DECENTRALIZED EXCHANGE, OUT-OF-EQUILIBRIUM DYNAMICS AND CONVERGENCE TO EFFICIENCY

SAYANTAN GHOSAL AND JAMES PORTER

ABSTRACT. In this paper, we study out-of-equilibrium dynamics with decentralized exchange (bilateral bargaining between randomly matched pairs of agents). We characterise the conditions under which out-of-equilibrium trading convergences to efficient allocations even when agents are myopic and have limited information. We show, numerically, that the rate of convergence to efficient allocations is exponential across a variety of different settings.

JEL C62, C63, C78. KEYWORDS: out-of-equilibrium, cautious, trading, efficiency, experimenting, computation

---

Date: December 15, 2010.

Ghosal: University of Warwick, Department of Economics, University of Warwick, Gibbet Hill Road, Coventry. CV4 7AL. Porter: University of Warwick, Complexity Science Doctoral Training Centre, University of Warwick, Gibbet Hill Road, Coventry. CV4 7AL. We would like to thank Peter Hammond and Alan Kirman for their comments. Porter would like to thank EPSRC for PhD funding via the Complexity Science Doctoral Training Centre at the University of Warwick.

1. Introduction

We study the limiting properties of out-of-equilibrium dynamics with decentralized exchange (bilateral bargaining between randomly matched pairs of agents)[1]. When agent's have perfect foresight, the equilibrium outcomes of decentralized exchange have been used to provide strategic foundations for competitive equilibria, see for example (Rubinstein & Wolinsky 1985, Gale 1986a, Gale 1986b, McLennan & Sonnenschein 1991, Gale & Sabourian 2005). In this paper, starting from an out-of-equilibrium scenario, we characterize the conditions under which out-of-equilibrium trading convergences to efficient allocations, and examine, numerically, the rate of convergence to efficient allocations.

In our set-up agents are myopic, have limited information about other agents and trading histories. Under assumptions on preferences that ensure pairwise optimal allocations are also Pareto optimal, we show that limit allocations must be efficient as long as traders trade cautiously (they propose and accept trades that improve their utility evaluated at their current holdings), the proposals made are drawn from a distribution that satisfies a minimum probability weight condition and the underlying trading process is connected (any pair of agents meet with positive probability after any history of matches). In an example we, then, show that trade may not converge to an efficient allocation if the minimum probability weight condition fails to be satisfied. Straightforwardly, our results extend to the case of production once Rader's principle of equivalence (Rader 1976) is invoked. Numerically, in economies where agent's preferences can be represented by Cobb-Douglass utility functions, we show that the rate of convergence to efficient allocations is exponential even as we vary both the number of agents and the number of commodities. We are also able to show, numerically, that the distribution of initial wealth and final wealth (initial and final endowments evaluated at limit prices) have a linear relationship.

Next, we turn to economies where multilateral exchange is essential for achieving gains from trade. An example of such a setting is the exchange economy studied by (Scarf 1959) with a unique competitive equilibrium that is globally unstable under tâtonnement dynamics. We show that in Scarf's example, if traders generate offers and accept proposals cautiously, trading fails to converge to efficient allocations. However, once trading is augmented to allow agents to experiment (that is accept proposals with lead to small utility loss relative to current holdings) and such experimentation is almost surely finite, we show that there is convergence to efficient allocations. Further, numerically,

[1]Our set-up could be interpreted as modelling exchange in barter economies where the underlying fundamentals are stationary.

we show that the speed of convergence remains exponential with Cobb Douglas utilities.

We, then, examine the role played by the connectedness assumption in obtaining our results. In the language of graph theory, we think of agents as vertices and potential trading links as edges that connect agents so that trade is restricted to only those agents who are connected by an edge. Straightforwardly, trade will converge to an efficient allocation as long as the underlying network of agents is connected implying that our earlier limiting results (which implicitly assumed that all agents were connected) are robust. Next, we examine numerically the impact of network structure on the rate of convergence and the link between the distribution of initial and final wealth. As long as the level of connectivity between agents is high, the average path of the economy is very similar. However, in a centralised network (a star network), we show that the speed of convergence is lowered and the link between initial and final wealth is random.

In our model of decentralised exchange, the map from action profiles to prices and allocations is well-defined both out-of-equilibrium and along the equilibrium path of play. Therefore, the properties of the out-of-equilibrium dynamics studied by us can be explicitly related to the behaviour of agents and the underlying structure of connections between agents. In contrast, classical approaches, for example (Arrow & Hahn 1971), whether tâtonnement (without explicit out-of-equilibrium trading) or non-tâtonnement (with explicit out-of-equilibrium trading) - suffer from the problem that the price adjustment and allocation dynamics isn't explicitly grounded on the behaviour of agents. Such conceptual problems have important consequences. For example, tâtonnement dynamics may not always converge. Moreover, to construct a convergent non-tâtonnement dynamics typically requires that the preferences of agents be known.

Various attempts have been made to model trade in decentralised economies. Early results (Feldman 1973, Rader 1976) characterise the conditions required for for a decentralised bilateral exchange economy to converge to a Pareto Optimal allocation. (Goldman & Starr 1982) derives generalised versions of these results for k-lateral exchange where exchange happens between groups of k agents. An alternative approach is the assumption of "zero intelligence" (Gode & Sunder 1993). Here there are a variety of computer agents, one form of which simply makes random offers subject to a budget constraint. They speculate that the "efficient" outcomes are due to the double auction market structure under investigation. Another angle is taken by Foley's work on statistical equilibrium, for example (Foley 1999), which models an economy via discrete flows of classes, that is homogeneous classes of, traders entering a market who have discrete sets of trades they wish to carry out. The result is probability distributions over trades, so as in our process agents with identical initial endowments may end up with different final allocations, but as in the many Walrasian frameworks, but unlike our approach, the trading process remains an unspecified black box. More recently (Gale 2000) has approached an out of equilibrium economy with a model with decentralized exchange in the special case with two commodities and quasi-linear utility functions.

Axtell (Axtell 2005) has explored decentralised exchange from a computational complexity perspective. He argues that the Walrasian auctioneer picture of exchange is not computationally feasible, while decentralised exchange is. While this adopts a somewhat decentralised (possibly bilateral perspective) it assumes a high level of information in the groups which are bargaining (essentially a Pareto optimal outcome for that group is directly calculated) and seems to sidestep the issue of coordinating the matching of these groups.

In a related contribution Fisher (Fisher 1981) studied a model of general equilibrium stability in which agents are aware they are not at equilibrium. In our paper, in contrast to Fisher (Fisher 1981) we do not require agents to hold their expectations with certainty and we allow for price setting by individual agents.

Gintis has looked at an agent-based model of both an exchange economy (Gintis 2006) and general equilibrium economy (Gintis 2007) although the dynamics in his models, driven by evolutionary selection, are limited to quite homogeneous agents (for example, in his exchange economy agents all have the same linear utility functions).

The remainder of the paper is structured as follows. The next section is devoted to the study of cautious trading. Section 3 presents numerical methods and results. Section 4 studies the effect of network topology on cautious trading. The last section concludes. Appendix B presents the key sections of the source code.

2. The model

We consider individuals who are aware they are in an out-of-equilibrium state and thus realise they may make mistakes if they were to attempt to condition their current trade based on their future expectations. In response to this agents may only accept trades which improve upon their current holdings or which disimprove in a limited way. We assume that the process is connected, that is at every time any given pair of agents will attempt exchange at some point in the future. We call this process, in a connected exchange economy, *cautious trading* and specify fully below.

2.1. Specification of model.

There are individuals $i \in I = \{1, \ldots, I\}$, commodities $j \in J = \{1, \ldots, J\}$ and endowments $e_i^j \in \mathbb{R}, e_i^j > 0$ of commodity j for individual i. Trade takes place in periods $t \in 1, 2, \ldots$ and we write the bundle of commodities belonging to individual i at time t as $\mathbf{x}_{it}$ and restrict these to positive bundles (you can only trade what you currently have). Agents have strictly increasing real valued utility functions $u_i(\mathbf{x}_{it})$ which are defined for all non-negative consumption bundles[2].

In each period t two agents are matched at random with equal probability that any particular pair will be selected. We will assume that once a pair is matched the two agents put up all their current holdings for exchange. One agent, the proposer, which without loss of generality is m, proposes a non-positive[3] trade $\mathbf{z}_t$ to a responder n such that:

$$x_{mt}^j > -z_t^j > -x_{nt}^j \quad \forall j$$

and

$$u_m(\mathbf{x}_{mt} + \mathbf{z}_t) > u_m(\mathbf{x}_{mt}).$$

The first condition is just that the trade would leave m and n with positive quantities of each good. The second condition is that the trade is utility increasing for m. The responder, n, will accept the trade if it improves his utility, that is

$$u_n(\mathbf{x}_{nt} - \mathbf{z}_t) > u_n(\mathbf{x}_{nt}).$$

but reject it otherwise (in which case no trade takes place). We will weaken this condition of mutual strict utility improvement shortly.

Note that the requirement that agents put up all their current holdings applies to a wide variety of cases. Firstly no agent is likely to have an incentive to conceal his holdings. An agent is free to reject any offer that is put on the table and by concealing some of his holdings the agent reduces the probability of generating a mutually improving trade[4]. Secondly a wide variety of conceptions of markets would make an agent's holding's public knowledge and this is a stronger requirement than the framework presented here which only requires that the proposer knows his current holdings, his utility function and the responder's holdings.

