Banks, Market Organization and Macroeconomic Performance: An Agent-Based Computational Analysis[1]

Quamrul Ashraf, Williams College

Boris Gershman, Brown University

Peter Howitt, Brown University

October 7, 2009

[1]First draft of presentation to Cleveland Fed conference, October 15-16, 2009

Abstract

This paper is an exploratory analysis of the role that banks play in supporting what Jevons called the mechanism of exchange. It considers a model economy in which exchange activities are facilitated and coordinated by a self-organizing network of entrepreneurial trading firms. Collectively these firms play the part of the Walrasian auctioneer, matching buyers with sellers and helping the economy to approximate equilibrium prices that no individual is able to calculate. Banks affect macroeconomic performance in this economy because their lending activities facilitate entry of trading firms and also influence their exit decisions. Both entry and exit have conflicting effects on performance, and we resort to computational analysis to understand how these conflicting effects are resolved. Our analysis suggests that banks normally improve macroeconomic performance but that when the economy is experiencing particularly bad times banks make the situation even worse. We also use the model to investigate the macro effects of bailouts, of the extent of recourse to which banks are entitled on their loans, and of various regulatory changes that have been proposed, such as restrictions on the size of banks, restrictions on loan-to-value ratios, increases in capital adequacy ratios and changes that would make capital adequacy ratios vary with the state of the economy. The results suggest two kinds of regulatory changes that might improve macro performance in bad times, namely (1) a restriction on loan-to-value ratios and (2) restrictions on bank dividends.

1 Introduction

How do banks affect the macroeconomy? If banks get in trouble how does that matter for various performance measures? Collective intervention to support the banking system in many countries over the past year or two has been predicated on the idea that without such support an economy is in danger of falling into a depression. This idea is roughly supported by the evidence of the United States in the 1930s, but not much support can be found for it in mainstream macroeconomic theory. There is a broad literature that studies the effects of financial development on long term growth through its effects on innovation, risk sharing, capital accumulation, the allocation of capital, and the screening and monitoring of investment projects.[1] But none of these effects seem likely to trigger a collapse of the sort that policy makers have been trying to avert. As for the DSGE models that are increasingly being used for policy purposes within central banks around the world, they are built on a foundation of perfectly functioning complete Arrow-Debreu contingent-claims markets, a foundation that leaves no useful role for banks to play. Introducing various financial frictions a la Bernanke-Gertler or Kiyotaki-Moore certainly makes the models more useful for many purposes but it does not change the fact that financial transactions in these models are not explicitly intermediated by banks, and thus there is no channel through which bank troubles impinge on the real economy.

In our view the critical importance of banks in a macroeconomy arises from the role they play in the economy's "mechanism of exchange," to borrow a phrase from Jevons. In any but the most primitive economic system, exchange activities are organized and coordinated by a network of specialist trading enterprises such as retailers, wholesalers, brokers, and various other intermediaries. These enterprises provide facilities for buying and selling at publicly known times and places, provide implicit guarantees of quality and availability of spare parts and advice, quote and advertise prices, and hold inventories that provide some guarantee to others that they can buy at times of their own choosing. In short, this network of intermediaries constitutes the economy's operating system, playing the role in real time that general equilibrium theory assumes is costlessly played in metatime by "the auctioneer," namely that of matching buyers with sellers and helping the economy to approximate the equilibrium vector of prices that no single person is able to calculate. Moreover, unlike the auctioneer, they provide the facilities and the buffer stocks that allow trading to proceed even when the system is far from an equilibrium.

The importance of this network of trading enterprises is attested to by Wallis and North (1986), who argue that providing transaction services is the major activity of

[1] For an introduction to this literature see Levine (2005) or Aghion and Howitt (2009, ch.6).

business firms in the US economy; more specifically, Wallis and North estimated that over half of measured GDP in the US economy consists of resources used up in the transaction process. And indeed, as everyday experience of any household will verify, almost all transactions in a modern economy are conducted with at least one side of the transaction being an enterprise that specializes in making similar transactions.

Banks and other financial intermediaries play a critical role in an economy's trading network, not just because they themselves are part of the network, intermediating between surplus and deficit units, but also because their lending activities influence the entry and exit of other intermediaries throughout the network. Entry of new facilities is not free and automatic. It requires entrepreneurship, which is not available in unlimited supply and which frequently needs finance. Likewise exit of existing facilities constitutes a loss of organizational capital that affects the system's performance, and exit activity is typically triggered by banks deciding when to cut off finance from a failing enterprise.

The purpose of this paper is to present a model that portrays this role of banks in the mechanism of exchange, and to explore ways in which banks influence the economy's macroeconomic performance through this role. Since a major part of our premise is that banks affect the economy's ability to approximate an equilibrium we obviously cannot use a rational expectations approach. Instead we present an agent-based computational analysis.

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.[2] 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, people are assumed to follow simple rules, whose interaction might or might not lead the system to approximate a coordinated equilibrium. 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.

More specifically, we use a modified version of the adaptive model developed by Howitt and Clower (2000) in which an economy's network of trade specialists is shown to be self-organizing and self-regulating. Howitt and Clower show that starting from an initial situation in which there is no trading network, such a network will often emerge endogenously, and that once it does emerge it will guide the economy to a stationary

[2] A survey of literature using the method in economics is provided by Tesfatsion and Judd (2006).

equilibrium in which almost all the gains from trade are fully exploited.

Here we modify the original Howitt-Clower model to allow for durable goods, fiat money and government bonds, to include an adaptive central bank that follows a Taylor rule with an explicit inflation target and a fiscal authority that adjusts tax rates in response to changes in the ratio of government debt to GDP, and to allow for banks that lend to the trade specialists in order to finance their inventories. The banks are subject to a number of regulatory influences, such as capital adequacy ratios and limits on loan-to-value ratios. We calibrate the model to US data and simulate it many times for many years under different parameter values to see how banks affect macro performance and how performance is affected by different dimensions of bank regulation.

1.1 Countervailing effects of entry and of exit

It might seem obvious that banks are crucial to the smooth functioning of an economy, and that keeping banks out of trouble allows them to play that role more effectively. But a glance at innovation-based endogenous growth theory suggests that the truth might be more complicated than this summary judgement would suggest. In particular, as Aghion and Howitt (1992) argue, innovation can sometimes be excessive, because of its rent-seeking aspect. A new entrant may internalize the social benefits of the new technology that she is putting in place but does not internalize the cost of losing the old technology that she will be replacing. Thus, to the extent that banks facilitate exchange by easing financial constraints on newly entering intermediaries they are not necessarily improving economic welfare.

Now the analysis below does not concern itself with "welfare analysis" of the usual sort. Instead it will be looking simply at the ability of the economy to approximate full capacity utilization. But even in that sense, the analysis of Howitt and Clower (2000), on which the current model is founded, also makes it clear that innovation can indeed be excessive. More specifically, Howitt and Clower argued that in order for entrepreneurship to play a useful role in creating and maintaining a well functioning trading network it must be made conditional on the state of the economy. That is, entry should be intensive when the economy is far from equilibrium but not so intensive when the economy is near equilibrium. Each act of entrepreneurial innovation disturbs the existing pattern of trade by diverting trade from existing intermediaries as well as by creating new trading opportunities for those that have not yet found a suitable intermediary. When the existing pattern of trade is dysfunctional it needs to be disturbed, but when people are already exploiting most of the potential gains from trade then entry typically does more harm than good by diverting resources away from current production (each new entrant incurs a fixed cost that does not contribute to final output) and by threatening the existence of otherwise viable incumbent trading enterprises.

In the Howitt-Clower model this modulation of entrepreneurship took place through a process of market research. Each potential entrant would survey some potential customers and potential suppliers before deciding whether or not to enter. When markets are disorganized these surveys will often signal that there are potential profits to be gained from entry, whereas when markets are already functioning well then the survey responses will tend more often to indicate that customers and suppliers will be hard to attract away from incumbent firms, which will discourage the entrepreneur from entering.

What these considerations suggest is that banks will indeed improve economic performance if they help to modulate entry appropriately, by facilitating entry when the economy is far from full capacity utilization but less so when it is close. However, bank lending often works in the opposite direction. That is, in good times banks will tend to have adequate capital to finance new entrants, whereas in tougher times they will tend more often to be in trouble and therefore less willing to finance new entrants.

The same considerations, taken together with the results of Howitt (2006) suggest that not just entry but also exit needs to be conditioned on the state of the economy, and that the way bank lending conditions exit may not be appropriate. More specifically, Howitt (2006) shows that, starting near an efficient pattern of trade, exit can have a cumulative multiplier effect that amplifies any deviation from that pattern. That is, failure by one shop can cause difficulties for other shops whose customers were deriving their income from the failed shop, thus causing a cumulative loss of organizational network capital. Banks can perhaps play a role in moderating this deviation-amplifying mechanism, by providing finance that keeps a firm going in the face of a loss of business caused by the failure of other shops. This consideration suggests that what might matter for macro performance is not just the overall level of financing provided by banks but also the timing. If banks lend freely during good times only to cut lending off when a cumulative downturn is underway they may end up exacerbating the most serious macroeconomic problems.

1.2 Looking ahead: Normal versus worst-case results

The model we will present exhibits an important nonlinearity. As we shall see, most of the time the evolving network of trade intermediaries performs reasonably well in counteracting macro shocks and keeping the economy in a neighborhood of full capacity utilization.

But in a significant fraction of runs the economy tends to spiral out of control. The model thus exhibits something like what Leijonhufvud (1973) called "corridor effects;" that is, if the system is displaced far enough from equilibrium its self-regulating mechanisms are liable to break down entirely. This could not happen in a stochastic linear model where expected impulse responses are independent of the size of displacement. The distinction between median results and worst-case results shows up dramatically in almost all the experiments we perform on the model.

Our analysis is at best a first step in the direction of introducing banks into an agent-based computational analysis. Bank lending is limited to financing inventory accumulation, and inventories in the model are typically fairly small in relation to GDP. Thus we will find that in normal times banks and bank regulation have a quantitatively small effect on the economy. We also find, however, that even when banks are restricted to this limited role they have a major impact on performance when the system is far from equilibrium. Moreover, the way banks and bank regulation affect median performance is often in directly the opposite direction to the way they affect performance in the worst times. Generally speaking it seems that the kind of banks that are normally favorable to good macroeconomic outcomes also increase the expected damage from having the system go out of control.

We also use the model to investigate the macro effects of bailouts, of the extent of recourse to which banks are entitled on their loans, and of various regulatory changes that have been proposed, such as restrictions on the size of banks, restrictions on loan-to-value ratios, increases in capital adequacy ratios and changes that would make capital adequacy ratios vary with the state of the economy. Some of these variables seem to have little systematic effect on macro performance even in the worst of times. But there are two regulatory changes that appear systematically to improve macro performance in bad times, namely (1) restrictions on loan-to-value ratios and (2) restrictions on bank dividends.

The next section describes our model in general terms. Section 3 provides details of the computer algorithm that implements the model. Section 4 describes the no-shock full-capacity-utilization equilibrium that the system approximates and discusses the ways in which entry, exit and bank lending affect this process. Section 5 describes how the model was calibrated and illustrates the difference between how the system behaves in normal times and how it behaves when things go wrong. Section 6 describes our results. Section 7 concludes.

2 The model

2.1 Preliminaries

Our agent-based model[3] is a variant of the adaptive model developed by Howitt and Clower (2000), as modified by Howitt (2006 and 2007). The model attempts to portray in an admittedly crude form the mechanism by which economic activities are coordinated in a decentralized economy. It starts from the proposition that in reality almost all exchanges in an advanced economy involve a specialized trader ("shopkeeper") on one side or the other of the market. We add several components to this model so as to make it considerably less stylized. In adding these new components we have tried to make the structure and its macroeconomic aggregates comparable to the baseline New Keynesian analysis (for example, Woodford, 2003) that is now commonly used by many central banks. That is, prices are set by competing firms acting under monopolistic competition, the rate of interest is set by a monetary authority following a Taylor rule, and consumer demands depend, *inter alia*, on current wealth. However, it is quite different in three important senses. First, we have introduced elements of search, in both goods (retail) markets and labor (wholesale) markets, whereas the canonical New Keynesian model has a Walrasian labor market and no search in the goods market.[4] Second, we assume that firms are subject to failure and that the process of replacing failed firms is a costly one, whereas the population of firms is fixed in the New Keynesian framework. Third, instead of the perfect and complete set of contingent financial markets assumed in the New Keynesian literature we assume that the only available financial instruments are non-contingent bank deposits, bank loans to shops, and government-issued money and bonds.

2.2 Disequilibrium analysis

Most importantly our analysis differs from New Keynesian analysis because we are investigating the out-of-equilibrium behavior of the system. The hypothesis that we are investigating is that banks can affect real economic performance by improving (or degrading) the system's ability to track a coordinated state of equilibrium. In a sense it is a continuation of a line of research into disequilibrium macroeconomics that reached its pinnacle in the Barro-Grossman (1976) book, which attempted to flesh out the details of what happens when people trade at prices that make their plans mutually incompatible.

[3] A similar model, but without private banks, is used by Ashraf and Howitt (2008) to investigate the effects of trend inflation on macroeconomic performance.

[4] See however Blanchard and Gali (2008).

That line of research ran into the problem that trading out of equilibrium in one market generated rationing constraints that affected traders in other markets, in ways that were hard to rationalize and hard to model. The Howitt-Clower model provides a systematic framework for modelling those out-of-equilibrium interactions through the agency that seems to manage the interactions in real-world markets, namely specialist trading firms that create and operate markets.

Moreover, agent-based computational analysis provides a way to deal with the complexities that stalled the earlier disequilibrium literature, not by sidestepping them through assumptions that limit their scope but rather by embracing them, putting them into the computer code, and using the power of modern computer technology to explore the macroeconomic patterns that emerge from seemingly chaotic microeconomic interactions.

2.3 The conceptual framework

There are N transactors and n goods. In addition to these goods there are four nominal assets, base money, bank deposits, bank loans and bonds. In addition to the N transactors there is a central government that issues the bonds and money, imposes a sales tax on retail transactions, acts as a lender of last resort to the banking system and facilitates the transfer of ownership of failed banks. Time is discrete, indexed by "weeks" $t = 1, ..., T$. Each bond is a promise to pay one unit of money ("dollar") next week.

Each transactor can eat only two of the goods (his "demand goods" 1 and 2) and is endowed each week with enough labor services to produce one unit of a particular good (his "supply good"), which is never one of his demand goods. A transactor whose supply good is i and whose first demand good is j is said to be of type (i, j), and his second demand good is $j + 1 \pmod n$. For each of the $n(n - 2)$ ordered pairs (i, j) of goods such that $i \neq j$ and $i \neq j + 1 \pmod n$, there is exactly one transactor of type (i, j). The population of the economy is thus $N = n(n - 2)$.

Because no transactor eats his own supply good, he must trade to eat. Trading is such a costly affair that it can only take place through an organized facility called a "shop." To trade with a shop a transactor must form a trading relationship with it. Each shop is capable of dealing in only one good, which it exchanges for money. Each transactor may have ongoing trading relationships with at most one shop (his "employer") that deals in his supply good, and at most one shop ("store") that deals in each of his demand goods. Each week each transactor delivers his unit endowment for sale to his employer if he has one, and visits his stores to buy demand goods from them. Goods are storable, but only by shops. Financial assets are storable by all agents.

Each shop posts a pair of prices for the good that it deals in: a wholesale price, or "wage," w and a retail price p. The shop agrees to buy all labor delivered to by each employee at the price w, provided that it has enough cash on hand at the time of delivery, up to a limit (described below) that depends on its inventory-sales ratio, and to sell whatever quantity is demanded by each customer at the retail price p provided that it has enough inventories on hand at the time the customer places the demand.

Every week, transactors that are without a complete set of shops search for information about possible trading relationships, through a combination of direct search and referral. The former consists in examining a randomly selected shop, and the latter consists in making inquiries with random people having the same demand or supply good. Given the results of such search activities a transactor will make his choice myopically – choosing the employer that is currently offering the highest wage and the shop that is offering the lowest retail price. In the case of a referral, however, the relevant price is not the actual price posted by the shop but rather the most recent effective price experienced by the referee, which will differ from the actual price if the shop has rationed the referee because of a shortage of goods or finance (see section 3 below). A shop will take on any given number of new customers but is not obligated to take on new employees in this process. So an employment relationship will form only if agreed to by both parties. Each trading relationship that is formed in this process is like those in standard search/matching theory – it will last until either the shop exits, the match dissolves for exogenous reasons, the transactor finds a preferable shop willing to match with him as the result of future search activities, or the employer lays off some of its workers.

