DEPARTMENT OF ECONOMICS

Working Paper

A Multi-Agent Systems Approach to Microeconomic Foundations of Macro

by

Bill Gibson

Working Paper 2007-10

UNIVERSITY OF MASSACHUSETTS
AMHERST

A MULTI-AGENT SYSTEMS APPROACH TO MICROECONOMIC FOUNDATIONS OF MACRO

BILL GIBSON†

Abstract. This paper is part of a broader project that attempts to generate microfoundations for macroeconomics as an emergent property of complex systems. The multi-agents systems approach is seen to produce realistic macro properties from a primitive set of agents that search for satisfactory activities, "jobs", in an informationally constrained, computationally noisy environment. There is frictional and structural unemployment, inflation, excess capacity, financial instability along with the possibility of relatively smooth expansion. There is no Phillips curve but an inegalitarian distribution of income emerges as fundamental property of the system.

1. Introduction

This paper is part of a broader project that attempts to generate a microeconomic basis for macroeconomics as an emergent property of complex systems[1]. It is illustrated with a simple agent-based model written in NetLogo, a Java-based code for multi-agent system modeling[2]. Agents are buyers and sellers of labor services, embedded in a landscape or grid, and negotiate deals that may or may not result in production[3]. Natural selection is at work for buyers in that if they raise their output prices, they increase the probability that they can be eliminated through competition. Sellers must also respect a survival constraint on total wealth. The model is not comprehensive and does not capture the full range of market activities in either the standard Walrasian or Keynesian economic frameworks. It does not, for example, include the formation of coalitions of agents in firms or employee organizations. Instead, the paper approaches the problem of microfoundations for macroeconomics from a more primitive perspective, that of agents searching for satisfactory activities, "jobs", in an informationally constrained, computationally

---

†Version: 1.02. September 2007; John Converse Professor of Economics, University of Vermont, Burlington, VT 045405 and Professor of Public Policy and Economics, University of Massachusetts, Amherst, MA 01003 USA 413-548-9448. wgibson@econs.umass.edu; http://people.umass.edu/wgibson.

Key words and phrases. Multi-agent system, agent-based models, microeconomic foundations, macroeconomics. JEL codes D58, D83, D30.

Thanks to Diane Flaherty for comments on an earlier version of this paper.

[1] See for example (Brock and Durlauf, 2005) and (Gatti et al., 2008).

[2] See (Railsback et al., 2007) for a review of various platforms available for Multi-Agent System modeling. This paper does not compare, however, ODML, the system designed by Lessor *et al.*, at the Multi-Agent Systems Laboratory at the University of Massachusetts at Amherst. NetLogo is from (Wilensky, 1998). The model of this paper is similar to the Sugarscape model of (Epstein and Axtell, 1996).

[3] See (Axelrod and Benett, 1997) for a discussion of landscape models.

noisy environment. The properties that emerge from the model are many of the standard macroeconomic features one would expect. There is frictional and structural unemployment, inflation, excess capacity and financial instability. It is further seen that an inegalitarian distribution of wealth is deeply embedded in even this primitive economy, since agents who find good jobs early and keep them are consistently wealthier than those who engage in costly search. This result seems to be independent of the initial distribution of wealth across agents.

The paper is organized as follows: Section 2 introduces the multi-agent systems approach, relating it to some of the broad themes of economics and science generally. Section 3 discusses the layout of the model, its adjustment mechanisms and the simulation procedure. The fourth section provides some numerical aggregates that are the model's macroeconomic properties. A concluding section summarizes the main findings of the paper[4].

2. The multi-agent systems approach

Multi-agent systems, also known as agent-based models, address one of the most fundamental problems in social science: the relationship between structure and individual agency (Axelrod, 1984). The long-run character of an economy is determined principally by agency, but in the short run, decisions made in the past present themselves in the form of structural constraints (Durlaf and Young, 2001). In economics, the multi-agent problem was originally framed in the language of game theory and many useful analytical results have been derived. But multi-agent analytical models are limited in institutional specificity and typically fail to capture combinatorily rich environments that even simple multi-agent systems can easily produce[5].

In science generally the approach has its roots in the late nineteenth century statistical mechanics of Gibbs, Boltzmann and Maxwell (Durlauf, 1999). Physics sought to explain macro properties of gases, temperature and pressure already described by the perfect gas law, in terms of the trajectories of individual molecules[6]. This statistical approach proved to be fundamental in the evolution of the quantum theory at the beginning of the next century, with the development of the idea of a probability wave to describe the position of an electron and beta decay of a neutron into a proton and neutrino. With the development of statistical thermodynamics, the macro approach to the theoretical behavior of molecular aggregates was all but abandoned.

---

[4] The NetLogo code is not included in this version of the paper, but is available from the author upon request.

[5] Stated formally, the question is: what is the joint probability distribution for the entire stochastic path that is compatible with the conditional probability distributions for each agent? Since the answer to this question is largely beyond conventional inductive and deductive approaches, researchers have turned to simulation methods, or what Axelrod calls a "third way of doing science" (Aexlrod, 1997), (Berry et al., 2002).

[6] Maxwell famously argued, by implication, that an individual atom could have no temperature. It was not a question of measurement, but rather derivative of the notion of heat as the energy released in collisions of individual atomic masses. Since one atom, or agent, could not collide with itself, it was meaningless to ask how hot it was. Writing to Wolfgang Pauli, Maxwell went further to argue that like temperature and pressure, time and space "are really only statistical concepts." (Johnson, 1995, p. 1).
See also (Durlauf, 1999) for a discussion of the relationship between statistical mechanics and social science.

More than a century later, economics is grappling with its own version of the unification of micro and macro principles. Keynesian macroeconomic theory, under sustained attack since the 1970s both on empirical and theoretical grounds, has also been largely abandoned. Much of what is considered to be the reason for the decline of Keynesian macroeconomics is the unstable relationship between employment and inflation that was supposedly exploitable along the Phillips curve

> ...the inflationary bias *on average* of monetary and fiscal policy [in the 1970s] should...have produced the lowest average unemployment rates for any decade since the 1940s. In fact, as we know, they produced the highest unemployment since the 1930s. This was economic failure on a grand scale (Lucas and Sargent, 1978, p 277).

The theoretical foundations of Keynesian theory were also weak. Sonnenschein, Mantel and Debreu convincingly demonstrated that the n-agent Walrasian model of general economic equilibrium could not provide microfoundations for one-sector macroeconomics[7]. Analysts seem to be left with the infelicitous choice between a workable but untethered macroeconomics and a coherent, but necessarily unrealistic and for the most part barren microeconomic alternative.

Since the early 1990s, the multi-agent systems or agent-based modeling literature has provided an alternative modeling framework for studying economies and forms of social organization. Multi-agent models are typically based on cellular automata in which agents interact with their immediate neighbors, defined by common edges or corners, and may use sophisticated learning algorithms derived from research in artificial intelligence (AI).

Wooldridge notes that multi-agent systems (1) are capable of suggesting results that can later be established analytically or experimentally; (2) can allow researchers to study events that are either dangerous, expensive or unethical to reproduce; (3) do not involve the loss of information from the use of "representative agents," the behavior of which averages important, underlying heterogeneity; (4) allow the study of problems that are beyond the grasp of either analytical, statistical or qualitative analysis (Wooldridge, 2002).

The clearest example of how these benefits present is also one of the earliest, due to Schelling (Schelling, 1971). Shelling's original neighborhood model envisioned a checkerboard, in which interior agents are surrounded by eight neighbors. Each agent's decision rule is simple; when a threshold of racial density is surpassed, an agent relocates to some randomly chosen location. This simple rule can lead this system to regular behavior. Although each agent prefers to live in a mixed neighborhood, after several iterations the neighborhoods are entirely segregated. This is an *emergent property* of the model, as when ice forms from the interaction of hydrogen and oxygen at zero degrees Celsius, "the product of the complex interaction of elementary particles that could not have been predicted on the basis of the properties of the particles alone" (Anderson, 1972).