For the analytical results presented in this section attention is focused on the following solution of $k-$wise optimality. An allocation $\mathbf{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_I)$ is *pairwise optimal* if there exists no way of redistributing bundles between any pair m, n that would make at least one strictly better off, while making the other at least as well off. This notion can be generalised to $k-$wise optimality in the obvious way. If $k = I$ (that is if it equals the total number of agents in the economy)

[2]Formally, the trading dynamics we study in this paper has the feature that agents do not consume till trade stops. However, following (Ghosal & Morelli 2004), note that a reinterpretation of our model so that agents trade durable goods that generate consumption flows within each period will allow for both consumption and trade.

[3]That is not simply proposing a gift: it must be an bilateral exchange.

[4]Analytically for the below convergence results we could work with a weaker condition, an upper bound on trade proposals, but the form presented here will turn out to be numerically convenient, something we will return to in section 3.

then one would be considering Pareto optimality. The formulation of these concepts used here are drawn from (Goldman & Starr 1982).

Let us assume that the proposals are drawn at random from the set of all such proposer's utility improving proposals, Z, such that there is a strictly positive probability of choosing a proposal within any open set $X \subset Z$. Furthermore we will assume that the random choice of a new proposal will satisfy the following *minimal probability weight condition*: there exists some $c \in (0, 1]$ such that for all periods t the probability of choosing a proposal from any open subset X of Z is greater than cp where p is the probability of choosing a proposal in X if we choose from a multivariate uniform random distribution over Z. We could actually use the weaker condition that for some strictly positive proportion of periods the original condition holds, however (at least with respect to analytical results) this would in effect mean ignoring the other periods.

The following example makes clear the crucial role of the *minimal probability weight condition* in obtaining convergence to pairwise optimal allocations.

Example 1. *Suppose there are two agents i and j. We will set up our example such that there is a non-zero probability that trade will never occur. Consider i's proposals to j; assuming that no trade occurs the set, Z, that these are drawn from will not vary with time. Furthermore we can partition the set of improving trades Z into Z_i the set of individually improving but not improving to j trades and $Z_{(i,j)}$ the set of mutually improving trades. Assume we are not at a pairwise (in this example trivially Pareto) optimal allocation and that Z_i is also non-empty. Now assume that i draws its proposals from a fixed probability distribution for each proposal in this particular state. That is it picks a $z \in Z_i \cup Z_{(i,j)}$. Now let p^i be the probability it picks a proposal in Z_i and $p^{(i,j)}$ be the probability it picks a proposal in $Z_{(i,j)}$. It has been assumed that there is a strictly positive probability of choosing a proposal within any open set $X \subset Z$, so this applies in particular to Z_i and $Z_{(i,j)}$.*

Now consider a new process where we transform the probability distributions over the disjoint sets Z_i and $Z_{(i,j)}$ by a constant scaling such that $p_t^i = (1 - \tau_t)p^i$ and $p_t^{(i,j)} = \tau_t p^{(i,j)}$ where τ_t is given by the sequence $\tau_t = \frac{1}{2^{t+1}}$ for time periods $t = 1, 2, \ldots$. We make no restrictions on the behaviour if we were to leave the initial state and claim that there is now a positive probability that trade will never occur so a fortiori we will not converge to a pairwise/Pareto optimal.

To see this consider the probability of at some point proposing a trade in $Z_{(i,j)}$, that is one which will be accepted. This is strictly less than $p^{(i,j)} \sum_t \tau_t = \frac{p^{(i,j)}}{2}$, which implies there is a non-zero probability that trade will never occur. Actually to complete this argument we need j to propose in the same way. If both agents are proposing in this fashion then there is a non-zero probability that trade, and hence any kind of *convergence, will never occur. Note that is is possible to generalise this to a larger number of agents by using the same weights on each distribution of proposals of i to any agent k.*

While this example is somewhat pathological it illustrates an important point. For cautious trade to work we can't have agents conditioning their actions on the period in a way which essentially rules out trade at all, or via a limiting process. [5]

A benchmark process that satisfies the above conditions and which will satisfy our solution concepts is a set of n agents, with Cobb-Douglas utility functions and interior endowments. In Lemma 1 this is proven for a more general set of scenarios and in section 3 further results are obtained numerically for this and more sophisticated cases.

So far a basic model of bilateral trade has been introduced; however there is one more major element to add to the model. Below it is shown how the process that has been defined along the lines proposed above may under certain circumstances fail to carry out *any* trades, even where it would benefit all individuals in such an economy to do so. Furthermore we would ideally like to include the possibility of making mistakes in the model. It turns out one can include both ideas in a straightforward way by including an "experimentation" process in the model.

One famous class of examples that show non-convergence and instability in a global competitive equilibrium is presented in (Scarf 1959). This example can be adapted for our model in a similar way to (Gintis 2007): the basic idea is that there are three classes of agents each of whom has a utility function which is the minimum of the good it has and one other; but no agent, at least initially, can find an agent with whom a mutually improving trade can take place. To specify precisely:

$$u_1 = \min(x^1, x^2) \text{ with endowment } e_1 = (1, 0, 0)$$
$$u_2 = \min(x^2, x^3) \text{ with endowment } e_2 = (0, 1, 0)$$
$$u_3 = \min(x^1, x^3) \text{ with endowment } e_3 = (0, 0, 1)$$

This means that in the model proposed above, and similar models, no trade will ever take place. However once one introduces a small probability ϵ of experimenting, that is proposing or accepting a disimproving trade (either deliberately or through making a mistake) then trade will take place and outcomes which are Pareto improvements over the initial state can be attained. In figure 8 in section 3 this is numerically illustrated.

One can incorporate a more general form of experimentation in a relatively straightforward way which can include both making mistakes

[5]Note that in this example we have assumed that agent i needs to know the utility function of agent j. One could weaken this to an assumption of the knowledge of the forms of utility functions over an economy as a whole.

and a heterogeneous, limited amount of experimentation on the part of agents. For each agent an experimentation function $f_i(t)$ from current period in $t \in \mathbb{N}$ to a probability in $[0, 1]$ is required. Furthermore the limit as $t \to \infty$ should be 0; this is the probability in a given period that experimentation will take place. Also required is a further function $h_i(t)$ which determines the loss in utility that is acceptable in a given period (that is loss in utility for each agent engaged in trade), subject to a similar condition that the limit a $t \to \infty$ is 0, that is no loss is deemed acceptable at the limit. For most of the analytical results the heterogeneity is unimportant. The realised loss by agent i in period t is ϵ_{it} where is there is a gain in utility by i we set $\epsilon_{it} = 0$.

Total experimentation is almost surely finite if the composite process described above leads to a total loss across periods t to all agents that is bounded with probability one. Formally, with probability one, the sum over losses $\sum_{i=1}^{i=I} \sum_{t=1}^{\infty} \epsilon_{it} \leq H$ for some finite H. Any form of experimentation which ceases in finite time will trivially satisfy the above.

To clarify these concepts consider the following examples:

(1) Experimentation is fixed at a level $\tilde{\epsilon}$ for all times, that is $h(t) = \tilde{\epsilon} \forall t$ and probability of experimentation is also fixed at some strictly positive value. In this case experimentation is not almost surely finite.
(2) $h(t) = 1/t$ and $f(t) = \tilde{p}$ is any function that satisfies the above definition. In this case experimentation is almost surely finite as the sum of the bounds on losses converges.

All of the key concepts have now been introduced. So to state in full, a set of n agents with strictly positive endowments carry out *Cautious Trading* if:

(1) In each period $t$two agents i, j are matched at random (with equal probability of any particular pair being selected).
(2) Agent i proposes a non-positive trade such that $x^j_{mt} > -z^j_t > -x^j_{nt} \ \forall j$.
(3) The trade either improves i's utility level or with probability $f_i(t)$ reduces it by at most $h_i(t)$.
(4) The trade is accepted if it either improves j's utility level or with probability $f_j(t)$ reduces it by at most $h_j(t)$.
(5) The experimentation process defined by the collection $\{f_i\}_{i \in I}$ and $\{h_i\}_{i \in I}$ collectively define an experimentation process which is almost surely finite.
(6) The proposals satisfy the minimal probability weight condition.

2.2. **Results.**

Lemma 1. *The Cautious Trading, with the limit for experimentation set at* 0 *(for all agents, for all time), converges in utility and the allocations converge to a set of pairwise optimal utility-identical allocations.*

Proof. We know that $u_i^{t+1} \geq u_i^t$ for any agent i as only mutually utility increasing trades will be made as experimentation is set at zero. Furthermore the sequence of utility values is bounded as the set of feasible allocations is compact (the sum of all goods must be the sum of the endowments) and a maximum utility value for each agent is the value when it has all of all goods. So for each agent i the sequence of utility values $u_i(\mathbf{x}_{it})$ converges to its supremum; call the vector of these $\bar{\mathbf{u}}$.

Now consider the sequence of allocations $\mathbf{X}_t$ generated by cautious trading. We claim that any limit points of such a sequence must be pairwise optimal allocations with utilities $\bar{\mathbf{u}}$. Suppose it wasn't then by definition there would exist a pair of agents i, j and trade vector $\mathbf{z}$ such that

$$u_i(\mathbf{x}_i^t + \mathbf{z}) > u_i(\mathbf{x}_i^t)$$

and

$$u_j(\mathbf{x}_i^t - \mathbf{z}) > u_j(\mathbf{x}_i^t)$$

But by assumption there is a strictly positive lower bound on the probability of picking a trade within every neighbourhood of $\mathbf{z}$ in every period. By continuity there exists some such neighbourhood of $\mathbf{z}$ which pairwise improves (there may in fact be additional regions of our allocation space where this holds) so we know that a trade will almost surely happen at some point in the future between these two agents and so this cannot be an allocation at $\bar{\mathbf{u}}$. □

It should be noted that the limit allocation is path dependent, there is no unique pairwise optimal allocation; though it would be possible to define a process which had such a feature, it would necessitate greatly constraining possible exchanges.

