Macroeconomics with Intelligent Autonomous Agents*

Peter Howitt

Brown University

June 1, 2007

*Paper presented at the festschrift in honor of Axel Leijonhufvud, UCLA, August 30–31, 2006.

1 Introduction

Axel Leijonhufvud has spent much of his distinguished career investigating how a decentralized economy coordinates economic activities. The question is basically the same as the one Keynes (1934) once posed - to what extent, and under what circumstances, particularly under what policy regimes, is the economy self-adjusting. In recent years he has advocated the use of agent-based computational economics as an approach to the problem.[1] The present paper discusses this methodology and describes an ongoing research agenda aimed at implementing it.

As described by Tesfatsion (2006), agent based computational economics is a set of techniques for studying a complex adaptive system involving many interacting agents with exogenously given behavioral rules. The idea motivating the approach is that complex systems, like economies or anthills, can exhibit behavioral patterns beyond what any of the individual agents in the system can comprehend. So instead of modelling the system as if everyone's actions and beliefs were coordinated in advance with everyone else's, as in rational expectations theory, the approach assumes simple behavioral rules and allows a coordinated equilibrium to be a possibly emergent property of the system itself. The approach is used to explain system behavior by "growing" it in the computer. Once one has devised a computer program that mimics the desired characteristics of the system in question one can then use the program as a "culture dish" in which to perform experiments.

Now the first reaction of many economists upon first hearing about this methodology is that all economic models with an explicit micro-foundation, which is to say almost all models that one sees in mainstream macroeconomic theory, are "agent-based". Some even have a multitude of heterogeneous agents (see Krusell and Smith, 1998 and Krebs, 2003, among others). So what's the big deal?

The big deal, as Tesfatsion has emphasized on many occasions, has to do with autonomy. An agent in a rational-expectations-equilibrium model has a behavioral rule that is not

[1] Leijonhufvud (1993, 2006)

independent of what everyone else is doing. In any given situation, her actions will depend on some key variables (prices, availability of job offers, etc.) or the rational expectation thereof, that are endogenous to the economic system. These variables will change when we change the agent's environment, and hence her behavior cannot be specified independently of the others'. The household, for example, in a market-clearing model of supply and demand cannot choose what quantity to demand until told what price will clear the market. Likewise the agent on a Lucas island (a Phelps Island with rational expectors) cannot choose how much to sell until informed of the stochastic process determining aggregate and relative demand fluctuations.

The problem with assuming non-autonomous agents is that it leaves the model incomplete, and in a way that precludes a deep analysis of the coordination problem. For if the model does not allow people to act without knowing the equilibrium value of some variable, then someone must have computed that equilibrium value a priori. In such a model there is no way to describe out-of-equilibrium behavior, and the problem of reconciling peoples' independently conceived plans is assumed to be solved by some unspecified mechanism that uses no scarce resources. Autonomy is thus essential to the problems that Axel has done so much to keep alive since the rational-expectations revolution.

Now under certain assumptions about common information, someone endowed with enough information could figure out on her own what the market-clearing price is going to be, or what the rational expectation of the price level is, and in this sense could act autonomously even in a rational-expectations equilibrium framework. But an economy full of agents that were autonomous in this sense would not be decentralized in the Hayekian sense, because no market would be needed to aggregate the diverse information of heterogeneous people, each of whom can do the aggregation in her head. Each would be capable of acting as the economy's central planner, although in this case the planner would not be needed. Moreover, such an economy would have no need for macroeconomists, because everyone would already know as much as could be known about the macroeconomy. The coordination problem would be trivial. So by "autonomous" agents I mean agents that are endowed with behavioral rules that can tell them what to do in any given situation, independently of each others' rules, even when no one has access to a correct model of the economy.

The literature on learning in macroeconomics, recently surveyed by Evans and Honkapohja (2001), specifies autonomous agents according to this definition. For example the least-squares learner in the Cagan model of hyperinflation has a sequence of observations on past inflation and acts according to a rule that tells her to demand the quantity of money whose log is a given linear function of the rate of inflation predicted by the OLS estimator. This function can be specified independently of what everyone else in the economy is doing. Of course there is the problem of specifying the time-series model correctly, and this will depend on the nominal demand process driving inflation, in a way that someone lacking a correct model of the economy could not know for sure. Much of the literature supposes that people are endowed with knowledge of how to specify the model to be estimated in such a way that it will eventually be correct if the economy converges to a rational-expectations equilibrium. But Evans and Honkapohja also cover the case of misspecified models, where the economy might or might not converge, and sometimes to an equilibrium that is not rational-expectations. In such a case, the agents are indeed autonomous.

Although I believe that much can be learned from the macroeconomic learning literature, and have discussed this at length elsewhere (Howitt, 2006b), that literature does not come to grips with what I now see as one the most salient aspects of how people behave in an economic environment that they don't understand, and that they know they don't understand. This aspect was an important theme of Keynes, who argued that people are not in a position to act according to the conventional theory of rational choice if they cannot attach numerical probabilities to all possible consequences of their decisions. Keynes argued that under such circumstances they tend to cope by falling back on custom and convention. They also devise institutions to insulate themselves from having to rely upon necessarily unreliable forecasts.

As Axel has argued, macroeconomic life is full of unquantifiable uncertainty. He once likened forecasting inflation to playing a game of chess refereed by someone who announces that "From now on bishops move like rooks and vice versa . . . and I'll be back with more later." (Leijonhufvud, 1981, p. 264). Hence a lot of economic decisions, particularly those involving risks of losses from inflation, are based on conventional or institutionalized rules of thumb.