Simple reactive systems were only the first step in the development of vastly more complex systems in which agents could compute decision trees, through formal Markov decision processes or informal heuristics, communicate and negotiate

---

[7] See (Debreu, 1974). For an interpretation of SMD theory, see (Rizvi, 1997) and (Rizvi, 1994).

among themselves to form coalitions, teams and blocs, formulate and test hypotheses about their environment and other agents' behavior, adapt to dynamically evolving environments, self-organize, develop error tolerance and resistance to attack[8]. In short, all the AI features that currently drive robotics, large-scale and complex organizational tasks, such as the air-traffic control system and distributed sensor and information retrieval networks can be built into multi-agent systems.

Generally, data structures representing individual agents are instantiated and then allowed to pursue their own objectives interactively, with incomplete information about the environment. System-wide regularities then emerge, despite the absence of governing equations for the social system as a whole (Axtell, 1999).

2.1. Structure and assumptions of the model.

In their survey of multi-agent organizational paradigms, Horling and Lesser note that "an organization of a multi-agent system is the collection of roles, relationships, and authority structures which govern its behavior" (Horling and Lesser, 2005). A market is one of the organizational structures considered in that paper, in which buyers and sellers agree to a binding exchange contract[9]. Classically, markets consist of real-time entities, individuals or groups, which coordinate the activities of participating groups of agents. More realistically, a market is simply a set of bilateral negotiations between *two* agents that may or may not be conclusive. If not, a new round of negotiations with the same, or indeed different, participants can follow *ad infinitum.* Markets are not, therefore, primitives in multi-agent systems but can arise as a form of multi-agent organization as Horling and Lesser have noted and others have worked out in detail (Cheng and Wellman, 1998). The specification of the negotiation process, how it begins and concludes, is logically prior to the particular organizational form that evolves from the interaction of agents. Similarly, firms need not be primitives. A production process operated by agents may be a more natural way to conceptualize production. Firms are then statistical aggregations of production processes that obtain as a result of decision-making at a more fundamental level, that of the agents themselves[10].

Like the definition of a firm, the traditional categories of classical political economy are not fundamental to multi-agent systems. Agents are not identified as "workers" or "capitalists," that is, with fixed social roles. The other extreme is the Walrasian "atomistic agents" who make decisions in isolated environments. Most of the multi-agent literature drops both of these extreme assumptions, however. Abdulla and Lesser, for example, show how agents can learn *through* runtime communication to form effective dynamic coalitions through self-organization (Abdallah and Lesser, 2007). Even Cheng and Wellman, who make every effort to replicate the Walrasian in their model have trading out of equilibrium undertaken by asynchronous agents[11] (Cheng and Wellman, 1998).

---

[8] The literature not already cited is immense. A good introduction is (Wooldridge, 2002) and the references cited therein. See also (Epstein and Axtell, 1996) for a good introduction.

[9] Markets emerge naturally in agent-based models and much of the early work in the field was devoted to how they come about. See (Arthur et al., 1997) and (Cheng and Wellman, 1998)

[10] For a thorough discussion of an agent-based model in which firms are endogenously generated, see (Axtell, 1999). There firms play an important role by preventing agents from defecting, thereby allowing the firm to move to Pareto superior allocation.

[11] Most multi-agent models celebrate the heterogeneity of their agents, as does this one. See in particular (Gatti et al., 2008).

Agents in multi-agent systems are best thought of as *computational entities* who make decisions based in an informationally constrained environment and with limited computational means in real time. In their review of the relationship between AI and economics, Boutilier, Shoham and Wellman note

> Although modeling computational entities as rational beings is standard AI practice we do not generally have the luxury of assuming rationality. It is our burden to explain how to realize approximately rational behaviors in operational computational terms (Boutilier et al., 1997, p 2).

Since computation itself requires real time, agents must cease their computational effort within an action frame of the model.

Many, if not most, real-world problems are computationally hard with no guarantee that an optimal solution can be found[12]. Sandholm and Lessor note that the problem of vehicle routing from only five dispatch centers becomes so large that an agent's rationality is limited by computational complexity[13]. Heuristic, satisficing algorithms replace optimal ones in multi-agent systems and it is the combined activity of imprecise calculations that define *bounded rationality.* In most multi-agent models sub-game perfect strategies, as are present in analytical games, are beyond the reach of agents.

2.2. **Model specifics.** Following Lesser and Horling, agents are first categorized according to their *role.* In the model of this paper, two roles, buyer and seller of labor services, are defined. Agents move from one location on a fixed grid. The role of the agent is to sell her labor services to a cell and the role of a cell is to combine labor and capital to produce output. Thus, cells are also agents, but to prevent terminological confusion, we restrict the use of the term "agent" to apply only to the mobile subset of agents. Cells in the grid will be referred to simply as cells and not agents.

The key decision is whether to produce. Cells and agents communicate information bilaterally in an effort to decide whether a deal can be struck that would allow output to be produced and agents to be compensated. The most important component of the cell-agent relationship is simply their collocation. Cells cannot produce without an agent present and agents cannot produce without cells. Moreover, to have the authority to produce cells must obtain financing from other agents. The aggregate supply of credit is determined by the sum of agent wealth. Agents as a whole thus own the cells as a whole but since there are many more cells than agents, cells compete for agents within this authority structure by which they are defined.

---

[12] Only the simplest games, such as tick-tack-toe and more recently checkers with its decision space on the order of magnitude of 5^{20} can be solved computationally and these only when no strategic error is made as the game is played in real time. See (Schaeffer et al., 2007).

[13] See (Sandholm and Lesser, 1997). Thus, the problem is $N-P$ complete, meaning that it cannot be solved by any known algorithm by a deterministic (i.e., normal) computer using resources (time and memory space) that are linked by way of a polynomial expression to a measure of the size of the input-data (here the five sensors). The traveling salesman problem is another example of an $N-P$ complete problem. There is no guarantee that we might not just "search forever" for the solution.

Mobile agents search for jobs that satisfy their consumption demands, the latter of which are randomly assigned. There is no traditional consumption function linking demand to income or wealth. Instead, agents adjust their consumption/wealth balance through their choice of economic activity. Some agents become wealthy by saving large fractions of their income. Other agents live beyond their means, spending down inherited or previously accumulated wealth. The latter move from job to job searching for a position that will support their lifestyles. This cannot go on forever; if wealth goes negative, agents leave the grid (die or emigrate). When an insolvent agent departs she must be replaced by another. A new entrant into the labor market comes with a fraction of some other agent's wealth. Consumption demand is also inherited from that "parent" and scaled to the level of inherited wealth received by the offspring. Over time, consumption for the society as whole is altered through a Darwinian selection process to become more consistent with the level and distribution of wealth. The model also allows for net population growth at an exogenously specified rate.

Agents' characteristics include job satisfaction, a variable that measures whether current consumption expenses are covered by wage income. Agents wish to accumulate wealth for retirement and to endow their offspring; hence, any job that does not cover current consumption expenses interrupts the accumulation of wealth and is judged unsatisfactory. An agent whose job satisfaction is zero will, most likely, move to search for a better job, one that will cover current expenses so that wealth accumulation can proceed. Moving to another job is risky however and must be supported by a minimum ratio of wealth to current expenses. Agents can therefore get "stuck" in unsatisfactory jobs because they cannot afford to search for better ones.

