Learning Algorithms in a Decentralized General Equilibrium Model☆

James Bruce*

Department of Economics, McMaster University, 1280 Main St. West, Hamilton, ON, Canada L8S 4M4

Abstract

A model is developed in which economic agents learn to make price-setting, price-response, and resource allocation decisions in decentralized markets where all information and interaction is local. Computer simulation shows that it is possible for agents to act almost as if they had the additional information necessary to define and solve a standard optimization problem. Their behaviour gives rise endogenously to phenomena resembling Adam Smith's invisible hand. The results also indicate that agents must engage in some form of price comparison for decentralized markets to clear--otherwise there is no incentive for firms to respond to excess supply by lowering prices. This suggests that agent-based models with decentralized interaction risk untenable results if price-response decisions are made without being first directed toward the most favourable local price.

JEL Classification: D83; C6; C7; B52; D5

Keywords: Complex adaptive systems; Agent-based computational economics; Adaptive search; Invisible hand

---

☆ The author would like to thank Ted Bergstrom, John Burbidge, Peter Howitt, Leigh Tesfatsion, and Michael Veall for the insights, comments, and/or encouragement.

* Tel.: + 1-905-525-9140 ext. 23211; fax +1-905-521-8232.

Email address: jbruce@sympatico.ca

1. Introduction

This paper presents an agent-based computational economic (ACE) general equilibrium model in which boundedly rational agents learn to produce and trade in such a way that the outcome resembles a competitive equilibrium guided by the invisible hand of Adam Smith. ACE models thus far have not overcome the difficulties imposed by having agents learning to propose and respond to prices while having the subsequent results guide agents' decisions about resource allocation. Such difficulties have either limited the scope of ACE models to partial equilibrium settings or required deviation from the methodology of having phenomena emerge from agents' behaviour.

ACE models employ autonomous interacting agents to simulate an economy or some aspect of one and as such treat the subject of study as a complex adaptive system (Tesfatsion 2001a). Any system-wide regularities emerge from the actions and decisions explicitly modelled at the agents' level rather than being imposed by the modeller. An example of this is Howitt and Clower's (2000) model in which through a process of search and trade agents create a form of money employing a good with low transaction costs. The assumption that agents have the necessary information and ability to define and solve an optimization problem is relaxed in an ACE model and agents typically employ learning schemes to guide their behaviour. Often the models are such that the agents' optimal behaviour can be determined and used as a benchmark from which to evaluate the results. For example Arifovic (1996) has agents in an overlapping generations framework use a genetic algorithm, formalized by Holland (1977), as a learning scheme to determine their consumption/savings decision and to allocate their savings between two different fiat currencies which allows the agents to (nearly) settle on the optimal savings/consumption decision.[1] The growing body of ACE literature has addressed issues in finance (LeBaron 2000), labour markets (Tesfatsion 2001b), R&D investment (Yildizoglu 2002), foreign exchange markets (Arifovic 1996,2001a, and 2001b), signalling (Arifovic and Eaton 1998 and DeVany and Lee 2001), auction design (Bower and Bunn 2001), development (Arifovic et al 1997), the demand for money (Howitt and Clower 2000), and industrial organization (Price 1997).

Despite the approach of explicitly modelling from the agent up and avoiding the imposition of system-wide controlling artifices with no real-world analogue, ACE models have not yet lived up to that potential with regard to the Walrasian Auctioneer. With the exception of Kirman and Vriend (2001) no ACE model has been developed in which agents learn to propose and respond to prices. They develop a model where wholesalers who learn strategies governing stocking, queue-ordering and pricing interact with retailers who learn strategies governing queue-choosing and price-response. The focus of their paper is not on pricing behaviour but on loyalty.[2] Their model has only one good so it is not clear if the price-related decisions are appropriate in a general equilibrium setting.

---

[1] The central result of the model is that the exchange rate fluctuates, whereas the rational expectations equilibrium of the equivalent model (Karaken and Wallace 1981) has a stable but indeterminate exchange rate.

[2] Although the agents' decision-making process does not allow them to formulate strategies in such terms, it appears that the retailers act as if they pay a price premium and exhibit loyalty to their chosen wholesaler in order to secure a guaranteed supply.

The majority of ACE models have agents learning to make non-pricing decisions and a Walrasian auctioneer subsequently determining the market clearing price.[3] Many other ACE models focus on issues in such a way that prices do not play a role and hence there is no call for price decisions nor a Walrasian auctioneer.[4] Other approaches taken involve having different agents propose different but immutable prices (Rouchier et al 2001), having pricing decisions follow an (unlearned) rule of thumb (Howitt and Clower 2000), bypassing a local bargaining process by having agents trade at a price determined as some defensible function of agent attributes (Epstein and Axtell 1996[5] and Dawid 1999), and having agents learn to set prices but respond to prices by optimally choosing a quantity (Dawid 2000).

The lack of an emergent replacement for, or equivalent of, the Walrasian auctioneer has not gone unnoticed in the literature. Taking the 'invisible hand' of Adam Smith to be the process by which optimal allocation of resources across sectors of the economy emerges from the self-interested actions of economic agents, Kochugovindan and Vriend (1998) argue that existing formal economic models treat it as a black box by relying on fictitious constructs, including the Walrasian auctioneer. They speculate that

---

[3] Examples of this include Arifovic (1995, 1996, 1998, 20001a, and 2001b), Arifovic and Ramazan (2000), Dawid and Kopel (1998), Duffy (2001), Price (1997), Tay and Linn (2001), Vriend (2000), and Yildizoglu (2002) as well as those modelling financial markets.

[4] Examples of this include Andersson and Sandholm (2001), Arifovic and Eaton (1998), DeVany and Lee (2001), and Tesfatsion (2001b).

[5] In chapter IV, Epstein and Axtell present a decentralized model with two goods that are traded at local prices. While they find that the mean price is that expected if agents fully optimized, the agents in their model do not learn but instead follow immutable rules of thumb.

the study of complex adaptive systems may yield insight into how a decentralized economy gives rise to this emergent property. Leijonhufvud (1999) makes a similar argument. Arifovic (2000 p241) surveys models with evolutionary learning algorithms. In her conclusion she discusses the state of affairs in the literature in so far as price determination is concerned and the hurdle it presents for more comprehensive ACE models:

> One of the challenges that the research in this area faces is the extension of evolutionary models to the general equilibrium type of economies with multiple markets. The main issue is the one of determination of prices. These models cannot take advantage of computing prices through simultaneous determination of agents' optimal decisions and market-clearing conditions. Instead, the calculation of prices has to be explicitly modeled by describing a bargaining process or some other equilibrating mechanism. This adds an extra layer of complexity on top of the dynamics that tend to be quite complicated anyway. However, this obstacle will have to be overcome if these models are to be more widely used in the general equilibrium setting.

Such a successful general equilibrium ACE model would, inherently, elucidate the invisible hand as an emergent property.

The purpose of this paper is to present and analyse such a model. In this model, agents must decide to produce one of two goods and then have the opportunity to engage in trade for the other good via a sequence of local interactions with other agents. Agents learn which good to produce as well as pricing and purchasing strategies.

The resultant dynamics mirror how the invisible hand is normally described as functioning: A shortage (surplus) of one good results in its price rising (falling) and agents respond by shifting resources towards (away from) the production of that good. This allocation however cycles around the optimal one. There is a lag between the introduction of a shortage (surplus) and agents' ability to perceive this and adjust prices in response. Thus when the shortage (surplus) is corrected, prices continue to be above (below) market-clearing levels for that optimal allocation, and resources continue to move, creating a surplus (shortage) of the good in question.

Additionally, agents must engage in price comparison—visiting a number of agents to collect price information and patronizing first those agents offering the most favourable prices—for the outcome to resemble a competitive equilibrium. This provides agents with an incentive to respond to a surplus of the good they produce by lowering prices.

The next section presents the model. Section three presents and analyses the results. Section four presents and analyses results of variations of the base model. Section five concludes.

2. The Base Model

There are n agents who live for T periods and have identical preferences over two goods. Table 1 gives the values used for the parameters. Every period each agent decides which of two goods to produce. Each agent then has a series of opportunities to engage in bilateral trade with a number of other agents. Any particular agent will, in some trades, act as a store by offering a specific exchange. In other trades, that same agent will act as a customer by responding to an offered exchange. Subsequent to the trades of a period, each agent revises its strategies by imitating the strategies of the most successful agents encountered during trade. Some of the agents then mutate some of their strategies through a stochastic process. A period ends with consumption.

Each period thus consists of five stages:

i) Production

ii) Trading

iii) Imitation

iv) Mutation

v) Consumption

which are explained in detail below.

2.1 The Production Stage

An agent is assumed to be able to produce only one of the two goods in any given period. Production behaviour is dictated for agent a by its production strategy $wp_a \in \{1,2\}$. If $wp_a = 1$, agent a produces e_1 units of good one; while if $wp_a = 2$, agent a produces e_2 units of good two[6].

2.2 The Trading Stage

A given agent a's trading behaviour is dictated by its exchange strategies $\{\{o_{1a}, r_{2a}, k_{1a}\}, \{o_{2a}, r_{1a}, k_{2a}\}\}$. These exchange strategies and the production strategy wp_a make up agent a's entire complement of strategies.

---

[6] Regarding notation: *a, b,* and *c* are used to denote agents; *i* and *j* are used to denote goods; *s* and *t* are used to denote time periods. Variables representing sets of agents and strategies have implicit subscripts denoting the period which are usually suppressed for readability where there is no loss of clarity. Occasionally, when it facilitates exposition, the agent subscripts are also suppressed.

At the beginning of every period, each agent a selects v different agents to visit that form a set Γ_a of agents. Γ_a is determined through an independent random process; in particular, it is independent of agent a's strategies, any Γ_a formed in previous periods, and any Γ_b where $a \neq b$. Every possible set of v of the $n-1$ other agents is equally likely to make up Γ_a. A pair of agents is assumed to not trade with each other if they both produced the same good. Thus a set $\gamma_a \subseteq \Gamma_a$ is formed by dropping from Γ_a agents producing the same good as agent a; γ_a consists of the agents that agent a will visit and potentially trade with, where agent $b \in \gamma_a$ acts as a store and agent a acts as a customer.

Agent b, acting as a store, is assumed to offer its good in a package of a size fixed for the period at o_{jb} where $j = wp_b$. Agent b additionally offers a price, also fixed for the duration of the period. The price agent a faces when visiting agent b is given by $\frac{r_{ib}}{o_{jb}}$; agent b offers o_{jb} units of the good it produced (good j) and in exchange requests r_{ib} units of the other good (good i, which agent a produced).

A proposed exchange is carried out if both agents involved have sufficient holdings of the good they produced available for trade. Each agent a keeps k_{ia} (where $i = wp_a$) of the good it produced for personal consumption, putting up the remaining $e_i - k_{ia}$ for trade. (Goods put up for trade but not actually traded are still available to the producing agent later for consumption.) Thus if agent a (the customer) visits agent b (the store), agent a gives to agent b r_{ib} units of good $i = wp_a$ and receives o_{jb} units of good $j = wp_b \neq i$ if and only if $x_{ia} - r_{ib} \geq k_{ia}$ and $x_{jb} - o_{jb} \geq k_{jb}$ where x_{ia} denotes agent a's current holdings of good i. Note that if agent a produces good one in some period, only $\{o_{1a}, r_{2a}, k_{1a}\}$ are used to determine its trade behaviour; while if $wp_a = 2$, $\{o_{2a}, r_{1a}, k_{2a}\}$ are the only active exchange strategies.

An agent a is assumed to go to the agents-as-stores in γ_a in an order determined by their offered prices, visiting first those agents-as-stores that offer more favourable prices (from agent a's point of view). γ_a is thus ordered into $\vec{\gamma}_a$ where $b \in \vec{\gamma}_a$ if and only if $b \in \gamma_a$, while additionally, if b is element q in $\vec{\gamma}_a$ and c is element $q+1$ in $\vec{\gamma}_a$ then $\frac{r_{ib}}{o_{jb}} \leq \frac{r_{ic}}{o_{jc}}$ where $wp_b = wp_c = j \neq wp_a$. Should agents $b, c \in \vec{\gamma}_a$ have strategies such that the prices they offer are equal, b is as likely to come before c in $\vec{\gamma}_a$ as the other way around and this is determined by a random process independent of any aspect of the state of the model.