Shops can be opened only by transactors that innovate. More specifically, each week a random subset of transactors ("entrepreneurs") are struck by an idea for opening a shop that trades the transactor's supply good. The entrepreneur will open such a shop if he thinks he can earn a profit higher than his current estimate of permanent income and if he can afford to pay the fixed overhead cost for at least the first week of operation (see below). In estimating the profit he can earn he is influenced by a random draw of "animal spirits" which determines his estimate of what volume of sales he can achieve, and by the nominal interest rate in the economy. Sales volume matters because the shop must cover its overhead cost each period before it can begin to sell its goods and therefore it must generate at least a minimal amount of revenue. The nominal interest rate affects the interest opportunity cost of holding cash to pay the wage bill. The bigger the sales estimate and the lower the nominal interest rate, the more likely he is to enter. Before entering the entrepreneur also engages in market research, as discussed in the preceding section and as explained in more detail in the following section.

Each shop posts a retail price that constitutes a fixed markup over its wage net of the sales tax. The wholesale price, or "wage," is adjusted periodically, at fixed intervals. As in an expectations-augmented Phillips Curve, the adjustment is proportional to expected inflation. Here and elsewhere we assume that transactors always believe the inflation-targeting central bank and hence always forecast an inflation rate equal to the target rate. Given the target rate of inflation the wage adjustment is also an increasing function of a measure of the shop's excess demand for employees; i.e., the ratio of its target employment level (see below) to its actual labor force. Whenever a shops wage is adjusted, its retail price is automatically marked up at the same time. We suppose that shops also adjust their prices at other times whenever their inventory/sales ratio falls above or below a band around the desired value of unity.

Each shop has a fixed overhead cost that must be incurred each period in the form of the shop's traded good. As just indicated, it also aims to keep an inventory level equal to one week's sales. When inventories fall below this target level it plans to increase them by a fixed fraction of this gap each period. Its target employment level each period is thus equal to its fixed cost plus the quantity demanded last week plus the planned inventory adjustment. Whenever the shop finds itself with an actual labor force in excess of its target employment level and also at the same time with an inventory/sales ratio that is above the critical band around the desired value of unity, the shop will lay off its excess workers. Otherwise it will continue to employ all of its labor force.

Each agent is assigned to one of m banking sectors, based on his supply good. We assume that banking is a natural monopoly, so that the number m of sectors is a policy variable. There is always one bank, owned by a single transactor, in each sector. Before trading with shops each week, transactors go to their banks, where they repay bank loans, take out new loans, make deposits and withdrawals from their bank accounts and plan their current expenditure levels for the current week. Before the banks open each week they are audited to see if they are solvent and if they have adequate capital. Capital adequacy means that the bank's equity is at least equal to a certain "capital adequacy ratio" times a risk-weighted sum of its assets. Banks that are insolvent are allowed to fail, at which time the government acts like the FDIC, finding a new owner for the bank and recapitalizing the bank with injections of high-powered money that restore capital adequacy. Banks that are found to have inadequate capital are not allowed to initiate new loans and their owners cannot receive any funds for personal use, until capital adequacy is restored.

Banks are assumed to grant all loans to shopkeepers up to a limit equal to the estimated value of the shop's inventories multiplied by a fixed loan-to-value ratio. The estimated value is equal to the current economy-wide average wage rate, which provides a rough approximation to the replacement cost, and this value times the loan-to-value ratio is what we call the "haircut price" of inventories. Shopowners who are unable to pay back their pre-existing loans even after rolling them over up to this limit imposed by the bank are declared bankrupt. The bank seizes all their assets and the shop will exit later that week. (In the no-recourse case studied below banks just seize all the bankrupt shop's inventories.) After all transactors have made their deposits and withdrawals, banks allocate all of their excess reserves to the purchase of short-term government bonds.

Banks pay the same interest rate on deposits as they receive on government bonds. But they charge a premium above the government bond rate for bank loans, to cover the risk of default. The premium is one of two values depending on the length of the relationship between the bank and the shop. Once the relationship becomes a year old the premium falls to the lower value. Thus one consequence of a bank failure is that all relationships become new ones and all loan rates become higher for a period of one year.

When allocating their portfolios, households take into account a cash-in-advance constraint. If they are shopkeepers they need cash in advance not only for their household purchases but also to pay their employees. In the baseline case where loans are with full recourse, shopkeepers with outstanding loans never hold deposits, because of the interest spread. When they are fully loaned up and do not have enough wealth to pay for their household planned expenditures and their estimated wage bill they are assumed to attach first priority to their wage bill. Households that are bank owners are free to take unlimited payments from the bank as long as it remains solvent. (We impose dividend restrictions in one of the experiments to be reported on below.)

The interest rate on government bonds is set by the government, who agrees to buy or sell whatever quantity banks wish to exchange at the government's target interest rate. This target interest rate is adjusted on the first week of each "month" (for ease of computation we assume 4 weeks in each month and 12 months in each "year") according to a Taylor Rule that responds positively to the difference between the most recent year-to-year inflation rate and the inflation target, and to the difference between the most recent estimate of the actual output gap and the government's estimate of its "natural" value. (Real GDP is the sum across all firms of input minus fixed costs.) The government is assumed not to know the economy's natural interest rate or the natural value of the output gap. Instead it has to estimate these values adaptively, following a procedure to be described in more detail below.

The government adjusts its tax rate once per year in response to any discrepancy between the actual level of outstanding debt and some fixed target level.

There is a "firesale market" that convenes just prior to trading in goods and labor markets each week. The sellers here are banks that hold seized collateral and former shopkeepers that are holding legacy inventories from their former shop. The buyers are shopowners trying to build up their inventories. Every shop with a shortfall of actual inventories below desired inventories will attempt to buy an amount equal to that shortfall. We assume that bargaining in the firesale market always results in a firesale price equal to half of the current economy-wide average wage rate, which in turn is half of (a rough approximation to) the replacement cost of the goods. The firesale price is always assumed to be no greater than the haircut price, so that as long as the bank in the sector containing all i-makers is not "in trouble" (i.e., deemed to have inadequate capital) it will be willing to finance the acquisition of firesale inventories. The firesale market is thus an important channel through which banks provide support to new entrants, who typically come into operation with few initial inventory holdings.

All of the rules of behavior that we have outlined above are simple intuitive rules that require no optimization or forecasting. However there is one place where the agents are assumed to be somewhat sophisticated, and that is in formulating their demands for consumption goods and their related demands for money and bonds. Each household plans to spend a fixed fraction of current wealth on its consumption goods. If he has just one shop then all of its expenditure will be allocated towards that shop. Otherwise it will be allocated to maximize a CES current utility function of the two consumption quantities.

The sophistication enters into this simple rule in the calculation of current wealth. We assume the transactor understands that wealth consists in not just current deposit and money holdings but also future income prospects. He calculates this second "human" wealth component in two stages. First he estimates "permanent income" according to a simple adaptive-expectations rule as in Friedman's original formulation of the permanent income hypothesis. Next he converts this permanent income into wealth using a capitalization factor that comes from a forecaster of interest rates, namely the central bank. Although each transactor has a different permanent income they all use the same capitalization factor. This factor in turn is consistent with the Taylor rule being used to set interest rates, under the assumption that the logs of output and the inflation factor follow an AR(1) stochastic process, which the forecaster estimates in an adaptive fashion described below. We believe that this sophisticated calculation of expected future interest rates is consistent with the philosophy of agent-based modelling because of the growing evidence that the transparency of inflation targeting has made interest rate forecasts relatively accurate and widely available. In any event, these forecasts are what allow monetary policy to have a non-trivial effect on demands in our model despite the fact that the monetary authority is constrained to operate only in very short-term bond markets and is assumed not to engage in deliberate interest-smoothing.

3 Algorithmic Details

The model sketched above has been implemented as a computer program, written in the C++ programming language. It represents the economy as an algorithm. The present section describes this algorithm in more detail and enumerates the free parameters on which behavior depends. The initial conditions each week consist of a certain number of established shops, the goods they trade, a historically given configuration of ongoing relationships between transactors, shops and banks, a set of targets, permanent income estimates and effective prices from last week, an array of money, deposit and inventory holdings, banks' balance sheets, a tax rate, an interest rate, a history of price and GDP levels, central bank's current monetary policy parameters, required capital adequacy ratio, a lagged measure of the economy-wide average wage rate, and a fire-sale price. The algorithm then proceeds to generate a new set of initial conditions for the following week in nine stages, each of which represents an important component of the workings of a decentralized economy. These nine stages of weekly activities, described below, are repeated for a total of T weeks.

3.1 Firm Entry

In the first stage of the algorithm representing weekly activities, each transactor who is not already a shop owner or a bank owner is considered as a potential innovator. Each one becomes an actual innovator with probability θ/N. The innovator is first assigned a random realization x of target sales (animal spirits) from a uniform distribution over $[1, \bar{x}+1]$, where the upper bound of the support represents the maximal target sales of a potential entrant, and a random markup μ drawn from a uniform distribution over $[0, 2\bar{\mu}]$. Given his target sales and markup, the innovator estimates his weekly target (after-tax) economic profit as a shop owner as[5]

$$\Pi = [(\mu - i_w)x - (1 + i_w)(F - 1)]w,$$

where F is a fixed cost parameter, i_w is the nominal interest factor per week, and w is the wage rate set by the innovator. The wage w is set as follows:

$$w = W(1 + \pi_w^*)^{\frac{1+\Delta}{2}},$$

where W is the economy-wide average wage rate for the previous week,[6] π_w^* is the weekly target inflation factor, and Δ is the length of the contract period in weeks. This wage-setting rule is designed so that the present value of the employee's wage in the middle of contract period (given the current target inflation factor) is equal to W. This implies that initially the wage is set to be higher than W to take into account the expected inflation, because wages are kept fixed during the contract period.

Next, the program checks if the innovator setting a wage rate w satisfies the *"hopelessness" condition*. A potential entrant is hopeless if the sum of his money holdings, deposits and the credit limit from his bank is not enough to cover the fixed costs of operating a shop, $(F - 1)w$, not counting own input. In this case there is no entry. On the other hand, a hopeful innovator goes on to check the *profitability condition*. This condition holds if the target profit exceeds the sum of the latest estimate of permanent income Y^p, updated during the financial market trading stage in the previous week, and the appropriately discounted value of his legacy assets LA (if any):

$$\Pi > Y^p + \text{LA} \cdot P_f / V,$$

where P_f is the fire-sale price and V is a capitalization factor, equal to the present value of a nominal income stream that grows each week at the constant weekly target rate of inflation, given the sequence of nominal interest rates that the central bank is projecting. A hopeful innovator that passes the profitability test proceeds to *market research*.

To simulate market research, the program identifies the innovator's supply good and

[5] This expression for target economic profit takes into account the opportunity cost of using money to pay for inputs. The target accounting profit of the shop owner is $[\mu x - (F - 1)]w$, where 1 is subtracted from F since the shop owner does not need to pay himself for the endowment. If, instead of producing, the agent put the money spent on inputs, $(x + F - 1)w$, in a bank, he would earn the interest $i_w(x + F - 1)w$. Subtracting this opportunity cost from the target accounting profit, we get target economic profit of the prospective shop owner.

[6] The average wage rate and the fire-sale price are updated weekly at the end of the labor and goods market trading stage and are known to all agents in the economy.

then looks for his potential employees and customers. First, a "comrade" (someone else with the same supply good as the innovator) is chosen at random and becomes a potential employee. Next, a potential customer (someone whose primary consumption good coincides with the innovator's supply good) is chosen at random. If the comrade's current effective wage is lower than the (inflation adjusted) wage $w/(1+\pi_w^*)$ offered by the innovator and the customer's effective retail price is lower than the (inflation adjusted) one offered by the innovator, $[(1+\mu)w]/[(1-\tau)(1+\pi_w^*)]$, where τ is the sales tax rate, the market research is successful and a new shop is created that trades the innovator's supply good. The innovator becomes a shop owner, the comrade becomes his actual employee with the respective effective wage, while the potential customer becomes his actual customer and is assigned a new effective price. The legacy inventories of the entrant (if any) become part of the shop's inventories, and the target input is set at the level $x + F + \lambda_I(x - I)$,[7] where λ_I is the weekly inventory adjustment speed and I is current inventories which for a new shop are just equal to the entrant's legacy inventories.

The key parameters introduced in this stage therefore include: θ, which determines the frequency of innovations; $\bar{x}$, governing the animal spirits of potential entrants, the average markup $\bar{\mu}$, fixed cost F, the length of the contract period Δ, and the weekly inventory adjustment speed λ_I.

3.2 Search and Matching

Next, the program gives each transactor an opportunity to search for possible trading relationships. This comprises both job search (for a shop that buys the transactor's supply good) as well as store search (for shops that sell either of his two demand goods). Each transactor who is not a shop owner engages in job search with probability σ. Job search consists in asking one randomly selected comrade what his effective wage is. If it exceeds the searcher's current effective wage, the searcher attempts to switch to the comrade's employer. The switch will be implemented if and only if the employer's current input level is less than its target input level. If so, the searcher's former employment relationship (if any) is severed and his effective wage is set equal to the comrade's.

Store search, on the other hand, is undertaken by every transactor. This type of search comprises not just referral-based but also direct search. First, the transactor asks a randomly selected "soulmate" (someone with the same two demand goods) for his effective retail prices. If either is lower than the searcher's, the searcher will switch to the corresponding store and set his effective retail price equal to the soulmate's. Then

[7] If the entrant's endowment, current cash holdings and deposits do not allow to finance this level of target input, the latter is set to the available maximum equal to entrant's wealth.

the transactor selects a shop at random. If the shop trades either of his demand goods and is posting a retail price lower than the searcher's effective retail price the searcher will switch to that store and set his effective retail price equal to the (inflation adjusted) store's posted price. Every time he switches, the transactor will sever any relationship with a store trading the same good.

The key parameter introduced at this stage of the weekly activities is the probability σ that a transactor will engage in job search.

3.3 Financial Market Trading

At this stage all the financial transactions take place. The balance sheet of each commercial bank looks as follows:

Assets	*Liabilities and Equity*
Commercial Loans	Deposits
Seized Collateral	Loans from CB
Government Bonds	Equity
Reserves	

On the assets side, commercial loans are loans given by banks to shop owners, seized collateral consists of inventories seized by the bank in from defaulting shops, valued at the current fire-sale price, government bonds are bonds held by the bank and reserves are holdings of high-powered money (possibly negative) resulting from the deposits and withdrawals of the banks' customers and from the issuance of new bank loans. The liabilities of a bank consist of deposits held by agents assigned to this bank and loans from the central bank. Equity is calculated as bank's assets minus its liabilities.

Before the financial market trading takes place, banks in all sectors are examined. Equity is updated after previous week's transactions. Banks with negative equity fail. When a bank fails, first, a government agency (FDIC) injects money to fully capitalize the new bank so that it fulfills the minimum capital requirement (see below). Then a new owner is chosen from the list of the failed bank's customers who do not own a shop. In particular, the richest of them (with the highest sum of cash and deposit holdings) becomes the new owner. Since its owner changes, the bank starts a new relationship history with all of its customers, i.e., the old customers who previously enjoyed the benefit of relationship lending get a higher interest rate on loans (see below). If the new bank owner has some legacy inventories, they are put on the bank's balance sheet (seized collateral account) and participate subsequently on the fire-sale market along with other foreclosed inventories that the bank has. The equity is updated to take into account possible additions to the balance sheet.

Next, all banks are checked for capital adequacy. In particular, the ratio of bank's equity to its risk-weighted assets must be greater or equal to the *capital adequacy ratio* κ which is set by the central bank at the monetary policy stage[8]:

$$\text{Equity} \geq \kappa \cdot (1 \cdot \text{Commercial Loans} + 1 \cdot \text{Seized Collateral} + 0 \cdot \text{Government Bonds}).$$

If this condition is violated, the corresponding bank becomes a "troubled" bank. Troubled banks are not allowed to increase loans, and their owners cannot get dividends.

Next, agents of all types do their *budget planning*. The total wealth of each agent is the sum of financial wealth, A_i, and the capitalized value of permanent income, Y_i^p. The *financial wealth* of transactors who don't own a shop or a bank is just the sum of their money holdings and bank deposits, plus the value of legacy inventories (if any). For bank owners the financial wealth is the sum of money holdings and bank's equity. For shop owners it is equal to the sum of money and deposit holdings, minus outstanding loans.

