

THE DYNAMICS OF GENERAL EQUILIBRIUM*

Herbert Gintis

The Walrasian general equilibrium model is the centrepiece of modern economic theory, but progress in understanding its dynamical properties has been meagre. This article shows that the instability of Walras' tâtonnement process is due to the public nature of prices, which leads to excessive correlation in the behaviour of economic agents. When prices are private information, a dynamic with a globally stable stationary state obtains in economies that are unstable in the tâtonnement process. We provide an agent-based model of a multi-sector Walrasian economy with production and exchange, in which prices are private information. This economy is dynamically well behaved.

In Walras' original description of general equilibrium (Walras, 1954 [1874]), market clearing was effected by a central authority. This authority, which has come to be known as the 'auctioneer', remains today because no one has succeeded in producing a plausible decentralised dynamic model of producers and consumers engaged in market interaction in which prices and quantities move towards market-clearing levels. Only under implausible assumptions can the continuous 'auctioneer' dynamic be shown to be stable (Fisher, 1983), and in a discrete model, even these assumptions (gross substitutability, for instance) do not preclude instability and chaos in price movements (Saari, 1985; Bala and Majumdar, 1992).¹ Moreover, contemporary analysis of excess demand functions suggests that restrictions on preferences are unlikely to entail the stability of tâtonnement (Sonnenschein, 1972, 1973; Debreu, 1974; Kirman and Koch, 1986).

It has been a half century since Debreu (1952) and Arrow and Debreu (1954) provided a satisfactory analysis of the equilibrium properties market economies, yet we know virtually nothing systematic about Walrasian dynamics. This suggests that we lack understanding of one or more fundamental properties of market exchange.

This article provides an agent-based model of the Walrasian economy. An agent-based model is a computer simulation of the repeated play of a game in which a large number of agents are endowed with software-encoded strategies governing both how they play the game and how they gather information and update their behaviour. The disequilibrium behaviour of agents in our agent-based models is governed by a *replicator dynamic* (Taylor and Jonker, 1978) in which, over time, successful agents tend in Darwinian fashion to increase in frequency at the expense of unsuccessful agents. We describe the process of shifting from lower to higher payoff strategies as 'imitation',

* I thank Kenneth Arrow, Robert Axtell, Samuel Bowles, John Geanakoplos, Peter Hammond, David Levine, Axel Leijonhufvud, Charles Plott, Sharon Harrison, Leigh Tesfatsion, two reviewers, and participants in seminars at the Santa Fe Institute, Columbia University, UCLA, and Trento for helpful comments, and the John D. and Catherine T. MacArthur Foundation for financial support. The computer programs used in this article are available from the author.

¹ Hirota (1981,1985,2003) shows that in an economy consisting of three goods, three consumers with generalised Leontief consumption functions and no production, the tâtonnement process is globally stable over 80% of the appropriate parameter space, but his methods apparently do not extend to more complex cases.

although this is indistinguishable from saying that unsuccessful agents die and are replaced by copies of successful agents. Agent-based modelling is effective in solving problems involving complex nonlinear dynamics that cannot be handled through standard optimisation techniques (Goldberg, 1989; Holland, 1975; Tesfatsion and Judd, 2006).

The results presented in this article suggest four general principles that are of clear relevance to market economies. First, a highly decentralised Walrasian economy, under a wide range of plausible conditions, has a unique, stable steady state in which the economy is reasonably close to Pareto efficient.

Second, the stability of a market system depends on the fact that prices are private information, in the sense that each agent, consumer, firm and worker possesses a set of reservation prices that are deployed to decide when and with whom to trade. These reservation prices are private information that each agent updates through time through trial and error, as well as by imitation. We call these *private prices*.

Third, when even a small fraction of agents are assumed to share the same price system and update in a coordinated manner, as suggested by the *tâtonnement* mechanism, the price system becomes highly volatile.

Fourth, a major mechanism leading to convergence of economic behaviour is *imitation* in which poorly performing agents copy the behaviour of better-performing agents. Under conditions of incomplete information, it can be shown that a positive level of imitation will always be sustained in equilibrium (Conlisk, 1988). This follows from the fact that if all agents engage in costly information-gathering and optimisation, under plausible conditions, a single agent can gain from copying the choice of others.

Each of these points is deserving of discussion and qualification. First, the uniqueness of equilibrium should be interpreted as follows. An agent-based model is a Markov chain of extremely high dimension. We introduce low frequency mutation in all agent parameters, so this Markov chain has no absorbing states, and hence by the ergodic theorem for Markov chains, it possesses a unique stationary distribution, independent of any other modelling assumptions (Feller, 1950). In general, this stationary distribution will be characterised by the existence of one or more 'attracting states' such that the Markov chain spends the bulk of the time in the neighbourhood of one of these states, although there can be orders of magnitudes of difference in the fraction of time it spends near one versus another attracting state (Freidlin and Wentzell, 1984). By a 'unique, stable steady state' I mean that there is likely only one such attracting state, and there is certainly only one when the model is initialised with parameters that are within an order of magnitude of their steady state values. I assert this as an empirical fact of the model, not something that I can prove analytically. The assertion of uniqueness 'under a wide range of plausible conditions' means that I have found no counter-example in the five or so years I have worked with various versions of the model. I can easily conceive of implausible and uninteresting conditions under which there would be multiple equilibria. For instance, suppose there is a set of goods that are highly complementary in consumption. Then there may be two attracting states, one in which members of the set are produced and consumed at a high level, and another in which they are not produced. The probability of moving from one attracting state to the other will be close to zero, requiring multiple, highly correlated mutations sustained over many periods.

It is worth noting that the very simple examples of multiple equilibria in the literature on multi-market economies refer to Scarf-type exchange economies without production, in which agents begin with endowments of these goods. For instance, Mas-Colell *et al.* (1995) provide an example of a two-good, two-agent economy (Edgeworth box, Exercise 15.B.6) in which, for particular initial endowments of each agent, there are four market-clearing price ratios. However, in an agent-based representation of this economy, using the same utility functions, when production is added to the system and agents supply labour and capital rather than holding an endowment of the final goods, there is only one steady-state solution. In this solution relative prices are dictated by relative production costs, and have no relationship to the four equilibria in the Scarf-type example.

Second, the violation of the Law of One Price is short-term only. Prices in the agent-based model are necessarily ergodic, the long-run historical average price equalling the equilibrium price. Moreover, the existence of significant price dispersion even in integrated markets, such as nation-states or the European Union, is confirmed by numerous empirical studies. For a recent overview and replication, see Allington *et al.* (2004).

Third, the fact that a highly diffuse set of private prices adds stability to the general equilibrium model when contrasted with the tâtonnement public price mechanism is explained by the fact that public prices act as a system-wide public correlation device that induces highly correlated behaviour on the part of economic actors, who all react to a price change in parallel fashion. An attempt by the auctioneer to dampen this effect by choosing very small price increments fails because so doing slows the rate of price change but not the path of change itself. It is thus the 'disorganisation' of private prices that contributes to its being a powerful equilibrating device. In effect, private pricing avoids the Draconian price adjustment process characteristic of tâtonnement, being rather a diffuse decentralised information processing mechanism that steers the economy, albeit slowly, towards long-run market clearing prices.

Fourth, imitation is considered second-rate learning in economics but it is the bedrock of evolutionary models. Genetic offspring inherit all their genes from the parents, with a very low rate of mutation. Cultural transmission is predominantly by imitation in biology, anthropology and sociology, and the importance of culture to a species depends largely on its capacity to imitate. This capacity is much more highly developed in humans than in any other mammalian species (Bandura, 1977; Tomasello, 1999). By contrast, individual learning by experience, the mechanism preferred in economic theory, is generally extremely slow and inefficient, except where volatility renders imitation ineffective (Conlisk, 1988).

1. Overview

To address the issue of private vs. public prices, I will use the highly simplified three-good exchange economy of Scarf (1960). In the Scarf economy, each agent produces one good and consumes some of his own good plus some of another good, in fixed proportions. Scarf showed that, if we label the goods X , Y , and Z , and if X -producers consume X and Y , Y -producers consume Y and Z , and Z -producers consume Z and X , then with a tâtonnement process of price adjustment, equilibrium prices follow closed paths in price space, and hence are not asymptotically stable. Hirota (1981) completely

characterised price paths for this economy. Recently, Anderson *et al.* (2004) have obtained experimental results with human subjects that indicate that this same economy also fails to converge in a laboratory setting where prices are set by a double auction.