Once the various vectors $\vec{\gamma}_a$ for $a = 1, 2, \ldots, n$ are formed, the whole of the n agents are randomly ordered into $\vec{N}_t$ (with t indicating the period). Each possible $\vec{N}_t$ is chosen with equal independent probability; in particular, $\vec{N}_t$ is independent of agents' strategies and the set of any prior such orderings $\{N_s\}_{s=1}^{t-1}$.

The first agent a in $\vec{N}_t$ visits the first agent in its list $\vec{\gamma}_a$. If they engage in a trade it is the end of agent a's 'turn'. If not, agent a moves to the second agent in $\vec{\gamma}_a$ and if it engages in a trade with this agent, it is the end of agent a's turn. If no trade occurs between agent a and the second agent in $\vec{\gamma}_a$, agent a moves on to the third agent in $\vec{\gamma}_a$ and similarly continues to move through $\vec{\gamma}_a$ until either it engages in a trade (ending its turn) or it has reached the end of $\vec{\gamma}_a$ (which also ends its turn). Agent a's turn being over, the second agent b in $\vec{N}_t$ similarly has a turn to go through $\vec{\gamma}_b$. Then the third agent in $\vec{N}_t$ has a turn and so on until the n th and final agent in $\vec{N}_t$ has had its turn. Call this process—the n agents each getting a turn to attempt to trade as a customer with an agent in its $\vec{\gamma}$ —a trading run.

The entire trading stage of a period consists of the formation of $\vec{N}_t$ and $\{\vec{\gamma}_a\}_{a=1}^{n}$ followed by trading runs which occur until no more trades are made (i.e. a trading run happens in which all n agents end their turns by failing to trade with any member of their respective $\vec{\gamma}$ s). Since a typical trading stage will involve hundreds of trading runs, agents who are stuck on the short side of an excess supply or demand outcome find themselves there more likely due to their strategies rather than their placement in the random order of movement, $\vec{N}_t$.

One way to conceive of the trading stage is to suppose that agents can shop and maintain their stores simultaneously, that a store sells discrete packages at a rate (packages per unit time per customer) that is constant across stores, and that the store is able to service an arbitrary number of customers simultaneously. Agents begin the trading period having already sampled some price information and going to their most preferred store. Whenever a store runs out of stock, agents at that store all disperse to their next-most preferred store and the process continues until all mutually agreeable trades between agents and the stores they visit have been exhausted.

2.3 The Imitation Stage

The utility obtained by agent a is given by $U_a = x_{1a}^{\rho} x_{2a}^{1-\rho}$.

Each agent has a set of agents met during the exchange stage including the v agents on its visitation list, the randomly determined number of agents that visited it, as well as itself. Thus if μ_a is the set of agents met by agent a then $b \in \mu_a$ if and only if $b \in \Gamma_a$, $a \in \Gamma_b$, or $a = b$. This set of agents met is divided into two groups based on which good they produced, with the imitating agent itself in both groups: $\mu_{1a} \subseteq \mu_a$ and $\mu_{2a} \subseteq \mu_a$ where $b \in \mu_{ia}$ if and only if $b \in \mu_a$ and $wp_b = i$ or $b = a$.

Agents are assumed to imitate their active exchange strategies from the agent met that obtained higher utility than any other agent met that produced the same good. Agents are assumed to imitate their inactive exchange strategies from the agent met that obtained higher utility than any other agent met that produced that good, but only if that highest-performing agent obtained more utility than the imitating agent. Otherwise it retains the same inactive exchange strategies. So, for $i = 1{,}2$ if during period t, $b \in \mu_{ia}$ and $\forall c \in \mu_{ia}: c \neq b$, $U_b > U_c$ then $\{o_{ia}, r_{ia}, k_{ia}\}$ during period $t+1$ will equal $\{o_{ib}, r_{ib}, k_{ib}\}$ during period t.[7]

With probability JCR, the agent looks to the agent encountered who obtained the highest utility (regardless of the good produced) and imitates wp from that agent. So if

[7] If, for some set of agents met, (μ_{1a} and μ_{2a} for exchange strategies or μ_a for the production strategy) there are more than one agent with the highest utility level, the agent to be imitated is determined as follows: If the imitating agent *a* is one of those with the highest utility it retains the strategies to be imitated i.e. imitates itself. Otherwise, if there are members of Γ_a with the highest level of utility, those members of Γ_a are each equally likely to be chosen for imitation by an independent random process. Otherwise all agents with the highest level of utility (none of who is agent *a* nor a member of Γ_a) are each equally likely to be chosen for imitation by an independent random process.

during period t, $b \in \mu_a$ and $\forall c \in \mu_{ia} : c \neq b$, $U_b > U_c$ then (with probability JCR) wp_a during period $t+1$ will equal wp_b during period t.

It is during the imitation stage that agents make use of more information than at any other stage in the model. In order for the imitation stage to be carried out, agents must know at least what each agent met obtained in goods at the end of the trading stage as well as the active exchange strategies of up to two other agents (perhaps best viewed as selected by the agent in question). To facilitate later discussion, agents' information sets are assumed to be greater than this bare minimum. Agent a's information set at period t, Ω_{at}, is assumed to include the goods obtained, and the strategies used, by all agents encountered during period t. Note that agents do not observe each others' utility functions but instead evaluate other agents' bundles of goods with respect to their own utility functions. Agents are assumed to have perfect recall of information gathered in the past as well as the ability to identify previously encountered agents (even though they make no use of this ability in the base model). Thus:

$$\Omega_{at} = \left\{ \left(X_{1sb}, X_{2sb}, wp_{sb}, ex_{sb} \right) : s \leq t, \left[b \in \Gamma_{sa} \vee a \in \Gamma_{sb} \right] \right\}$$

where X_{ita} represents agent a's holdings of good i at the end of the trading stage of period t while ex_{ta} represents agent a's six-member set of exchange strategies in period t.

2.4 The Mutation Stage

The six exchange strategies are all subject to mutation; the production strategy, wp, is not. View the imitative learning process as a search algorithm that combs through the set of sensible exchange strategies $\left[0, e_1\right]^3 \times \left[0, e_2\right]^3$ as well as the set of production strategies $\{1,2\}$. Mutation of exchange strategies is necessary for a comprehensive search. So long as both goods are being produced, mutation of the production strategy is not, and so is neglected in order to avoid subjecting the results to any more noise than necessary.

Each agent perturbs each of its exchange strategies separately with independent probability MR. If a strategy is to be mutated and its initial value is x, its new value will be independently drawn from a uniform distribution over the interval $[(1-\delta)x,(1+\delta)x]$.

2.5 The Consumption Stage

In the consumption stage, each agent a consumes all of its holdings of the two goods and experiences utility $U_a = x_{1a}^{\rho} x_{2a}^{1-\rho}$. The goods are assumed to be perishable, so by the end of the consumption stage, $x_{ia} = 0 \ \forall i \in \{1,2\}$ and $a \in \{1,2,\ldots,n\}$.

2.6 The Initial State

The strategies in period 1 are determined by random and independent draws for each agent and each strategy. Strategies o_1, r_1, and k_1 are drawn from a uniform distribution over $[0,e_1]$ while o_2, r_2, and k_2 are drawn from $[0,e_2]$. Note that a value for any active strategy that is above this interval would preclude any trade, while a negative value is equally unacceptable. The fact that $o_{ia} \geq e_{ia} - k_{ia}$ would also preclude trade if $wp_a = i$ is not used to further restrict the range that o_{ia} is drawn from. Instead agents must, and do, learn to set $o_{ia} < e_{ia} - k_{ia}$. wp is initialized at 1 or 2 with equal probability.

2.7 The Bargaining Structure and Homo Economicus

The bilateral bargaining encounters in the base model are ones in which one agent b proposes a price-quantity pair (p_b, q_b), while the other agent a's response is determined by some function $f: R^2 \rightarrow \{accept, reject\}$. The base model uses $(p,q)_{BM} = \left(\frac{r_{ib}}{o_{jb}}, o_{jb}\right)$ if $x_{jb} - o_{jb} \geq k_{jb}$ and $(p,q)_{BM} = (\cdot, 0)$ otherwise. The response function used is $f_{BM}\left(\frac{r_{ib}}{o_{jb}}, o_{jb}\right) = accept$ if $x_{ia} - r_{ib} \geq k_{ia}$ and $f_{BM}\left(\frac{r_{ib}}{o_{jb}}, o_{jb}\right) = reject$ otherwise. The dictum that agents should do as well as they can under the circumstances ostensibly suggests that agents should instead use $f^*(p_b, q_b) = accept$ if acceptance of the proposed exchange would cause agent a's utility to rise and $f^*(p_b, q_b) = reject$ if it would cause utility to fall.

However, agent a has sufficient information to make its problem more complicated than simply choosing from $\{accept, reject\}$ to maximize U_a at the end of the exchange. The perfectly rational agent should be making its response based on $\{p_b, p_a, \Omega_t\}$ where p_b is the price it faces now as a customer and p_a is the price it offers as a store in any future bilateral bargaining situations to come during the current period t. In particular, the response decision is complicated by the fact that when p_a is a more favourable price for agent a than p_b (which the results show is usually the case), the purchases that agent a now makes limit the uncertain quantity of future trading at a more favourable price. (Note that if agents know the initial conditions and other agents' preferences, the expected utility of each response in $\{accept, reject\}$ is computable and so agent a's problem is well-defined if exceedingly difficult.[8])

As an optimization problem, the agents' situation is better viewed as imperfectly choosing a quantity to purchase based on a distribution of prices to be faced. More accurately, agent a's problem involves choosing a quantity to attempt to spend ($e_i - k_{iat}$) and a price-quantity pair, for its store behaviour, based on $\Omega_{a,t-1}$. Last period's information set is enough to compute the distribution of other agents' strategies this period and define the expectation of U_{at} as a function of strategies. The agents do not perform any optimization, so if they manage to settle on a nearly optimal outcome, this is due entirely to the learning algorithm. This would open up the possibility that the learning algorithm can be used in situations where the agents' problem is more complicated.

3. Results of the Base Model

The purpose of this model is to subject bargaining decisions (price setting and price response) to a learning algorithm in an information-poor general equilibrium environment to gauge their ability to generate behaviour leading to an optimal (in this case synonymous with competitive) outcome. Since the economic process in the model is subject to noise, assessing the degree of optimality of the outcome is not straightforward.

The most obvious criterion for evaluation would be implemented by assessing the observed welfare outcomes and comparing them to what would be realized by a perfect information general equilibrium outcome. However any particular outcome (in terms of

[8] Assuming that all other agents follow the simple imitation and mutation rules.

resource allocation and consumption patterns) can be characterized as arbitrarily close or far from the optimal outcome using utility functions that are monotonic transformations of each other. Utility can be appropriately used to rank the ability of different variations of the model to achieve the optimal outcome. In an attempt to impose some rigour on this, only utility functions which were homogeneous of degree one were used. When reported, utility will be as a percentage of the optimal outcome's level.

Another criterion for consideration is the degree to which the aggregate price level comes close to the competitive equilibrium level. But since the model will produce price dispersion, it is problematic to declare that some given level of price variance is compatible with being near an optimal outcome, even if the mean price is at that level.[9]

The third criterion used is resource allocation. When the resource allocation is near to optimal, the imitative learning algorithm is deemed to work in the sense that it nearly gives rise to the competitive allocation that full information and rational behaviour would produce.

While none of these criteria is entirely satisfactory since they do not indicate how near to optimal the results need to be to declare the outcome nearly optimal, in practice it turns out that for the various versions of the model analysed, it is almost always straightforward to tell when the decision-making algorithm has produced (or failed to produce) a nearly optimal resource allocation.

The model is run for 1000 periods. 10 such runs were made and the results were qualitatively similar. During the first 500 periods rho (the preference parameter) is set at

[9] In section four, one variation of the base model produces a mean price at the competitive level but arbitrarily large price dispersion.

0.6. For the last 500 periods it is 0.3. With the parameter values listed in table 1, the competitive equilibrium would involve 1000ρ agents producing good 1 (with the other $1000(1-\rho)$ agents producing good 2), and trades occurring at a price of 1 so that each agent consumes ρ units of good 1 and $1-\rho$ units of good 2 (resulting in a mean utility of 100%).