Corollary 1. *If after some finite time an exchange process begins cautious trading as in Lemma 1, then it will converge to a pairwise optimal allocation.*

So we could have any kind of initial experimentation process or trading conditioned on future expectations based on empirical distribution of trades and still obtain the same result if eventually cautious trading with zero experimentation commences.

Proposition 1. *Cautious trading converges with probability one to a set of Pairwise optimal allocations.*

Proof. For a particular realisation let $\mathbf{x}_i^t$ be the current allocation of agent i at time t, u_i^t the utility of agent i at time t. Let ϵ_{it} be the loss in utility to agent i in period t. (If no experimentation occurs in period t for agent i then $\epsilon_{it} = 0$ as before.) By assumption the total amount of experimentation of all agents is almost surely finite, so for any particular agent the sum of ϵ_{it} is also finite.

Let the sequence v_i, indexed by t, be given by $v_i^t = u_i^t + \sum_{k=1}^{t} \epsilon_{it}$. Then this new sequence v_i is increasing. It is also bounded as it is the sum of two bounded sequences. Therefore it converges to a limit, say $\tilde{v}_i$. But this implies that u_i also converges to some limit $\tilde{u}_i$.

Now consider once more allocations at this limit $\tilde{u}_i$. They must be pairwise optimal as if they weren't then a pairwise improving trade would be made at some point in the future, even without experimentation to perturb the state. □

The following proposition shows that under the assumptions made cautious trading will get arbitrarily close to the Pareto frontier in finite time.

Proposition 2. *(i) If the utility functions are continuously differentiable on the interior of the consumption set a Pairwise optimal allocation is Pareto optimal. (ii) If indifference surfaces through the interior of the allocation set do not intersect the boundary of the allocation set then if one agent has some of all goods and others have some of at least one good then cautious trading without experimentation converges to a Pareto optimal.*

Proof. Let $\mathbf{X}$ be a pairwise optimal allocation in the interior of the allocation set. If we are in the interior of the allocation set then by assumption marginal rates of substitution exist for each agent i and for each pair of goods m, n. These must be equal for every pair of agents i, j otherwise a pairwise improvement would be possible. So they must be equal for all agents which implies that the allocation $\mathbf{X}$ is Pareto optimal.

From proposition 1 we know that the sequence converges to a set of Pairwise optimal states, so under the extra conditions imposed above it converges to a set of Pareto optimal states.

Now consider the case where one agent has some of all goods, without loss of generality let this be agent 1 and others have some of at least one good, without loss of generality less this be good 1. We need to establish that the process reaches the interior of the the allocation set then the result follows by the above argument.

Consider an agent $i \neq 1$ and set non-empty set $M = \{m \in J | x_i^m = 0\}$. As the indifference curves through the interior of the allocation set for all agents do not intersect the boundary of the set it is always in the agent's interest to accept a trade away from the boundary, that is a trade z such that for each $m \in M$, $z_i \neq 0$. When paired with agent 1 there exists an open set of trades Z which leaves him with some of all goods and improves the utility of agent 1. So eventually such at trade will happen. This argument trivially extends to all agents on the boundary, so with probability one in finite time we will reach an allocation in the interior of the allocation set. □

Even if we augment the trading process with the possibility that agents may trade to boundary allocations, subject to the conditions of continuity and strict monotonicity this will never occur under the conditions specified below.

Corollary 2 (First Welfare Theorem for Cautious Trading)**.** *If utility functions are continuously differentiable and indifference curves in the interior of the allocation set do not intersect the boundary, the process of Cautious Trading will with probability one both*

(1) *not go to an allocation on the boundary*
(2) *and will converge to a set of Pareto Optimal allocations.*

Proof. To go to an allocation on the boundary with the above conditions an agent must in effect accept an infinite loss in utility; but with probability one this will not occur as it is assumed that the total amount of experimentation is almost surely finite, so the loss in any particular period must also be bounded.

If the utility functions are continuously differentiable then any pairwise optimal allocation is a Pareto optimal allocation as the marginal rates of substitution of goods for each agent must be equal. By proposition 1 the process converges to a set of pairwise optimal allocation allocations, so with the additional assumption this is Pareto optimal. □

2.3. **Extension to Production.** Our convergence results for exchange can be extended to economies with production using the process described in (Rader 1964, Rader 1976). Formally an exchange economy is an array $\{(u_i, \mathbf{e}_i, \mathbb{R}^J_+)\colon\ i \in I\}$. An economy with production is an array $\{(u_i, \mathbf{e}_i, \mathbb{R}^J_+) : i \in I; (Y^f) : f \in F, \theta_{if} : f \in F, i \in I\}$ where $f \in F = \{1 \dots F\}$ is the set of firms and θ_{if} is individual i's share in firm f with $\sum_i \theta_{if} = 1, \forall f$. Assume that the production set Y^f of firm f is convex, non-empty, closed, satisfies the no free lunch condition ($Y^f \cap \mathbb{R}^J_+ \subset \{0\}$), allows for inaction (that is $0 \in Y^f$), satisfies free disposal and irreversibility (that is if $y \in Y^f$ and $y \neq 0$ then $-y \notin Y^f$). We can convert an economy with production to an economy with household production by endowing each individual i with a production set $\tilde{Y}_i = \sum_f \theta_{if} Y^f$. Next, by using Rader's principle of equivalence (Rader 1976), an economy with household production can be associated with an equivalent economy with pure exchange with indirect preferences defined on trades. The conditions under which pairwise optimality implies Pareto optimality with such indirect preferences follow directly from Theorem 2 and its applications, also in (Rader 1976).

3. Numerical Results

While we have shown that the sequence of allocations will converge to a Pareto optimal set, this does not answer the question of how long such a process will take to get close to Pareto optimal. This section examines this question via a numerical approach, showing that for a common class of utility functions, the average speed of convergence is, in a sense to be specified shortly, good[6]. This section also examines the question of how the cautious trading process effects the wealth of agents.

3.1. Numerical Model.

Attention is focused on sets of heterogeneous agents with Cobb-Douglas preferences and random initial endowments as a benchmark case. We can represent the preferences by utility functions:

$$u_i(\mathbf{x_i}) = \sum_j \alpha_i^j \ln(x_i^j).$$

One can of course represent Cobb Douglas utilities by $u_i(\mathbf{x_i}) = \prod_j (x_i^j)^{\lambda_i^j}$. However, the logarithmic representation is preferred for numerical work because it has a considerably lower computational cost. We have initial endowments, $\mathbf{e}_i^j$, of each commodity drawn from a uniform distribution over $(0, 1]$ and parameters α_i^j of the functions are again drawn from $(0, 1]$ uniformly, then normalised such that the sum, $\sum_j \alpha_i^j = 1$. They are normalised to a fixed value so as to make talking about global utility as the sum of agent's utilities more meaningful; this does not change the preferences which they represent.

As before trades are restricted to the set of all trades which leave both proposer i and responder j with positive quantities of each good, that is:

$$-x_{mt}^j < z_t^i < -x_{nt}^i$$

as to actually implement the trading process it is necessary to fix some boundary values[7]

[6]This section has been written so as to be as accessible as possible to the non-programmer. Those with experience of programming may wish to skim this section, while consulting the source code directly, the key sections of which are included in appendix.

[7]An alternative, and in some ways more satisfying alternative (as it limits required information), might be to restrict trades to within the total endowment of the economy. While analytically we would obtain the same asymptotic results, numerically it would simply lead to many rejected proposal and vastly longer running times if these were simulated directly. One could try and simulate the proposal process indirectly if one could formulate joint probability distributions over improving offers, over improving proposals and over agent pairing. However for anything other than trivial economies this is extremely difficult due to the number of dimensions and changing state when proposals accepted.

The key objects we need in our computational model is an agent and a collection of agents. The former implements agents with Cobb-Douglas utility functions as specified above, random initial endowments and importantly specifies the actual mechanics of trade proposals, acceptance or rejection and trades. The later creates a collection of these agents and carries out realisations of the economy. A schematic representation of these classes can be found in figure 1. Utilising these we can obtain various numerical results via processes like that illustrated in figure 2.

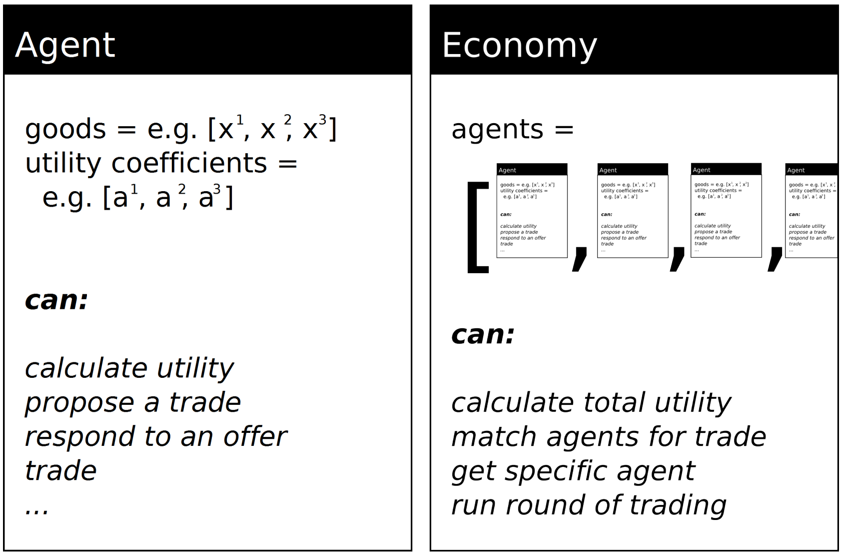

FIGURE 1. An outline of the main attributes and methods of the *Agent* and *Economy* objects.

