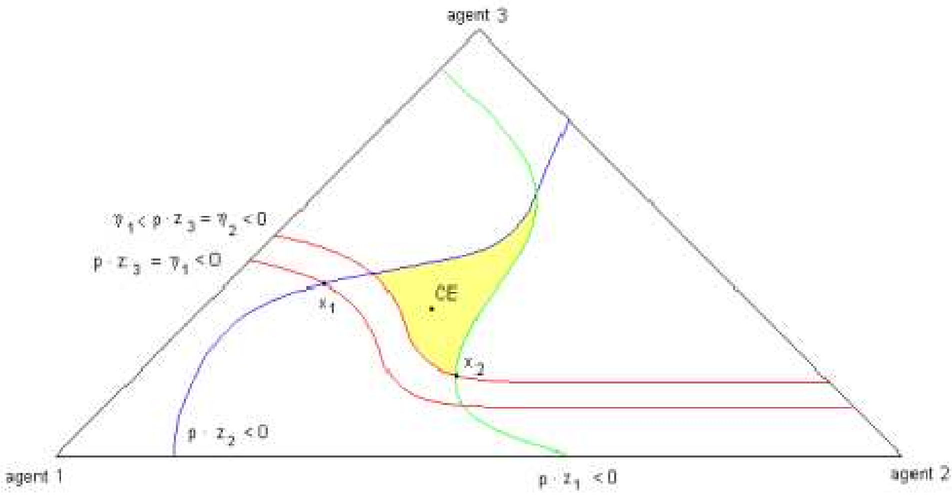

Figure 4: Subsidization Manifolds are Nested

sets corresponding to two different allocations at which different pairs of agents in the same 3-agent economy are subsidizing the remaining agent. The yellow shaded region shows the set of acceptable remaining allocations in the set of Pareto optima.

To deal with the issue of critical y_i, note that these can occur when the wealth level associated with a set of allocations reaches a relative maximum or minimum. In this case, we will find two branches of the manifold corresponding to y_i converging on each other and merging when y_i reaches the critical value. The presence of critical wealth values doesn't alter the functioning of the learning process, although it can end up generating discontinuous changes in the set of PO allocations a subsidizing agent will accept. We illustrate this in Figure 5. In the top diagram, the yellow shaded region shows the set of acceptable allocations in the space of Pareto optima for agent 3 given that she is subsidizing at the level y_i. The lower diagram indicates that as we move in a northeasterly direction in the Pareto set, the value of the allocations first increases to a local maximum, then decreases to a local minimum before increasing again to become positive. The yellow region in the first diagram corresponds to the regions above the line at y_i in the lower diagram. Note that if we were to increase y_i continuously until it was above the value at the local maximum in the lower diagram, we would see the left region in the upper diagram shrink and then disappear. For our purposes, the critical feature of the $R_i(y_i)$ manifolds is that they partition the set of Pareto optima into disjoint regions in which allocations either have larger or smaller value than y_i.

Recall that in period $t+1$, χ_i^t is the portfolio that generated agent i's current subsidization constraint (or $\chi_i^t = \omega_i$ if no binding subsidization constraint has been formed by period t for agent i). Since these constraint portfolios are only updated when the current loss is less negative than the constraint loss, it must be the case that for $\tau \geq 0$, $R_i\left(\chi_i^{\tau+1}\right) \geq R_i\left(\chi_i^{\tau}\right)$ for all i, with this relationship holding with strict inequality for at least one agent. But then $\bigcap_{i=1}^{M} \breve{R}_i\left(\chi_i^{\tau+1}\right) \bigcap CS(\Omega) \subset \bigcap_{i=1}^{M} \breve{R}_i\left(\chi_i^{\tau}\right) \bigcap CS(\Omega)$. ■

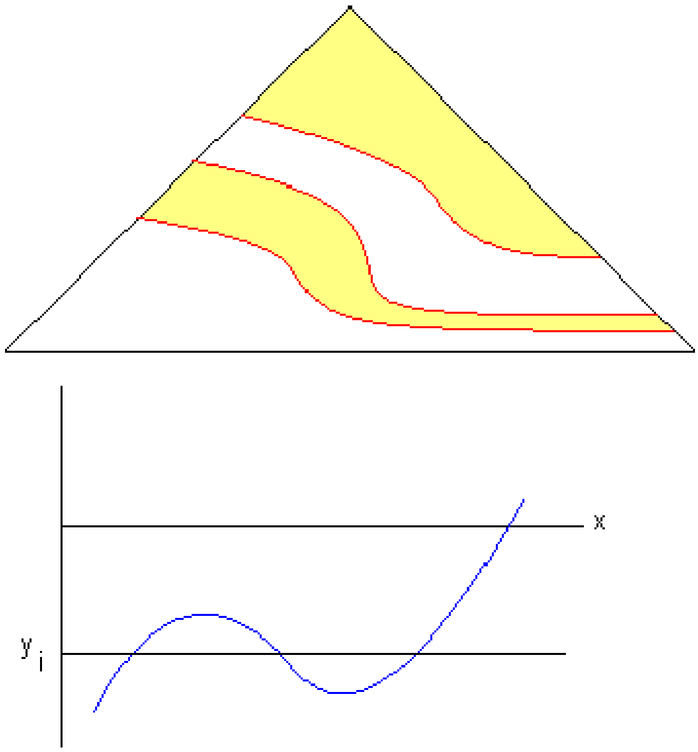

Figure 5: Critical Subsidization Levels

Given this nestedness feature, we can show that exchange must converge to competitive equilibrium as follows. Consider the sequence of period-ending allocations that lie in the contract set and meet the restriction in equation (1) in all periods. As before, there are two possibilities. If exchange hits the competitive equilibrium in finite time, we are obviously done. So assume an infinite sequence of allocations is generated. Since the Pareto set is compact, the sequence clusters. Since we know the sequence can't cycle, the cluster points must be limit points. Consider the allocation for an arbitrary agent i and form the sequence

$$\sigma_i = \left[p\left(\widehat{\mathbf{X}}^t\right) \cdot \widehat{z}_i^t \right]_{t=1}^{\infty}.$$

Split the sequence into non-negative and strictly negative subsequences, σ_i^+ and σ_i^-, and consider σ_i^-. One of two things must occur. Either σ_i^- is finite, in which case the sequence σ_i is eventually non-negative, or σ_i^- has infinitely many elements. In this case, the elements of σ_i^- are monotonically increasing (by construction and no cycling), so that σ_i^- converges to a limit of 0. In either case, then, σ_i must be asymptotically non-negative. Since this is true for all agents, the limit point of the allocation sequence has no agent subsidizing any other agent, so we are at a competitive equilibrium. This completes the proof of convergence. This argument becomes a proof of existence of competitive equilibrium once we demonstrate that at the limit prices, agents will actually choose the limit allocation as the solution to their budget constrained utility maximization problems. This will obviously be the case for the environment we study in which agents have strictly convex preferences.

