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

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

April 23, 2010

[1]Rough draft to be presented at Columbia University, April 27, 2010. The authors are grateful for suggestions from Blake LeBaron, Ned Phelps, Bob Tetlow and various seminar participants at the Federal Reserve Bank of Cleveland, the Marshak Seminar at UCLA, the Center on Capitalism and Society at Columbia and Brandeis University, none of whom bears any responsibility for errors or should be taken to endorse the analysis.

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 sheds new light on the conflict between micro prudential bank regulation and macroeconomic stability; under some circumstances the economy performs better when bank regulation pays less attention to micro prudence (ie when capital adequacy ratios are lower and allowable loan-to-value ratios are higher). Related to this, the analysis draws an important difference between "normal" performance of the economy and "worst-case" scenarios; the micro prudence conflicts with macro stability only in the worst-case scenarios.

1 Introduction

How do banks affect the macroeconomy? If banks get in trouble how does that matter for various performance measures? In this paper we explore one possible channel through which banks might either help or hinder the macroeconomy, namely their role in what Jevons called the "mechanism of exchange."

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 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 helping the economy to track a coordinated state of general equilibrium. In a sense our work 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. That line of research ran into the problem that trading out of equilibrium in one market generates rationing constraints that affects traders in other markets, in complicated ways that are hard to analyze. Because of these complications, analysis of general disequilibrium became highly intractable, as did the optimization problems facing transactors. To deal with these complexities we have chosen to model the mechanism of exchange using the methodology of 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.[1] 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 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.

The model we present is no more than a first pass at introducing banks into a short-

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

run macro model. It is certainly too stylized to be take seriously for policy purposes. But it does produce two results of general methodological interest. First, the model seems to provide a framework for guiding policy during "rare disasters." 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 small 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. We also find, that banks have their biggest impact on performance in those rare cases when the system is far from equilibrium.

Our second result is that the way banks and bank regulation affect median performance of the macroeconomy is often in conflict with conventional notions of micro-prudential bank regulation. Specifically, we find that performance in the worst decile of runs is significantly improved if banks are allowed to use higher loan-to-value ratios and are subject to lower capital adequacy ratios.

The next section attempts to place the paper's contribution in the literature on banks and the macroeconomy. Section 3 discusses the basic elements of our model. Section 4 discusses the behavioral rules that we are imputing to the various actors in the model. Section 5 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 6 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 7 describes our results. Section 8 concludes.

2 Previous literature

There is a broad literature that studies the effects of financial intermediation 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.[2] But none of these effects seem likely to trigger a collapse of the sort that policy makers have been trying to avert. Financial frictions a la Bernanke-Gertler or Kiyotaki-Moore are

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

currently being introduced into New Keynesian DSGE models, but 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. General equilibrium models in which intermediation plays a role in amplifying shocks have been provided, for example, by Williamson (1986) and Holmstrom and Tirole (1996), but these are not models of economic depressions.

[TO BE COMPLETED]

3 The model

3.1 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 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 governments. 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 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.

[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] Search and matching is now being introduced into labor markets in New Keynesian models by such authors as Gertler, Sala and Trigari (2008) and Blanchard and Gali (2008).

3.2 The conceptual framework

3.2.1 Transactors, goods and labor

There is a fixed number N of transactors, a fixed number n of different durable goods and the same number n of different types of labor. Labor of type i can be used only to produce good i. Time is discrete, indexed by "weeks" $t = 1, ..., T$. There are 48 weeks per "year." In addition to the n goods there are four nominal assets: fiat money, bank deposits, bank loans and bonds. Each of the last three is a promise to pay one unit of money ("dollar") next week.

Each transactor has a fixed type (i, j), where $i \neq j$ and $i \neq j + 1 \pmod n$, meaning that each week he is endowed with one unit of labor of type i (his "production good") and can eat only goods j and $j + 1 \pmod n$ (his two "consumption goods"). We assume that there is exactly one transactor of each type. Thus the population of the economy is $N = n(n - 2)$.

3.2.2 Shops, trading, production and storage

Because no transactor can eat his own production good, he must trade to eat. Trading can take place only through facilities called "shops." Each shop is a combined production/trading operation. There are n different types of shop. A shop of type i is capable of buying type i labor with money, selling good i for money, and converting type i labor into good i. The number of shops of each type will evolve endogenously.

To trade with a shop a transactor must form a trading relationship with it. Each transactor may have a trading relationships with at most one shop (his "employer") that deals in his production good, and at most one shop ("store") that deals in each of his consumption goods. Each transactor's trading relationships will evolve endogenously.

Each shop of type i has a single owner, whose production good is i. Operating the shop entails a fixed overhead cost of F units of type i labor per week and a variable cost of one unit of type i labor per unit of good i produced. When the shop is first opened the owner also incurs a setup cost; that is, he must invest S units of either of his consumption goods into the shop's fixed capital.

All trade with a shop takes place at prices that are posted in advance by the shop. Specifically, each shop posts a retail price p and a wholesale price w, which it adjusts periodically.

There is no depreciation or other physical storage cost. Goods produced but not sold in a week are kept in inventory. Fixed capital cannot be used for any other purpose until the shop exits. Former shopowners that still hold fixed capital cannot consume it but can continue to hold it in hopes of selling it to another shop in special "firesale" markets, to be described below. Likewise they can continue to hold the former shop's inventory until sold in a firesale market. The fixed capital and inventory still held by former shopowners are referred to as "legacy capital."

3.2.3 Banks

There is also a fixed number m of bank "sectors," where m is a divisor of the number of goods n. Agents are assigned to sectors depending on their production good, with the same number of agents in each sector. There is one bank in each sector, which is is owned by a single transactor in the sector.

A transactor who does not own a bank can deal only with the bank in his sector. He can hold financial wealth in three forms: money, bank deposits or shop equity. A transactor who is a shopowner can take out bank loans, which are made with full recourse but are also collateralized by inventory and fixed capital. Banks cannot lend to non-shopowners. A bankowner holds financial wealth in the form of money and bank equity.

A shopowner that exits must repay his bank loan. If the loan exceeds his money and deposit holdings the bank may seize the shop's inventory and fixed capital, crediting the debtor with a "haircut price" P_h for each unit seized. A shopowner who at any time is unable to pay his bank loan is declared bankrupt and his bank seizes all his assets, including his shop's inventory and fixed capital. The shop of a bankrupt owner must exit in the week of bankruptcy.

In addition to loans and seized collateral, banks can hold money and government bonds. They also have access to a lender of last resort facility from the government.

3.2.4 The government

There is also a government, which raises taxes, sets interest rates, lends to banks, regulates banks, and insures deposits. It does not purchase goods or labor but it does issue money and bonds, and services the interest on bonds through a sales tax on every retail transaction. It adjusts the ad valorem tax rate τ once per year.

The government pegs the interest rate i on its bonds by buying or selling whatever quantity the banks wish to hold at that rate. It adjusts this rate every 4 weeks. The rate it charges on advances to banks is $i + s_d$, where s_d is a fixed premium.

As bank regulator the government requires each bank to maintain equity at least equal to a fixed fraction κ of the value of its bank loans and seized collateral (i.e the value of its "risky assets"), where κ its the "capital adequacy ratio." A bank that is does satisfying this constraint is declared to be "troubled," and is forbidden to initiate any new loans or to pay its owner a dividend.

A bank with negative equity is forced into failure. The government seizes all the owner's wealth and injects it into the bank, injects enough extra money to make the bank no longer troubled and finds a new owner among the bank's depositors. The new owner adds any legacy capital that he might be holding to the bank's holdings of foreclosed capital, and adds his own deposits to the bank's equity. The recapitalized bank immediately reopens under the new owner, with all previous loans and deposits (except for the new owner's) unchanged.

4 Behavior

Each week the economy proceeds through the following sequence of events:

1. **Entry**. A random subset of transactors has an opportunity to set up a shop, and may realize that opportunity.
2. **Search**. Transactors search for new trading relationships with shops.
3. **Finance**. The government audits each bank, sanctioning those in trouble and reorganizing those that fail; transactors form spending and portfolio plans for the current week, and financial transactions take place through banks.
4. **Labor and goods trading**. First the firesale markets meet. Then everyone with an employer shows up to deliver the labor and receive the wage and everyone with stores visits them to execute their trading plans.
5. **Interest rate setting**. Every 4th week the government resets its interest rate and announces its forecasts of real interest rates.
6. **Match breakups**. A random subset of people dissolve all their trading relationships.
7. **Fiscal policy**. Every 48th week the government resets the tax rate.
8. **Shop closing**. Some shop owners may decide to exit, in which case all trading relationships with the shop are dissolved.
9. **Wage and price setting**. Shops periodically reset their posted wages and prices.

