A Simple Decentralized Institution for Learning Competitive Equilibrium

Sean Crockett
Graduate School of Industrial Administration
Carnegie Mellon University

Stephen Spear
Graduate School of Industrial Administration
Carnegie Mellon University

Shyam Sunder
School of Management
Yale University

January 2004

Abstract

The epsilon-intelligent competitive equilibrium algorithm is a decentralized alternative to Walras' tatonnement procedure for markets to arrive at competitive equilibrium. We build on the Gode-Spear-Sunder zero-intelligent algorithm in which random generation of bids and offers from agents' welfare-enhancing opportunity sets generates Pareto optimal allocations in a pure exchange economy. We permit agents to know if they are subsidizing others at such allocations, and to veto such allocations, restricting the subsequent iterations of the algorithm only to those trades that are both Pareto-improving and provide strictly greater wealth, and ultimately utility, for such agents. In this simple institution actions of minimally intelligent agents based on local information can lead the market to approximate competitive equilibrium in a larger set of economies than the tatonnement process would allow. This helps address one of the major shortcomings of the Arrow-Debreu-McKenzie model with respect to the instability of tatonnement in an open set of economies. It also addresses the behavioral critique of mathematically derived equilibria for the inability of cognitively-limited humans to maximize. The proof of convergence of the algorithm presented here also provides a way of showing the existence of competitive equilibrium for monotonic, convex exchange economies with heterogeneous agents and many goods without application of a fixed-point theorem.

Keywords: Learning competitive equilibrium, minimal rationality, allocative efficiency, 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 more 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 [10] has argued that the ADM model and its descendents are overly mathematical, that the underlying assumptions on agents' information processing abilities are unrealistic, and, as a result, the models lack predictive power. Simon's contention that neoclassical economics should be abandoned in favor of a more individual-oriented "satisficing" approach has garnered numerous supporters, particularly outside of mainstream economics.

Criticism of the ADM model has not been limited to those who have concluded that it should be abandoned. Alan Kirman [5] argues 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 attainment of equilibrium. 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 *tatonnement* 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.

The Walrasian *tatonnement* can be implemented, at least conceptually, in a decentralized environment by positing that sellers hold inventories, and when inventories accumulate above "normal" levels, sellers should lower prices, while if inventories drop, then sellers should raise prices. Mathematically, this procedure can be modeled (once we assume conditions sufficient to generate a differentiable excess demand function) as a simple differential equation in the prices of the form

$$\dot{p} = z(p).$$

This price adjustment process increases the prices of any good whose excess demand is positive, and decreases the price of any good whose excess demand is negative. The adjustment stops when $z(p) = 0$.

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's [9] famous example demonstrated the existence of an open set of economies having a unique equilibrium which was unstable under the Walrasian *tatonnement*. This surprising result spawned a substantial literature on stability and came to the conclusion that it was always possible to construct a *tatonnement* 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 *tatonnement* procedure (see, for example, Smale [13]). 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 can hardly be characterized as a decentralized market. 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 *tatonnement* process. When preferences and endowments are such that all goods are gross substitutes, the Walrasian *tatonnement* 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. Therefore, the potential instability of *tatonnement* 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 [11] in which price 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 [2] 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 *tatonnement* 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 *tatonnement* will always converge to some competitive equilibrium.

Kumar and Shubik [6] 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 *tatonnement* literature.

In this paper, we explore an alternative to the idea of *tatonnement* processes, by postulating that economic equilibria, unlike those of physical systems, must be learned (or discovered) by the agents in the system. In physical systems, equilibria occur as natural "rest points" in dynamic processes based on fixed laws of motion of the system. Economic equilibrium, on the other hand, involves not only a physical "rest point" condition (market-clearing), but also a psychological condition (satisfaction of needs or wants) interacting with an artificial construct (prices) derived from exhaustion of the opportunities set. Indeed, it is easy to find market-clearing allocations: any bully can do it very effectively. It is less easy to find Pareto optimal allocations, although well-defined systems of property rights together with enforcement mechanisms for ensuring that contracts are honored makes implementing Pareto improving trades possible, which can lead under very simple search procedures to optimal allocations. But, as we already know, not every Pareto optimum is a competitive equilibrium. Thus, it seems likely that if the concept of the competitive equilibrium is to be useful in economic analysis, we need a mechanism for explaining how agents in the economy may act to generate competitive prices.

Gode and Sunder [3] took a very different approach to the problem of implementing the competitive equilibrium. In their approach, they analyzed a partial equilibrium environment involving a single market with many interacting agents, based on the standard double auction supply and demand experiments pioneered by Vernon Smith and Charles Plott. In the experimental version of this market, one group of agents play the role of buyers, the other the role of sellers. Buyers can purchase one unit of the good, and this one unit is worth a given reservation price to them. Hence, if they buy the good for a price at or below their reservation value, they earn a profit. Sellers can each sell up to one unit of the good. If they sell their unit, they incur a given production cost. Hence, if they sell at a price above their cost, they make a profit. It is well established in the literature on experimental markets that human traders in this environment eventually end up trading the competitive amount of the good at prices that closely approximate the competitive equilibrium (i.e., transactions take place according to the price and aggregate quantity specified by the intersection of the supply and demand schedules for the experiment). We note, however, that agents in this experiment generally require several rounds of trading before they learn what the relevant equilibrium prices are, so that the data generated in such experiments exhibit a convergence of prices and quantities to the predicted competitive equilibrium prices and quantities, rather than an abrupt and direct implementation of the equilibrium.

Gode and Sunder asked whether this process of finding the right prices and allocations requires sophisticated learning, or whether it could be implemented with "zero intelligence" search procedures. They proceeded to replicate 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. Thus, buyers were restricted to bid below their reservation prices, while sellers were restricted to offer above their costs. 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. If we interpret the prices at which the inframarginal trades occur as the limit of the pricing sequence generated in the simulation, then the zero intelligence procedure is also capable of finding the competitive equilibrium prices, as well.

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.

Follow-on work to the Gode and Sunder research by Gode, Spear and Sunder [4] showed that, at least in the context of a two agent, two good exchange economy, simple random 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 research undertaken here focuses on the question of how much additional "intelligence" 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 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 agent's 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. An algorithm to learn equilibrium is developed by combining directed random search for Pareto optima with an application of the welfare theorems that makes use of the learned gradient.

The First Welfare Theorem implies that every competitive equilibrium is a Pareto optimum. Therefore, the algorithm 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. Using the common normalized utility gradient to price a Pareto optimal allocation, these prices can be used to determine the implied wealth redistribution associated with the allocation. Obviously, if the agents have reached a Pareto optimum that is not a competitive equilibrium allocation, some agents (although weakly happier) will have less implied wealth under the current allocation than with their initial endowments, i.e. these agents will be subsidizing the consumption of other agents at the observed prices. In such a situation, it is plausible to posit that these agents would wish 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. We implement this idea as an algorithm operating over repeated rounds of trading which can be viewed either as actual repeated trading in an otherwise stationary environment, or as repeated rounds of fictitious trade as is standard in the *tatonnement* literature. In each round of trade, agents move in small random steps from the endowment allocation to something in the Pareto set. At a Pareto optimal allocation, agents will agree on relative prices and can determine whether they are subsidizing other agents at these prices or not. If no one is subsidizing anyone else, we are at a competitive equilibrium. If some agents are providing subsidies, a new round of trading occurs. In this round, agents who were providing subsidies in the pervious round accept only trades which have them paying strictly smaller subsidies at the prices determined from the previous round. We then demonstrate two convergence results showing that a mild strengthening of this process converges to competitive equilibrium, while the original process does so with probability one.

