MACROECONOMICS IN A SELF-ORGANIZING ECONOMY*

Quamrul Ashraf
Williams College
Boris Gershman
American University
Peter Howitt
Brown University

This paper emphasizes the importance of considering the mechanisms that coordinate economic transactions in a decentralized economy, namely the role played by a self-organizing network of entrepreneurial trading firms, for theories aimed at guiding macroeconomic policy. We review a research program that aims to understand how, and how well, trading activities are coordinated in various circumstances by employing agent-based computational (ACE) models of stylized economies where these activities take place in a self-organizing network of markets created and operated by profit-seeking business firms. We discuss how such a research program can yield important policy-relevant insights, beyond those that can be offered by conventional dynamic stochastic general equilibrium (DSGE) models, into several macroeconomic phenomena including the emergence of monetary equilibria in a decentralized economy, the microfoundations of the multiplier process, the costs of a higher trend rate of inflation, and the role of the banking system in economic crises.

Keywords: Self-organizing trade networks, Market coordination, Agent-based computational economics, Emergence of monetary equilibria, The multiplier process, Costs of inflation, Bank regulation, Economic crises.

* We thank participants of the REPLHA International Conference, Milan, October 13, 2011, for helpful comments. The research reported on here was partly funded by the NSF.

This paper is largely about macroeconomic theory, but it is motivated by some of the most important policy challenges that we are confronted with in the extraordinary world we are living in today. Dealing effectively with these challenges requires a conceptual framework that focuses on those parts of the economic system that matter most for the question at hand. We are not alone in thinking that such a conceptual framework is not provided by the theories that guide macroeconomic policy in most countries today. What we aim to do in this paper is to describe a line of research that we have been pursuing, one that goes back to joint work between Howitt and his former teacher and co-author, Robert Clower, and is aimed at providing a more appropriate conceptual framework for thinking about some important macroeconomic policy issues.

Our starting point is one of the oldest and most important ideas in economics, going back at least to Adam Smith, namely the idea that a decentralized economic system is self-organizing. It is capable of “spontaneous order,” in the sense that a globally coherent pattern of transactions can result from purely local interactions, without the intervention of a central coordinator. Indeed, like an anthill, a free market economy can organize transactions into patterns that are beyond the comprehension of any of its individual participants. We would like to understand how this self-organization takes place. Specifically, what is the process that coordinates the exchange activities of millions of independent transactors in a decentralized economy?

The reason why these questions are critical for understanding macroeconomic policy is that an economy’s coordination mechanism works better at some times than others. Even Smith and Hayek recognized that the automatic workings of the decentralized economy could sometimes be improved by collective intervention. Consider, for example, the increase in unemployment that takes place during a deep recession. Unemployed workers who used to be employed are just as willing and able to work as before, the fall in aggregate output that accompanies recession has enhanced the scarcity value of the output they could potentially produce if employed, and yet the market for their services has somehow disappeared. The coordination mechanism that had previously allowed them and those with a taste for their output to realize their potential gains from mutually advantageous exchange is no longer allowing them to do this, even though those gains are if anything larger than before.[1] Macroeconomic policy to deal with unemployment thus amounts to fixing a mechanism that has malfunctioned, and a highly complex mechanism at that. And attempts to fix this broken mechanism without first understanding how it is supposed to work normally are likely to be as successful as medieval medicine was in treating bacterial infections.

The main premise of our research has been that the role of coordinating transactions in a decentralized economy is performed, for better or worse, by a self-organizing network of business firms that seek profit from creating and operating the markets through which others transact. To use a phrase that Clower once coined, business firms are the visible fingers of the invisible hand. Economics has a long tradition of regarding exchange as a do-it-yourself affair, in which people with goods and services to sell trade directly with the ultimate demanders of those goods and services. But a little reflection on the experience of daily life is enough to persuade most people that exchange in a market economy is not a do-it-yourself affair. People are not like the actors in a typical monetary search model, who when hungry go wandering aimlessly in hopes of randomly encountering someone with surplus food. They go to a grocery store or to a restaurant. When in search of clothing they visit a tailor or a clothing store. They lend surplus funds through the intermediation of a bank, arrange for long-distance travel by using facilities provided by a travel agent, and so on and so forth. Most of us also sell our labor services to an economic entity, either a private business or a government agency (the latter of which would not exist in a purely decentralized economy) whose primary purpose is

1. Of course there are some macroeconomists who would claim that unemployment can only arise when productivity falls or tastes change toward more leisure, in which case there is no malfunction; it's not that the gains from trade are unexploited but rather that the gains have disappeared. But this is a view that we do not accept. Recessions are not periods of technological regress or of contagious laziness. Although there may be shifts in demand from one sector to another, or changes in the overall level of aggregate demand, there is nothing inherent about such shifts that would imply that the gains from trade have shrunk in all sectors of the economy, yet in the typical recession unemployment rises in all sectors of the economy. Instead, it seems clear to us that some kind of market failure takes place in recessions.

to purchase various such services, organize them into production, and sell the resulting output. So to understand how, and how well, exchange activities are coordinated in a decentralized economy, the first place to look is to this self-organizing network of firms that constitutes the institutional structure through which we all conduct our daily business.