After financial wealth has been calculated, the agent's *permanent income* is adjusted according to the following adaptive rule:

$$\Delta Y_i^p = \lambda_y \left(Y_i - Y_i^p \right),$$

where Y_i is the actual last period's income, and λ_y is the weekly permanent income adjustment speed. Here, Y_i is equal to last period's profit for shop owners and effective wage rate for all other agents. Then, the agents update Y_i^p again to adjust for estimated weekly inflation assuming that inflation is taking place each week at the target rate, i.e., multiply it by $(1 + \pi_w^*)$.

We assume that each agent wants to spend a fixed fraction v of total wealth on consumption goods during the current week:

$$E_i = v \cdot (A_i + V \cdot Y_i^p),$$

where V is the same capitalization factor as in section 3.1, i.e., the present value of a nominal income stream that grows each week at the constant weekly target rate of inflation, given the sequence of nominal interest rates that the central bank is projecting.

[8]This formulation mimics Basel I capital accord. The assigned risk weights (1 for loans and seized collateral and 0 for government bonds) come directly from Basel I recommendations. In one of our policy experiments we will allow the capital adequacy ratio κ to vary depending on the central bank's estimate of the current economic situation (see section 6.3).

Note that this is precisely the expenditure function that would apply if the transactor knew for certain what future incomes and interest rates would be and were choosing E_i so as to maximize a standard intertemporal additive logarithmic utility function with a weekly rate of time preference $\rho_w = v/(1-v)$. We will use this interpretation of the above expenditure function when calibrating the model, and will calibrate it in terms of the annual rate of time preference, ρ, defined by $(1+\rho) = (1+\rho_w)^{48}$.

Having decided on the desired planned expenditure the agents choose the amount of cash M_i taking into account the constraints they face. Consider first the transactors that don't own a bank or a shop. If $E_i < A_i$, they set $M_i = E_i$ and put the rest, $A_i - E_i$, on the deposit account in their bank. Otherwise, they withdraw all of their deposits and so, their actual planned expenditure and money holdings are equal to the total wealth $M_i = A_i$. The idea here is that the agents will need to have E_i in the form of money when they visit their demand stores. But they do not know whether they will be paid their income before or after shopping for goods, so they carry E_i out of the financial market to ensure against being unable to fulfil expenditure plans.[9]

Next, consider a bank owner. If he owns a troubled bank, i.e., the minimum capital requirement is violated, he cannot receive dividends and his expenditure is bounded by current money holdings. If the latter exceed desired planned expenditure, the remaining part goes into the bank. If the bank is not troubled, then the owner can receive dividends but only up to the full value of its equity. At this stage, the owner of an untroubled bank does not worry about meeting the capital requirement and consumption expenditures are a priority.

Finally, consider a shop owner. His desired money holding equals not just enough to pay for his planned goods expenditure but also enough to pay his target wage bill, equal to his current posted wage times his current target input minus one (since he does not have to pay himself). First, the shop owner evaluates how much money will be available this period from the bank. The credit limit is set equal to the current haircut price, P_h, times the amount of inventories the shop has, i.e., inventories are used as collateral and are evaluated at the haircut price. This is the maximum loan the shop owner can take from his bank if the latter is not troubled. So, in this case the resulting financial constraint of the shop owner is $A_i + \mathrm{CL}_i$, where $\mathrm{CL}_i = P_h \cdot I_i$ is the available credit limit. If the bank is troubled, then it cannot increase the size of loans. Given his financial

[9] This motivation for a precautionary demand for money is similar to the "stochastic payment process" that Patinkin (1965) used to rationalize putting money in the utility function. In this case we are using it to justify what looks like a conventional cash-in-advance constraint.

constraint, the shop owner updates target input and the target wage bill.[10]

Based on the shop owner's financial situation, the following cases are possible:

1. If $A_i + \mathrm{CL}_i < 0$, the shop goes bankrupt. In this case the actual planned expenditure is set to zero, the bank seizes the collateral (inventories) and, in the case of non-recourse loans, all the cash and deposit holdings of the shop owner. The shop owner's loans are voided, and the bankrupt firm waits until the exit stage to leave the market.

2. If $A_i + \mathrm{CL}_i > 0$ but is not enough to pay the target wage bill, the shop owner sets the actual planned expenditure to zero, withdraws all the deposits and borrows as much as he can from the bank. The cash holdings are $M_i = A_i + \mathrm{CL}_i$.

3. If $A_i + \mathrm{CL}_i$ is large enough to pay the target wage bill but not the total desired consumption expenditure, the shop owner withdraws all deposits, borrows as much as he can from the bank and sets the actual planned expenditure to $A_i + \mathrm{CL}_i$ minus the target wage bill, i.e., his priority is to have enough cash to cover the target payroll. As before, $M_i = A_i + \mathrm{CL}_i$.

4. If the shop owner can afford to finance the entire wage bill and desired consumption expenditure, but cannot pay off the whole outstanding loan, he pays off as much of the loan as he can and holds enough cash to pay for the target wage bill and planned goods consumption.

5. Finally, shop owner may be able to cover the entire wage bill, desired consumption expenditure and the whole outstanding loan. Here again we make a distinction between the case of recourse and non-recourse loans. In the former case, the shop owner pays off the whole loan, holds enough cash to pay for the target wage bill and planned consumption, and puts the excess into his deposit account. In the latter case we also assume that the shop borrows up to the credit limit and then deposits this amount.[11]

Finally, each bank's weekly lending rate is recalculated along with the actual amount of outstanding loans due next period by the shop owners. The base rate is the weekly nominal interest rate on government bonds (equal to the weekly deposit rate). The

[10] The target input is calculated in the same way as in section 3.1, but the target sales are set equal to the previous period's actual sales.

[11] We make this simple assumption about the shop owner's "opportunistic" behavior to emphasize the riskiness of the non-recourse loan regime: with non-recourse loans, the deposits holdings cannot be seized by the bank.

spread over the deposit rate differs, however, for old and new customers. In particular, we assume that if the history of bank-shop relationship exceeds some fixed period of time, l, the spread is smaller than otherwise. In doing this we follow the literature on relationship lending. The normal (annual) loan spread, s, and the extra spread for new customers, $\tilde{s}$, are two of the model's parameters.

After all agents are done with budget planning, banks adjust their portfolios. They update their deposits and loans given the budget planning decisions, government bonds are redeemed, and the central bank debt is charged[12]. If the bank does not have enough funds to clear all financial operations, it has to borrow the shortfall from the central bank. If, in contrast, there is surplus, it is invested in government bonds.

The key parameters introduced in this stage are the annual rate of time preference ρ, the adjustment speed of permanent income λ_y, the length of the initiation period for bank's new customers l, the normal loan spread s and the extra spread $\tilde{s}$, and the discount rate premium s_d.

3.4 Labor and Goods Market Trading

This stage starts with the fire-sale market trade. All active shops that are not bankrupt can buy input good on the fire-sale market if they need it, i.e., if their target sales exceed the amount of available inventories. The difference between target sales and inventories is the amount they wish to purchase. If the shop's bank is in trouble, the shop owner cannot place an order that, evaluated at the fire-sale price, exceeds his deposit holdings.

If the desired amount of input is positive, the shop owner is matched to the first seller of his good in the queue (if any) with his order. Sellers at the fire-sale market can be of two types. First, these are bank owners who hold foreclosed inventories. Second, they can be agents who were previously shop owners and hold some legacy inventories after having exited the market. If the first seller in the queue cannot fulfill the whole order, he sells what he has and the turn goes to the next seller in the queue, and so on after the order is fulfilled or the queue runs out of sellers. The payment for inventories bought at the fire-sale market is in dollars due at the beginning of next week. If this amount is less than shop owner's current deposit holdings, the latter are decreased by a respective amount. Otherwise, the shop owner not only runs out of deposits but also has to take an express loan for the amount that he lacks.

Next, for each transactor, the program simulates trading in both the labor market (i.e., with the transactor's employer) and the goods markets (i.e., with his stores). With

[12] The rate paid on the central bank borrowing (the discount rate) is equal to the nominal weekly interest rate plus the discount rate premium, s_d.

probability 1/2 he first executes his labor market trading; otherwise, he first executes his goods market trading.

Labor market trading proceeds as follows. If the transactor is a shop owner (i.e., is self-employed), he simply uses his unit endowment as input. If he is not a shop owner, then he trades his endowment for money to the shop owner at the posted wage, subject to the constraint that, if the employer's money holding is less than the posted wage, they trade the amount that just exhausts that money holding, so that the transactor is rationed. The transactor then sets his effective wage equal to the amount of money he has just received. All input in excess of the fixed cost turns into inventories and is subsequently sold to final consumers. If the agent gets no money, he is effectively unemployed, since his employer is broke, but the relationship with the shop remains in place. The relationship with the shop is severed, however, if the worker is laid off. This happens if the employer is overstocked, that is, his actual input is greater than the target input if the unit offered by the current transactor is bought and the ratio of inventories to target sales exceeds a certain threshold, IS. To recap, there are 4 possible reasons for a worker to become unemployed in the model: 1) when he gets laid off by an overstocked shop owner; 2) if a random breakup of the match happens and he has to quit the job (see section 3.6); 3) when the employer exits the market; 4) when the employer is broke, i.e., does not have enough cash to pay the wage bill.

Goods market trading happens in the following manner. Given the amount of cash and the total amount of planned expenditure determined in the previous stage, the transactor determines his planned expenditure on each of his consumption goods by maximizing a two-good CES utility function

$$c_1^{\varepsilon/(\varepsilon+1)} + c_2^{\varepsilon/(\varepsilon+1)}$$

with a "demand parameter" ε. If the transactor has established relationships with stores for both of his consumption goods he trades with both of them, which meet his demand up to the point where their inventories are exhausted. The transactor then sets his effective retail price for each good equal to the actual posted retail price over the fraction of his demand that was satisfied. If the transactor has a customer relationship with only one shop, he goes through the same routine attempting to spend his entire actual planned expenditure on the corresponding consumption good.

At all stages of trading, the program adjusts inventories and money holdings. In goods market trading stage the program also makes sure to deduct the shop's tax liability, equal to the prevailing tax rate τ times the value of all executed retail transactions. Also, the weekly real and nominal GDP is computed after all trading takes place as well as monthly averages of real and nominal GDP subsequently used by the central bank at the monetary policy stage.

The stage ends with updating of the average weekly wage rate, fire-sale and haircut prices. The average wage rate W is just the total wage bill paid by the shops this week divided by total employment. The fire-sale price for next week is computed as follows:

$$P_f = \frac{1}{2} \cdot W \cdot (1 + \pi_w^*).$$

The haircut price for next week is set equal to

$$P_h = h \cdot W \cdot (1 + \pi_w^*),$$

where h is the loan-to-value ratio, i.e., the discount that banks use to value the inventories used as collateral.

The key parameters introduced in this stage of the weekly activities is the demand parameter ε, determining the proportion of the weekly planned expenditures allocated to each of a transactor's consumption goods, the critical ratio of inventories to target sales which triggers a layoff IS and the loan-to-value ratio h.

3.5 Monetary Policy

Next the program comes to the monetary policy stage at which the central bank sets the nominal interest rate i. First, it checks whether this is a fixed action date, which is true every fourth week. If not, this stage is skipped and the interest rate remains unchanged. If it is a fixed action date, then the central bank calculates average real GDP per week (the sum of each shop's input in excess of its fixed cost, over the past month, divided by 4, the number of weeks per month) and the current price level (GDP deflator). The central bank also keeps track of the values of year-to-year inflation factors, price levels, and average real weekly GDP for the last 12 months. Besides that, the central bank computes the 3-months moving averages for the weekly average GDP series and subsequently uses this information to estimate the actual output gap.

The central bank sets the per annum rate of interest i according to the following Taylor rule:

$$\ln(1 + i) = \ln(1 + i^*) + \gamma_\pi[\ln(1 + \pi) - \ln(1 + \pi^*)] + \gamma_q[q - q^*], \tag{1}$$

where γ_π and γ_q are fixed coefficients, $1+\pi$ is the inflation factor over the past 12 months, π^* is the fixed inflation target, q is the central bank's estimate of the current log output gap[13], and $r^* \equiv i^* - \pi^*$ and q^* are evolving targets for the long-run real interest rate and the log output gap, which are adjusted over time.[14] We also assume that the central bank respects the zero lower bound on nominal interest rate. The weekly interest rate is determined according to $1 + i_w = (1 + i)^{1/48}$.

The targets q^* and r^* are reset during the monetary policy stage in the following manner. The target log output gap q^* is adjusted according to a simple rule:

$$\Delta q^* = \eta_q (q - q^*),$$

where η_q is a fixed target output gap adjustment coefficient. The interest rate target is adjusted according to:

$$\Delta r^* = \eta_r (\pi - \pi^*) \cdot f(\pi, r^*), \quad f(\pi, r^*) \equiv \frac{r^*}{\sqrt{\eta_r^2 (\pi - \pi^*)^2 + (r_0^*)^2}},$$

where η_i is a fixed target interest rate adjustment coefficient, r_0^* is the initial target real interest factor, and $f(\pi, r^*)$ is a "squasher." So, the interest rate target is increased if current inflation rate exceeds the central bank's target and decreased, otherwise. The squasher makes sure that this change is symmetric (S-shaped) around the point $\pi = \pi^*$ and does not exceed the current target, r^*, in absolute value. Around the point $(\pi, r^*) = (\pi^*, r_0^*)$, $\Delta r^* \approx \eta_r (\pi - \pi^*)$, i.e., the adjustment of the target interest rate is roughly proportional to the deviation of the actual inflation from its target.

At this point the central bank also announces its projections of future inflation and future nominal interest rates. It does this by assuming that q and $\ln(1+\pi)$ will approach their long-run targets in yearly movements according to simple autoregressive processes:

$$\begin{aligned} &q_{+12} - q^* = \hat{\lambda}_q (q - q^*) + \xi_q, \\ &\ln(1 + \pi_{+12}) - \ln(1 + \pi^*) = \hat{\lambda}_\pi (\ln(1 + \pi) - \ln(1 + \pi^*)) + \xi_\pi, \end{aligned}$$

where $\hat{\lambda}_q$ and $\hat{\lambda}_\pi$ are current estimates of the corresponding evolving autocorrelation coefficients which are updated each month, and ξ_q and ξ_π are random i.i.d. shocks.[15] These parameters are recalculated in the following way. Let $\hat{\lambda}_{q,t-1}$ be the central bank's

[13] The estimate of the current gap is computed as the ratio of the current 3-months moving average of real GDP per week to capacity GDP, $y^{\text{cap}} = n(n - 2 - F)$, which is the weekly level of output when everyone is working or operating a shop and there is exactly one shop per good.

[14] We assume that the central bank begins to adjust these targets after some initial learning period, T_{cb}.

[15] Monthly re-estimation of these coefficients also starts after some initial learning period T_{cb}, as in the case of interest rate and output gap targets.

estimate in month $t-1$. Then, in period t, the target change in estimates is calculated as[16]

$$\Delta_{\lambda,t} = \frac{(q_t - q_t^*)(q_{t-1} - q_{t-1}^*) - (q_{t-1} - q_{t-1}^*)^2 \hat{\lambda}_{q,t-1}}{\sum_{k=1}^{t}(q_k - q_k^*)^2},$$

Next, the new estimate is set equal to

$$\hat{\lambda}_{q,t} = \hat{\lambda}_{q,t-1} + \Delta_{\lambda,t} \cdot f(\Delta_{\lambda,t}, \hat{\lambda}_{q,t-1}), \quad f(\Delta_{\lambda,t}, \hat{\lambda}_{q,t-1}) \equiv \frac{1 - |\hat{\lambda}_{q,t-1}|}{\sqrt{\Delta_{\lambda,t}^2 + (1 - |\hat{\lambda}_{q,0}|)^2}},$$

where $\hat{\lambda}_{q,0}$ is the initial estimate of the respective autocorrelation coefficient, and "squashing" the change guarantees that the new estimate does not exceed 1 in absolute value. The coefficient $\hat{\lambda}_{\pi}$ in the inflation equation is adjusted in the same way.

The central bank's projections are subsequently used in computing the capitalization factor V employed by transactors when formulating their expenditure plans.

The key parameters introduced in this section are the Taylor rule coefficients γ_π and γ_q, the target inflation rate π^*, the interest rate and output gap target adjustment coefficients η_r and η_q, and the capital adequacy ratio κ.

3.6 Match Breakups

The program next simulates the random breakup of established trading relationships in the economy. In particular, each transactor in the economy who does not own a shop is subjected to a probability δ of quitting the labor and goods markets, which entails the unconditional severance of all current trading relationships by the transactor with his employer as well as his consumption stores.