3.3 Connection to Negishi's proof of existence of CE

The Negishi approach to showing existence of CE in convex exchange economies first maximizes a social welfare function of the form

$$\sum_{i=1}^{M} \alpha_i u_i (x_i)$$

with $\alpha_i > 0$ for all i and $\sum \alpha_i = 1$ to obtain a Pareto optimal allocation. Let $\widehat{\mathbf{X}}(\alpha)$ be the optimal allocation obtained for weight vector α, define each agent's surplus or deficit $s_i(\alpha) = p\left[\widehat{\mathbf{X}}(\alpha)\right] \cdot \widehat{z}_i(\alpha)$, and let $s(\alpha) = [s_1(\alpha), \ldots, s_M(\alpha)]$. Negishi uses a fixed point argument to find a zero for the mapping from welfare weights to surpluses. To see the connection between what we do and the Negishi approach, note that Negishi's surplus mapping $s(\alpha)$ takes values on the subspace of $\mathbb{R}^M$ on which $\sum_i s_i(\alpha) = 0$ (by the feasibility constraint). Since this is the tangent space to the $M-1$ dimensional simplex in $\mathbb{R}^M$, the mapping defines a vector field. For our purposes, we will work with essentially the same vector field , but define it not on the space of welfare weights, but on the space PO identified with the $M-1$ utility possibilities simplex $\Delta \subset \mathbb{R}^M_+$. Given any Pareto optimal allocation $\widehat{\mathbf{X}} \in \Delta$, we associate with $\widehat{\mathbf{X}}$ the vector

$$\nu(\widehat{\mathbf{X}}) = \left[p(\widehat{\mathbf{X}}) \cdot \widehat{z}_1, \ldots, p(\widehat{\mathbf{X}}) \cdot \widehat{z}_M\right] \in T_{\widehat{\mathbf{X}}}\Delta = \left\{(w_1, ..., w_M) \in \mathbb{R}^M | \sum w_i = 0 \right\}.$$

The sets $R_i(y_i)$ defined above consist of the loci of allocations in PO on which $p(\widehat{\mathbf{X}}) \cdot \widehat{z}_i = y_i$ for y_i fixed. Now, to implement Negishi's proof of existence in this framework, we note that on the boundary of Δ, the vector field $\nu(\widehat{\mathbf{X}})$ points out of the simplex. This can be verified by considering first any vertex of the simplex. An allocation at a vertex gives everything (subject to individual rationality) to the agent corresponding to that vertex. Such an agent is necessarily subsidy-receiving, while all other agents are subsidy providing, since they will be receiving an allocation on the indifference surface through their endowment, and this allocation necessarily minimizes the cost of anything at least as good as the endowment at the associated supporting prices. For this allocation, then $\nu(\widehat{\mathbf{X}})$ is positive only in the direction of the subsidy-receiving and negative in all other directions and hence points out of the simplex (when we translate the vector field onto Δ) through the vertex. Next, consider any $(M-1)$-dimensional face of the simplex on which one agent i (corresponding to the vertex at e_i opposite the face) receives minimal utility. This face is contained in the coordinate plane corresponding to $u_i = \underline{u}_i$. Since i is a subsidizer in this case, the vector field will have it's i^{th} component negative, at any allocation in the face, in which case the vector itself lies on the negative side of the coordinate plane, and hence, the field points out of the simplex along the face. Since this is true of all the boundary faces, it follows that the vector field points out. We illustrate this in Figure 6 for a three agent economy. Existence of a zero for the vector field now follows from the Poincare-Hopf theorem.

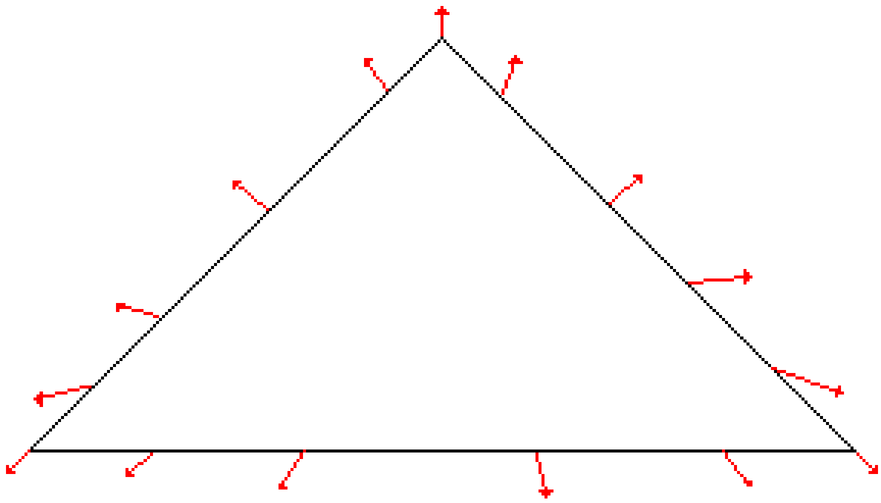

Figure 6: Subsidization Vector Field Points Out of Utility Possibilities Simplex

4 Intermediate exchange

We now turn our consideration to a class of intermediate processes of exchange, that is, exchange processes that do not necessarily involve exchange directly from the endowment to the Pareto set. The processes we describe are intuitive and behaviorally plausible, but it turns out that in higher dimensions they can fail to converge to the Pareto set, leaving open the question of how Pareto optima that survive subsidization constraints might be implemented in an informationally decentralized system.