4.1 Entry

In the first stage, each transactor who is not already a shop owner or a bank owner is considered as a potential innovator. Each one becomes a potential entrant with probability θ/N. To enter the market the innovator undergoes a series of "tests" and does market research. First, each entrant needs to pay the real setup cost of establishing a new shop, S, which is incurred in units of either consumption good of the innovator. This cost (fixed capital) may be covered using various sources (and at different prices), namely: 1) own legacy capital, 2) firesale markets, 3) stores from which the innovator buys consumption goods. First, the innovator checks the availability of capital from all sources and evaluates its nominal cost, S_N. If there is enough capital to cover the setup cost, the innovator moves on to the financial viability test. This test checks whether there is enough financial wealth to cover the setup cost, S_N, and the fixed cost of operating the shop during the first month, $4(F-1)w$, where F is a fixed cost parameter, w is the shop's wage, and 1 is subtracted since the shopkeeper does not need to pay himself for own unit endowment. 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,[5] 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 expected inflation, because wages are kept fixed during the contract period.

The financial wealth consists of money holdings, deposit holdings, and the credit limit provided by the bank. To evaluate the latter each potential entrant first applies for a credit line in his bank. The credit line is granted with probability P_{CL} which depends on the financial condition of the bank (see section 4.3 for details). If the loan application is approved, the credit limit is determined as the haircut price, P_h, times the available collateral which for the entrant consists of legacy inventories (if any), LI, and the fixed capital he will need to purchase to enter the market: $CL = P_h \cdot (LI + S)$. If the financial viability test is passed, a business plan is initiated.

The innovator is assigned a random realization x of target sales (animal spirits) from a uniform distribution over $[1, n]$, and a random markup μ is drawn from a uniform

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

distribution over $[0, 2\bar{\mu}]$, where $\bar{\mu}$ is a parameter measuring the average percentage markup over variable costs. Given his target sales and markup, the innovator estimates his weekly target (after-tax) economic profit as a shop owner as[6]

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

where i_w is the nominal interest factor. Given this target profit, a financially viable innovator takes the profitability test. This test checks if the target profit (converted to last week's dollars) 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 total legacy assets LA and fixed capital:

$$\Pi > Y^p + (LA \cdot P_f + S_N)/V,$$

where P_f is the firesale 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.

An innovator that passes all of the above tests moves on to market research. To simulate market research, the program identifies the innovator's production good and then looks for his potential employees and customers. First, a "comrade" (someone else with the same production good as the innovator) is chosen at random and is considered as potential employee. Next, a potential customer (someone whose primary consumption good coincides with the innovator's production 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 production 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 necessary amount of fixed capital is actually purchased. The legacy inventories of the entrant become part of the shop's inventories, and the target input is set at the level

[6] 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 Π of the prospective shop owner.

$x + F + \lambda_I(x - I)$, 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.

4.2 Search and Matching

Next, each transactor is given an opportunity to search for possible trading relationships. This comprises both job search (for a shop that buys the transactor's production good) as well as store search (for shops that sell either of his two consumption 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 consumption 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 the transactor selects a shop at random. If the shop trades either of his consumption 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.

4.3 Financial Market Trading

At this stage all the financial transactions (beyond entry assistance) 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 and fixed capital seized by the bank from defaulting shops, valued at the current firesale 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 and the entry stage. 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. If the new bank owner has some legacy assets, they are put on the bank's balance sheet (seized collateral account) and participate subsequently on the firesale market along with other foreclosed assets that the bank has. 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 κ, the capital adequacy ratio[7]:

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

If this condition is violated, i.e., equity is less than the required capital, the corresponding bank becomes a "troubled" bank. Troubled banks are not allowed to provide loans, and their owners cannot get dividends. If the bank is not troubled, the probability of loan approval, P_{CL} is determined as

$$P_{CL} = l \cdot \left(\frac{\text{Equity}}{\text{Required capital}} - 1 \right),$$

where l is the slope parameter of the bank's loan approval schedule. This means that a loan is granted with probability 1 if the required capital as a fraction of equity does not exceed $l/(1+l)$. Thus, as the financial position of a bank deteriorates, that is, the ratio of equity to required capital falls, the probability of loan approval decreases linearly in that ratio.

[7] 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).

Next, agents of all types do their budget planning. The total wealth of each agent is the sum of financial wealth, A, and the capitalized value of permanent income, Y^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 assets (if any). For bank owners the financial wealth is the sum of money holdings and bank's equity after subtracting required capital. 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:

$$\Delta Y^p = \lambda_p \left(Y - Y^p \right),$$

where Y is the actual last period's income, and λ_p is the weekly permanent income adjustment speed. Here, Y is equal to last period's profit for shop owners and effective wage rate for all other agents. Then, the agents update Y^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:

$$E = v \cdot (A + P_f \cdot LA + V \cdot Y^p),$$

where V is the same capitalization factor as in section 4.1, 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. 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 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 taking into account the constraints they face. Consider first the transactors that don't own a bank or a shop. If $E < A$, they set $M = E$ and put the rest, $A - E$, 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 $M = A$. The idea here is that the agents will need to have E 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 out of the financial market to ensure against being unable to fulfil expenditure plans.[8]

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 value of its equity after covering the capital requirements.

Finally, consider a shop owner. His desired money holding is the amount enough to cover not only his planned goods expenditure but also 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. Again, the line of credit is granted with probability P_{CL}, as described above. If it is granted, the credit limit is set equal to the current haircut price, P_h, times the amount of inventories and fixed capital the shop has, i.e., inventories and fixed capital 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. In this case the resulting financial constraint of the shop owner is $A+CL$, where $CL = P_h \cdot (I+S)$ is the available credit limit. If the bank is troubled, it cannot provide loans. The bank's weekly lending rate is determined as the weekly nominal interest rate on government bonds (equal to the weekly deposit rate) plus a fixed spread s.

Given his financial constraint, the shop owner updates target input and the target wage bill.[9] Based on the shop owner's financial situation, the following cases are possible:

1. If $A+CL < 0$, the shop goes bankrupt. In this case the actual planned expenditure is set to zero, the bank seizes the collateral, cash and deposit holdings of the shop owner. We allow for a real cost of bankruptcy, C_b, which means that the bank only gets a fraction $1 - C_b$ of collateral. The shop owner's loans are voided, and the bankrupt shop waits until the exit stage to leave the market.

2. If $A+CL > 0$ but is not enough to pay for 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 $M = A + CL$.

[8] 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.

[9] The target input is calculated in the same way as in section 4.1, but the target sales are set equal to the previous period's actual sales.

3. If $A + CL$ is large enough to pay for 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 $A + CL$ minus the target wage bill, i.e., his priority is to have enough cash to cover the target payroll. As before, $M = A + CL$.

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. In this 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.

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 government debt is charged[10]. If the bank does not have enough funds to clear all financial operations, it has to borrow the shortfall from the government. If, in contrast, there is surplus, it is invested in government bonds.

4.4 Labor and Goods Market Trading

This stage starts with the firesale market trade. All active shops that are not bankrupt can buy input good on the firesale 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 or rejects to open the credit line, the shop owner cannot place an order that, evaluated at the firesale 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 firesale 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 assets 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

[10] 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.

is fulfilled or the queue runs out of sellers. The payment for inventories bought at the firesale 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 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.[11] 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. In general, 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) when the employer is broke, i.e., does not have enough cash to pay the wage bill; 3) if a random breakup of the match happens and he has to quit the job (see section 4.6); 4) when the employer exits the market (see section 4.8).

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

[11] The fixed cost is covered by each shop from own input, inventories and, if needed, from employed inputs. If the shop is unable to cover the fixed cost of operation, it is not allowed to sell anything during that week.

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 government at the monetary policy stage.

The stage ends with updating of the average weekly wage rate, firesale 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 firesale 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 collateral.

4.5 Monetary Policy

Next is the stage in which the government 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 government 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) and the current price level (GDP deflator). The government also keeps track of the values of year-to-year inflation factors, price levels, and average real weekly GDP for the last 12 months.

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

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

where γ_π and γ_y are fixed coefficients, $1+\pi$ is the inflation factor over the past 12 months, π^* is the fixed inflation target, y is the current 3-months moving average for the weekly average log GDP, $\tilde{y}$ is the government's evolving estimate of log potential output, and $i^* \equiv r^* + \pi^*$, where r^* is the evolving target for the long-run real interest rate. We also assume that the government 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 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_r 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 government'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.