Paul Davidson (1989, pp.15-17) has also argued that money, in its role as a unit of account and standard of deferred payment, is an institution through which people cope with uncertainty without having to rely upon necessarily imperfect predictions of the future. In a money-using economy, firms and households are concerned not just with their "real" economic profits, but also with their cash-flow, for no matter what happens to the value of money they can at least stay out of the bankruptcy court as long as inflow exceeds outflow. Historical cost accounting helps firms to keep track of their cash flow better than would an indexed system, and nominal, nonindexed debt contracts allow them to insulate their cash flow from unpredictable fluctuations in the price level, especially in a world where their customers demand some assurance of predictable nominal prices. Controlling real cash flow would be the ideal objective if it were possible, but controlling nominal cash flow, with the aid of such devices as nominal debt and historical cost accounting, is at least a useful objective with the advantage of being reasonably attainable.

Likewise, one of the central themes of David Laidler's work is that money is a device for economizing on the costs of processing information. People use it as a buffer stock that automatically absorbs unforeseen changes in income and expenses without the need for deliberation. They also use it as a unit of account, measure of value and standard of deferred payment because it is convenient to use, conventional and easily understood, even if this seems to introduce biases and inefficiencies into their decision making and even if economists can think of better measures and standards.[2]

From the point of view of mainstream macroeconomics, people that follow rules of behav-

[2] For example, Laidler (1974, 1984).

ior that have no obvious rationalization in terms of optimal choice appear to have "bounded rationality". Economists are socialized to be skeptical of any theory that relies on such limits to human intelligence. But the case can be made that using simple rules that seem to work reasonably well is actually a more intelligent way to arrange one's affairs in an uncertain world than the more conventional Bayesian alternative of fabricating a model of the world and choosing a plan that would be optimal under the incredible assumption that the model was a true representation of reality. Under the Bayesian alternative, not only would specification error be likely to lead the decision-maker astray, but Bellman's curse of dimensionality would render the task of implementing the strategy unworkable, since it would involve computing a solution to a very complicated model, unless the model was made tractable by other convenient heuristic devices, like assuming all agents identical, only two possible states of the world, Cobb-Douglas aggregate production functions, and so forth.

This view as to what constitutes intelligent behavior is explained vividly in the recent book by Andy Clark (1998) which I first learned about at Axel's 2006 Trento summer school. Clark's book is an account of recent developments in cognitive science and artificial intelligence, developments exemplified by the idea of "neural networks." His thesis is that human intelligence is not to be thought of as an abstract reasoning capability joined to a memory bank of facts, but rather as a device for controlling the body's varied set of adaptive behaviors in a way that helps the body cope with the particular environment it finds itself in. Clark calls this view "embodied, environmentally embedded cognition." It portrays intelligence not as a central computer program solving a well defined maximization problem but as a decentralized network of autonomous neurons that interact with each other, often sending conflicting messages, and often competing to execute the same task. Intelligence emerges not from the capacity to solve planning problems but from simple chemical reactions that reinforce the neural processes that have been associated with improvements in the body's well being and weaken the processes that have not.

Clark points out that the process of human adaptation is not guided by an internal model of the world which the brain takes as defining the constraint set for optimization, but rather it consists of simple rules for acting in ways that cope quickly and effectively with environmental hazards such as the presence of predators, the need for food, and so on. These rules may sometimes make use of internal representations but typically they need to operate much faster than the construction and use of any such representation would allow – people typically need to make economic decisions faster than would be possible by the use of dynamic programming. The intelligent individual is not the one capable of solving big analytical problems but the one that has learned useful tricks and strategies with which to act quickly in a way that is well adapted to one's environment.

In short, if we are to study the economy's coordination mechanisms we must specify autonomous agents, and if we are to endow these autonomous agents with what commentators from Keynes to Clark have argued is the essence of intelligence, then we are driven to assume that they act according to simple rules. The key to agent-based modeling is not to make use of the classical distinction between estimation and optimization, but to find good, robust behavioral rules that map situations directly into actions. The distinction between these two modelling strategies is basically the same as that made in the literature on learning in games, between learning/optimization and stimulus/response. Although someone that first estimates a model and then optimizes subject to the estimated model will end up with a rule of the same general form as the person that adapts the rules directly to success in achieving her goals in the given environment, the former strategy has proven in artificial intelligence applications to produce brittle outcomes; behavior that makes no sense whatsoever when circumstances change in a way that violates the exclusion restrictions of the estimated model.

Axel once remarked that the real problem of macroeconomics was to understand how order can arise from the interactions of people following simple rules to cope with a complex environment, and contrasted that with much of mainstream macroeconomics which postulates people using complex decision procedures to deal with a simple environment.[3] This is the challenge I wish to take up. How to model the macroeconomy as a human anthill; one that organizes individuals' activities into patterns more complex than the individuals can fully comprehend, performs collective tasks that the individuals are hardly aware of, and adapts to shocks whose consequences none of the individuals can predict. Accounting for spontaneous order as an "emergent property" of a complex system is one of the main themes of agent-based modeling, and the research agenda that Axel has been advocating is to apply this idea to the study of macroeconomic coordination, by modeling people not as rational maximizers but as intelligent autonomous agents.