3.7 Fiscal Policy

Next comes the stage where the retail sales tax rate τ is adjusted. This happens only once a year, in the last week of the year. In all other weeks this stage is bypassed and τ remains unchanged. When deciding on the new tax rate the government first calculates the size of the government debt (normalized by the price level) relative to annual capacity GDP times the target output gap. It then sets the tax rate equal to a value τ^* which is the value that would leave the debt/GDP ratio undisturbed in the unshocked equilibrium to be described below, plus an adjustment factor that is proportional to the difference

[16] This procedure is motivated by a recursive formula for updating the OLS estimates of the autocorrelation coefficients.

between the actual and the target debt to GDP ratio b^*:

$$\tau = \tau^* + \lambda_\tau \cdot \left(\frac{B}{(1+i_w)P(48 \cdot y^{\text{cap}} \cdot e^{q^*})} - b^* \right),$$

where B is the total stock of issued government bonds, P is the current price level, and λ_τ is the adjustment coefficient.

Thus, there are two new key parameters introduced in this section, the target debt-GDP ratio b^* and the fiscal adjustment speed λ_τ.

3.8 Firm Exit

Firms can exit for several reasons. First, each unprofitable shop owner has an opportunity to close down his shop with probability ϕ. The condition of unprofitability is met if, given current wealth and inventories, the shop owner is better off closing down his business and working for the economy-wide aggregate wage rate W. The precise conditions under which a firm is unprofitable[17] will be spelled out below. Second, a shop owner can decide to exit the market if he is hopeless about his business prospects, i.e., if, even borrowing the allowed maximum, he won't be able to cover the fixed cost of operating a shop. Finally, a firm exits if it is bankrupt.

Once a shop exits for any of the above reasons, all trading relationships (with both employees and customers) are severed and the shop owner has to settle with the bank. Consider first the recourse loan case. If the total wealth of the shop owner is positive, he pays off the whole loan to the bank, and, if the deposit holdings are lower than the outstanding loan, cash is used to cover a part of the loan. If the shop owner is not able to cover the whole loan, the bank estimates the value of the shop's inventories at the current firesale price. If the total wealth plus the value of inventories is enough to cover the whole loan, the bank seizes the needed amount of inventories. Otherwise, it seizes everything it can from the shop owner. In the non-recourse loan case the bank can not seize any assets beyond the collateral, which was pledged against the loan at the haircut price.

In case if there are some legacy inventories left after settling with the bank, the exiting shop owner is added to the fire-sale market queue corresponding to his supply good. If the bank seizes part or all of the inventories from the shop owner, the bank is also added to the fire-sale market queue if it is not already there with the same type of supply good (if the bank is there already, it keeps the previous position but the size of the order

[17] Clearly, whether we are in the case of recourse or non-recourse loans affects the calculations of the shop owner. Non-recourse loans make the exit more likely because of the limited liability of the borrower.

increases by the amount of seized inventories).

3.9 Updating of Targets, Wages, and Prices

In the final stage of weekly activities, the program first updates each shop's target sales setting them equal to current period's actual sales. Then it proceeds to update the shop's wages but does so only if the current week happens to be a "wage-updating" week for that shop, i.e., if Δ weeks have elapsed since the shop's most recent wage update or if the current week is the shop's very first wage-updating week, which is a random realization from the discrete uniform distribution over $[1, \Delta]$ assigned when the shop first opens. The random assignment of initial updating weeks implies that wage setting will be "staggered," and the fraction $1/\Delta$ of all wages will change in a average week.

Given that the current week is indeed a wage-updating week for the shop being considered, its wage and price updates are performed in the following manner. First, a desired annual real wage adjustment figure is calculated as a proportion β of the percentage deviation in the shop's target employment from its actual employment, based on all current relationships with employees regardless of whether they were paid or not.[18] The parameter β hence indexes the degree of wage and price flexibility in the economy. This annual wage adjustment figure is then rescaled to obtain a corresponding adjustment for the length of the contract period, taking into account the shop owner's estimate of inflation during this period as derived from the weekly target inflation rate. This rescaled wage adjustment figure is then applied to the current wage offered by the shop and its retail price (net of taxes) is correspondingly altered to retain the markup over the wage.

Independently of whether the week is a wage-setting period, each shop can apply a lower than usual markup (put on a "sale") or a higher than usual markup (engage in "gouging") depending on the ratio of its current inventories to this period's actual sales. If this ratio is higher than the layoff threshold, IS, the shop cuts its normal price (multiplying it by $0 < \delta_p < 1$) to the sale price. If this ratio is less than the threshold level $1/IS$, the shop increases the normal price (dividing it over δ_p), i.e., starts gouging.

The key parameters introduced in this final stage of weekly activities include the speed of wage adjustment, β and the sale price discount δ_p.

[18] In computing this expression we use the maximum of actual employment and the shop's fixed cost F to avoid division by zero when actual employment falls to zero.

3.10 Simulating the model

The entire run of the algorithm over T weeks is then repeated for R different runs, where a run always starts off near the flexible-price, no-shock equilibrium of the economy (see below). Each run, however, is unique in the initial seeding of the computer's random number generator for that run. This allows us to exploit the randomness that is built into the system by ultimately enabling us to examine the average performance of the system across different realizations of the economy's behavior over T weeks.

4 The workings of the model

4.1 Equilibrium with price flexibility and no shocks

As the preceding discussion has made clear, all shocks in this economy are individual shocks. Unlike in the standard New Keynesian framework, we have postulated no exogenous shock process impinging on aggregate productivity, price adjustment, aggregate demand, monetary policy or fiscal policy. Nevertheless, the individual shocks that cause matches to break up and shops to enter or leave particular markets do have aggregate consequences because there is only a finite number of agents. So in general the economy will not settle down to a deterministic steady state unless we turn off these shocks. However if we do turn off all shocks, there is a deterministic equilibrium that the economy would stay in if left undisturbed by breakups and entry if the inflation target π^* were equal to zero and the output gap were also equal to zero. Moreover, if the contract period Δ were equal to one week, the economy would remain in this equilibrium for any positive rate of inflation. This equilibrium will serve as an initial position for all the experiments we perform below, and a brief description of it helps to illustrate the workings of the model.

The equilibrium is one in which all the potential gains from trade have been exhausted. Each transactor is matched with one employer and two stores. There are n shops, one trading in each of the n goods. Each shop begins each week with actual, potential and target input all equal to $n-2$, which is the number of suppliers of each good, and with actual and potential sales equal to inventory holdings equal to its actual output: $n-2-F$. So the economy's total output equals full capacity: $y^* = n\,(n-2-F)$.

Each shop begins each week with a wage rate equal to $W = (1+\pi_w^*)\,W_0$, which is the same for all shops, where W_0 was the common wage rate last week, and with a retail price equal to $P = W(1+\mu)/\,(1-\tau)$ where the tax rate τ equals:

$$\tau^* = 1 - (1 + \pi_w^*)(1 - 48\rho_w b^*) / \left(1 - \pi_w^* \frac{n-3}{(n-2-F)(1+\mu)}\right).$$

In this no-shock equilibrium all shops are self-financing, and banks are just conduits, converting deposits into government bonds. The initial outstanding stock of bonds is

$$B = b^* (1 + i_w) 48 y^* P_0,$$

where $P_0 = P/(1 + \pi_w^*)$ is last week's price level and where the weekly interest rate i_w is given by:

$$1 + i_w = (1 + \rho)^{1/48} (1 + \pi_w^*).$$

The money supply at the start of the week is

$$M = W_0 (N - n) + (1 - \tau) P_0 y^*,$$

which is the sum of all wage receipts of non-shopowners, and all sales receipts (ex taxes) of shopowners, from last period.

Each agent starts the period with an effective supply price equal to W_0 and both effective demand prices equal to P_0. The owner of each shop starts with a permanent income equal to last period's profit: $(1 - \tau) P_0 (n - 2 - F) - W_0 (n - 3)$ and with money holding equal to last period's revenue: $(1 - \tau) P_0 (n - 2 - F)$. Each of the non-shopowner transactors begins with money holding equal to permanent income, which in turn is equal to last period's wage income W_0.

The aggregate bond supply B is assumed to be initially distributed across agents in proportion to their initial money holdings.

The initial history is one in which the output gap has been equal to zero for the past 12 months and inflation has equaled its target rate for the past 12 months.

It is straightforward to verify that if prices are changed every period this configuration will repeat itself indefinitely, except that all the nominal maginitudes – money and bond holdings, prices and permanent incomes – will rise at the weekly target inflation rate π_w^*.

4.1.1 Entry, exit, and systemic performance

As we shall see below, the economy is able achieve on average about 90% of the potential GDP attained in the no-shock equilibrium. GDP goes down whenever a shop that was satisfying some consumers goes out of business or a customer loses a store because of a random breakup. GDP also goes down whenever a new shop enters and diverts workers from old shops that were satisfying some customers, because some of these workers' efforts will be used up in deferring the fixed cost of the new shop rather than producing goods that can be eaten by customers of the old shop.

These events that reduce GDP are constantly being offset to some degree by the growth of new firms that are able to satisfy customers in markets where there had previously been no viable firm, and by the exit of firms that were using up fixed costs but not producing enough to satisfy their customers. Thus both entry and exit are critical to the system's ability to approximate full capacity utilization. However, as described in the introduction to the paper, and as should be clearer now that the details of the model have been described, although entry of new firms is useful in markets where there are no incumbents or where the incumbents are not hiring all the potential workers because of layoffs or because of financial problems that prevent them from meeting their payroll, entry can be harmful in cases where incumbent firms were hiring most of the potential workers and satisfying most of the potential customers. Likewise, although exit is important in cases where the firm has ceased to play an active intermediation role, either because financial difficulties or a surfeit of inventories prevent it from hiring many workers or because its high markup has driven customers away to neighboring markets, exit can be very harmful in cases where the incumbent was previously doing well, because it can cascade across markets causing a cumulative loss of output.

In this model banks play an indirect role in promoting entry, in the sense that the successful entrant will typically not have inventories and may not have enough personal wealth to purchase inventories without the aid of a bank loan, and a shop without inventories cannot generate revenue. Thus in the absence of bank financing, new entrants would be more prone to fall quickly into hopelessness and thus to disappear before having a chance to grow.

The exit decision can be described simply as follows. (1) Bankruptcy will force a shop into quitting. (2) Likewise if a shop can no longer afford to hire enough labor to cover its fixed cost next period it will quit with certainty (the hopeless case). (3) If the following condition is satisfied the firm will decide that it is no longer profitable to stay in business and thus will quit with probability $\phi > 0$. In the baseline case where all loans are with full recourse, the condition is:

$$\begin{array}{llllll} \text{(a)} & A + P_f I \geq 0 & \text{and} & WV + P_f I & > & \Pi^e V, \text{ or} \\ \text{(b)} & A + P_f I < 0 & \text{and} & WV & > & \Pi^e V + A \end{array}$$

where, as above, A is the financial wealth of the shopkeeper, I is the shop's inventory, W is the current economy-wide average wage, P_f is the firesale price, V is the capitalization factor and Π^e is the shopkeeper's current permanent income (expected profit from staying in business).[19] Under no recourse the condition changes to:[20]

$$\begin{array}{llllll} \text{(a)} & P_h I \geq L & \text{and} & WV + P_f I & > & \Pi^e V - L\left(1 - \frac{P_f}{P_h}\right), \text{ or} \\ \text{(b)} & P_h I < L & \text{and} & WV & > & \Pi^e V - L \end{array}$$

Banks affect these exit decisions in different ways. First, by allowing shops to borrow against inventories the bank will delay the time when the shop can no longer afford to hire enough workers to cover its fixed costs. Note that through this channel, a bank that falls into trouble because of capital inadequacy is likely to induce an increased hazard of shop failure in its sector because it will no longer make new bank loans to keep shops away from hopelessness.

Second, in the case of no recourse banks actually encourage exit because they allow the shopkeeper to escape the burden of a loan without having to sacrifice her personal financial wealth. (Hence the negative terms in L on the right hand side of the second pair of inequalities above.) Indeed it is easy to see that in any given situation, nonprofitability under recourse implies nonprofitability under non-recourse but not vice versa. The fact that shops have a bigger incentive to exit under non-recourse may help to explain some of the results to be presented below.

5 Calibration

Although the model has many agents we have imposed a great deal of ex ante symmetry. We have done this partly so that we can fully characterize the unshocked equilibrium that serves as a reference point, which will facilitate our analysis of what is generating our experimental results, and also partly so that we can keep the number of parameters small enough to calibrate them to US economic data. This section describes our calibration procedure.

[19] In case (a), the shopkeeper would be able to repay her loan in full, either with her money and deposit holdings or with giving up some of her inventories. So in the event of failure her tangible plus human wealth would go from $A + \Pi^e V$ to $A + P_f I + WV$ on the assumption that she could get a job with wage W. She will quit with probability ϕ if the latter exceeds the former. In case (b) under full recourse she would be unable to repay her loan in full, so upon quitting her tangible plus human wealth would go from $\Pi^e V + A$ to WV, so again she will quit with probability ϕ if the latter exceeds the former.

[20] In case (a) the loan can be repaid by having the bank seize the collateral, which is just L/P_h, so total wealth on quitting would go from $\Pi^e V + A$ to $WV + A + L + P_f (I - L/P_h)$ (the first term in L on the RHS comes from being relieved of the loan while the second one is part of the firesale value of the exiter's legacy inventory). In case (b) there will be no legacy inventory so total wealth goes from $\Pi^e V + A$ to $WV + A + L$.

There are a total of 34 parameters, which we have categorized as shop parameters, transactor parameters, bank parameters, and government parameters. These are listed in Table 5 along with their assigned values.

Our calibration of these parameters took place at three different levels. At the first level, one subset of parameter values was chosen to match empirical counterparts in the US data. At the second level, the values of other parameters were chosen so as to be internally consistent with average simulation outcomes. At the third level the values of the remaining parameters, for which we could find no convenient empirical counterparts, were chosen to make the median simulation outcomes match (loosely) certain properties of the US data.

5.1 First level of calibration

5.1.1 Shop parameters

Smets and Wouters (2007) estimate the average duration of wage contracts in the US to be about a year. This is consistent with evidence from other studies cited in Amano et al (2009, section 4). Accordingly we set the length of contract period Δ to 48 weeks, which in our model is one year.

Estimates of the degree of returns to scale in the US economy vary from 0 to about 30 percent. This is commonly measured as the ratio of average to marginal cost (minus unity). In our model the typical firm in a steady state with input equal to x and sales equal to $x - F$ would thus have a degree of returns to scale equal to

$$\frac{AC}{MC} = \frac{Wx/(x-F)}{W} = \frac{1}{1-F/x}$$

If the economy was operating with a 5 percent average unemployment rate then the typical firm would have

$$x = 0.95 \cdot (n-2) = 45.6$$

so by setting the fixed cost F equal to 3.5 we get a typical degree of returns to scale equal to 8.3 percent.

The inventory adjustment speed $\lambda_I = 0.17$ corresponds to the estimate by Durlauf and Maccini (1995) of a monthly adjustment speed equal to approximately 0.5 ($\approx 1-(1-0.17)^4$).

Roberts (1995) estimated aggregate expectations-augmented Phillips relations with a coefficient on detrended output between 0.25 and 0.334 using annual data. A linear approximation to our wage-adjustment equation yields a same relation if we assume

TABLE 1

	Shop parameters	
$\bar{x}$	Maximal target sales	40
Δ	Length of the wage contract period (in weeks)	48
$\bar{\mu}$	Average percentage markup over variable costs	0.15
ϕ	Failure rate of unprofitable shops (weekly)	0.01
F	Fixed cost	3.5
λ_I	Inventory adjustment speed (weekly)	0.17
β	Wage adjustment coefficient (annual)	0.3
IS	Critical inventory/sales ratio triggering a layoff	3.0
δ_p	Size of price cut (old price/new price)	1.02
	Transactor parameters	
ε	Demand parameter	7.0
λ_y	Permanent income adjustment speed (weekly)	0.2929
ρ	Rate of time preference (annual)	0.04
θ	Frequency of innovation (weekly)	100
δ	Quit rate (per week)	0.001
σ	Job search probability	0.5
	Bank parameters	
l	Length of initiation period for new customers (in weeks)	48
s	Normal (annual) loan spread	0.0175
$\tilde{s}$	Extra (annual) spread for new customers	0.005
h	Loan-to-value ratio	0.5
	Policy parameters	
	Fiscal Policy	
b^*	Target Debt-GDP ratio	0.33
λ_τ	Fiscal adjustment speed (annual)	0.054
	Monetary Policy (Taylor rule)	
π^*	Annual target inflation factor	0.03
γ_π	Inflation coefficient in Taylor rule	1.5
γ_q	Output gap coefficient in Taylor rule	0.5
q_0^*	Initial target logarithmic output gap	0.055
r_0^*	Initial target real interest factor	0.035
η_q	Adjustment speed of evolving gap target	0.015
η_r	Adjustment speed of evolving real rate target	0.0075
$\hat{\lambda}_{\pi,0}$	Inflation autocorrelation factor (initial estimate)	0.29
$\hat{\lambda}_{q,0}$	Output gap autocorrelation factor (initial estimate)	0.66
T_{cb}	Number of years before central bank learning begins	15
	Bank regulation	
m	Number of banks	5
κ	Required capital adequacy ratio	0.08
s_d	Premium on Fed's discount rate	0.005

that actual/capacity output ratio is proportional to the target/potential input ratio. Accordingly we chose $\beta = 0.3$ to lie near the midpoint of Roberts' range of estimates.