The government is modeling the year-to-year behavior of log GDP in the following way:

$$y_t = \alpha_y + \lambda_y y_{t-1} + \xi_{yt},$$

where $\tilde{y} \equiv \alpha_y/(1 - \lambda_y)$ is defined as potential log GDP, and ξ_{yt} is an i.i.d. disturbance term. That is, the government assumes that the level of log GDP will approach its potential level in yearly movements according to a simple autoregressive process. Given the initial estimates $(\hat{\alpha}_{y0}, \hat{\lambda}_{y0})$, the government re-estimates these parameters monthly[12] using a recursive OLS scheme on annual data:

$$\begin{pmatrix} \hat{\alpha}_{yt} \\ \hat{\lambda}_{yt} \end{pmatrix} = \begin{pmatrix} \hat{\alpha}_{yt-1} \\ \hat{\lambda}_{yt-1} \end{pmatrix} + \begin{pmatrix} \overline{y_{t-1}^2} - y_{t-1}\overline{y}_{t-1} \\ y_{t-1} - \overline{y}_{t-1} \end{pmatrix} \cdot \frac{y_t - \hat{\alpha}_{yt-1} - \hat{\lambda}_{yt-1} y_{t-1}}{t\overline{y_{t-1}^2} - t\left(\overline{y}_{t-1}\right)^2},$$

where $\overline{y}_{t-1}$ is the average lagged log GDP in the learning period sample, and $\overline{y_{t-1}^2}$ is the average squared lagged log GDP. The new estimate of potential log GDP, equal to $\hat{\alpha}_{yt}/(1 - \hat{\lambda}_{yt})$, is the one subsequently used in the Taylor rule.

Similarly, the CB is modeling the annual evolution of inflation:

$$z_t = \lambda_\pi z_{t-1} + \xi_{\pi t},$$

where $z_t \equiv \ln(1 + \pi_t) - \ln(1 + \pi^*)$, and $\xi_{\pi t}$ is an i.i.d. shock. Again, given the initial

[12] We assume that the central bank begins to adjust all the estimates and targets after T_{cb} years.

estimate $\hat{\lambda}_{\pi 0}$, the CB recursively re-estimates this parameter on a monthly basis after the learning period is over:

$$\hat{\lambda}_{\pi t} = \hat{\lambda}_{\pi t-1} + \frac{\overline{z_t z_{t-1}} - \hat{\lambda}_{\pi t-1} \overline{z_{t-1}^2}}{t \overline{z_{t-1}^2}}.$$

The government's projections of inflation and GDP, based on the above equations, are then used in computing the capitalization factor V employed by transactors when formulating their expenditure plans.

4.6 Match Breakups

Next there are random breakups of established trading relationships. 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.

4.7 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 estimated potential GDP. 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 between the actual and the target debt to GDP ratio b^*:

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

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

4.8 Shop Closing

Now each shop has an opportunity to exit. Exit can occur for any of the following reasons:

1. Each shop exits for exogenous reasons with exogenous probability δ.

2. Each unprofitable shop exits with exogenous probability ϕ

3. Each shop that lacks the financial resources to pay for next the coming week's fixed overhead cost $(A + CL < w(F-1)$ exits.

4. Each shop whose owner is bankrupt exits.

Once a shop exits for any of the above reasons, all trading relationships (with both employees and customers) are dissolved and the shop owner has to settle with the bank. 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 and fixed capital at the current firesale price. If the total wealth plus the value of inventories and fixed capital is enough to cover the whole loan, the bank seizes the needed amount of assets. Otherwise, it seizes everything it can from the shop owner.

The profitability condition in case 2 takes into account the opportunity cost of the owner's labor services, which could be earning a wage, and the interest-opportunity cost of maintaining the shop's fixed capital and inventories. Specifically, the shopowner will decide that the shop is unprofitable if either

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

where, as above, A is the financial wealth of the shopkeeper, I is the shop's inventory, S is the setup cost (fixed capital), 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). In both cases the shop will exit if the owner's tangible plus human wealth would increase as a result, under the assumption that he could get a job paying the wage W. In case (a), he would be able to repay his loan in full, although perhaps allowing some some inventories and fixed capital to be seized, so in the event of exit his tangible plus human wealth would go from $A + \Pi^e V$ to $A + P_f(I+S) + WV$. In case (b) he would be unable to repay his loan in full, so upon exit his tangible plus human wealth would go from $\Pi^e V + A$ to WV.

In case if there are some legacy assets left after settling with the bank, the exiting shop owner is added to the firesale market queue corresponding to the legacy assets he has. If the bank seizes part or all of the assets from the shop owner, the bank is also added to the firesale market queue if it is not already there with the same type of good (if the bank is there already, it keeps the previous position but the size of the order increases by the amount of seized collateral).

Upon exit, the former shopowner resets his permanent income to W.

4.9 Updating of Targets, Wages, and Prices

In the final stage of weekly activities, each shop first updates its target sales y^{trg}, setting them equal to the 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 is set equal to:

$$w = w_0 \cdot \left[\left(1 + \left(\beta \cdot \left(\overline{trInput}/\overline{Input} - 1\right)\right)\right) * (1 + \pi^*)\right]^{\Delta/48}$$

where w_0 is the preexisting wage, $\overline{trInput}$ is the average target input over the past Δ weeks and and $\overline{pInput}$ is the average potential input over the same period (ie the number of transactors having an employment relationship with the shop, even if they were laid off or refused to work because they were not paid).[13] The parameter β hence indexes the degree of wage and price flexibility in the economy. This annual wage adjustment figure anticipates inflation over the coming contract period at an annual rate equal to the government's target rate π^*.

Every week each shop has an opportunity to revise its retail price. Its "normal" price is

$$p^{normal} = (1 + \mu)\, w / (1 - \tau)$$

which would equate its after-tax price to its wage times its desired markup. It will choose this normal price unless its inventories are far from the desired level, namely its target

[13] In computing this expression we use the maximum of $\overline{pInput}$ and the shop's fixed cost F to avoid division by zero when potential employment falls to zero.

sales. Specifically it will set:

$$p = \begin{cases} p^{normal} \cdot \delta_p & \text{if } I > IS \cdot y^{trg} \\ p^{normal}/\delta_p & \text{if } I < IS \cdot y^{trg} \\ p^{normal} & \text{otherwise} \end{cases}$$

Thus the frequency of price change will be endogenous. A shop will change p almost certainly twice a year, when its wage is changed and when the tax rate τ changes, because in both cases its normal price will change. Beyond that it will only change when its inventory/sales ratio passes one of the critical thresholds IS and $1/IS$. When the ratio rises above the upper threshold the shop cuts its price by the factor δ_p. When the ratio falls below the lower ratio it raises its price by the factor $1/\delta_p$.

4.10 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.

5 The workings of the model

5.1 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:

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

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

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

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 magnitudes – money and bond holdings, prices and permanent incomes – will rise at the weekly target inflation rate π_w^*.

5.1.1 Entry, exit, and systemic performance

As we shall see below, the economy is able to achieve on average about 93% 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 shops that are able to satisfy customers in markets where there had previously been no viable shop, and by the exit of shops 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 shops 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 shops were hiring most of the potential workers and satisfying most of the potential customers. Likewise, although exit is important in cases where the shop 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.

Banks have two influences on this cumulative process of shop failures. One is the familiar "lending channel" of the financial accelerator. That is, when a shop fails this may cause its bank to get into trouble, or at least to get close to being in trouble, because some of its loans will go bad. This will cause it to lend less readily, either by reducing its probability of approving credit line applications or in the extreme case by having the government forbid it from making any new loans. This makes it more likely that other firms may fail for want for financial capacity.

The other influence that banks have is a more salutary one of ameliorating the cumulative process. That is, the cascade of shop failures will be dampened if another shop quickly replaces each failed shop. This is more likely to happen when banks are willing to lend because bank lending makes it easier to pass the financial viability test of entry in stage 1 each week.

As we shall see it seems that the latter role of banks, that of facilitating entry, appears to play a more important role than the more traditional former role.

6 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.

There are a total of 33 parameters, which we have categorized as shop parameters, transactor parameters, bank parameters, and government parameters. These are listed in Table 1 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.

6.1 First level of calibration

6.1.1 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.