We make one further major assumption: each agent makes one proposal per round, irrespective of the size of the economy. An alternative way to approach implementing this model might be to fix some n, perhaps $n = 1$ as the total number of proposals per round, with agents drawn at random in each round. However, if one takes seriously the decentralisation of the economy, then one should assume that the agents actions are unconstrained by the size of the global economy.

3.2. **Results.** Attention was focused on estimated convergence in average global utility to assess the performance of cautious trading. To estimate this we calculate global utility by summing across utility for all agents in the economy, then take an average over many runs as the process is stochastic. One can then use the final value as an estimate of limiting utility and calculate how far away earlier values are. The last few hundred values are discarded as for them this estimate of limiting utility is not, relatively speaking, as good. The analysis depends on the increasing nature of sequences of utility values for agents this analysis to make sense.

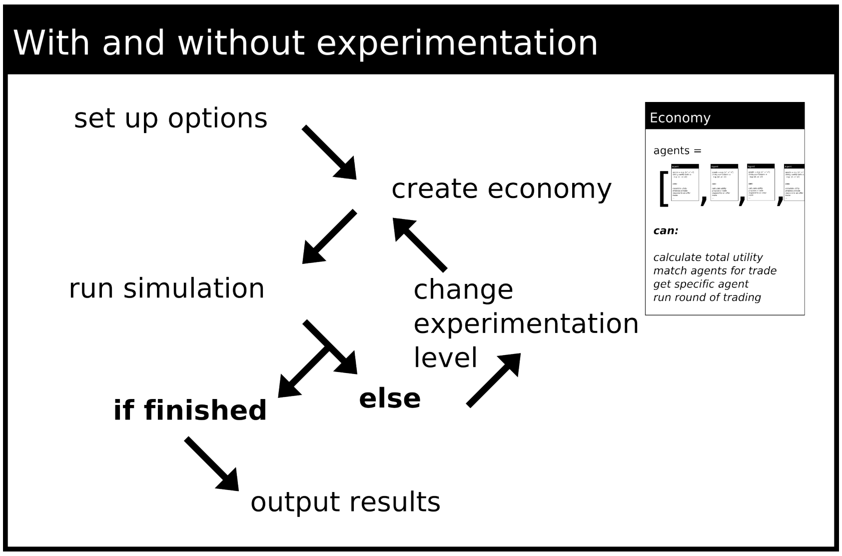

FIGURE 2. An example algorithm of a numerical simulation of Cautious Trading. The precise details vary depending the experiment being carried out but this example gives an overview of the kind of algorithm used to generate the data for most of the figures in this paper.

In figure 3 one can see how varying the total number of agents effects the average speed of convergence. As one can see there is in fact very little qualitative effect. There is some increase in time taken, however when one plots the log of average convergence as in figure 4 one can see that we get a close approximation to a straight line after an initial faster period; suggesting an exponential speed of convergence, at least over the these time periods.

We also examined the effect of the number of goods via similar analysis. In figure 5 one can see how the speed of convergence varies with the total number of goods in the economy. There is a similar result of little qualitative change. This is more surprising as we have the same number of proposals taking place as before over larger increasing numbers of goods. When one examines the the log plot in figure 6 one gets the same kind of result as for varying agents.

One can fit an exponential function, via regression on the log of the values, to these average utility paths in order to obtain a numerical estimate for the average speed of convergence. In tables 1 and 2 we present such results for a range of model sizes. The important point to note is the approximately exponential convergence in global utility for a range of sizes of economy, both in terms of number of goods and number of agents, rather than the actual fitted parameters. Note that the p-values for the regressions are less than 0.0005 indicating an extremely high level of confidence in the fit of the model.

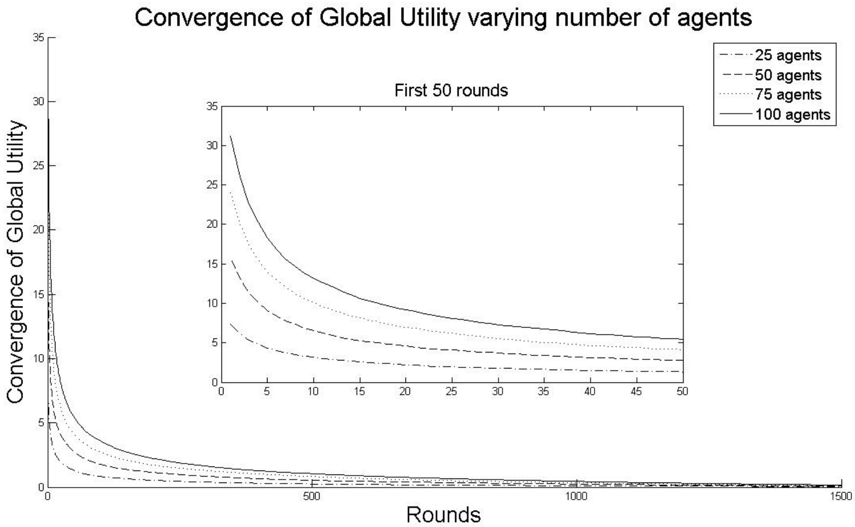

FIGURE 3. Average over many runs of global utility convergence when varying the total number of agents in the economy. *Parameters: 5 goods, 25-100 agents.*

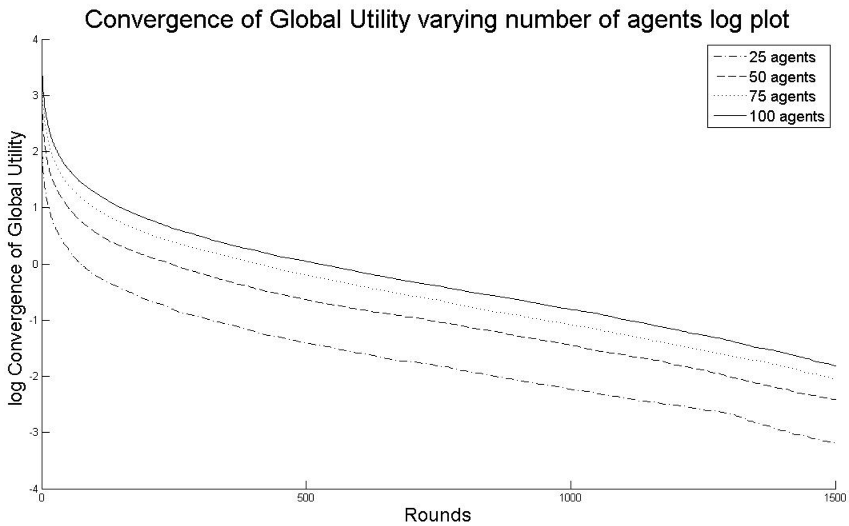

FIGURE 4. Log of average global utility convergence when varying the total number of agents in the economy. *Parameters: 5 goods, 25-100 agents.*

Another aspect of the exchange process we can analyse numerically is that of wealth dynamics or change. If we had a single set of prices or relative valuations $\mathbf{p}$ in the economy, then we could obtain the wealth of an agent i, simply be calculating $\mathbf{p}\mathbf{x}_i$. In our out-of-equilibrium scenario there is no single set of prices, however given that the marginal rates of substitution converge, this implies a convergence to a uniform

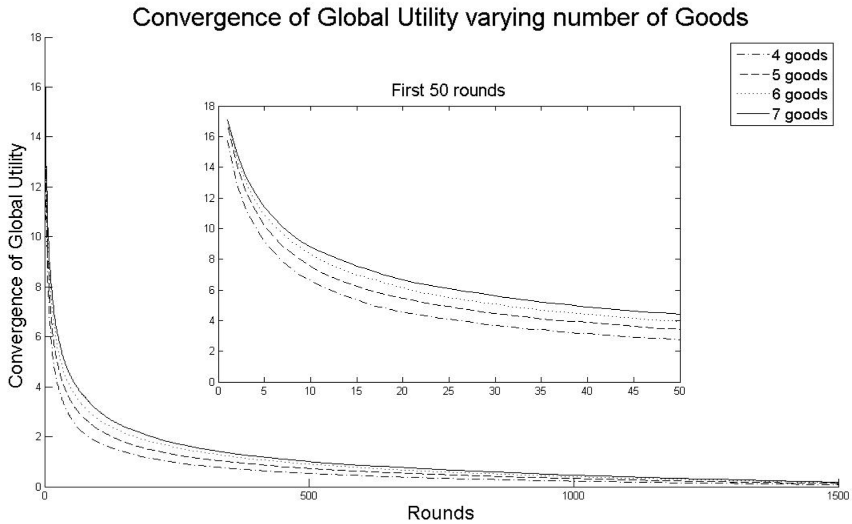

FIGURE 5. Average global utility convergence when varying the total number of goods in the economy. *Parameters: 4-7 goods, 50 agents.*

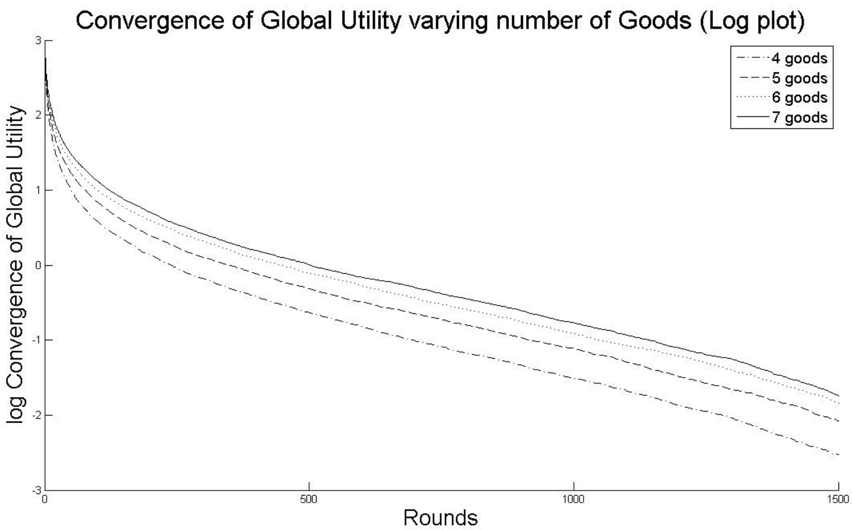

FIGURE 6. Log of average global utility convergence when varying the total number of goods in the economy. *Parameters: 4-7 goods, 50 agents.*

set of relative evaluations (in terms of changes in utility); in effect a common set of prices. From this set of prices we can calculate a value for each agent's bundle. But if we know their original endowment we can calculate their initial wealth using these prices, so we can obtain a set of *ex-post* wealth values.