Cells are best conceived as *blueprints for technologies* that can be activated by the presence of an agent. A cell blueprint specifies the amount of capital to be used in conjunction with the present agent and gives the details of the output, marginal product of capital and labor. Since only one agent can produce on one cell at a time, the amount of capital called for by the blueprint is also the capital-labor ratio for that cell. Price, wage and profit are determined by way of cell-agent communication, following a Nash-bargaining procedure, as discussed in detail below. To render the problem more precisely, we say that blueprints for technologies are initially distributed over the cells in a random fashion and are characterized by

$$q_i = q_i(k_i, \beta_i)$$

with cell index $i = 1, 2, ..., n$. Here n is the default setting of the software, 2601. The q function denotes the technology, k_i the capital stock and β_i a share parameter that controls marginal productivity. Thus, q and k are per-capita levels of output and capital respectively. Consistent with the quantity of capital assigned randomly to the cell is a *max-wage* rate, $\bar{w}$. This is a wage that is equal to the marginal product of labor

$$\frac{\bar{w}_i}{p_i} = \frac{dq}{dL} = (1 - \beta)q \tag{2.1}$$

for a Cobb-Douglas production function. The initial price is set at unity, but changes as described below. The wage offer by cell is initially set randomly according to

$$w_i \leq \bar{w}_i$$

The capital market is endogenous to the model. The supply of capital is simply the sum of real wealth owned by agents as the result of previous production and consumption decisions. The demand for capital is the sum of capital specified by all the blueprints of cells with agents present. Those cells that are the most creditworthy obtain financing and produce; those that cannot obtain financing are eliminated through a process of natural selection, *outsourced* for short. Agents whose jobs have been outsourced will move from that cell to another in the next action frame of the simulation.

If an agent is present on a cell and there is sufficient finance to enable production, cells do not have the authority to refuse to produce. Some cells have very low levels of capital, however, and the production processes they support may pay a wage that does not meet the agent's reservation wage, defined as the wage that covers their expenses. Since it is not within the cell's authority to shut down, it will agree to the agent's wage demand and raise its price so that the product wage equals the marginal product specified in its blueprint. In effect, cells and agents play an *ultimatum game* imposed by the runtime clock (Güth et al., 1982). Cells face a payoff of zero if they do not produce, so if the maximum wage offer, $\bar{w}$, is less than the agent's reservation wage, the cell meets the wage demand but raises its price until equation 2.1 holds. System-wide inflation results and the average real wage for all agents then falls. In the following period, all agents revise their reservation wages to take into account the higher aggregate price level.

Wealth maximizing agents are always better off to accept any wage since their option is zero income for that period and a decline in their wealth by the value of consumption expenses. Experimental results in ultimatum games show that not all agents behave according to the principles of classical rationality however (Henrich et al., 2004). Thus, it is within the authority of an agent to refuse a valid wage offer. Some agents, referred to here as "militants", punish the cell for what the agent deems as an "unfair" division of output. The agent does not know its marginal product and therefore cannot determine what the cell can afford to pay. From the punisher's perspective, the $\bar{w}$ agreement is only marginally acceptable and some fraction will reject the offer and move on. The cell then forfeits its financing and does not produce[14]. Note that it may well be that militant behavior leads to further search and possibly a better job that would not have otherwise obtained.

It is a precarious strategy for a cell to raise its price from the perspective of its own survival. Aggregate wealth is a constraint on total capital formation and the algorithm that allocates finance takes into account profitability in additional to a random error term. A cell that pays the full marginal product of labor thus risks cell death since it will be the least creditworthy. If the minimum wage demand increases because of induced inflation or life-cycle factors discussed below, the cell must then attract net investment (gross investment less depreciation). If the cell is successful, its marginal product of labor will rise with the new capital and it will be able to meet the higher wage demand. If not, the agent will likely move to take a better job offer elsewhere.

---

[14] Militants serve a progressive social purpose, however, preventing technologically weak cells from using scarce capital. Unless capital is abundant, cells whose blueprints do not allow a wage that would be acceptable to most agents, become prime candidates for outsourcing. Militants serve an additional social purpose in keeping inflation down since firms do not then raise their prices to accommodate a wage demand. On the other hand, had the militant accepted the job, total wealth would have increased and more capital would have been available for the next period.

2.3. Education, mobility and collisions.

Agents move in this model for five reasons: (1) the agent is dissatisfied with the job; (2) the job has been outsourced; (3) a better offer arises; (4) the agent is not dissatisfied, but wishes to experiment; (5) the agent quits or completes school. This section deals with the third motivation.

At any point in time, an agent in any one cell is immediately surrounded by at most eight other technologies, four of which are on cardinal headings and four that are not. Depending on the initial parameters, good jobs, those that offer a wage greater than current consumption expenses, may be abundant or not. If there are not many good jobs, some of the eight cells may well be empty.

Assume for the moment that an agent can only see four of the technologies that immediately surround her, those on the cardinal headings. Thus, an agent might be close to a good job, but not be able to "see" it. The agent decides whether to move as a result of computing the highest wage of these four cells it can see. If the highest wage is greater than the current wage, and the agent's wealth to expense is above an exogenously set minimum, then the probability of movement increases. There is also the possibility that the agent will not match the skill or educational requirements for the job. The agent does not know what the employer "really wants" in an applicant. If a degree is required, for example, then an under-educated agent will be turned away and the cell will be unable to produce.

An agent's ability to see jobs ahead is a function of educational. At birth, agents are assigned a random variable for their level of education. After four years of education, agents graduate and can compete for higher paying jobs. With two years of education, agents can see into the second ring of cells around their position and with three jobs three rings and so on[15]. Graduate status confers an additional benefit on an agent that enables her to see jobs on a randomized non-cardinal search heading.

Since education enhances the ability of an agent to find a good job, agents may elect to return to school for more education if their current job is unsatisfactory. The choice involves an opportunity cost of foregone wages as well as a direct school cost that is set parametrically as a fraction of agent's wealth[16]. Only those agents who can afford to return to school do so. An agent must meet an exogenously set minimum ratio of wealth to expenses before returning to school and this ratio is a function of the years remaining to graduation. If an agent's rate of return falls while in school, it is possible that she would have to leave[17].

---

[15] As the distance to a job becomes greater, the cardinal heading restriction reduces the probability of seeing a job significantly. Four of the eight immediate cells are visible, but of the sixteen cells surrounding the immediate eight, only four of those are visible and thus the probability drops from fifty-fifty for the first eight surrounding cells to 0.25 for the second ring and so on.

[16] Note that there is a social cost of education in the model that might not be immediately apparent. Since the sum of agent wealth determines the amount of capital that can be financed, a decision by an agent to return to school causes a reduction in both aggregate wealth and capital stock. The capital stock is immediately reduced since the cell cannot produce when the agent is in school. Thus if the percentage of wealth of the agent that goes to finance her education is greater than the capital stock of the cell which has been abandoned, aggregate capital decreases (through outsourcing) and production falls. When agents with low paying, under capitalized jobs decide to return to school and wealth decreases, the capital market can tighten and the unemployment can spread as some processes are outsourced. On the other hand, education benefits the economy as a whole, by possibly enabling technologies that, while available and would pay high wages, cannot be used because of a shortage of graduates.

[17] She can also "flunk out" with an exogenously set probability.

If agent density is high, more than one agent can arrive on a cell at any given time in the simulation. If there is a collision, the cell will employ the most highly skilled agent. Skill is a quadratic function of age, experience, local experience (in a given cell) and education. The functional form insures that aging workers eventually lose their advantage, for the same level of education, to younger workers, as both grow older. Local experience in the current cell is also specified. When an agent moves, local experience is set to zero; but if an agent remains in a job, her skill variable accumulates rapidly. From the cell's point of view, the agent is increasingly suited to her job relative to a competitor.

2.4. **Nash bargaining.** When agents move, they move to the cell with the best advertised offer, where the best offer is determined by how far an agent can see from its starting cell. Agents then Nash bargain with cells over the wage payment. The agent's bargaining strength relies on the cell's marginal benefit of having an agent present. The agent knows that if he refuses the wage-offer, the cell will have no other option than shutting down. Conversely the cell is aware that if the agent refuses the offer, there will be no production and the agent's wealth will decline by his consumption expense for the period.