TABLE 1

	Shop parameters	
Δ	Length of the wage contract period (in weeks)	48
$\bar{\mu}$	Average percentage markup over variable costs	0.131
ϕ	Failure rate of unprofitable shops (weekly)	0.01
F	Fixed cost	3.5
S	Setup cost	15
λ_I	Inventory adjustment speed (weekly)	0.16
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)	1.017
	Transactor parameters	
ε	Demand parameter	7.0
λ_p	Permanent income adjustment speed (weekly)	0.4
ρ	Rate of time preference (annual)	0.04
θ	Frequency of innovation (weekly)	100
δ	Quit rate (per week)	0.0007
σ	Job search probability	0.5
	Bank parameters	
l	Slope of loan approval schedule	9
s	Annual loan spread	0.0175
h	Loan-to-value ratio	0.5
C_b	Cost of bankruptcy	0.1
	Policy parameters	
	Fiscal Policy	
b^*	Target Debt-GDP ratio	0.33
λ_τ	Fiscal adjustment speed (annual)	0.054
	Monetary Policy	
π^*	Annual target inflation factor	0.03
γ_π	Inflation coefficient in Taylor rule	1.5
γ_y	Output gap coefficient in Taylor rule	0.5
r_0^*	Initial target real interest factor	0.032
η_r	Adjustment speed of evolving real rate target	0.0075
$\hat{\lambda}_{\pi 0}$	Inflation autocorrelation factor (initial estimate)	0.29
$\hat{\alpha}_{y0}$	Intercept in output equation (initial estimate)	2.595
$\hat{\lambda}_{y0}$	Output autocorrelation factor (initial estimate)	0.66
T_{cb}	Number of years before central bank's learning begins	10
	Bank regulation	
m	Number of banks	5
κ	Required capital adequacy ratio	0.08
s_d	Premium on Fed's discount rate	0.005

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 shop 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

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

If the economy was operating with a 6 percent average unemployment rate then the typical shop would have

$$x = 0.94 \cdot (n - 2) = 43.24$$

so by setting the fixed cost F equal to 3.5 we get a typical degree of returns to scale equal to 8.8 percent.

The inventory adjustment speed $\lambda_I = 0.16$ corresponds to the estimate by Durlauf and Maccini (1995) of a monthly adjustment speed equal to approximately 0.5 ($\approx 1 - (1 - 0.16)^4$).

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 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.

6.1.2 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[14]. 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.

[14] 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.

6.1.3 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.

We set the cost of bankruptcy equal to 0.1 which falls in the 2% to 20% range suggested by Bris et al. (2006).

6.1.4 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 $\hat{\lambda}_{y0}$ 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_y 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 government 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.

6.2 Second level of calibration

6.2.1 Government targets – Finessing Wicksell

The two government targets – the target real interest rate r^* and log potential output $\tilde{y}$ – are chosen adaptively by the government. We chose their initial values (respectively, .032 and 7.5, so that the output gap is .075) close to the steady state values in our calibration simulation, although since government 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 governments estimate r^* and $\tilde{y}$ 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 potential output, and hence that controlling the rate of interest to the neglect of the money supply risks aggravating volatility or even engendering a cumulative inflation.

6.2.2 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 $\tilde{y}$ – by internal consistency. In our baseline simulations the median markup is about 13.14 percent when shops choose from a distribution whose mean $\bar{\mu}$ is 13.1 percent.

6.3 Third level of calibration

This leaves 13 parameters still to be determined, namely ϕ, S, θ, δ, λ_p, σ, IS, δ_p, l, h, m, T_{cb} and η_r, were chosen by searching (manually) for parameter values that would generate values of thirteen different indicator variables, that approximated their counterparts in US data. More specifically, we ran 10,000 simulations of 60 years. Each simulation started near the no-shock equilibrium and continued for 20 years before we started calculating the average value of each variable across the remaining 40 years of that simulation. For each variable we then computed the median across all simulations of these cross-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	3.2
Unemployment rate	6.1	5.9
Unemployment duration	14	10
Job loss rate	.69	.62
Volatility of output gap	2.0 to 3.2	2.8
Volatility of inflation	1.3	.69
Autocorrelation of gap	20 to 76	43
Autocorrelation of inflation	16	19
Average markup	10 to 20	13
Exit rate	46	42
Price change frequency	4.0	4.2
Annual bank failure rate	0.51	1.00

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 1969 to 2009. The two numbers given for the volatility of the output gap are the standard deviation of linearly-detrended or HP-filtered, respectively, 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 AR(1) 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 shops 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.[15] 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

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

3. However, in our model shops 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. Specifically, our model significantly underpredicts the duration of unemployment, the volatility of inflation, the job loss rate and the exit rate. We do not mind that the model overpredicts bank failures, as bank mergers (which cannot occur in our model) often take the place of reported failures.

6.4 Approximating a steady state

As Figure 1 below shows, our 20 year adjustment period 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 downward trend in the average gap over the first half of the final 40 year period.

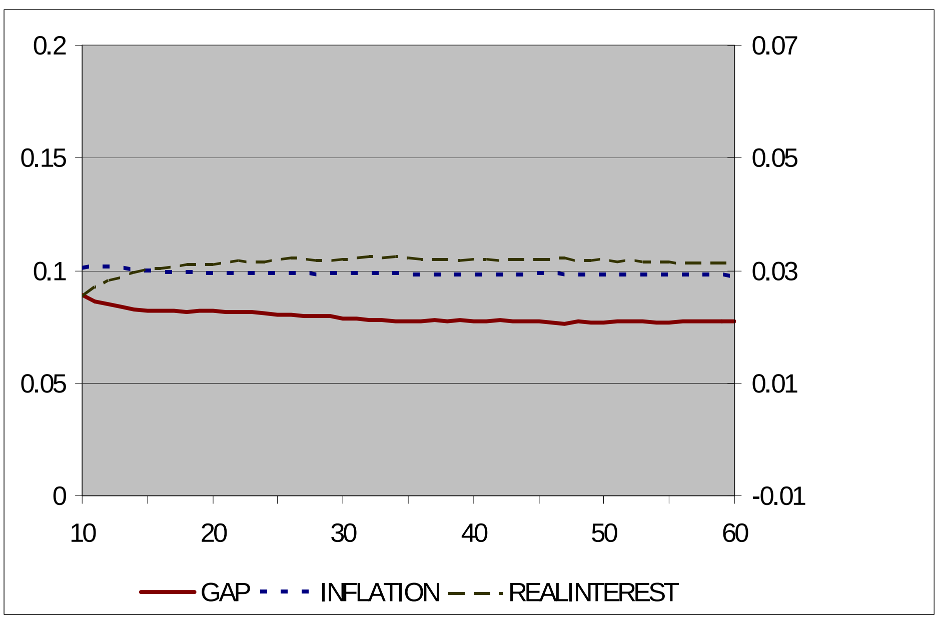

Figure 1: Average simulated values in baseline calibration

7 Results

As indicated in the introduction to the paper, our results indicate that banks do matter for macroeconomic performance, especially when times are extremely bad, and in bad times they are especially helpful when they are riskier from the perspective of micro-prudential regulation, that is when they have higher loan-to-value ratios and lower capital adequacy ratios.

7.1 Normal times and bad times

Figure 1 above describes the average time path of 10,000 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.[16] A randomly chosen simulation would depict similar characteristics. There are times when the output gap rises for a few years but these times are soon followed by recovery.

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

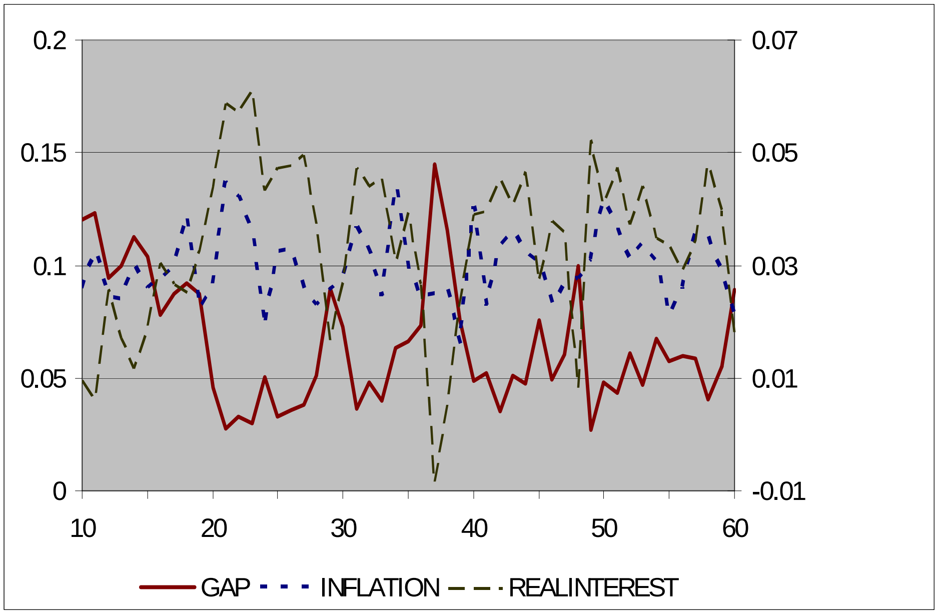

Figure 2: A normal run (rnseed=11)