2 Modeling the coordination process

Another question raised by the idea of autonomous agents is how anyone could possibly be autonomous when it takes two to tango. Surely in a complex economy such as ours each person's behavior and well-being are inextricably linked with the others' behavior. What this question highlights is that assuming intelligent autonomous agents is just the starting point for studying the coordination problem. Economics being a social science, the really important issue is not so much how people behave as how they interact. An autonomous agent has simple rules for finding other people, sending communications to them, and responding to their communications. What sort of exchange patterns emerge from the interaction between these rules, how orderly and stable they are, and how they evolve over time is the quaesitum of the research agenda Axel has proposed.

Several years ago, Robert Clower and I took this question up using what I now recognize as agent-based computational economics. We started from the observation that in the real world trades are coordinated by a self-organizing and self-regulating network of trade specialists

[3] In (1993, pp.1-2) he quoted Daniel Heymann as having "remarked that practical men of affairs, if they know anything about economics, often distrust it because it seems to describe the behavior of incredibly smart people in unbelievably simple situations", and then went on to suggest "asking how believably simple people cope with incredibly complex situations."

– shops, brokers, middlemen, banks, realtors, lawyers, accountants, employers and so forth. Modern economic life is largely a sequence of exchanges with people outside one's own family and social circle. Almost every such exchange involves a specialized trader ("shopkeeper") on one side or the other of the market. Shopkeepers are the agents that undertake to match buyers and sellers, arrange terms of exchange and bear the costs of adjustment in the face of day-to-day imbalances between spending plans and productive capacity; in short, they are the real-world counterparts of the fictional auctioneer of general equilibrium theory.

In our (2000) paper we attempted to model in a crude way how such a network of coordinating shopkeepers might emerge spontaneously, from elementary interactions between people following simple opportunistic rules of behavior that represented what we considered to be the most salient activities of a decentralized economy. The idea was not to say this is how economic organization emerged but rather that this is a model of economic organization that passes the minimal test of being self-organizing. If the organizational structure was not there, it would quite likely arise from the interaction of intelligent autonomous agents. Moreover, we showed that the organizational structure that emerged exhibited one of the most common features of real-world exchange mechanisms - their monetary structure. That is, whenever a stable network of shops emerged that supported a stable pattern of exchange activities, one of the tradeable objects would always emerge as a universal medium of exchange, being traded in every shop and involved in every act of exchange. The fact that this model can generate what we believe to be the most salient features of real-world coordination mechanisms gives us confidence that we can potentially use it as a platform with which to study various issues that interact with the coordination problem. What follows in this section is a sketch of that model and a summary of how it works.

The basic idea underlying the model is that trade is a useful but costly activity. People are motivated primarily by the goal of maximizing consumption, and they exhibit a tendency to wander and encounter each other, as in any random matching model. From time to time it occurs to someone that there are gains from trading with a randomly encountered stranger.

However, the probability of meeting a stranger with whom there is a double coincidence of wants is so small, and the probability of meeting that stranger again in the absence of any institutionalized trading arrangements is so small, that a negligible amount of trade takes place until someone gets the idea to set up a trading facility that can easily be located, along with an idea for how to operate that facility. Sometimes such an idea does occur to someone, and this person (entrepreneur) perceives that acting on the idea might be more profitable than continuing to wander, because she could set a spread between her buying and selling prices. People on each side of the market she would be creating might willingly pay for this spread in order to secure a reliable source of consumption rather than continuing to wander about looking for a random trading partner. As such facilities start to open, trading patterns start to form.

The economy has N perishable commodities and a discrete number of transactors, each one being ex ante identical except for type. A type (i, j) transactor is endowed with commodity i and can consume only commodity j. There is the same number of each type for each ordered pair of distinct commodities. The model focuses on five activities: entrepreneurship, search, exchange, business failure and price setting. The computer program (written in C++) to implement the model goes through a sequence of periods, in each of which these activities take place in the following sequence.

Entrepreneurship

Each transactor in turn has a probability μ of receiving an innovation - an idea for setting up and operating a shop. If the transactor is of type (i, j) then the shop will be one where commodities i and j can be traded for each other. There is however a fixed cost of setting the shop up, so before doing so the potential shopkeeper surveys a number of other transactors and inquires whether they would patronize the shop if it opened at prices determined according to the full-cost formula described below. If she finds any interest at all on either side of the market she will pay the setup cost and open the shop. This "market research" allows a form of simulated annealing (see Sargent, 1993 for a simple description) with an endogenous degree of cooling. This is important because it allows for lots of innovation to shake up the system when there are plenty of gains from trade still unexploited but not to keep disturbing a system that has already exploited most of the gains.

The full-cost rule for price-setting takes into account the fixed setup cost, as well as a fixed operating (overhead) cost, so it will depend on the amount of business the shopkeeper expects to be doing. After the shop has become established, these expectations will be revised in the light of experience, but at first the entrepreneur simply picks a number based on a parameter $xMax$ representing "animal spirits". More specifically, the initial estimate of how much will be delivered to the shop of each commodity it trades is chosen from a uniform distribution on the interval from 1 to $xMax$. The higher this initial guess the more she will offer in exchange for quantities delivered to the shop.

