

Banks, Market Organization, and Macroeconomic Performance: An Agent-Based Computational Analysis*

Quamrul Ashraf[†] Boris Gershman[‡] Peter Howitt[§]

First version: May 11, 2011
This revision: September 1, 2016

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 reach prices at which peoples’ trading plans are mutually compatible. Banks affect macroeconomic performance in this economy because their lending activities facilitate the entry and influence the exit decisions of trading firms. Both entry and exit have ambiguous effects on performance, and we resort to computational analysis to understand how they are resolved. Our analysis draws an important distinction between normal and worst-case scenarios, with the economy experiencing systemic breakdowns in the latter. We show that banks can provide a “financial stabilizer” that more than counteracts the familiar financial accelerator, and that the stabilizing role of the banking system is particularly apparent in worst-case scenarios. In line with this result, we also find that under less restrictive lending standards banks are able to more effectively improve macroeconomic performance in the worst-case scenarios.

Keywords: Agent-based computational model, Market organization, Banking system, Macroeconomic stability, Financial stabilizer.

JEL classification codes: C63, E00, E63, G20, G28.

*We acknowledge valuable suggestions from the editors, Nick Vriend and Mauro Napoletano, and several anonymous referees, in addition to those provided by Blake LeBaron, Ned Phelps, Bob Tetlow, and various seminar participants at the Federal Reserve Bank of Cleveland, the Marschak Colloquium at UCLA, the Center on Capitalism and Society at Columbia, Brandeis University, Columbia Business School, the Federal Reserve Bank of Dallas, the National Bank of Austria, Brown, and MIT, none of whom bears any responsibility for errors or should be taken to endorse our analysis. The C++ code for our computational model is available online at <https://drive.google.com/file/d/0Bw0xeVVEG5XRHdNb19fYUkxY1E/view?usp=sharing>. It was compiled and run as a 64-bit Windows application.

[†]Williams College, Department of Economics, 24 Hopkins Hall Drive, Williamstown, MA 01267 (email: Quamrul.H.Ashraf@williams.edu).

[‡]American University, Department of Economics, 4400 Massachusetts Avenue NW, Washington, DC 20016 (email: boris.gershman@american.edu).

[§]Brown University, Department of Economics, 64 Waterman Street, Providence, RI 02912 (email: Peter.Howitt@brown.edu).

1 Introduction

How do banks affect the macroeconomy? The vast literature on finance and macroeconomics has provided numerous answers to this question, many of which are common knowledge. This paper develops an agent-based computational model to examine a relatively unexplored channel through which banks can influence macroeconomic performance, namely, their role in what Jevons called the “mechanism of exchange.”

In any but the most primitive economic system, exchange activities are organized by a network of specialized enterprises, not just the firms that produce goods and services but also retailers, wholesalers, brokers, and various other intermediaries. These enterprises provide facilities for trading 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 assurance to others that they can buy at times of their own choosing. In short, they play the role in real time that general equilibrium theory assumes is played in meta time by “the auctioneer,” namely that of matching buyers with sellers and establishing prices that coordinate peoples’ trading plans. Moreover, unlike the auctioneer, these intermediaries provide facilities and buffer stocks that allow trading to proceed even when individual plans are imperfectly aligned.

The importance of this network of trading enterprises is attested to by Wallis and North (1986), who show that providing transaction services is the major activity of business firms in the U.S. economy; they estimate that over half of measured GDP in the U.S. consists of resources used up by the transaction process. Indeed, as everyday experience of any household will verify, almost all transactions in a modern economy involve enterprises that specialize in making these transactions.

Banks 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 enterprises that make up the network. Entry of new facilities is neither free nor 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 is often triggered by banks deciding when to cut off finance from a failing enterprise.

The present paper describes a model that portrays this role of banks in helping to coordinate the economy. In a sense, our work is a continuation of a line of research into disequilibrium macroeconomics that began with Patinkin (1956, ch. 13) and Clower (1965), and reached its pinnacle in the Barro-Grossman (1976) book. That line of research ran into the problem that the failure of one market to clear generates rationing constraints that affect traders in other markets in complicated ways that are analytically intractable. To deal with such complexities, we have chosen to model the mechanism of exchange in an agent-based computational framework.¹

More specifically, we use a modified version of the model originally developed by Howitt and Clower (2000), in which an economy’s network of trade specialists was shown to be self-organizing

¹A wide-ranging survey of the literature using this method in economics is provided by the various contributors to Tesfatsion and Judd (2006).

and self-regulating. Howitt and Clower find that starting from an initial situation in which there is no trading network, such a network will often emerge endogenously and will, in the absence of shocks, also guide the economy to a stationary state in which almost all the gains from trade are fully realized. Here, we extend the model to allow for durable goods, fiat money, and government bonds, to include monetary and fiscal authorities, and to incorporate banks that lend to the trade specialists.² We additionally introduce various random shocks that prevent the system from ever settling into anything like Howitt and Clower’s fully coordinated stationary state. We calibrate the model to U.S. data and simulate it many times under different parameter values to see how banks and various dimensions of their lending behavior affect macroeconomic performance.

Although our model is admittedly too stylized to be used for policy-making purposes, it does produce three interesting results. First, it provides a framework for understanding “rare disasters.” Most of the time the evolving network of trade intermediaries performs reasonably well in counteracting shocks and keeping the economy in the neighborhood of full capacity utilization, but in a small fraction of simulation runs, the economy spirals out of control. The model thus exhibits what Leijonhufvud (1973) called “corridor effects” – that is, the system’s self-regulating mechanism is unable to counteract shocks beyond a certain point. The distinction between normal and worst-case scenarios shows up dramatically in almost all the experiments we perform on the model.

Our second result concerns the stabilizing influence of banks on macroeconomic performance, especially in worst-case scenarios. It is generally accepted that although finance may help promote economic growth and development, this long-run benefit comes at a cost of increased short-run volatility. This notion is embodied in the basic idea of the financial accelerator that Williamson (1987), Bernanke and Gertler (1989), Holmstrom and Tirole (1997), Kiyotaki and Moore (1997), and others have shown can amplify the effects of macroeconomic shocks because of the endogenous nature of collateral. Our model, however, shows that banks that make collateralized loans to business firms also provide an important “financial stabilizer,” which can potentially be more powerful than the financial accelerator. In particular, when a negative macroeconomic shock is accompanied by firm failures, the presence of banks that can finance replacement firms and sustain other existing firms will often dampen the effects of the shock by ameliorating or even averting a secondary wave of failures that would amplify the drop in output and employment.

Finally, related to the role of banks as financial stabilizers, our third result is that under less restrictive lending standards, due to either higher loan-to-value ratios or lower capital requirements, banks are able to more effectively improve macroeconomic performance in the worst-case scenarios. Thus, in bad states of the world, there exists a conflict between micro-prudential bank regulation and macroeconomic stability.

The next section contains a brief literature review. Section 3 discusses the basic elements of our model. Section 4 describes the protocol by which agents interact in the model, as well as the behavioral rules that we are imputing to them. Section 5 describes a full capacity utilization

²Some of these extensions were incorporated earlier in Howitt (2008). A model similar to the one developed here, but without private banks, is used by Ashraf et al. (2016) to investigate the effects of trend inflation on macroeconomic performance.

stationary state that the system approximates in the hypothetical absence of shocks, and it also discusses the ways in which entry, exit, and bank lending affect the economy’s performance. Section 6 describes how we calibrate our model and discusses its ability to match various empirical facts. Section 7 reveals our main results. Section 8 considers the robustness of our results to alternative assumptions regarding the banking system, and section 9 concludes. Additional modeling details and supplementary results, including a sensitivity analysis, are relegated to appendix sections.