Our agent-based model of the Scarf economy assumes private prices, so the law of one price does not hold. Private prices and imitation alone ensure that this economy converges to the appropriate steady state from arbitrary initial conditions (see the previous Section for the meaning of 'steady state'). This suggests that the problem with the tâtonnement and double auction processes is that they assume that prices are public information.

To support this interpretation, I study a model in which the fraction of agents who use public prices can be experimentally varied. When this fraction is zero, we observe strong convergence to steady state prices. When the fraction is 10%, there is significant long-run price volatility, when the fraction is 40%, prices are highly volatile and when the fraction is 100%, the system behaves precisely in the manner described by Scarf (1960).

Following the treatment of the Scarf economy, I analyse an agent-based Walrasian economy with a fixed number of market sectors (10 in this model), a constant number (5,000) of agents who consume, work and own financial assets, a single financial asset that is also a factor of production, and a variable number of firms in each market sector (averaging about 14 firms per sector). The equilibrium conditions for this model are market clearing in all sectors. However, because prices are private information and agents and firms gather information through search processes that have stochastic elements, disequilibrium is the general state of the system. In this model, as in real life, all agents regularly engage in out-of-equilibrium trades and production, so both prices and quantities respond to conditions of excess supply. Moreover, agents maximise utility by searching for favourable goods prices and employment opportunities. Their optimisation techniques include both experimentation (e.g., firms vary their product prices, wage rates, capital demand, and other firm operating parameters) and imitation (failing firms copy the operating characteristics of more successful firms).

2. The Scarf Economy with Public Prices

Our model is that of Scarf (1960), except that we follow Anderson *et al.* (2004) in choosing parameters so that equilibrium relative prices are extremely unequal. We assume there are three goods, X , Y , and Z and one agent for each good. The X -producer is endowed with $O_x = 10$ units of X , the Y -producer is endowed with $O_y = 20$ units of Y , and the Z -producer is endowed with $O_z = 400$ units of Z . The X -producer consumes X and Y in proportion $x/O_x = y/O_y$, so his utility is given by

$$u_x(x, y, z) = \min\left(\frac{x}{O_x}, \frac{y}{O_y}\right).$$

Similarly, for the other two agents, we have

$$u_y(x, y, z) = \min\left(\frac{y}{O_y}, \frac{z}{O_z}\right)$$

and

$$u_z(x, y, z) = \min\left(\frac{x}{O_x}, \frac{z}{O_z}\right).$$

An agent who produces good $H \in \{X, Y, Z\}$ and consumes H and $W \in \{X, Y, Z\}$ optimises by solving the follow pair of equations, where P_h is the price of H and p_w is the price of W , h_o is the amount of H he consumes and w_o is the amount of W he consumes:

$$\frac{h_o}{O_h} = \frac{w_o}{O_w} \quad \text{and} \quad p_h(O_h - h_o) = p_w w_o,$$

which give the agents final demand (h_o, w_o) :

$$h_o = \frac{O_h^2 P_h}{O_h p_h + O_w p_w}.$$

This equation allows us to calculate total excess demand E_g for each good $g \in \{X, Y, Z\}$ as a function of the prices of the three goods.

It is easy to check that the market-clearing prices, normalising $p_z = 1$, are given by $p_x^* = O_z/O_x$, $p_y^* = O_z/O_y$, and $p_z^* = 1$. Numerically, this becomes $p_x^* = 40$, $p_y^* = 20$, and $p_z^* = 1$.

Suppose we start with disequilibrium prices $p_x = p_x^* + 3$, $p_y = p_y^* - 2$, $p_z = p_z^* = 1$. The auctioneer broadcasts these prices, and each of the agents registers his demand for the three goods as a function of these prices. The auctioneer then updates prices by the equations

$$p'_x = p_x + E_x/100$$

and

$$p'_y = p_y + E_y/100.$$

We calculate that the new prices are $p'_x = 43.0026$ and $p'_y = 18.0036$. The process of calling out prices, collecting excess demand amounts c and updating prices is then repeated indefinitely. Here, as throughout the simulation, prices change in any period by less than 0.02%. The result after 5,200 rounds is shown in Figure 1, and perfectly replicates the analytical results of Scarf (1960).

The performance of the Scarf three-good economy with public prices is in fact considerably worse than that depicted in Figure 1 if we add a small amount of noise to the demand functions of the three agents in each period. In Figure 2, we multiplied each public price for each agent by a random variable uniformly distributed on the interval $[0.8, 1.2]$ before calculating that individual's demand and supply. The Figure shows that the resulting price trajectory moves monotonically at a positive rate away from the equilibrium. Lowering the amount of noise slows but does not otherwise alter this result.

3. The Scarf Economy with Private Prices

For our private price model we maintain all of the above assumptions, except now there are 1,000 traders of each of the three types, each trader is endowed at the beginning of a run a with a set of private prices randomly drawn from the uniform distribution on




Fig. 1. *Non-convergence of Price to Steady State in a Three-good Scarf Economy with Public Prices*
 Note: Prices move in an ellipse in the counter-clockwise direction. Here, (P_x^o, P_y^o) is the initial relative price ratios




Fig. 2. *Shows the Increase Over Time in the Maximum Positive Deviation From the Equilibrium*
 Note: Relative Price of the X Good When a Small Amount of Noise is Added to Each Public Price for Each Agent for Good in Each Period

(0, 1). The general structure of an agent-based model is depicted in Appendix I. There are 2,500 generations and 10 periods per generation. At the start of each period, each agent’s inventory is re-initialised to O_h units of the production good H, and zero units of the other goods. Each agent in turn is then designated a trade initiator and is paired with a randomly chosen responder, who can either accept or reject the proposed trade. Each agent is thus an initiator exactly once and responder on average once per period. After a successful trade, agents consume whatever is feasible from their updated inventory.²

² The number of trades per production period can be increased without altering the behaviour of the economy, except that convergence is faster.

In the reproduction stage, which occurs every ten periods, 5% of agents are randomly chosen either to copy a more successful agent or to be copied by a less successful agent, where success is measured by total undiscounted utility of consumption over the previous ten periods. Such an agent is chosen randomly and assigned a randomly chosen partner with the same production and consumption parameters. The less successful of the pair then copies the private prices of the more successful. In addition, after the reproduction stage, each price of each agent is mutated with 1% probability, the new price either increasing or decreasing by 10%.

The trade procedure is as follows. The initiator offers a certain quantity of one good in exchange for a certain quantity of a second good. If the responder has some of the second good, and if the value of what he gets exceeds the value of what he gives up, according to his private prices, then he agrees to trade. If he has less of the second good than the initiator wants, the trade is scaled down proportionally.

The initiator's trade ratios are given by his private prices. Which good he offers to trade for which other good is determined as follows. Let us call an agent's production good his *P*-good, the additional good he consumes his *C*-good, and the good which he neither produces nor consumes his *T*-good. Note that agents must be willing to acquire their *T*-good despite the fact that it does not enter their utility function. This is because *X*-producers want *Y*, but *Y*-producers do not want *X*. Only *Z*-producers want *X*. Since a similar situation holds with *Y*-producers and *Z*-producers, consumption ultimately depends on at least one type of producer accepting the *T*-good in trade, and then using the *T*-good to purchase their *C*-good.

If the initiator has his *T*-good in inventory, he offers to trade this for his *C*-good. If this offer is rejected, he offers to trade for his *P*-good. This may sound bizarre, but it conforms to the general rule of agreeing to trade as long as the value of one's inventory increases. If the initiator does not have his *T*-good but has his *P*-good, he offers this in trade for his *C*-good. If this is rejected, he offers to trade half his *P*-good for his *T*-good.³ If the initiator had neither his *T*-good nor his *P*-good, he offers his *C*-good in trade for his *P*-good, and if this fails he offers to trade for his *T*-good. In all cases, when a trade is carried out, the terms are dictated by the initiator and the amount is the maximum compatible with the inventories of the initiator and responder.

The results of a typical run are exhibited in Figure 3. Each of the curves in Figure 3 is given by $(p - p^*)/p^*$, where p^* is the equilibrium relative price. Because initial prices are generated by uniform distributions on the unit interval, the initial mean price of the *Z* good is 40 times its equilibrium value, and the price of the *Y* good is 20 times its equilibrium value. Nevertheless, in sharp contrast to the Scarf economy with public prices, convergence to a steady state is rapid and complete. Moreover, the standard deviation of prices falls monotonically to near zero.