The remainder of the paper outlines the theoretical basis for this research, provides a proof of the convergence of the algorithm to the competitive equilibrium, and reports on computer simulations of the search algorithm in a number of well-known examples.

2 An ϵ-intelligent Implementation of Competitive Equilibrium

The economic model is one in which a finite number of agents trade a finite number of goods and services in a deterministic, pure exchange environment. By way of notation, we index agents as $i = 1, ..., M < \infty$ and goods as $j = 1, ..., \ell < \infty$. Each agent is completely characterized by his preferences and endowments. Preferences are represented by a utility function,

$$u_i : \mathbb{R}_+^\ell \to \mathbb{R}$$

which we assume is strictly increasing, strictly quasi-concave, and at least twice differentiable. An allocation in this economy is defined in the usual way, as an

m-tuple of commodity vectors $[x_1, ..., x_M]$ with $x_i \in \mathbb{R}^\ell_+$ for $i = 1, ..., M$, and $x_{i\ell}$ is agent i's allocation of good ℓ. Let $\omega_i \in \mathbb{R}^\ell_+$ be the initial allocation for agent i. Feasible allocations satisfy

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

An allocation is Pareto optimal if it is feasible, and there exists no other feasible allocation which makes no agent worse off, and makes some agent better off. An allocation $[\widehat{x}_1, ..., \widehat{x}_M]$ is said to be a competitive equilibrium if it is Pareto optimal, and such that for each agent

$$p \cdot \widehat{x}_i = p \cdot \omega_i$$

where the price vector p is defined as

$$p = \widehat{Du}_i\left(x_i\right) = \frac{1}{\sum_{i=1}^{\ell} \partial u_i\left(x_i\right) / \partial x_i} Du_i\left(x_i\right).$$

The ϵ-intelligence algorithm for implementing the competitive equilibrium operates in stages. Each stage is broken down into a sequence of steps, the end result of which is an allocation which is approximately Pareto optimal. Once an approximately optimal allocation is found, agents use their knowledge of their own preferences to price the allocation, and compare the value of the allocation at these prices to the value of their endowments. If an agent learns that the value of her allocation is less than that of her endowment, we impose the constraint during all future stages that this agent receive an allocation of greater value, and that the limit allocation of the next round of trading yield the agent an allocation which makes her strictly better off.

Stage 0:

Comment on Coordinates: Since we wish to work in the generalized Edgeworth box context, we will adopt a coordinate system consisting of the allocations received by the first $M-1$ agents, the last agent's allocation then being defined as the difference between the total resources and what every other agent gets.

For notational symmetry, we will denote the initial endowment allocation as

$$y_0^0 = \left[\left(x_i^0\right)_{i=1}^M\right] = \left[\left(\omega_i\right)_{i=1}^M\right].$$

Subsequent allocations generated within stage zero will be denoted by

$$y_n^0 = \left[\left(x_i^n\right)_{i=1}^M\right]$$

where n denotes the iteration of the algorithm (or step number) within stage zero. In later stages t, we will denote these allocations as y_n^t.

At stage 0, we simply generate a near Pareto optimal allocation using the following procedure. Fix a number $\varepsilon > 0$ and a number $0 < r < \infty$. For each agent $i = 1, ..., M-1$, let Q_i^{ε} be the (boundary of the) cube of side ε centered at x_i^0. Let

$$S_i^0 = Q_i^{\varepsilon} \cap \left\{ x \in \mathbb{R}_+^{\ell} \mid u_i(x) \geq u_i(\omega_i) \right\}.$$

Now, choose a vector z_i^0 from Q_i^{ε} at random, taking the probability measure on Q_i^{ε} as (normalized) Lebesgue measure. Since the at-least-as-good-as set $\left\{ x \in \mathbb{R}_+^{\ell} \mid u_i(x) \geq u_i(\omega_i) \right\}$ has open interior (as long as the initial endowment is not Pareto optimal), there is a positive probability that $z_i^0 \in S_i^0$. If we take a sequence of independent draws from the distribution on Q_i^{ε}, the Borel-Cantelli lemma implies that we will realize a vector in S_i^0 with probability one. Since the process of taking independent draws from a uniform distribution on Q_i^{ε} will generate vectors which are either in S_i^0 or in $Q_i^{\varepsilon} \backslash S_i^0$, we can apply standard formulas for determining waiting times for events drawn according to a binomial distribution to estimate how many realizations it takes on average to obtain a vector in S_i^0. Let $q = prob\left(Q_i^{\varepsilon} \backslash S_i^0\right)$. Then, the probability that it takes more than r tries to obtain the first vector in S_i^0 is given by

$$prob(r) = q^r.$$

We can use this formula to determine how many trials we need, given q, in order to ensure that the probability of not obtaining a vector in S_i^0 in r tries is less than any desired probability. Letting $\widehat{\rho}$ be the desired probability, we find

$$r = \frac{\ln \hat{\rho}}{\ln q}.$$

For the case where $q = 0.5$ and we want the waiting probability to be $\hat{\rho} = 0.00001$, this yields a value of $r \doteq 16$. For $\hat{\rho} = 0.000001$, we get $r \doteq 20$. Note that as $q \to 1$ (so that the improving region for this individual becomes vanishingly small), the expected waiting times will diverge. We will return to this issue later in our discussion of the simulations.

Once we generate the vectors $\left\{ z_i^0 \right\}_{i=1}^{M-1}$, we construct the allocation $y_1^0 = \left[\left(x_i^1 \right)_{i=1}^{M} \right]$ by taking

$$x_i^1 = z_i^0$$

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

$$x_M^1 = \sum_{i=1}^{M} \omega_i - \sum_{i=1}^{M-1} x_i^1$$

If $u_i\left(x_i^1\right) \geq u_i(\omega_i)\ \forall i = 1, ..., M$ and is strictly greater for at least one agent, the new allocation y_1^0 is a Pareto-improvement and is adopted. If the new allocation does not Pareto-improve upon the old, we set $y_1^0 = y_0^0$. We note, from this definition of the reallocation, it could occur that some agents might receive negative quantities of some good. If all agents have strictly positive endowments, then by restricting the ε-cube from which we select the reallocation to be small enough, we can guarantee that all agents will receive strictly positive. If endowments are not strictly positive, then we modify the algorithm so that if any agent receives a negative allocation, the search procedure is re-started.