Search

Each transactor in turn visits one location and encounters a small sample of other transactors. If she finds a shop on that location or finds someone else who has formed a trading relationship with a shop, she learns of those shops, the commodities they trade, and their currently posted prices. At this point she has a sample of shops which include the ones she has just learned of and the ones with which she already has a relationship. From this sample she can choose to have a relationship with at most two shops. Each relationship will last until the shop exits or the transactor herself chooses to sever the relationship to form a different one. In choosing which relationships to form she chooses the shop, or pair of shops, that would maximize attainable consumption at currently posted prices

Exchange

At this stage, each person in turn can visit the shops with which she has a trading relationship. A type (i, j) transactor having a trading relationship with a shop trading i and j will deliver her endowment of i to that shop (everyone has one unit of endowment per period) for p_i units of j, where p_i is the shop's offer price for i. Alternatively if she has a trading relationships with two shops, one trading i for some third commodity c and the other trading c for j, she will deliver her i to the first shop and then deliver the proceeds to the second shop, thus allowing her to consume $p_i p_c$ where p_i is the first shop's offer price for i and p_c is the second shop's offer price for c. In this paper we evaded the stockout problem by assuming that the shopkeeper was capable of engaging in negative consumption when the amount of any commodity demanded by customers plus the amount needed to defray operating costs exceeded the amount delivered by suppliers. (This assumption captures, as best we could do in a model with no durable commodities, the idea that one of the services provided by specialist traders is the availability of ready stocks.)

Business failure

At this point anyone who has set up a shop in the past has a chance to exit, and thereby avoid having to pay the fixed operating cost of a shop. Each shop will follow a rule that says exit with probability θ if its operating surplus in either of the commodities it trades (deliveries minus the quantity paid out for deliveries of the other commodity minus the amount needed to defray the fixed operating cost) is negative; otherwise stay in business with certainty.

Price setting

Each shop that remains in business now sets its prices for the following period. The rule it follows for price setting is a variant on full-cost pricing. First estimate the quantity that will be delivered of each of its traded commodities, and then set prices that would allow it just to break even in each commodity; that is, to have an operating surplus just large enough to provide what the shopkeeper deems an appropriate compensation for having incurred the setup cost of establishing the shop. The estimation of deliveries is done using simple adaptive expectations.

2.1 Emergence of organization

Clower and I showed that, provided that animal spirits are not too large and that the fixed setup and operating costs are not too large, the model has several absorbing states - that is, arrays of shops, prices, trading relationships and shopkeeper expectations that would persist indefinitely once established. One such absorbing state is a "barter steady state", in which there are $n \cdot (n-1)/2$ shops, one for each unordered pair of distinct commodities, and in which each person who is not a shopkeeper trades her endowment directly for her consumption good each period. Another set of absorbing states consists of "monetary steady states," in which one commodity c has emerged as a universal medium of exchange and there are $n-1$ shops, one for each of the other commodities, trading the other commodity for c, each person that consumes or is endowed with c trades directly with one shop each period, and each other person trades with two shops each period using c as a medium of exchange.

A monetary stationary state is much like the monetary equilibrium studied by Starr and Stinchcombe (1998, 1999) in their version of the Shapley-Shubik trading post model, It constitutes a Pareto efficient allocation of resources given that all trading must occur through shops. Aggregate GDP in the economy is total consumption. Capacity GDP is the sum of all endowments minus the operating costs of the shops. The smallest number of shops consistent with everyone trading is $n-1$, so that capacity GDP equals $N-(n-1)f$, where f is the operating cost of each shop. This is achieved in an equilibrium because all endowments are delivered to a shop and either used to pay the operating cost, or paid out to a customer who consumes it, or consumed by the shop's owner to defray his setup cost.

We ran experiments on the model by simulating it repeatedly, starting from an initial situation of autarky; that is, a situation in which no shops exist and hence no trading relationships exist. We found that very often the economy converged to a stable situation in which every agent was either a shopkeeper or else had a profitable trading relationship with one or two shops. Moreover, when it did converge, it always converged to a monetary stationary state, except in the limiting case where the cost of operating a shop was zero.

Our explanation for the emergence of this monetary structure is based on the network externality created by the fixed costs of shops. During the early stages of a simulation, by chance one commodity (say "wheat") will come to be traded in enough shops that the survival chances of such shops is much greater than those of shops not trading this commodity. This is because a shop that opens trading apples for wheat can attract not just people with a double coincidence of wants - making wheat and consuming apples or vice versa - but also anyone with a single coincidence who can engage in indirect exchange - someone making apples and having a relationship already with a shop trading wheat for her consumption good or consuming apples and having a relationship already with a shop trading wheat for her endowment good. Attracting more customers makes the shop more likely to survive because it makes it easier to cover the fixed operating cost of the shop. Thus once a "standard" medium of exchange has emerged randomly it will tend to snowball; the more shops trade wheat the greater the survival probability of a wheat-trading shop compared to one that does not trade wheat.

3 The multiplier process

The Keynesian multiplier process is an example of a positive feedback loop, or what Axel has called a deviation-amplifying process, in which an initial departure from full-employment equilibrium cumulates instead of being corrected. The existence of some such positive feedback loop in actual economies is attested to by the typical hump-shaped impulse response pattern of GDP to a random shock in estimated times-series models. For example, Chari, Kehoe and McGrattan (2000) report that quarterly movements in the log of detrended US GDP are well approximated by the following AR2 process:

$$y_t = 1.30 y_{t-1} - 0.38 y_{t-2} \tag{1}$$

according to which a negative shock that reduces GDP by 1 percent this quarter is expected to reduced it by 1.3 percent next quarter, and by 1.31 percent the following quarter.