The public price dynamics in this economy are determined as follows. In each period, given current prices, excess demands for all goods are calculated and prices are increased or decreased by an amount proportion to excess demand for the good in question. Prices are then normalised so the price of the *Z* good is unity. The proportionality coefficient is set so prices in any one period move no more than by 1% to avoid cobweb-type phenomena.

³ The fraction 1/2 is not optimal, but the model works sufficiently well that allowing this fraction to evolve endogenously was unnecessary.




Fig. 3. *Convergence of Price to Steady State in a Three-good Scarf Economy with Private Prices*

To show that public prices destabilise the Scarf economy, we reintroduce the tâtonnement process, where a fraction f of the population uses public prices, and the auctioneer updates public prices in each period by aggregating the excess demands of the agents who use public prices. However, we retain the assumption that agents trade and consume in each period even though prices are not in equilibrium. Also, agents who use private prices can imitate the price structure of a public price agent who has been more successful over the past ten periods. Such an agent, however, remains a private price agent.

Figure 4 shows the result of this run of the agent-based Scarf economy with 10% (top pane) and 40% (bottom pane) of agents using public prices. We see that even 10% public prices entails considerable initial price instability but this wears off eventually. With 40% public price agents, there are recurrent price explosions in which prices deviate manifold from their equilibrium values.

4. An Agent-based Walrasian Economy

The standard Walrasian economy has many goods, with many firms producing each good. Firms rent capital and hire labour to produce a single good. The returns to labour and capital accrue to individuals who use their income to purchase the goods produced by the firms. In competitive equilibrium, each agent maximises utility subject to an income constraint, all markets clear and firms earn zero excess profits. The general equilibrium existence theorem says that if certain technical conditions are satisfied, then for every distribution of ownership of factors of production, there is a price vector giving rise to a competitive equilibrium. The Fundamental Theorem of Welfare Economics states that every equilibrium is Pareto-efficient, and every efficient distribution of utility among agents can be attained as a competitive equilibrium for some initial distribution of ownership, again provided certain conditions are satisfied (Arrow, 1951; Debreu, 1952; Arrow and Debreu, 1954).

In our version of the Walrasian model, demand is decentralised, each agent having a private utility function with individualised parameters (disutility of labour and discount rate), giving rise to a supply of labour function reflecting an individual's trade-off




Fig. 4. *Scarf Price Dynamics with Public Price and Private Price Agents*

Each of the curves is given by $(p - \bar{p}) / \bar{p}$, where \bar{p} is the equilibrium relative price and p is the average observed relative price. Equilibrium corresponds to zero on the y-axis. The top pane represents 10% public prices, and the bottom pane 40% public prices

between income and the disutility of effort. Effort is visible to the firm, so an enforceable contract can be written specifying the wage paid and the effort received. Also, each agent is endowed with a fraction of the total capital stock from the rental of which the agent derives non-labour income.

We include one centralised institution, which we call the Monetary Authority, whose existence is required by the fact that the economy must have a monetary system. Firms sell their product on markets in exchange for money and make monetary payments to factors of production that use the money to purchase products. The Monetary Authority has the power to create money but does so only under two conditions. First, a new firm is loaned enough money for one period of production and sales, and firms that lose money are loaned enough to continue in business until their poor performance leads to their dissolution through bankruptcy (in each reproduction period, the least well performing 5% of firms are forced into bankruptcy). Conversely, if a firm makes positive profits, this is taxed away by the Monetary Authority. Second, unemployed workers are given unemployment insurance by the Monetary Authority, although this is recouped by the excess profits tax and a tax on wages. Since the Monetary Authority is not obliged to run a balanced budget, the money supply need not be constant. As we shall see, after an initial several hundred periods of volatility, excess profits are virtually zero and the wage tax fully covers unemployment compensation expenditure, so the money supply

is constant, as are long-run average prices. Indeed, consumer wealth (which is an accumulation of money in private hands) runs less than 1% of national income (which we define as wage income plus the return to capital), and a tendency for firms to run a net positive profit is almost exactly offset by the losses of bankrupt firms. Infrequently, however, the economy experiences considerable volatility, in which case the liquidity provided by the monetary sector is an important stabilising influence.

In each period, consumers sample a fraction of the firms in a sector (a *sector* is the set of firms supplying a particular good) to assess the price structure, and then contract with particular firms in each sector. If current prices appear too high to the consumer, or if firms run out of goods, consumers can carry their purchasing power over to the next period. Consumers live forever, never reproduce (although imitating others is a form of death and rebirth), and never retire from the labour force (although they may be unemployed, in which case they receive unemployment insurance from the Monetary Authority). Thus, there is no incentive to accumulate wealth except to keep some purchasing power in reserve in case favourable purchases cannot be effected in a particular period. Moreover, all goods are infinitely divisible, so there is no role for consumer credit or consumer saving for large purchases.

Agents who are unemployed stochastically explore employment opportunities and contract with firms when an acceptable job offer is encountered. A job offer is considered acceptable if the utility flowing from the job is greater than the agent's discount rate times the expected present value of remaining unemployed (which in turn is a function of the expected length of unemployment, the expected future wage and length of employment). This expected present value, which we call the worker's *fallback position*, is determined as described in Section 4.1. Results are similar if we force the fallback to be zero, so workers always accept a job when the wage exceeds the disutility of labour, except that in this case the level of volatility of prices, employment and quantities are reduced considerably.

All firms in the same sector produce the same good, but each firm sets the price of the good independently from other firms, and sells to consumers who happen to discover, and accept, the exchange that it is currently offering. The number of firms is endogenously determined, depending on cost and demand conditions. In the results presented, we initialise the number of firms to $n = 10$, but the steady state number is about 14.3. Firms can carry unsold merchandise over to succeeding periods, but there is a positive depreciation rate. We endow firms in the same sector with the same production function, thus allowing us to assess how close production and pricing is to Pareto-optimality. However, firms have incomplete knowledge of the conditions of demand and the characteristics of their labour force. Therefore, firms maximise profits by choosing and adjusting certain *operating characteristics*, including product price, target labour force size, wage offer and labour effort requirement, adjusting these over time in response to excess supply in product and labour markets, inventory size and realised profits (see Table 1).

It should be emphasised the firms do not know either the unemployment rate or the degree of excess supply or demand in their own (or any other) sector. Firms do know whether their inventory is increasing or decreasing, and whether their workforce is increasing or decreasing. Firm wage, price, labour demand and production target depend on these data alone. Given these characteristics, firms can estimate the marginal product of capital, thus determining their demand for capital. Moreover, firms

Table 1
Firm Operating Characteristics

Operation Characteristic	Initialisation
Product Price	2 to 6
Wage Offer	2 to 5
Effort Demanded	0.8 to 0.99
Target Employment	5 to 60
Price Adjustment Coefficient	0.95 to 0.99
Wage and Effort Adjustment Coefficients	0.95 to 0.99

The characteristics are initialised using a uniform distribution on an interval given in column 2.

that have high profits tend to have their choice of operating characteristics copied by less successful firms.

4.1. *Workers*

We assume m agents, who are both suppliers of factors of production and consumers of sectoral output. In the results presented, we assume there are ten sectors, so $m = 500$ and $n = 5,000$. For simplicity, we assume the agent's utility function is separable into a supply side, as worker, and demand side, as consumer. Each worker i has a labour utility function of the form

$$u_i = \frac{w - a_i}{1 - e}, \quad (1)$$

where w is the wage, $e \in [0, 1)$ is the effort supplied to the employer, and a_i is a constant reflecting the disutility of effort. The shape of the function is not crucial. Note that a_i is the cost of showing up for work, even if one exerts zero effort, and as effort approaches unity, the disutility of effort becomes infinite. Effort e is set contractually. Note that the wage is nominal. The disutility of effort in effect 'pegs' the price system for the economy, since it is the only non-price that has a value in terms of a price (the wage rate). The derivation of the worker fallback position is presented in Appendix 2.

We initially endow each worker i with a personal discount rate ρ_i drawn from a uniform distribution on [3.5%, 4.5%]. Each worker is also endowed with a value of a , the disutility of labour, drawn from a uniform distribution on [0.015, 0.025]. The personal discount rate ρ_i and the disutility of labour a_i characterise the worker and remain fixed throughout the run. The range of worker parameters is not critical to the operation of the model. If we raise the average disutility of labour, then average effort level will fall, and if we raise the discount rate, workers will take less desirable jobs, as expected from standard choice theory.