First, we assume that for all i, equation (1) is modified to apply to any consumption bundle x_i and not simply bundles in the contract set. Thus, exchange must take place within the intersection of the upper half-spaces defined by the subsidization hyperplanes; call the set of allocations that satisfy this condition after period t has ended S^t. Further, we assume that all exchange must be feasible and Pareto-improving, so that if $\mathbf{X}$ is the current allocation during period t, then $\mathbf{Y}$ can be subsequently adopted in the period only if $\mathbf{Y} \in S^{t-1} \bigcap IR(\mathbf{X})$. Finally, we impose a non-intransigence condition,[2] as follows. Fix $\varepsilon > 0$, and consider the intersection of an ε-ball centered at $\mathbf{X}$ with $S^{t-1} \bigcap IR(\mathbf{X})$. For any subset of allocations in this intersection with positive Lesbesgue measure, the probability that $\mathbf{Y}$ is drawn from this subset must be greater than zero. This class of intermediate exchange processes can implement any individual rational allocation that survives all subsidization constraints. The [Gode and Sunder 1993] zero intelligence process and a special case of the [Hurwicz et al. 1975] B-process, amended with subsidization constraints, are examples of allowable processes.

[2]We use "non-intransigence" because agents are not permitted to be deterministically obstinate with respect to the likelihood of accepting any utility-improving opportunity. If all agents are stochastically non-instransigent, this implies the assumption in Section 2 that any subset of the contract set with positive Lesbesgue measure which survives all subsidization constraints can be reached with positive probability.

Under our restrictions on preferences it is straightforward to demonstrate that in period 1 (that is, when subsidization constraints are not binding), exchange must converge to the contract set. Minus the sum of individual utilities at each adopted allocation is a Lyapunov function (the function is continuous in its arguments, Pareto-improvement guarantees strict monotonicity, and feasibility provides boundedness), so the process is globally stable. If it converges to some $\mathbf{X}$ that is not Pareto optimal, non-intransigence and the Borel-Cantelli Lemma imply that some $\mathbf{Y} \in IR(\mathbf{X})$ will eventually be adopted during the period with probability one, violating global stability.

4.1 Convergence with 2 agents, ℓ goods

When there are only 2 agents in the economy (or 2 types of agents who behave identically), we can prove convergence to $\widehat{\mathbf{X}}^t \in CS^{t-1}(\Omega)$ for all $t \geq 1$, in which case convergence across periods to competitive equilibrium has already been established. Noting that global stability is still implied by the Lyapunov function defined above even with subsidization constraints, suppose that exchange in period t converges to $\bar{\mathbf{X}}^t \notin CS^{t-1}(\Omega)$. Since exchange is required to take place within S^{t-1}, it must be the case that $\bar{\mathbf{X}}^t \notin PO$. Then at least one agent must be on her constraint hyperplane at $\bar{\mathbf{X}}^t$; otherwise convergence is violated by Borel-Cantelli as above.

Without loss of generality, suppose agent 1 is on her constraint hyperplane. Under constraint price ρ_1^{t-1}, by virtue of Pareto optimality it must be the case that the corresponding constraint allocation χ_1^{t-1} is optimal for the agent at that price. Since $\rho_1^{t-1}\left(\bar{x}_1^t - \chi_1^{t-1}\right) = 0$ (because she is assumed to be on her constraint hyperplane) and $\bar{x}_1^t - \chi_1^{t-1} = \left(R - \bar{x}_2^t\right) - \left(R - \chi_2^{t-1}\right)$ (by feasibility), then $\rho_1^{t-1}\left(\bar{x}_2^t - \chi_2^{t-1}\right) = 0$. Again by Pareto optimality, it must be the case that $R - \chi_1^{t-1}$ is the optimal consumption bundle available for agent 2 at price ρ_1^{t-1}. By convexity of preferences, therefore, there exists α such that $0 < \alpha < \varepsilon$ and $u_1\left(\bar{x}_1^t + \alpha\left[\chi_1^{t-1} - \bar{x}_1^t\right]\right) > u_1\left(\bar{x}_1^t\right)$ and $u_2\left(\bar{x}_2^t + \alpha\left[R - \chi_1^{t-1} - \bar{x}_2^t\right]\right) > u_1\left(\bar{x}_2^t\right)$. Thus, there exists a Pareto-improvement from $\bar{\mathbf{X}}^t$ that satisfies agent 1's constraint, and furthermore by strict convexity of preferences there exists an open set of Pareto-improving allocations that satisfy 1's subsidization constraint within the ε-ball at $\bar{\mathbf{X}}^t$. Unless agent 2 is also on her subsidization constraint, then $\bar{\mathbf{X}}^t$ cannot be stable.

Finally, note in this 2-person economy that agent 2 cannot have formed her subsidization constraint in the same period as agent 1, or they both would have been subsidizers in the same period, an impossibility. Therefore, without loss of generality assume that the subsidization constraint for agent 2 was updated more recently than the one for agent 1. Then $R - \chi_2^{t-1}$ must satisfy agent 1's subsidization constraint. But by argument analogous to the one made above, then we know there exists an open set of Pareto-improving allocations along the subsidization hyperplane for agent 2 that satisfies 1's subsidization constraint. Thus $\bar{\mathbf{X}}^t$ cannot be stable. Therefore, we have established that $\widehat{\mathbf{X}}^t \in CS^{t-1}(\Omega)$ when $M = 2$, and thus an intermediate exchange process that meets the requirements above will converge across periods to a competitive equilibrium.