The procedure is then repeated by generating y_2^0 and either adopting it if it is a Pareto-improvement or setting it equal to y_1^0 if it is not. We then repeat the process until we find a near-optimal allocation. To determine near-optimality, we use the fact that, given our assumptions on the utility functions, agents utility gradients are colinear if and only if the underlying allocation is Pareto optimal. We use this to define the convergence criterion by first normalizing each agent's utility gradient. We then calculate the mean (i.e. the centroid) of the normalized gradients, and compute the total squared deviation of all normalized gradients from the mean. When this total deviation is less than a small, pre-specified number that determines the accuracy of the approximation, the search process stops. (Details of this construction are outlined in the Appendix.)

Stage t+1

Given a near-Pareto optimal allocation from stage t, y^{t+1}, we price this allocation using the common normalized gradients

$$p^t = \widehat{Du}_i\left(\widehat{x}_i^t\right) = \frac{1}{\sum_{i=1}^{\ell} \partial u_i\left(\hat{x}_i^t\right)/\partial \hat{x}_i^t} Du_i\left(\widehat{x}_i^t\right)$$

where $\widehat{x}_i^t$ is agent i's final allocation at the end of stage t.

Now, define the i^{th} agent's *gain* at this allocation as

$$\lambda_i^t = p^t \cdot \left(\widehat{x}_i^t - \omega_i\right).$$

If $\lambda_i^t < 0$, then we say that agent i is *subsidizing other agents*. We note that if no agent in the economy is providing any subsidies, then we must be at a competitive equilibrium, since $\lambda_i^t \geq 0$ for all i implies

$$p^t \cdot \sum_{i=1}^{M} \left(\widehat{x}_i^t - \omega_i\right) \geq 0.$$

From our assumptions on utility functions, we know that $p^t \gg 0$, while, from feasibility, $\sum_{i=1}^{M} \left(\widehat{x}_i^t - \omega_i\right) = 0$. We infer, then, that in fact $p^t \cdot \left(\widehat{x}_i^t - \omega_i\right) = 0$ and we are at a competitive equilibrium.

Fix a small δ. If every agent in the economy has $\left|\lambda_i^t\right| < \delta$, then we will say that the algorithm has converged to a δ-competitive equilibrium, and the algorithm halts. If some agent has $\left|\lambda_i^t\right| > \delta$, then we repeat our construction of the near Pareto optimal allocation, but with the added set of constraints that

$$0 \geq p^t \cdot \left(x_i^{t+1} - \omega_i\right) \geq \lambda_i^t + \eta_i$$

where η_i is small and positive, for any i such that $\lambda_i^t < 0$. These constraints guarantee that any agent who was providing subsidies at stage t will be providing smaller subsidies at stage $t+1$.

Note that in passing from one stage to the next, we must always carry along the subsidization constraints, even if we move to a new allocation in which an agent who was providing a subsidy in the previous stage is receiving a subsidy in the current stage. If we "forget" the past subsidization constraint, then the algorithm could go back to an allocation in which this agent was again making losses, possibly larger than in the previous stage. Hence, at each stage t, the data required for each agent is $[\widehat{x}_i^\tau, \lambda_i^\tau, p^\tau]$, i.e. the allocation, price and loss in the last stage τ at which agent i incurred a loss. We illustrate this in Figure 1, for the 2-by-2 economy. Here, at some earlier allocation marked A, agent 1 incurred a loss, so subsequent generations of optimal allocations must lie on or above the line defined by $p^t \cdot z_1 = \lambda_1 + \eta_1$, where z_1 is agent 1's net trade. Similarly, at a different stage, agent B incurred a loss, so subsequent allocations must lie on or below the line defined by $p^{t+1} \cdot z_2 = \lambda_2 + \eta_2$. This restricts the set of potential reallocations from which the ε-intelligent mechanism can draw to those in the grey shaded area. In the absence of the loss constraints, the set of potential reallocations would be the full lens-shaped region between the pair of indifference curves containing the endowment point.

The diagram, as drawn, basically assumes that the competitive equilibrium always lies in the region defined by the intersection of the set of Pareto improving allocations and the subsidization constraints, so we turn next to the question of whether (or when) the algorithm converges to competitive equilibrium.

3 Convergence Results

To examine convergence, we let PO denote the set of Pareto optimal allocations in the generalized Edgeworth box consisting of all allocations $\mathbf{x} = (x_1, ..., x_M)$ such that

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

Under our standing assumptions on preferences, it is well-known that generically, PO is the intersection of an $M-1$ dimensional manifold in the space of allocations, with the set of feasible allocations with fixed resources (see, for example, Smale [12] or Balasko [1]). Since this set is bounded while any manifold is closed, it follows that PO is compact. From the specification of the algorithm, one of two things must happen at any stage: either the reallocation at that stage hits a near-competitive equilibrium, in which case the algorithm halts, or the algorithm continues to the next stage. Obvoiusly, if we attain a near-competitive equilibrium, we are done. So, we need to deal with the case

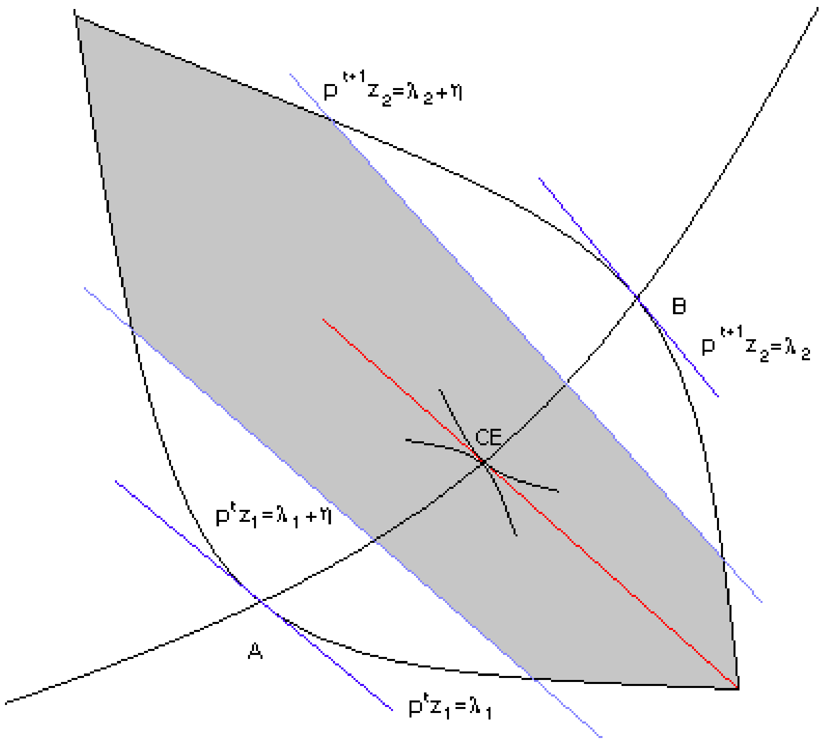

Figure 1:

where the algorithm continues to run (and for this, we assume we are trying to implement an exact competitive equilibrium), and ask under what conditions the sequence of allocations generated will have the competitive equilibrium as its limit.