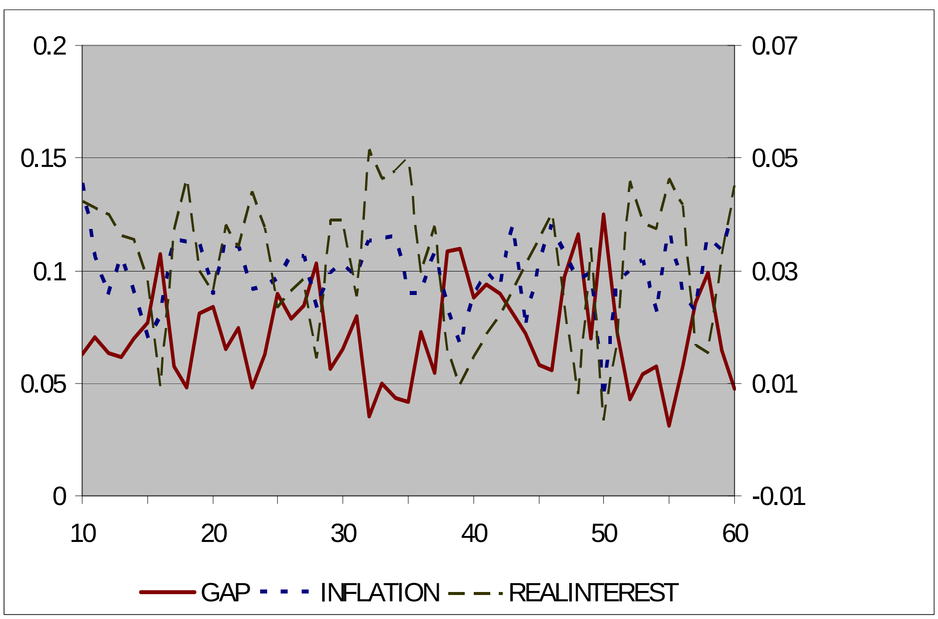

Figure 3: A normal run (rnseed=13)

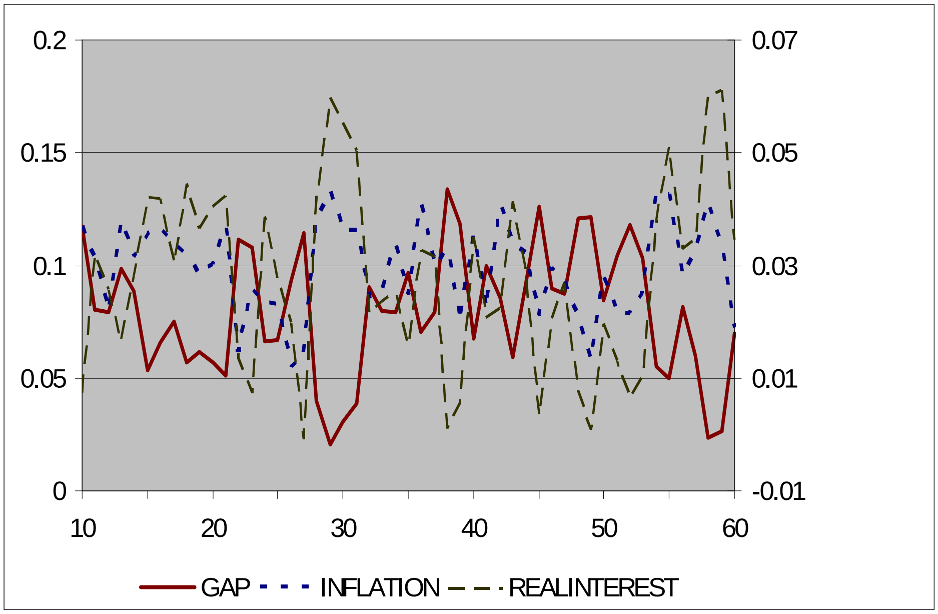

Figure 4: A normal run (rnseed=15)

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.

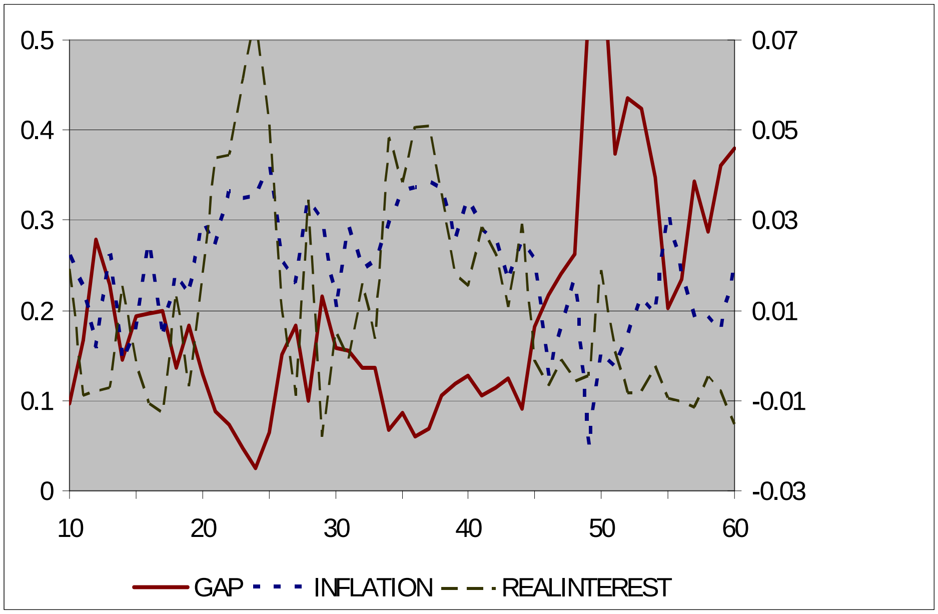

Figure 5: A collapse (rnseed=419)

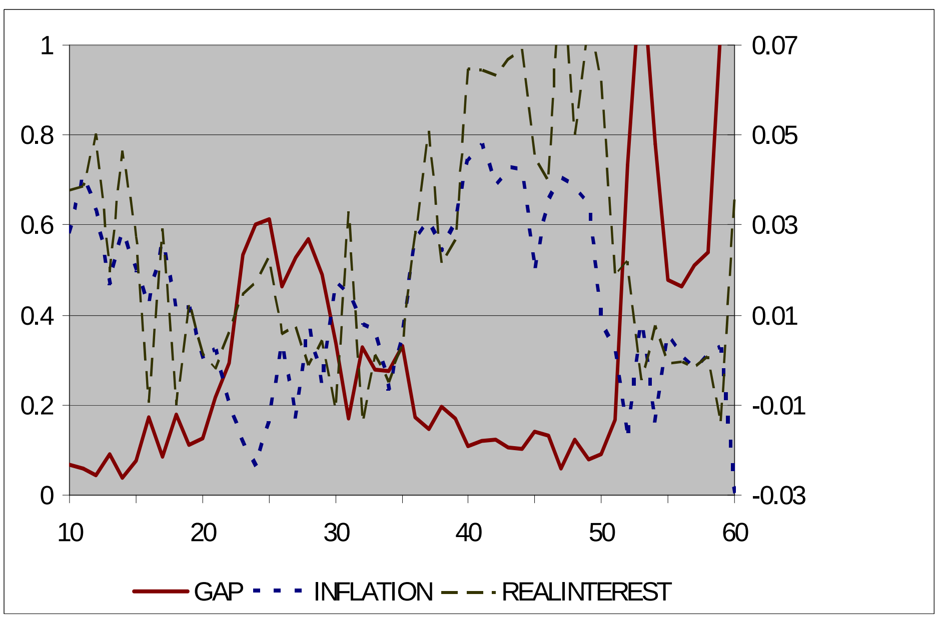

Figure 6: A collapse (rnseed=61)

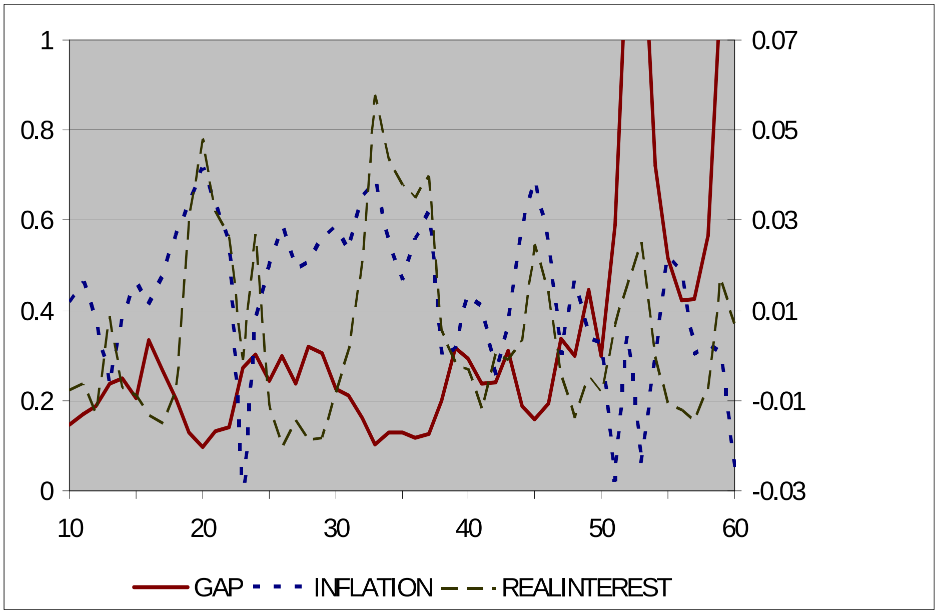

Figure 7: A collapse (rnseed=1061)

These collapses suggest the presence of something like Leijonhufvud's (1973) "corridor;" it seems as if the economy is capable of absorbing shocks up to some limit but not beyond.