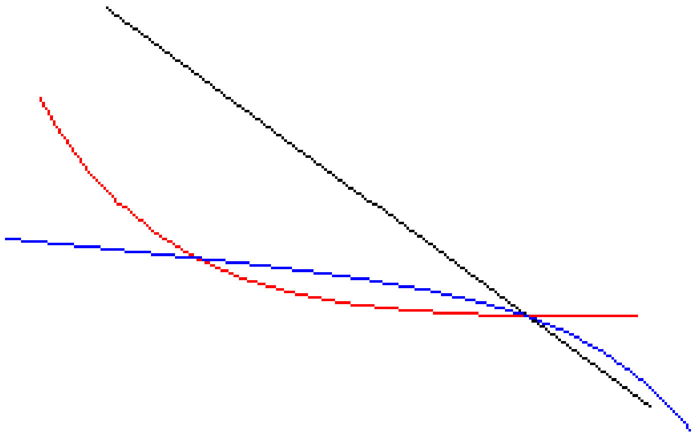

Figure 7: Possbility of non-convergence

4.2 Convergence with more than two agents?

It is instructive to begin consideration of convergence when $M > 2$ under a regime of randomized bilateral exchange. Suppose an exchange process randomly matches two agents at a time, and proposes a feasible reallocation between them drawn from a uniform random distribution on an ε-cube centered on their current (pairwise) allocation. We call behavior in such an institution (bilateral) *ZI+* when the proposal is accepted if and only if it is Pareto-improving and survives all subsidization constraints. Bilateral ZI+ behavior will converge to Pareto optimal allocations when no agent is subsidy constrained, a fact that has been known since the early 1970's; see Feldman [1973], for example. In the presence of subsidization constraints, it is possible for an agent to reach a subsidization constraint at a non-PO allocation where there are no possible pairwise improvements with other agents that do not violate his subsidization constraint. In this situation, unconstrained agents will typically move to a Pareto optimum relative to the set of unconstrained agents, at which point ZI+ cannot get to a fully multilateral Pareto optimal allocation.

This situation is illustrated in Figure 7 for a three agent, two good economy. Here, agent 1 (whose origin is in the lower left-hand corner) is at an allocation which lies on a previously established subsidization constraint. Agent 2's marginal rate of substitution at the allocation is smaller than the slope of agent 1's constraint line, so all possible Pareto improving trades between them necessarily violate agent 1's constraint. A similar picture holds for the pairing of agents 1 and 3. If agents 2 and 3 trade to a pairwise Pareto optimum, where there marginal rates of substitution are smaller than the slope of agent 1's constraint line (but not so small that they would wish to trade good 2 for good 1 with agent 1), bilateral ZI+ will become stuck at a non-Pareto optimal allocation. As will be discussed in the next section, we actually observed this phenomenon in bilateral ZI+ simulations of Scarf's example. It is important to recognize that even a fully multilateral intermediate exchange process as outlined at the beginning of this section of the paper would be stuck at such an allocation, so convergence of intermediate exchange cannot be guaranteed with more than two agents.

While this counter-example is robust under bilateral ZI+, we will shortly demonstrate this is not necessarily the case under fully multilateral exchange. The fact that bilateral search processes can systematically fail in higher dimensional settings suggests that there is a role for intermediaries in facilitating the search for subsidy constrained optima, since avoiding the problem of someone getting stuck on a subsidization constraint may necessarily involve multilateral trade. Hence, we turn our attention next to an algorithm capable of implementing multilateral trade. The basic idea behind the algorithm is to posit a neutral central planner who searches for Pareto improving allocations by implementing a steepest ascent algorithm based on a weighted sum of utilities objective function, with the weights restricted so as to generate Pareto improvements which satisfy subsidization constraints. To develop this algorithm, we first consider the case where no agent is subsidy constrained.

From standard optimization theory, we know that maximization requires that we move in the direction of the gradient of the function being optimized. For our purposes, the function we want to maximize will be a weighted sum of agents' utilities evaluated at a feasible allocation, which we take to be of the form

$$W = \sum_{i=1}^{M-1} \lambda_i u_i(x_i) + \lambda_M u_M \left(r - \sum_{j=1}^{M-1} x_j \right).$$

The gradient of this objective function with respect to x_i is

$$\lambda_i Du_i - \lambda_M Du_m.$$

Since the gradient of the objective function moves us in directions that increase the value of the objective, we use this as the starting point for our search for weights, and let

$$z_i = \lambda_i Du_i - \lambda_M Du_m$$

for $i = 1, ..., M-1$ and

$$z_M = -\sum_{i=1}^{M-1} z_i.$$

Since we want to restrict ourselves to using weights which generate Pareto improvements at each step, we impose constraints on what the weights can be of the form

$$Du_i \cdot z_i \geq 0 \text{ for } i = 1, ..., M-1$$

and

$$Du_M \cdot z_M \geq 0.$$

The first set of conditions requires that

$$\lambda_i \geq \lambda_M \frac{Du_i \cdot Du_M}{\|Du_i\|^2}$$

while the last condition requires that

$$\lambda_M \geq \frac{1}{M-1} \sum_{i=1}^{M-1} \frac{\lambda_i Du_i \cdot Du_M}{\|Du_M\|^2}.$$

So, consider setting the first $M-1$ weights to

$$\lambda_i = \lambda_M \frac{Du_i \cdot Du_M}{\|Du_i\|^2} + \lambda_M \varepsilon_i$$

for small ε_i. The constraint on agent M's weight then requires that

$$\lambda_M \geq \frac{1}{M-1} \sum_{i=1}^{M-1} \left[\lambda_M \frac{Du_i \cdot Du_M}{\|Du_i\|^2} + \lambda_M \varepsilon_i \right] \frac{Du_i \cdot Du_M}{\|Du_M\|^2}$$

or

$$1 \geq \frac{1}{M-1} \sum_{i=1}^{M-1} \left[\frac{(Du_i \cdot Du_M)^2}{\|Du_i\|^2 \|Du_M\|^2} + \varepsilon_i \frac{Du_i \cdot Du_M}{\|Du_M\|^2} \right].$$

Since $\frac{(Du_i \cdot Du_M)^2}{\|Du_i\|^2 \|Du_M\|^2}$ is always less than or equal to one, and equal only when the gradients are colinear, it follows that as long as we are not at a Pareto optimal allocation, we can construct suitable weights by first normalizing $\lambda_M = 1$, and then selecting the remaining weights as

$$\lambda_i = \frac{Du_i \cdot Du_M}{\|Du_i\|^2} + \varepsilon_i$$

for sufficiently small random ε_i's.