As originally formulated by Kahn and Keynes, and as described in most undergraduate textbooks, the multiplier process involves a coordination problem arising from non-price interactions between decentralized transactors. In a world of perfect price flexibility, a drop in planned spending would cause wages and prices to adjust instantaneously so as to keep aggregate demand fully coordinated with productive capacity. But when prices are slow to adjust, one person's drop in spending causes a drop in other people's incomes, causing a drop in their spending, and so on, resulting in a cumulative increase in the gap between demand and capacity.

The theoretical foundation of this multiplier process is still not well understood. Clower (1965) showed how such a process could arise in a Walrasian general equilibrium setting if price adjustment takes place in real transaction time; when labor is in excess supply, unemployed workers will not present their notional consumption demands to the auctioneer but will instead present demands that are constrained by realized sales income. These ideas were pursued at length in the literature on disequilibrium analysis that followed Clower's original contribution and culminated in the book by Barro and Grossman (1976). But this literature raised more questions than it answered, largely because it offered no explicit account of a decentralized market economy's coordination mechanisms. Instead, it modeled price adjustment as if it takes place just the same as in the idealized world of Walrasian eonomics, where it is led by a fictitious centralized auctioneer, and supposed that while the auctioneer is groping towards equilibrium, transactors are constrained to trade according to rationing rules that are imposed from outside the system by a process that was never even discussed.

One of the supposed advantages of the rational-expectations-equilibrium approach that quickly displaced disequilibrium analysis from its dominant position on the frontiers of macroeconomic theory in the early 1970s was that it did not have to deal with the thorny details of disequilibrium adjustment. Instead it was based on the premise that one can restrict attention exclusively to equilibrium states, in which everyone's beliefs and actions have somehow been coordinated with the beliefs and actions of everyone else. But by adopting this premise, the approach has taken the coordination problem out of macroeconomics, and has denied the very existence of the Keynesian multiplier process, a process which has to do with disequilibrium adjustment rather than with equilibrium behavior (see Leijonhufvud, 1968 and Patinkin, 1976).

In Howitt (2006a) I re-examined the foundations of the multiplier process making use of the same agent-based model of the coordination mechanism that I have just described. This paper investigated the real-time dynamics of the model in the face of disturbances, under the assumption that a monetary stationary state has already been reached and with a particular commodity (again, "wheat") having been established for long enough as the economy's medium of exchange that no entrepreneur ever considers opening a shop that does not trade wheat. I showed that these dynamics contain within them a multiplier process that produces a hump-shaped impulse-response pattern very similar to that of equation (1) that characterizes the US economy.

The multiplier process takes place because of an institutional factor not usually considered in the macroeconomics literature, namely the exit of trading facilities. A shock that disrupts normal trading relationships can cause some of the businesses that coordinate trades to fail, inducing people who formerly had employment (supplier) relationships with those businesses to curtail their expenditures for lack of money, which forces closing of other businesses in cumulative fashion.

In this paper I again simulated the model many times, this time starting from an initial position of a stationary monetary equilibrium and disturbing it in period 1 by a shock that induces some fraction of the population to switch from eating one good to another. To preserve the aggregate structure I supposed that the total number of each type remains constant, so that for every i-eater that becomes a j-eater there is a j-eater that switches to

i. At the time of this shock, each switcher is suddenly without a consumption-shop ("store"), and her former store loses a customer. The switcher may continue to sell her manna to her employer (endowment-shop) but she does not spend her wages. GDP falls because of the reduced goods consumption of the switchers that no longer show up to their former stores, and because of the reduced wheat consumption of the entrepreneurs whose operating surplus in wheat suddenly falls.

Because their revenues have fallen, the former stores of switchers will reduce both their wages (offer price for the non-wheat commodity) and their retail prices (the inverse of their offer price for wheat). The fall in wages will help to offset their profit shortfall, but it will spread the shortfall to other shops, some of whose customers will now deliver less wheat because their wages have fallen. Meanwhile, the fall in wages and prices will do little by itself to raise GDP, which will stay below capacity until the switchers find new stores.

During this process, the luck of the draw may result in particularly large shortfalls for some shops. A shop whose wheat surplus has fallen below zero will be at risk of failure. If that happens then all of the former suppliers of the failed shops will be without an employer, and their sudden drop of wage income will result in a sudden drop in revenues to their respective stores, who may also now become at risk of failure. In this way, the collapse of shops can be self-reinforcing, leading to a cumulative fall in GDP as in the more familiar multiplier process of textbook macroeconomics.

Of course whenever a shop fails, new entrepreneurs will start entering, and employment relations will start to form again. But because of fixed costs, and because a lot of firms may enter the same market, there will be a "shakeout period" which not all new entrants will survive. Thus the process of shop failures is likely to continue for some time before a new stable pattern of shops re-emerges and the economy begins to recover from the cumulative downturn.

4 Upgrading the model

There are many features of the agent-based model described above that make it difficult to take to real data. The latest version (available at http://www.econ.brown/fac/Peter_Howitt) differs from the above in the following ways: First, to recognize that loss of a trading relation with one's employer (i.e., unemployment) is typically a more significant event than loss of a trading relationship with a store, I allow everyone to have two distinct consumption goods, and a CES utility function over those two goods.