4.2. *Consumers*

A consumer's income is the proceeds from his activity in that period in the labour market, plus the return on the capital he owns. Each consumer has a utility function with various parameters drawn from uniform distributions, so agents are intrinsically

heterogeneous as consumers; i.e., consumers never 'adjust' their consumption function. The specification of the consumption functions is described in Appendix 3.

Hirota (2003) shows that a general equilibrium model having 3 consumers with CES consumption functions over three goods and a Walrasian price adjustment process converges to equilibrium with very high probability. This suggests that there might be something about CES functions that promotes general equilibrium stability. I have not been able to test this idea, since I have not been able to find tractable alternatives to CES. However, I did check to see if our hybrid CES utility functions satisfy the Weak Axiom for Excess Demand Functions, in which case under plausible regularity conditions satisfied by our model, assuming no production, there could be only one equilibrium (Mas-Colell *et al.*, 1995, Proposition 17.F.2, p. 609). My estimates show that for a small number of consumers, the failure rate of the Weak Axiom (searching uniformly over a phase space including the consumption parameters, the distribution of holdings, and prices) is quite high (over 27%) but, for a large number of consumers, the failure rate falls to about 1%. This means that our model, without production, has a unique global steady state with very high probability, and it is unlikely that adding production changes this conclusion, since supply functions generally cannot have the unruly behaviour of demand functions. Of course, this says nothing about stability, which is the focus of this article.

In each period, each consumer chooses one firm in each sector to supply his desired consumption of each good. The choice process is as follows. Each consumer i is endowed with an integer k_i , drawn from the uniform distribution on $\{3, 4, 5, 6\}$, representing the number of firms in a sector whose prices the consumer will inspect before deciding on a supplier. In each period, consumers successively (in random order) inspect suppliers. Consumer i inspects k_i firms in each sector (or all firms, if k_i is larger than the number of firms in the sector, and in each sector chooses the lowest price firm that has a positive amount of the good remaining). Using this price structure, the consumer maximises utility to determine how much of each good to buy. In rare cases, the chosen supplier will not have enough of the product to fulfil the consumer's request, in which case the consumer purchases only the supplier's remaining stock. In rare cases, none of the firms inspected by the consumer will have a positive amount of the good remaining, in which case the consumer does not purchase this good, adding the purchasing power reserved for this good to his next-period income. It is rather more likely that a consumer will be at the tail of the queue and will face relatively high prices. In this case (specifically, when the price of a good is more than 20% higher than in the previous period), the consumer simply saves the money allocated to that good, and lives off the reserves. In all other cases, the consumer purchases the good, the selling firm's inventory is reduced and the seller's profits are increased accordingly.

4.3. *Initial Endowments*

The capital stock of the economy is fixed at a nominal 100 units per sector at the start of the run. Ownership is assigned as follows. Each agent is assigned a draw from the uniform distribution on $[0, 1]$. Each agent is assigned a share of the total capital stock in proportion to this draw.

4.4. *Firms*

If both the firm and all workers shared the same estimate of the fallback v_w , and if the firm knew the workers' utility function and the final demand for its products, the firm could choose the wage w to maximise profits. However, the workers are not homogeneous, the firm does not know the workers' estimates of v_w , and it does not know final demand. We thus assume that firms maximise profits by updating their operating characteristics. Initially, firms are endowed with operating characteristics drawn from uniform distributions, as described in Table 1. Operating characteristics change over time as described below. A firm entering the sector after the initial period starts with the average values of these operating characteristics for the incumbent firms, except with a small probability ($\mu = 0.0005$) the firm is a 'mutant' that chooses a random value for a random subset of operating characteristics.

Firm output per unit of effort q is a function of the number of employees, k , of the firm, and the amount of capital K that the firm hires. Maximum output per unit of effort occurs at some $k = k^*$, and q has an inverted-U shape, symmetric about k^* , with $q = 0$ when $k = 0$. The following is a simple q schedule reflecting these considerations.

$$q(k, K) = \sqrt{1 + 0.1K}k(2k^* - k)/(k^*)^2 \quad (2)$$

As a function of k , this is an inverted U-shaped curve with maximum $\sqrt{1 + 0.1K}$ at k^* . For specificity, we set $k^* = 35$. Thus if firms produced at their productivity-maximising levels, and if all goods had the same demand function (which must be approximately true, since demand functions are randomly generated) there would be, on average $500/35 \approx 14.3$ firms per sector. Of course, it would be easy to specify distinct production functions for different sectors but this would not alter the dynamics of the economy, and the current specification implies that prices in all sectors should be equal in steady state.

The firm's profits in a production period are given by

$$\pi(w, k, K) = pkekq(k, K) - wk - rK, \quad (3)$$

where p is the price of the good, e is the effort level of workers, κ is the fraction of production that is sold, and r is the rental price of capital.

Firms maximise profits by adapting to market conditions. Firms that perform poorly copy the operating characteristics of randomly encountered others that have higher profits. To determine the copying process, we use a replicator dynamic, which implies that a firm that has above-average profits is copied with a probability proportional to its profitability, and a firm that has below-average profits copies another with a probability proportional to its degree of failure.

4.5. *Calculating the Demand for Capital*

If r is the rental price of capital, and if all output is sold, (3) and (2) give

$$\pi(w, k, K) = pek\sqrt{1 + 0.1K}k^2(2k^* - k)/(k^*)^2 - wk - rK.$$

The first order condition for capital demand K is then

$$K(r) = \frac{k^2 p^2 e^2}{40r^2} - 10, \quad (4)$$

assuming $k = k^*$ and $\kappa = 1$. We assume $K(r)$ is the firm's demand for capital, and we choose r to equate the supply and demand for capital in each period.

This model of the capital market leads to considerable macroeconomic stability in comparison with alternatives that assume less firm information concerning the production function. An alternative is to add capital demand to the firm's operating characteristics and have this evolve dynamically in the same manner as other firm characteristics. For instance, we can define the firm's demand for capital as

$$K(r) = \frac{k^2 p^2 e^2}{c_1 r^2} - c_2$$

and permit the parameters c_1 and c_2 to evolve endogenously. Not surprisingly, a model implementing this alternative exhibits more output and profit volatility than the one used in the body of this article.

5. The Agent-based Algorithm

After firms and agents are initialised, as described above, the follow processes occur serially until the end of the run, specified by a total number of period elapsed (3,000 periods in our runs).

- 1 *Production and Consumption*: Each firm produces according to its operating characteristics. Consumers receive their income from labour and capital rental. They search among firms for suppliers, and spend all of their income, if they can.
- 2 *Firms Copy Successful Behaviour*: A fraction (5%) of firms change their operating characteristics by copying the behaviour of other firms. The probability of copying or being copied is proportional to the firm's profitability.
- 3 *Firm Entry and Exit*: When average profits are positive in a sector, the Monetary Authority finances a new firm, which enters the sector. The firm is endowed with operating characteristics that are the average for the sector, except target firm size (preferred number of workers) is adjusted up or down by one unit, and other characteristics are adjusted up or down by a factor depending on the adjustment characteristics of the firm. The new firm then attempts to hire a labour force of target size k^* . The firm randomly samples currently unemployed workers and a worker agrees to join the firm if the worker's utility with the new firm will be higher than that of remaining unemployed. If a firm cannot reach its target labour force in this manner, it then randomly samples employed workers and a worker agrees to switch firms if he will have higher utility in the new firm. Note that if a firm's wage offer is too low, or its demands for effort are too high, the firm may attract few or no workers. On the other hand, the current implementation assumes no labour search costs on the part of firms, so the allocation of labour among firms is likely to be unrealistically efficient.