This base model is also run without agents engaging in price comparison i.e. where an agent a does not, at the beginning of the exchange stage, order the agents it will visit into $\vec{\gamma}_a$ according to the prices offered by those agents in $\vec{\gamma}_a$; the order in which those agents will be visited is, instead, random.

Figure 1 shows the number of agents producing good 1 for the base model with and without agents engaging in price comparison. The straight lines, at 600 agents for the first 500 periods and 300 agents for the last 500 periods, show the competitive equilibrium resource allocation. Figure 2 shows the mean price of good 1 at which trades occurred for the two versions of the base model (again with a line at 1 showing the competitive result) and figure 3 shows the mean utility agents realized each period. For all three graphs, only every fifth observation is shown[10] since otherwise they are indecipherable. Furthermore, prices of good one greater than two are shown as two since values could be as high as sixteen. Prices above two were rarely observed in the runs with price comparison. When observed they are attributable to the initial conditions; the latest they were observed was the 31st period.

[10] This is also true of most of the subsequent graphs.

With agents engaging in price comparison, the price hovers near the competitive equilibrium price. This allows for resource allocation approximating that of a competitive equilibrium, and a correspondingly higher mean utility.

Note that price and resource allocation cycle around the competitive level. When less than the competitive level of good one is being produced, the mean price rises above the competitive level, inducing more agents to produce it. This process occurs until competitive resource allocation levels are reached, at which point the price has not yet fallen to reflect this. Agents continue to switch to producing good one and it is then over-produced. Eventually the price begins to reflect this and agents start switching to production of good two. When competitive resource allocation is again achieved, the price has yet to reflect this, and so the economy goes back to overproducing good one.

This process's dynamics have two driving forces to be explored:

1) When the price of a good is below the competitive level, agents switch away from producing it.

2) When the economy produces more than the competitive level of a good, that good's mean price falls below the competitive level.

Each of these is examined in turn:

Ignoring the impact of the decision of how much of one's product to reserve for personal consumption (or assuming it is optimal), agents are equally well off regardless of which good they produce if all trades occur at the competitive equilibrium price. Thus if the average price of good one is above the competitive level, the average agent producing good one is better off than the average agent producing good two, and so an agent considering switching its production good is more likely to imitate an agent producing good one. The reverse, of course, is true if the average price of good two is above the competitive level. This reasoning suggests that if, in some period, the price of good one is above the competitive level we should see production of good one rise and vice versa. This turns out to be the case in over 83% of periods in ten runs for which it was examined. Furthermore, almost all of the periods during which production does not move in the direction suggested by mean price are clustered around the peaks and valleys of the resource allocation cycle when the price is generally closest to competitive levels.

The other dynamic force moves the price of good one above competitive levels when it is under-produced and below when it is overproduced. A more restrictive sufficient phenomena would be if the price tends to move towards the market clearing price (i.e. the competitive price given the existing resource allocation which itself is not always at the competitive level). This seems to be roughly the case. Figure 4 shows the market clearing price and the observed mean price of good one for a typical run. Figure 5 shows the market clearing price and the observed mean price (every tenth observation shown) when the model is altered so that agents do not change the good they produce: the same 600 agents are all producing good one. The preference parameter ρ moves back and forth between .6 and .3 and the resultant market clearing prices are 1 and 2/7 respectively. While it is clear that agents trade at prices near to the market clearing level, it is not obvious why they do so given that their behaviour is a simple imitative algorithm and they lack information about the current resource allocation needed to determine what that price would be.

Consider that a particular pricing strategy propagates through agents more readily if it delivers for its users more utility. If we imagine a much more clever agent faced with the same problem these agents face for selecting a price, that more clever agent has both benefits and costs to a marginal increase in the price it posts. The obvious benefit is that the agent gets a better deal from the customers (visiting agents) it does deal with.[11] The cost of raising the price offered is that it increases the likelihood that the agent will find itself too low on its visitors' ordered lists, too few agents visit to trade, and the agent finds itself at the end of the period having engaged in an insufficient number of trades. Note that without the price comparison process, the cost of raising the price disappears and so the imitative algorithm selects higher prices. The implication of this is that the costs and benefits of raising the price must balance at or near the market clearing price.

To test this four different cases were run in which all but ten agents had the same fixed pricing strategy, while the remaining ten agents, all of whom produced good one, had their pricing strategy fixed at some other level. Furthermore, the production strategies of all agents was fixed to maintain a stable market clearing price. The first case had $n - 10$ agents employing strategies giving a market clearing price (They offered .1 of their production good in exchange for .1 of the other good.) while the other ten agents offered exchanges at a higher price (offering .099 of their production good for .1 of the other good). The second case also had $n - 10$ agents offer a market clearing price with the other ten agents offering a lower price (.101 of their good for .1 of the other good). The third case had 10 agents offering a market clearing price while the other $n - 10$

[11] An additional potential benefit is that by posting a high price, visiting agents are less likely to have the agent near the top of their sorted lists, in which case the agent has more opportunity to trade at the best price it finds among the agents it visited without a danger that it runs out of its stock of marketed product by selling it to visitors. As it turns out, this is not a consideration at actualized strategies since agents generally propose prices more favourable to themselves than the ones they encounter.

agents offered the higher price. The fourth case had 10 agents offering a market clearing price while the others offered a lower price. Table 2 summarizes the results; the figure reports the percentage of periods in which the group of agents employing the competitive pricing strategy have a higher mean utility than the other group of agents.

The higher these figures, the more it suggests the costs and benefits of price movement balance at the market clearing price and so the more stable it is as a phenomena resulting from agent behaviour. For example, suppose initially all agents offer exchanges at the competitive price and a small number of them, due to strategy mutation, begin to offer a higher price. This describes the first case. The agents maintaining the competitive pricing strategy are on average usually doing better. Thus we expect both that these agents will not adopt the higher price and that the ones that did will abandon their new high price by imitating the majority of agents offering competitive prices--which suggests stability around the competitive price. For three of the four cases we have results suggesting movement towards the competitive price. This is consistent with general stability at the competitive price since in the case where a minority of agents offering a lower than competitive price do better, the resulting dynamics would lead to something more resembling the case where the majority offer a low price and do less well than the minority still offering a competitive price.

To see why it is that these costs and benefits balance near the market clearing price we need to consider the role of k_i—the decision of how much of its product an agent keeps for personal consumption and how much it makes available for trade. If all of the agents are initially employing strategies leading to a competitive outcome and then those producing good two were to switch to a higher k_2, the agents producing good one will respond by offering a lower price on good one—they now face a positive probability of being unable to make enough trades and so those that avoid that by offering a slightly lower price can on average outperform those that do not. Another way of viewing the effect of k_i is to note that by increasing k_2, the marginal cost of a price increase for producers of good one is increased because the probability of engaging in too few trades has increased, so it is by way of k_2 that producers of good two prevent producers of good one from charging an arbitrarily high price. Table 3 shows the results of a version of the base model designed to test this—k_2 for all agents is fixed at 3 levels: one run each with k_2 below, at, and above the competitive level. (.3, .4, and .5 respectively) The production strategy *wp* was also held fixed at the competitive level, otherwise it soaks up the effect of k_2 on price by shifting resource allocation until the price returns to the competitive level. The mean observed price of transactions and the mean price of the strategies employed by producers of good one are shown as well as the 5th and 95th percentiles. As expected, the higher k_2, the lower the price of good one both in terms of strategies employed by producers of good one and realized trades.

In the base model agents do not collude to restrict supply; they choose their strategies imitatively. So despite the advantage of affecting price by jointly restricting supply, agents choose individually beneficial levels of k producing an approximately competitive price.

Given some distribution of k_i strategies, if every agent posts the market clearing price, the marginal cost of raising one's price is almost zero: the probability of being unable to make enough trades due to a higher price is close to zero since the market is nearly clearing. Thus if all agents learn to set a price where the marginal costs and benefits of a price change are equal, barring a large discontinuity in the costs when the probability of unrealized trades becomes positive, agents will set a price that is more favourable to themselves than the market clearing price. The results of the base model support this reasoning. Figure 6 shows the mean price of good one posted by producers of good one, the mean price of good one posted by producers of good two, and the market clearing price. The mean price posted by producers of good one is usually greater than the market clearing price (true in 80.1% of periods in the ten runs) which in turn is usually greater than the mean price of good one posted by producers of good two (true in 85.2% of periods in the ten runs).

The overall dynamics of the base model are similar to the story of Adam Smith's invisible hand. Agents post prices near to market clearing to avoid a large probability of being caught absorbing excess supply (which works because agents engage in price comparison). On average the observed price is close to the market clearing level because agents with different production strategies are posting prices on opposite sides of market clearing and because price comparison will likely mean agents with prices closer to market clearing realize more trades. Meanwhile agents shift into the more advantageous production strategy, which reduces that advantage. Lags between actual and perceived advantages (due to non-instantaneous price adjustment in the information-poor environment) cause agents to keep shifting into an industry even after it has become disadvantageous to do so, leading to a cyclical outcome around the competitive one.

4. Variants of the Base Model

This section discusses the robustness of the base model. For an agent-based general equilibrium model to produce an outcome approximately that of a competitive equilibrium while also portraying economic agents as pursuers of self-interest, it needs three features:

1) The set of possible actions the n agents are able to undertake (arising from the set of possible strategies fed through the economic environment) must include nearly optimal behaviour at the individual level and something near a competitive outcome at the economy level. i.e. It must be possible for agents to choose optimal actions and it must be possible for agents to choose a competitive outcome.

2) If each agent chooses its optimal strategy at time t given its information set, the outcome needs to be approximately a competitive outcome.

3) The learning algorithm agents use finds something close to the optimal strategy.

A number of variants of the base model are considered and analysed in terms of their impact on the above three features. When describing the variations, only the differences from the base model are outlined. Table 4 presents summary statistics of these variations. As in the base model, ten runs of one thousand periods were generated for each variant.

4.1 Altering the strategy-action relationship

The inability of the base model without price comparison to result in the competitive outcome can be attributed to agents not trading at the most favourable available price first, since the strategies could deliver such action only by coincidence. Here I present two variants on the base model designed to build a favourable-price-first action into the strategies subject to learning.

Variant 1.1—The formation of $\{\bar{\gamma}_a\}_{a=1}^{n}$ from $\{\gamma_a\}_{a=1}^{n}$ is governed by strategies subject to the learning algorithm. Each agent a has an additional exchange strategy D_a

consisting of the integers one through ten in some order. If the first integer in D_a is q, the first agent visited for trade (i.e. the first agent in $\bar{\gamma}_a$) is the one met producing the appropriate good (i.e. from γ_a) with the q th lowest price. If the second integer is r, the second agent visited is the one with the r th lowest price, and so on. This strategy is imitated from the most successful agent encountered. It is mutated with independent probability MR for every agent. If this strategy is mutated for an agent, then two entries are chosen at random and swapped.

Agents in variant 1.1 were able to learn to find a competitive outcome but not as effectively as in the base model. Figure 7 shows resource allocation and figure 8 the mean price.[12] Agents obtained a lower mean utility, 79.8 over ten runs, as opposed to 94.2 in the base model, although this is an improvement over the base model without price comparison which had mean utility of 45.2. Figure 9 shows the path of the mean of the first three entries in this ordering strategy for a typical run. Optimally they would have values of one, two, and three respectively. The first entry tended to be low—it had a mean value of 3.48 over the ten runs and exhibited a pattern of staying just above one with periodic spells at higher values. On average 48.1% of the agents had it equal to one in any period over the ten runs. The other nine entries had mean values greater than five and did not appear to differ much from each other overall. Thus the price comparison that did occur consisted of most agents visiting the best price first and the remaining

[12] Some of the values for the mean price are as high as 24, so to preserve the graph's legibility, values greater than 2 are shown as 2. Such values are attributable to the initial conditions; no values greater than 2 were observed after the 60th period in any of the ten runs.

visits being carried out in random order. This was, however, significantly more successful than no price comparison.

Variant 1.2—Agents do not engage in price comparison but do have additional strategies rp_1 and rp_2: reserve prices. The conditions for a proposed trade to be carried out become: i) Each agent will remain holding at least as much as its relevant k strategy dictates and ii) The implicit price of the trade does not exceed the visiting agent's relevant reserve price.