To extend the ascent algorithm to the case where some (or all) agents may be subsidy constrained, note that if some agent becomes subsidy constrained at an allocation which is not Pareto optimal, requiring that any reallocation not violate the subsidization constraints essentially replaces the gradient Du_i with the support price p_i at the constraint. So, to cover the possibility that some subset of agents becomes constrained, let $q_i = Du_i$ if i is unconstrained, and $q_i = p_i$ if the agent is constrained. We then replace the normalized gradient in the weight inequalities with the q's. As with the unconstrained algorithm, this algorithm will generate Pareto improvements which satisfy the subsidization constraints as long as

$$1 > \frac{1}{M-1} \sum_{i=1}^{M-1} \left[\frac{(q_i \cdot q_M)^2}{\|q_i\|^2 \|q_M\|^2} \right].$$

In this case, however, the algorithm can get stuck at an allocation which is not Pareto optimal. To see this, suppose we are at a Pareto optimum at which agent 1 has just established a new subsidization constraint. Assuming that all other agents have income expansion paths at the new subsidization support price p_1 which are linearly independent, we can generate a reallocation along each agent's expansion path (either utility increasing or decreasing) such that when agent 1 takes the opposite side of the reallocation she remains on her subsidization constraint. The end result is a non-Pareto optimal, feasible allocation at which all q_i's are colinear with p_1, which violates the inequality above. Hence, in this case, there is no way to generate a Pareto improvement which satisfies agent 1's subsidization constraint. Note also that we can further perturb this allocation along the $M-1$-wise Pareto set in a direction such that the resulting allocation leaves the constrained agent stuck on her constraint with no further Pareto improvements possible. The direction of the perturbation needs to be such that in any pairwise Edgeworth box with one unconstrained agent and the constrained agent, the orthogonal projection of the unconstrained agent's utility gradient onto the constrained agent's constraint hyperplane lies on the side of the intersection of the constraint hyperplane and the support hyperplane determined by the unconstrained agent's utility gradient on which the constrained agent's utility is decreasing. In this situation, the only utility increasing directions for the unconstrained agent which don't violated the constrained agent's subsidization constraint necessarily reduce the constrained agent's utility. (For the 3-agent, 2-good model, this is just the scenario illustrated in the diagram above.) Note, however, that the requirement that we be at an $M-1$-wise Pareto optimal allocation is necessary, since if we perturb to something which is not Pareto optimal for the unconstrained agents, then by making a Pareto improvement for these agents, and then taking a small amount of all goods away from each of the unconstrained agents and giving it to the constrained agent, we can generate a full Pareto improvement for all agents which doesn't violate the constrained agent's subsidization constraint.

The important thing to note about these counter-examples, though, is that under multilateral exchange the situations described are not generic. The set of allocations at which all agents have their gradients colinear with some agent's subsidization constraint is closed and of measure zero in the set of feasible allocations. Hence, the probability of hitting such an allocation during the search for Pareto improvements is zero. Similarly, the set of $M-1$-wise Pareto optima are also closed and of measure zero in the space of feasible allocations. In the example we considered in Figure 7, we note that there are many multilateral reallocations that can move the agent in question off of his constraint, just no bilateral ones. As long as the unconstrained agents are not at a pairwise Pareto optimum, the (Lesbesgue) measure of the set of Pareto-improving reallocations that satisfy agent 1's subsidization constraint is strictly greater than zero, while the measure of allocations at which convergence to (fully multilateral) PO will fail is zero.

These arguments also extend to the cases where any subset of k agents, with $1 < k < M-1$, are subsidy constrained, since a Pareto improvement for all agents is possible unless the remaining $M-k$ agents are at a relative Pareto optimum of the type described above. The two cases where $k = M-1$ or $k = M$ are somewhat problematic, however. To analyze these cases, we consider two subcases. Suppose first that the number of commodities ℓ is strictly greater than $M-1$. Under this assumption, the intersections of the $M-1$ or M constraint sets will be a hyperplane of dimension at least $\ell-(M-1) > 0$. If we restrict preferences to this subspace, then generically there will exist a Pareto improvement within the constraint set. The only time such an improvement is not possible is if we are at a restricted Pareto optimum, which is not generic. Given such a Pareto improvement, we can also generate an improvement which moves all agents off of their constraint hyperplanes. For second case where $M-1 \geq \ell$, we can envision situations where the algorithm could get

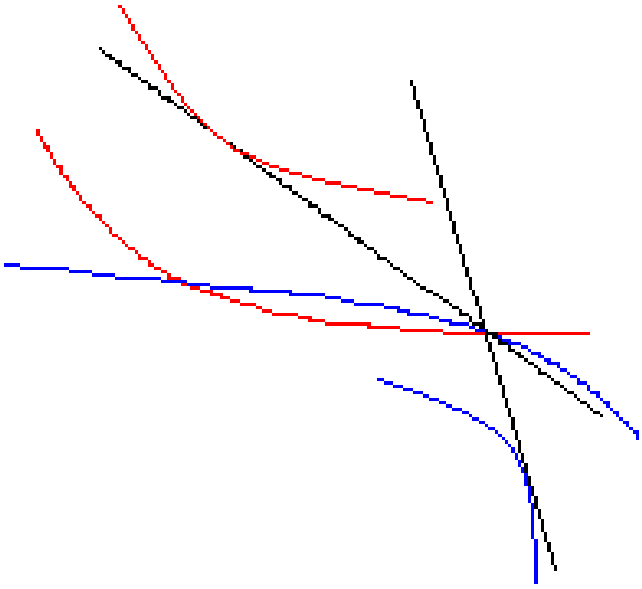

Figure 8: Another example of non-convergence

stuck, as, for example in the 3-agent 2-good economy diagramed in Figure 8: We can show for this situation that Pareto improvements are always possible as long as we are close to a Pareto optimal allocation. Hence, the situation in the diagram must occur away from a Pareto optimum, and we have not been able to rule out this possibility. At the same time, we have not been able to construct a counter-example for the 3-agent, 2-good model whereby we can actually get to such a situation, so we conjecture that the algorithm will work in this case as well. Thus, there is hope that a generic convergence result for multilateral intermediate exchange processes might be obtained in future research, although we have no such result in the present paper.