7.1.1 The anatomy of bad times

Figure 8 below indicates the average behavior of various macro performance measures by decile of runs, where the runs are ordered in terms of their cross-year average output gap. As this figure indicates, there is a sharp deterioration in all these measures, except the job loss rate, in the tenth decile.

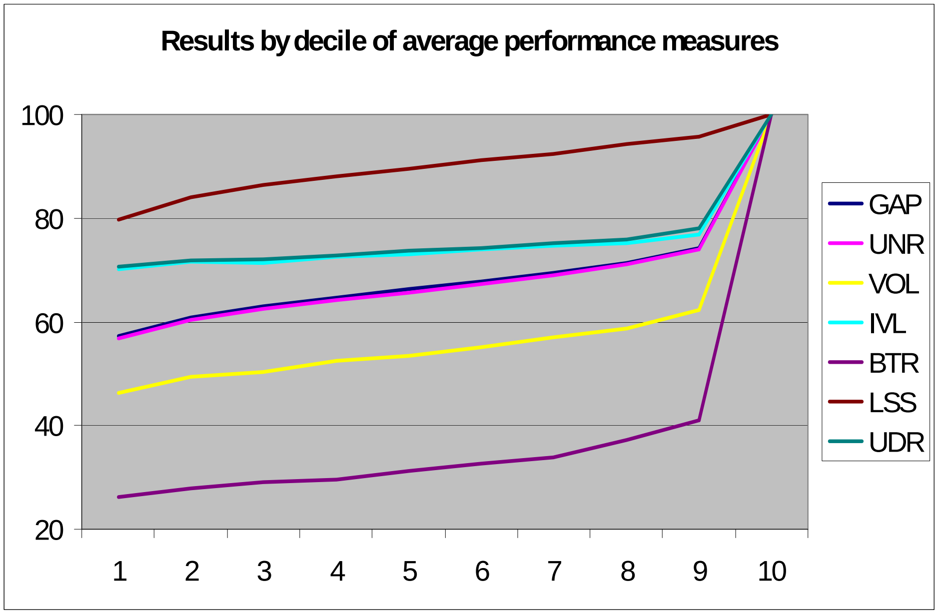

Figure 8: Macro performance by decile

To some extent this deterioration is just what one would expect from a more standard New Keynesian DSGE model that respected the zero lower bound on nominal interest rates.[17] Figure 9 below shows that the fraction of times in which the zero lower bound is hit rises sharply from xx to xx between the 9th and tenth decile of runs.

[17] See for example Eggertsson and Woodford (2003) or Levin et al. (2010).

Figure 9: The zero lower bound on nominal interest rates

However, there is more going on. Figure 10 reports the results of redoing the 10,000 runs of our basic calibration but allowing the nominal interest rate to be determined by the Taylor rule even when this rule would make it negative. All other behavioral rules in the model were left unchanged. Of course this experiment begs the question of why people would aim to have their cash-in-advance constraints binding when the nominal interest rate was below zero, but it does nevertheless show that there is more going on in the model than the zero lower bound. Figure 10 below provides the same decile-by-decile report of various performance measures as Figure 8 for this experiment with the zero lower bound suppressed. It shows that although the deterioration is less severe than before when going from the ninth to the tenth decile of runs there is still a kink in the relationship much like before.

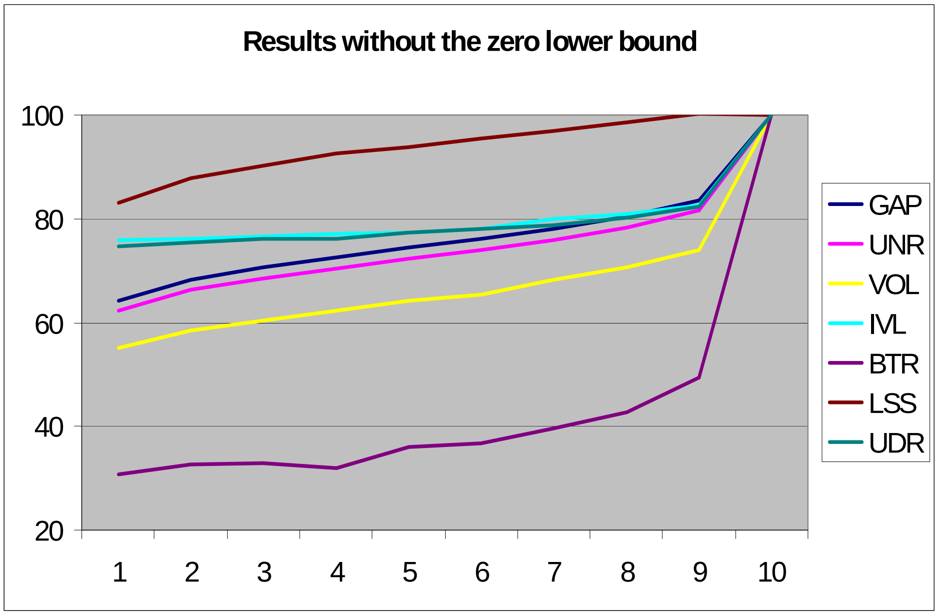

Figure 10: Performance measures decile by decile without the zero lower bound

Instead of a more conventional zero lower bound effect, we believe that what is happening here is an amplification of recessions that works through a bank-lending channel. Specifically, when the economy experiences a bad run banks fall into trouble. Indeed according to both Figures 8 and 10 the sharpest deterioration of all reported performance measures going from the ninth to the tenth decile is in the fraction of banks in trouble. In the baseline case this fraction rises monotonically from decile to decile but it rises by more than three times as much from the ninth to the tenth as it did from the first to the ninth.

Having banks in trouble means that they stop granting new loans. This discourages entry of new shops and makes exit of existing shops more likely. Figure 11 below shows that in the baseline case this is associated with a sharp reduction in the average number of shops and a sharp increase in the fraction of entrepreneurs who fail the financial viability test when contemplating entry.

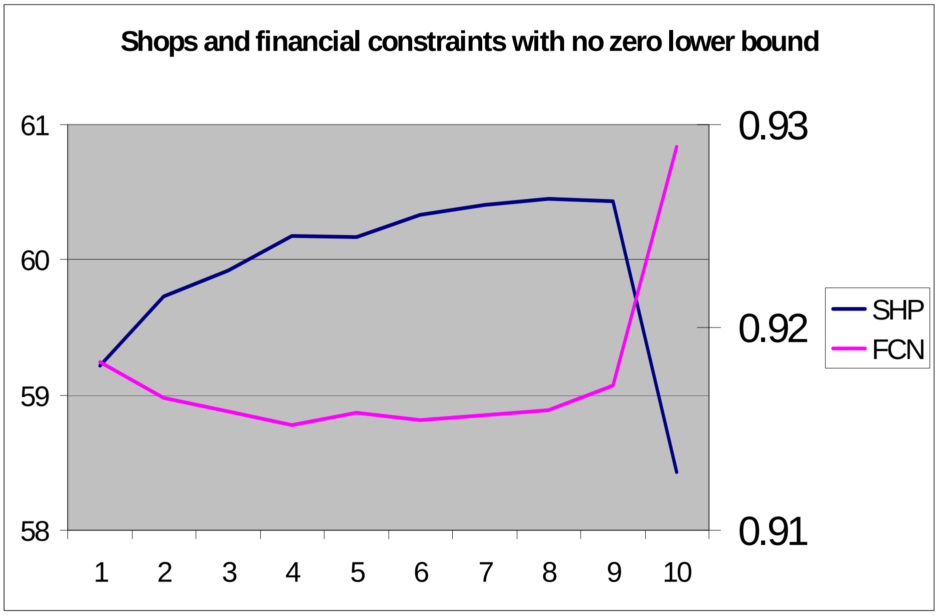

Figure 11: The lending channel at work

7.2 Banks, no banks and "risky" banks

In order to gain some more insight into how 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 marked deterioration in median performance when banks are shut down.