The first issue we need to deal with is the fact that in it's current form (i.e. with subsidization comparisons made using previous round prices), the algorithm can cycle. While it should be clear from the geometry of the 2×2 economy illustrated in Figure 1 that the algorithm will not cycle, in higher dimensional settings, it may. To see this, consider a 2-good, 3-agent economy and a given Pareto optimal allocation in which agent 1 consumes the bundle x_1. Given agent 1's consumption, the remaining resources define an Edgeworth box for the remaining 2 agents. At the Pareto optimal allocation, agent's 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. 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 the simple version of the algorithm considered so far 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.

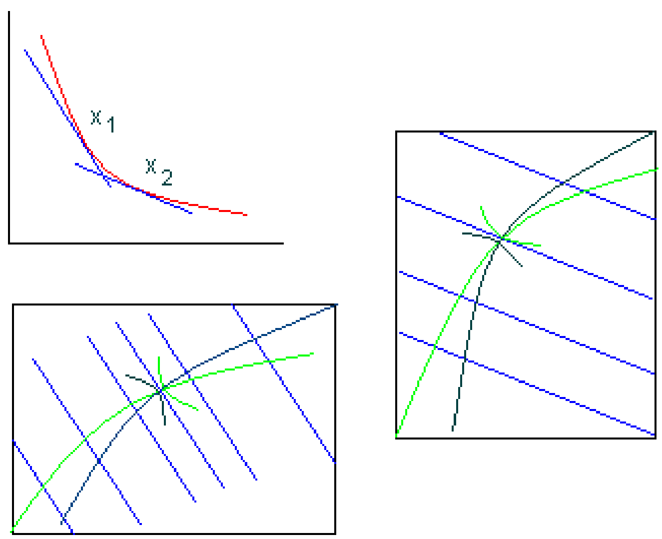

Figure 2:

Since the algorithm can cycle, we will discuss convergence in two steps. First, we will show that a strengthening of the algorithm to require comparison of subsidies not at the old prices, but rather at the prices determined by the new Pareto optimum, does in fact converge. In looking at this result, it will also become apparent that the modified algorithm is closely related to Negishi's proof of existence of competitive equilibrium, so that the modified algorithm can be viewed as an algorithmic implementation of the Negishi approach. Following this, we turn to the question of how likely the original algorithm is to cycle indefinitely, and will show that in fact, cyclic behavior must be transient, so that the original algorithm will converge to the competitive equilibrium with probability one.

3.1 The Negishi algorithm

The basic algorithm is as before, except that we change the subsidy domination requirement as follows.

1.
   - Let x^t be the Pareto optimal allocation at round t, and x^{t+1} the PO allocation at round $t+1$. Round t prices are given by $p(x)$, which are the common supporting prices corresponding to allocation x, for $x = x^t, x^{t+1}$.
   - The modified subsidy domination condition for each agent, then, requires that if

$$p\left(x^t\right) \cdot z_i^t = \eta_i^t < 0$$

then agent i will reject allocation x^{t+1} unless

$$p\left(x^{t+1}\right) \cdot z_i^{t+1} > \eta_i^t.$$

We need to show that the sequence of allocations so generated converges to competitive equilibrium. Let

$$R_i(y_i) \subset PO \cap IR$$

be the set of individually rational Pareto optimal allocations such that $p(x) \cdot z_i = y_i$. We will call y_i for which $R_i(y_i)$ is non-empty *admissable* values of y_i. Now, let

$$\hat{\Delta} = \left\{ p \in \Delta \mid p = \widehat{Du}_i(x_i) \text{ for } (x_1, ..., x_M) \in PO \cap IR \right\}.$$

We also let

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

be the set

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

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

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

Then,

$$R_i(y_i) = \bar{R}_i(y_i) \cap PO \cap IR$$

Proposition: $R_i(y_i)$ is non-empty for admissable y_i.

Proof: Since y_i is admissable, there exists an allocation x in $PO \cap IR$ such that $p(x) \cdot z_i = y_i$, so that x_i lies on the hyperplane determined by $p(x)$ corresponding to wealth $p(x) \cdot \omega_i + y_i$. Since this hyperplane supports x_i, it follows that $x_i = f_i[p(x), p(x) \cdot \omega_i + y_i]$ and $x_i \in \hat{R}_i(y_i)$. It now follows that $x \in R_i(y_i)$ and hence, $R_i(y_i)$ is non-empty.■

Proposition: $\bar{R}_i(y_i) \pitchfork PO$ 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) 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 Balasko's Theorem 2.4.1 (*Foundations of the Theory of General Equilibrium*), 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.■

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,

$$R_i(y_i) = \bar{R}_i(y_i) \cap PO \cap IR$$

will be a codimension 1 submanifold of $PO \cap IR$ with boundary (making the usual identification of PO with the $M - 1$ dimensional simplex) for regular, admissable y_i. Hence, we have that for every non-critical y_i, $R_i(y_i)$ separates $PO \cap IR$ locally into two half spaces. On one side, allocations are such that

$p(x) \cdot z_i < y_i$, while on the other $p(x) \cdot z_i > y_i$. Figure 3 shows the $R_i(y_i)$ sets fo $y_i = 0$ for a 3-agent economy having a unique competitive equilibrium. Figure 4 shows the y_i 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.

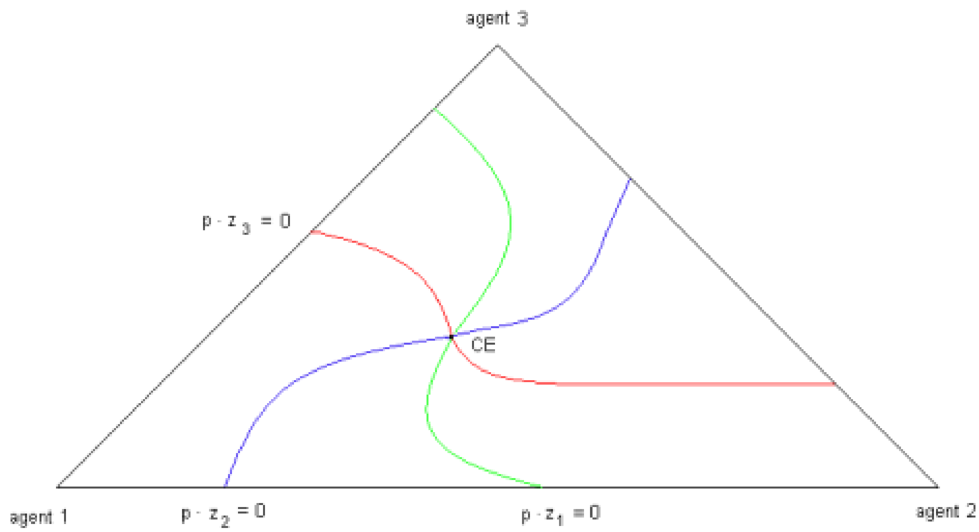

Figure 3:

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 algorithm, 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

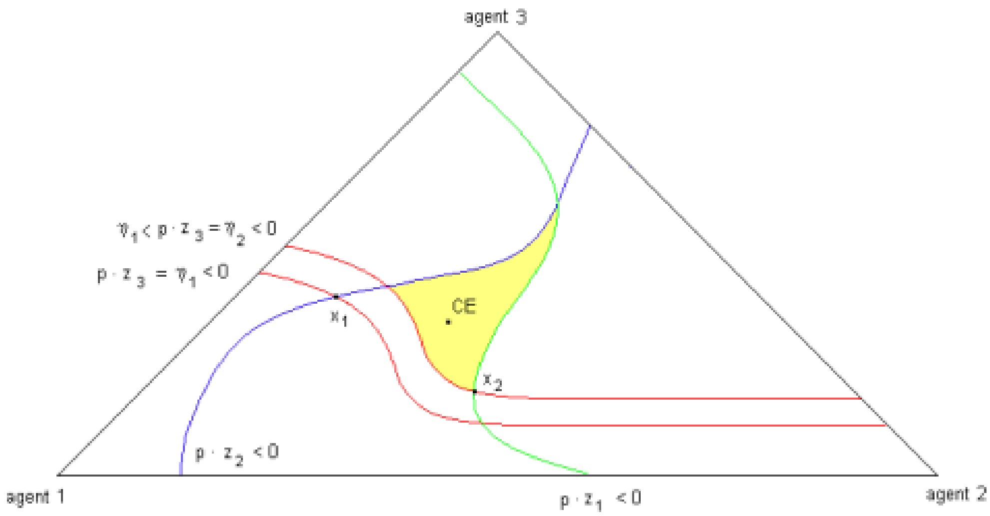

Figure 4:

the set of Pareto optima into disjoint regions in which allocations either have larger or smaller value than y_i. This feature implies that the algorithm cannot cycle.

Given the non-cycling of the algorithm, we can show that it must converge to competitive equilibrium as follows. Consider the sequence of allocations generated by the algorithm. As before, there are two possibilities. If the algorithm hits the competitive equilibrium in finite time, we are obviously done. So assume the algorithm generates an infinite sequence of allocations. 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

$$\sigma_i = \left[p\left(x^t\right) \cdot z_i^t\right]_{t=1}^{\infty}.$$

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 the fact that the algorithm can't cycle), 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

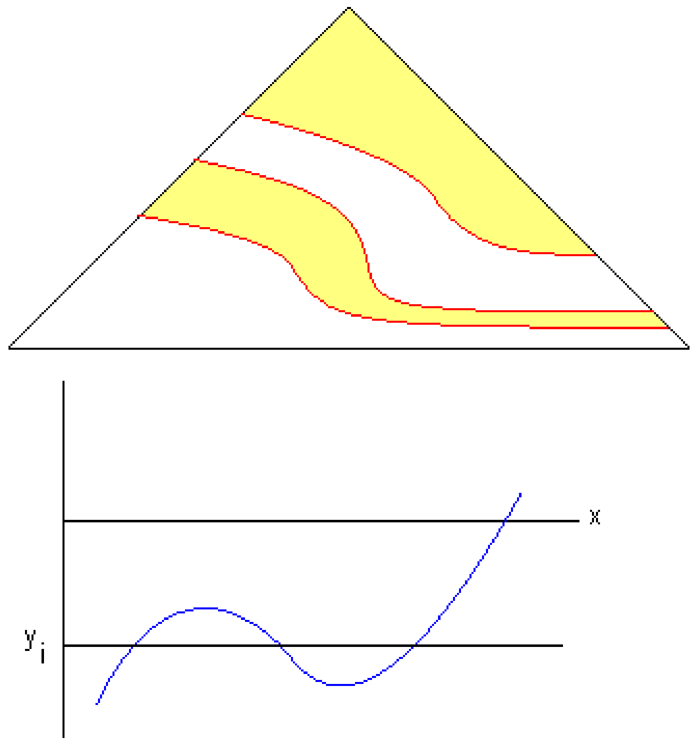

Figure 5:

environment we study in which agents have strictly convex preferences.

3.2 Connection to Negishi's proof of existence of CE

The so-called 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 $\sum \alpha_i = 1$ to obtain a Pareto optimal allocation. Let $x(\alpha)$ be the optimal allocation obtained for weight vector α, and define each agent's surplus or deficit $s_i(\alpha) = p[x(\alpha)] \cdot z_i(\alpha)$ and let $s(\alpha) = [s_1(\alpha), ..., s_M(\alpha)]$. Negishi then 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 $x \in \Delta$, we associate with x the vector

$$\nu(x) = [p(x) \cdot z_1, ..., p(x) \cdot z_M] \in T_x\Delta = \left\{ (w_1, ..., w_M) \in \mathbb{R}^M \mid \sum w_i = 0 \right\}.$$

The sets $R_i\,(y_i)$ defined above consist of the loci of allocations in PO on which $p(x) \cdot z_i = y_i$ 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 $\nu(x)$ 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 $\nu(x)$ 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

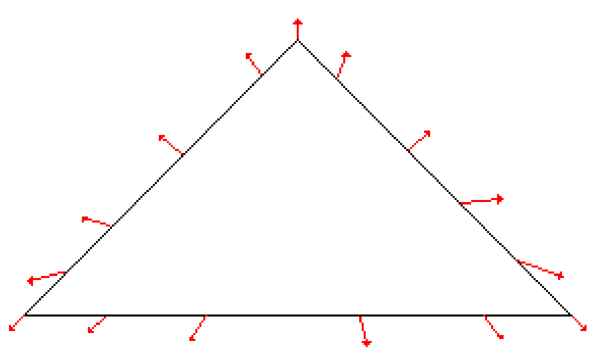

Figure 6: Negishi vector field

follows from the Poincare-Hopf theorem.

3.3 Convergence in probability

We turn now to the question of how likely the original algorithm is to cycle, given that we generate new Pareto optimal allocations randomly. To this end, consider a given Pareto optimal allocation, and look at agent i's consumption bundle, which we denote by x_i^0. Generate the offer surface for agent i through this bundle. This surface is the locus of consumptions for agent i which have value (at the associated supporting prices) just equal to the value at these prices of x_i^0. Note next that in agent i's consumption set, we can bound the agent's consumption from above with the total resources of the economy, and from below by the indifference curve through the agent's endowment (by individual rationality). Also, as long as there are more than 2 agents, there will exist a subset of consumptions having dimension greater than one which correspond to some set of Pareto optimal allocations. This follows from the analysis of the cycling case above, since with more than 2 agents, it is possible to move along some agent's indifference curve starting from a component of some Pareto optimal allocation, and adjust other agent's consumptions so as to maintain Pareto optimality.