Finally, we note that for the cases where we can show the algorithm can't get stuck generically, we can show that, with probability one, the algorithm converges using the same Lyapunov-type argument we used for the two agent economy.

4.3 Simulations of Scarf's example

Scarf's example is a three agent, three good pure exchange economy in which Walras' simple *tâtonnement* procedure fails. We modify the economy so that the preferences conform to our monotonicity assumptions by specifying utility functions and endowments as

$$
\begin{aligned}
u_1(x_{11}, x_{12}, x_{13}) &= -\left[\frac{b^3}{x_{11}^2} + \frac{1}{x_{12}^2} + \frac{d^3}{x_{13}^2}\right], \ \omega_1 = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}^T \\
u_2(x_{21}, x_{22}, x_{23}) &= -\left[\frac{b^3}{x_{22}^2} + \frac{1}{x_{23}^2} + \frac{d^3}{x_{21}^2}\right], \ \omega_2 = \begin{bmatrix} 0 & 1 & 0 \end{bmatrix}^T \\
u_3(x_{31}, x_{32}, x_{33}) &= -\left[\frac{b^3}{x_{33}^2} + \frac{1}{x_{31}^2} + \frac{d^3}{x_{32}^2}\right], \ \omega_3 = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^T,
\end{aligned}
$$

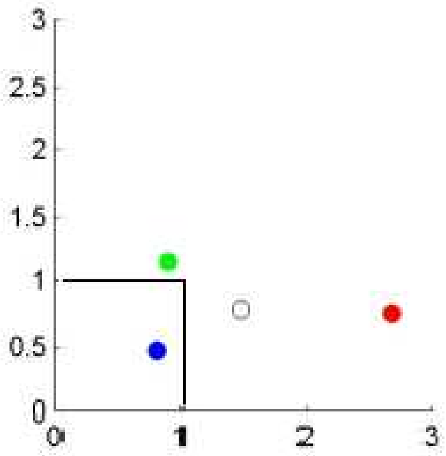

Figure 9: Early Period in Convergence to near-PO

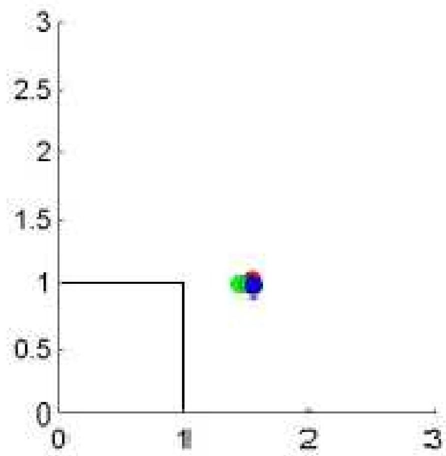

Figure 10: Final Period in Convergence to near-PO

where d is a small positive number. We recover the original example when $d = 0$. We show in the Appendix that $\phi_1 = \phi_2 = \phi_3 = 1$ is the unique equilibrium price for this extended model when the price of one good is normalized to one, and that this equilibrium is unstable under the Walrasian *tâtonnement* as long as $d < \frac{2}{3}$. One can easily demonstrate that the initial endowments as specified are not Pareto optimal.

With this extension of Scarf's example, we simulated a multilateral version of the ZI+ algorithm outlined above. Loosely speaking, we generated a sequence of random (feasible) proposal reallocations within a fixed ε-cube about the current allocation, and proposals were adopted if and only if they were Pareto-improving and satisfied all subsidization constraints. Importantly, as in any simulation, we had to "settle" for approximate convergence to Pareto optimality within-period and approximate convergence to CE across periods. Convergence to PO was determined by using the approximate colinearity of utility gradients. Convergence to CE was established when the absolute value of the subsidization constraint for each agent was less than a "small" value, which we varied between 0.001 and 0.25.

During these simulations, it became apparent that as the allocations approached the Pareto optimal set, the waiting times for improving allocations to be generated became quite long. Because of this, we modified the proposal generation process to direct the search toward directions that yielded Pareto improvements. This procedure essentially involved generating a social welfare function with random coefficients and then using the gradient of this function to direct the search toward allocations which yielded Pareto improvements (while we don't report the details of this modification here, they are available on request). The need for enhancing the search process suggests that in real economic environments, search intermediaries (brokers) play an important role in helping the market exhaust gains from trade.

Figures 8 and 9 each present a snapshot of the normalized gradients of each agent (with the clear circle representing the centroid of the normalized gradients), the former at the start of the search for a Pareto optimum, and the latter at the

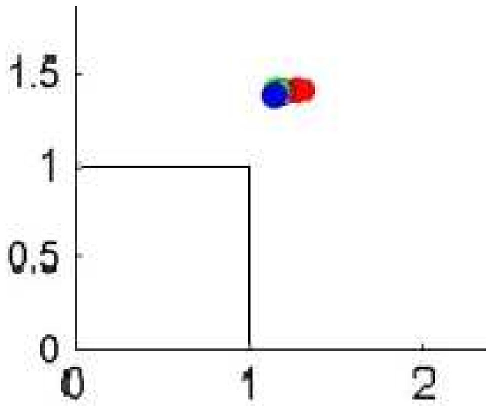

Figure 11: Early Period in Convergence to near-CE

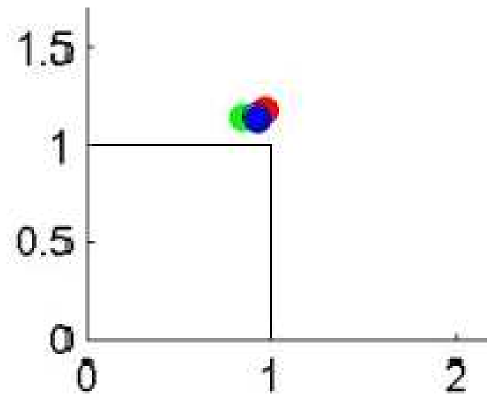

Figure 12: Final Period in Convergence to near-CE

completion. The price of good 1 has been normalized to one, the price of good 2 is represented on the horizontal axis, and the price of good 3 is represented on the vertical axis. The vertical and horizontal lines at the point $(1, 1)$ indicate the competitive equilibrium vector of prices.