- 4 *Bankruptcy*: In each period, if average profits are negative in a particular sector, the least profitable 5% of firms are retired, and its workers become unemployed. Note that both firm entry and exit depend on aggregate sectoral conditions, and this is the only behaviour in the model that is not based upon purely local information. We may think of the Monetary Authority having this aggregate information, agreeing to finance a new firm when profits are positive in a sector, and forcing a firm into bankruptcy when aggregate profits are negative.
- 5 *Firms Adjust Their Workforce*: Once every ten periods, each firm adjusts the composition of its workforce as follows. First, if a firm has zero inventory and its workforce is less than its target employment level, the firm attempts to bring its workforce up to its target level and, if it does not succeed, it randomly revises its wage offer up by a firm-specific, endogenously determined, percentage, or it revises its effort demand down by a similarly a firm-specific, endogenously determined, percentage.
- 6 *Workers Adjust Employment Status*: Each worker, once every ten periods, checks that the present value of his current employment, which could have changed because the firm altered its wage rate or effort demand, exceeds his (estimate of the) fallback. If not, he quits and joins the ranks of the unemployed. Then, each employed worker inspects one other firm in the economy at random and, if that firm is hiring, and if the firm is offering a better deal than the current employer, the worker switches to the new firm.
- 7 *Mutation*: In each period, for each firm with probability 0.0005, the firm's operating characteristics are altered as follows. With probability 0.9 the firm's wage is altered by a firm-specific, endogenously determined, percentage, with an increase or decrease equally likely. With probability 0.1 the firm's wage is drawn anew from the distribution used to initialise firms at the start of the run. A similar operation is performed on the firm's price and effort demand, and other non-discrete operating characteristics. Then with probability 1/2 the firm's target size is increased by one unit and, with probability 1/2, it is decreased by one unit, so long as target size is greater than two.
- 8 *Price Adjustment*: Each firm, if it has positive inventory, lowers its price by a firm-specific, endogenously determined, percentage and, if it has zero inventory, it raises its price in a parallel manner. In addition, the firm adjusts its price up or down by a small amount to be more competitive within the sector, depending on whether or not it sold its complete inventory in the previous period.

6. The Walrasian Model: Results

We find that there is a unique pattern of long-term prices and quantities. There is considerable short-term price volatility, but relative prices closely approximate their equilibrium values. In our agent-based model, all firms have the same production technology and quantities are normalised so that all equilibrium relative prices are unity.⁴ Thus, all prices are equal in equilibrium. The actual dispersion of relative prices

⁴ No result reported in this article, except unit relative prices in equilibrium, depends on this choice of technologies.




Fig. 5. *Sectoral Prices*

Since conditions of production are the same in all sectors, relative prices are unity in equilibrium




Fig. 6. *Deviation of Sectoral Prices from Equilibrium Values*

Prices are volatile but the standard deviation of prices is rarely more than 10% of average prices and the mean standard deviation of prices is 5.8% of average prices

is shown in Figure 5, and the standard deviation of prices is displayed in Figure 6. The standard deviation of prices in all sectors in our results averages 5.8% of mean price levels, and is rarely greater than 10%.

We also find that there is some volatility in supply of each good, and excess supply is positive in almost all periods. This is because firms are allowed to carry inventories (which depreciate at the rate 10% per production/consumption period) and profit maximisation in an uncertain environment entails positive inventories. However, the extent of excess supply is small – about 8% of average supply in each sector (see Figure 7).

The steady state profit rate is close to its equilibrium value of zero. Average excess profits, after paying labour and capital rental costs, are usually slightly positive, about 0.9% of total sales on average, as shown in Figure 8. This is offset by the losses of bankrupt firms, which are not included in the Figure. Similarly, the long-run average excess demand for labour by firms is small but consistently positive, amounting to about 1% of the demand for labour (see Figure 9).

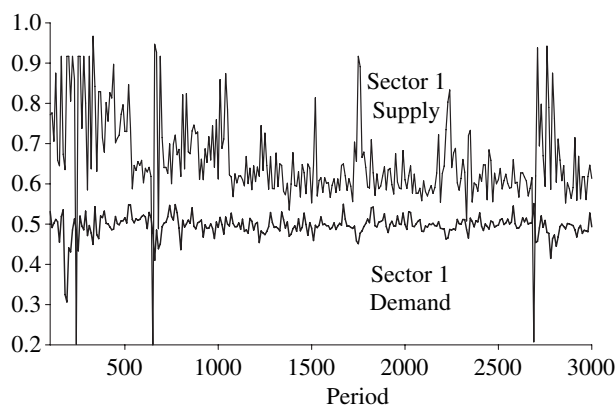


Fig. 7. *Demand and Supply in Sector 1*

Note that there is considerable period-to-period volatility. Excess supply averages about 8% of average supply




Fig. 8. *Excess Profits*

There is a slight tendency for excess profits to be positive, by an amount averaging less than 1% of total income




Fig. 9. *Excess Demand for Labour in Sector 1*

The excess demand for labour by firms is volatile, and almost always positive, with a mean of less than 1% of demand for labour. The demand for labour is close to its equilibrium value of 35 employees per firm (not shown)

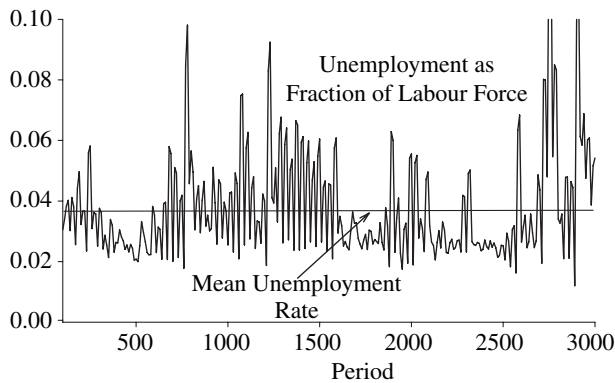


Fig. 10. *Unemployment Rate*

Unemployment averages about 3.8% of the labour force, but exhibits considerable volatility

Out of equilibrium, there can be both an excess demand for labour by firms, as discussed above, and an excess supply of labour, because of the frictions induced by a decentralised process of firms searching for new employees and employees seeking out better employment opportunities in other firms. The long-run behaviour of the unemployment rate is depicted in Figure 10. In our agent-based model, workers accept employment only if the wage/effort pair offered is sufficiently high that the present value of being employed is greater than the present value of continuing job search, which we call the worker's *fallback position*. The fallback position is a personal characteristic of each worker, depending on the agent's subjective discount rate as well as the objective probability of being hired from the unemployment pool and the objective probability of moving from employment to unemployment. The fallback is an expected future value, but it is estimated by workers in the model as an average of labour market behaviour over the past several periods. Unemployment exhibits considerable volatility but averages only about 3.8% of the labour force. Because the economy does not converge strongly to equilibrium values, we have both positive excess demand for labour, shown in Figure 9, and positive unemployment, shown in Figure 10.⁵

Since the model has many agents and sectors, there are no correlated shocks and there is no coordinated decision making, this volatility reflects intrinsic properties of markets where prices are private information out of equilibrium. One might speculate that a larger model more accurately reflecting empirical industrial structures, including a larger number of firms in a sector, might exhibit less volatility. I have not found the latter to be the case using as many as twenty sectors. Reducing optimum firm size does reduce volatility by increasing the average number of firms per sector.⁶

⁵ Note that the cycle of unemployment is about 35 periods. I have not attempted to calibrate this agent-based model to a real economy, but as an approximation, a period in the model represents a month. This is the minimum time in which a firm can enter or leave a market or change prices, it is the time interval in production and payment of suppliers to the firm, and it is the minimum period in which workers become employed or unemployed. With this approximation, the unemployment cycle is about three years, which is close to that of many real market economies.

⁶ All initialisation of firm operating characteristics take the form of draws from uniform distributions. Mutations during the course of the run are not constrained to lie within the limits given by these uniform distributions.

We cannot calculate the equilibria of this model explicitly, for reasons explained above. However, we can estimate the efficiency of the simulated economy as follows. By definition, an efficient equilibrium has all firms operating at the point at which labour efficiency is a maximum, which means $k = k^* = 35$ employees per firm (2). There are 500 consumers per sector, so if consumption is about equal for all sectors, there will be about $500/35 \approx 14.3$ firms per sector. The total capital stock is set to 100 per sector, which is about $100/14.3 \approx 7$ per firm. If workers were homogeneous in the disutility of effort a , from (2) we see that optimal effort e^* maximises

$$e\sqrt{1 + 0.1(7)} - a/(1 - e).$$

In fact, a is uniformly distributed on the interval $[0.015, 0.025]$ but, if we use the average value $a^* = 0.02$, we find $e^* \approx 0.98$. Average worker income is then equal to average worker output, given that prices are unity and workers have equal claims on firm output, which follows from the fact that they are also claimants on the capital returns. Worker income assuming unit prices is thus $y^* = e^*\sqrt{1.7} \approx 1.28$. This allows us to calculate optimal consumption for each consumer. We define the *efficiency* of an agent-based model in a given period as the ratio of actual consumption to optimal consumption, averaged over all consumers. Figure 11 depicts the efficiency of the run, which is greater than 75%.

Finally, Figures 12 and 13 show that consumers on average hold a small amount of money over from period to period, and excess profits are virtually zero. The latter result is due to the fact that the losses of bankrupt firms offset the slightly positive average profits of firms that stay in business.