5.1.2 Transactor parameters

We set the annual rate of time preference ρ equal to 0.04 as is standard in the real business cycle literature. We chose the demand parameter ε to equal 7, which implies that in a no shock equilibrium with all shops charging the same price the elasticity of demand facing each shop would be $1+\varepsilon/2 = 4.5$. This lies in the range of values typically found in New Keynesian DSGE models[21]. (The elasticity of demand faced by a shop out of equilibrium, when he has rivals selling the identical good, will however be larger than 4.5 because raising the retail price may induce a loss of all demand from any customer that finds another shop during the matching process.)

5.1.3 Bank parameters

The value of normal loan spread, s, was set equal to 0.0175 which is the average spread between lending and deposit rates for all commercial and industrial loans during the period 1986–2008. This figure comes from the Survey of Business Lending Terms conducted by the Federal Reserve.

To take into account the value of bank-firm relationship in our analysis we introduce the extra spread on commercial loans for new customers. In doing this we rely on the empirical literature on relationship lending. Berger and Udell (1995) show that banks offer lower lending rates to firms with which they have longer relationship. In particular, a firm with an 11-year banking relationship can expect to pay a loan premium 48 basis points less than a firm that has only a 1-year relationship, controlling for other characteristics. There is no agreement about how exactly the terms of lending depend on the length of the relationship. To formulate the behavior rule for banks, we follow Cole (1998) who reports that there is no incremental effect from pre-existing relationships of longer duration than 1 year, and set l, the length of initiation period for new customers, equal to 48 weeks. Using the estimate of Berger and Udell (1995) as a baseline, we assume that new customers pay an extra premium of $\tilde{s} = 0.005$ per year on their loans.

[21] Demand elasticity is set equal to 3 in Midrigan (2009), 4 in Nakamura and Steinsson (2008), 7 in Golosov and Lucas (2007), and 11 in Yun (2005). Calibration in Burstein and Hellwig (2008) yields the elasticity of 4.4.

5.1.4 Policy parameters

Fiscal Policy. The target Debt-GDP ratio b^* was set equal to 0.33 because this is the average ratio of federal marketable debt to GDP in the US on average between 1969 and 2005. The fiscal adjustment speed λ_τ was estimated at 0.054 by Bohn (1998).

Monetary Policy. The initial estimates of the autocorrelation factors $\hat{\lambda}_{\pi,0}$ and $\hat{\lambda}_{q,0}$ were taken from estimates of univariate AR(1) processes on inflation and on (linearly-detrended) log per-capita GDP using annual data for the US over the 1984-2006 period. The coefficients γ_π and γ_q are Taylor's original specification. In the calibration exercise to be described immediately below we took the inflation target π^* to equal 3%, which is the average in the US over the period from 1984-2006.

Bank Regulation. We set the required capital adequacy ratio κ equal to 0.08 which corresponds to Basel I bank regulation. The central bank discount rate premium s_d is set to 0.005, which approximates the typical spread of the Federal Reserve's primary credit discount rate over the Fed Funds rate.

5.2 Second level of calibration

5.2.1 Government targets – Finessing Wicksell

The two government targets – the target real interest rate r^* and the target output gap q^* are chosen adaptively by the central bank. We chose their initial values (.035 and .055 respectively) close to the steady state values in our calibration simulation, although since central bank adaptation was relatively quick and our learning period was quite long, our results were not sensitive to the choice of initial values. Note that our procedure of having central banks estimate r^* and q^* forces it to deal with the danger that writers from Wicksell to Friedman and up through Orphanides have warned of - the danger that no one knows the economy's natural rate of interest or the steady-state level of the output gap, and hence that controlling the rate of interest to the neglect of the money supply risks aggravating volatility or even engendering a cumulative inflation.

5.2.2 Markups and the Lucas Critique

In early trials with the model we assumed that all shops applied the same markup. But we found that the results of many experiments were highly sensitive to the assumed size of the markup. Awareness of the Lucas critique prompted us to revise the model in favor of the current assumption, namely that each shop picks its markup at birth. This variant allows the economy-wide average markup to respond endogenously to the policy environment, through the evolutionary selection mechanism implicit in the exit process. We chose the mean of the distribution from which markups are drawn in the same way that we decided the initial values of the government targets r^* and q^* – by internal consistency. In our baseline simulations the median markup is about 16 percent when shops choose from a distribution whose mean $\bar{\mu}$ is 15 percent. (We will calibrate this a little finer in the next draft.)

5.3 Third level of calibration

This leaves 13 parameters still to be determined, namely $\overline{x}, \phi, \theta, \delta, \lambda_y, \sigma, IS, \delta_p, h, m, T_{cb}, \eta_q$ and η_r, were chosen by searching (manually) for parameter values that would generate values of thirteen different indicator variables, across 5,000 simulations of 40 years each, that approximated their counterparts in US data. More specifically, we started each simulation near the no-shock equilibrium and then allowed for 30 years to pass before calculating the average value of each variable across all the years of that simulation. For each variable we then computed the median across all simulations of these across-year averages.

The thirteen indicator variables are listed in Table 2 below, along with their data values and the median values in our fully calibrated model averaged over 5,000 simulations.

TABLE 2

Variable	Data	Model
Inflation	3.0	2.9
Real interest rate	1.8	3.3
Unemployment rate	5.0	6.1
Unemployment duration	14	11
Job loss rate	.69	.59
Volatility of output gap	1.6	2.8
Volatility of inflation	1.2	.74
Autocorrelation of gap	64	43
Autocorrelation of inflation	9	20
Average markup	10 to 20	16
Exit rate	46	40
Price change frequency	4.0	4.1
Annual bank failure rate	0.51	.50

All numbers are expressed in percent, except for unemployment duration which is expressed in weeks, and price change frequency which is expressed in numbers per year. The real annual interest rate is computed as difference between the annual interest rate on 3-month T-bills (monthly data from the Federal Reserve) and CPI inflation rate (monthly data from the U.S. Bureau of Labor Statistics) averaged over the period 1984-2006. The unemployment rate is the average over the period from 1984 to 2006. The volatility of the output gap is the standard deviation of linearly-detrended log per-capita GDP over the same period and the volatility of inflation is the standard deviation of annual US CPI inflation over the same period. Autocorrelations of these two variables are computed by estimating an AR1 process over the same time period. Golosov and Lucas (2007) indicate that estimates of the percentage markup vary between 10 and 20 percent. The exit rate is the fraction of all firms found operating in a given industry in one census year that are found still operating in that industry in the next census year (five years later), which Dunne, Roberts and Samuelson (1988) report to be 46.4 percent. The job-loss rate is the weekly rate of job loss that would give rise to the number reported by Hall (1995), namely that 71.8% of people working at a given date have been working continuously for the same employer for the past year.[22] Bils and Klenow (2004) find an average price-change frequency of 16 weeks, which in our model would imply an average annual price change frequency of 3. However, in our model firms will almost always change prices every time there is a change in the sales tax, which is once per year, whereas in reality sales tax changes are very infrequent, so we aimed to match a price change frequency of 4. According to FDIC (Historical Statistics on Banking), for the period 1984-2006 the average commercial bank failure rate was about 0.51 percent per year.

As the above table shows, we were at best partially successful in mimicking the data with these 13 parameters. We do not mind that the model overpredicts output volatility and unemployment, however, since the data period to which we are comparing is that of the "great moderation," which in retrospect is looking increasingly like an abnormally quiescent period. The model's volatility is not excessive in comparison to the US economy over the longer period from 1969 to 2006, over which period the standard deviation of the detrended gap was 2.9, exactly the same as in the median run of our model. In the next draft we will attempt to come somewhat closer to the data, and will conduct sensitivity analysis on these thirteen variables to see to what extent our results depend on some of our calibration choices.

[22] That is, $(1-\alpha)^{48} = 0.718$ if $\alpha = 0.00688$.

5.4 Approximating a steady state

Although each of the data points described in the above table is an average over 5,000 40-year (1920-week) runs, we actually allowed each run to proceed for 70 years (3360 weeks), and ignored the first 30 years of data when computing the averages. This was in order to give the economy time to approximate a steady state, having been started in the above-described initial position which was a steady state only under the assumption of no shocks and weekly price adjustment by all firms. As Figure 1 below shows, 30 years was indeed enough for the cross-run average values of the real interest rate, the output gap, the inflation rate and the markup to become more or less constant, except for the slight upturn of the average gap towards the end of the 70 year period, a feature that we discuss in the next section.

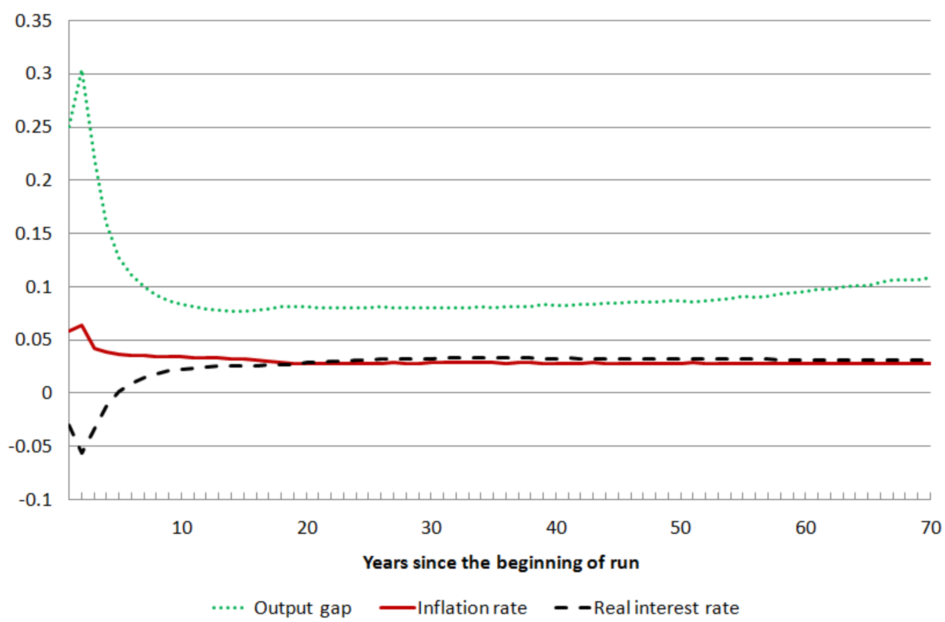

Figure 1: Average simulated values in baseline calibration

5.5 Normal times and bad times

Figure 1 above describes the average time path of 5,000 simulations, which seems to hone in nicely to a steady state. But in fact this is a weighted average of a lot of "normal" runs that exhibit strong homeostatic tendencies and a few "pathological" runs in which the market makers appear to have lost control of the system. As indicated above in our introductory remarks, this is a crucial nonlinear feature of the model, which seems to behave in a qualitatively different manner in bad times as compared to normal. To convey some idea of this qualitative difference, each of Figures 2 through 4 below depicts the actual time path of three major macroeconomic variables in one of the many "normal" runs.[23] A randomly chosen simulation would depict similar characteristics. There are times when the output gap rises for a few years but these times are eventually followed by a recovery.

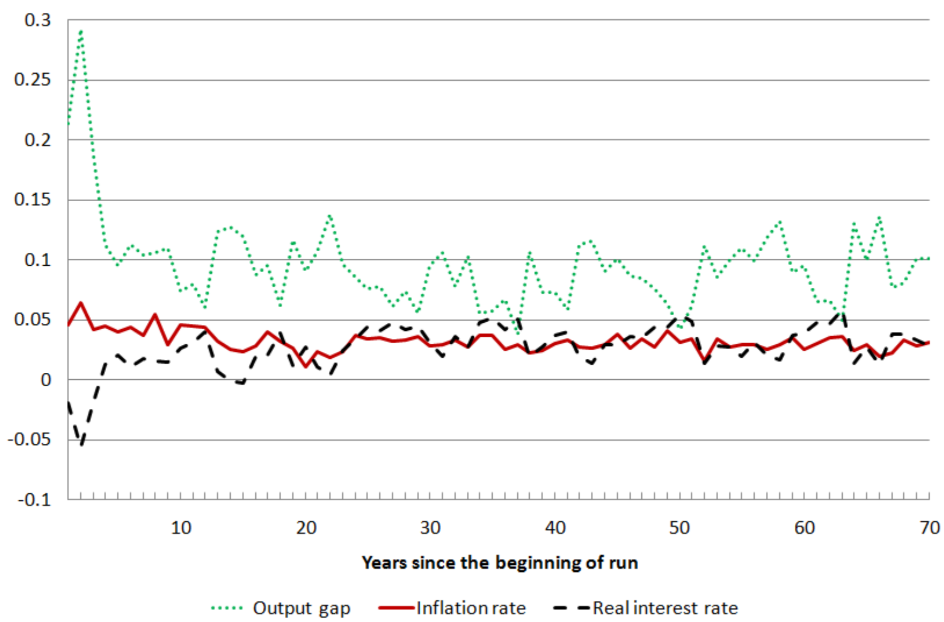

Figure 2: A normal run (rnseed=11)

[23]The rnseed number in each Figure caption indicates the initial seed value of the random number generator for that run.

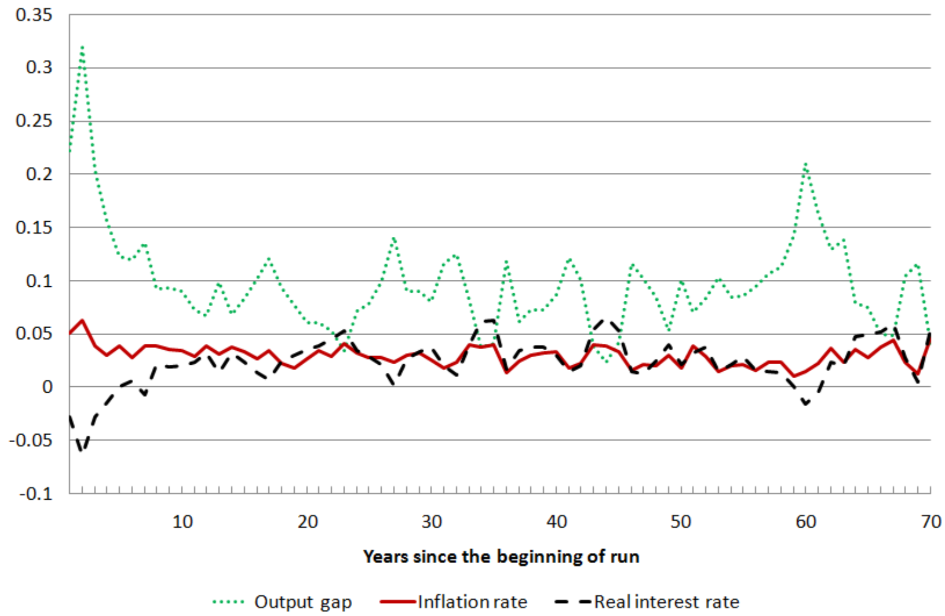

Figure 3: A normal run (rnseed=13)

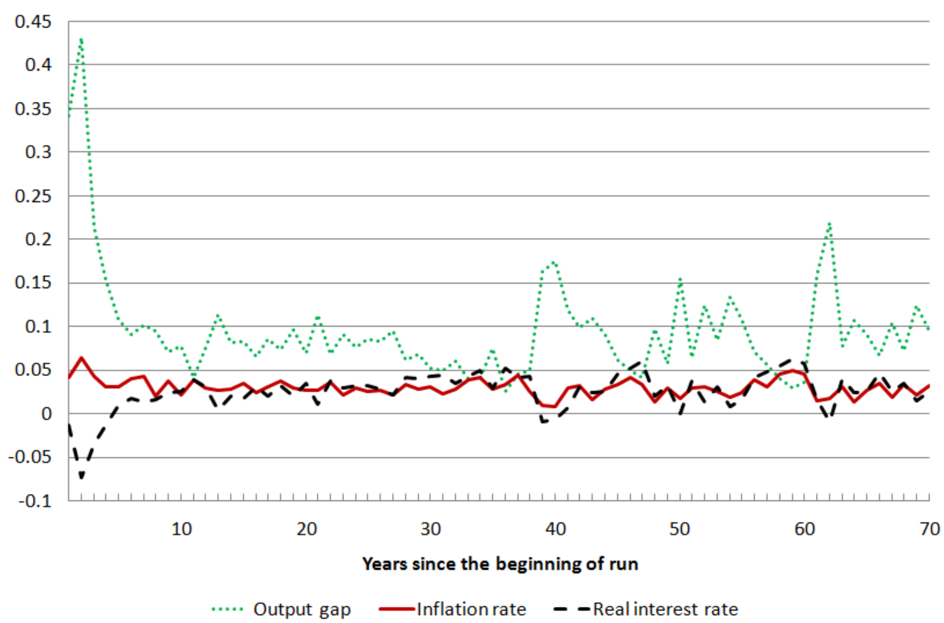

Figure 4: A normal run (rnseed=19)

In contrast with these normal runs, Figures 5 through 7 show what happens in some of the worst decile of runs. Again there are times when the output gap rises for a year or two, but beyond some point it ceases to return to a normal value and instead diverges with no apparent tendency to return. Although we have not investigated the matter in any rigorous statistical sense we doubt that such behavior could be produced by any known linear macro model.