Next, I assume that all goods are perfectly storable instead of perishable. This allows me to deal directly with the stockout issue instead of evading it through the device of negative consumption by shopkeepers. In this model shopkeepers target a level of stocks equal to some multiple of target sales, and they execute all buy orders if their stocks permit. The fact that not all orders might be executed introduces a new consideration into peoples' search decisions. Specifically, an agent might have to choose between an employer offering a low wage and another offering a high wage but who has not been executing all sell orders, or between a store offering a high retail price and another offering a lower price but who has not been executing all buy orders. To deal with this complication I assume that each agent keeps track of her employer's "effective" wage and of her stores' "effective" inverse-retail prices, where in each case the effective price equals the actual price multiplied by the fraction of the agent's last order that was executed at that shop. During the search process all choices are made on the basis of effective rather than actual prices.

To make the monetary structure of the model look a little more like that of a modern real-world economy I suppose that the money commodity is a pure token, being neither produced nor consumed by the private agents. There is now a government sector regulating the stock of fiat money. The government also issues perpetual bonds, paying one dollar per period, whose price is P_b, and levies taxes, at a proportional ad valorem rate τ on all purchases from a shopkeeper. There is still no private debt, but people can go to a bond market every period, before trading with shops begins, to trade money for bonds. The government regulates the nominal interest-rate $r = 1/P_b$ according to a Taylor-like rule that makes P_b adjust gradually towards a target value which is an increasing function of the current level of nominal GDP. The government also regulates the tax rate according to a fiscal rule that makes τ adjust gradually towards a target value which is an increasing function of the outstanding value of government debt and of the current rate of change of government debt.

Because there are durable assets, there is now a non-trivial consumption/saving decision to be made. Accordingly I suppose that people spend a fixed fraction of current wealth on consumption each period, as they would in a conventional lifetime-utility maximization model with logarithmic preferences, where the fixed fraction would be the rate of time preference. This expenditure is allocated between the two consumption goods in such a way as to maximize the current period CES utility function if the agent has relations with a shop for each good, or else allocated all to the consumption good traded in the store with which she does have a relationship. In this consumption function, current wealth is calculated as the current market value of bond holdings plus a fixed fraction of the capitalized value of permanent income, discounted at the current nominal rate of interest.

Because they can save in two different forms (money and bonds) people also now have a non-trivial portfolio allocation problem to solve. Here I adopt a buffer-stock approach to the demand for money, following Laidler (1984), and implementing it in much the same way as Akerlof and Milbourne (1980). In this approach, people allow money holdings to absorb unexpected shocks to income and expenditure, thereby economizing on frequent small portfolio adjustments. Specifically, each agent has a target money stock, equal to 4 times current planned expenditure (I think of a period as a week). When current money-holdings exceed twice this target the agent visits the bond market and spends the excess on bonds. When current money holdings fall below current planned expenditure the agent visits the bond market and sells enough bonds to restore money holdings to the target; if not enough bonds are being held to do this, all bonds are sold. If the agent still does not have enough cash to pay for current planned expenditure, all money is allocated to current expenditure.

In all other cases the agent stays away from the bond market.

Instead of the full-cost pricing assumption of the earlier version I treat wages and retail prices separately. For wages, I take seriously the concern for fairness emphasized by writers like Akerlof and Yellen (1990) and Bewley (1999) by supposing that whenever an employer's wages fall below 80 percent of the economy-wide average there is a confrontation between the employers and its work force that results in an increase all the way up to the economy-wide average. Likewise if the employer's permanent income rises to 120 percent of the economy-wide average wage her posted wage is raised to the economy-wide average.

On top of these relatively infrequent corrections, each period the wage is adjusted by a percentage amount that is proportional to the percentage gap between target input and its virtual input. Target input is the amount of labor needed to produce the firm's expected sales (formed at first by animal spirits and then by simple adaptive expectations) plus a gradual adjustment towards its target inventory, and virtual input is the number of suppliers with which the shop has a relationship. (Actual input can differ from virtual if the shop runs out of money with which to buy labor).

Prices are set as a fixed markup over current marginal cost, where the latter includes wages and taxes. The markup fraction is chosen by a shop, when it enters, from a uniform distribution that is centered on the markup that would be optimal if firms' demand functions were given by the constant-elasticity function implied by its customers' CES utility function in a symmetrical equilibrium with a constant number of customers. A profit-maximizing equilibrium is thus given a chance to emerge if the economy converges.

In this version, firms are not forced to take on redundant workers during the search process as they were in the earlier versions. That is, a searching agent cannot choose to form a relation with a shop trading her endowment good if that shop's virtual input exceeds already exceeds its target input.

Finally, I allow each relationship to break up with a fixed probability and each firm to exit with a fixed probability even if still profitable, and I set animal spirits high enough that from time to time firms will enter with overly optimistic sales expectations. The effect of all three of these modifications is to keep the system constantly subject to disturbances and to destroy the existence of any absorbing state to the stochastic process generated by the model. (These modifications were deliberately not made in the earlier version, in order to be able to tell with certainty when the computer simulation was converging and what equilibrium it was converging to.)

5 The effects of wage-flexibility

Is increased wage-flexibility stabilizing? This is the question that Keynes raised in Chapter 19 of his (1936) *General Theory*, where he argued that unemployment was not attributable to the failure of wages to respond, and that in fact if wages were more flexible the unemployment problem would be even worse, because of adverse distributional effects and the disruption of debt-deflation. As Axel's (1968) book showed us, this aspect of the economics of Keynes was in sharp contrast to what became called, and for the most part still is called, Keynesian economics, which from Modigliani to Taylor has attributed unemployment to wage and/or price rigidity.