2 Previous literature

There is a large literature on the effects of financial intermediation on long-term growth.³ Our paper focuses on fluctuations rather than growth, and it therefore speaks more to the empirical literature on the effects of financial development on stability (e.g., Easterly et al., 2001; Braun and Larrain, 2005; Raddatz, 2006; Loayza and Raddatz, 2007). In particular, our notion of banks as financial stabilizers is consistent with the results of Braun and Larrain (2005) and Raddatz (2006) to the effect that financial development tends to reduce macroeconomic volatility. Our findings are also in line with those of Coleman and Feler (2015) who document that Brazilian localities with greater access to government banks showed better macroeconomic outcomes during the 2008–2009 financial crisis, relative to localities with a lower share of government banks. Importantly, the authors argue that this pattern could be explained by the fact that the former localities received substantially more counter-cyclical credit, as government banks were better-positioned to provide such support during the downturn.

As mentioned in the introduction, financial accelerators arising from endogenous collateral have been modeled by various authors and have also been introduced into New Keynesian DSGE models (e.g., Gertler and Kiyotaki, 2010) with a view to understanding their role in credit crises.⁴ As we shall see, in the present model, there is a financial accelerator that operates through a bank-lending channel. Specifically, the failure of one firm can impose losses on the bank that has lent to the firm, thereby making it more likely that the bank will further restrict its lending to other firms, thus increasing their likelihood of failure. Since failed firms are not immediately replaced, each failure has an adverse effect on aggregate output, and the above mechanism will tend to amplify this negative impact. Nevertheless, we find that in worst-case scenarios these effects are more than offset by the above-mentioned financial stabilizer.

More generally, our paper adds to the growing literature on agent-based macroeconomic modeling. This literature traces its roots back to the micro-simulation models of Orcutt et al. (1976), Eliasson (1977), and Bennett and Bergmann (1986), which aimed at capturing heterogeneous microeconomic behavioral rules operating in various sectors of an otherwise standard Keynesian macro model. More recently, such authors as Basu et al. (1998), Howitt and Clower (2000), Howitt (2006), Deissenberg et al. (2008), Dosi et al. (2010, 2013, 2015), Delli Gatti et al. (2011), Dawid et al. (2012), and Ashraf et al. (2016) have moved the literature closer to the ideal of modeling the

³For an introduction to this literature, see Levine (2005) or Aghion and Howitt (2009, ch. 6).

⁴See Chinazzi and Fagiolo (2015) for a general survey on the effects of finance on economic stability.

economy as a system of interacting “autonomous” agents, that is, agents capable of behavior in an unknown environment.⁵

What distinguishes our approach from these other agent-based macro models is mainly our focus on the self-organizing capabilities of the system of markets through which agents interact. In our model, shocks are intermediated by agents that spontaneously open new markets, while the same shocks are also capable of disrupting markets by inducing or forcing such agents to close their operations. In other agent-based models, consumer goods are typically homogeneous and failing producers are automatically replaced by new ones, so the entry and exit of firms do not play the key role in the transmission of macro shocks as they do in our model. We do not, however, attempt to integrate in our analysis of fluctuations an explicit model of the innovation process underlying long-term economic growth as do Dosi et al. (2010, 2013, 2015). Nor do we have the detailed analysis of heterogeneous physical and human capital as in Deissenberg et al. (2008) or Dawid et al. (2012).

Several agent-based models have focused on the role of finance in business fluctuations. For instance, Delli Gatti et al. (2005, 2010) explore the implications of financially fragile firms in a world with a financial accelerator, modeled along the lines of Greenwald and Stiglitz (1993). Dosi et al. (2013) and Assenza et al. (2015) expand this analysis into a full-blown closed agent-based macro model and explore various policy issues. Cincotti et al. (2010) simulate the macroeconomic implications of firms’ dividend policy in the EURACE model, which is also a closed agent-based macro model with credit rationing and financially fragile firms. Generally speaking, these papers do not investigate the implications of having a more risky banking system or having a banking system at all, as we do.⁶ Furthermore, of all these papers, only Dosi et al. (2015) allows banks to fail, and none of them allows banks to be sanctioned by regulators when they violate their capital adequacy requirements. As we discuss below, our modeling of banks’ lending decisions is quite different from all these papers. Popoyan et al. (2016) modify an earlier version of our model to include a more detailed regulatory structure and dependence of lending behavior on firms’ financial situations, and they study the effects of alternative macro prudential and monetary policies.

3 The model

Our model portrays, in an admittedly crude form, the “mechanism of exchange” through which economic activities are conducted in a decentralized economy. At the heart of the model is a self-organizing network of firms that coordinate all production and trading activities. Macroeconomic fluctuations in the model arise, in part, from disruptions to this network due to firm turnover, resulting in the break-up of established trading relationships in both labor and goods markets. Since exit and entry decisions depend on the availability of credit, banks can mitigate fluctuations by alleviating the financing constraints faced by existing firms and potential entrants.

⁵Leijonhufvud (1993, 2006) has long been a forceful advocate of such an approach to macroeconomic modeling.

⁶Dosi et al. (2013) do investigate the consequences of a regulatory policy that restricts the amount of bank credit.

In some ways, the structure of our model and its macroeconomic aggregates are largely comparable to the canonical New Keynesian analysis of Woodford (2003). Specifically, prices are set by firms acting under monopolistic competition, the rate of interest is set by a monetary authority following a Taylor rule, and individual consumer demand depends, *inter alia*, on current wealth. Nonetheless, our model differs significantly from the standard New Keynesian framework in other elements of its structure and in its behavioral assumptions.

First, both goods and labor markets in our model are characterized by costly search, which we consider to be a critical part of the coordination process. Second, firm turnover, which is the principal source of macroeconomic fluctuations in our model, is absent from the canonical New Keynesian framework that postulates a fixed population of firms. Third, our model assumes that the only available financial instruments are non-contingent bank deposits, bank loans to shops, and government-issued money and bonds, whereas the standard New Keynesian setup assumes a complete set of contingent financial markets. Fourth, we assume heterogeneous preferences and endowments, thus emphasizing the complexity of the coordination problem in the real world. Fifth, there are no aggregate shocks in our model, because with a finite number of agents, the economy does not need such shocks to exhibit persistent aggregate fluctuations. Finally, instead of deriving behavioral rules from utility maximization under precoordinated rational expectations, we postulate them as primitive elements of our model, following the common practice in agent-based computational frameworks.

3.1 Modeling principles

3.1.1 Behavioral rules

In contrast to rational-expectations-equilibrium macroeconomics, the underlying principle of the agent-based approach is that people interact in complex environments using simple behavioral rules. Of course, there are many simple rules that one could attribute to any actor in an agent-based model. In choosing between alternatives, we have favored rules that are familiar to macroeconomists, in the interest of making our analysis more accessible to such an audience. We have also favored rules that appear to generate behavior that is roughly consistent with empirical evidence at the micro level. Where neither criterion seemed applicable, we have favored rules that appear plausible to us. Our strategy is thus not very different from that of most other agent-based macro models.