Now some would argue that a good economic theory involves abstraction, and that if we want to model the transactions process of a modern economy we should maybe abstract from business firms, by assuming that people who work for businesses trade their output directly with those having a taste for the workers' output. This is the stance taken by search theoretic models of money, which typically assume that trade takes place directly between ultimate suppliers and ultimate demanders, who meet in random non-repeated encounters without the aid of any intermediary. Presumably the rationale for this way of looking at the transactions process is the same as the rationale for abstracting from money in the theory of value. On the surface, what we see is people trading goods and services for money, but the deeper underlying reality that we see once we pierce the veil of money is that people are ultimately trading goods and services for other goods and services, with money acting only as a device for executing these ultimate exchanges.

The analogy between money and firms is a useful one. But as John Stuart Mill once observed, there is nothing more insignificant than money, except when it goes wrong. By the same token, the fact that people find it convenient to trade through shops rather than directly with one another is perhaps of little significance for understanding the long-run structure of relative prices. But when something goes wrong with the network of firms that people normally rely upon, then abstracting from the existence of such firms is as unhelpful as it would be to ignore money when trying to understand inflation.

Moreover, a good case can be made that, by recognizing that production takes place in firms but not recognizing their role in coordinating transactions by creating and operating markets, we are ignoring their most important activity, as measured by the value of resources they use. Wallis and North (1986), for instance, have shown that more than 50 percent (by value) of the primary resources used up in the course of economic activity in the United

States are devoted not to production activities but to providing transactions services, services that would be of no use to a Robinson Crusoe with no trading partners. The value added in Finance, Insurance, and Real Estate, for example, is typically much more than that of the entire manufacturing sector, as is the value added in Retail and Wholesale Trade. Moreover, much of the input to the manufacturing sector is best construed as being used up in the production of transactions that help people realize gains from trade rather than being used up in transforming inanimate objects. We have in mind the inputs of lawyers, sales people, those engaged in personnel, marketing, and advertising, and so forth, all of whom are undertaking activities whose main purpose is to facilitate transactions.

We made reference at the start of this paper to the theories now guiding macroeconomic policy. We were referring, of course, to the broad class of rational-expectations equilibrium models generally known as DSGE, for dynamic stochastic general equilibrium. DSGE started out 4 decades ago as a reaction against the Keynesian economics that had been the dominant paradigm of the profession in the 1950s and 60s, a reaction that was first expressed by new classical economists, like Lucas and Sargent, and later by real business cycle theorists. The early proponents of DSGE argued that an equilibrium model built on a slightly modified Walrasian conceptual framework, in which markets clear everywhere and always, could account for important short-run as well as long-run macroeconomic phenomena. But soon Keynesian economists began developing their own versions of DSGE, which consisted of rational expectations equilibrium models in which not all prices were perfectly flexible, and by now DSGE has become the dominant paradigm agreed upon by all sides of the great macroeconomic debates.

Of course there are many serious criticisms one might make of DSGE, and many of them have been made in the literature. The criticism we consider most important for present purposes is that existing DSGE models, even those with imperfectly flexible prices, are built on a conceptual foundation that pays little or no attention to the way in which economic transactions are organized. To borrow a phrase from Jevons, they ignore the institutions and processes that make up the mechanism of exchange.

When we examine DSGE models, looking for what else might go wrong with the market mechanisms that coordinate economic transactions, we find that in most of them there is no such mechanism. In models with perfect competition, the setting of prices is left to a mysterious outside agent called the auctioneer, whose behavior is left largely unexplained. But that is just the tip of the iceberg. There is no description of how trades are arranged. Even if we accept that the auctioneer can provide everyone with a price vector such that the sum of desired demands equals total supply for each tradable object, there is no account of how buyers and sellers are matched up with one another and how the trades that people have planned will be executed. When demand and supply are not equal, the theory offers no guidance as to who gets to trade how much and with whom, no indication of how people learn about trading opportunities, about who creates and maintains the shops and other facilities through which they trade, about how bids and offers are transmitted, and so on and so forth.

The canonical model of Woodford (2003), which forms the basis of the estimated New Keynesian DSGE models, now used in central banks around the world, makes less use of the mysterious auctioneer, inasmuch as many prices are set by a given set of monopolistically competitive firms who are explicitly motivated to maximize their shareholders' wealth. So far, so good. But, there is no account of where these price-setting monopolists come from, how they maintain their monopolies against the threat of entry, how people decide to trade with one set of firms rather than another, how firms manage to coordinate with their suppliers and customers, what happens when one of them goes out of business in a recession, and so forth. Instead, all transactors are in continuous touch with each other through the intermediation of these firms, whose existence is merely assumed, and who take care of enough details of the transactions process that the other people in the model are connected only through the market prices that they take as given from the firms. As a result, there is nothing that can go wrong in the transactions process other than some mistake in price-setting.