This new strategy and trade condition could apparently function in a way similar to price comparison—it potentially punishes agents who post high prices—but it could not however deliver on this. Over the ten runs the mean level of utility was 47.8, only slightly better than the base model without price comparison. There was no typical resource allocation result; the chief similarity between runs was that most cycled around some level that was not the competitive allocation, but that level varied from run to run. Figure 10 shows resource allocation for three of the runs.

The above two variants were designed to endogenize price comparison as a learned behaviour and met with mixed success. The price-ordering strategy is rather contrived. It conditions learning on information of which the proper use should be obvious. One possible extension would be to have agents learn to revisit agents with favourable prices based on previous encounters, a set-up that has been employed by Howitt and Clower (2000) and Kirman and Vriend (2001).

4.2 Alterations in Agents' Optimal Strategy

One variation was considered that significantly altered the agents' optimal strategy.

Variant 2.1—Agents' exchange proposal strategies o_1, o_2, r_1, and r_2, are replaced by pricing strategies p_1 and p_2. Agents post a price of p_{wp} for the good they produce. Visiting agents respond by purchasing up to the optimal quantity, limited by the visited agent's relevant k strategy.

Note that if two agents producing different goods are involved together in their first interactions of a trading stage, the visiting agent will spend the same amount of its production good regardless of the price: if the visiting agent produces good 1, it spends $(1-\rho)e_1$, while if it produces good 2, it spends ρe_2. Thus, in such an encounter, the visited agent's subsequent utility rises with its posted price. Price comparison was not able to prevent offered prices from climbing geometrically—presumably as fast as mutation and propagation via imitation allowed. Figure 11 shows the natural log of the mean level of strategies p_1 and p_2. Figure 12 shows the resource allocation and figure 13 shows the mean price.

Given the growth of posted prices, the outcome for any individual agent is dependent upon whether it is visited (and receives goods nearly for free) before visiting or vice versa (in which case it gives goods nearly for free). Hence the observed price will reflect the resource allocation by being above one if more agents produce good one and hence are more likely to be visited first, otherwise below. This suggests that if ρ were 0.5, the mean price would be 1 and resource allocation would also be near the competitive level. This was borne out in the results. With ρ set equal to 0.5 throughout the 1000 periods of a run, the mean price was 0.995 over 10 runs and resource allocation stayed near 500 (with a mean value of 501) agents producing each of the goods. Figures 14 and 15 show the results of typical run. Despite nearly competitive price and resource allocation, mean utility was only 71.5 over ten runs and every individual trade had a price growing arbitrarily far from the competitive level. This suggests that summary statistics in agent-based models have the potential to be misleading.

One other variant was considered that had the potential to alter the agents' optimal strategy.

Variant 2.2—An additional criterion was added to the base model for a proposed trade to be realized. If utility went down for either agent, the trade was rejected. Qualitatively the outcome was the same as the base model and so the results are not presented. Mean utility was slightly lower at 92.6 vs. 94.2 in the base model.

4.3 Variations in the Learning Algorithm

Three minor variations in the learning algorithm were analyzed.

Variant 3.1—Agents imitate all exchange strategies from the most successful agent met.

Variant 3.2—Agents do not imitate their inactive exchange strategies. The active exchange strategies are imitated from the most successful agent met who produced the same good.

Variant 3.3—Agents imitate probabilistically with more successful agents encountered more likely to be imitated than less successful agents. Agents encountered who produced good one are ranked from most successful to least successful. If agent a is ranked q and agent b is ranked $q+1$, then agent a is twice as likely to be imitated (for its active exchange strategies) as agent b. Probabilities are scaled so that some agent is imitated. Similarly, agents who produced good two are ranked and one is chosen from whom its active exchange strategies are imitated. All agents met are ranked and with probability JCR an agent is chosen for imitation of the production strategy.

None of the above three variants produced results qualitatively different from the base model, so the results are not presented.

The final variants on the learning algorithm employed genetic algorithms (GAs) for learning. Formalized by Holland (1977), the GA is the most common learning algorithm in the ACE literature. Each agent has its strategy represented as a binary string $6\Lambda + 1$ bits long written over the alphabet $\{0,1\}$. Each of the six exchange strategies is represented by an Λ bit section of the string. If S is the substring for one of those strategies and $S_\lambda \in \{0,1\}$ is the λ th entry in that substring, then the value for that strategy is given by $e_i \sum_{\lambda=1}^{\Lambda} 2^{-S_\lambda}$ where $i = 1$ for, o_1, r_1, and k_1, while $i = 2$ for o_2, r_2, and k_2. The $6\Lambda + 1$th bit encodes wp, the production decision.

The strategy strings were subject to the standard crossover and mutation operators. Regarding crossover, agents are grouped randomly into $\frac{n}{2}$ pairs with each of the n agents in one pair. With an independent probability of 0.6, a pair of agents will swap some elements of their strategy string. If so, a number $c \in \{1,2,\ldots,6\Lambda\}$ is randomly chosen and the agents swap the first c elements of their strategy string. Regarding mutation, each element of each agent's strategy string, with an independent probability of 0.033 toggles its value—to zero if one and to one if zero.

The third genetic operator used, reproduction, was employed in two forms: tournament reproduction and proportional selection. Tournament reproduction was applied by randomly selecting two agents and adding the one which obtained the higher level of utility to the pool of strategy strings to exist in the next period. The process was implemented n times to determine the strategies of the n agents in the next period. Each of the n agents was equally available to be selected for each of the n tournaments. Proportional selection was applied by making n independent random selections of the agents' strategies to populate the subsequent period, where an agent's probability of being selected (during each of the n selections) was proportional to its obtained utility.

Variant 3.4—The strategy imitation and mutation regimes are replaced by a GA consisting of tournament reproduction, crossover, and genetic mutation, implemented in that order.

Variant 3.5—The strategy imitation and mutation regimes are replaced by a GA consisting of proportional selection, crossover, and genetic mutation, implemented in that order.

Both variants with GAs had strategy strings initialized where each bit of each strategy string had an independent even chance of starting as a zero or one.

The GAs' performance was mixed. It appears to find something near the competitive outcome, but relative to the other learning algorithms that did so, the GAs obtained the lowest mean utility: 81.5 with tournament reproduction and 80.2 with proportional selection. The resource allocations were biased towards a fifty-fifty split, with tournament reproduction less biased. Figure 16 shows resource allocation with the GAs and figure 17 shows mean price.

4.4 Spatially Placed Agents

One last pair of variants considered affects all three of the above mentioned features. Agents were placed on a lattice—as if occupying the squares of a checker board fifty spaces across and twenty high.

Variant 4.1—Agents interacted only with ten of their nearest neighbours—the eight adjacent agents, the agent two spots directly east, and the agent two spots directly west.

Variant 4.2—Agents visited ten other agents chosen randomly each period from the twenty-four nearest neighbours.

In both variants, for the purpose of providing agents on the edges of the lattice with a full complement of neighbours, the lattice is reflected on an infinite plain—equivalently, the lattice is a projection of the surface of a torus.

A priori it would seem that this alteration should have two effects. One is that by observing and imitating neighbours, an agent is choosing from strategies that interacted (last period) with agents that the agent in question is more likely to encounter (rather than some random subset of the entire population as in the base model). This may have the effect of making the search algorithm more effective, at least from the individual agents' point of view. Working against that, the spatial arrangement limits the flow of new, superior strategies (obtained by mutation) relative to the mix-and-match encounters of the base model. The agents in variant 4.1 obtained less utility than in variant 4.2 (86.9 vs. 93.1) presumably because the propagation of strategies was more restricted. The resource allocation was, however, similar in both variants to the base model, and so is not shown.

5. Conclusion

I analysed a model in which, due to the local nature of information and interaction, the agents' problem is extremely complex. With global information it is straightforward to ascertain the (optimal) competitive equilibrium.

Each agent chooses what to produce, what price to post and package size to use for its product, and a response to prices encountered. These decisions are guided by an imitative learning algorithm. Agents' behaviour at the global level appears to be guided by Smith's invisible hand, but this is entirely generated by and emerges from local interactions. Agents shift resources towards producing the good that is trading at a price above its long run competitive level. The price tends toward the market clearing level (although with a lag due to the local nature of information) because suppliers respond to excess supply by lowering prices.

Due to the lag in price adjustment, the economy cycles around the competitive equilibrium. If good one is initially overproduced, its market clearing price is greater than the long run competitive level. Prices move towards that market clearing level, eventually becoming greater than the long run competitive level. Agents start shifting resources into producing good one, but when it is eventually produced at the long run competitive level, prices have yet to reflect this. Agents thus continue shifting resources towards producing good one and it becomes overproduced. Consequently its price eventually falls below the long run competitive level and agents shift resources away from producing it. Because of the price lag, agents again overshoot the competitive level and return to under-producing good one.

A key aspect of the model is that agents engage in some price comparison prior to trading, and attempt to trade first with the most favourable price found. Without this price comparison, agents have no incentive to reduce prices in the face of excess supply--the costs of such a market disequilibrium are distributed randomly rather than to those whose actions (setting high prices) are creating (or maintaining) the disequilibrium. Agents can learn to practice price comparison, but not with great precision.

Additionally, a case was found where summary statistics of an agent-based model can be highly misleading—resource allocation and mean price were at competitive levels but every actual trade was being carried out at prices growing arbitrarily large.

The genetic algorithm (GA) proved unable to search effectively over strategies relative to the more directed learning algorithm employed which presumably must make better use of the information available to agents. This result should not be treated as condemning use of the GA in agent-based GE models. Dawid and Kopel (1998) showed that the outcome of a GA can be sensitive to the method by which a strategy is encoded as a binary string. Their result seemed to be driven by a significant amount of behavioural information being encoded in one particular bit of the strategy string—a feature shared here by the bit encoding the production decision.

Thus an agent-based general equilibrium model with an emergent invisible hand has been formulated. The assumptions behind the exchange environment do not closely resemble specific existing trading institutions, but provide an economic setting in which to evaluate learning schemes and their resultant dynamics. It remains to be seen if the insights garnered—that agents in the model need to be able to engage in price comparison and may need to employ a more focussed learning algorithm than the GA—can be applied in models with more realistic trading institutions.

References

Andersson, M. R.., Sandholm, T. W., 2001. Leveled commitment contracts with myopic and strategic agents. Journal of Economic Dynamics and Control 25, 615-640.

Arifovic, J., 2000. Evolutionary Algorithms in Macroeconomic Models. Macroeconomic Dynamics 4, 373-414.

Arifovic, J., 2001. Evolutionary Dynamics of Currency Substitution. Journal of Economic Dynamics and Control 25, 395-417.

Arifovic, J., 1995. Genetic Algorithms and Inflationary Economies. Journal of Monetary Economics 36, 219-43.

Arifovic, J., 2001. Performance of Rational and Boundedly Rational Agents in a Model with Persistent Exchange-Rate Volatility. Macroeconomic Dynamics 5, 204-24.

Arifovic, J., 1998. Stability of Equilibria under Genetic-Algorithm Adaptation: An Analysis. Macroeconomic Dynamics 2, 1-21.

Arifovic, J., 1996. The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies. Journal of Political Economy 104, 510-41.

Arifovic, J., Bullard, J., Duffy, J., 1997. The Transition from Stagnation to Growth: An Adaptive Learning Approach. Journal of Economic Growth 2, 185-209.

Arifovic, J., Eaton, C., 1998. The Evolution of Type Communication in a Sender/Receiver Game of Common Interest with Cheap Talk. Journal of Economic Dynamics and Control 22, 1187-1207.

Arifovic, J., Gencay, R., 2000. Statistical Properties of Genetic Learning in a Model of Exchange Rate. Journal of Economic Dynamics and Control 24, 981-1005.

Bower, J., Bunn, D., 2001. Experimental analysis of the efficiency of uniform-price versus discriminatory auctions in the England and Wales electricity market. Journal of Economic Dynamics and Control 25, 561-592.

Dawid, H., 1999. On the convergence of genetic learning in a double auction market. Journal of Economic Dynamics and Control 23, 1545-1567.

Dawid, H., 2000 On the emergence of exchange and mediation in a production economy. Journal of Economic Behavior and Organization 41, 27-53.