3.1.2 Banks

Our choice of how to model bank behavior is especially constrained by considerations of simplicity. The banks in our model do not speculate on their own account. The only financial assets they invest in are government bonds and fully collateralized commercial loans. Banks do not engage in maturity transformation, and they do not compete with one other (each serves a distinct group of

customers). There is no interbank market.⁷ Furthermore, as described in section 3.2.4 below, no bank is “too big to fail.”⁸

In our baseline model, each bank is assumed to accede to all its customers’ loan requests, subject to an exogenously fixed loan-to-value ratio, except when the bank is at risk of violating its capital adequacy requirement. Thus, our results should be interpreted as saying how banks might affect fluctuations under ideal circumstances if they could be induced and/or permitted to lend freely, even when the economy is in a downturn.⁹ In section 8.2 below, we explore a somewhat more realistic structure in which banks become less willing to lend (that is, they lower the loan-to-value ratio) when economic activity slows down. This investigation suggests that our main results continue to hold under more realistic assumptions with respect to bank behavior.

3.1.3 Expectations

Expectations concerning macroeconomic aggregates play less of a role in our model than in standard macro theory, primarily because we wish to emphasize disruption to trading networks, rather than expectational dynamics, as a source of fluctuations. Whenever agents need to forecast inflation (for example, when setting wages), we assume that they use the government’s targeted rate of inflation as their forecast. This simple assumption accords with the claim made by many inflation-targeting central banks that inflation targets anchor peoples’ expectations.

There is, however, one area in which agents in our setup need to make somewhat more sophisticated forecasts than would normally be consistent with the spirit of agent-based modeling; namely, in computing their perceived wealth, which is assumed to be the sole determinant (modulo a finance constraint) of each agent’s demand for goods. Because monetary policy acts on a very short-term interest rate, some expectational mechanism is needed for it to have a quantitatively significant effect on agents’ perceived wealth.

In standard DSGE macro, this effect comes through each person’s expectation of the future path of real interest rates. In our model, it would be incongruous to assume that agents attempt to form such expectations. Accordingly, we suppose that the government makes the predictions for them. Specifically, like most real-world central banks, the monetary authority forecasts the future time paths of inflation and aggregate output, feeds those forecasts into its Taylor rule, and thereby arrives at a forecast of the future path of real interest rates. It then publicly disseminates an estimated capitalization factor V , which indicates the estimated present value of a constant future real flow of income equal to one dollar at current prices. Each agent forms a forecast of

⁷The robust-yet-fragile property of an interconnected financial system, discussed extensively by Chinazzi and Fagiolo (2015), suggests that adding an interbank market would make our economy’s banking sector more stable up to a certain point.

⁸In all these respects, Dosi et al. (2015) model banks in a similar fashion as we do. In contrast to our setup, however, which assigns the same number of customers to each bank, they use a random assignment to match the size distribution of banks. They also model bank lending behavior quite differently from us.

⁹As discussed in section 2, Coleman and Feler (2015) provide a real world example in which government banks have played a crucial role in this regard.

his future wage income (or profit, in the case of a shop owner) and applies the publicly available capitalization factor to that forecast when calculating perceived wealth.

The public dissemination of a capitalization factor in our model is our way of capturing the “forward guidance” provided by many inflation-targeting central banks. When central banks signal a higher likelihood of future increases in interest rates, the intention is to trigger a downward revision in asset values and wealth, which is what happens in our model when the government reduces its publicly announced capitalization factor.

3.2 The main components of the model

3.2.1 People, goods, and labor

There is a fixed number N of people, 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” $t = 1, \dots, T$. 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 (one “dollar”) next week.

Each person has a fixed type (i, j) , where $i \neq j$ and $i \neq j + 1 \pmod{n}$, meaning that he is endowed every week with one unit of labor of type i (good i is his “production good”) and can consume only goods j and $j + 1 \pmod{n}$ (his primary and secondary “consumption goods”).¹⁰ We assume that there is exactly one person of each type. Thus, the population of the economy is $N = n(n - 2)$. In everything that follows, we set the number of goods n equal to 50, implying a population of $N = 2400$ people.

3.2.2 Shops, production, and trading

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, converting type i labor into good i , and selling good i for money. The number of shops of each type will evolve endogenously.

To trade with a shop, a person must form a trading relationship with it. Each person may have a trading relationship with at most one shop that deals in his production good, in which case that shop is his “employer” and he is one of its “employees,” and with at most one shop that deals in each of his consumption goods, in which case that shop is one of his “stores” and he is one of its “customers.” Each person’s trading relationships will evolve endogenously.

Each shop of type i has a single owner, whose production good is i .¹¹ To ensure that shops fail often enough, we assume that each active shop incurs a fixed cost. Specifically, it faces a fixed

¹⁰This fragmented structure (compared with the more conventional assumption that everyone consumes all goods) captures what we believe to be an important aspect of real-world trading networks, namely that most people trade with only a small fraction of all business firms.

¹¹We make the assumption that the shop’s type must coincide with the shop owner’s production good merely to simplify the analysis. Without such an assumption, we would have to model each potential entrant’s choice of what kind of shop to open.

overhead cost of F units of type i labor per week, in addition to a variable cost of one unit of type i labor per unit of good i produced. To ensure that entry is costly, we also assume that when the shop is first opened, the owner 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 takes place at prices that are posted in advance by the shop. Specifically, every shop posts a price p and a weekly wage rate w , each of which may be adjusted periodically.

There is no depreciation or other physical storage cost. Goods produced but not sold in a week are held in inventory. Fixed capital cannot be used for any purpose (other than serving as collateral for loans) before the shop exits. Former shop owners that still hold fixed capital cannot consume it but continue to hold it in the hopes of selling it to another shop in special “firesale” markets to be described below. Likewise, they continue to hold the former shop’s inventory until sold in a firesale market. The fixed capital and inventory still held by former shop owners are referred to as “legacy capital.”

3.2.3 Banks, deposits, and loans

There is 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, and the number of agents is the same in each sector. There is one bank per sector, owned by a single person in the sector. Each agent can only deal with the bank in his sector.

The two main functions of a bank are to accept deposits from agents in its sector and to give out loans to shop owners in its sector. Loans are made with full recourse but are also collateralized by inventory and fixed capital. Each bank applies a haircut to collateral, as determined by its fixed loan-to-value ratio h . Specifically, for each unit of inventory or fixed capital it accepts as collateral, the bank will lend h times the value of the collateral (see section 4.4 below). Banks cannot lend to consumers who are not shop owners.

Since loans are made with full recourse, if a shop owner is unable to repay his bank loan, the bank seizes all of his money and deposits, as well as the shop’s inventory and fixed capital, until the value of what has been seized equals the amount owed. Foreclosure entails a resource cost. All seized assets remain on the bank’s balance sheet until sold in the firesale markets.