Now, consider the set of consumptions for agent i which lie below the offer surface, but above the hyperplane generated by the support prices at x_i^0. (See Figure 7.) Consumption bundles in this region are such that at their associated support prices, the cycling relationship holds. To see this, consider the diagram in Figure 8, which shows a blow-up of the offer curve near x_i^0. For any point on the offer curve different from x_i^0, the line from x_i^0 to the curve will be tangent to an indifference curve at the point on the offer curve. If we move up this indifference curve on the lower side of the offer curve, the support prices will change so that at such a bundle, x_i^0 has higher value at these support prices that the new bundle, while the new bundle has higher value at the support prices for x_i^0, as illustrated by the green support line in the diagram. If we intersect this region with the set of posssible Pareto optimal consumptions (and with the resource and individual rationality bounds), we see that the algorithm can cycle in this region. No cycling is possible, however, for allocations above the offer curve. We can also obviously extend this argument to higher dimensional settings as well, although if there are more goods than agents in the economy, the set of allocations for any agent i compatible with Pareto optima will be a lower dimension set in $\mathbb{R}_+^\ell$, and not the region with non-empty interior indicated in the diagram.

Note next that when we restrict our attention to the stochastic process the original algorithm generates for choosing Pareto optimal allocations, this process is Markovian. Indeed, given a Pareto optimal allocation x^t at some time t, this allocation completely determines the support of the possible next period allocations

$$\left\{(x_1, ..., x_M) \in PO \cap IR \mid p\left(x^t\right) \cdot z_i \geq \eta_i^{t+1},\ i = 1, ..., M\right\}$$

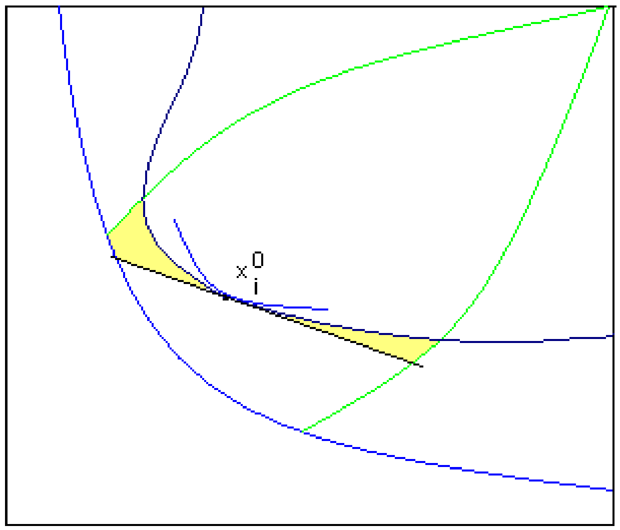

Figure 7: Cycling Region

where η_i^{t+1} is agent i's (ratcheted) subsidization constraint active in period $t+1$. Given the support, the algorithm then generates essentially a uniform probability of selecting an allocation in this support. From the discussion above, then, we can split the support into cycling and non-cycling components for each agent. Since there are always relative open sets of allocations above any agent's offer surface in the support, it follows that the probability of hitting any agent's cycling component is strictly less than one. Hence, with probability 1, given x^t, the algorithm will select an allocation which lies in the non-cycling region of every agent's offer surface through x_i^t. Once this event occurs, the algorithm can never return to the period t cycling regions. Hence, the cycling allocations are transient, and to determine the long-run behavior of the algorithm, it is sufficient to consider how it works on allocations which are non-cycling at each step.

We can now show that the weaker algorithm converges to CE with probability one by arguing as before using the monotonicity of the ratcheting procedure.

Both versions of the algorithm improve significantly on the *tatonnement* results, to the extent that each requires only that agents be able to price Pareto optimal allocations, a process which requires information only about the common normalized utility gradient at the current allocation point, rather than information about both first- and second-derivatives of all agents' utility functions. The well-known fact that the supporting prices for a Pareto optimal allocation are also unique means that this price vector will be unambiguously defined (at least in the complete markets setting). The procedure also provides a more realistic foundation for actually implementing the competitive equilibrium

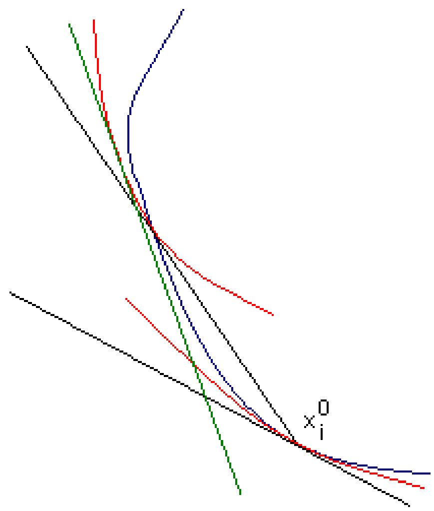

Figure 8:

than does the fictitious price-adjusting auctioneer, based on standard bargaining theory. In this framework, once agents learn that they are in fact subsidizing other agents, they may use the threat of refusing to trade with agents who are benefitting from this subsidization to extract concessions (in the form of trades which reduce the degree of subsidization). This threat is made credible by the fact that in a large economy, basic core-equivalence results imply that there are always subsets of agents who can improve on any non-CE allocation by trading solely among themselves.

4 Scarf's Example

In this section, we apply the simple version of the algorithm to the problem of finding the unique competitive equilibrium in Scarf's example, a case in which the original *tatonnement* procedure fails.

Scarf's example is a three agent, three good pure exchange economy. Utility functions and endowments are specified as:

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

The agents' demand functions are obtained as usual by solving the budget-constrained utility maximization problems

$$\max_{(x_{i1},x_{i2},x_{i3})} u_i\left(x_{i1},x_{i2},x_{i3}\right)$$
subject to
$$p\left(\omega_i - x_i\right) = 0$$

Solving first order conditions and making appropriate substitutions gives the demand functions

$$\begin{aligned}
x_{11} &= \frac{bp_1^{2/3}}{bp_1^{2/3}+p_2^{2/3}},\ x_{12} = \frac{p_1}{bp_1^{2/3}p_2^{1/3}+p_2} \\
x_{22} &= \frac{bp_2^{2/3}}{bp_2^{2/3}+p_3^{2/3}},\ x_{23} = \frac{p_2}{bp_2^{2/3}p_3^{1/3}+p_3} \\
x_{33} &= \frac{bp_3^{2/3}}{bp_3^{2/3}+p_1^{2/3}},\ x_{31} = \frac{p_3}{bp_3^{2/3}p_1^{1/3}+p_1}
\end{aligned}$$

We make the change of variables $\phi_\ell = p_\ell^{1/3}$, and by Walras' Law limit ourselves to looking at excess demands for goods x_1 and x_2 (i.e., if markets for both of these goods clear, then so will the market for good x_3):

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

Clearly, $\phi_1 = \phi_2 = \phi_3 = 1$ is an equilibrium for this economy. It is not difficult to show that this equilibrium is also unique for values of $b \geq 3$, but we shall not do so here. As previously mentioned, Walrasian *tatonnement* is unstable in this economy; the solution trajectories of $\dot{p} = z\left(p\right)$ converge to a limit cycle.