The surplus in cell i, S_i, is defined as the difference between the marginal product of labor there and the consumption expenses, c_j, of the present agent j

$$S_i = \frac{dq_i}{dL} - c_j.$$

If this surplus is positive then Nash bargaining proceeds, but if S_i is zero or negative, then the cell's only option is to meet the wage demand, c_j. If the agent is militant the offer is rejected; if not, then it is accepted.

If the surplus is positive, the standard solution to the Nash bargaining problem has the share of the surplus determined by the impatience of the agent, with impatience measured by the agent's discount rate. Here we substitute the agent's wealth-to-expense ratio under the assumption that if wealth will be consumed quickly, agents are more willing to accept a low offer and vice-versa. The effect is moderated by first taking the natural log of the wealth-to-expense ratio and then computing the share of the surplus, σ_i, by way of a logistic function L, such that

$$\sigma_i = L(W_i/c_i)$$

where W_i is the wealth of the *ith* agent. The share function is shown in figure 1

The Nash bargaining scheme is a mechanism to determine the final wage paid and may or may not be realistic. Skillman notes, for example, that the process unrealistically implies that outside offers *continuously* affect the bargaining outcome. This is inconsistent with the data for most countries at the aggregate level and challenges a sense of microeconomic realism. Skillman's solution is to introduce "endogenous termination" in which any participant can terminate the bargaining at any moment[18]. Endogenous termination is only effective when the outside offer *improves* on the solution for the agent and therefore smooths wage determination.

Note that in the multi-agent systems approach, problems arising from analytical simplification, such as those identified by Skillman, are resolved naturally by the sequential nature of the model. The code is constructed such that at the end of each

---

[18] Exogenous termination requires that bargaining comes to end in an exogenously specified time period.

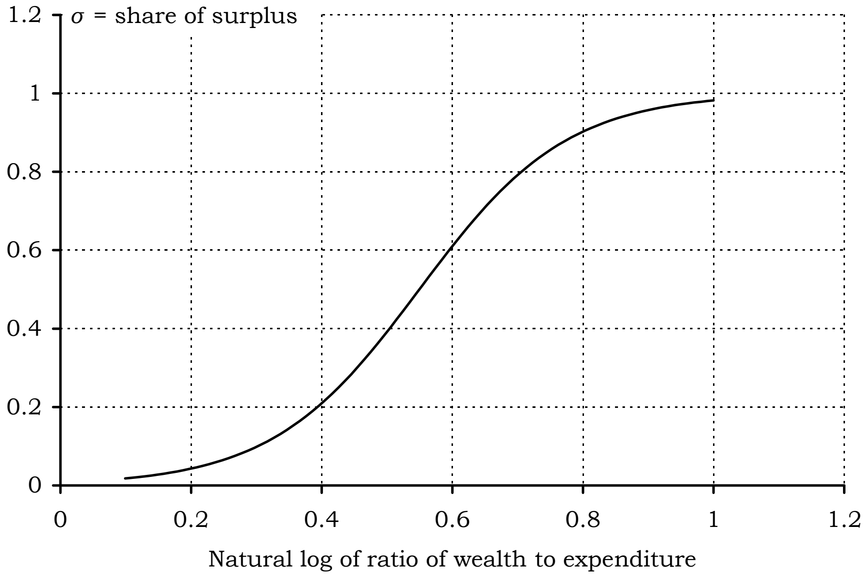

Figure 1. Share of surplus in the Nash bargaining problem

period every agent must bargain anew with its cell for a share of the production of the coming period[19]. Prior to bargaining, however, the agent is given an opportunity to move to a new cell. If the agent remains on its current cell, again by definition the outside offers are irrelevant since they are lower. On the other hand if the agent moves, outside offers were clearly relevant and endogenously terminated the bargaining on the previous cell. In any given period, then, agents Nash bargain without continuous interference from outside offers[20].

2.5. **Life cycle.** Agents live a randomly assigned number of years between the minimum life-span and the maximum. Both are set exogenously and after reaching the end of their given life-span, agents die. Population grows exogenously in the model with new agents appearing discretely as the "children" of an existing agent, randomly selected[21]. Parents transfer a randomized percentage between 20 and 80

---

[19] That agents must Nash bargain at the beginning of each period is not as unrealistic as it may at first appear. It does not, for example, generate noisy wages. Consider an agent who bargains over the wage for a second year of employment with a cell. If the agent's wage just covers expenses, wealth would remain fixed and ignoring for the moment any change in expenses, the share of surplus σ would remain constant. The wage would then follow capital accumulation. Collisions might cause wages to change abruptly, but since skill depends on "experience here," it is difficult for one agent to unseat another even if the interloper is more broadly experienced or more educated.

[20] Imagine that an agent moves from a cell in a neighborhood with little capital to a cell on the fringes of an economic opportunity zone. The assumption implies that a superior environment will not affect the agent's bargained share of the surplus until the *following period.*

[21] Children enter the world at 15 years old with a randomly assigned level of education and no job experience. Consequently, their bargaining share is low, especially when they inherit an extravagant lifestyle from their parent.

percent of their wealth their offspring. The rest of their wealth is lost to consumption in the period of transfer, spent on transaction costs for services, funeral, legal and the like. These services are provided by existing technologies and no other adjustment need be made.

If the agent is educated, both human and physical capital can be lost at the time of agent death. Output can fall if the replacement agent is poorly educated, has no degree or must compete for a job with another agent who is more skilled or experienced. If the agent cannot find work, aggregate wealth falls and rate of growth of the capital stock declines.

2.6. **Reinforcement learning and mobility.** If wage income falls below consumption expenses, agents become unsatisfied with their current jobs and may begin a search for a new one. But agents may also move even if they are satisfied with their current jobs, but wish to explore the map of the environment and its associated rewards[22]. In simple reinforcement learning models, the *value function* is an aggregator of rewards as they accrue to a specific agent. The decision to move increases the agent's knowledge of the value function. The simplest response is that if the value function is rising, the probability that an agent will move in the next period, even if satisfied with her job, increases. Agents update the value function using the simple algorithm

$$V_{t+1} = \delta_m \left(V_{t+1} - V_t\right) + V_t$$

where V is the value function and δ_m is 1 if the agent moves and 0 if not. Here the value function is the surplus of wage income over living expenses. If the agent's move is associated with improvement in the value function, then the probability of moving in the future is increased. Agent mobility, M, is the given by the logistic function of the current value.

2.7.

$$M_t = L(V_t)$$

Dynamics. There is no investment function in the model as traditionally defined, but, as noted, the sum of value of the capital stock must be equal to total agent real wealth. Cells operating processes more profitable than average are able to find financing to expand their production while those that are less profitable are less likely. If cells cannot attract capital, depreciation will cause the maximum wage to fall, lowering the surplus and creating the conditions for agents to depart. As noted above, cells have an initially assigned demand for capital stock level consistent with their technology. If activated by the presence of an agent, the capital stock must be financed out of the total wealth of agents. If total wealth is less then the value of the capital stock, some processes are outsourced. The choice of which cell is to shut down and go abroad is made by ranking profitability along with a random error term.

If total wealth exceeds the sum of the activated capital stock, investment takes place on the *ith* producing cell according to the Gibbs-Boltzmann distribution of

---

[22] In the AI literature, agents are either a "greedy" or "exploratory" in their choices (Sutton and Barto, 1998). Note the subtle difference in terminology relative to the way the word greedy is used in economics. For the AI literature "greedy" clearly implies short-sightedness.

shares, σ_{I_i}

$$\sigma_{I_i} = \frac{e_i^{\pi_i/\tau}}{\sum e^{\pi/\tau}} + \gamma\epsilon_i$$

where π is the current rate of profit and τ is a parameter of the distribution. A uniformly distributed random error term, ϵ_i, with mean of zero and weight γ is added so that the pattern of investment is not strictly determined by profitability. Nevertheless, more profitable cells receive, on average, higher shares of available financing.