In essence, these New Keynesian DSGE models are providing the same diagnosis that economists have given for generations; unemployment rises because wages and prices are slow to adjust to shifts in demand and supply. This is the answer provided by classical economists from Hume through Marshall. It is still the answer offered by modern Keynesian economics. Indeed it is now even the answer that has been finally accepted by most proponents of the real-business-cycle school of macroeconomics, who admit the need for wage-and-price stickiness to account for various features of the business cycle.[2]

The problem with this time-honored tradition of blaming wage-and-price stickiness is not that the assumption of stickiness is factually incorrect. On the contrary, the stickiness of wages and prices is one of the most well-documented facts of macroeconomics. Instead, as Leijonhufvud (1968) forcefully pointed out, the problem is that, first, the experience of the Great Depression in the United States shows clearly that the downturn that started in 1929 did not come to an end until wages and prices started to rise, that is, until the reflation that was clearly a deliberate policy move on the part of the Roosevelt administration started to take place. If lack of wage and price flexibility had caused the downturn, then it would have taken deflation rather than reflation to cure the unemployment problem. Second, as Keynes argued in Chapter 19 of the *General Theory*, and as Fisher had already argued in *Debt Deflation Theory of Depressions*, there are many reasons for believing that wage and price flexibility would actually make fluctuations in unemployment larger rather than smaller.

So when unemployment rises in a recession, something has gone wrong with the process by which economic transactions become organized, something that goes beyond the mere stickiness of wages and prices, something that we think can only be discovered by investigating simple stylized models of economies in which trading activities take place, in and out of equilibrium, in a self-organizing network of markets that are created and operated by profit-seeking business firms, and by asking how, and how well, those activities are coordinated in various circumstances. What we would like to do in this paper is to give the reader an idea of what kind of model that research has led us to construct, and why we think this class of models provides a more solid framework for analyzing certain policy questions than does any DSGE model currently in use.

2. See Chari and Kehoe (2006) for example.

It turns out that this research agenda is one for which Agent-Based Computational Economics (ACE) is particularly well suited for two main reasons. First, by endowing each agent with a set of relatively simple adaptive behavioral rules that allow the agent to operate intelligently in an unknown environment, an ACE model gives the system a chance to achieve some semblance of order without giving anyone the kind of systemic knowledge that would allow him to act as a central coordinator. A rational expectations equilibrium might or might not emerge from the interaction of these rules. If it does, then we have discovered something about at least one possible mechanism that produces that kind of spontaneous order whereas, if the system fails to approximate a rational expectations equilibrium, we will have discovered something about the conditions under which a spontaneous order is likely to require some kind of collective intervention. The second reason for using ACE is that models of spillovers between multiple markets that are not in supply-demand equilibrium are notoriously difficult to analyze. The attempts by Barro and Grossman, Benassy, Malinvaud, and others in the 1970s to understand what some called "general disequilibrium analysis," building on the original contributions of Clower and Patinkin, made little progress largely because the problem quickly became analytically intractable. ACE can deal with this kind of intractability by substituting simulation and Monte-Carlo results for unattainable analytical results.

1. Self-organization of trading firms

At the heart of all our work is a parable concerning the spontaneous emergence of a more-or-less self-regulating network of markets operated by profit-seeking business firms. The details of this parable were first laid out in the form of an ACE model by Howitt and Clower (2000). The rest of this section briefly describes the Howitt-Clower model of a self-organizing economy.

Time comes in discrete "weeks" indexed by t. There are n non-storable goods and m households. In the computer simulations, $n = 10$ and $m = 2160$. Each household is of a type (i, j), where i and j are distinct goods; such an agent receives a weekly endowment of one unit of good i and wishes to consume only good j. There is a symmetric distribution of types in the population, with the same number b (= 25 in the simulations) of each type of household.

It is assumed that households can trade only through organized trading facilities called "shops." A new shop is created whenever a household chooses to act on a random opportunity to enter business (*i.e.*, to become an "entrepreneur"). Each shop trades only two goods, which we assume must be the endowment and consumption goods of the "owner" (the entrepreneur who opened the shop). Once opened, a shop of type (i, j) will make weekly postings of two offer prices—an amount of i offered for each unit of j delivered by anyone who chooses to do so, and an amount of j offered for each unit of i delivered. Those households seeking to trade any good for which they do not own a shop must form a "trading relationship" with such a shop. They form such relationships by searching, both directly through the space in which the shops are located and indirectly by querying randomly met people about the shops with which they have relationships.

Each week the computer simulation takes the households through a fixed routine involving 5 stages, each of which represents what we take to be an important aspect of the trading process. First, there is an entry process in which a random selection of households is faced with an opportunity to become an entrepreneur. There is a fixed setup cost of opening the shop, and there will be a fixed weekly overhead cost of maintaining the shop in operation if opened. So before deciding whether or not to open, the potential entrepreneur will conduct "market research" by querying two other households—one that might want to trade i for j and one that might want to trade j for i, asking each if they would choose to form a trading relationship with the proposed shop if it were to open, at the posted prices that the entrepreneur has decided on (more on these prices below). Both households will answer the query using the same criterion (to be described below) they will use when searching for better trading opportunities in the next stage of the program. If both answer affirmatively, then the potential entrepreneur will indeed open. Otherwise, the opportunity will be allowed to lapse.

In choosing his shop's prices, the potential entrepreneur uses a simple form of full-cost markup pricing. First, an estimate is formed of how much of each good will be delivered by the shop's suppliers/ customers. Each of these initial estimates is taken from a uniform distribution between *1* and *X*, where *X* represents the state of "animal spirits." The entrepreneur then calculates, given the fixed costs of operation and the margin required to compensate for the sunk cost of setting up the shop, the combination of offer prices that would just allow the shop to break even in terms of each of the goods it trades if its estimates were correct. The bigger the delivery estimates, the bigger those breakeven offer prices will be because of the economy of scale implicit in the fixed costs. These breakeven prices are the prices the shop will post.