To apply the ϵ-intelligent algorithm to this model, we first modify it so that the preferences conform to our monotonicity assumptions by specifying utility functions and endowments as

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

where d is a small positive number. Note that 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 for this extended model, and that this equilibrium in unstable under the Walrasian *tatonnement* as long as $d < \frac{2}{3}$.

4.1 Results for Scarf's Example

With this extension of Scarf's example, we simulated the ε-intelligent algorithm to search for the competitive equilibrium in this environment, specifying a range of small values for δ (between 0.001 and 0.25) and using the approximate colinearity of the utility gradients to determine convergence to Pareto optima. (We note that for these simulations, we normalized utility gradients by dividing all marginal utilities by the marginal utility of the first good.) 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 algorithm 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 algorithm suggests that in real economic environments, search intermediaries (brokers) play an important role in helping the market find its way to the competitive equilibrium.

Figures 9 and 10 show two views of the normalized gradients of each agent (with the clear circle representing the centroid of the normalized gradients) at the start of the search for a Pareto optimum, and at the completion. The vertical and horizontal lines at the point (1,1) indicate the competitive equilibrium.

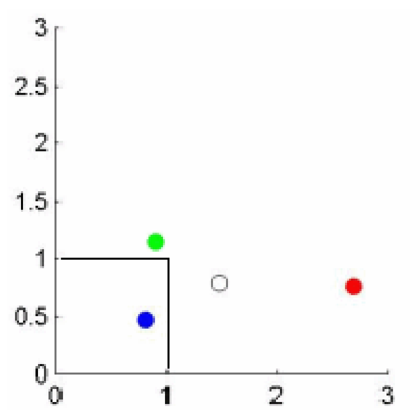

Figure 9: Early stage in convergence to PO

Figures 11 and 12 show, respectively, an early round in the search for competitive equilibrium, in which we have attained a Pareto optimum, but one

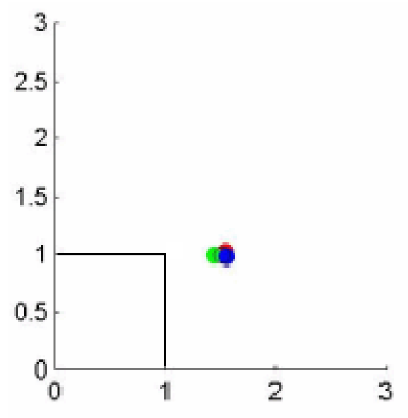

Figure 10: Final stage convergence to PO

which is not the competitive equilibrium, and the final round where the prices (defined by the normalized gradients) have converged to within a neighborhood of the competitive equilibrium of the model. Again, the vertical and horizontal lines at (1,1) show the position of the competitive equilibrium price. Animated gifs of the convergence process can be found online at

http://econ.gsia.cmu.edu/spear/po_scarf.gif13http://econ.gsia.cmu.edu/spear/po_scarf.gif

(for the convergence to Pareto optima) and

http://econ.gsia.cmu.edu/spear/scarf.gif 13http://econ.gsia.cmu.edu/spear/scarf.gif

or

http://econ.gsia.cmu.edu/spear/scarf2.gif13http://econ.gsia.cmu.edu/spear/scarf2.gif

(for the price convergence process).

5 Conclusion and Extensions

The zero- (or ε-) intelligent implementation of competitive equilibrium raises some interesting theoretical issues that we hope to explore in later research. What would the algorithm select in models with multiple equilibria? In the simple 2×2 economy, we can illustrate what may happen geometrically (see Figure 10).

In the diagram, CE_1 denotes one of the competitive equilibria of the model, while the point labeled A denotes a non-equilibrium Pareto optimum. The line

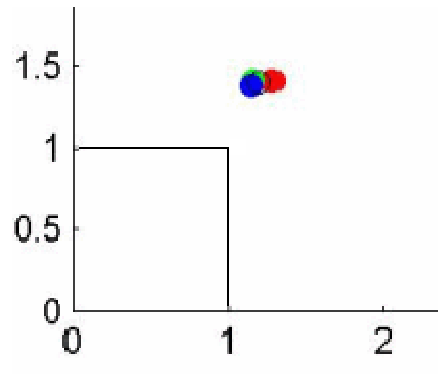

Figure 11: Convergence to CE, initial round

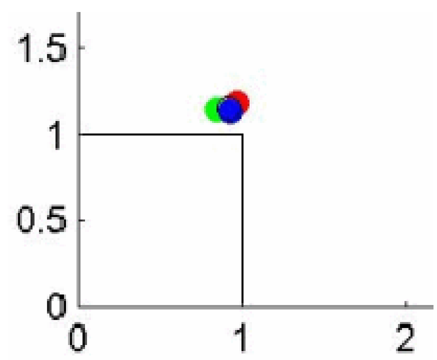

Figure 12: Convergence to CE, final round

through A supports the allocation at A, and, as drawn, implies that agent 2 is subsidizing agent 1. Under the ε-intelligent algorithm, we would seek a new Pareto optimum lying below the support line through A (i.e. we would move in the direction indicated by the downward arrow). This moves us away from the equilibrium at CE_1. But in this case, there must exist another equilibrium along the branch of the contract curve between A and the point B (which is the worst individually rational allocation agent 1 can be given). The line through B supports the allocation at B and necessarily has agent 1 subsidizing agent 2, since B minimizes the cost to agent 1 of buying any allocation at least as good as the endowment. Given convexity (and the fact that the endowment cannot be Pareto optimal if there are to be multiple equilibria), it follows that agent 1 is strictly subsidizing agent 2 at B. Hence, from this allocation, the algorithm would move us along the contract curve in the direction indicated by the upward arrow. By continuity, then, there must exist a second competitive

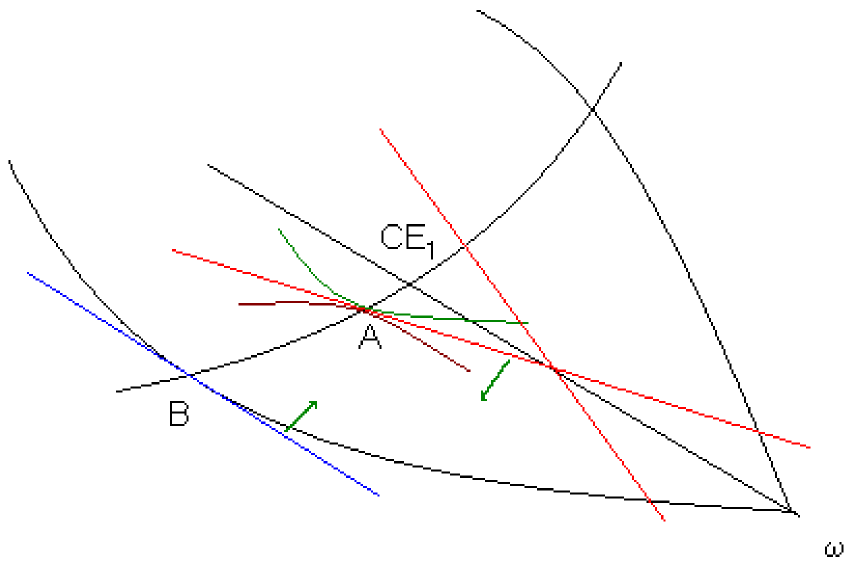

Figure 13:

equilibrium between A and B. We have run simulations of the algorithm on the well-known Blair example for the 2×2 economy (with CES preferences), and find that (as the geometry above suggests) for this economy, the two equilibria which are stable for the Walras tatonnement are the ones selected (with equal probability) by the ε-intelligent algorithm. How the algorithm would behave in a model with more goods and agents having multiple equilibria is, at this stage, an open question.