High profitability on cells results from the combination of a large capital stock together with a low reservation wage. Cells with a minimal endowment of capital stock cannot pay high wages and therefore must either raise their prices or hope to find an agent with modest consumption expenses. Since this is sometimes possible, it follows that there can be discontinuous departures from any implicit balanced growth path, departures in which small but highly profitable cells occasionally experience explosive growth.

The model can easily produce unstable dynamics as a result. If a cell's share in new finance is less than depreciation, the capital stock will fall and eventually disappear. Any resident agent would most likely leave before the capital stock was exhausted, but then other agents could arrive and force the cell to raise its price. Since such agent-cell combinations produce no net increase in total wealth, yet draw on capital, they can easily lead to outsourcing. This in turn reduces wealth further, along an unstable branch.

2.8. **Conflict.** Nothing in the model prevents two or more agents from arriving on the same cell[23]. The cell queries the agents in residence and makes a list of qualified applicants, depending on whether a graduate is required. The cell then selects the most skilled of these agents and Nash bargains over the wage rate. If the surplus is negative, the cell would still prefer to produce and raise its price, as discussed above. But in the case in which more than one qualified applicant is available, the cell assigns a "rejected" code to the agent with a negative surplus and then Nash bargains with the runner up, the next most skilled. This is an improved survival strategy for the cell since the chance that it will be a victim of outsourcing is reduced. If the Nash bargaining fails with all qualified applicants, the cell searches for the cheapest of the rejected candidates, pays the reservation wage and raise the price to compensate. If the candidate is a punisher, however, the cell may offer the reservation wage and find that it is rejected. In this case the cell, fails to produce, just as above[24].

2.9. **Program flow.** The program is structured as follows. Agents are endowed with a given amount of wealth and cells are given blueprints. At the opening bell, every agent must strike a deal with the cell on which it resides. If the deal is successfully concluded, there is production and income is distributed. Each agent's age, wealth, experience, education, mobility, skill, probability of going back to school, impatience and job satisfaction are then updated. Cell capital stock is

---

[23] If n agents occupy the same cell, and $m < n$ make a decision to return to school, the program prevents the cell from becoming a school and thereby depriving the $n - m$ agents of a job. The m agents must move, thereby postponing their decision to return to school by one period.

[24] Currently, the cell gives up and does not try to negotiate with any other qualified candidates.

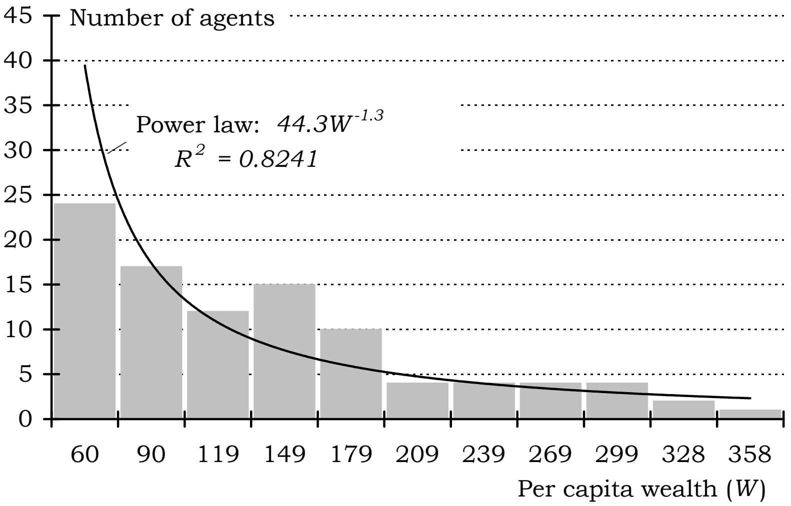

FIGURE 2. Initial wage distribution (n = 2601)

then adjusted according to the balance of wealth and activated capital. If the net capital stock rises, the maximum wage increases with the marginal product and vice-versa. The surplus is then updated to reflect these changes and adjustments to the reservation wage, if any. On the basis of this new distribution of job opportunities, agents decide whether to move or stay in their current location. The model then repeats.

3. SIMULATION RESULTS

The model is initialized with a population of 48 (and a growth rate of 1.5%), depreciation of 4%, maximum education of 4 and a distribution of good jobs that resembles figure 2[25]. School cost is 6% of wealth per year and cells without agents raise their wage offers by 1% per period until either they attract an agent or the maximum wage is reached. The initial distribution is set such that 25% of the cells are defined as good jobs, with a wage offer equal to its maximum wage. These are diffused to create economic opportunity zones[26]. Initially all agents have the same wealth and their living expenses average about 1.5 with a standard deviation of 0.34. There are enough productive technologies to ensure job satisfaction for all agents although agents may not have sufficient education to see them. Indeed, randomly placing 48 agents on 2601 cells yields a probability of job-satisfaction of less than a third. Life expectancy is set at 72 with minimum of 50. Expenses grow

---

[25] Both a power law and exponential distribution were fitted for comparison but the exponential seems to produce a better fit. Axtell shows that the distribution of firms is closer to a power law than an exponential (Axtell, 1999).

[26] After the initial distribution of wages, 25% of the wage is spread to a neighbor. This is repeated five times and the result is distribution of wage shown in figure 2.

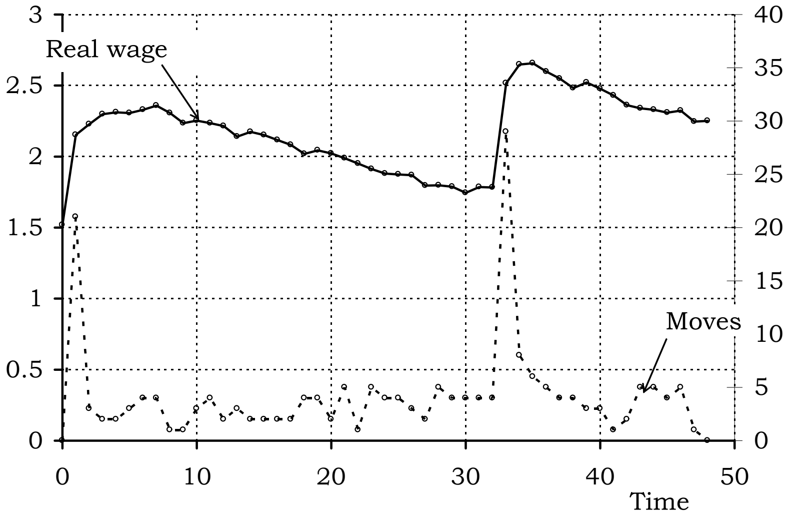

Figure 3. Real wage and number of moves

with age with an elasticity of one to maximum at retirement age of 61 and then decline[27].

3.1. **Unemployment, inflation and job satisfaction.** Unemployment in the model is mostly frictional in that there are no cyclical effects that are imposed on the model from outside. The standard explanation for cyclical unemployment is that the demand for labor increases, driving down the rate of profit and thus investment. Aggregate demand falls, unemployment rises and output contracts, causing a multiplier effect on aggregate demand that reduces it even further. Unemployment slows wage increases and thus inflation. Thus, inflation and cyclical unemployment should not coexist, as noted above, but rather trade-off along the conventional Phillips curve. Figure 3 depicts the relationship between the average real wage and number of agent moves per period.

There are no cycles here because aggregate demand as such plays no independent role in this model. Inflation arises in the model because there is a mismatch between wage demands and the marginal product of labor in a given cell. The conflict between wage expectations and the ability to meet those wage demands drives inflation, but the overall level of unemployment has no effect on the matching process. Thus job dissatisfaction, as the root cause of inflation, is *positively* correlated with unemployment and so would be inconsistent with the standard Phillips curve.

Structural unemployment is certainly a factor in the model, but it too is rooted in job dissatisfaction[28]. Mismatching causes output and income to fall as prices rise and vulnerable technologies are outsourced[29]. This can have a cumulative effect

---

[27] Retirement age is set at the average of life-expectancy and minimum life-expectancy.