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 and are subject to a required capital adequacy ratio.¹²

3.2.4 Government and its policies

In our model economy, the government conducts fiscal policy, monetary policy, and bank regulation. Specifically, the government collects taxes, sets interest rates, lends to banks, enforces capital adequacy requirements on banks, and insures bank deposits. It does not purchase goods or labor

¹²In the baseline calibration below, the lender-of-last-resort facility was hardly ever used; specifically, out of 10,000 runs of 2,880 weeks each, there were only 130 weeks during which a bank used the facility.

but does issue money and bonds, and it services the interest on bonds through a sales tax on every goods market transaction. It adjusts the ad valorem tax rate τ once per year.

In setting interest rates, the government acts like the central bank of standard New Keynesian theory. More specifically, it pegs the interest rate i on government bonds by buying or selling whatever quantity the banks wish to hold at that rate. It adjusts this interest rate every 4 weeks according to a Taylor rule. In addition, like many real-world central banks, the government in our model forecasts future GDP, inflation, and interest rates, and it releases these predictions to the public.

As bank regulator, the government requires each bank to maintain equity of at least a fixed fraction κ of the value of its “risky assets” (bank loans and seized collateral), where κ is the required capital adequacy ratio. A bank that is not satisfying this constraint is declared to be “troubled,” and it is forbidden to initiate any new loans or to pay its owner a dividend.

A bank with negative equity is forced into failure. We model bank failures after the pattern that became common among small banks in the United States following the 2008–2009 financial crisis, namely, the FDIC closes the bank, injects new capital, finds a new owner, and allows the bank to continue operating under the new ownership with essentially no loss of deposits. Thus, in our model, the government seizes the wealth of a failed bank’s owner and injects it into the bank, adds 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 have to the bank’s holdings of foreclosed capital, and he also 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) remaining unchanged.

4 Protocol and behavioral rules

Every week, the actors in our model proceed sequentially through the following nine stages: 1) entry, 2) search and matching, 3) financial market trading, 4) labor and goods market trading, 5) monetary policy, 6) match breakups, 7) fiscal policy, 8) exit, 9) wage and price setting. The rest of this section describes the protocol and behavioral rules governing agents’ actions at each stage.

4.1 Entry

Following Baumol (2010), we assume that entrepreneurship is a scarce resource. Accordingly, in the first stage, each person who is not already a shop owner or a bank owner becomes a potential entrant, or “entrepreneur,” with probability θ/N , where the parameter θ represents the supply of entrepreneurship.

Setup cost. In line with Fonseca et al. (2001), who show that start-up costs constitute an important barrier to entry, we assume that an entrepreneur of type (i, j) has the opportunity to open a shop of type i if he is able to immediately defray the setup cost of S units (in total) of his two consumption goods. This fixed capital can be obtained from the entrepreneur’s own legacy

capital (if any), from firesale markets, or from stores with which the entrepreneur has a trading relationship. The entrepreneur is allowed to find out from the firesale markets and from his stores how much capital is available and at what price.¹³

Credit line. Each entrepreneur may apply for a line of credit to his bank, and the bank will either approve or deny the application before the final entry decision is made. The bank’s rules for credit line approval and maximum loan size determination will be made evident in section 4.3.

An entrepreneur who has discovered that there is enough fixed capital available for him to enter and has also learned whether or not his bank approved his credit line application next decides whether to enter the market or let the opportunity lapse. We assume that he goes through the following decision process.

Business plan. The entrepreneur first computes the nominal setup cost S_N , which is the minimum cost of acquiring fixed capital from all available sources, and he then formulates a business plan for his prospective shop, consisting of a wage rate w , a markup μ , and a sales target y^{trg} . The business plan is just an initial setting, made with little knowledge of actual market conditions, and subject to change in response to experience.

If the entrepreneur enters, he will have to keep his wage fixed for Δ weeks, where the length Δ of the wage contract is an exogenous parameter. Realizing that he will have to compete with other firms for workers, and taking into account the need to offer a “fair” wage (e.g., Akerlof and Yellen, 1988), the entrepreneur sets his initial wage rate equal to his estimate of the economy-wide average wage rate that will prevail halfway through the contract period: $w = W(1 + \pi_w^*)^{\frac{\Delta+1}{2}}$, where π_w^* is the government’s exogenously fixed target weekly inflation rate and W is last week’s (publicly known) average wage.

He picks his markup randomly from a uniform distribution over $[0, 2\bar{\mu}]$, where $\bar{\mu}$ is a parameter measuring the average percentage markup over variable costs. This initial markup will always be his “normal” markup if he enters, but as we shall see, he will sometimes charge a higher or lower than normal markup depending on the inventory-to-sales ratio.¹⁴ The initial sales target y^{trg} is chosen randomly from a uniform distribution over $[1, n]$. Upon entry, this target will be altered each week in response to actual sales.

Financial viability and profitability. Many studies have shown that financial strength and profitability are important determinants of successful entrepreneurship.¹⁵ Accordingly, we assume that the entrepreneur tests for both of these factors.

¹³Specifically, each of his stores’ inventories are available at their respective posted prices, whereas in the firesale market, capital is sold at a publicly known firesale price P_f , to be discussed in section 4.4. Each firesale market has a queue of sellers who are either banks offering to sell foreclosed capital or former shop owners offering to sell their legacy capital.

¹⁴We have opted for random markups rather than having all entrepreneurs charge the same markup, in order to make the model evolutionarily adaptive. When conditions change, Darwinian competition will alter the average markup of surviving firms. Alternatively, we could have allowed each firm to adaptively learn its own “normal” markup, but this seemed more complicated to us.

¹⁵For studies at the firm level, see, e.g., Austin and Rosenbaum (1990) or Khemani and Shapiro (1986). For aggregate evidence on the importance of finance for entrepreneurship, see Levine (2005).

Given his business plan, the entrepreneur now conducts a financial viability test. Specifically, if he cannot raise enough finance to pay for the setup cost and the fixed cost of operating the shop during the first month, $S_N + 4(F - 1)w$, then he will allow the entry opportunity to lapse.¹⁶ The amount of finance he can raise at this stage is his personal wealth (money holdings plus deposit holdings) plus the credit limit provided by the bank. The latter equals 0 if the entrepreneur has not been granted a line of credit, and $P_h(S + \text{LI})$ otherwise, where LI is the entrepreneur's stock of legacy inventories (if any).¹⁷ That is, as mentioned in section 3.2.3, the loan is collateralized by fixed capital and inventory, to which a haircut P_h is applied by the bank. The rule for setting P_h will be stated in section 4.4.

If the financial viability test is passed, the entrepreneur moves on to a profitability test. Specifically, he allows the entry opportunity to lapse unless $\Pi > Y^p + (P_f \cdot \text{LC} + S_N)/V$, where Π is the flow of profit implied by the business plan, and the right-hand side is an estimate of the weekly cost of entering, namely, the entrepreneur's current estimate of his permanent income Y^p , as described in section 4.3 below, plus the equivalent permanent flow of income that could be achieved from selling his legacy capital LC on the firesale markets at price P_f and putting the proceeds, along with the money S_N that he saves by not incurring the setup cost, in his bank account.¹⁸ Further, as will become evident in section 4.5, V is the capitalization factor provided by the government.