7. Relaxation Properties: Reaction to a Technology Shock

Under plausible regularity conditions, the equilibria in the Walras-Arrow-Debreu model are locally unique (Debreu, 1970). There is, however, no guarantee of global

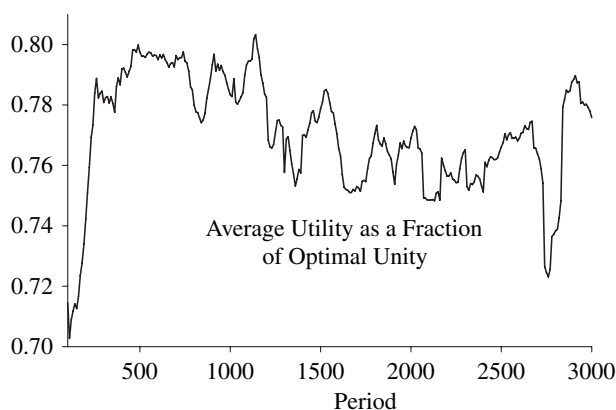


Fig. 11. *The Efficiency of the Simulated Economy*

Note that after only 100 periods, the system has attained nearly 79% efficiency. After 100 periods, average utility appears to be a random walk between 72% and 80% (the apparent declining trend in the Figure is an artifact of truncation to 3,000 periods)




Fig. 12. *Average Monetary (wealth) Holdings of Consumers*

Note that only a fraction of a percentage of national income is held over as savings from period to period




Fig. 13. *Average Excess Profits of Firms*

uniqueness, or even that there are a small number of locally stable equilibria with large basins of attraction. But our experience with real economies gives little support to the notion that there are multiple equilibria. With the hybrid CES consumption functions and no production, however, the likelihood of multiple equilibria is very low – an issue discussed above (see Section 4.2). With production,

heterogeneous capital intensity can lead to ‘reswitching’ and hence multiple equilibria (Garegnani, 1970).

As explained above, an agent-based model is a Markov chain with no absorbing states, so the ergodic theory for Markov chains ensures that our model has a stationary distribution. There is no theoretical tool, however, for ascertaining how many periods of operation are typically required for the effects of the initial conditions to ‘wear off’. Moreover, Markov chains can mirror multiple equilibrium behaviour by spending long periods of time near each of several ‘quasi-steady states’ located in diverse parts of the state space. These time intervals may be so long that a single run will never reveal more than one quasi-steady state. When this is the case, the initial state of the system determines the first quasi-steady state that the system will visit, so successive runs of the agent-based model will exhibit quite different characteristics when multiple quasi-steady states exist.

In dozens of runs with initial conditions within an order of magnitude of their long-run values, I encountered only the steady state described in this article. Even if the system is initialised with parameters that are extremely far from their high-efficiency values, such as prices one hundred times and wages one hundredth of their high-efficiency values, the economy converges to its high-efficiency long-run steady state within one thousand periods.

This analysis suggests that when the simulated economy experiences a system-wide shock of moderate proportions, it should return to its long-run state after a certain number of periods, which we may call the *relaxation time* of the dynamical system. This is in fact the case. For instance, I simulated a four-sector economy with 2,000 agents, and ran the economy for 3,000 periods. After each 500 periods, the firms in the economy were all subjected to a technological shock, taking the form of the optimal firm sized k^* falling from 35 to 14. This shock persisted for 10 periods, after which the original value of k^* was re-established. Figure 14 suggests that the economy recovers its high efficiency price structure after a few hundred rounds. Figure 15 shows that efficiency is severely compromised by the shocks, but is restored within two hundred rounds.




Fig. 14. *Relaxation Time in a Four Sector Economy Subject to Macro Shocks*

Every 500 periods, the economy sustains a shock whereby each firm’s optimal size is reduced from 35 to 14. The shock lasts for 10 periods, after which the original optimal firm size is restored. Note that goods prices stabilise after a few hundred periods, and are approximately equal across all runs




Fig. 15. *Efficiency in a Four Sector Economy Subject to Macro Shocks*
The shock parameter are as in Figure 14

8. Reaction to Permanent Structural Change

As a test of the ability of the simulated economy to adjust to macro-structural changes in its operating parameters, I used the same four-sector model described in Section 7. In this case, I subjected a four-sector economy to a one-time, permanent increase in the labour force from 2,000 to 3,000 agents, the newly-created workers initially assigned to the unemployment pool. Figure 16 shows the effects on the efficiency of the economy. Perhaps surprisingly, neither average prices or their dispersion are affected by the change, but efficiency takes a considerable hit, and is restored only slowly over time, not achieving its original level after 1500 periods. The long-run rental price of capital is higher after the shift, as would be expected, since the labour/capital ratio has increased by 50%.

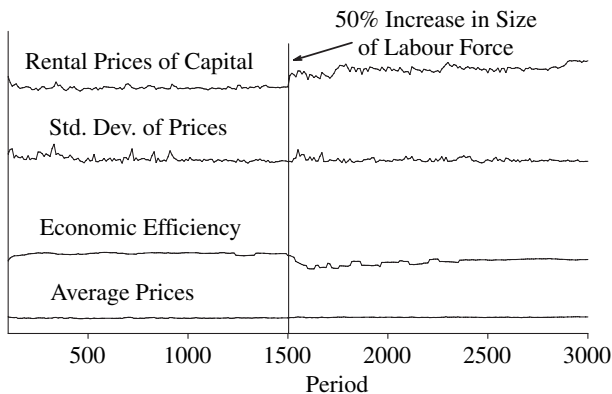


Fig. 16. *Efficiency in a Four Sector Economy Subject to a 50% increase in the Labour Force in Period 1500*
Note that neither average prices nor dispersion are affected, but economic efficiency drops abruptly and is restored only after a long interval, and the rental price of capital increases, reflecting the 50% increase in the capital/labour ratio for the economy




Fig. 17. *Unemployment and Wages after a Large Labour Force Increase in Period 1500*
 Note that the wage rate returns to its long-run value quite slowly (the vertical axis is relevant for the unemployment rate alone). The unemployment rate returns quickly to its long-run average but with increased volatility for more than 1,000 periods

Figure 17 sheds light on the reasons for the decline in economic efficiency. The unemployment rate increases strongly at first but attains low levels after a couple of hundred periods, although for the rest of the run the unemployment rate shows much higher volatility than prior to the shock. The wage rate declines sharply in response to the 50% increase in the supply of labour and its pre-shock value is approached but never re-attained by the end of the 3,000-period run.

Figure 18 shows the effect of the labour force increase in Sector 2, which has been chosen because it is quite well behaved. The number of firms in the sector adjusts rapidly, and demand and supply also adjust rapidly to the new macroeconomic conditions. Sector 1, which is not shown, exhibits considerably more excess supply volatility after the shock than before, with no tendency towards a return to 'normal' by the end of the run. Doubtless it is this volatility that accounts for the somewhat reduced efficiency of the system at the end of the 3,000 periods.

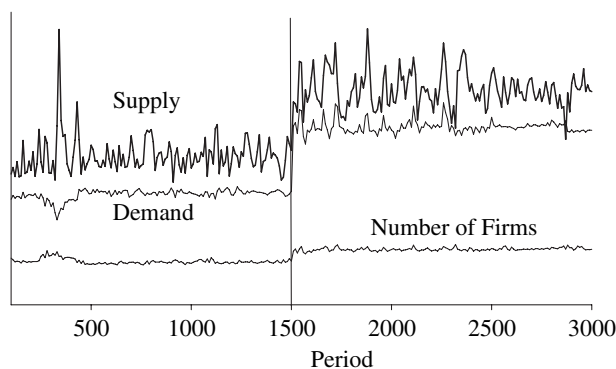


Fig. 18. *Supply, Demand, and Number of Firms in a Sector after a Large Labour Force Increase*
 As expected, the steady state number of firms is higher, as are demand and supply per sector

9. Relation to the Literature

There have been notable analytical contributions to general equilibrium dynamics, including (Arrow and Hurwicz (1958), Arrow *et al.* (1959), Scarf (1960), McKenzie (1960), Gale (1963), Smale (1976) and Foley (1994). Nevertheless, Franklin Fisher's assessment (Fisher, 1983) remains valid: we have no plausible analytical model of multi-sector dynamics with heterogeneous agents. The article presents the first general, highly decentralised, agent-based model of the dynamics of general equilibrium.