A second theoretical issue concerns the relationship between subsidization and the core. It would be interesting to know if there is any relationship between the subsidies we observe during the convergence process, and core allocations for the economy. One of the problems one typically encounters in trying to decide whether an allocation is in the core of a particular economy is that the number of coalitions grows exponentially as the number of agents in the economy increases. If there is a systematic relationship between the degree of subsidization that occurs away from the competitive equilibrium and whether an allocation is in the core, consideration of the subsidies might provide a way of calculating core allocations.

More generally, the above results provide important support for the viability of the concept of competitive general equilibrium in the context of decentralized economic institutions populated by self-seeking agents, even if they have only limited cognitive abilities. There is a reasonably simple way for the economy as a whole to arrive in arbitrarily close proximity to competitive equilibrium through individual actions of agents who seek to improve their own welfare within the constraints of the social institutions and their information and resource endowments. Although economists identify such equilibria through demonstrably non-descriptive maximization assumptions, the feasibility of arriving at, and maintaining such equilibria through the simple process of negotiating trades provides critical support for the mathematical approach. These mathematical models tend to charaterize the equilibria without specifying the decision making and institutional processes through which the outcomes might be mapped from individual actions. By combining specification of a social institution with minimally rational agents to arrive at theoretical equilibria, our results take the sting out of the behavioral critique of the neo-classical paradigm (Sunder [14]).

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 an important one. As it stands, the results we obtain here would be applicable in a setting such as the one described by Radford [8] 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 algorithm 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 simulations of Scarf's example, simple (i.e. uniform) random search for Pareto optima becomes very slow as we get closer to the Pareto set. Our resort to the gradient ascent method applied to a social welfare function with random weights strongly suggests an important role for intermediaries, particularly as the number of agents involved in trade becomes large. Some of our earlier efforts to speed up the algorithm in Scarf's example involved looking at pairwise trades between randomly selected agents. While it is easy for such pairs to come up with trades that Pareto improve for themselves, this process generally resulted in some agent getting "stuck" against a subsidization constraint while the remaining agents traded to a pairwise Pareto optimum. Avoiding this problem required simultaneous consideration of trade among all three agents in the example. 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 studying market microstructure will be important in understanding how markets get to equilibrium.

References

[1] Balasko, Y., *Foundations of the Theory of General Equilibrium*, Boston: Academic Press, 1988.

[2] Chatterji, S., and S. Ghosal, "Local Coordination in a Competitive Market", *Journal of Economic Theory*, forthcoming.

[3] Gode, D. K. and S. Sunder, "Allocative Efficiency of Markets with Zero Intelligence Traders: Market as a Partial Substitute for Individual Rationality," *Journal of Political Economy*, 101, 1 (February 1993), 119-137.

[4] Gode, D. K., S. Spear, and S. Sunder, "Convergence of Double Auctions to Contract Curve in an Edgeworth Box", GSIA Working Paper #2001-E28.

[5] Kirman, A., "The Intrinsic Limits of Modern Economic Theory: The Emperor Has No Clothes", *The Economic Journal*, Vol. 99, Issue 395, 126-139 (1989).

[6] Kumar, A. and M. Shubik, "Variations on the Theme of Scarf's Counter-Example", Cornell University Working Paper.

[7] Negishi, T., "Welfare Economics and Existence of an Equilibrium for a Competitive Economy", *Metroeconomia*, 12 (1960), 92-97.

[8] Radford, R.A., "The Economic Organisation of a P.O.W. Camp", *Economica*, Vol. XII (1945).

[9] Scarf, H., "Some Examples of Global Instability of the Competitive Equilibrium", *International Economic Review*, 1 (1960), 157-172.

[10] Simon, H., "Rationality as Process and as Product of Thought", *American Economic Review,* Vol. 68, Issue 2, 1-16 (May 1978).

[11] Shapley, S. and M. Shubik, "Trading Using One Commodity as a Means of Payment", *Journal of Political Economy*, (1977) 84, 937-968

[12] Smale, S., "Global Analysis and Economics III: Pareto Optimal and Price Equilibria", *Journal of Mathematical Economics*, 1, (1974), 1-7-117.

[13] Smale, S., “A Convergent Process of Price Adjustments and Global Newton Methods”, *Journal of Mathematical Economics*, 3,1 , (1976), 1-14.

[14] Sunder, S., ”Markets as Artifacts: Aggregate Efficiency from Zero Intelligence Traders,” in M.E. Augier and J.G. March, eds., *Modesl of Man: Essays in Memory of Herbert A. Simon*, Cambridge, MA: MIT Press, 2004.

6 Appendix

6.1 Gradient Colinearity Criterion

The near-colinearity criterion used to determine when the search for Pareto optimal allocations should be terminated is determined as follows. At each new allocation, we compute each agent's utility gradient, and then normalize it by dividing through by the marginal utility of the first good (the monotonicity assumption on preferences guarantees that this is non-zero). This determines an agent specific personal price of the form

$$p_i = \begin{bmatrix} 1 \\ \vdots \\ \frac{\frac{\partial u_i(x_i)}{\partial x_{ij}}}{\frac{\partial u_i(x_i)}{\partial x_{i1}}} \\ \vdots \\ \frac{\frac{\partial u_i(x_i)}{\partial x_{i\ell}}}{\frac{\partial u_i(x_i)}{\partial x_{i1}}} \end{bmatrix} \text{ for } i = 1, ..., M.$$

To determine near-equality of the normalized gradients, we compute the standard deviation of prices for each agent from the centroid of the agent specific prices

$$\bar{p}_j \in \mathbb{R}_+ = \frac{1}{M}\sum_{i=1}^{M} p_{ij}.$$

This yields a measure

$$\sum_{i=1}^{M}\sqrt{\frac{(p_{i2}-\bar{p}_2)^2+(p_{i3}-\bar{p}_3)^2+...+(p_{iM}-\bar{p}_M)^2}{M-1}}$$

If the sum of these standard deviations (summing across agents) is less than some small number, the routine determines that the new allocation is at a near Pareto optimum and halts. We then take $\bar{p} \in \mathbb{R}^{\ell}_{+}$ as our market price and use this to determine the implied subsidies for the allocation.

6.2 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:

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

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 tatonnement, 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.