Market research. The entrepreneur also conducts “market research” before deciding whether to enter. In particular, he sends messages to two people, one of whom is a “comrade” (a randomly chosen person with the same production good as the entrepreneur) that does not own a shop and the other of whom is a prospective customer (a randomly chosen person whose primary consumption good is the same as the entrepreneur's production good). The message to the comrade declares the wage rate w that the entrepreneur will post if he decides to open shop, and the message to the prospective customer declares the price he plans to post: $p = w(1 + \mu)/(1 - \tau)$. This price represents an after-tax markup of μ over the marginal cost w . The comrade will consent to forming an employment relationship with the new shop either if he is unemployed or if the shop's offered wage rate exceeds his current effective wage. Similarly, the prospective customer will consent to forming a customer relationship with the new shop if its declared price is lower than the effective price that he currently faces for the same production good as the entrepreneur's. The definitions for effective wages and prices are given in section 4.4.

Entry decision. The entrepreneur then makes his entry decision. He decides to enter if: 1) he can raise the necessary amount of fixed capital to cover the setup cost; 2) his business plan passes financial viability and profitability tests; and 3) both the comrade and the prospective customer respond affirmatively to his invitation to form a relationship with the new shop. If any of the above three conditions is violated, the entrepreneur decides not to open a shop, and the opportunity to do so lapses. Otherwise, he moves on to purchase enough fixed capital to cover his setup cost, part

¹⁶Since the shop owner supplies his own unit of labor, the weekly fixed cost of operating the shop is $(F - 1)w$.

¹⁷Note that, in contrast to legacy fixed capital, legacy inventories cannot be part of S , because inventories are units of the person's production good, whereas S must be defrayed using his consumption goods.

¹⁸See appendix A.1 below for the accounting details of how Π is defined.

of which, up to the credit limit, may be covered by the bank loan if his credit line application has been approved. He then opens the shop with a posted wage w , a posted price p , a markup μ , a sales target y^{trg} , and an inventory level equal to the entrepreneur’s legacy inventory LI . In accordance with the well-established principle of production smoothing (see, e.g., Ramey and West, 1999), the entrepreneur also sets an “input target:” $x^{\text{trg}} = y^{\text{trg}} + F + \lambda_I(y^{\text{trg}} - I)$, where λ_I is the exogenous weekly inventory adjustment speed and I is current inventories, which for a new shop are just equal to the entrant’s legacy inventories LI .¹⁹ Implicit in this rule is a desired inventory level equal to one week’s sales.

The comrade and the potential customer, surveyed during market research, becomes the new shop’s actual employee and actual customer, respectively. The entrepreneur becomes a shop owner and terminates any employment relationship he might have had.

4.2 Search and matching

Next, each person is given an opportunity to search for possible trading relationships, comprising both job search and store search. We take the approach of Stigler (1961), who modeled the search process as a search for stores with which to engage in repeated rather than one-off exchange.²⁰

Job search. Each person who is not a shop owner engages in job search with an exogenous probability σ . Job search consists in asking one comrade what his effective wage is. If the comrade is not a shop owner and his effective wage exceeds the searcher’s, the latter will ask the comrade’s employer for a job. If the comrade is a shop owner (i.e., is an employer himself), the searcher will ask him for a job if and only if the latter’s effective wage is lower than the latest posted wage at the comrade’s shop. In either case, the employer accepts the searcher as a new worker if and only if the amount of labor input he used last period is less than or equal to his current input target. If this condition holds, the searcher terminates his pre-existing employment relationship (if any) and forms a new one, changing his effective wage either to his comrade’s (if the comrade is not a shop owner) or to the latest posted wage at the comrade’s shop.

Store search. Store search is undertaken by every person. It comprises referral-based and direct search. First, the person asks a “soulmate” (a randomly chosen person with the same primary consumption good) for his effective prices. The searcher decides to switch to a soulmate’s store if and only if the soulmate’s effective price for that good is less than the searcher’s. Second, the person engages in direct search, which consists in asking a randomly chosen shop if it trades either of his consumption goods and, if so, what the posted prices are. The searcher decides to switch if and only if his effective price exceeds the shop’s latest posted price. If the searcher switches for either of his consumption goods, he terminates the pre-existing customer relationship with a shop of

¹⁹Our “input target” rule is similar to the labor demand function posited by Dosi et al. (2010).

²⁰This is also Phelps and Winter’s (1970) “customer market” view of product markets. It means that the exit of a shop involves the destruction of organizational capital in the form of what would otherwise have been a long-term customer relationship. Dosi et al. (2010) adopt the “customer market” approach as we do, but they do so without an explicit matching process.

the corresponding type and forms a new one, revising his effective price either to his soulmate’s (in the case of referral-based search) or to the shop’s latest posted price (in the case of direct search).

4.3 Financial market trading

At this stage, all financial transactions take place except for loans acquired to finance the purchase of fixed capital by entrants, as described in section 4.1, and the purchase of working capital in firesale markets during the trading stage that comes next.

Bank regulation and lending policy. Each bank’s balance sheet looks as follows:

<i>Assets</i>	<i>Liabilities and Equity</i>
Commercial loans	Deposits
Seized collateral	Loans from central bank
Government bonds	Equity
Reserves	

On the assets side, commercial loans are loans made by the bank to shop owners, measured as the dollar value of principal and interest payable this week; seized collateral consists of inventories and fixed capital seized by the bank from defaulting shops, valued at the firesale price; government bonds are bonds held by the bank, evaluated as the amount due this week; and reserves are holdings of high-powered money, that is, cash and money on deposit with the central bank wing of the government. 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 the bank’s assets minus its liabilities.

At the beginning of this stage, each bank updates its equity after the previous week’s transactions and the entry stage. Before any financial market trading takes place, banks in all sectors are examined by the regulatory branch of the government. Banks with negative equity fail, in which case the government acts much like the FDIC in the U.S. First, it injects money to fully recapitalize the new bank so that it fulfills the minimum capital requirement, discussed further 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 amongst them (i.e., the one with the highest sum of cash and deposit holdings) becomes the new owner.²¹ If the new bank owner has some legacy capital, it is placed on the bank’s balance sheet (seized collateral account). Equity is again updated to take into account these changes to the balance sheet. Note that no deposits are ever destroyed during the bank failure process, because the government’s policy of injecting new capital and finding a new owner serves to fully insure all deposits.

Next, all banks are checked for capital adequacy. In particular, as under Basel I, the ratio of bank equity to risk-weighted assets must be greater than or equal to κ , the required capital adequacy ratio: $\text{equity} \geq \kappa \cdot (\text{commercial loans} + \text{seized collateral}) \equiv \text{required capital}$.²² If this

²¹We assume the richest customer is given the bank because otherwise, given the limited ability of anyone but shop owners to raise capital, there might be an artificial fragility to the system. We can report, however, that even if the new owner is chosen at random from among the bank’s customers, our baseline results are largely unaffected.

²²The assigned risk weights (1 for loans and seized collateral, and 0 for government bonds and reserves) come directly from Basel I recommendations.

condition is violated, that is, if equity is less than the required capital, the corresponding bank becomes troubled. As recommended by Basel I, troubled banks are sanctioned; in particular, they are not allowed to provide new loans (although they are allowed to roll over outstanding loans), and their owners cannot receive dividends.

The bank now chooses the probability P_{CL} with which it approves any new applications for a credit line this week and during the entry stage next week. Its choice of P_{CL} is motivated by the desire to maintain a buffer above the required minimum capital, in line with the evidence presented by the Bank for International Settlements (1999). If its equity is less than its required capital, it sets P_{CL} equal to zero. Otherwise, P_{CL} increases linearly in the percentage deviation of equity from required capital, with a slope parameter equal to l , until it hits unity.