Figures 10 and 11 show, respectively, an early period in the search for competitive equilibrium, in which we have attained a Pareto optimum but one which is not the competitive equilibrium, and the final period, where the prices (defined by the normalized gradients) have converged to within a neighborhood of the competitive equilibrium of the model. Animated gifs of the convergence process can be found online at `http://econ.tepper.cmu.edu/spear/po_scarf.gif` (for convergence to Pareto optima) and `http://econ.tepper.cmu.edu/spear/scarf.gif` or `http://econ.tepper.cmu.edu/spear/scarf2.gif` (for price convergence to CE).

These simulations of Scarf's example, each of which converged to a predetermined neighborhood of competitive equilibrium, demonstrate that while general convergence fails under such an intermediate processes, there may potentially be a numerical story of convergence in important examples where *tâtonnement* deterministically does not converge.

5 Conclusion

Because economists identify equilibria through demonstrably non-descriptive maximization assumptions, the feasibility of arriving at and maintaining such equilibria through a plausible process of exchange can provide critical support for the mathematical approach. In this paper we borrow an intuition from both game theory and experimental economics: Coordination is facilitated by many repetitions of a static environment. We assume that markets are generally able to facilitate the exhaustion of gains from exchange, an observation supported by a host of economics experiments.[3] We then consider repetitions of a market from a fixed endowment, and posit an additional intuitive restriction on where the market can coordinate in the future, which eventually forces it to converge to a competitive equilibrium. A nice property of this restriction is that it is immanently testable in the laboratory, or perhaps even in a carefully controlled field experiment (we only need assume that period-ending allocations are Pareto optimal and then estimate a normalized utility gradient from transaction information at the end of the period). Of course, it remains an open question what sort of institutional and behavioral specifications can guarantee subsidy-constrained convergence to the contract set in each period.

On the subject of economic institutions, the results raise several issues. One such issue (which will be of concern to macroeconomists and others doing applied general equilibrium theory) is the pace at which learning about competitive equilibrium occurs. Dynamic economic models typically focus on the "slow" dynamics by which market equilibrium evolves in response to external shocks from one period to the next. Behind the slow dynamics are a "fast" dynamic by which markets are assumed to attain competitive equilibrium. The reference time periods for the slow dynamics (i.e. whether a period represents a week, a month, a quarter, or a year) determines the way in which these kinds of models are calibrated to the data. It is almost universal in such models to assume that the fast dynamics of market equilibration occur quickly enough that calibration is not influenced. Obviously, though, if learning the competitive equilibrium takes significant amounts of time, then short-period calibrations may not be justified. It should also be apparent that environments which exhibit substantial non-stationarity (or complicated though stationary dynamics) will require more time for learning about equilibrium than is required in simpler, stationary environments.

The issue of incorporating production in the model is also important. As it stands, the results we obtain here would be applicable in a setting such as the one described by Radford [1945] on the economic organization of a German prisoner of war camp. We conjecture that a modification of the Negishi approach to demonstrating existence of equilibrium, combined with our notion of reducing the degree of subsidization in moving from one Pareto optimum to the next, would allow the learning process we propose to be extended to production economies as well.

The issue of intermediation (which also arises in the context of the POW camp economy described by Radford) and intermediary institutions is also an important one. As we saw in our discussion of the constrained ascent algorithm, and observed in our simulations of Scarf's example, uniform random search processes become very slow near Pareto optima, and simple bilateral search processes, while fast, can become stuck when agents face subsidization constraints. Our need to resort to the constrained ascent algorithm reflects the fact that in the presence of subsidization constraints, multilateral trade is necessary. Coordinating such trade in economies with more than two agents clearly requires some sort of intermediary institution to facilitate extraction of full gains from trade. This, of course, suggests that

[3] In a partial equilibrium setting Smith [1962] begat a host of other works. In a general equilibrium setting see Williams, Smith, Ledyard and Gjerstad [2000], Anderson, Plott, Shimomura and Granat [2004], Crockett [forthcoming], Crockett, Smith and Wilson [2006], and Noussair, Plott and Reizman [1995], among others.

studying market microstructure will be important in understanding how markets get to equilibrium.

References

Anderson, C. M., C. Plott, K.-I. Shimomura and S. Granat (2004). "Global instability in experimental general equilibrium: The Scarf example." *Journal of Economic Theory*, **115**(2), 209–249.

Balasko, Y. (1988). *Foundations of the Theory of General Equilibrium.* Academic Press.

Bewley, T. (1973). "Edgeworth's conjecture." *Econometrica*, **41**(3), 425–454.

Chatterji, S. and S. Ghosal (2004). "Local coordination in a competitive market." *Journal of Economic Theory*, **114**(2), 255–279.

Cliff, D. and J. Bruten (1997). "Less than human: Simple adaptive trading agents for CDA markets." *Technical Report HP-97-155, Hewlett Packard Research Laboratories, Bristol.*

Crockett, S. (forthcoming). "Learning competitive equilibrium in laboratory exchange economies." *Economic Theory.*

Crockett, S., V. Smith and B. Wilson (2006). "Spontaneous specialization and exchange." Working Paper.

Feldman, A. (1973). "Bilateral trading processes, pairwise optimality, and Pareto optimality." *Review of Economic Studies*, **40**(4), 463–73.

Gode, D., S. Spear and S. Sunder (2004). "Convergence of double auctions to Pareto optimal allocations in the edgeworth box." Working paper.

Gode, D. and S. Sunder (1993). "Allocative efficiency of markets with zero intelligence traders: Market as a partial substitute for individual rationality." *Journal of Political Economy*, **101**(1), 119–137.

Green, J. (1974). "The stability of Edgeworth's recontracting process." *Econometrica*, **42**, 21–34.