In (1986) I argued that wage flexibility could increase the volatility of real output in a fairly conventional rational-expectations equilibrium model because of the depressing effect on aggregate demand of a reduction in the expected rate of inflation.[4] A fall in demand would give rise to a fall in expected inflation through a Phillips-type relationship, which, under a given money-supply rule, results in a rise in the real rate of interest and hence an amplification of the fall in aggregate demand, as in a conventional IS-LM system where the IS curve depends on the real rate of interest and the LM depends on the nominal rate. The greater the slope of this Phillips relationship, the more the rational expectation of inflation will fall with any given negative shock to aggregate demand, and hence the stronger will be this deviation-amplifying mechanism.

[4] A similar analysis was produced independently by DeLong and Summers (1986).

The analysis assumed a given path for the money supply, as was conventional in macro theory until recently. But this assumption does not correspond to the way monetary policy is actually conducted in most countries, where the rate of interest is the central bank's instrument and the money supply is typically not given much attention. Moreover, under the more realistic alternative assumption of an interest-rate rule the mechanism that I identified in my 1986 paper would not be at work unless the rule violated the famous "Taylor principle" in an expectational sense. This principle requires that any change in expected inflation result in more than a one-for-one change in the rate of interest. Thus if a fall in aggregate demand generated a fall in expected inflation, the interest-rate rule would cause the nominal interest rate to fall by enough so as to cause a *decrease* in the real rate of interest. (If the interest-rate rule also responded directly to the fall in output caused by the demand shock then the real rate of interest would fall by even more.) Thus wage flexibility, by strengthening the reaction of expected inflation to a demand shock, would be unambiguously stabilizing.

The present model can deal with both of these issues. We have assumed that the monetary authority follows a Taylor-like rule in setting interest rates, but there is no one forming expectations of inflation. Instead, the effect of changing the degree of wage flexibility would be on the efficacy of the system's coordination mechanism. And as far as this is concerned, there is clearly a reason for having some flexibility but not too much. That is, without some wage flexibility there would be no way for shops to eliminate excess demands or supplies. On the other hand, as pointed out in the previous section, the propagation mechanism that I identified in my multiplier analysis is not something that is necessarily helped by having more wage flexibility. Indeed, by introducing extraneous movements in relative prices as the economy adjusts to shocks, increased wage flexibility could easily weaken the coordination mechanism, by inducing a lot of extraneous switching of trading relationships in pursuit of transitory wage and price differences, thus perhaps inducing too many business failures.

I have run the model under alternative speeds of wage adjustment, as measured by the coefficient in the wage-adjustment equation of the percentage gap between target and virtual input. The results as shown in the figure below indicate that indeed macroeconomic stability is achieved at an intermediate level of flexibility. This figure shows the median value of the

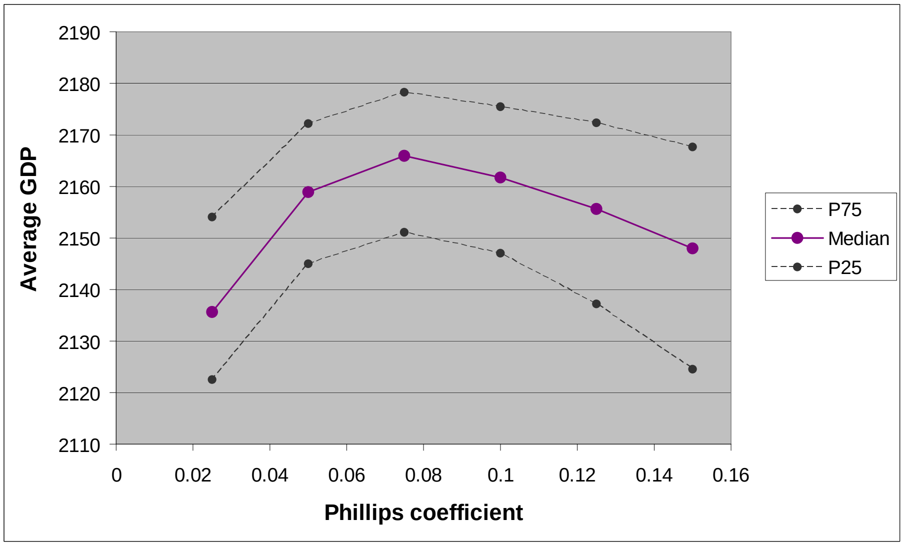

Figure 1: The effects of wage flexibility

average end-of-year GDP when I ran the model for 100 years (5,000 periods), 1000 times for each of 6 different values of the coefficient, along with the interquartile range.[5] So, for example, the P75 line indicates, for each coefficient value, the average end-of-year GDP over the 100 years of the run in which this average was the 250th largest out of the 1000 runs using that coefficient value. The figure indicates that as flexibility goes up, average GDP first rises and then falls.

6 Conclusion

The work I have outlined above is just the beginning of a research agenda aimed at characterizing the behavior of an economic system inhabited by intelligent autonomous agents, a system in which economic activities are coordinated by a self-organizing network of specialist traders. The ultimate goal of this agenda is to understand how, as Axel has said on many

[5]In each simulation there were 50 distinct commodities and 2400 separate transactors.

occasions,[6] a large complex economic system like that of the United States is capable of exhibiting such a high degree of coordination most of the time, and yet is also capable from time to time of departing drastically from such a coordinated state. Of course much work remains to be done. In particular, there are many free parameters in the latest version of the model, and as yet no attempt has been made to calibrate them to actual data. But the results so far indicate that the model is capable of addressing some of the questions that Axel has raised, and which a rational-expectations equilibrium approach is incapable of addressing, such as how a multiplier process can cause cumulative deviations from equilibrium and how the degree of wage flexibility affects the ability of an economy to coordinate activities.