There are many agent-based models described in the economics literature (Tesfatsion and Judd, 2006) but none, to my knowledge, deals with several sectors, disaggregated to the firm level, and many heterogeneous agents. The closest to the approach in this article is Chen *et al.*, who use the Swarm software (Stefansson, 1997) to model a one-sector economy using labour, the only endogenous variable being the single product price. Close in spirit is Epstein and Axtell (1997), who model a highly complex multi-agent system but have one good ('sugar') and no firms. Also close in spirit is Lian and Plott (1998), who implement a laboratory experiment with human subjects, with one good, labour and *fiat* money, and Anderson *et al.* (2004) who implement a similar experiment with three goods and three consumers with Leontief consumption functions, based Scarf (1960). Weddepohl (1997) simulates a global tâtonnement process in an economy with two goods plus labour, one firm producing each good, and three consumers who also supply a fixed amount of labour. Other than these papers, I have been able to find only examples of sector-level models with aggregate demand, and equilibrium-computing simulations of highly aggregated models, e.g., Taylor and Uhlig (1990).

10. Conclusion

The major finding of this article is that a plausible dynamics exists in which prices and quantities converge to their market clearing values with a stochastic error term that exhibits moderately large excursions from zero at irregular intervals, even in the absence of global shocks to the system. The economy generally exhibits a high level of efficiency but again with periodic low-efficiency excursions. This behaviour provides some justification for the importance placed upon the Walrasian model in contemporary economic theory. Moreover, for the first time, we know something substantive about the dynamic properties of the Walrasian system. They are nothing like tâtonnement. The knowledge gained should aid in modelling the market economy better as a complex adaptive system, and in finding improved tools for dampening the stochastic behaviour of market economies.

A major attraction of the Walrasian economy is that the only information an individual need have is his personal preferences and endowments, as well as the prices of all goods, and the only information a firm need have is its production function and the prices of its inputs and outputs. This article shows that these assumptions are both too strong and too weak. They are too strong because the dynamic properties of the system are improved if we assume that the economic actors have no public information whatever but rather each agent has a private set of relative prices that he updates through experience. Similarly, firms have private and imperfect knowledge of their supply and demand conditions. However, we have assumed the Monetary Authority will

finance new firms and retire poorly performing firms based on aggregate sectoral profits. Moreover, we have assume that workers change their fallback positions optimally to reflect prevailing labour market conditions.

The general equilibrium assumptions are too weak because they do not take into account that agents can learn from one another's successes and failures. We have seen that agent-based models allowing traders, consumers, workers and firms to imitate successful others leads to an economy with a reasonable level of stability and efficiency.

Agent-based modelling is not an alternative to analytical modelling but rather an empirical investigation of the characteristics of a complex system in a controlled laboratory setting. The findings may serve as a basis for formulating analytical models more accurately reflecting these characteristics. Such modelling may be inspired by physics, where only the tiniest systems are analytically tractable but where statistical mechanics, simulated annealing, percolation theory and other powerful techniques, are deployed to specify the macro behaviour of a system with many degrees of freedom (Albin and Foley, 1992; Foley, 1994).

The desirability of using public prices in modelling the market economy does not survive our analysis. Economic theory assumes public prices without justification. Public prices do not generally exist and equilibrium public prices cannot even be calculated in an economy of any appreciable size (I assumed all sectors in my agent-based model have the same production function, so equilibrium prices are unity.)

This article suggests the fruitfulness of additional work in the area of agent-based modelling of the market economy. However, the model presented in this article has many limitations. There is no inter-industry trade and there is only one financial asset. Consumers do no life-cycle saving and labour is homogeneous. There is no retail trade. Consumers all have hybrid CES consumption functions. The only source of asset heterogeneity is agent capital holdings, which are fixed throughout the run. There are no 'key' goods that form a large part of aggregate consumption or enter in a consistent manner into the production of most or all goods. What we call 'firms' should really be called 'plants'. More realistic firms should have a capital asset structure, a set of managers, and the ability to acquire and divest itself of many 'plants'. All contracts are assumed complete and costlessly enforced by a third party, despite the lack of realism of this assumption (Bowles and Gintis, 1993; Gintis, 2002).

Appendix

1. *The Method of Agent-based Modelling*

Figure 19 shows the programming structure of a typical agent-based model. In the Figure, 'Game Parameters' refer to the specifics of the stage game being modelled, including the payoffs and the probabilities with which various events occur. The 'Number of Generations' specifies how many periods of reproduction will take place. This may be as small as 10 or as large as 10,000,000. The 'Number of Periods/Generation' refers to the speed of play as compared to the speed of reproduction.

The 'Agents Reproduce' box is expanded in Figure 20. First we set various parameters, including the rate of mutation of new agents and the death rate of old agents. We then replace the appropriate number of unsuccessful agents with (possibly mutated) copies of the high success agents.




Fig. 19. Structure of an Evolutionary Agent-based Model

2. Calculation of Worker Fallback Position

An unemployed worker should accept employment only if the present value of being employed at the wage and effort condition of the firm offering a position exceed the present value of continuing job search. We call the latter the worker's *fallback position*. Each worker estimates his fallback as follows. Let p_u be the fraction of unemployed workers who found employment in the previous period and let p_e be the fraction of employed workers in the last period who remained employed in the current period. Let (w, e) be the expected wage and expected effort requirement from any job offer. Then if the worker's discount rate is ρ , so his discount factor is $\delta = 1/(1 + \rho)$, then the present value v_u of being unemployed is given by

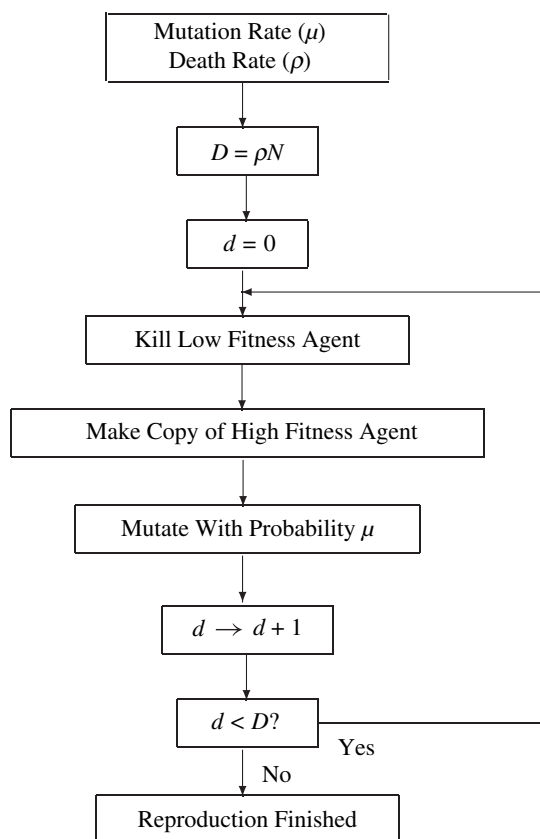


Fig. 20. Structure of Reproduction Process

$$v_u = \delta[p_u v_e + (1 - p_u)v_u],$$

where v_e is the present value of being employed. Similarly,

$$v_e = w - a/(1 - e) + \delta[p_e v_e + (1 - p_e)v_u].$$

Solving simultaneously, we have

$$v_u = \frac{p_u \delta [w - a/(1 - e)]}{(1 - \delta)[1 + \delta(p_u - p_e)]}. \quad (5)$$

In our agent-based model, we assume that the worker solves this equation to find v_u , where the values of the parameters are the average values over the previous five periods. A fraction of new entrants firms also solve this equation in deciding their wage offer.

This method of determining the fallback rather violates the spirit of agent-based modelling, because it assumes agents know certain macroeconomic variables, and are capable of solving complex sets of equations. I have experimented with alternative fallback determination algorithms. The closest to the above is to have each worker sample the market rather than use the population averages. The effect is to increase the long-run unemployment rate and introduce autocorrelation in to the unemployment rate (a labour-generated business cycle). More radical is to take the fallback as a personal characteristic of each worker, and use a replicator dynamic to

govern the evolution of the fallback, unsuccessful workers adopting the fallback of more successful workers. This algorithm does not change the qualitative performance of the agent-based model, but it introduces even more autocorrelation and higher mean unemployment levels (on the order of 6%).