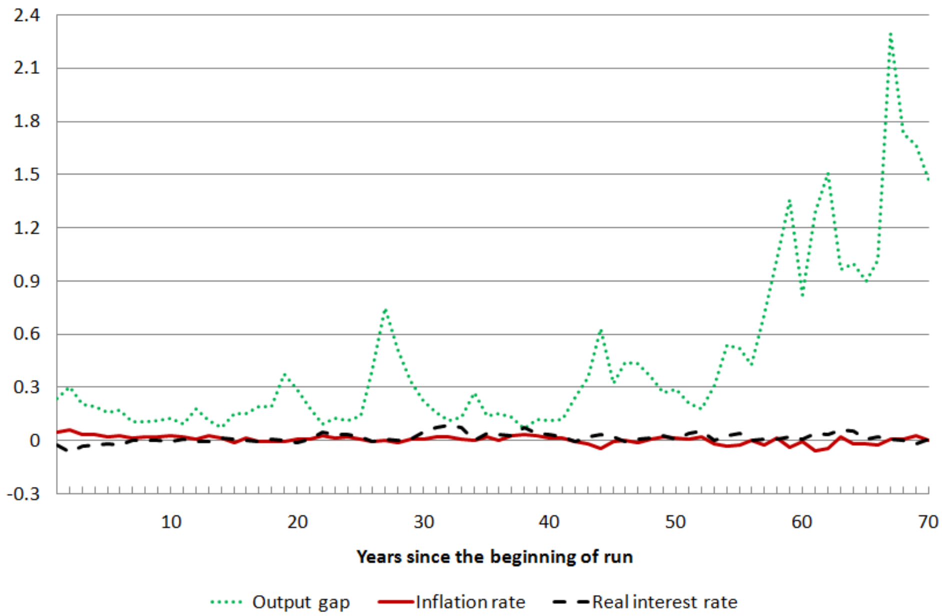

Figure 5: A collapse (rnseed=114)

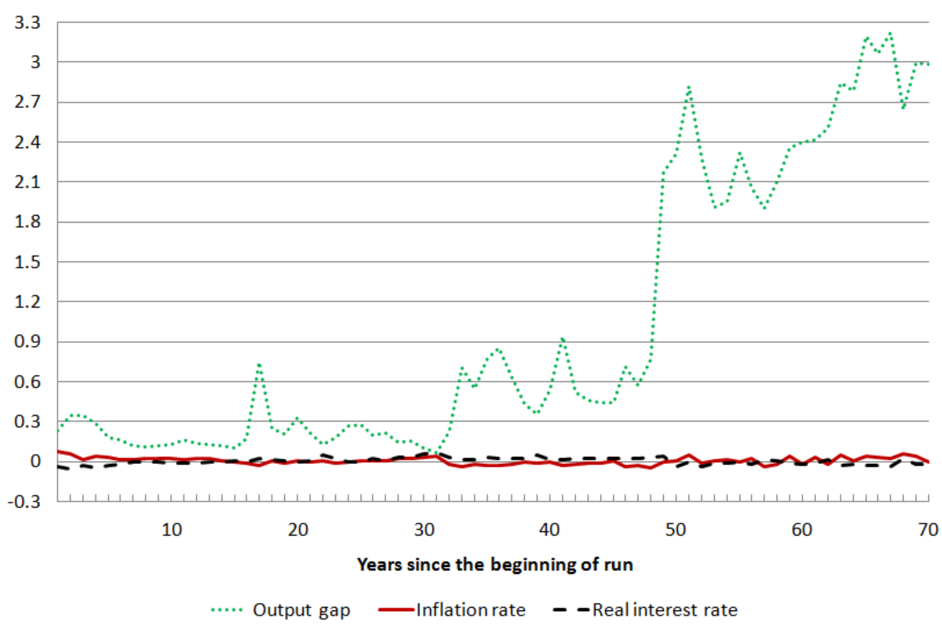

Figure 6: A collapse (rnseed=443)

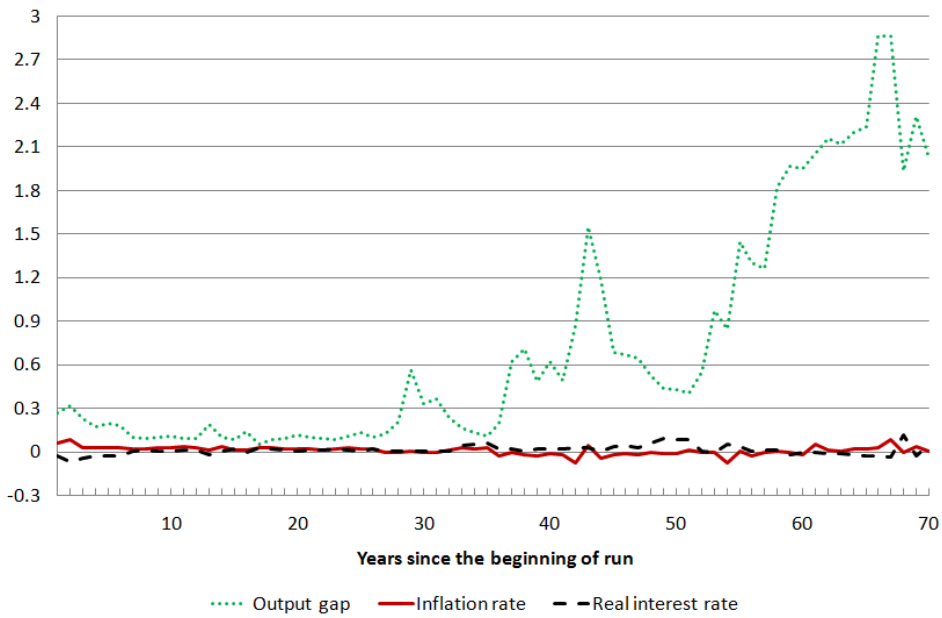

Figure 7: A collapse (rnseed=627)

The connections to Leijonhufvud's (1973) "corridor" are suggested by the fact that in all these cases of bad runs, the eventual collapse was preceded by one or two very large shocks, with the output gap rising above 0.6, much larger than any of the shocks seen in normal runs. In all the bad runs we have investigated so far the economy appears to be recovering from the large shock but then another shock comes and then an almost total collapse occurs. It seems as if the economy is capable of absorbing shocks up to some limit but not beyond.

It seems likely to us that the slight upturn in the average GDP gap seen Figure 1 above is the result of outliers like the runs depicted in the previous three figures, where the logarithmic gap tends to rise sharply in the second half of the run. In the next draft we will check this out by plotting the time series of the median gap, which we suspect will not show the same kind of uptick as the average gap.

6 Results

As indicated in the introduction to the paper, our results indicate that banks do matter for median economic performance, but not nearly as much as they matter when times are extremely bad, and not in the same direction as when times are extremely bad. We report below on the simulation experiments that give rise to this conclusion, on some related experiments that attempt to trace the source of these results, and on some experiments that suggest policies for alleviating the most harmful effects of bank lending in bad times, specifically (1) a restriction on loan-to-value ratios and (2) a restriction on dividend payments by banks.

6.1 Banks, no banks and "risky" banks

In order to see how much banks affect median performance we performed an experiment of shutting all banks down and rerunning the simulations underlying the above described calibration. To do this we simply changed the banks' behavior in the financial market stage of each week so that they always imposed a credit limit of zero on all customers. This turns banks into mere conduits for the private holding of government debt. In this experiment a bank's equity will always be equal to the deposit holdings of its owner, and its risk-weighted assets (loans and seized collateral) will always be zero, so the bank will never fail and will never be in trouble.

The results of this experiment can be seen in the first two columns of Table 3 below, which reports the median across simulations of the average across years of different performance indicators. As can be seen, all indicators show a deterioration in median performance, or at best no change, when banks are shut down. These median effects of closing down banks appear to be fairly small, far from the disastrous consequences that policy makers have been trying to avert during the present crisis, but they seem substantial in light of the limited role that we have assigned to banks in this exploratory analysis, namely that of financing trade inventories, which on average are no more than about one week's GDP.

TABLE 3

Median results

	Banks	No banks	Risky banks
Inflation	2.9	2.9	2.9
Output gap	7.6	8.4	7.9
Unemployment rate	6.1	6.7	6.2
Unemployment duration	11	12	11
Job loss rate	.59	.60	.61
Volatility of output gap	2.8	3.3	2.9
Volatility of inflation	.74	.92	.83
Annual bank failure rate	.50	0	13
Fraction of banks in trouble	3.1	0	56

We can get another measure of the importance of banks by simulating the economy's response to various exogenous shocks. Figure 8 below shows the average response of the economy to a shock in which one of the five banks is forced to be in trouble for one year. Specifically, at the beginning of year 30 in the simulation (that is, at the end of the adjustment period to a steady state) we impose on the first of the five banks the restriction that it cannot lend and cannot pay dividends to its owner, independently of its capital adequacy. We repeat this simulation 2 thousand times. For a counterfactual we perform the same two thousand simulations, with the same sequence of seeds for the random number generator, but without the shock that forces a bank into trouble. The dashed ("safe") line in Figure 8 indicates the average difference in log GDP between the shocked and counterfactual simulations on a monthly basis for ten years following the shock. Again the effect is small, never exceeding a third of one percent of the average counterfactual value, but it is persistent, and again its size must be considered in relation to the quite limited role that we have assigned to banks so far in our analysis.

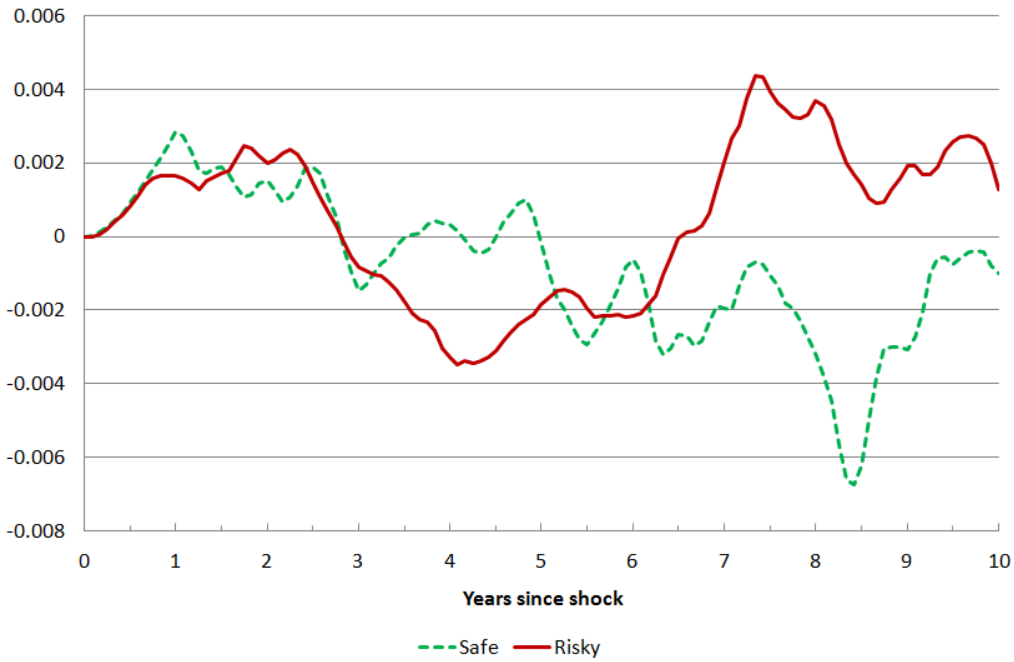

Figure 8: Response of GDP gap to troubled bank shock

6.1.1 Risky banks

To gauge how much difference bank regulation makes, we consider an alternative "risky" situation which differs from our baseline calibration in three senses. First, loans are assumed to be made without recourse. Given that without recourse shopkeepers can protect their personal wealth from creditors by minimizing their equity position, we suppose that during the financial market stage each week shopkeepers always borrow up their credit limit. Second, when the government seeks a new owner for a failed bank we assume that instead of choosing the richest non-shopowning customer of the bank it simply chooses such a customer at random to take over. Third, instead of limiting its customers with a haircut price equal to the firesale price (i.e., with a loan-to-value ratio of 0.5) each bank in this risky scenario allows a loan-to-value ratio of 0.9.

As the third column of Table 3 above indicates, macro performance under the risky scenario is somewhat worse than under our baseline safe scenario, although except for the fraction of banks in trouble and the hazard rate of bank failures the difference is quite small between the safe and risky scenarios. Likewise, as the solid line in Figure 8 above indicates, the response of GDP to an exogenous increase in troubled banks is not much different under the two scenarios.

6.1.2 Median versus worst-decile average performance

We have seen that the model implies rather small but qualitatively not surprising effects of banks, and of prudential regulations, on median macroeconomic performance. It turns out, however, that the effects on macroeconomic performance in hard times are quite large and qualitatively surprising, at least to us. Table 4 below reproduces the results of Table 3 but this time each number indicates not the median across all runs but rather the average over the worst ten decile of runs; that is, over the 500 runs with the highest cross-year average GDP gap.

TABLE 4

Worst-decile average results

	Banks	No banks	Risky banks
Inflation	2.1	2.3	2.4
Output gap	22	17	16
Unemployment rate	15	12	12
Unemployment duration	17	16	15
Job loss rate	1.2	.94	.88
Volatility of output gap	9.6	7.8	7.0
Volatility of inflation	1.4	1.4	1.3
Annual bank failure rate	1.0	0	19
Fraction of banks in trouble	27	0	68

As this table indicates, under the baseline (safe) scenario average macro performance in the worst decile of runs is worse than if banks were entirely shut down, in terms of all indicators except the mean and standard deviation of inflation. Clearly the worst-decile average inflation rate is better with banks simply because the central bank's Taylor rule is trying to fight the high output gap.

Moreover, average performance in the worst decile of runs is much better under the risky scenario than under the scenario which by conventional standards would seem much more prudent.

6.2 Experimenting with bad times

In search of some understanding for why banks make bad times worse and why safe banks make them worse than risky banks, we conducted an experiment in which we varied the loan-to-value ratio used by banks, from 0.5 (the "safe" value equal to the firesale discount) up to 0.9 (the "risky" value) by increments of 0.1. As the dotted line marked with diamonds in Figure 9 below indicates, in the baseline (safe) scenario, allowing banks to make much larger loans led to a dramatic deterioration of performance. The worst-decile average output gap increased from 21.8% to 31.4%. Figures 10 through 12 show the same sort of deterioration in terms of the unemployment rate, the volatility of output and the volatility of inflation.