The second stage of weekly activities is one in which households search for trading relationships. Each household can have only two relationships at a time—one with a shop (outlet) trading the household's endowment good and the other with a shop (source) trading its consumption good. In some cases, the same shop can serve as both source and outlet (double coincidence). The household wants to maximize weekly consumption. In the double coincidence case, weekly consumption good will be the shop's offer price for the household's endowment good. Otherwise, if the household has a source and outlet both trading the same complementary good, then it can engage each week in indirect exchange using the common complementary good as an exchange intermediary, and weekly consumption will be the product of the outlet's offer price for the endowment good and the source's offer price for the complementary good. In all other cases, weekly consumption will be zero. During the search stage, a household always forms a new relationship with any shop that would allow weekly consumption to be raised at currently posted prices. Whenever a new source or outlet is chosen, the relationship with the old source or outlet must be severed.

The third stage of the Howitt-Clower model is a trading stage, in which each household, in random order, visits the shops with which it has a relationship, first delivering its endowment to the outlet and then, if possible, using the entire sales proceeds to buy the consumption good from the source. During this stage it is assumed that all planned trades can be executed by the shops, regardless of the shop's inventory position.

The fourth stage of this process is the exit stage. During the course of trading, a shop trading good *j* will have had an amount *x* of good *j* delivered to it and will have had to pay out a certain amount of good *j* to its suppliers of the other good it trades. In addition, it will have to pay out some fixed amount to cover its overhead cost. The remainder is available for the owner's consumption. If this is negative, then the owner will have had to engage in home production of the good just to stay in business. This home production, or negative consumption, represents a loss that the shopowner does not want to incur indefinitely. Any shop that has incurred a loss in either of the goods it trades will exit at this stage with some fixed probability.

The fifth and final stage of weekly activity is where expectations and prices are updated by each shop. Delivery expectations are adjusted using the simplest form of adaptive expectations. That is, the expectation for each delivery amount is adjusted by some fixed fraction of the gap between what was actually delivered this week and what had been expected this week. Prices are then updated using the same full-cost markup procedure used when the shop opened, but with the newly revised delivery estimates.

The model sketched above can be described as a stochastic process. Provided that animal spirits are not too large and that fixed costs are not too large, the process will possess several absorbing states. These absorbing states correspond to stable shop networks that provide a coherent pattern of trading activity throughout the system.

One such absorbing state is a stationary barter equilibrium, in which there exists in operation exactly one shop for each unordered (*i, j*) pair. Each household of type (*i, j*) either owns the shop of that type or has a trading relationship with it, and delivers its unit endowment to it each week in exchange for consumption. Delivery estimates of each shop equal *b*, the number of households of each type, and prices are constant at the values that allow the shops to break even at those estimates.

Other absorbing states are stationary monetary equilibria, in which one good *k* has emerged as a universal medium of exchange, being traded in each shop, and in which there is exactly one shop trading each other good. Every non-shopowner who is not endowed with *k* and does not consume *k* engages in indirect exchange using *k* as an exchange intermediary. The others are able to consume using just one shop. Delivery estimates have adjusted to equal actual amounts, and prices are constant.

In order for both barter and monetary absorbing states to exist, it is necessary that animal spirits not be too large. This ensures that the state is never disturbed by new entrants who pass the market research test during entry because their delivery estimates are so optimistic that they can undercut the prices charged by incumbents.

Howitt and Clower show that the monetary equilibrium in which fixed costs are lowest (a shop's fixed costs are assumed to vary according to the identity of the goods traded in the shop) allows the maximal feasible total consumption. As in the barter stationary equilibrium, since each household has an outlet, all endowment is used either for consumption or for defraying a shop's fixed costs. The total number of fixed costs is minimized in a monetary equilibrium since it uses only n-1 shops as opposed to the $n(n-1)/2$ shops needed to support the barter equilibrium, and using the least cost exchange intermediary obviously allows the most to be consumed of the money commodity.

Howitt and Clower show that this model is capable of self-organization under a wide set of parameter values. Specifically, they report the results of 6,000 simulation runs. Each run starts in autarky, with no shops and hence no trading relationships in place, and continues for 20,000 weeks or until a monetary equilibrium has been reached. They find that in almost all runs an absorbing state is reached unless fixed costs were set too high or too low.[3] Furthermore, they find that the only absorbing states that ever emerge are monetary equilibria. This latter result arises from the "network externality" implicit in the above sketch. That is, once a few shops trading the same complementary good have emerged, the survival chances of other shops are greatly enhanced if they also trade that complementary good since this allows them to attract more suppliers/customers and hence makes it easier for them to defray the fixed cost while offering competitive prices. The fact that the same model that was capable of "growing" market organization also happened to exhibit a feature of all economies of record that orthodox theories have trouble accommodating without artificial

3. The problem with low fixed costs appears to be that they weaken the network externality that helps to promote the achievement of the efficient monetary equilibria.

contrivances, namely monetary exchange, adds to the plausibility of this model's account of self-organization.