TABLE 3

Median results

	Banks	No banks	Risky banks
Inflation	2.9	2.5	3.0
Output gap	7.5	14	7.8
Unemployment rate	5.9	12	5.9
Unemployment duration	10	18	8.4
Job loss rate	.62	.69	.73
Volatility of output gap	2.8	6.3	2.5
Volatility of inflation	.69	1.1	.61
Annual bank failure rate	1.0	0	3.5
Fraction of banks in trouble	3.8	0	1.3

We can get another measure of the importance of banks by simulating the economy's response to various exogenous shocks. Figure 12 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 20 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 10,000 times. For a counterfactual we perform the same 10,000 simulations, with the same sequence of seeds for the random number generator, but without the shock that forces a bank into trouble. The upper ("safe") line in Figure 12 indicates the trimmed average difference in log GDP between the shocked and counterfactual simulations on a monthly basis for twenty years (24 months) following the shock. That is, in order to avoid the noisiness of runs where either the shocked or counterfactual simulation spun out of control we calculated the average each month between the 25th and 75th percentile of all these differences. The effect is not huge, never exceeding a one and a half percent of the average counterfactual value, but it is persistent.

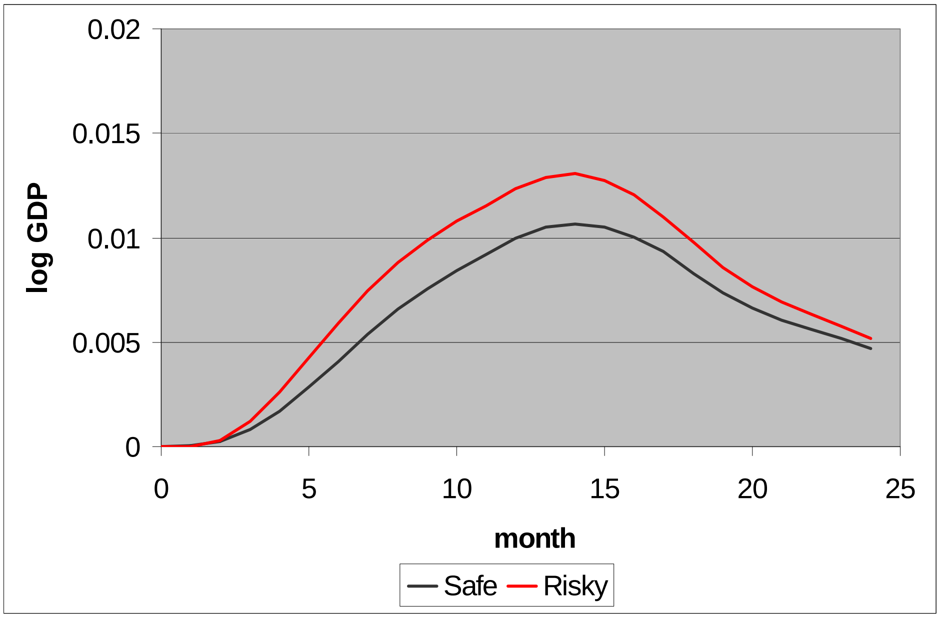

Figure 12: (Trimmed) average response to a troubled bank shock

Another experiment we ran was to force a shop to fail. We did this both with and without banks. Figure 13 below shows that in either case there was a humped-shape impulse response of log GDP, which is much stronger in the economy without banks (the upper curve) than in the baseline ("safe") scenario. As explained by (Howitt 2006) in a much simpler model based on the same foundations, the exit of a shop can cause a cumulative fall in output by depriving former suppliers of the shop of their source of income and thus endangering other shops from loss of demand. Each time a shop fails, aggregate output falls automatically by an amount equal to that shop's employment less its fixed overhead cost, unless it was already employing too few to cover its fixed cost. The process continues until the former employees of failed shops find new employers. Often this will require new shops to enter. As Figure 13 shows, banks play an important role in mitigating the multiplier process by facilitating this new entry.

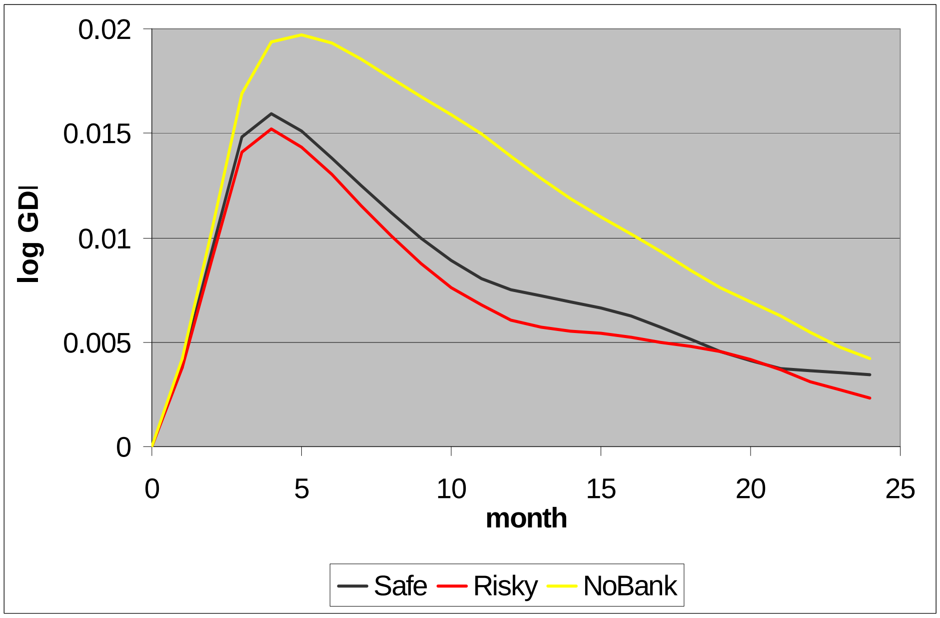

Figure 13: Response to a shop failure

This last experiment suggests that banks should make an even bigger difference in bad times than in normal times. For, as have seen, bad times are exacerbated by a loss of shops, which loss can be mitigated by bank lending that facilitates entry, whereas, as argued above, entry is less critical when the economy is already near a situation of full capacity utilization. Table 4 below reveals that banks do indeed make a bigger difference in bad times. It shows the average across the worst decile of runs of various performance measures, with and without banks. In almost all cases the difference in worst-decile averages is considerably larger than the difference in medians reported in Table 3. For example, banks reduce the median unemployment rate from 12 to 5.9 percent, but they reduce the worst-decile average unemployment rate from 25 to 8.9 percent.

TABLE 4

Worst-decile average results

	Banks	No banks	Risky banks
Inflation	2.7	1.5	2.9
Output gap	11	39	9.5
Unemployment rate	8.9	25	7.1
Unemployment duration	13	26	9.3
Job loss rate	0.69	1.7	.81
Volatility of output gap	5.2	20	5.0
Volatility of inflation	0.94	1.7	0.67
Annual bank failure rate	1.4	0	4.3
Fraction of banks in trouble	13	0	1.8

7.2.1 Risky banks

To gauge how much difference bank regulation makes, we consider an alternative “risky” situation which differs from our baseline calibration in two senses. First, 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. Second, we suppose that the minimum required capital adequacy ratio is 0.02 instead of 0.08.

As the third column of Table 3 above indicates, macro performance under the risky scenario is roughly comparable to performance under our baseline safe scenario. Median unemployment is the same, the median output gap is a little higher but its volatility is a little lower. As expected, the incidence of bank failures is much higher, but because trouble is defined more loosely the median fraction of banks in trouble is much lower.

However, as Table 4 indicates worst-case performance is much better under the risky scenario than with safe banks. This is clearly because the critical role that banks play in mitigating bad times is to facilitate the entry of new shops by their willingness to lend. "Risky" banks will perform this role better than "safe" banks because their higher loan-to-value ratio makes them automatically willing to lend more when they are not in trouble and also because their lower capital asset ratio means that they are less often in trouble and hence less often unable to lend.

These insights are corroborated by Figure 13 above, which shows that the hump-shaped impulse response of GDP to a shop failure is more moderate with risky banks than safe. Figure 12 shows on the other hand that a troubled-bank shock has a larger output effect under the risky scenario, which is hardly surprising because the bank that falls into trouble cuts its lending from a larger average amount when risky than when safe. But this amplification of the troubled bank shock is mitigated by the fact that banks are less often in trouble in the risky scenario.