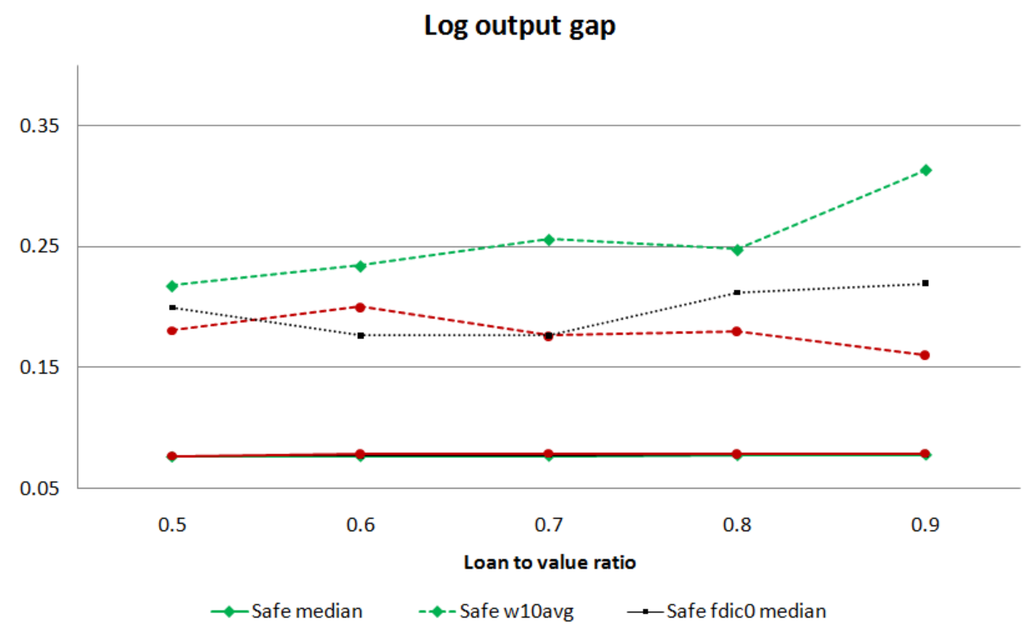

Figure 9: Effect of loan-to-value ratio on output gap

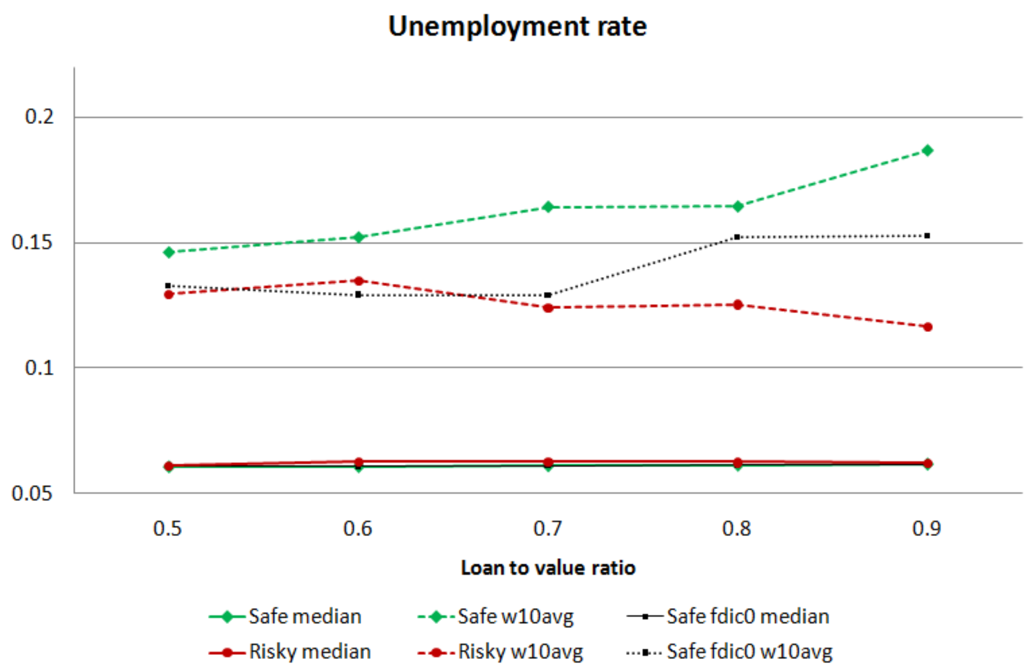

Figure 10: Effect of loan-to-value ratio on unemployment rate

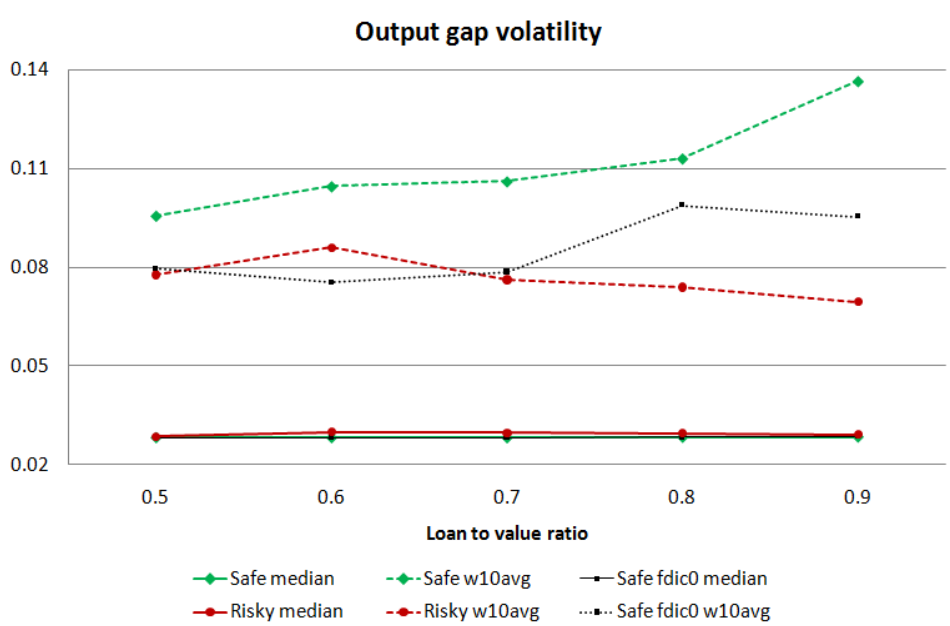

Figure 11: Effect of loan-to-value ratio on output gap volatility

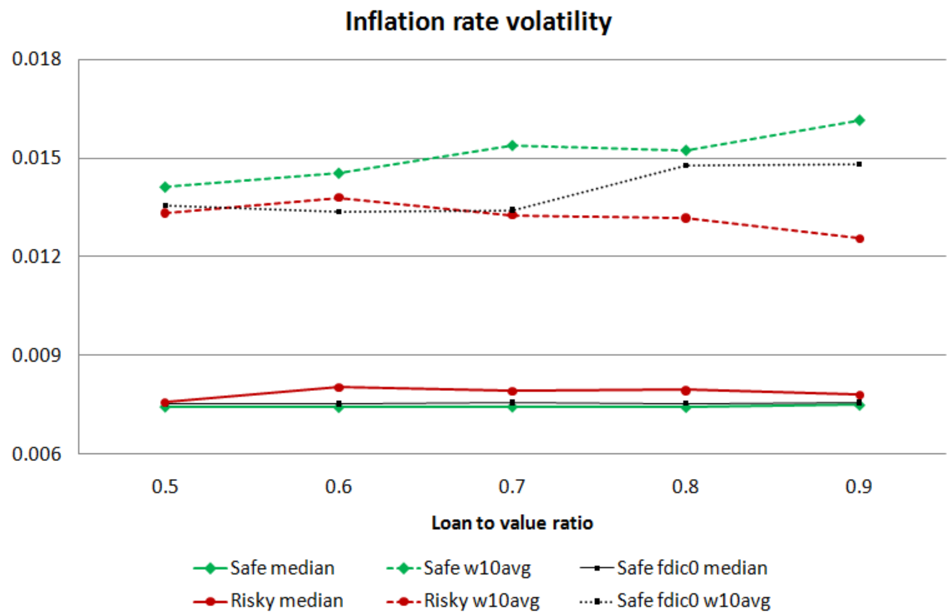

Figure 12: Effect of loan-to-value ratio on inflation volatility

Other than the higher loan-to-value ratio, the risky scenario differs from our baseline safe scenario in two other respects: loans are without recourse and new owners of failed banks are chosen at random rather than among those most likely to survive. To get a better idea of what explains the difference in hard-time performance between the two scenarios we ran the same experiment first under the risky scenario (indicated in Figures 9 through 12 by the lines marked with circles) and then under a variation of the safe scenario in which new owners of failed banks are chosen at random (indicated by the lines marked with squares). Looking at the high end of the loan-to-value range we see that holding loan-to-value constant about half the remaining difference between the safe and risky scenarios is attributable to the difference in recourse and the other half attributable to the care with which the government chooses new owners of failed banks. In both dimensions, what might seem like the safe option makes bad times even worse.

As indicated above, the loan-to-value ratio itself seems to work in the direction one might expect even in bad times, at least in our baseline safe scenario. But interestingly enough a higher loan-to-value ratio seems to make bad times less bad in the risky scenario.

6.2.1 Interpreting bad-time results

At this point we can offer only a tentative explanation for these seemingly paradoxical results, which runs in terms of the sensitivity of lending conditions to the state of the macroeconomy versus the average state of lending conditions. When times are bad banks tend to get into trouble and suddenly cut off their lending just when it is needed the most. Comparison of Tables 3 and 4 above indicates that although risky banks are in trouble more frequently in both the median situation and in the average worst-decile situation, the increase in the incidence of bank trouble going from the median to the worst-decile is much greater for safe banks (from 3.1 percent to 27 percent) than for risky banks (from 56 percent to 68 percent). Thus it appears that what makes bad times even worse is not so much the absence of lending as the cutting off of loans that had been granted before the banks got into trouble.

This line of reasoning helps to explain why careless choice of successors to failed banks seems to improve worst-case performance in the safe scenario. Weak successors are likely to be chronically in trouble, not just in the worst of times. This is indicated by Figure 13 below which shows that the median fraction of safe banks in trouble is almost 20 percent when successors are chosen randomly.

This line of reasoning is consistent with another result, namely that in the baseline scenario an increase in the loan-to-value ratio matters for bad times not because it increases the likelihood that a bank will become troubled (Figure 13 indicates that this likelihood is little affected) but because it makes things more difficult for the shops whose lending has been cut off, who ex post would often be better off if they had not borrowed as much. Having borrowed more they are more likely to exit when borrowing is cut off, and also more likely to lay workers off. Figures 14 and 15 below show that in the baseline scenario both job losses due to layoffs and job losses due to exits increase with the loan-to-value ratio.

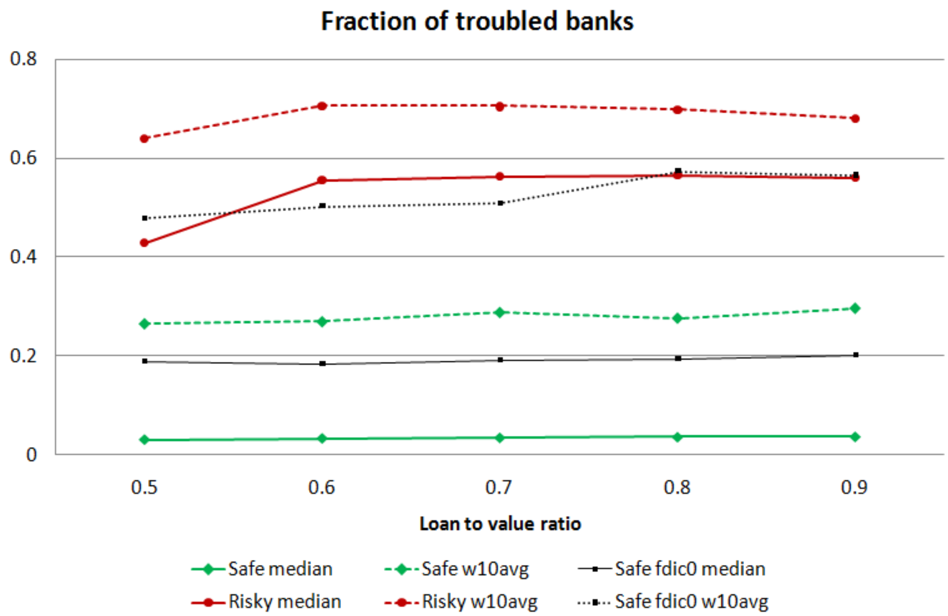

Figure 13: Effect of loan-to-value ratio on fraction of troubled banks

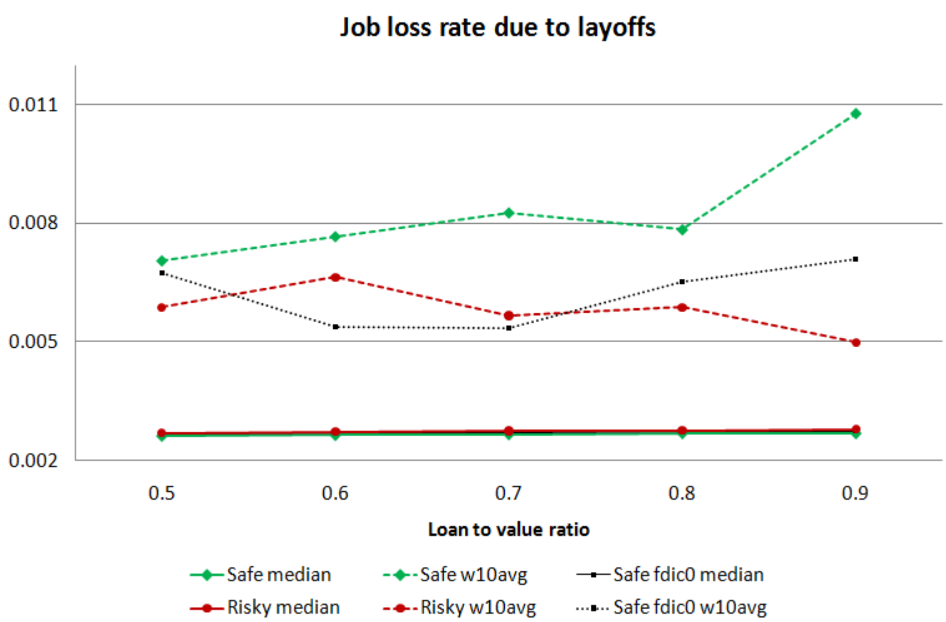

Figure 14: Effect of loan-to-value ratio on job loss due to layoffs

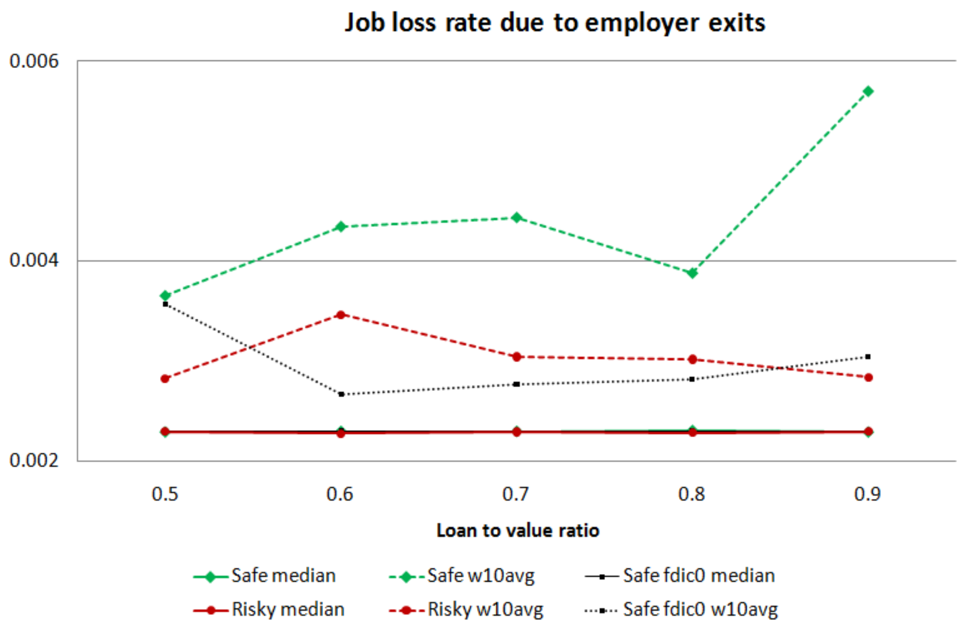

Figure 15: Effect of loan-to-value ratio on job loss due to exits

6.3 Policy experiments

The above results suggest that safe banks typically result in better macroeconomic performance than more risky banks, as usually defined by prudential regulations. What sorts of policies might work so as to improve their performance in bad times? The results of the previous subsection indicate that restrictions on loan-to-value ratios can be useful in limiting the damage that can be caused by having normally safe banks fall into trouble. In this section we experiment with other possible changes. The only one we have found that seems to work is restrictions on bank dividends.