Once a credit line has been granted, the bank will accept all fully collateralized loan applications from the customer in the future, unless the bank becomes troubled. Thus, unlike in other agent-based macro models, loan decisions in our model are made without regard to the financial situation of the borrower, and the only way in which the bank's financial condition affects its loan decisions is through the dependency of P_{CL} on the bank's capital adequacy ratio. The fact that all loans in our model are fully collateralized, however, minimizes risk faced by the bank and makes any residual risk largely independent of the borrower's financial situation.²³

Each bank also sets the interest rate i_L on all new loans that will be issued this week and the interest rate on all deposits i_D held at the end of this stage. The bank always sets $i_D = i_w$ and $i_L = i_w + s/48$, where i_w is the weekly nominal interest rate set by the monetary authority on government bonds (as discussed in section 4.5) and s is a fixed annual spread, common across all banks. We suppose that these choices are constrained by regulations that prevent a bank from exploiting the monopoly power implied by the limitation of one bank per sector in our setup.

Budget planning. Next, people of all types do their budget planning. Since the only reason for holding money is to spend it during the next stage, each person must first decide on his planned consumption expenditures for the week. In preparation for this decision, everyone first adjusts his permanent income according to the following adaptive rule: $\Delta Y^p = \lambda_p(Y - Y^p)$, where Y is actual income from the previous period and λ_p is the weekly permanent income adjustment speed. Here, Y is equal to last period's profit for shop owners and the effective wage rate for all other people. Then, each person adjusts Y^p for estimated weekly inflation, assuming that inflation is taking place each week at the target rate, that is, he multiplies Y^p by $(1 + \pi_w^*)$.

In line with theoretical and empirical analyses that show wealth to be an important determinant of consumption, we assume that each person sets his planned consumption expenditures this week equal to a fixed fraction v of his total wealth.²⁴ The total wealth of each person is the sum of

²³As we shall see, in our baseline simulations, if a borrower fails, his seized collateral can be sold on "firesale markets" for a real price that is expected to equal the haircut price at which the loan was originally made.

²⁴Consumption is also dependent on wealth in the EURACE model, as well as in Delli Gatti et al. (2011). Dosi et al. (2010, 2013, 2015) assume that consumption equals current labor income. For a brief survey of the empirical literature on wealth and consumption, see Poterba (2000).

financial wealth A and the capitalized value of permanent income. Thus,

$$E = v \cdot (A + V \cdot Y^p), \quad (1)$$

where V is the capitalization factor that converts expected income into wealth. The financial wealth A of people who don't own a shop or a bank is the sum of their money holdings and bank deposits, plus the firesale value of their legacy capital (if any). For bank owners, financial wealth consists of money holdings and, if the bank is not troubled, the bank's equity after subtracting required capital. For shop owners, it is equal to the sum of money and deposit holdings minus outstanding loans.²⁵ In appendix A.4, we describe the ideal circumstances under which the expenditure function (1) would be the optimal expenditure of a life-time utility maximizer possessing logarithmic preferences and a weekly rate of time preference equal to $v/(1 - v)$.

Financial transactions. Next, people of all types update their portfolios of financial assets. Having chosen E , each person chooses the amount of cash M , taking into account the constraints he faces. Consider first a person who does not own a bank or a shop. He enters this stage owning \bar{M} in cash and \bar{D} in deposits, and he must choose M and D with which to leave this stage, subject to neither being negative and to $D = (\bar{M} + \bar{D} - M)(1 + i_D)$. We use the convention of measuring D as the amount owed by the bank at the next stage of financial market trading. If $E \leq \bar{M} + \bar{D}$, the person sets $M = E$, leaving the rest in his bank deposit account. Otherwise, he withdraws all of his deposits and revises his planned expenditures so that $E = M = \bar{M} + \bar{D}$. The idea here is that he will need to have E in the form of money when he visits his stores, but he does not know whether he will be paid his income before or after shopping for goods, so he plans to carry E out of the financial market to ensure against being unable to fulfill his expenditure plans.²⁶

Next, consider a bank owner. If he owns a troubled bank, that is, if the minimum capital requirement is violated, then he cannot receive dividends and his expenditure is bounded by current money holdings \bar{M} . If the latter exceeds E , the remaining $\bar{M} - E$ goes into the bank; otherwise, he sets $E = M = \bar{M}$. If the owned bank is not troubled, then he employs his financial wealth A in the following way. If $E \leq A$, he sets $M = E$ and leaves the surplus $A - E$ in bank equity; otherwise, he sets $E = M = A$.

Finally, consider a shop owner. A shop owner can hold money and deposits, and he can also take a bank loan L (measured as the amount owing next week) up to his credit limit. If the shop has already been granted a credit line earlier, his credit limit is set equal either to the haircut value of his eligible collateral (fixed capital and inventories), in case his bank is not troubled, or to the minimum of the pre-existing loan \bar{L} and the haircut value of his eligible collateral if he has a line of credit with a troubled bank.²⁷ If the shop does not yet have a credit line, the shop owner applies

²⁵We omit fixed and working capital from the definition of a shop owner's wealth, because otherwise, we would be double-counting; that is, the shop owner's income Y includes the flow of profit that this capital helps generate.

²⁶This motivation for a precautionary demand for money is similar to the "stochastic payment process" that Patinkin (1956) invoked 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.

²⁷Recall that collateralized rollover of existing loans is always allowed, even if the bank is troubled.

for one at this stage and obtains it with probability P_{CL} , as described above, if his bank is not in trouble (otherwise, his application is rejected). If the credit line is granted, his credit limit CL is set to $P_h(S + I)$, where I is his shop's current inventory holdings and P_h is the haircut price; otherwise, his CL is set to 0.

If the shop owner is unable, even after borrowing up to his credit limit, to repay his pre-existing loan \bar{L} in full, his shop is declared bankrupt. In this case, his bank seizes all his cash, fixed capital, and inventories, and it nullifies his deposits. Moreover, the shop must close in the exit stage this week. A bank that has seized capital adds the fraction $1 - C_b$ to its holdings of seized collateral, and it joins the queue for the firesale market in each of those goods (if not already there), where C_b is the exogenous resource cost of foreclosure.

The shop owner's desired money holdings is the amount needed to cover not only his planned consumption expenditures but also his target wage bill WB^{trg} , equal to his current posted wage times his current input target x^{trg} minus one (since he does not have to pay himself for his own endowment). Before computing WB^{trg} , the shop owner updates his input target using the rule discussed in section 4.1 above.

The shop owner now borrows, if necessary, to cover his planned consumption expenditure and his target wage bill.²⁸ If his credit limit is not enough to finance these planned expenditures, he borrows up to the limit and reduces planned expenditure to the amount that can be afforded, setting the wage bill as his first priority and decreasing it only after E has been reduced to zero.²⁹ After making the needed adjustments to E and WB^{trg} , he then holds the amount $M = E + \text{WB}^{\text{trg}}$ in cash. Appendix A.2 provides additional details.