[28] One can see how important job satisfaction is in this model by raising the percentage of good jobs in the initial period. Inflation virtually disappears from the system.

[29] Outsourcing takes place with or without inflation, so long as financing is not available. Hence, the correlation with inflation in the model results is not strong.

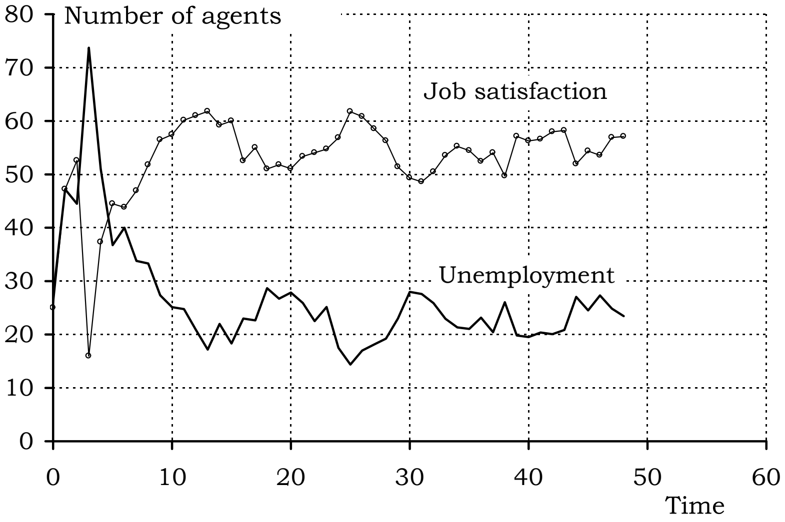

FIGURE 4. Unemployment and job satisfaction

on the model and it is easy for the model to enter an unstable trajectory, with outsourcing causing a decline in income which then leads to more outsourcing as wealth declines. Foreign borrowing could no doubt prevent this in reality, but it is not built into the model.

Figure 4 shows the path of unemployment over time in a characteristic run of the model along with the level of job satisfaction. It is clear from this diagram that the two series are highly negatively correlated (-0.795). It is evident that it is not high wages that cause unemployment, but rather low wages that lead to job-dissatisfaction. Since there is no guarantee that an agent will find another job if she leaves one, the rate of unemployment can easily rise as a result.

Observe that the model more likely predicts "stagflation" than a profit squeeze on the level of investment. A profit-squeeze is possible, however, in the following sense. At the existing wage, let expenses rise to cause less job satisfaction, say as a result of a demographic bubble. More movement comes about and output falls as agents seek more remunerative jobs but are not uniformly successful. As a result, unemployment rises and aggregate wealth falls, which in turn increases the rate of cell death. Unemployment then increases structurally, initiated by a profit squeeze.

There is an additional factor at play in the model, a skill mismatch. If an agent is a graduate, yet the cell does not require a graduate, she may experience job dissatisfaction despite that her income exceeds expenses. If she moves as a result, she may not find another job and the unemployment rate might increase temporarily. Moreover, when good jobs are not plentiful, wealthy but less educated agents can make a decision to return to school and may earn a degree[30]. If the demand for graduate labor does not increase in the meantime, the new graduate may become unemployed during an extended search for a satisfactory job. Aggregate

---

[30] The probability of not finishing any given year is currently set to 20% and the wealth constraint is set at (6 - education) × wealth).

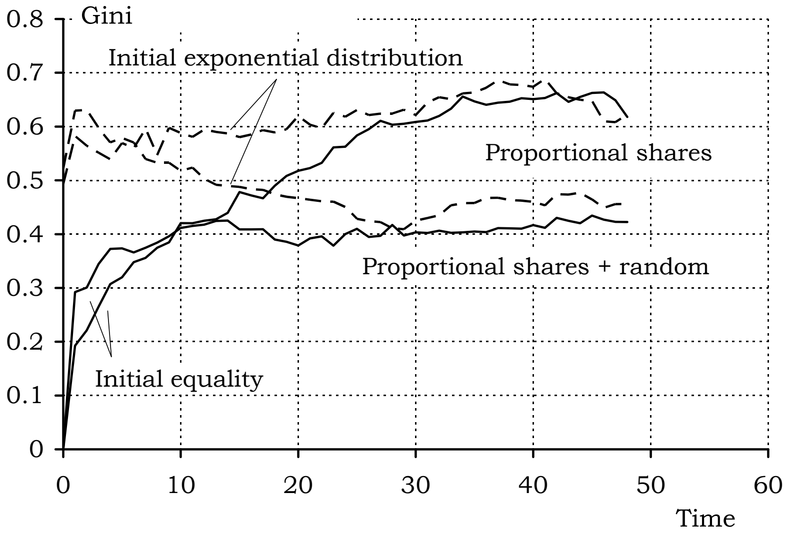

Figure 5. Gini coefficient versus time

wealth will fall, in turn, and cause the economy to contract or at least to slow its rate of growth.

3.2. **Income distribution and growth.** It appears that an inegalitarian distribution of income is an emergent property of the model, as it is in many agent based models of income distribution (Axtell et al., 2001). In this section, we explore the effect of the initial distribution of income and also risk in the rate of return to accumulated wealth.

In the same run that generated figure 4, the Phillips curve is vertical with little inflation. The initial distribution of good jobs is the same as in Figure 2 above and by the end of the 50-year period, only 57% of the population is in a satisfying job. The economy grows rapidly with per capita income increasing by 2.1% per year. Figure 5 shows four paths of the Gini coefficient. The two solid sequences have an initially egalitarian distribution of wealth, but for the solid upper sequence, economy-wide total profit is distributed in proportion to wealth. In the lower, total profit is distributed in proportion to wealth plus a random term with approximately equal weights[31]. The same is true of the upper and lower dashed sequences.

Note first that the initial conditions matter to the path of the Gini, but quickly dampen out. For both, the Gini coefficients paths are similar after about 25 years. Initial equalities disappear after only a generation. The random factor in the rate of return seems to be more important.Figure 6 shows a typical run of the model. There the "poor" are defined as agents with less than 20% of the average wealth of the "rich", in turn defined as the top 20% of the population. Note that at time 0, all agents have the same total wealth and so all 48 are rich. The number of poor is increasing fairly dramatically in this simulation, suggesting rising relative

---

[31] Shares are determined by randomly assigning each agent a number between -0.15 and 0.85 and then normalizing by the sum. The agent's share is then determined by an exogenously weighted sum of their wealth share and this randomly assigned share.