Dawid, H., Kopel, M., 1998. On economic applications of the genetic algorithm: a model of the cobweb type. Journal of Evolutionary Economics 8, 297-315.

De Vany, A., and Lee, C., 2001. Quality signals in information cascades and the dynamics of the distribution of motion picture box office revenues. Journal of Economic Dynamics and Control 25, 593-614.

Duffy, J., 2001. Learning to speculate: Experiments with artificial and real agents. Journal of Economic Dynamics and Control 25, 295-319.

Epstein, J. M., Axtell, R., 1996. Growing artificial societies: Social science from the bottom up. MIT Press, Cambridge, MA.

Holland, J.H., 1977. Adaptation in Natural and Artificial Systems. University of Michigan Press, Ann Arbor, MI.

Kirman, A. P., Vriend, N. J., 2001. Evolving Market Structure: An ACE Model of Price Dispersion and Loyalty. Journal of Economic Dynamics and Control 25, 459-502.

Kochugovindan, S., Vriend, N. J., 1998. Is the Study of Complex Adaptive Systems Going to Solve the Mystery of Adam Smith's 'Invisible Hand'? Independent Review 3, 53-66

LeBaron, B., 2000 Agent-based computational finance: Suggested readings and early research. Journal of Economic Dynamics and Control 24, 679-702.

Leijonhufvud, A., 1995. Adaptive Behavior, Market Processes and the Computable Approach. Revue Economique 46, 1497-1510.

Price, T. C., 1997. Using co-evolutionary programming to simulate strategic behaviour in markets. Journal of Evolutionary Economics 7, 219-254.

Rouchier, J., Bousquet, F., Requier-Desjardins, M., Antona, M., 2001. A multi-agent model for describing transhumance in North Cameroon: Comparison of different rationality to develop a routine. Journal of Economic Dynamics and Control 25 527-559.

Tay, N., Linn, S. C. 2001. Fuzzy inductive reasoning, expectation formation and the behavior of security prices. Journal of Economic Dynamics and Control 25, 321-361.

Tesfatsion, L., 2001. Introduction to the Special Issue on Agent-Based Computational Economics. Journal of Economic Dynamics and Control 25, 281-93.

Tesfatsion, L., 2001. Structure, Behavior, and Market Power in an Evolutionary Labor Market with Adaptive Search. Journal of Economic Dynamics and Control 25, 419-57.

Vriend, N.J., 2000. An Illustration of the Essential Difference between Individual and Social Learning, and Its Consequences for Computational Analyses. Journal of Economic

Dynamics and Control 24, 1-19.

Yildizoglu, M., 2002. Competing R&D Strategies in an Evolutionary Industry Model. Computational Economics 19, 51-65.

Table 1: Parameter Values

Parameter	**Parameter Name**	**Value Used**
n	Number of agents	1000
T	Number of periods	1000
e_1	Production level for good one	1
e_2	Production level for good two	1
v	Number of agents visited per period	10
ρ	Preference parameter	0.6 and 0.3
JCR	Job change rate	0.01
MR	Mutation rate	0.1
δ	Mutation radius	0.1

Table 2

Majority Pricing Strategy		**Minority Pricing Strategy**		**Stability Percentage**
Competitive	1	Higher than Competitive	101/100	75.6
Competitive	1	Lower than Competitive	99/100	7.9
Higher than Competitive	101/99	Competitive	1	98.8
Lower than Competitive	99/100	Competitive	1	85.5

Table 3: Prices with k_2 and wp held constant.

k_2	Implicit Price of Strategy			Price Observed in Trades		
	5th Percentile	Mean	95th Percentile	5th Percentile	Mean	95th Percentile
.3	1.68	1.86	2.04	1.16	1.19	1.22
.4	1.47	1.63	1.80	1.03	1.12	1.20
.5	1.07	1.09	1.12	0.805	0.812	0.818

Table 4 Summary the Base Model and Variants

Variant	Description	Mean Utility	Resource Allocation 600[13]	Resource Allocation 300	Mean Price
Base Model		94.2	599	312	0.989
1.1	Price ordering is a strategy	84.6	623	314	1.314
1.2	Reserve price	47.8	640	258	0.980
2.1	Respond 'optimally' to price	61.3	499	454	0.958
2.2	Utility falling cancels trade	92.6	582	378	0.985
3.1	Imitate only from best	83.2	634	376	1.023
3.2	Imitate only active strategies	92.2	588	330	1.000
3.3	Imitate probabilistically	96.2	600	324	0.992
3.4	GA Tournament reproduction	81.5	549	388	1.004
3.5	GA Proportional Selection	80.2	534	422	1.004
4.1	Lattice, visiting same always	86.9	617	269	0.963
4.2	Lattice, visiting 10 of 24	93.1	595	298	0.973

[13] Resource Allocation 600 is the mean number of agents producing good one when rho = .6 and so the optimal number of such agents is 600. Resource Allocation 300 is obtained similarly for rho = .3.