Hirota, M. (1981). "On the stability of competitive equilibrium and the patterns of initial holdings: An example." *International Economic Review*, **22**(2), 461–467.

Hurwicz, L., R. Radner and S. Reiter (1975). "A stochastic decentralized resource allocation process: Part I." *Econometrica*, **43**(2), 187–222.

Kirman, A. (1989). "The intrinsic limits of modern economic theory: The emperor has no clothes." *The Economic Journal*, **99**, 126–139.

Kumar, A. and M. Shubik (2004). "Variations on the theme of Scarf's counter-example." *Computational Economics*, **24**(1), 1–19.

Negishi, T. (1960). "Welfare economics and existence of an equilibrium for a competitive economy." *Metroeconomia*, **12**, 92–97.

Noussair, C., C. Plott and R. Reizman (1995). "An experimental investigation of the patterns of international trade." *The American Economic Review*, **85**(3), 462–491.

Radford, R. (1945). "The economic organisation of a P.O.W. camp." *Economica*, **12**.

Scarf, H. (1960). "Some examples of global instability of the competitive equilibrium." *International Economic Review*, **1**, 157–172.

Serrano, R. and O. Volij (2002). "Mistakes in cooperation: The stochastic stability of Edgeworth's recontracting." Institute for Advanced Studies Working Paper 29.

Shapley, S. and M. Shubik (1977). "Trading using one commodity as a means of payment." *Journal of Political Economy*, **84**, 937–968.

Simon, H. (1978). "Rationality as process and as product of thought." *The American Economic Review*, **68**(2), 1–16.

Smale, S. (1974). "Global analysis and economics III: Pareto optimal and price equilibria." *Journal of Mathematical Economics.*

Smale, S. (1976). "A convergent process of price adjustments and global newton methods." *Journal of Mathematical Economics*, **3**(1), 1–14.

Smith, V. (1962). "An experimental study of competitive market behavior." *Journal of Political Economy*, **70**(2), 111–137.

Williams, A., V. Smith, J. Ledyard and S. Gjerstad (2000). "Concurrent trading in two experimental markets with demand interdependence." *Economic Theory*, **16**, 511–528.

6 Appendix: Instability of Scarf's Example for Small d

To show that the basic instability result of Scarf's example continues to hold when preferences are perturbed to make them monotonic, we begin with the excess demand functions for the perturbed economy (by Walras' Law limit ourselves to looking at excess demands for goods x_1 and x_2):

$$\begin{aligned} z_1 &= \frac{b\phi_1^2}{b\phi_1^2+\phi_2^2+d\phi_3^2}+\frac{d\phi_2^3}{b\phi_2^2\phi_1+\phi_3^2\phi_1+d\phi_1^3}+\frac{\phi_3^3}{b\phi_3^2\phi_1+\phi_1^3+d\phi_2^2\phi_1}-1 \\ z_2 &= \frac{\phi_1^3}{b\phi_1^2\phi_2+\phi_2^3+d\phi_3^2\phi_2}+\frac{b\phi_2^2}{b\phi_2^2+\phi_3^2+d\phi_1^2}+\frac{d\phi_3^3}{b\phi_3^2\phi_2+\phi_1^2\phi_2+d\phi_2^3}-1, \end{aligned}$$

where $\phi_\ell = p_\ell^{1/3}$. It is straight-forward to verify that $\phi_1 = \phi_2 = \phi_3 = 1$ is an equilibrium. Uniqueness of equilibrium for small d follows by continuity from the uniqueness result for $d = 0$. To show that for these values of d the equilibrium is unstable under the Walras tâtonnement, we need to compute the Jacobian matrix of this system around the steady-state. For this, we normalize prices such that $\phi_3 = 1$. This yields excess demands

$$\begin{aligned} z_1 &= \frac{b\phi_1^2}{b\phi_1^2+\phi_2^2+d}+\frac{d\phi_2^3}{b\phi_2^2\phi_1+\phi_1+d\phi_1^3}+\frac{1}{b\phi_1+\phi_1^3+d\phi_2^2\phi_1}-1 \\ z_2 &= \frac{\phi_1^3}{b\phi_1^2\phi_2+\phi_2^3+d\phi_2}+\frac{b\phi_2^2}{b\phi_2^2+1+d\phi_1^2}+\frac{d}{b\phi_2+\phi_1^2\phi_2+d\phi_2^3}-1. \end{aligned}$$

A straight-forward but tedious differentiation of these expressions with respect to ϕ_1 and ϕ_2, evaluated at $\phi_1 = \phi_2 = 1$, yields

$$D_\phi z = \frac{1}{(1+b+d)^2}\begin{bmatrix} b(1+d)-2d-3 & (1+b)d+3d^2-2b \\ (1-2d)b+d+3 & (1+d)b-3-2d-3d^2 \end{bmatrix} = \frac{1}{(1+b+d)^2}\mathbf{M}.$$

Letting

$$\mathbf{D} = \frac{1}{(1+b+d)^2}$$

the characteristic equation for $D_\phi z$ is given by

$$ch(\lambda) = \mathbf{D}^4\lambda^2 - \mathbf{D}^2\mathrm{tr}(\mathbf{M})\lambda + \det(\mathbf{M}).$$

The real parts of the roots of this equation are given by

$$r(\lambda) = \frac{\mathbf{D}^2\mathrm{tr}(\mathbf{M})}{2\mathbf{D}^4} = \frac{\mathrm{tr}(\mathbf{M})}{2\mathbf{D}^2}.$$

Now,

$$\mathrm{tr}(\mathbf{M}) = 2b(1+d) - 4d - 3d^2 - 6.$$

For this to be positive requires that

$$2b(1+d) - 4d - 3d^2 - 6 > 0$$

or

$$3d^2 + (4 - 2b)\, d < 2b - 6.$$

Since Scarf's example requires $b \geq 3$, it is sufficient to set $b = 3$ and require that

$$3d^2 - 2d < 0$$

or, since $d \geq 0$

$$d < \frac{2}{3}.$$

Hence, for small values of d, the instability result will continue to hold.