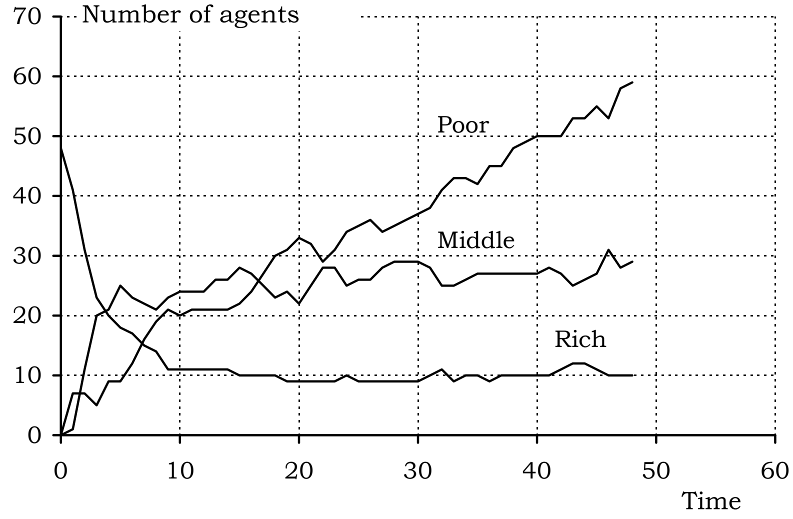

Figure 6. Income classes

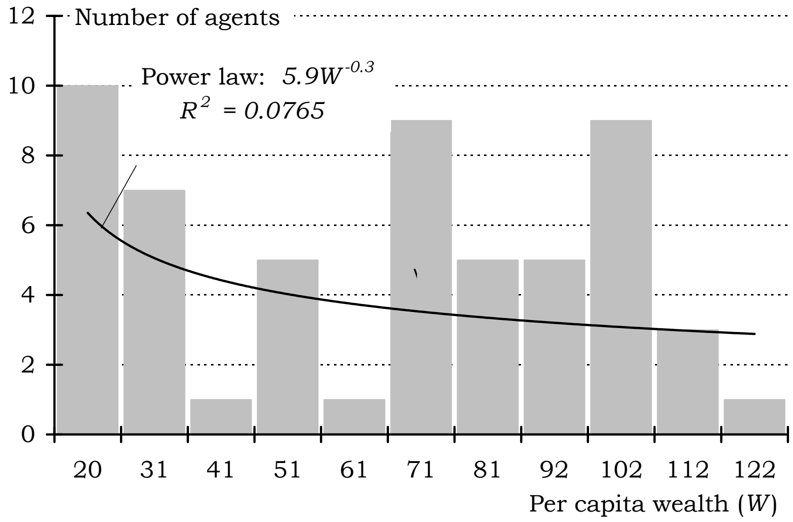

Figure 7. Wealth distribution ($t = 10$)

poverty. The distribution of wealth is shown in figures 7, 8 and 9 after 10, 30 and 50 periods. Over time, the distribution of wealth approaches the Pareto or power law distribution. It is clear from how the R^2 increases that the power law does not apply in the early stages of the development of the distribution of wealth but emerges toward the end of the simulated period.

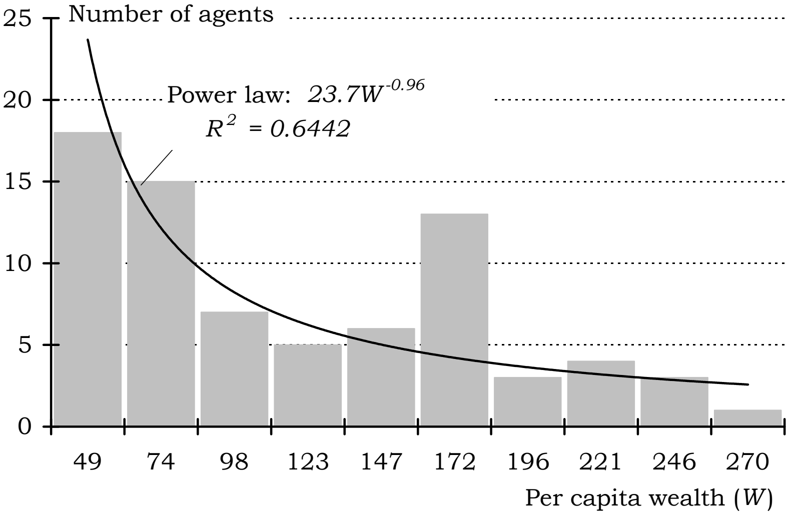

FIGURE 8. Wealth distribution ($t = 30$)

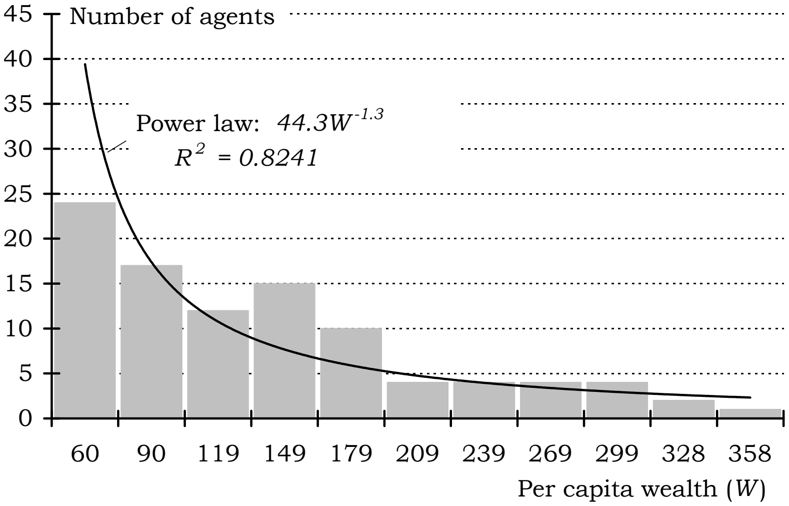

FIGURE 9. Wealth distribution ($t = 50$)

Why does the distribution of income deteriorate in this economy so definitively? Table 1 provides some preliminary answers. Analyzing the cross-sectional wealth correlations at the beginning, middle and end of the simulated period, it is evident that the strongest correlation between wealth is for property income. Correlation coefficients when $t = 10$, with agent population that increased from 48 to 55 are shown in the first column. When property income is distributed in strict proportion

TABLE 1. Wealth correlations

	Correlation coefficients[1]		
	$t = 10$ $(n = 55)$	$t = 30$ $(n = 74)$	$t = 50$ $(n = 97)$
Age	0.77	0.67	0.56
Property income	1	0.85	0.96
Skill	0.69	0.55	0.54
Graduate	-0.09	0.06	0.09
Experience	0.47	0.47	0.46
Mobility	-0.18	-0.05	-0.07
Education	-0.1	-0.05	-0.05
Expenses	0.43	0.46	0.29
School prob	-0.45	-0.28	-0.27
Share	0.43	0.19	0.14
Wage income	0.4	0.4	0.24
Experience here	0.54	0.49	0.53

Source: Model computations. 1. Property income distributed in proportion to wealth.

to wealth, the correlation of age and wealth in this case is 0.77. Property income is initially distributed in proportion to initial wealth and the run has not had enough time to upset this balance; thus, the correlation between the two rounds to one. Note that the skill and experience variables are all relatively highly correlated with wealth accumulation, while the mobility and educational variables are not. Indeed, mobility is negatively correlated, as is education and school probability.

With a more randomized rate of return on accumulated wealth, the outcome is slightly different. Property income is no longer as strongly correlated with wealth, but *experience* is. This suggests that the most important factor in wealth accumulation is obtaining a good job and keeping it. A good job is, of course, one that covers the agent's expenses while at the same time generating sufficient profit to attract capital to expand.

The reinforcement learning model driving mobility shows that moving from job to job in this model is not a wealth enhancing strategy. Agent mobility decreases over time as a result. Neither, surprisingly, is education correlated with wealth accumulation. Here agents choose to return to school as a reactive strategy, when they become dissatisfied with their jobs. School then is costly to wealth, both directly and as an opportunity cost. Obtaining a job once one has completed school is also difficult since while the job search is enhanced, those who remain at work accumulate experience that puts them into a strong competitive position. In the longer run, of course, it is still better to be a graduate than not. Militants are not wealthy in the model.

Growth in this model depends on technological change and capital accumulation, as in virtually any economic model. Growth comes from resources not consumed

Table 2. Wealth correlations (con't)

	Correlation coefficients[1]		
	$t = 10$ ($n = 55$)	$t = 30$ ($n = 74$)	$t = 50$ ($n = 97$)
Age	0.70	0.84	0.77
Property income	0.45	0.52	0.42
Skill	0.56	0.55	0.76
Graduate	-0.22	-0.15	0.13
Experience	0.44	0.65	0.66
Mobility	-0.07	-0.04	-0.19
Education	-0.32	-0.23	0.05
Expenses	0.22	0.09	0.36
School prob	-0.27	-0.28	-0.37
Share	0.33	0.38	-0.06
Wage income	0.5	0.23	0.49
Experience here	0.46	0.45	0.74

Source: Model computations. 1. Property income distributed in proportion to wealth with a uniformly distributed random error.

and is directed to technologies with labor suited to it. Outsourcing, which is endogenous in the model, is the result of slow technological progress or consumption patterns inconsistent with technological change.