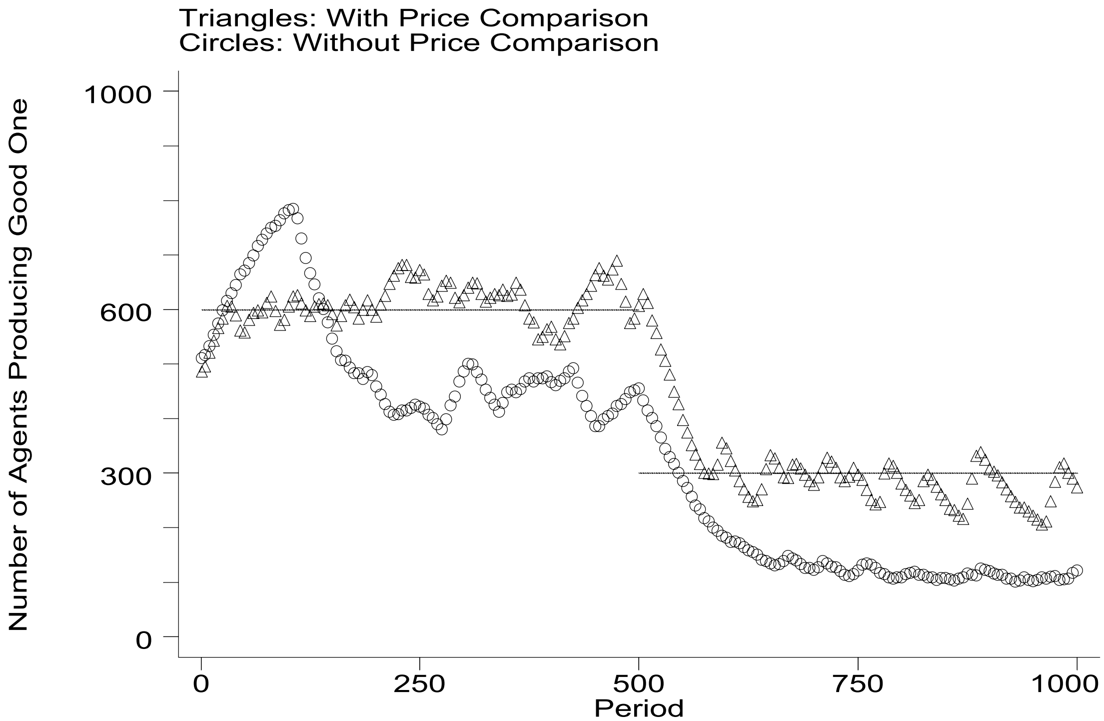

Figure 1: Resource Allocation with and without Price Comparison

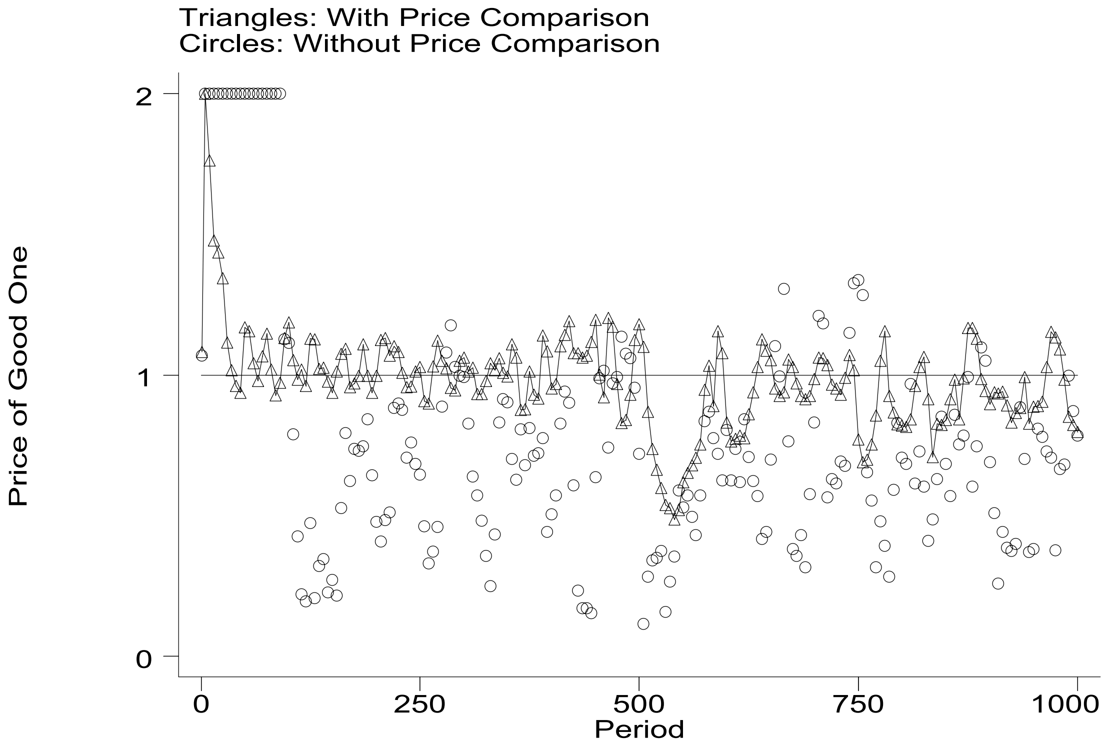

Figure 2: Mean Price of Good One with and without Price Comparison

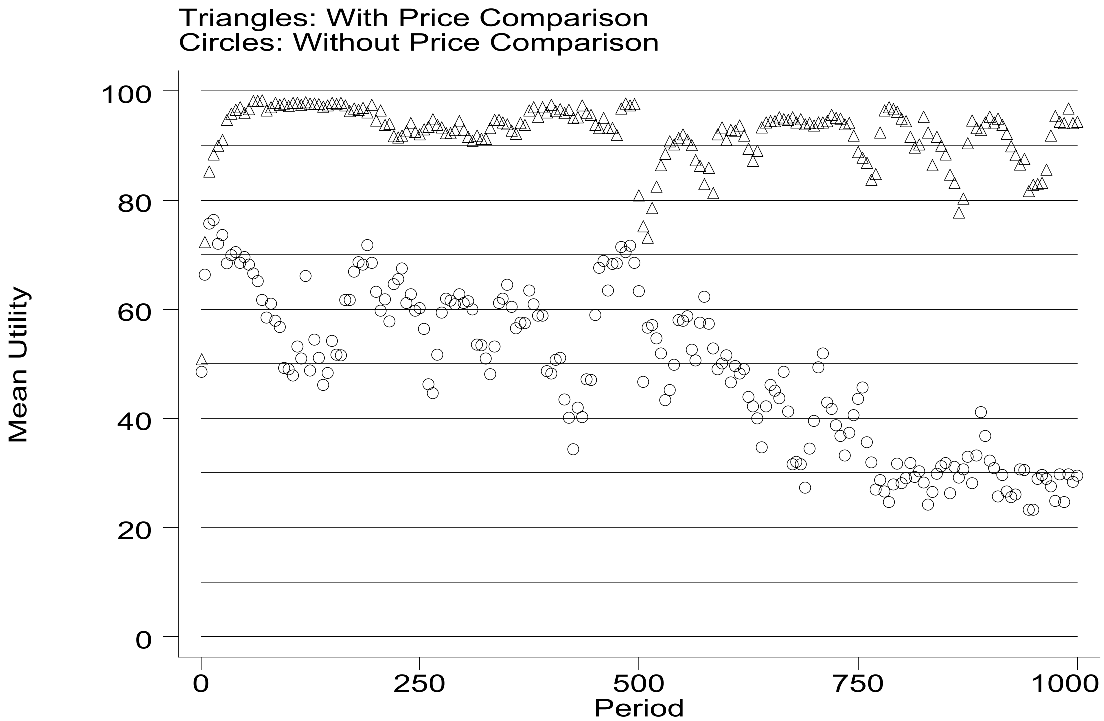

Figure 3: Mean Utility with and without Price Comparison

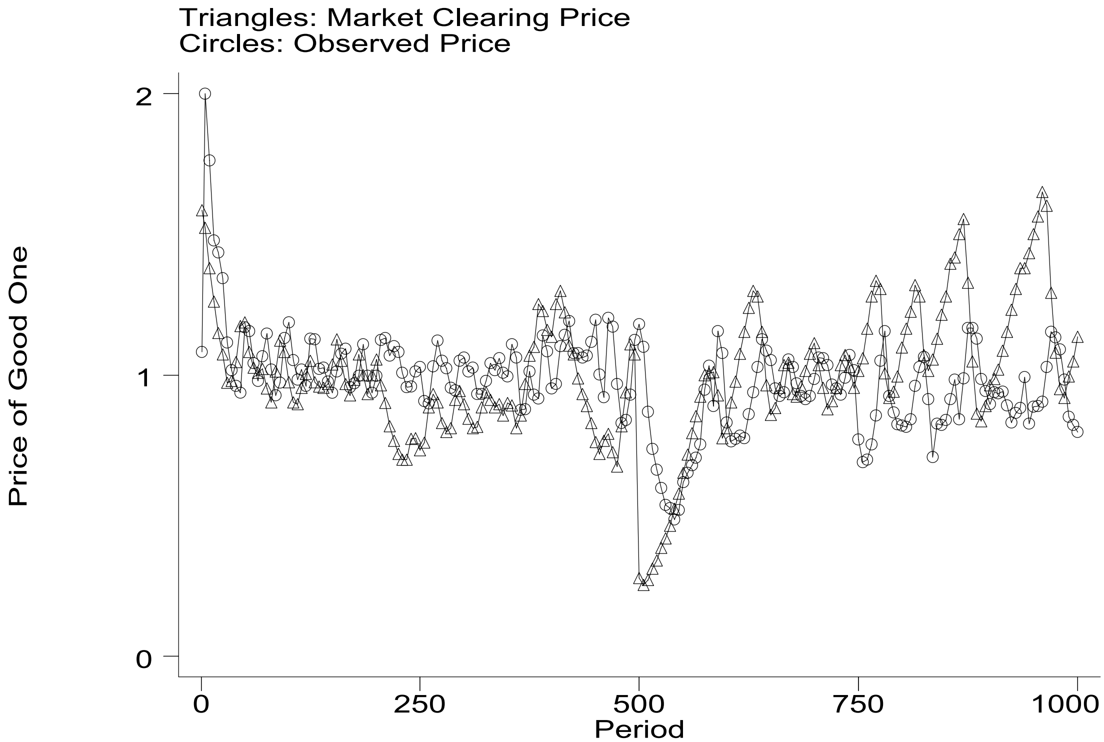

Figure 4: Market Clearing and Observed Prices

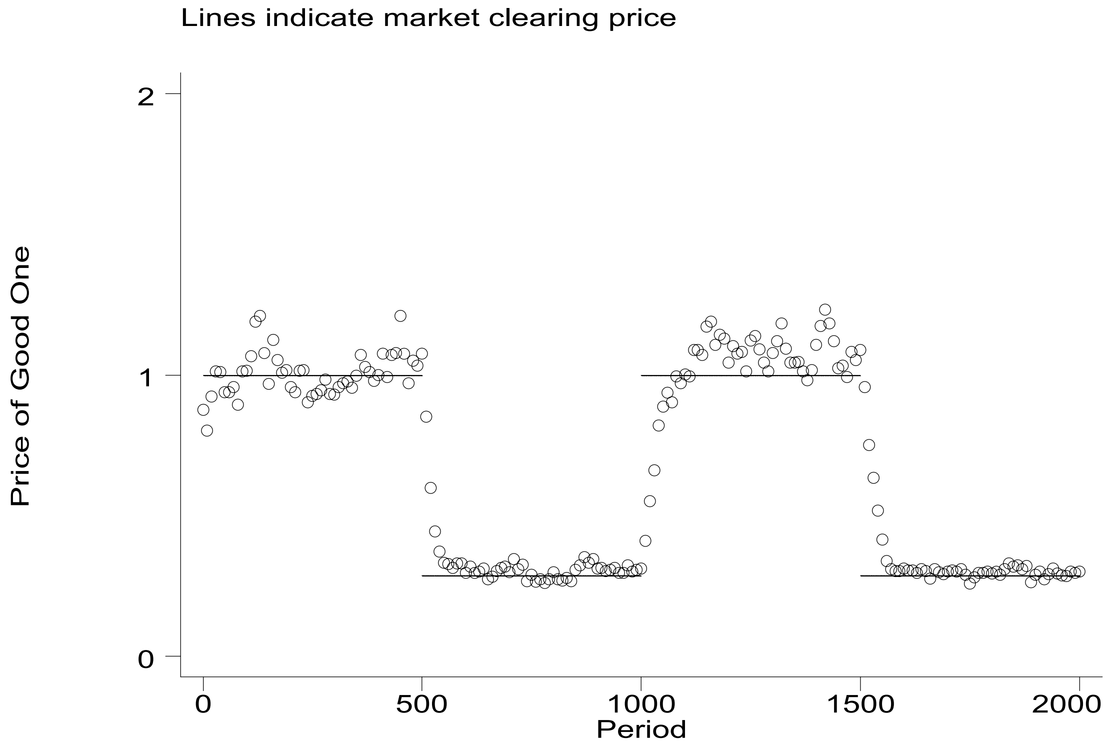

Figure 5: Observed Mean Price with Exogenous Resource Allocation

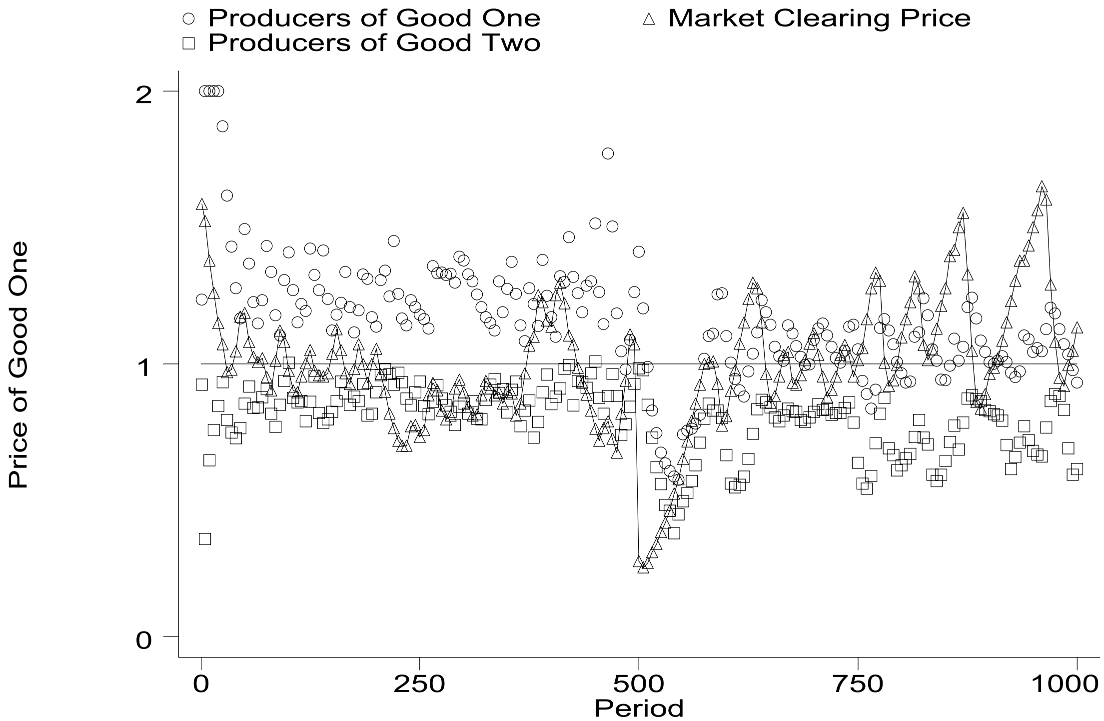