2. The multiplier process

One implication of the above account of self-organization is that there is more that can go wrong with the trading process than just having disequilibrium prices. In particular, the failure of a business firm will always disrupt the economy at least locally for some time, and there is nothing much that speedy price adjustment can do to compensate for this shock because the disappearance of a shop constitutes a loss of organizational capital, which can only be rectified by the successful entry of a new replacement shop regardless of the prices that are charged by the surviving shops.

Indeed, the random nature of the entry process can lead to a shakeout period that makes failure more likely for other firms, causing a cumulative contraction or multiplier process in the economy. This is because a supplier of shop A that disappears will also be the customer of some other shop B. If several of those suppliers suddenly lose their source of income, then shop B will be faced with a drop in demand that threatens its existence. If it was on the verge of making losses, then this might be enough to put it under as well.

Howitt (2006) showed more systematically how this can happen in a slightly modified version of the Howitt-Clower model. In this modified model, it was assumed that a convention had already been reached that good *1* was the universal medium of exchange. Thus an entrepreneur of type (*i, j*) can only open a shop of type (*i, 1*). The model was started in a monetary equilibrium and subjected to a reallocative shock of the sort that occurs when people reduce their demands for some products, without immediately signaling to anyone what they are planning to demand instead of these products. This is the classic coordination problem that Keynes wrestled with. Consumers may decide to spend less than their income, but this does not amount to a specific demand for future consumption. Instead, their future demands remain latent, and entrepreneurs must somehow discover them through trial and error. Likewise, unemployed workers' notional demands remain undiscovered until some entrepreneurs find it in their interest to employ the workers and thereby provide them with the means of making their demands effective.

To portray such a shock in the above system, Howitt supposed that at a certain date some fraction of the population switches from consuming one good to another. To preserve the aggregate structure, he supposed that the total number of each type remains constant, so that for every *i*-consumer that becomes a *j*-consumer there is a *j*-consumer that switches to *i*. At the time of this shock, each switcher is suddenly without a source, and his former source loses a customer. The switcher may continue to sell his endowment to his outlet but he does not spend his sales proceeds. GDP falls because of the reduced goods consumption of the switchers that no longer show up to their former sources, and because of the reduced good-1 consumption of the entrepreneurs whose operating surplus in good *1* suddenly falls.

Figure 1 below shows the average impulse response, over 10,000 runs, of aggregate real GDP relative to the full-capacity level achieved in the initial equilibrium, in the case where 12 percent of the population made the switch. The blue curve in this figure represents the actual impulse-response of HP-filtered GDP under the univariate AR(2) model estimated by Chari, Kehoe, and McGrattan (2000), assuming 50 weeks per year. As it turns out, this simple model, where deviation amplification comes just from the cumulative process of shop failures, does a good job of tracking the actual hump-shaped impulse response pattern of US GDP!

Howitt verified that this hump-shaped impulse response pattern arises from the cumulative process of shop failures by showing that if shop exits were eliminated (by modifying the code to set the exit probability of a loss-making firm equal to zero), then the economy always snapped back into equilibrium within a short number of weeks. Moreover, as Figure 2 demonstrates, there was a strong positive correlation between the size of the displacement in GDP and the number of shop failures in the 5 years following the shock. He also demonstrated that allowing for greater wage-and-price flexibility by having a higher speed of adaptation in delivery estimates did nothing systematically to reduce the amplitude of the impulse response.

Figure 1. Impulse response of real GDP to a 12 percent shock

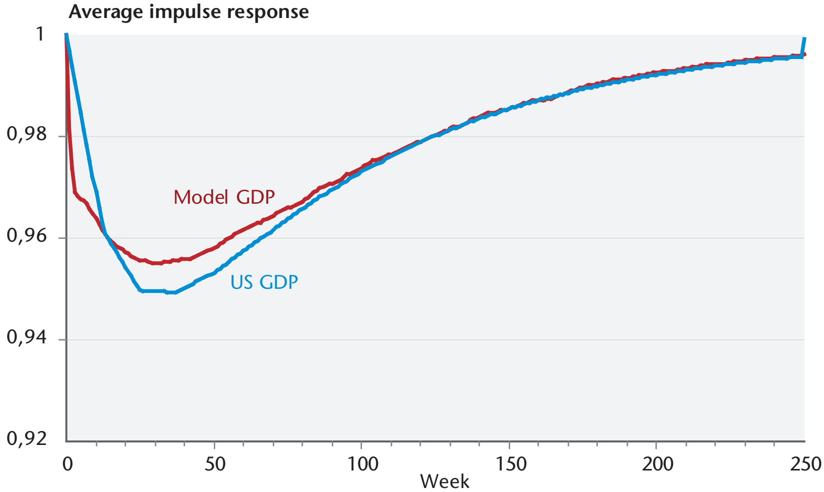

Source: Howitt, 2006.