Goods	4	5	6	7
linear coefficient:	-0.0026	-0.0019	-0.0021	-0.0021
constant term:	6.3642	5.6688	5.9079	6.3833

Table 1. Fitting exponential function to average convergence for varying numbers of goods. The values given are for linear fit of log of convergence. For every fit the p-values for the regressions are less than 0.0005 indicating an extremely high level of confidence in the fit of the model.

Agents	25	50	75	100
linear coefficient:	-0.0023	-0.0022	-0.0022	-0.0023
constant term:	4.4728	4.3750	5.3190	5.2299

Table 2. Fitting exponential function to average convergence for varying numbers of agents. The values given are for linear fit of log of convergence. For every fit the p-values for the regressions are less than 0.0005 indicating an extremely high level of confidence in the fit of the model.

In figure 7 the density of wealth change is plotted. We see a linear relationship of final to original wealth, but with a very high level of noise, as one might expect. It should be emphasised that in all cases utility for every agent increases over time, however 'wealth' may change in either direction.

In figure 8 we take the example outlined above in section 2.1 and examine what happens numerically, introducing a small probability ϵ of making a mistake, that is proposing or accepting a disimproving trade. If no experimentation takes place no trade ever happens and global utility remains at 0. As we increase the level of experimentation short term global utility improves (rises more steeply) at the cost of a lower level of long term convergence. In cautious trading form nothing happens, but with experimentation trade happens.

For high values of experimentation faster initial improvement than low values, but longer term global utility is slightly lower and the economy more volatile. This suggests that in selecting the level of experimentation there is a trade off between convergent level of utility and speed of convergence.

Above an example adapted from Scarf was presented which showed how experimentation could lead to a better outcome than before, however, this example is a very special case. An interesting question we can ask numerically is how experimentation effects the speed of convergence in a larger, more heterogeneous example such as the Cobb Douglas

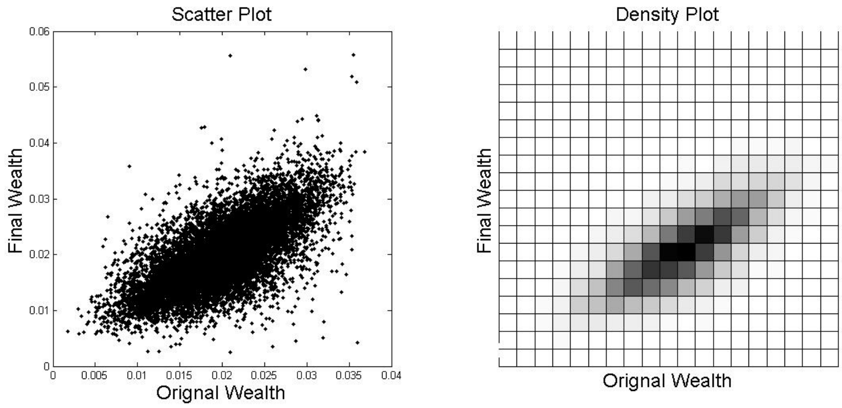

FIGURE 7. This uses the average of the final marginal rates of substitution to obtain an estimate for initial and final wealth. The darker the square the more likely such a transition from initial to final wealth in that square is; this heatmap was included as the scatter plot is difficult to read due to the high number of samples with similar changes in wealth *12500 samples (or 250 realisations of 2000 periods, with 50 agents; normalised per realisation.*

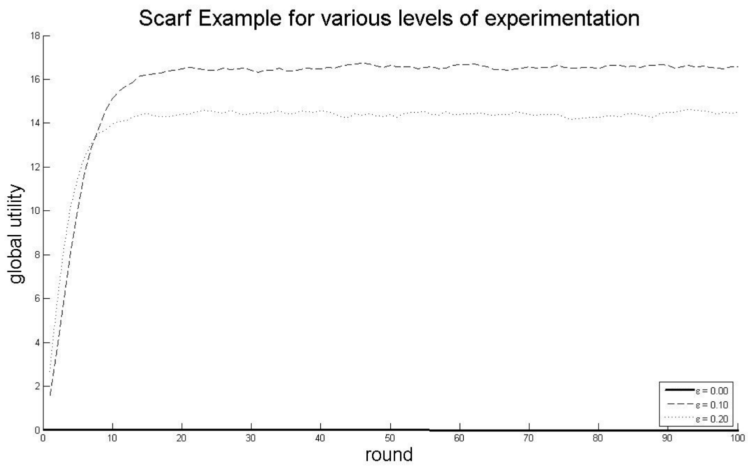

FIGURE 8. Without experimentation no trade takes place in this model adapted from a model of Scarf. Notice how long term utility appears lower for a higher level of experimentation.

utility function economy we looked at previously. In fact there is qualitatively similar long term behaviour when experimentation is included as can be seen in figure 6 where experimentation is introduced into the original model from section 3. For certain values of experimentation we even see slightly better overall performance with experimentation.

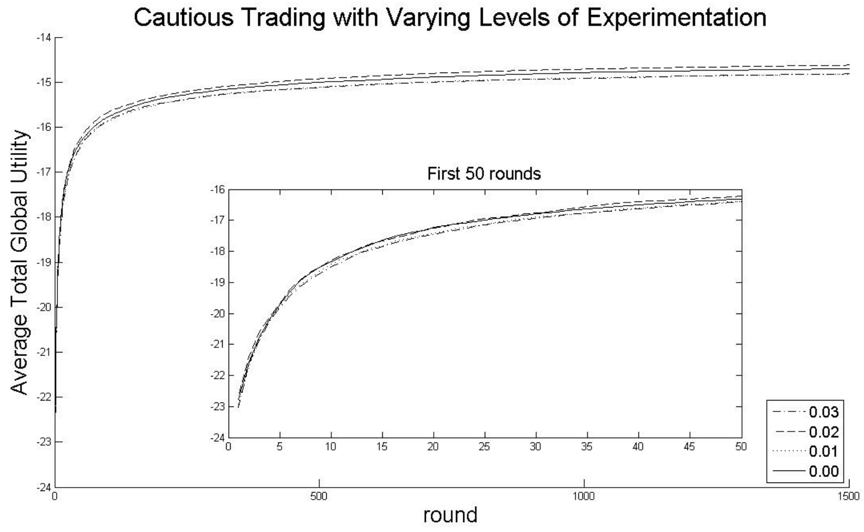

FIGURE 9. Low levels of experimentation have little effect in the Cobb Douglas economy we looked at before.

So we have seen that cautious trading allows trade to occur when it would not have otherwise happened. Furthermore, in economies where experimentation is not required, such as the Cobb-Douglas economy in figure 9, experimentation does not appear to have a qualitatively detrimental effect.

4. The Role of Connectedness

So far we have assumed a anonymous, fully connected economy, with trading partners picked at random from all agents in the economy. But an attempt to investigate decentralised economies would be incomplete without a consideration of if and how the structure of that economy effects outcomes. The natural way to think about this is in terms of a network of agents with edges representing potential trading partners[8]. So we have a undirected graph $G = (V, E)$, where V is the set of vertices (agents) and E the set of edges (potential trading links).

4.1. **Cautious Trading on Networks.** Agent i has endowment e_i as before, however we now restrict offers and trade to pairs of agents connected by an edge $e \in E$. We call this *Networked Cautious Trading.* We can use the same formulations for proposals and acceptance as before, however our *"local"* optimality will have to be redefined as follows: an allocation $\mathbf{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_I)$ is *connected-pairwise optimal* if there exists no way of redistributing bundles between any connected

[8]We consider only static networks, however if one is conceiving of an exchange economy where the cautious trading process is one step repeated with changes to the fundamentals of the economy at each step, then there is no reason why the network topology couldn't be considered a fundamental to be altered at each step.

pair m, n that would make at least one strictly better off, while making the other at least as well off. If we have a fully connected[9] graph then we would be considering *Pareto* optimality.

4.2. **Results.** We can reformulate the above analytical results for Cautious Trading in the context of networks, as summarised in the following propositions. Proofs are omitted as they can be obtained by replacing the notion of pairwise optimality with connected-pairwise optimality and taking account of networks specific issues such as requiring the network be connected[10].