3. Consumption Functions

The utility function of each agent is the product of powers of CES utility functions of the following form. For each consumer, we partition the n consumer goods into k segments (k is chosen randomly from $1 \dots n/2$) of randomly chosen sizes m_1, \dots, m_k , with $m_j > 1$ for all j , and $\sum_j m_j = n$. We randomly assign goods to the various segments, and for each segment, we generate a CES consumption with random weights and an elasticity randomly drawn from the uniform distribution on $[0.3, 2]$. Total utility is the product of the k CES utility functions to random powers f_j such that $\sum_j f_j = 1$. In effect, no two consumers have the same utility function.

For example, consider a segment using goods x_1, \dots, x_m with prices p_1, \dots, p_m and (constant) elasticity of substitution s , and suppose the power of this segment in the overall utility function is f . It is straightforward to show that the agent spends a fraction f of his income M on goods in this segment, whatever prices he faces. The utility function associated with this segment is then

$$u(x_1, \dots, x_n) = \left(\sum_{l=1}^m \alpha_l x_l^\gamma \right)^{1/\gamma}, \quad (6)$$

where $\gamma = (s-1)/s$, and $\alpha_1, \dots, \alpha_m > 0$ satisfy $\sum \alpha_l = 1$. The income constraint is $\sum_{l=1}^m p_l x_l = f_i M$. Solving the resulting first order conditions for utility maximisation, and assuming $\gamma \neq 0$ (i.e., the utility function segment is not Cobb-Douglas), this gives

$$x_i = \frac{M f_i}{\sum_{l=1}^m p_l \phi_{il}^{1/(1-\gamma)}}, \quad (7)$$

where

$$\phi_{il} = \frac{p_i \alpha_l}{p_l \alpha_i} \quad \text{for } i, l = 1, \dots, m.$$

When $\gamma = 0$ (which occurs with almost zero probability), we have a Cobb-Douglas utility function with exponents α_b , so the solution becomes

$$x_i = \frac{M f_i \alpha_i}{p_i}. \quad (8)$$

Submitted: 20 August 2005

Accepted: 22 September 2006

References

- Albin, P. and Foley, D. K. (1992). 'Decentralized, dispersed exchange without an auctioneer: a simulation study', *Journal of Economic Behavior and Organization*, vol. 18, pp. 27–51.
- Allington, N. F. B., Kattuman, P. A. and Waldmann, F. A. (2004). 'One market, one money, one price? Price dispersion in the European Union', mimeo Corpus Christi College, University of Cambridge.
- Anderson, C. M., Plott, C. L., Shimomura, K.-I., and Granat, S. (2004). 'Global instability in experimental general equilibrium: the Scarf example', *Journal of Economic Theory*, vol. 115, pp. 209–249.
- Arrow, K. and Hurwicz, L. (1958). 'On the stability of the competitive equilibrium, I', *Econometrica*, vol. 26, pp. 522–52.

- Arrow, K., Block, H. D. and Hurwicz, L. (1959). 'On the stability of the competitive equilibrium, II', *Econometrica*, vol. 27, pp. 82–109.
- Arrow, K. J. (1951) 'An extension of the basic theorems of classical welfare economics', in (J. Neyman, ed.) *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pp. 507–32, Berkeley: University of California Press.
- Arrow, K. J. and Debreu, G. (1954). 'Existence of an equilibrium for a competitive economy', *Econometrica*, vol. 22(3), 265–90.
- Bala, V. and Majumdar, M. (1992). 'Chaotic Tatonnement', *Economic Theory*, vol. 2, pp. 437–45.
- Bandura, A. (1977). *Social Learning Theory*, Englewood Cliffs, NJ: Prentice Hall.
- Bowles, S. and Gintis, H. (1993). 'The revenge of homo economicus: contested exchange and the revival of political economy', *Journal of Economic Perspectives*, vol. 7(1) (Winter), pp. 83–102.
- Chen, Y., Su, X., Rao, J. and Xiong, H. (2000). 'Agent-based microsimulation of economy from a Complexity Perspective', mimeo, Renmin University.
- Conlisk, J. (1988). 'Optimization cost', *Journal of Economic Behavior and Organization*, vol. 9, pp. 213–28.
- Debreu, G. (1952). 'A social equilibrium existence theorem', *Proceedings of the National Academy of Sciences*, vol. 38, pp. 886–93.
- Debreu, G. (1970). 'Economies with a finite set of equilibria', *Econometrica*, vol. 38(3) (May), pp. 387–92.
- Debreu, G. (1974). 'Excess demand function', *Journal of Mathematical Economics*, vol. 1, pp. 15–23.
- Epstein, J. and Axtell, R. (1997). *Sugarscape*, Cambridge, MA: MIT Press.
- Feller, W. (1950). *An Introduction to Probability Theory and Its Applications*, Vol. 1, New York: Wiley.
- Fisher, F. M. (1983). *Disequilibrium Foundations of Equilibrium Economics*, Cambridge: Cambridge University Press.
- Foley, D. (1994). 'A statistical equilibrium theory of markets', *Journal of Economic Theory*, vol. 62 (2), (April), pp. 321–45.
- Freidlin, M. I. and Wentzell, A. D. (1984). *Random Perturbations of Dynamical Systems*, New York: Springer.
- Gale, D. (1963). 'A note on global instability of competitive equilibrium', *Naval Research Log Quarterly*, vol. 10, pp. 81–7.
- Garegnani, P. (1970). 'Heterogeneous capital, the production function and the theory of distribution', *Review of Economic Studies*, vol. 37, pp. 407–36.
- Gintis, H. (2002) 'Some implications of endogenous contract enforcement for general equilibrium theory', in (F. Petri and F. Hahn eds.) *General Equilibrium: Problems and Prospects*, pp. 176–205. London: Routledge.
- Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*, Reading, MA: Addison-Wesley.
- Hirota, M. (1981). 'On the stability of competitive equilibrium and the patterns of initial holdings: an example', *International Economic Review*, vol. 22, pp. 461–67.
- Hirota, M. (1985). 'Global stability in a class of markets with three commodities and three consumers', *Journal of Economic Theory*, vol. 36, pp. 186–92.
- Hirota, M. (2003). 'On the probability of the competitive equilibrium being globally stable: the C.E.S. example', Caltech Social Sciences Working Paper 1129.
- Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems*, Ann Arbor: University of Michigan Press.
- Kirman, A. P. and Koch, K. J. (1986). 'Market excess demand in exchange economies with identical preferences and collinear endowments', *Review of Economic Studies*, vol. 53, pp. 457–63.
- Lian, P. and Plott, C. R. (1998). 'General equilibrium, markets, macroeconomics and money in a laboratory experimental environment', *Economic Theory*, vol. 12, pp. 21–75.
- Mas-Colell, A., Whinston, M. D. and Green, J. R. (1995). *Microeconomic Theory*, New York: Oxford University Press.
- McKenzie, L. W. (1960). 'Stability of equilibrium and value of positive excess demand', *Econometrica*, vol. 28, pp. 606–17.
- Saari, D. G. (1985). 'Iterative price mechanisms', *Econometrica*, vol. 53, pp. 1117–31.
- Scarf, H. E. (1960). 'Some examples of global instability of competitive equilibrium', *International Economic Review*, vol. 1, pp. 157–72.
- Smale, S. (1976). 'Dynamics in general equilibrium theory', *American Economic Review*, vol. 66, pp. 288–94.
- Sonnenschein, H. (1972). 'Market excess demand functions', *Econometrica*, vol. 40, pp. 549–63.
- Sonnenschein, H. (1973). 'Do Walras' identity and continuity characterize the class of community excess demand functions?', *Journal of Economic Theory*, vol. 6, pp. 345–54.
- Stefansson, B. (1997). *Swarm: An Object Oriented Simulation Platform Applied to Markets and Organizations*, New York: Springer-Verlag.
- Taylor, J. and Uhlig, H. (1990). 'Solving nonlinear stochastic growth models: a comparison of alternative methods', *Journal of Business and Economic Statistics*, vol. 8, pp. 1–19.
- Taylor, P. and Jonker, L. (1978). 'Evolutionarily stable strategies and game dynamics', *Mathematical Biosciences*, vol. 40, pp. 145–56.

- Tesfatsion, L. and Judd, K. L. (2006). *Handbook of Computational Economics II: Agent-Based Computational Economics*, Amsterdam: North-Holland.
- Tomasello, M. (1999). *The Cultural Origins of Human Cognition*, Cambridge, MA: Harvard University Press.
- Walras, L. (1954). *Elements of Pure Economics*, London: George Allen and Unwin, [1874].
- Weddepohl, C. (1997). 'Simulating tatonnement in a production economy', *Economic Theory*, vol. 10, pp. 551–8.