Figure 2. Maximum GDP gap and shop failures following the shocks

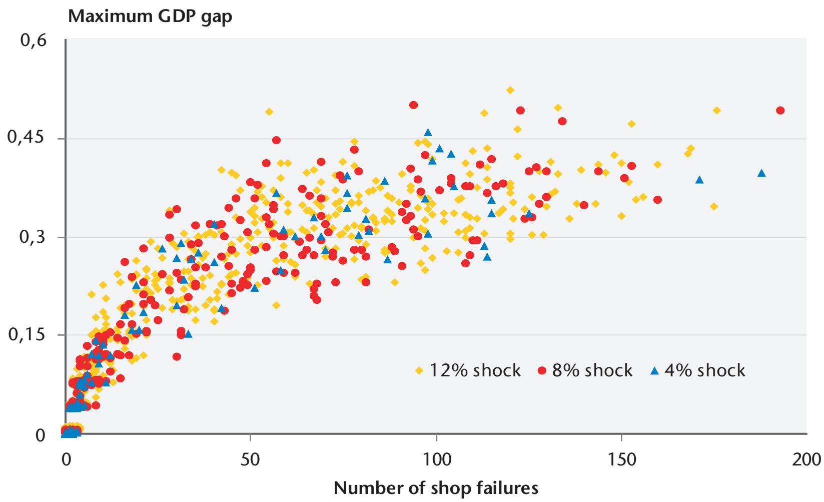

Source: Howitt, 2006.

3. Costs of inflation

Ashraf, Gershman, and Howitt (2012) have built a somewhat less stylized version of the Howitt-Clower model to analyze one of the important policy issues that has proven particularly intractable for orthodox theory. Specifically, they address the question of the extent to which an economy's macro performance is enhanced by having a lower trend rate of inflation.

According to conventional theory, the answer to this question is "not much." The case for low inflation in modern macro theory comes from various DSGE studies that have confirmed the optimality of Friedman's rule, which is to reduce inflation to the point where the nominal rate of interest equals zero. But the saving that would arise in principle from following this rule consists of the elimination of a tax on non-interest-bearing money holdings, a saving that almost all published research estimates to be a trivial fraction of GDP because the base of this tax is just a tiny fraction of total wealth in any advanced economy.

New Keynesian DSGE models, in which money as a means of exchange and store of value plays no essential role, offer another possible reason for targeting low inflation, namely the inefficiency that comes from having a wider dispersion of relative prices for no reason other than the fact that different sellers are at different stages of the price-change cycle; those with more recent price changes will tend to have higher relative prices because they have made the most recent adjustment to inflation. In these models, the optimal trend rate of inflation is clearly zero, except possibly for second-best public finance reasons (Phelps, 1972) or risk-sharing considerations (Levine, 1991) that might argue for a positive rate.

Howitt and Milionis (2007) show that this argument is especially dependent on the Calvo pricing model that everyone agrees is particularly unconvincing. Specifically, in a conventionally calibrated model, once the trend inflation rate reaches 10 percent, over 35 percent of aggregate output is produced by the 0.3 percent of firms that are selling at a price below marginal cost! These firms would certainly want to either raise their price or curtail production if it were not for the fact that they have not recently been visited by the Calvo fairy, but the model requires them anyway to produce however much is demanded at their obsolete prices. Replacing the

Calvo model by a Taylor model with as much as a 7 quarter lag between price changes gets rid of this counterintuitive feature of the model and has no firms selling below marginal cost, but it also reduces the cost of a 10 percent rate of inflation to about 1.5 percent of aggregate consumption. Moreover, if one keeps the assumption of Calvo pricing but reinserts lagged inflation in the Phillips Curve, as central bank DSGE models typically do, by invoking the usual indexation story that price setters not visited by the Calvo fairy adjust their prices as a function of lagged inflation, then the cost of inflation in DSGE models is almost entirely eliminated, because indexation greatly reduces the extent to which inflation raises price dispersion (Billi, 2011).

Despite the failure of conventional theory to account for significant costs, central bankers around the world continue to attach the highest priority to maintaining very low inflation. Before concluding that central bankers are being wrongheaded, we believe that one needs to explore non-conventional theoretical reasons why inflation might be costly. In particular, we have used a much-expanded version of the Howitt-Clower model to explore the suggestion by Heymann and Leijonhufvud (1995) to the effect that inflation impedes the coordination mechanism that Howitt and Clower focus on and that conventional theory takes as functioning perfectly at all times at no cost.

In particular, the results of the preceding section concerning the cumulative process of shop failures suggest an unconventional mechanism through which inflation might really matter for macro performance. Specifically, the higher is the rate of inflation the more difficult it is for the firms that operate markets to remain in business, because of the well-known tendency of inflation to induce noise into the price system. Thus, an environment with higher inflation is likely to have a higher incidence of such cumulative contractions and, hence, a worse overall macro performance.

The changes made by Ashraf, Gershman, and Howitt (2012) include making the goods durable, allowing each household to have two consumption goods, introducing fiat money instead of commodity money, having staggered price setting (but making it state dependent instead of having a Calvo-type Poisson process delivering price change opportunities), having government bonds and a central bank that conducts open market operations using a Taylor rule, having a retail sales tax at a rate adjusted annually to pay the interest on the government debt, and having continual shocks and high enough animal spirits that there is no absorbing state for the economy. Moreover, parameter values were (roughly) calibrated to US economic data. In this model, peoples' endowments of any good *j* are interpreted as a type of labor services capable of producing good *j*, so that people without an outlet can be interpreted as unemployed.