Once all agents have completed their portfolio adjustments, a bank's reserves could potentially become negative. We suppose that the bank is able to get cash from the government instantaneously to honor all withdrawals but that its deposit account with the central bank is debited accordingly and can become overdrawn. At this point, the government pays the amount owing on its bonds to the banks that are holding them. If a bank still has negative reserves, it must borrow enough from the government to make its reserves nonnegative again. The annual rate charged on advances to banks is $i + s_d$, where s_d is a fixed premium (the "discount rate"). The bank invests any positive reserves in new government bonds.

4.4 Labor and goods market trading

Firesale markets trading. The next stage starts with trade in the firesale markets. Each shop of type i whose owner has a line of credit and whose bank is not troubled can place an order for any amount Q of good i . He will pay for his purchases with deposit holdings, at least until his

²⁸Running down deposits before taking a loan is consistent with pecking order theory in corporate finance (Myers and Majluf, 1984). For supporting empirical evidence, see Shyam-Sunder and Myers (1999).

²⁹The assumption that the shop owner gives wage payments first priority helps to avoid shop failures that might otherwise result from spending the wage bill on consumption goods.

deposits are exhausted. After that he can borrow using inventory as collateral.³⁰ Any other shop owner is subject to the finance constraint: $P_f Q \leq D$.

If $Q > 0$, the shop owner is matched to the first seller (if any) in the i^{th} queue. If the first seller cannot fulfill the whole order, he sells what he has and the turn goes to the next seller in the queue, and so on, until either the order is fulfilled or the queue runs out of sellers.³¹ A shop's "notional" demand for inventory in this market is $Q^* = \min\{y^{\text{trg}} - I, 0\}$, that is, its desired inventory is its sales target. So it orders an amount equal to Q^* , unless it is subject to the above finance constraint, in which case it orders $\min\{Q^*, D/P_f\}$.

Next, each shop's inventory is reduced by the amount $\min\{F, \bar{I}\}$ to cover the fixed cost, where \bar{I} is the inventory after the firesale transactions. If $F > \bar{I}$, then all input held by the shop must be used to defray the fixed cost, while inventory remains equal to zero, until the fixed cost has been fully covered. After that, any input received adds one for one to inventory.

Labor market trading. Next, each person engages in labor market trading (with his employer) and goods market trading (with his stores), starting with the store that trades his primary consumption good. With probability $1/2$ he either first trades with his employer or with his stores. Labor market trading proceeds as follows. If the person is a shop owner (i.e., if he is self-employed), he simply uses his unit endowment as input. If the person has no employer, then nothing happens when it is his turn to engage in labor market trading. Otherwise, if he has an employer with positive money holdings, he offers to trade his endowment in exchange for the effective wage $w^{\text{eff}} = \min\{w, M\}$, where w is the employer's posted wage and M is the employer's money holdings just prior to this trade. If $M = 0$, then the employee is not required to deliver his input, but his employment relationship with the shop remains intact.

The employer is not obligated to accept an employee's offer to trade. He can instead choose to lay the worker off, in which case no input is delivered, no payment is made, and the employment relationship is severed. We assume that the employer chooses to do so if and only if: (a) the shop's labor input is already higher than its input target, and (b) the ratio of its inventory to its sales target exceeds a critical threshold value $IS > 1$. In general, there are four ways for a person to become unemployed in our model economy: 1) his employer is broke; 2) he is laid off; 3) his job match breaks up randomly (see section 4.6); and 4) his employer exits the market (see section 4.8). Note that the employment relationship is retained only in the first case.

Goods market trading. Goods market trading happens in the following manner. First, the customer learns the price currently posted by both of his stores. Then, when a person visits a store with positive inventory I , he can place an order for some amount c , subject to the cash-in-advance constraint $pc \leq M$, where p is the posted price and M is the person's money holdings just prior

³⁰This makes sense in our calibrated model because we never assume that the loan-to-value ratio h is less than $1/2$ (except possibly in section 8.2, where we allow h to vary procyclically). Therefore, the cost of a unit of inventory (the firesale price P_f) is less than the amount the shop owner can borrow against that unit (the haircut price P_h). The setting of the firesale and haircut prices is discussed below.

³¹The payment for inventories bought at the firesale market is expressed in dollars due at the beginning of next week.

to the visit. The store then gives the person an amount $c^{\text{eff}} = \min\{c, I\}$ in exchange for pc^{eff} of money. The person’s effective price for that good is defined as $p^{\text{eff}} = pc/c^{\text{eff}}$. If $I = 0$, or the person does not have a store for that good, nothing happens at this stage of goods trading, except that the person’s effective price is set to $+\infty$. For every dollar spent by a customer, the shop owner receives $1 - \tau$ dollars, with the rest going to the government to cover the sales tax.

We assume that each customer chooses his desired consumption bundle (c_1, c_2) to maximize his utility function $c_1^{\varepsilon/(\varepsilon+1)} + c_2^{\varepsilon/(\varepsilon+1)}$ (where $\varepsilon > 0$ is a “demand parameter”), subject to his budget constraint. If he has established relationships with stores for both of his consumption goods, the budget constraint is $p_1c_1 + p_2c_2 = E$, where E is the planned expenditure determined in the previous stage and p_1 and p_2 are the posted prices. If he has an established relationship with only one store with posted price p , then he orders the amount E/p from that shop.

Average wage, firesale price, and haircut price. At this point, the government computes the average wage rate W , which is the employment-weighted average across all shops, and makes this information public. All firesale prices are then set for next week at the value $P_f = 0.5 \cdot W \cdot (1 + \pi_w^*)$, where π_w^* is the government’s weekly inflation target. Our rationale here is that P_f splits the difference between the value of inventory to a shop owner (approximately the marginal cost of producing an extra unit) and its value to a seller in the firesale queue for whom the goods have no direct use. The target inflation factor is added in anticipation of next week’s increase in the marginal cost.³²

Finally, each bank sets its haircut price P_h for next week. It is set proportionally to the estimated marginal cost of production: $P_h = h \cdot W \cdot (1 + \pi_w^*)$, where h is the loan-to-value ratio employed by all banks, and the value is measured as the replacement cost of the collateral to the borrower. All banks will therefore have the same haircut price at any given time. For most of the analysis, we take h to be an exogenous parameter, representing the risk tolerance of banks. In section 8.2, however, we will examine the case in which h varies procyclically.

4.5 Monetary policy

Interest rate setting. Next is the stage in which the government (acting as central bank) sets the nominal interest rate i . In this respect, it operates in a similar fashion as the U.S. Federal Open Market Committee, which does not act every week and whose behavior is approximated by a Taylor rule (Taylor, 1993). 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. Otherwise, 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” (4 weeks), 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.

³²The fact that the firesale prices are not directly dependent on the amounts offered for sale implies that firesale markets in our model do not play the deviation-amplifying role that they might in reality. When we suppress the firesale markets and rerun our baseline simulations, we find that our economy actually performs somewhat worse in terms of almost all macroeconomic indicators.

The annual rate of interest i is set according to the Taylor rule:

$$\ln(1+i) = \max\{\ln(1+i^*) + \gamma_\pi [\ln(1+\pi) - \ln(1+\pi^*)] + \gamma_y [y - \tilde{y}], 0\},$$

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-month moving average for the weekly average log GDP, \tilde{y} is the government's evolving estimate of weekly log potential output, and $i^* \equiv r^* + \pi^*$, where r^* is its evolving estimate of the “natural” real interest rate. The Taylor rule, as specified above, respects the zero lower bound on nominal interest rates. The weekly interest rate is determined according to $1+i_w = (1+i)^{1/48}$.