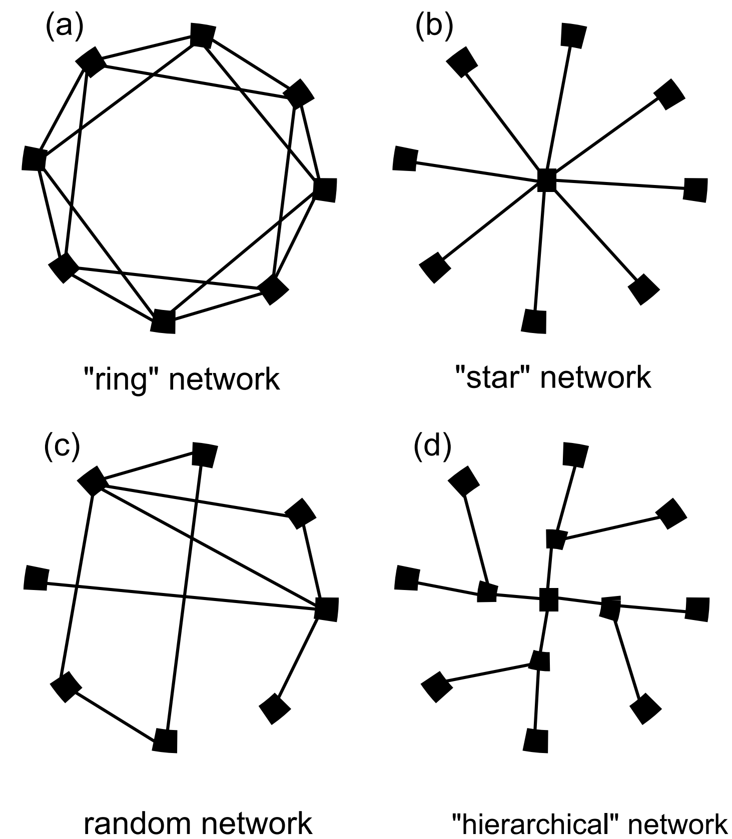

FIGURE 10. This figure illustrates some network structures. Note that simulations were run on networks of much larger size.

(1) The Networked Cautious Trading process with zero experimentation converges in utility and the allocations converge to a set of connected-pairwise optimal utility-identical allocations.
(2) If we also have that:

[9]There is an edge linking every agent to all all other agents.
[10]There exists a path from every vertex to all other vertices.

   (a) the utility functions are continuously differentiable and the graph G is connected and
   (b) indifference surfaces through the interior of the allocation set do not intersect the boundary of the allocation set,

   then then on the interior of the consumption set a pairwise optimal allocation is Pareto optimal and if one agent has some of all goods and others have some of at least one good then cautious trading converges to a Pareto optimal.
(3) If total experimentation is almost surely finite then networked cautious trading converges with probability one to a set of connected-Pairwise optimal allocations.
(4) If furthermore utility functions are continuously differentiable, indifference curves in the interior of the allocation set do not intersect the boundary and G is connected the process of Cautious Trading will with probability one, both:
   (a) not go to an allocation on the boundary
   (b) and will converge to a set of Pareto Optimal allocations.

So asymptotically Cautious Trading on a network has similar properties to anonymous, pairwise matching. However, there may be significant differences in properties like initial convergence which we are able to examine numerically. In figure 10 some different structures of networks are illustrated: a ring network, where each vertex has an edge joining it to each of its nearest neighbours; a star network where one vertex is joined to all others; a random network and a hierarchical network which has a hierarchy of star-like components. All these networks are connected, but far from fully connected as was assumed to be the case for the original formulation of cautious trading. In any case if an economy were not connected then we would really be dealing with two or more economies.

The key issue seems to be the level of connectivity as illustrated in figures 11 and 12. In the first figure extremely similar results are obtained for the original cautious trading process, cautious trading on a "ring" network and cautious trading on a random network. In figure 12 the results are quite different for networks with lower levels of connectivity. In summary, if connectivity is low, as is the case for a "star" network the speed of convergence is reduced; if it is sufficiently high (for the sizes of networks examined in this paper this means around four edges per vertex) then the results are similar to the fully connected scenario, with the actual structure of connection having little effect.

We looked at varying levels of clustering using the Watts and Strogatz model for small world graph generation (Watts & Strogatz 1998)); basically we start off with a ring network (agents are connected to their nearest neighbours) and rewire each edge with a fixed probability β. We obtained similar results to figure 11, that is there was little effect at a global level if we kept connectivity constant (as the Watts-Strogatz model does by construction).

Examining wealth change in a centralised economy (a star network) there is a substantial contrast with the original cautious trading model, now final wealth appears to be more random, with only a slight correlation with initial wealth as can be seen in figure 13. Again there would appear to be a high level of noise in the system, with many outliers.

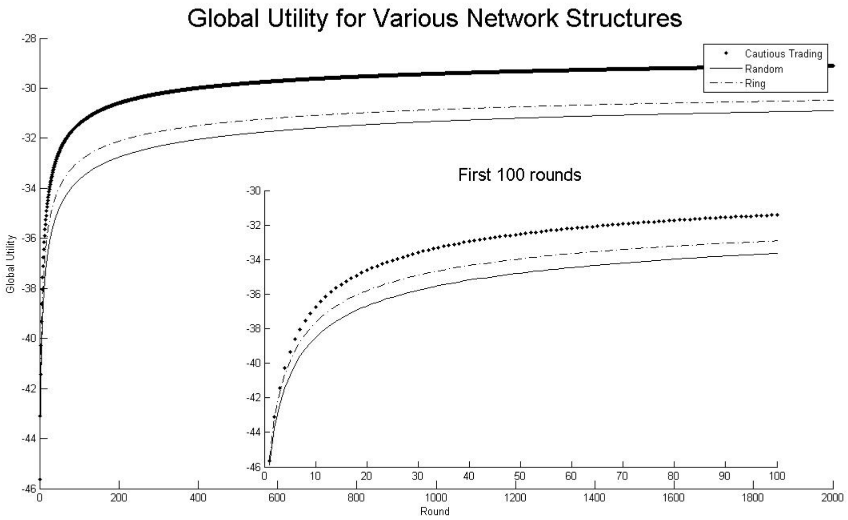

FIGURE 11. Here we have three quite different structures for our economy; however the average path of the economy if very similar. We have the original Cautious Trading, a uniformly randomly connected graph (with an average of four edges per vertex) and a ring graph (each agent is connected to it four nearest neighbours). These all have quite different properties but seemingly due to the reasonably high level of connectivity result in the same average behaviour. *Averaging over* 2000 *realisations,* 100000 *proposals per realisation,* 50 *agents.*

5. Conclusions

Even with "zero information" an exchange economy with typical assumptions will converge to a Pareto optimal outcome purely through bilateral exchange among uninformed partners. It is possible to numerically examine the speed of convergence which turns out to be exponential for a typical class of utility function. Augmenting this process with experimentation leads to both convergence in some examples where it did not previously occur and potentially faster convergence in cases which did converge previously.

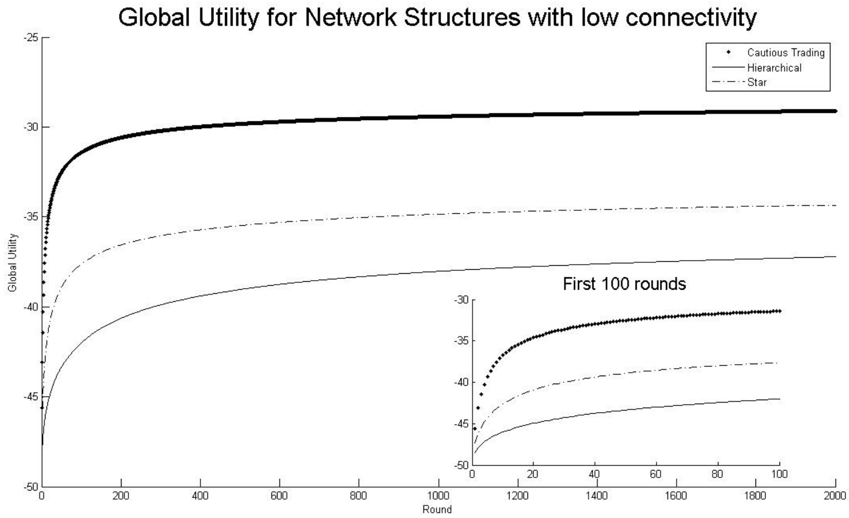

Figure 12. In contrast to figure 11 we see quite different average behaviours for these networks versus the original cautious trading. These networks are a star network in which one agent is connected to all others (an idealisation of a central market in which all agents exchange goods) and a hierarchical network (where there is a 'central market' which is connected to a small number of other 'markets', in turn connected to all the other agents in the economy). These networks have low levels of connectivity (roughly one edge per agent) and this seemingly restricts the speed of converge and lowers global utility. *Averaging over* 2000 *realisations,* 100000 *proposals per realisation,* 50 *agents.*

One can conceive of this "zero information" as a worst case assumption. In "real" markets one presumably has more to work with but almost never the kind of complete information that is typically assumed in comparable models of exchange. The dual discipline of having to deal with decentralisation and its resultant lack of information (not simply uncertainty over a small number of possible states of the world) and having to explicitly implement the models for numerical investigation has proved useful. A possible next step would be to examine out-of-equilibrium dynamics in asset trades.

References

K. Arrow & F. Hahn (1971). *General Competitive Analysis.*

R. Axtell (2005). 'The Complexity of Exchange'. *The Economic Journal* **115**(504):F193–F210.

A. Feldman (1973). 'Bilateral Trading Processes, Pairwise Optimality, and Pareto Optimality'. *The Review of Economic Studies* **40**:463–473.

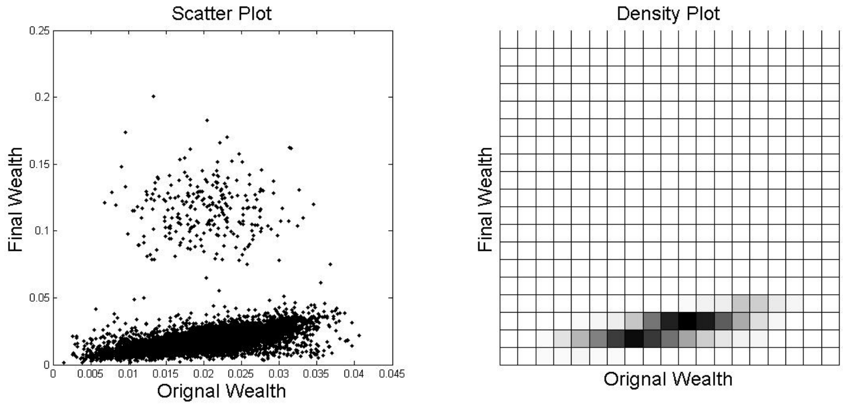

FIGURE 13. These plots show the change in wealth, using the average of the final marginal rates of substitution to obtain an estimate of changes in wealth for a networked (star network) economy *12500 samples (or 250 realisations of 2000 periods, with 50 agents; normalised on a per realisation basis such that the total wealth sums to one.*

F. M. Fisher (1981). 'Stability, Disequilibrium Awareness, and the Perception of New Opportunities'. *Econometrica* **49**(2):279–317.