Figure 6: Mean Price of Good One in Offers

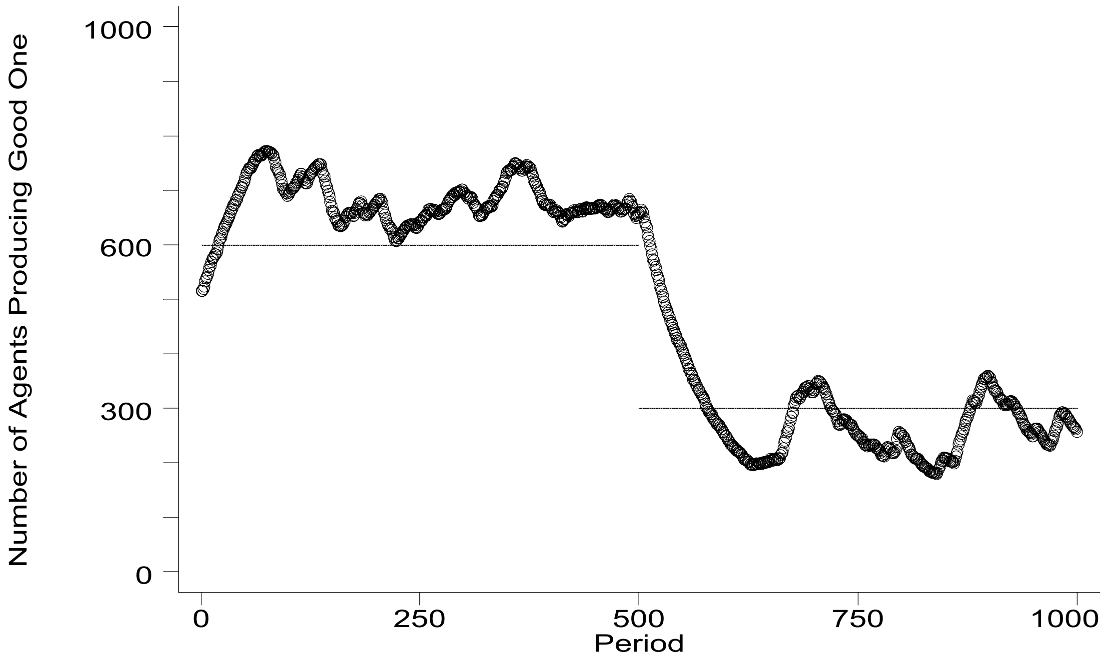

Figure 7: Resource Allocation—Variant 1.1

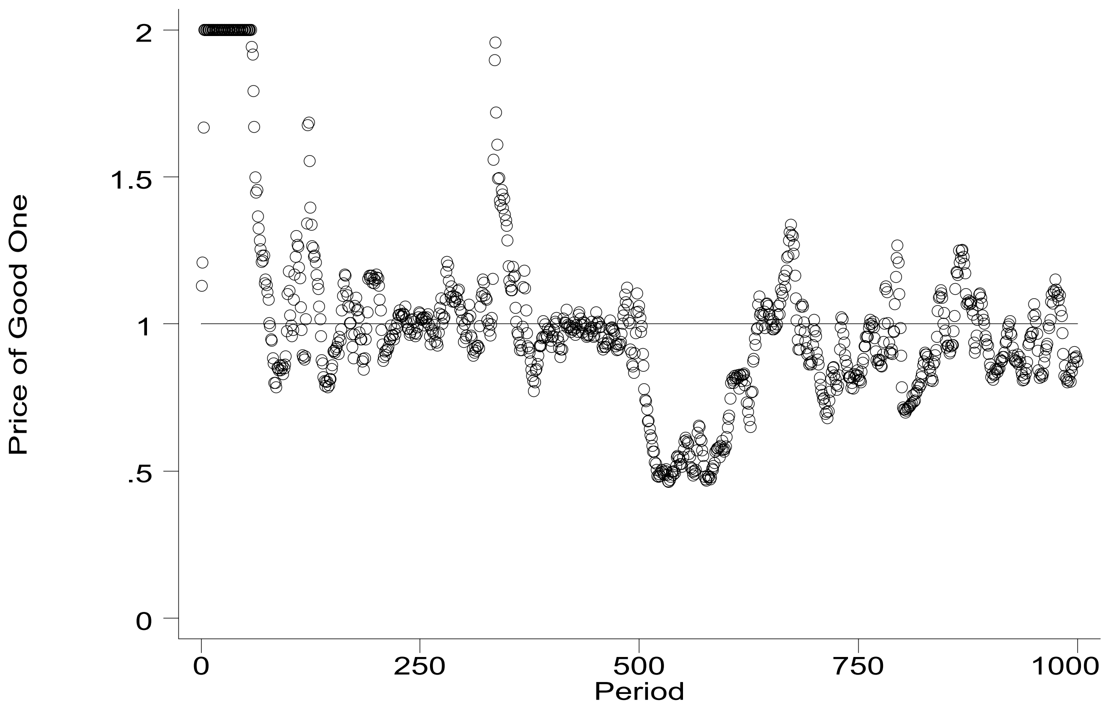

Figure 8: Mean Price—Variant 1.1

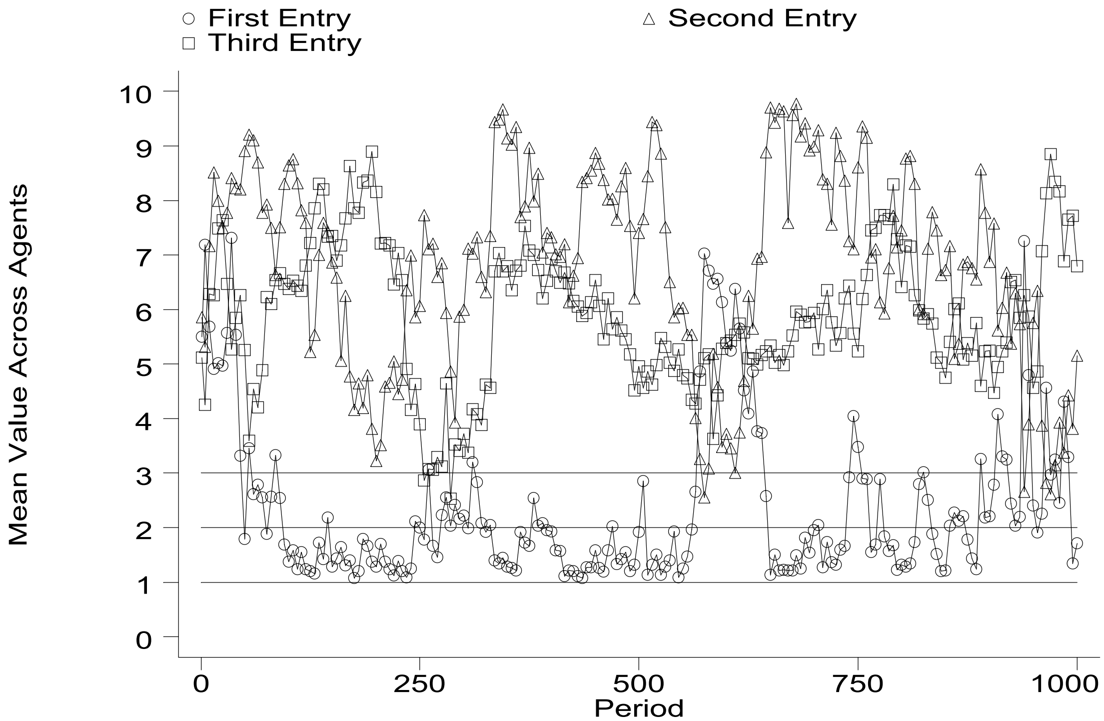

Figure 9: First, Second, and Third Entries in Ordering Strategy

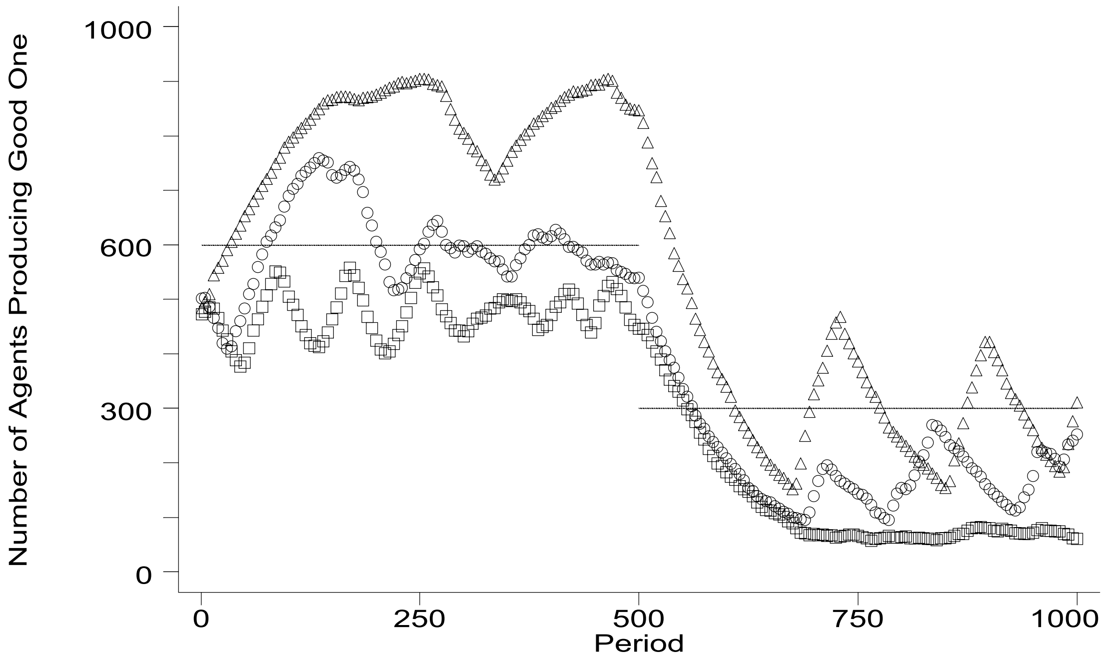

Figure 10: Resource Allocation—Variant 1.2

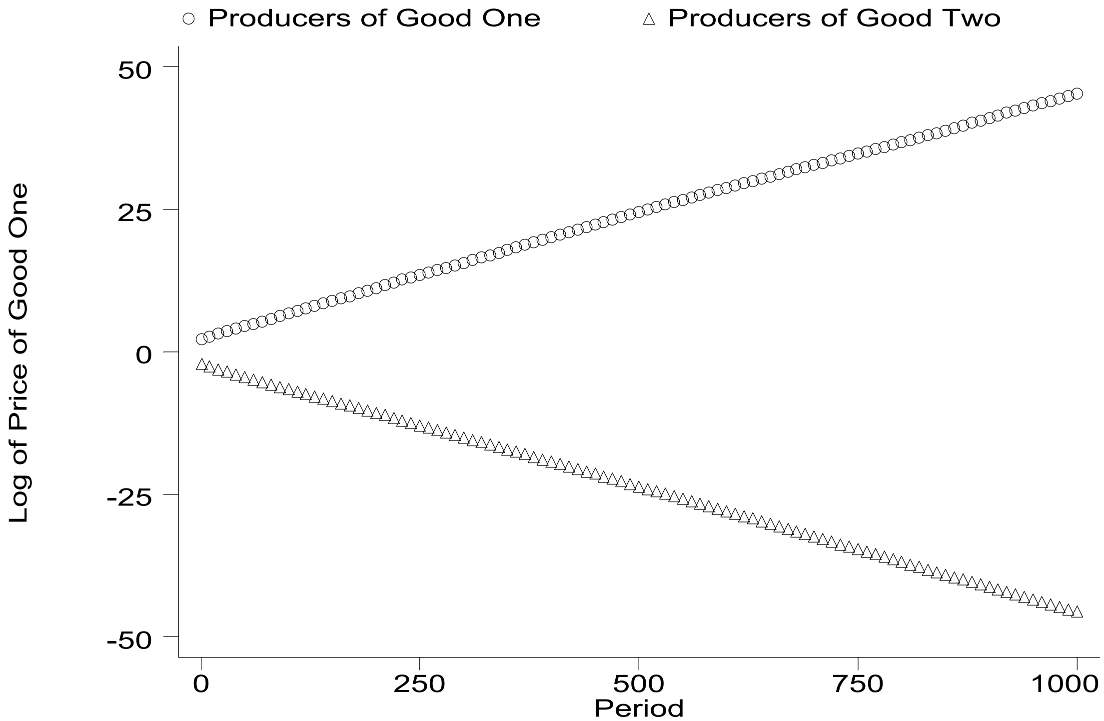

Figure 11: Mean Posted Price by Production Strategy—Variant 2.1

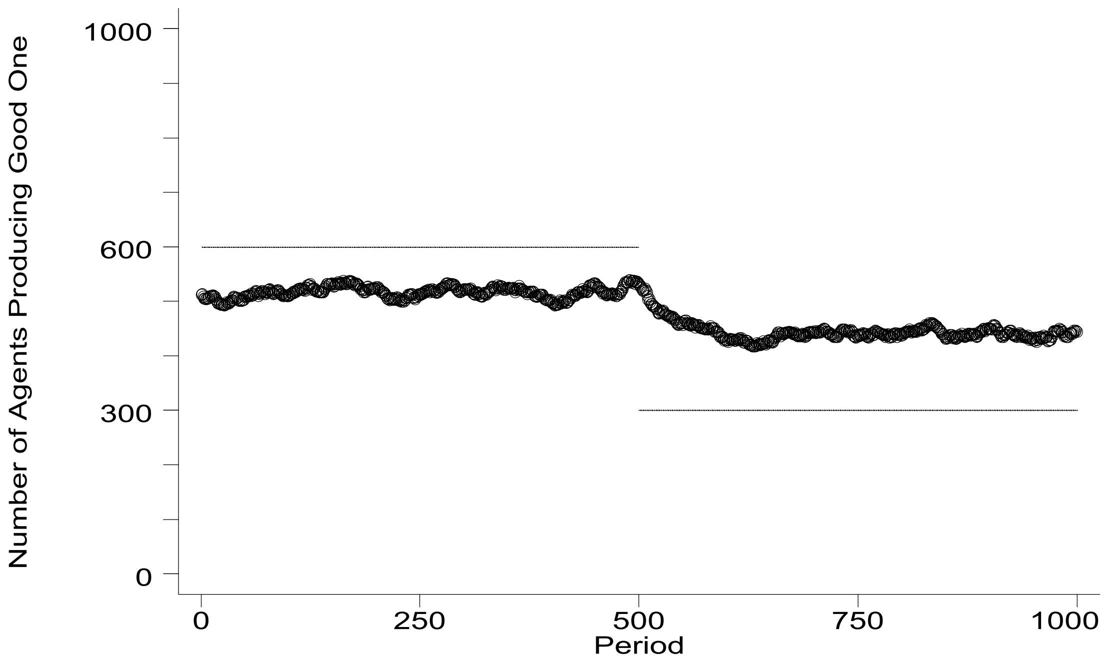