4. Conclusions

The question addressed in this paper is whether a multi-agent system generates recognizable macroeconomic properties. Since the work of Debreu and others in the 1970s, it is has been widely recognized that the Walrasian system cannot be trusted to provide analytical microfoundations for Keynesian macroeconomics. One cannot assume that atomistic agents with full information will necessarily generate an aggregate excess demand function that is downward sloping, even though the demand curves for individual agents are. In response, economists have moved away from both the Keynesian and Walrasian framework in search of more robust and realistic models (Colander, 2003),(Gibson, 2003). Agent-based models are an alternative, but nothing constrains these systems to look like a real macroeconomy when aggregated. There are no demand or supply curves here, no markets in the traditional sense and no auctioneer as in the standard model. There are only technologies available and agents must operate them successfully or face economic extinction. Whether agents look for and find better technologies depends on whether they are satisfied with what they are doing at the moment[32]. Operating a

---

[32] Consider this letter to the editor in the *New York Times.*

Yesterday I was offered a tech-suuport job for a publicly traded company. It pays a paltry \$9 to \$10 hour. When I squawked that this is not a livable wage, I was hesitantly offered \$11.75...no benefits...I added up the cost of the two-hour-a-day commute, a mortgage on an average home, health insurance...home and auto insurance and utilities. The break-even point was \$10.35. Take into account laundry, groceries, clothing and other basic expenses, and I am working at a deficit...

technology requires financing, available from a financial system that channels savings from agents to cells. Rates of return on savings can be more or less risky, and wealthy agents generally do a better job at finding high-paying jobs, perpetuating inequality.

Broadly speaking, the multi-agent system is capable of generating persistent unemployment and inflation as one typically sees in macroeconomies. It also occasionally becomes unstable. When agents demand higher wages than technologies can support, the conflict generates inflation. But unemployment is the result of job dissatisfaction, not a shortage of effective aggregate demand simply because there is no aggregate demand in the model of which to be short. If unsold goods remain on the shop shelves, the Keynesian model shows an increase in unintended inventory accumulation. The same takes place here, but since there are no firms to react in any way, these unsold goods get recycled into other technologies and are resold in future periods.

There is excess technological capacity in the sense that the economy could produce more with a better fit between agents and their blueprints, but there is no excess financial capacity. If there is inadequate financing, processes disappear; if there is a surplus, agents find a way to use it.

The emergence and persistence of high Gini coefficients is traced here to a simple rule similar to what drives the Schelling neighborhood model. If computation is costly in terms of foregone earnings on the part of both agents and cells, if follows that those who are able to minimize search time by finding a good job early and keeping it will be better off in the long run. Policy that does not address this fundamental fact of modern market economies will find that solutions to distributional inequality are difficult to implement and will have disappointing results.

References

Abdallah, S. and V. Lesser (2007, May). Multiagent reinforcement learning and self-organization in a network of agents. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, Honolulu.

Aexlrod, R. (1997). *The Complexity of Cooperation: Agent Based Models of Competition and Collaboration.* Princeton Studies in Complexity. Princeton, NJ: Princeton University Press.

Anderson, P. W. (1972, August). More is different. *Science 177*(4047), 393–396.

Arthur, W. B., S. Durlauf, and D. Lane (1997). *The Economy as an Evolving Complex System,*, Volume 2. New York: Addison-Wesley.

Axelrod, R. (1984). *The Evolution of Cooperation.* New York: Basic Books.

Axelrod, R. and D. S. Benett (1997). A landscape theory of aggregation. In R. Axelrod (Ed.), *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*, Chapter 4, pp. 73–94. Princeton, N J: Princeton University Press.

Axtell, R. (1999, June). The emergence of firms in a population of agents. http://www.brookings.edu/es/dynamics/papers/firms/firmspage.htm.

Pay workers sufficiently and they will be loyal and dependable. Stiff them with low wages and they immediately begin to look for something better.

(September 4, 2007).

Axtell, R., J. Epstein, and P. Young (2001). The emergence of classes in a multi-agent bargaining model. In S. Durlauf and P. Young (Eds.), *Social dynamics*, pp. 191–212. Cambridge, MA: MIT Press.

Berry, B. J. L., L. D. Kiel, and E. Elliott (2002, May). Adaptive agents, intelligence, and emergent human organization: Capturing. *Proceedings of the National Academy of Sciences of the United States of America 99*(10), 7187–7188.

Boutilier, C., Y. Shoham, and M. P. Wellman (1997). Economic principles of multi-agent systems. *Artificial Intelligence 94*(1-2), 1–6.

Brock, W. A. and S. Durlauf (2005). Social interactions and macroeconomics. http://www.ssc.wisc.edu/econ/archive/wp2005-05.pdf.

Cheng, J. and M. P. Wellman (1998). The walras algorithm: A convergent distributed implementation of general equilibrium outcomes. *Computational Economics 12*(1), 1–24.

Colander, D. (2003). Post walrasian macro policy and the economics of muddling through. *International Journal of Political Economy 33*(2), 17–35.

Debreu, G. (1974). Excess demand functions. *Journal of Mathematical Economics 1*, 15–21.

Durlaf, S. N. and H. P. Young (2001). *Social Dynamics*, Volume 4 of *Economic Learning and Social Evolution*. Cambridge, MA: MIT Press for the Brookings Institution.

Durlauf, S. (1999). How can statistical mechanics contribute to social sciences? Prepared for the Proceedings of the National Academy of Sciences.

Epstein, J. and R. Axtell (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. Washington, DC:: Brookings Institution Press.

Gatti, D. D., E. Gaffeo, M. Gallegati, G. Giulioni, and A. Palestrini (2008). *Emergent macroeconomics*. New Economic Windows. Frankfort: Springer.

Gibson, B. (2003). Thinking outside the walrasian box. *International Journal of Political Economy 33*(2), 36–46.

Güth, W., R. Schmittberger, and B. Schwarze (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization 3*(4), 367–388.

Henrich, J., R. Boyd, S. Bowles, C. Camerer, E. Fehr, and H. Gintis (2004). *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*. Oxford: Oxford University Press.

Horling, B. and V. Lesser (2005). A survey of multi-agent organizational paradigms. *The Knowledge Engineering Review 19*(4), 281–316.

Johnson, G. (1995). *Fire in the Mind*. New York, NY: Alfred A. Knopf, Inc.

Lucas, R. E. and T. J. Sargent (1978). After keynesian macroeconomics. In *After the Phillips Curve: Peristence of High Inflation and High Unemployment*, pp. 49–72. Boston, MA: Federal Reserve Bank.

Railsback, S. F., S. L. Lytinen, and S. K. Jackson (2007). Agent-based simulation platforms: Review and development recommendations. *Simulation 82*(9), 609 – 623.

Rizvi, S. A. T. (1994). The microfoundations project in general equilibrium theory. *Cambridge Journal of Economics 18*, 357–377.

Rizvi, S. A. T. (1997). Responses to arbitrariness in contemporary economics. *History of Political Economy 29 Supplement*, 273–88.

Sandholm, T. W. and V. R. Lesser (1997, July). Coalitions among computationally bounded agents. *Artificial Intelligence 94*(1-2), 99 – 137.

Schaeffer, J., N. Burch, Y. Bjornsson, A. Kishimoto, M. Muller, R. Lake, P. Lu, and S. Sutphen (2007, 19 July). Checkers is solved. *Science 317*(5844), 1518 – 1522.

Schelling, T. (1971, July). Dynamic models of segregation. *Journal of Mathematical Sociology 1*, 143–186.

Sutton, R. S. and A. G. Barto (1998). *Reinforcement Learning:*. Cambridge, Massachusetts and London, England: MIT Press.

Wilensky, U. (1998). Center for connected learning and computer-based modeling. http://ccl.northwestern.edu/netlogo/.

Wooldridge, M. (2002). *MultiAgent Systems*. West Sussex: John Wiley and Sons, LTD.