The trend rate of inflation in the model is the inflation target implicit in the central bank's Taylor rule. In the baseline calibration, this inflation target was set equal to 3 percent, and the economy achieved almost exactly 3 percent inflation on average across all runs and all years. To address the issue of the macro consequences of trend inflation, the paper simulated the artificial economy 10,000 times for each integer value of the inflation target from 0 to 10 percent. The main results are depicted in Figure 3 below. As shown in this figure, the median performance of the economy deteriorates steadily by all reported measures when trend inflation rises above 3 percent. It also shows that this deterioration is highly significant in economic terms when trend inflation reaches the 10 percent level. We know of no conventional analysis that provides such powerful support for the idea that a central bank can improve an economy's performance simply by choosing a low inflation target.

Ashraf, Gershman, and Howitt use the model to explore the reasons for this effect of inflation, and find strong evidence that the above mentioned link between inflation, price dispersion, and shop failure was at work. In particular, as Figure 4 shows, increases in trend inflation produce monotonic increases in price dispersion and monotonic decreases in the median number of shops in existence.

They also show that the effects work even if the zero lower bound on nominal interest rates is suppressed and even if efforts are made to take into account the Lucas critique by allowing critical parameters in peoples' decision rules, like the markup parameter in price setting, to vary systematically with the inflation rate. In short, it seems that, according to this particular ACE model, inflation does create a big macroeconomic cost by impairing the self-organizing capacity of the economic system.

Figure 3. Economic performance and target inflation

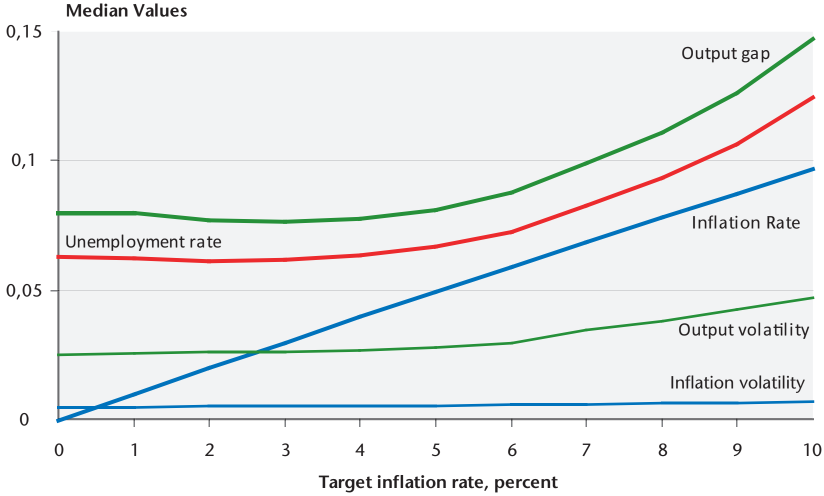

Source: Ashraf, Gershman and Howitt, 2012.

Figure 4. Price dispersion, number of shops, and target inflation

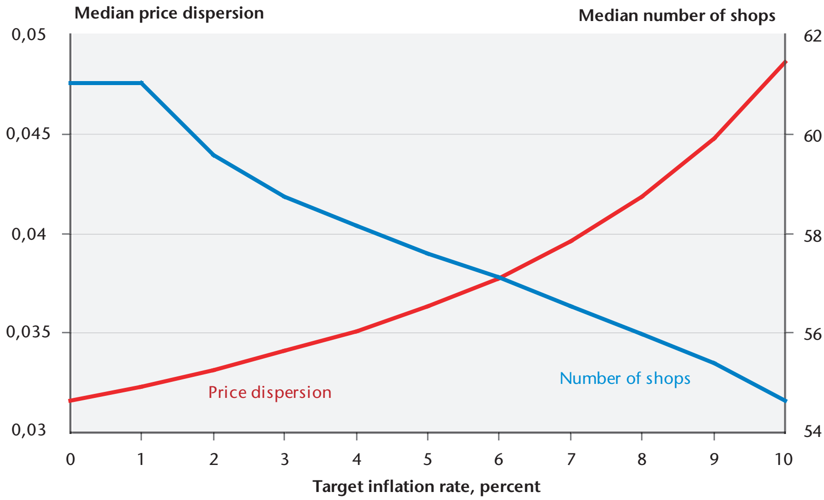

Source: Ashraf, Gershman and Howitt, 2012.

4. Banks and economic crises

In Ashraf, Gershman, and Howitt (2011), we explored the role of banks in a self-organizing economy. Banks are of course a kind of shop in themselves, being intermediaries between buyers and sellers of credit. But they also play a critical role in the creation and destruction of the other shops in an economy's trading network by providing or withholding finance. In this paper, we explored the effect that banks have on macro performance, much as we did in the other paper where we explored the effects of trend inflation.

The model used in this paper was a further extension of the one used to study inflation. We added a fixed number of banks, who make commercial loans to shops, with full recourse, secured by collateral in the form of fixed capital and inventories. The banks charge a fixed spread over their deposit rate, and always allow customers to borrow up to a regulatory maximum loan-to-value ratio. The banks also invest in government bonds. Households hold their non-money wealth in the form of bank deposits, which bear the same interest rate as government bonds. Banks are also subject to capital requirements, and are sanctioned when their capital is inadequate. The government in this extended model acts not just as central bank and fiscal authority but also as bank regulator and deposit insurer. When a bank is found to be insolvent, the government injects enough capital to make all deposits good and to restore its capital adequacy, and finds a new owner for the bank, much as the FDIC now does routinely for small banks that fail in the United States.