Specifically, suppose we force banks not to pay dividends in excess of a certain fraction of their equity. In the following experiment we set that fraction equal to the rate of time preference, which is .04 per annum. Table 5 below shows that in terms of the output gap, unemployment and the volatility of the gap and of inflation this restriction led to a marked improvement over the baseline scenario, in terms of both median and worst-decile averages, although performance in the worst deciles was still not as good as in the risky bank scenario. The table also shows that this restriction reduced the increase in the likelihood of banks being in trouble between median and worst-case averages to zero, in line with the tentative explanation provided in the previous subsection.[24]

[24] The fact that W10 performance is still worse than with no banks is not however accounted for by our tentative explanation, since in this restricted case there was no increase in the fraction of troubled

TABLE 5

Results with dividend restriction

	restricted		unrestricted	
	median	w10	median	w10
Output gap	7.6	20	7.6	22
Unemployment rate	6.0	13	6.1	15
Volatility of output gap	2.8	8.3	2.8	9.6
Volatility of inflation	.73	1.3	.74	1.4
Fraction of troubled banks	0	0	3.1	27

We also experimented with various other policy interventions. First, Figure 16 below shows that varying the capital adequacy ratio κ between 0.02 and 0.16. The figure shows no systematic pattern other than that the lowest ratio seems to be the worst in terms of worst-decile average output gap.

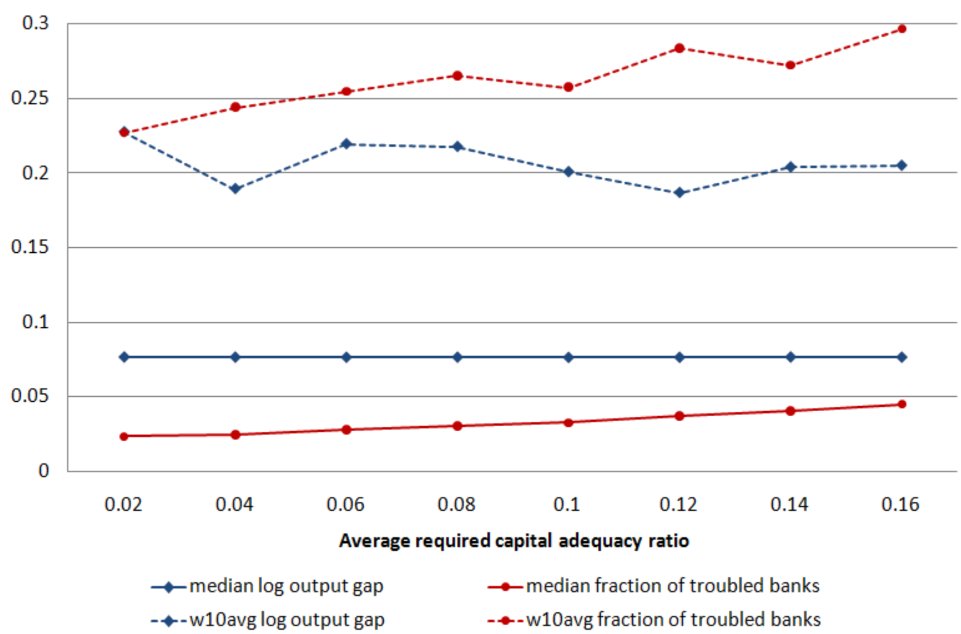

Figure 16: Effects of the capital adequacy ratio

Several commentators have called for making capital adequacy requirements procyclical. Figure 17 below shows the results of simulating such a policy. In these simulations, we supposed that the capital adequacy rate is tied to the difference between the output gap target and the estimate of the actual output gap. The adjustment is similar to that

banks!

of the target interest rate and also employs a squasher:

$$\kappa = \bar{\kappa} + \eta_\kappa (q - q^*) \cdot f(q, \kappa), \quad f(q, \kappa) \equiv \frac{\bar{\kappa}}{\sqrt{\eta_\kappa^2 (q - q^*)^2 + \bar{\kappa}^2}},$$

where η_κ is a fixed capital adequacy ratio adjustment parameter and $\bar{\kappa}$ is the average required capital adequacy ratio which we kept equal to 0.08. The rate was adjusted according to this formula once a month, at the time of the government's interest setting decision. As Figure 17 shows, the policy has almost no effect on the median output gap or median fraction of troubled banks, and its effects on the worst-decile average output gap are too variable to see any systematic pattern.

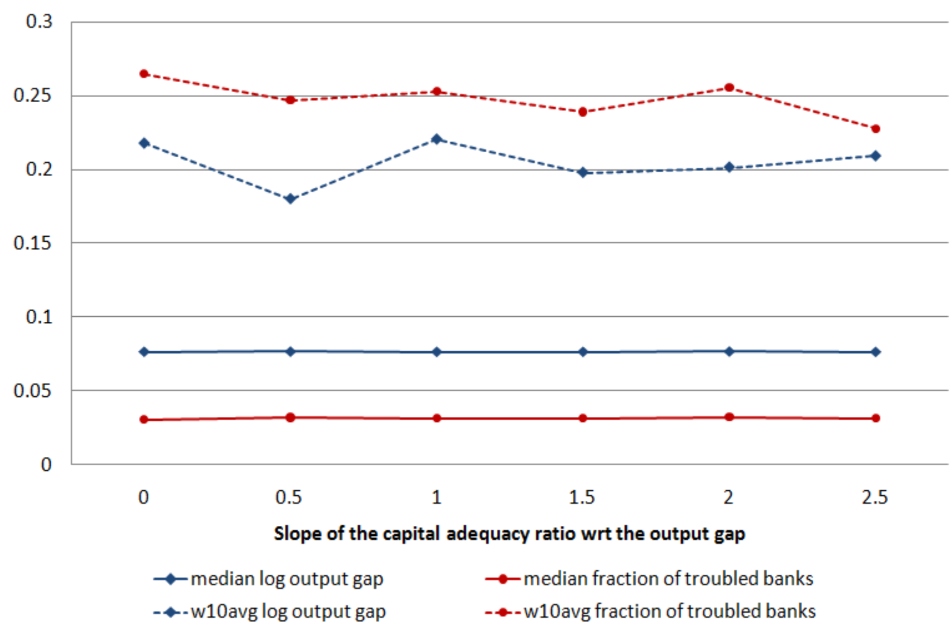

Figure 17: Effects of the slope of capital adequacy rate with respect to log GDP

Another policy measure relating to banks that could possibly affect macro performance is liberalization or deliberalization of geographic restrictions that would serve to reduce or increase the number of banks. Fewer banks means more diversified banks and hence more stable banks. But fewer banks also means that each bank is larger and hence the effect of any given bank getting into trouble will probably be larger. In a nonlinear world like this, large shocks that could take the economy outside the corridor of stability are a danger. Figure 18 below shows the results of varying the assumed number of banks from 1 to the maximum value equal to the number of goods, 50. We tried all values that were divisors of 50 between these limits so that each banking sector is of the same size.

As Figure 18 shows, worst-decile performance generally gets worse when the number of banks increases above 2. Roughly speaking, it seems that in terms of worst-case outcomes the danger of less diversified banks is generally greater than the danger of having a large bank get into trouble.

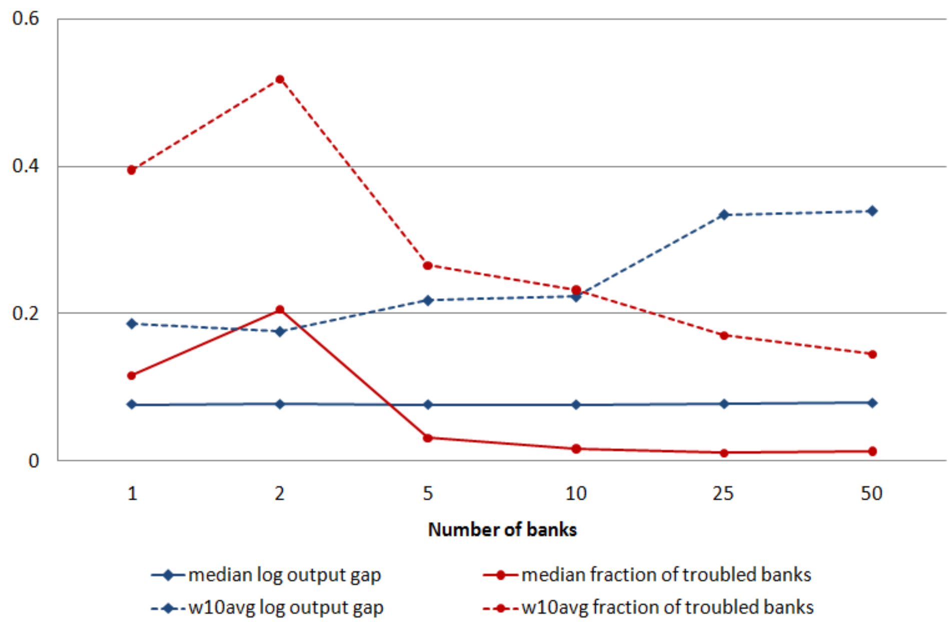

Figure 18: Effects of the number of banks

For our last policy experiment we asked what would happen if the government had a policy of bailing banks out when their equity fell below zero. Specifically, each week during financial market trading when the bank examination reveals negative equity we suppose that the government injects enough cash into the bank to bring its equity back up to zero, provided that the required injection is no greater than Ω/m, where Ω is the maximal fraction of nominal weekly GDP that the government is prepared to spend on bailouts and m is the number of banks. If equity falls below $-\Omega/m$ the bank is liquidated as in our baseline scenario. The parameter Ω is a measure of the government's willingness to bail out banks. In this experiment we allowed Ω to vary between 0 and 0.1. As Figure 19 below shows, bailing out banks is generally good for economic performance in worst-case scenarios. This is despite the fact that allowing a bank to fail results in a new owner being found who will at least initially be willing to resume lending, whereas bailing the bank out keeps it in the "zombie" state where it refuses to grant new loans to shops.

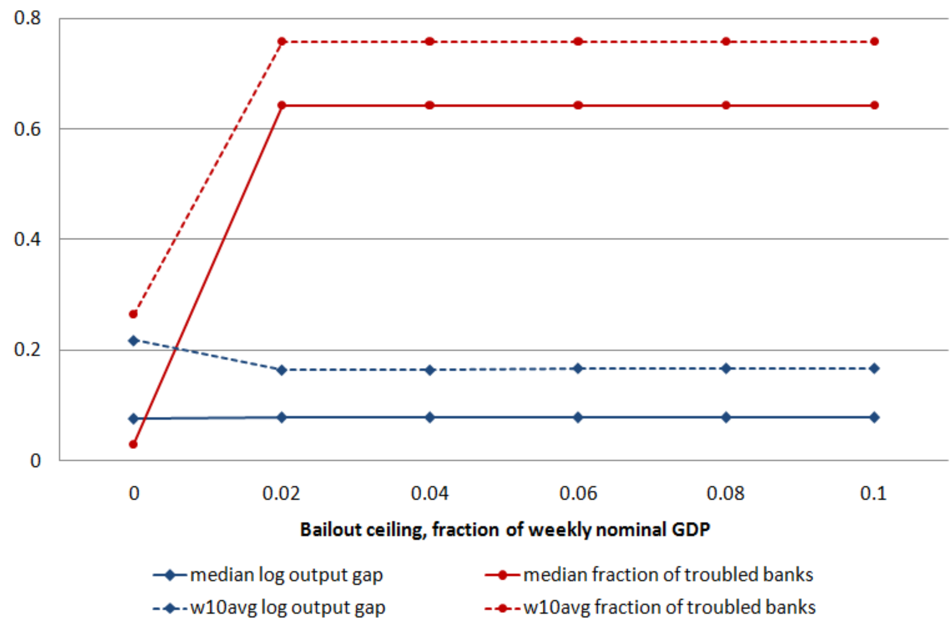

Figure 19: Effects of the government's willingness to bail out banks

7 Conclusion

We conclude by reiterating the tentative and exploratory nature of our present investigation. As indicated above we have taken nothing more than a small first step towards investigating the role of banks in the mechanisms that normally make a free-market economic system self-regulating. This preliminary investigation suggests that banks normally improve the economy's performance and that prudential bank regulation normally is good for performance, but that these same factors which improve performance in normal times also expose the economy to a greater risk of collapse, and thus have the effect of making bad times even worse. Exactly why this is the case, and whether this general conclusion will be robust to realistic extensions of the model is the subject of our ongoing investigations.

References

Aghion, Philippe, and Peter Howitt. "A Model of Growth Through Creative Destruction." *Econometrica* 60 (March 1992): 323–51.

Aghion, Philippe, and Peter Howitt. *The Economics of Growth.* Cambridge, MA: MIT Press, 2009.

Amano, Robert, Kevin Moran, Stephen Murchison, and Andrew Rennison. "Trend Inflation, Wage and Price Rigidities, and Productivity Growth." *Journal of Monetary Economics* 56 (April 2009): 353–64.

Ashraf, Quamrul, and Peter Howitt. "How Inflation Affects Macroeconomic Performance: An Agent-Based Computational Investigation." Unpublished, Brown University, 2008.

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

Berger, Allen N., and Gregory F. Udell. "Relationship Lending and Lines of Credit in Small Firm Finance." *Journal of Business* 68 (July 1995): 351–81.

Bils, Mark, and Peter J. Klenow. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* 112 (October 2004): 947–85.

Blanchard, Olivier, and Jordi Gali. "Labor Markets and Monetary Policy: A New-Keynesian Model with Unemployment." NBER Working Paper 13897, March 2008.

Bohn, Henning. "The Behavior of U.S. Public Debt and Deficits." *Quarterly Journal of Economics* 113 (August 1998): 949–63.

Burstein, Ariel, and Christian Hellwig. "Welfare Costs of Inflation in a Menu Cost Model." *American Economic Review: Papers and Proceedings* 98 (May 2008): 438–43.

Cole, Rebel A. "The Importance of Relationships to the Availability of Credit." *Journal of Banking and Finance* 22 (August 1998): 959–77.

Dunne, Timothy, Mark J. Roberts, and Larry Samuelson. "Patterns of Firm Entry and Exit in U.S. Manufacturing Industries." *RAND Journal of Economics* 19 (Winter 1988): 495–515.

Durlauf, Steven N., and Louis J. Maccini. "Measuring Noise in Inventory Models." *Journal of Monetary Economics* 36 (December 1995): 65–89.

Golosov, Mikhail, and Robert E. Lucas Jr. "Menu Costs and Phillips Curves." *Journal of Political Economy* 115 (April 2007): 171–99.

Hall, Robert E. "Lost Jobs." *Brookings Papers on Economic Activity* 26 (1995): 221–73.

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

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

Howitt, Peter. "Macroeconomics with Intelligent Autonomous Agents." Unpublished, Brown University, 2007.

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

Levine, Ross. "Finance and Growth: Theory and Evidence." In *Handbook of Economic Growth*, edited by Philippe Aghion and Steven N. Durlauf, 865–934. Vol. 1. Amsterdam: Elsevier, 2005.

Midrigan, Virgiliu. "Menu Costs, Multi-Product Firms, and Aggregate Fluctuations." Unpublished, New York University, 2009.

Nakamura, Emi, and Jón Steinsson. "Monetary Non-Neutrality in a Multi-Sector Menu Cost Model." NBER Working Paper 14001, May 2008.

Patinkin, Don. *Money, Interest, and Prices.* 2nd ed. New York: Harper and Row, 1965.

Roberts, John M. "New Keynesian Economics and the Phillips Curve." *Journal of Money, Credit and Banking* 27 (November 1995): 975–84.

Smets, Frank, and Rafael Wouters. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97 (June 2007): 586–606.

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.

Tesfatsion, Leigh, and Kenneth L. Judd, eds. *Handbook of Computational Economics.* Vol. 2. Amsterdam: North-Holland, 2006.

Wallis, John J., and Douglass C. North. "Measuring the Transaction Sector in the American Economy, 1870–1970." In *Long Term Factors in American Economic Growth*, edited by S.L. Engerman and R.E. Gallman, 95–162. Chicago: University of Chicago Press, 1986.

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

Yun, Tack. "Optimal Monetary Policy with Relative Price Distortions." *American Economic Review* 95 (March 2005): 89–109.