Figure 12: Resource Allocation—Variant 2.1

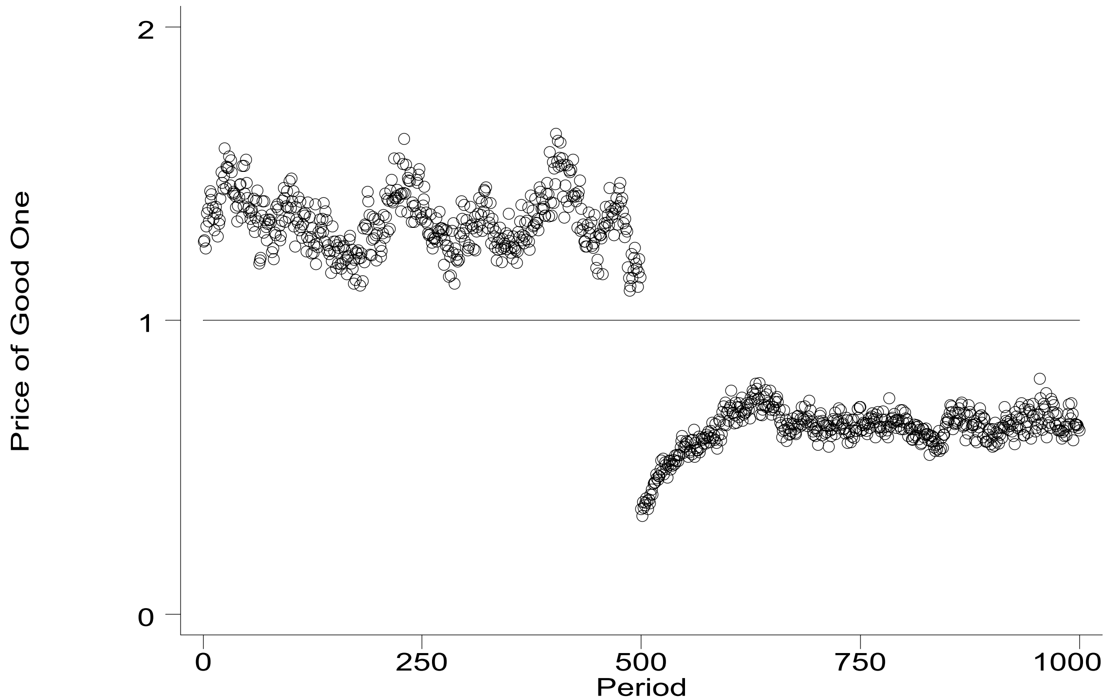

Figure 13: Price of Good One—Variant 2.1

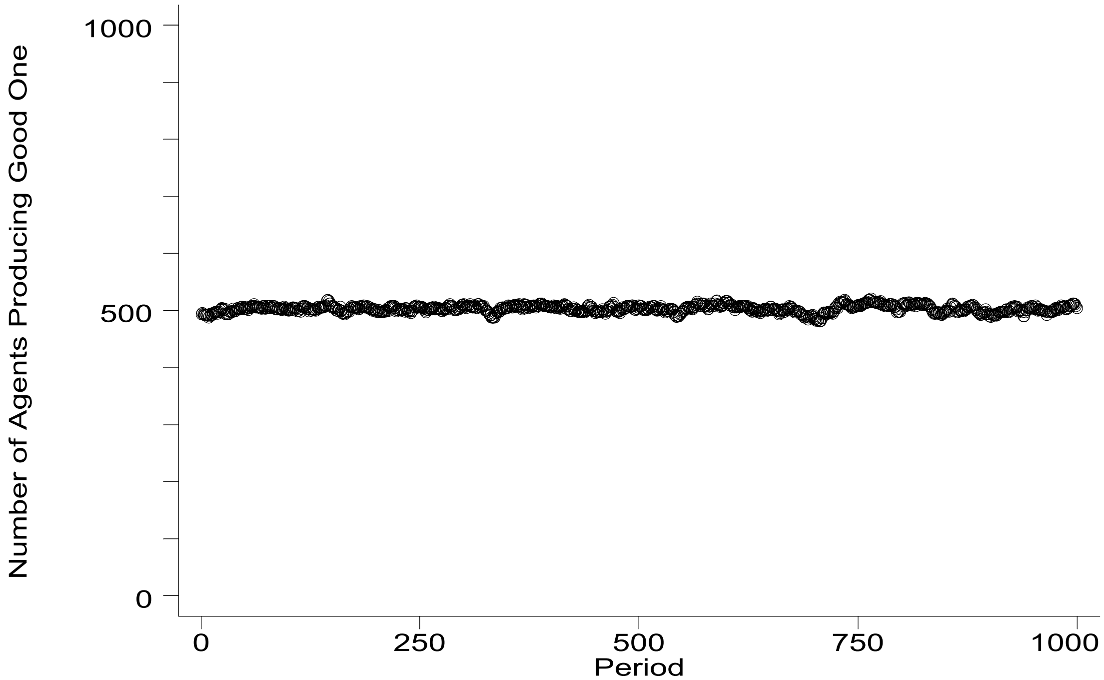

Figure 14: Resource Allocation—Variant 2.1, Rho Constant at 0.5.

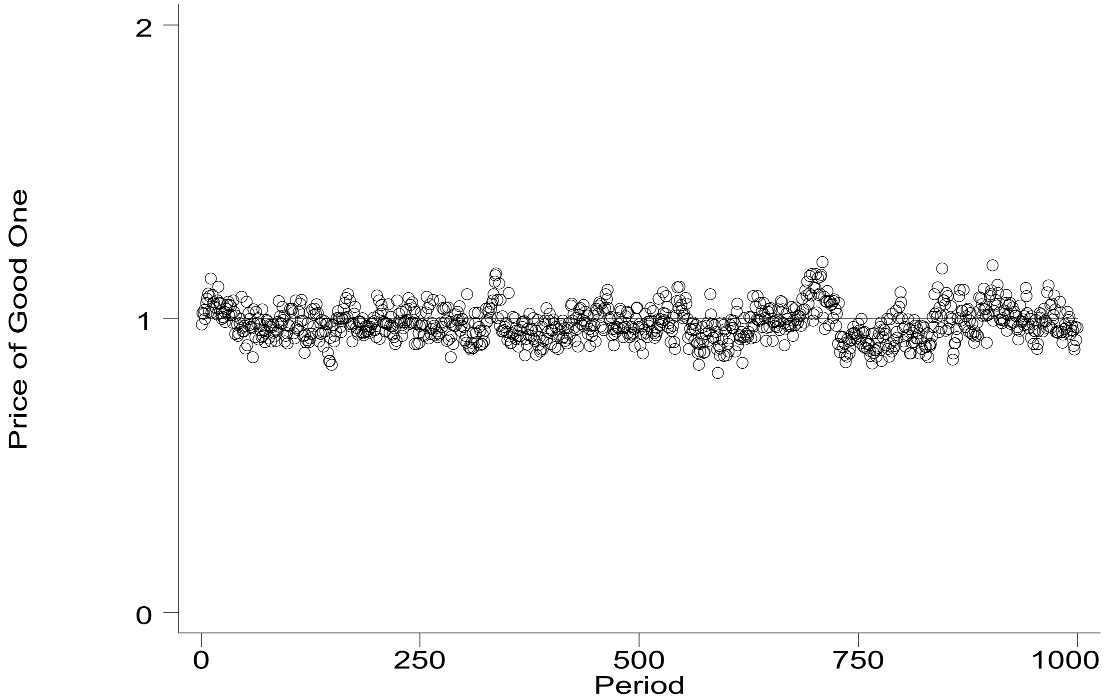

Figure 15: Price of Good One—Variant 2.1, Rho Constant at 0.5.

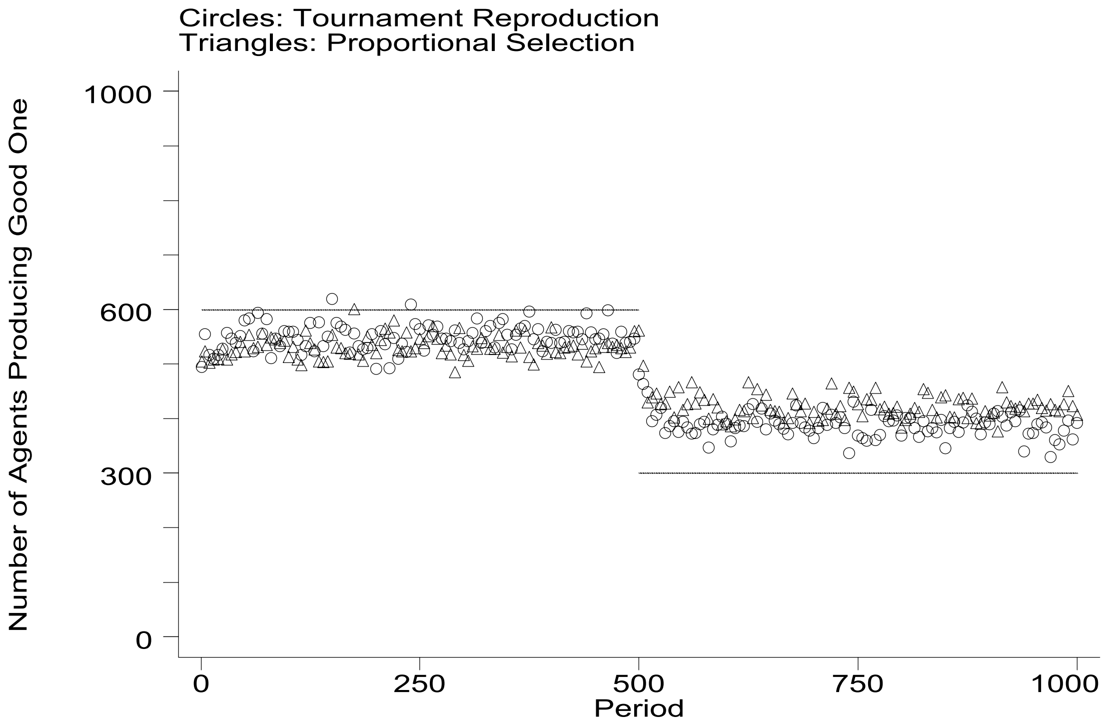

Figure 16: Resource Allocation with Genetic Algorithms

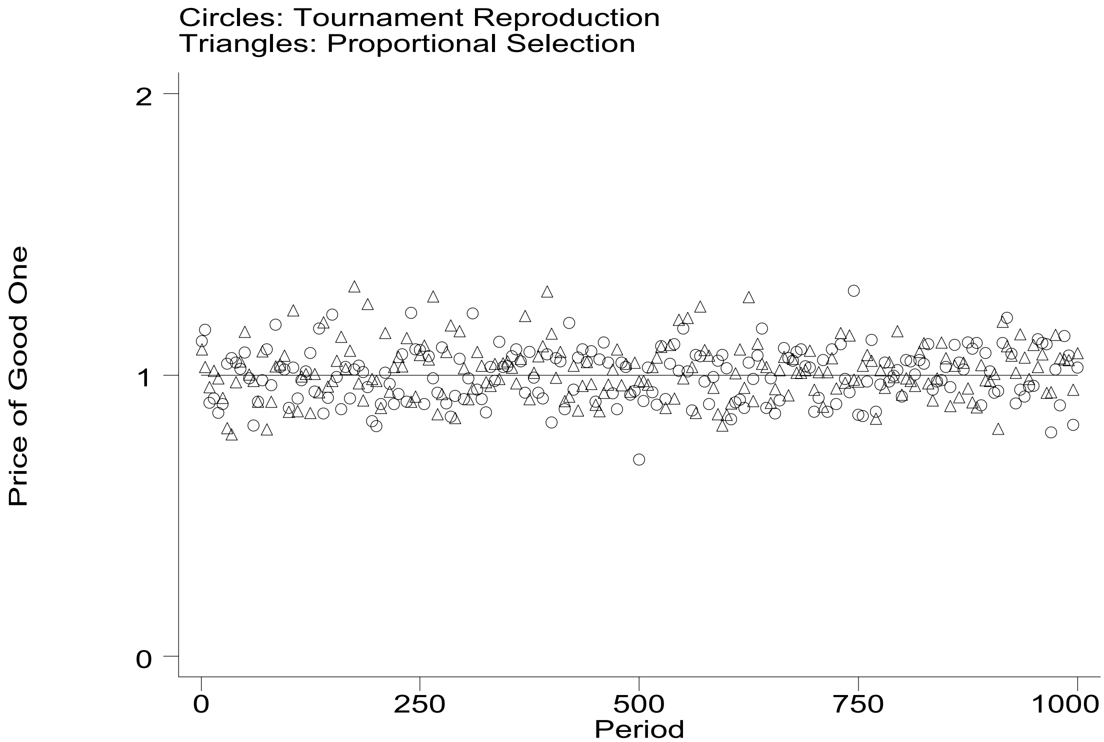

Figure 17: Price of Good One with Genetic Algorithms