D. K. Foley (1999). 'Statistical Equilibrium in Economics: Method, Interpretation, and an Example' .

D. Gale (1986a). 'Bargaining and Competition Part I: Characterization'. *Econometrica* **54**:785–806.

D. Gale (1986b). 'Bargaining and Competition Part II: Existence'. *Econometrica* **54**:807–818.

D. Gale (2000). *Strategic Foundations of General Equilibrium.* Cambridge University Press.

D. Gale & H. Sabourian (2005). 'Competition and Complexity'. *Econometrica* **73**:739–769.

S. Ghosal & M. Morelli (2004). 'Retrading in Market Games'. *Journal of Economic Theory* **115**:151–181.

H. Gintis (2006). 'The Emergence of a Price System from Decentralized Bilateral Exchange'. *Contributions to Theoretical Economics* **6**(1):1302–1302.

H. Gintis (2007). 'The Dynamics of General Equilibrium'. *Economic Journal* **117**(523):1280–1309.

D. K. Gode & S. Sunder (1993). 'Allocative Efficiency of Markets with Zero Intelligence Traders: Market as a Partial Substitute for Individual Rationality'. *The Journal of Political Economy* **101**(1):119–137.

S. M. Goldman & R. M. Starr (1982). 'Pairwise, t-wise and Pareto Optimalities'. *Econometrica* **50**(3):593–606.

A. McLennan & H. Sonnenschein (1991). 'Sequential Bargaining as a Noncooperative Foundation for Walrasian Equilibrium'. *Econometrica* **59**(5):1395–1424.

T. Rader (1964). 'Edgeworth Exchange and General Economic Equilibrium'.

T. Rader (1976). 'Pairwise Optimality, Multilateral Optimality and Efficiency, With and Without Externalities'. In *Economics of Externalities.* Academic Press.

A. Rubinstein & A. Wolinsky (1985). 'Equilibrium in a Market with Sequential Bargaining'. *Econometrica* **53**:1133–1150.

H. E. Scarf (1959). 'Some Examples of Global Instability of the Competitive Equilibrium' (79).

D. J. Watts & S. H. Strogatz (1998)). 'Collective dynamics of 'small-world' networks'. *Nature* **393**:440–442.

Appendix A. Numerical Implementation

The numerical model was implemented using the Java programming language. The implementation explicitly models individual agents via an Agent class. Each instance of the class stores the agent's current bundle of goods, the parameters of its utility function and its current marginal rates of substitution. Each agent can make or consider offers, carry out trades and reset itself for another realisation of trading.

The following subsections outline the details of the models and implementation. The key source files are contained in appendix B below.

A.1. Cautious Trading.

Two files contain the key parts of the implementation of Cautious Trading: the *Agent* and *CautiousEconomy* classes. The former implements agents with Cobb-Douglas utility functions, random initial endowments and specifies the mechanics of trade proposals and trades. The later creates a collection of these agents and carries out simulated runs of the economy.

A.2. Scarf Example.

A modified version of the *Agent* and *CautiousEconomy* classes was created to study the behaviour of an economy which in many settings may not converge. The implementation is broadly similar to the original, the main changes being to the endowments and utility functions.

A.3. Experimentation.

The *CautiousEconomy* has been augmented with the possibility of experimentation. Essentially the *ExperimentingEconomy* class adds experimentation to *CautiousEconomy* via a scaling parameter to proposed trades. To be more precise an initial level of allowable experimentation is selected and the allowable level decreases linearly until it ceases. The probability of experimentation is fixed at an initial level and this too decreases over time.

Appendix B. Source Code

This section contains the key source code files; many more were actually used to model the cautious economy. The code is arranged into four distinct levels: agent, economy, experiment and simulation. The first two play obvious roles, the experiment code provides general code to investigate the cautious economy and the simulation code runs experiments and does some processing of results. Figures in this report were then produced using Matlab.

B.1. Agent Code.

The below code is for the basic form of the Agent class.

Listing 1. Agent class source code

```
package ac.decentralisedExchange.cautious.model;

import java.util.Random;
```

```java

public class Agent {

  protected double goods[];
  protected double originalGoods[];
  protected double exponents[];
  protected double originalExponents[];
  protected double currentUtility;
  protected int nGoods;
  protected Random gen;

  public Agent(int nGoods){
    this.nGoods = nGoods;
    gen = new Random();

    goods = new double[nGoods];
    exponents = new double[nGoods];

    originalGoods = new double[nGoods]; //initial endowment stored for restart
    originalExponents = new double[nGoods]; //initial exponents stored for restart
       initializeRandomly();//actually initialise these arrays

    update();
  }

  /**
  * Initialise the agent with a random set of exponents and goods; then normalising the exponenents
  */

  protected void initializeRandomly() {
      for(int i=0; i < nGoods; i++){
          goods[i] = gen.nextDouble();
          originalGoods[i] = goods[i];

          exponents[i] = gen.nextDouble();
          originalExponents[i] = exponents[i];
      }
      normalise();
  }

  /**
   * Reset the agent, i.e. generate new endowments and exponents
   */
  public void reset() {
      initializeRandomly();
    update();
  }
```

```

  /**
   * Restart the agent, i.e. restore endowments and
       exponents
   */
  public void restart() {
    restore();
      update();
  }

  /**
  * Update utility when this is necessary. Should add any
      other update
  * actions here.
  */
  protected void update(){
    updateUtility();
  }

  /**
   * Restore original state of agent
   */
  protected void restore() {
      for(int i=0; i < nGoods; i++){
        goods[i] = originalGoods[i];
        exponents[i] = originalExponents[i];
      }
      normalise();
  }

  /**
     * Normalise the utility function so exponents sum to
         1.
     */
  protected void normalise() {
    double sum = 0.0;
    for(int i=0; i < goods.length; i++){
      sum += exponents[i];
    }
    assert sum != 0.0;
    for(int i=0; i < goods.length; i++){
      exponents[i] = exponents[i]/sum;

    }
  }

  /**
     * Return utility of an agent (Cobb-Douglas)
     */
  public double utility() {
    return utility(this.goods);
  }
```

```

  /**
     * Assess the utility of bundle (Cobb-Douglas)
     */
  public double utility(double[] bundle) {
    double u = 0.0;
    for(int i = 0; i < goods.length; i++){
      u += (Math.log(bundle[i]) * exponents[i]);
    }
    assert(! Double.isNaN(u));
    return u;
  }

  /**
  * Assess the utility if give a (i.e. subtract a) bundle.
  */
  protected double utilityIfGive(double change[]) {
    double[] temp = new double[nGoods];
    for(int i=0; i<nGoods; i++){
      temp[i] = this.goods[i] - change[i];
    }
    return utility(temp);
  }

  /**
  * Assess the utility if get a (i.e. add a) bundle.
  */
  protected double utilityIfGet(double change[]) {
    double[] temp = new double[nGoods];
    for(int i=0; i<nGoods; i++){
      temp[i] = this.goods[i] + change[i];
    }
    return utility(temp);
  }

  /**
     * Update the current utility level - call this if you
         change the bundle or exponents
     */
  protected void updateUtility() {
          currentUtility = utility();
  }

  /**
       * An agent makes a proposal to another Agent other.
       * @param The agent to propose offer to
       * @return Whether a trade took place
       */
  public boolean propose(Agent other) {
    if(other != null){
      double proposal[] = getProposal(other);

```

```
      if (other.consider(proposal)){
        trade(other, proposal);
        return true;
      }
      else{
        return false;
      }
    }else{
      return false;
    }
  }

  /** Consider a trade of change, return true if improving, false otherwise*/
  public boolean consider(double change[]) {
    if(utilityIfGet(change) > currentUtility){
      return true;
    }
    else{
      return false;
    }
  }

  public boolean propose(Agent other, double allowable_experimentation) {
    double[] proposal = getProposal(other);
    if (other.consider(proposal, allowable_experimentation)){
      trade(other, proposal);
      return true;
    }
    else{
      return false;
    }
  }

  protected double[] getProposal(Agent other){
    double proposal[] = new double[nGoods];
    boolean improving = false;
    int j = 0;
    while(!improving){
      for (int i = 0; i < nGoods; i++) {
        proposal[i] =  goods[i] - gen.nextDouble()*(goods[i] + other.goods[i]);
        assert(proposal[i] < goods[i]);
      }
      if(utilityIfGive(proposal) > currentUtility){
        improving = true;
      }
      j++;
    }
    return proposal;
```

```
  }


  /**
   * Consider a trade of change, return true if improving, false otherwise
   */
  public boolean consider(double change[], double allowable_experimentation) {
    if(utilityIfGet(change) > currentUtility - allowable_experimentation){
      return true;
    }
    else{
      return false;
    }
  }

  /**
   * Agent gets the bundle of goods change (some or all components may be negative i.e. they lose this)
   */
  public void get(double change[]) {
    for (int i = 0; i < nGoods; i++) {
      goods[i] += change[i];
    }
    update();
  }

  /**
   * Agent gives the bundle of goods change (some or all components may be negative i.e. they gain this)
   */
  public void give(double change[]) {
    for (int i = 0; i < nGoods; i++) {
      goods[i] -= change[i];
    }
    update();
  }

  /**
   * Trading procedure: parameters: another Agent other and the trade to take place change.
   */
  public void trade(Agent other, double change[]) {
    give(change);
    other.get(change);
  }

  /** The exponents of the agent are shocked via a normalised
  * Gaussian scaled via the shockSize parameter
```

```
   * @param shockSize The scaling to be applied to a
       normalised Gaussian
   * */
  public void shockGaussian(double shockSize) {
    for (int i = 0; i < nGoods; i++) {
      exponents[i] += this.gen.nextGaussian()*shockSize;
    }
  }

  public double[][] getMRS(){
    double[][] mrs = new double[nGoods][nGoods];
    for (int i = 0; i < nGoods; i++) {
      for (int j = 0; j < nGoods; j++) {
        mrs[i][j] = (goods[i]*exponents[j]) / (goods[j]*
            exponents[i]);
        assert mrs[i][j] != 0.0;
      }
    }
    return mrs;
  }
}
```

B.2. **Cautious Economy Code.** The below code presents the basic abstract economy class, which all economies subclass.