Another way to see the difference between safe and risky banks is to conduct a policy experiment in which we start with the safe scenario and increase the loan-to-value ratio gradually from 0.5 to the riskier value of 0.9, by increments of 0.1. For each of the five indicated values of the loan-to-value ratio the model was again simulated 10,000 times. The same set of experiments was then repeated for the risky scenario. Figure 14 below shows the results of this experiment. It shows that the increase in loan-to-value ratio resulted in a moderate deterioration of median performance (as measured by the output gap) under each scenario but a marked improvement in worst-decile average performance. Again this suggests that willingness to lend makes relatively little difference in normal times but helps substantially when the economy is in trouble, especially in the "safe" case where the economy suffers more from financial restraint.

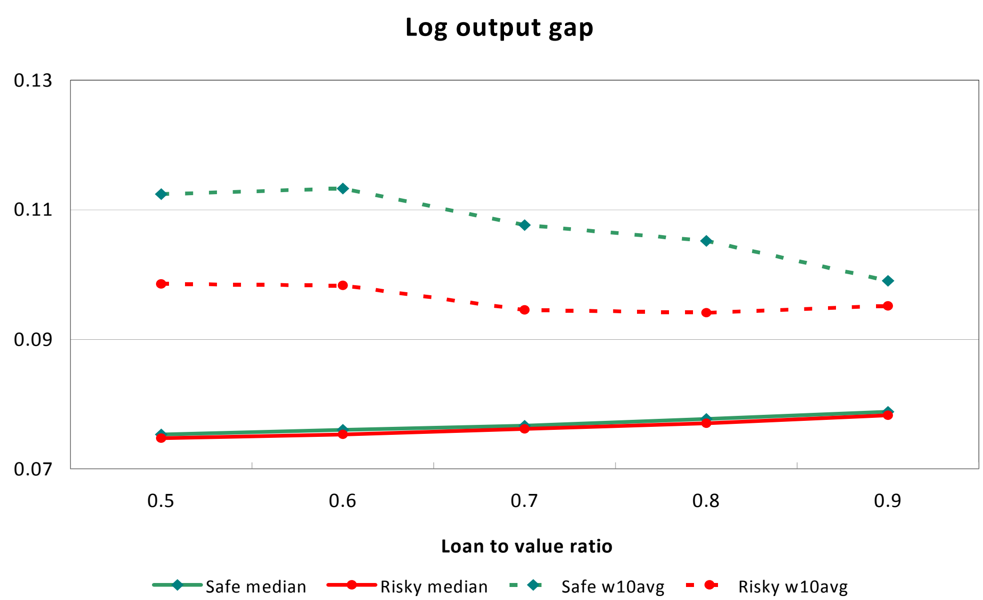

Figure 14: Effect of loan-to-value ratio on output gap

We also experimented with other policy interventions. Figure 15 below shows that varying the capital adequacy ratio κ between 0.02 and 0.12. The figure shows again that there is little effect on median performance but also that worst-case performance is significantly degraded by this tightening of regulation, especially in the “safe” scenario where banks are already overly[18] restrained in their lending.

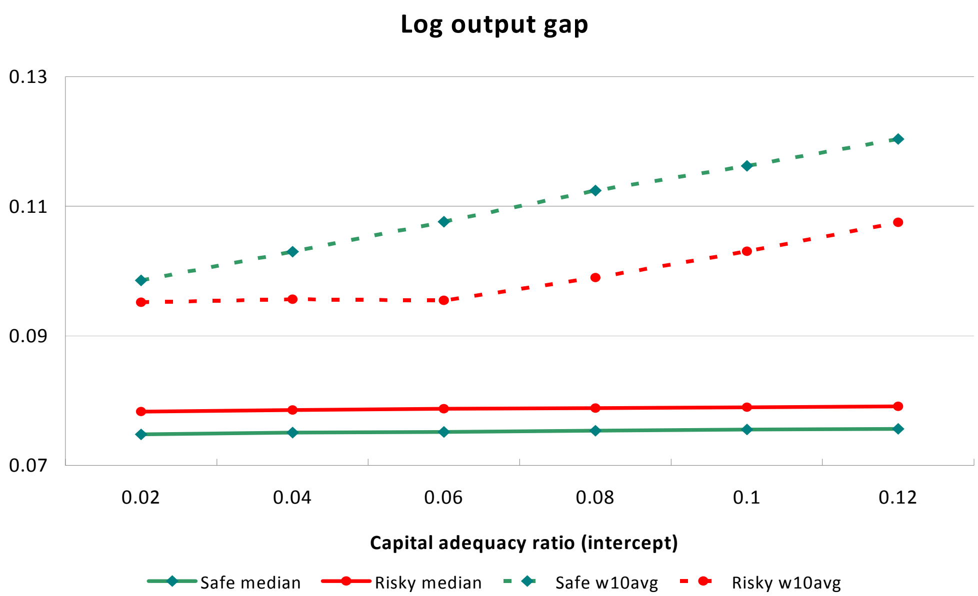

Figure 15: Effect of capital adequacy ratio on output gap

Several commentators have called for making capital adequacy requirements procyclical. Figure 16 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 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 16 shows, the policy has almost no effect on the median output gap or on the worst-case average performance under the risky scenario, but it does result in significantly better worst-case performance under the safe scenario.

[18] "Overly" in the sense of resulting in a higher worst-case average output gap than would be achieved in the alternative scenario.

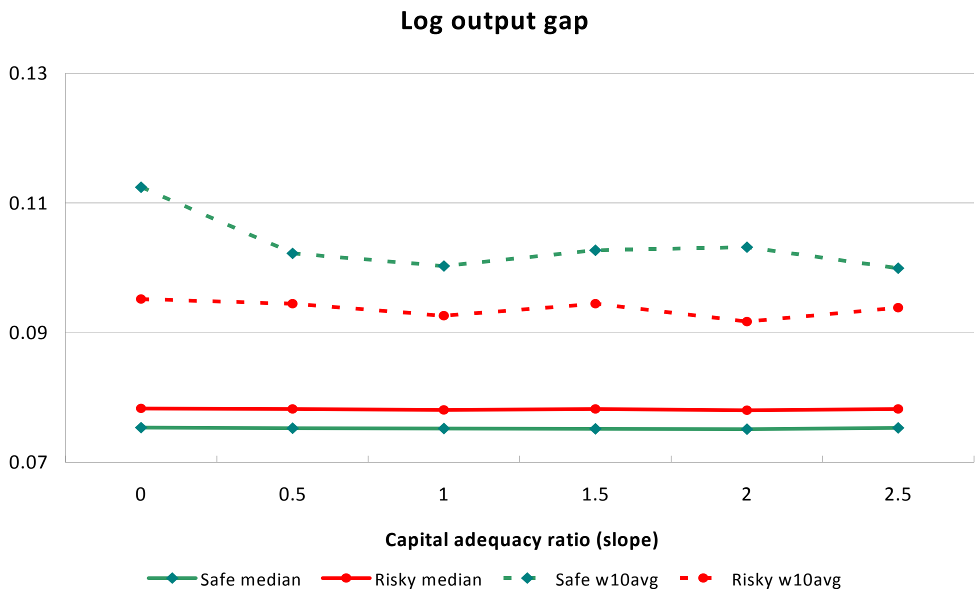

Figure 16: Effect of the slope of capital adequacy ratio with respect to log GDP

8 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 improve the economy's performance by facilitating the entry of shops that organize economic transactions. It also suggests that this improvement is most noticeable in worst-case scenarios. Finally, it suggests that prudential bank regulation in normal times makes little difference for macro performance but in worst-case scenarios it makes the economy perform significantly worse by suppressing the lending activity that is especially needed in bad times.

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.

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.

Bris, Arturo, Welch, Ivo, and Ning Zhu. "The Costs of Bankruptcy: Chapter 7 Liquidation versus Chapter 11 Reorganization." *Journal of Finance* 61 (June 2006): 1253–1303.

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.

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.

Eggertsson, Gauti B. and Michael Woodford. "Optimal Monetary Policy in a Liquidity Trap." *Brookings Papers on Economic Activity* (2003:1): 139–233.

Gertler, Mark, Luca Sala and Antonella Trigari. “An Estimated Monetary DSGE Model with Unemployment and Staggered Nominal Wage Bargaining.” *Journal of Money, Credit, and Banking* 40 (December 2008): 1713–64.

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.

Levin, A., D. Lopez-Salido, E. Nelson and T. Yun. “Limitations on the Effectiveness of Forward Guidance at the Zero Lower Bound.” *International Journal of Central Banking* 6 (March 2010): 143–89.

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.