[6] For example, Leijonhufvud (1976, 1981, 1993).

References

Akerlof, George A., and Ross D. Milbourne. "Irving Fisher on His Head II: The Consequences of the Timing of Payments for the Demand for Money." *Quarterly Journal of Economics* 95 (August 1980): 145-57.

Akerlof, George A., and Janet L. Yellen. "The Fair Wage-Effort Hypothesis and Unemployment." *Quarterly Journal of Economics* 105 (May 1990): 255-83.

Barro, Robert J., and Herschel I. Grossman. *Money, Employment and Inflation.* New York: Cambridge University Press, 1976.

Bewley, Truman F. *Why Wages Don't Fall During a Recession.* Cambridge, MA: Harvard University Press, 1999.

Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan. "Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem?" *Econometrica* 68 (September 2000): 1151-79.

Clark, Andy. *Being There: Putting Brain, Body and World Together Again.* Cambridge, MA: MIT Press, 1998.

Clower, Robert W. "The Keynesian Counter-revolution: A Theoretical Appraisal." In *The Theory of Interest Rates*, edited by Frank H. Hahn and Frank P. R. Brechling. London: Macmillan, 1965.

Davidson, Paul. "Keynes and Money." In *Keynes, Money and Monetarism: The Eighth Keynes Seminar held at the University of Kent at Canterbury, 1987* , edited by Roger Hill, 2-26. London: Macmillan, 1989.

De Long, J. Bradford, and Lawrence H. Summers. "Is Increased Price Flexibility Stabilizing?" *American Economic Review* 76 (December 1986): 1031-44.

Evans, George, and Seppo Honkapohja. *Learning and Expectations in Macroeconomics.* Princeton, NJ: Princeton University Press, 2001.

Howitt, Peter. "Wage Flexibility and Employment." *Eastern Economic Journal* 12 (July-September 1986): 237-42.

Howitt, Peter. "The Microfoundations of the Keynesian Multiplier Process." *Journal of Economic Interaction and Coordination* 1 (May 2006a): 33-44.

Howitt, Peter. "Monetary Policy and the Limitations of Economic Knowledge." In *Post Walrasian Macroeconomics: Beyond the Dyanamic Stochastic General Equilibrium Model*, edited by David Colander. New York: Cambridge University Press, 347-67. 2006b.

Howitt, Peter, and Robert Clower. "The Emergence of Economic Organization." *Journal of Economic Behavior and Organization* 41 (January 2000): 55- 84.

Keynes, John Maynard. "Poverty in Plenty: Is The Economic System Self-Adjusting?" In *The Collected Writings of John Maynard Keynes*, 485-92. Vol. XIII. London: Macmillan, 1934. For The Royal Economic Society.

Keynes, John Maynard. *The General Theory of Employment, Interest, and Money.* London: Macmillan, 1936.

Krebs, Tom. "Growth and Welfare Effects of Business Cycles in Economies Idiosyncratic Human Capital Risk." *Review of Economic Dynamics* 6 (2003): 846-68.

Krusell, Per, and Anthony A. Smith Jr. "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy* 106 (October 1998): 867-96.

Laidler, David. "Information, Money and the Macroeconomics of Inflation." *Swedish Journal of Economics* 76 (March 1974): 26-41.

Laidler, David. "The 'Buffer Stock' Notion in Monetary Economics." *Economic Journal* 94 (Supplement 1984): 17-34.

Leijonhufvud, Axel. *On Keynesian Economics and the Economics of Keynes: A Study in Monetary Theory.* New York: Oxford University Press, 1968.

Leijonhufvud, Axel. "Schools, 'Revolutions', and Research Programmes in Economic Theory." In *Method and Appraisal in Economics*, edited by Spiro J. Latsis. Cambridge, UK: Cambridge University Press, 1976.

Leijonhufvud, Axel. *Information and Coordination.* New York: Oxford University Press, 1981.

Leijonhufvud, Axel. "Towards a Not-Too-Rational Macroeconomics." *Southern Economic Journal* 60 (July 1993): 1-13.

Leijonhufvud, Axel. "Agent-Based Macro." In *Handbook of Computational Economics*, edited by Leigh Tesfatsion and Kenneth L. Judd, 1625-37. Vol. 2. Amsterdam: North-Holland, 2006.

Patinkin, Don. *Keynes' Monetary Thought: A Study of Its Development.* Durham, NC: Duke University Press, 1976.

Sargent, Thomas J. *Bounded Rationality in Macroeconomics.* Oxford: Clarendon Press, 1993.

Starr, Ross M., and Maxwell B. Stinchcombe. “Monetary Equilibrium with Pairwise Trade and Transaction Costs.” Unpublished, University of California San Diego, September 3, 1998.

Starr, Ross M., and Maxwell B. Stinchcombe. “Exchange in a Network of Trading Posts.” In *Markets, Information and Uncertainty: Essays in Economic Theory in Honor of Kenneth J. Arrow*, edited by Graciela Chichilnisky, 217-34. Cambridge UK: Cambridge University Press, 1999.

Tesfatsion, Leigh. “Agent Based Computational Economics: A Constructive Approach to Economic Theory.” In *Handbook of Computational Economics*, edited by Leigh Tesfatsion and Kenneth L. Judd, 831-80. Vol. 2. Amsterdam: North-Holland, 2006.