Listing 2. Cautious Economy source code

```
package ac.decentralisedExchange.cautious.model;
import java.io.*;
import java.util.ArrayList;
import ac.decentralisedExchange.util.*;

/** The class Economy consists of a collection
 * of independent Agents, who trade
 * via Cautious Trading.
 */
public abstract class Economy{

  public ArrayList<Agent> agents;

  public int size;
  public int nGoods;
  public int trades, period, round;

  /** Reset the economy i.e. give each agent a random
   * allocation and utility function.
   */
  public void reset(){
    for (Agent a: agents) {
      a.reset();
    }
    resetCounters();
  }
```

```

  /**
     * Restore original state of economy
     */
    public void restart(){
    for (Agent a: agents) {
      a.restart();
    }
    resetCounters();
  }

    protected void resetCounters(){
        trades = 0;
        period = 0;
        round = 0;
    }

  /**
     * Return the total utility of all Agents in the Economy.
     */
  public double totalUtility(){
    double total = 0;
    for (int i = 0; i < this.agents.size(); i++) {
      total += agents.get(i).currentUtility;
    }
    return total;
  }
  /**
    * Attempt one exchange per member of the economy
  */
  public abstract void round();

  /**
   * Carry out multiple rounds of trading
   * @param n Number of rounds to run
   */
  public void runRounds(int n){
    for (int i = 0; i < n; i++) {
      round();
    }
  }
//
//     public void outputTotalUtility(FileWriter writer){
//          try{
//               writer.write("Total Utility: " + totalUtility());
//          }
//          catch(IOException e){
//               e.printStackTrace();
//          }
//     }
```

```

  /**
   * @param periods The number of periods
   * @param repetitions The number of realisations to average over
   * @throws IOException
   * */
  public double[] averageUtility(int periods, int repetitions){
    double results[] = new double[periods];
    Processing.initialiseArrayToZero(results);

    for (int r = 0; r < repetitions; r++) {
      for (int i = 0; i < periods; i++) {
        round();
        results[i]+=totalUtility();
      }
      restart();
    }
    for (int i = 0; i < results.length; i++) {
      results[i] /= repetitions;
    }
    return results;
  }

  /**
   * @param periods The number of periods
   * @param repetitions The number of realisations to average over
   * @throws IOException
   * */
  public double[][] manyUtility(int rounds, int repetitions){
    double results[][] = new double[rounds][repetitions];
    Processing.initialiseArrayToZero(results);

    for (int r = 0; r < repetitions; r++) {
      for (int i = 0; i < rounds; i++) {
        round();
        results[i][r] = totalUtility();
      }
      reset();
    }
    return results;
  }

  /**
   * Returns average MRS. If economy has converged sufficiently
   * this is a proxy for prices
   * @return Average MRS values
   */
```

```
  public double[][] getsAverageMRS(){
    double[][] mrs = new double[nGoods][nGoods];
    Processing.initialiseArrayToZero(mrs);

    for (Agent a : agents) {
      Processing.add2dArrayInPlace(mrs, a.getMRS());
    }

    Processing.normalise2dArrayInPlace(mrs,(double)size);
    return mrs;
  }

  public double[] estimateWealth(double[][] goodsList,
      double[][] mrs){
    double[] wealth = new double[size];
    for (int i = 0; i < size; i++) {
      for (int j = 0; j < nGoods; j++) {
        //Add to wealth the amount of good j multiplied by
            mrs with good 1
        wealth[i] += mrs[j][0]*goodsList[i][j];
      }
    }
    return wealth;
  }

    public double[] estimateWealth(double[][] goodsBundles)
        {
        return estimateWealth(goodsBundles, getsAverageMRS
            ());
    }


    public double[] estimateCurrentWealth(){
        return estimateWealth(getAllGoodsBundles(),
            getsAverageMRS());
    }

    public double[] estimateOriginalWealth(){
        return estimateWealth(getAllOriginalGoodsBundles(),
            getsAverageMRS());
    }

  /**
   * Get the current allocation bundle of goods
   * @return current allocation bundle of goods
   */
  public double[][] getAllGoodsBundles(){
    double[][] goods = new double[size][nGoods];
    for (int i = 0; i < size; i++) {
      for (int j = 0; j < nGoods; j++) {
        goods[i][j] = agents.get(i).goods[j];
      }
```

```
    }
    return goods;
  }

  public double[][] getAllOriginalGoodsBundles(){
    double[][] originalGoods = new double[size][nGoods];
    for (int i = 0; i < size; i++) {
      for (int j = 0; j < nGoods; j++) {
        originalGoods[i][j] = agents.get(i).originalGoods[j
            ];
      }
    }
    return originalGoods;
  }
}
```

B.3. **Experimenting Economy Code.** The below code shows how the above economy has been expanded to include the idea of experimentation. We were able to utilise much of the functionality of the CautiousEconomy superclass. The key changes are to the round method and to the counters which are now of type double for efficiency purposes as we would otherwise need to cast integers to doubles to calculate experimentation scaling in each round.

Listing 3. Experimenting Economy source code

```
package ac.decentralisedExchange.cautious.model;

/** The class Economy consists of a collection
 * of independent Agents, who trade
 * via Cautious Trading.*/
public class ExperimentingEconomy extends CautiousEconomy {
  public double acceptable, propensity, decay;
  public double doubleEndDecay, doubleRoundCount;
  public double baselineExperimentation;
  public int endDecay;

  /**An Economy is of size no. of agents each
   * of whom deal with nGoods no. of Goods.
   * @param size Number of agents in economy
   * @param nGoods Number of goods in economy
   * @param acceptable_loss_proportion The proportion of
       average initial
   * absolute utility that is initially
   * acceptable to lose in a trade.
   * This declines until 0 at endDecay. Obviously there are
       schemes which are
   * more analytically satisfying, this one is a compromise
       between this and ease of computation.
   * @param propensity_to_experiment How often to
       experiment
```

```
   * @param endDecay The point at which experimentation stops
   * */
  public ExperimentingEconomy(int size,int nGoods, double acceptable_loss_proportion, double propensity_to_experiment, int endDecay)
        throws IllegalArgumentException{

        super(size,nGoods);
        //Check values of parameters
        if ( 0.0 > acceptable_loss_proportion || acceptable_loss_proportion > 1.0
    || 0.0 > propensity_to_experiment || propensity_to_experiment > 1.0
    || 0 > endDecay){
      throw new IllegalArgumentException("Values must be in range (0,1] for Acceptable, Propensity and positive integer for endDecay");
        }

        this.propensity = propensity_to_experiment;
        this.endDecay = endDecay;
        this.doubleEndDecay = (float) endDecay;

        this.acceptable = acceptable_loss_proportion;

        this.baselineExperimentation = calculateBaselineExperimentation(acceptable_loss_proportion);
  }

  /**
   * No experimentation version of Economy, should perform as Cautious Economy
   * @param size
   * @param nGoods
   * @param acceptable
   * @param proportion
   * @throws IllegalArgumentException
   */
  public ExperimentingEconomy(int size,int nGoods) throws IllegalArgumentException{
    this(size, nGoods, 0.0,1.0, 0);
  }

  protected double calculateBaselineExperimentation(double acceptable_loss_proportion){
    return acceptable_loss_proportion * calculateAverageAbsoluteUtility();
  }

  protected double calculateAverageAbsoluteUtility(){
```

```
    double total = 0.0;
    for (int i = 0; i < size; i++) {
      total += Math.abs(agents.get(i).currentUtility);
    }
    return total /(double)size;
  }

    @Override
  public void round(){
    double allowable_experimentation = this.
        baselineExperimentation *
                         (1.0 - doubleRoundCount/
                              doubleEndDecay);
        int r;

        for (int i = 0; i < size; i++) {
          r = gen.nextInt(size);

          //get another agent at random
          while(r == i){
                r = gen.nextInt(size);
            }

            if(this.round < endDecay && gen.nextDouble() <
                propensity * (1.0 - doubleRoundCount/
                doubleEndDecay)){
            if(agents.get(i).propose(agents.get(r),
                allowable_experimentation)){
              trades++;
            }
            }
          else{
            if(agents.get(i).propose(agents.get(r))){
                trades++;
              }
          }
          period++;
        }
        round++;
        doubleRoundCount++;
    }

  @Override
  protected void resetCounters(){
    super.resetCounters();
        doubleRoundCount = 0.0f;
    doubleEndDecay = 0.0f;
  }
}
```

The code for networked economies can be found in the source files; but it is omitted here as it functions more or less as the above code.