What this paper shows is that banks matter a great deal for economic performance. The baseline model is calibrated to US data and simulated 10,000 times, first with and then without banks. The difference is remarkable. With banks, the median simulation run had an annual average unemployment rate of 5.9 percent. Without banks it was 11 percent. Similarly, the volatility of the output gap was 2.8 percent with banks and 6.2 percent without.

We also explored in this paper one of the aspects of economic performance that almost always escapes orthodox DSGE analysis, namely the prospect that an economy can perform reasonably well most of the time but can, on occasion, go completely out of control. On average, across all 10,000 runs, the rate of inflation and the output gap were roughly constant after a 20-year initialization period. In the case of most individual runs, the output gap never deviated by more than 20 percent from its average value. But in a small fraction of runs (about 5 percent) the economy at some point diverged radically from this pattern and exhibited wild fluctuations of the output gap, with GDP falling in some cases even to zero.

In an attempt to contribute to the literature on "macroprudential" regulation, we used this model to explore the ways in which bank regulation affects the economy. In particular, we ran experiments in which we varied either the maximal loan-to-value ratio (LTV) or the capital adequacy ratio (CAR). What we found surprised us somewhat in two senses. First, in terms of median results, neither of these regulatory parameters seemed to matter. For example we compared the median results of our baseline model, with LTV = 0.5 and CAR = 0.08, to those from a "risky" scenario in which LTV = 0.9 and CAR = 0.02. The median results were quite similar across scenarios in terms of all macro indicators except bank failures. Thus, it seems that in "normal times" the economy is not much affected by bank regulations.

But when we sorted the 10,000 runs by average GDP over the 40 years of each run, we found another surprise. In the worst decile of runs (the 10 percent of runs with the lowest average GDP), we found that the regulatory parameters mattered a lot, and that in fact these "bad times" were a lot *less bad* under the "risky" scenario! Thus, what would normally seem like risky bank regulation from a microprudential point of view turned out to alleviate the problems that arise in those small number of cases where the economy was spinning out of control. For example, the average unemployment rate in the worst decile of runs was 7.2 percent in the "risky" scenario *versus* 8.9 percent in the baseline, and output volatility was 3.4 percent *versus* 5.2 percent.

Exploring the source of this last surprise, we discovered that it worked the way it did because banks provide the self-organizing economy with not just a financial accelerator, as emphasized by the literature started by Williamson (1987) and Bernanke and Gertler (1989), but also with a financial stabilizer. That is, when the economy is starting to become disorganized it has a critical need for entrepreneurship to replace shops that have failed. Bank finance is essential for facilitating this entrepreneurship. But when banks are heavily regulated they are less able to provide this finance just when it is most needed, that is, just when firms are running short of unpledged collateral and banks are finding themselves short of capital. The more of their capital they are able to devote to financing entrepreneurship and the less collateral entrepreneurs are required to put up, the more banks can play this essential role of averting a cumulative collapse to the economy's trading network.

Of course, all of these results must be heavily qualified by noting in particular that we have postulated banks that do not engage in proprietary trading, and are not influenced by the moral hazard issues associated with too-big-to-fail. Nevertheless, the results do illustrate the potential for properly regulated banks to help prevent an economy from leaving what Leijonhufvud (1973) calls the "corridor" of stability. The results also illustrate the new perspective that can be had from looking at economic fluctuations from the point of view of a self-organizing economy.

References

Ashraf Quamrul, Boris Gershman & Peter Howitt, 2011, *Banks, Market Organization and Macroeconomic Performance: An Agent-Based Computational Analysis*. Unpublished, Brown University.

———, 2012, *How Inflation Affects Macroeconomic Performance: An Agent-Based Computational Investigation*. Unpublished, Brown University.

Bernanke Ben & Mark Gertler, 1989, "Agency Costs, Net Worth, and Business Fluctuations." *American Economic Review,* 79: 14-31.

Billi Roberto M., 2011, "Optimal Inflation for the U.S. Economy." *American Economic Journal: Macroeconomics,* 3: 29-52.

Chari V. V. & Patrick J. Kehoe, 2006, "Modern Macroeconomics in Practice: How Theory is Shaping Policy." *Journal of Economic Perspectives,* 20: 3-28.

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

Heymann Daniel, & Axel Leijonhufvud, 1995, *High Inflation.* Oxford: Clarendon Press.

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

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

Howitt Peter, & Petros Milionis, 2007, *On the Robustness of the Calvo Price Setting Model*. Unpublished, Brown University.

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

Leijonhufvud Axel, 1973, "Effective Demand Failures." *Swedish Journal of Economics* 75: 27-48. (now *Scandinavian Journal of Economics*).

Levine David K., 1991, "Asset Trading Mechanisms and Expansionary Policy." *Journal of Economic Theory* 54: 148-64.

Phelps Edmund S., 1972, *Inflation Policy and Unemployment Theory*. New York: Norton.

Wallis John J., & Douglass C. North, 1986, "Measuring the Transaction Sector in the American Economy, 1870-1970." In *Long-Term Factors in American Economic Growth*, edited by Stanley Engerman and Robert Gallman. Chicago: University of Chicago Press.

Williamson Stephen D., 1987, "Financial Intermediation, Business Failures, and Real Business Cycles." *Journal of Political Economy* 95: 1196-1216.

Woodford Michael, 2003, *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.