The government's estimate \tilde{y} of log potential output is the steady state of an estimated AR(1) process, which the government updates once per year. Its estimate r^* of the natural interest rate is changed each period as an increasing function of the difference between actual inflation and the target level π^* , on the grounds that an above-target inflation rate might be due to an underestimated natural interest rate. For details on this estimation process, see appendix A.3.

Capitalization factor. The government also estimates an AR(1) process for inflation that it uses together with its estimated process for log GDP to forecast the future values of these series. It then feeds these values of inflation and log GDP into the Taylor rule to compute forecasts of the rate of interest, under the assumption that the coefficients and the values of \tilde{y} and i^* will remain unchanged (see appendix A.3). The government employs these forecasts to calculate the capitalization factor V , which is publicly announced so that agents can use it to estimate the present values of their incomes. Specifically,

$$V = \sum_{t=1}^{\infty} \left(\frac{1}{1+\pi_{wt}} \right) \prod_{k=1}^t \left(\frac{1+\pi_{wk}}{1+i_{wk}} \right), \quad (2)$$

where π_{wk} and i_{wk} are the inflation rate and the interest rate projected for week k . As we explain in appendix A.4, this is precisely the capitalization factor that, under ideal circumstances, would make the second factor on the right-hand side of the expenditure function (1) equal to the right-hand side of an agent's lifetime budget constraint, even in a standard DSGE model.

4.6 Match breakups

Established trading relationships may break up randomly. In particular, as in the standard theory of search and matching, each agent in the economy who does not own a shop is subjected to an exogenous probability δ of quitting the labor and goods markets, entailing the unconditional severance of all current trading relationships by the agent with his employer and his consumption stores.

4.7 Fiscal policy

Next, if it is the last week of the year, the government adjusts the sales tax rate τ . It 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 τ^* , which is the value that would leave the debt-to-GDP ratio undisturbed in the no-shock stationary state to be described in section 5.1, plus an adjustment factor that is proportional to the difference between the actual debt-to-GDP ratio and an exogenous target ratio b^* : $\tau = \tau^* + \lambda_\tau \cdot [B/(P(1 + i_w) \cdot 48 \cdot e^{\tilde{y}}) - b^*]$, 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 Exit

Now, each shop has an opportunity to exit. Consistent with empirical evidence (e.g., Altman, 1968; Beaver et al., 2005; Wiklund et al., 2010), we assume that financial shortfall and low profitability are the main determinants of exit. In particular,

1. Each shop whose owner is bankrupt must exit;
2. Each remaining shop exits for exogenous reasons with probability δ ;
3. Each shop can choose voluntarily to exit.

We assume that a shop will voluntarily exit with certainty if it cannot afford to pay for the coming week's fixed overhead cost, that is, if $A + CL < w(F - 1)$, where A is the shop owner's financial wealth, computed as in section 4.3 above, and CL is his credit limit.

The only other situation in which a shop will voluntarily exit is if it is unprofitable, in which case it will exit with an exogenous probability ϕ . In computing profitability, the shop owner takes into account the opportunity cost of his labor services and the interest-opportunity cost of maintaining the shop's fixed capital and inventory. Specifically, the shop owner will decide that the shop is unprofitable if one of the following two cases applies:

$$\begin{aligned} \text{(a)} \quad & A + P_f(I + S) \geq 0 \quad \text{and} \quad V \cdot W + P_f(I + S) > V \cdot \Pi^e; \\ \text{(b)} \quad & A + P_f(I + S) < 0 \quad \text{and} \quad V \cdot W > V \cdot \Pi^e + A, \end{aligned}$$

where I is the shop's inventory, S is his 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 shop owner's permanent income (expected profit from staying in business). Each case implies that the owner's tangible plus human wealth would increase as a result of exit, under the assumption that he could get a job that pays the wage W . In case (a), he would be able to repay his loan in full, although perhaps allowing some inventory and fixed capital to be seized, so in the event of exit, his tangible plus human wealth would go from $V \cdot \Pi^e + A$ to $V \cdot W + P_f(I + S) + A$. In case (b), he would be unable to repay his loan in full, so upon exit, his tangible plus human wealth would go from $V \cdot \Pi^e + A$ to $V \cdot W$.

Once a shop exits for any reason, all trading relationships (with both employees and customers) are dissolved and the shop owner has to repay his bank loan to the extent possible (bankrupt shops have already settled their loans). If the sum of shop owner’s money and deposits exceeds the bank loan, he repays the whole loan to the bank. Otherwise, the bank seizes capital as described in section 4.3 above, evaluating the capital at firesale prices. Banks with seized capital and former shop owners with non-seized capital now join the appropriate firesale queues if not already there. Upon exit, the former shop owner resets his permanent income to W .

4.9 Wage and price setting

In the final stage of weekly activities, each shop has an opportunity to update its posted wage and price. Wage and price changes are not communicated automatically to the shop’s employees and customers until the start of next period’s trading. In line with the standard theory of search and matching (Pissarides, 2000) and the empirical literature on the Phillips curve (e.g., Roberts, 1995), we recognize that labor market tightness and expected inflation are the key determinants of wage setting by firms.

Each shop first updates its sales target y^{trg} , setting it equal to the current week’s actual sales. It then proceeds to update its wage, but only if the last wage change was Δ weeks ago, where Δ is the length of the contract period (recall section 4.1 above). Given that the current week is indeed a wage-updating week for a shop, it adjusts its wage according to a firm-level expectations-augmented Phillips-type equation, as shown in appendix A.5. The size of this adjustment depends on the shop’s estimate of its excess demand for labor over the previous contract period and on the exogenous parameter β that captures its sensitivity to this estimate.

For price adjustments, we take into account evidence that applying a markup over cost is the main consideration in price setting (Fabiani et al., 2006). We also recognize that price changes often represent temporary adjustments to the markup, motivated mainly by market conditions (Bils and Klenow, 2004). Accordingly, a shop adjusts its price using its markup (selected at entry) whenever the after-tax marginal cost $w/(1 - \tau)$ changes. It will also raise its price, relative to its normal level, by the fixed factor δ_p whenever the inventory-to-sales ratio falls below a critical value $1/IS$, and it will reduce its price by the same factor whenever the inventory-to-sales ratio rises above a critical upper value IS (see appendix A.5 for details).

5 The workings of the model

5.1 A full-capacity stationary state

In our model economy, GDP is the sum of all goods produced.³³ As shown in the similar model of Howitt and Clower (2000), GDP will be maximized when there is exactly one shop trading each good (which economizes on fixed costs) and everyone either has an employer or is one. In this

³³The implicit assumption that all relative prices are unity is rationalized by the ex ante symmetry that we have imposed across goods.

case, GDP will equal full-capacity output Y^* , which is total labor input minus the total amount needed to cover fixed costs: $N - nF$. There exists at least one hypothetical stationary state in which GDP would equal its potential level if all the shocks associated with entry, exit, and match breakups were turned off (i.e., $\theta = \delta = 0$) and wages were adjusted every period ($\Delta = 1$). We use this hypothetical stationary state as the initial position of the economy in all our model simulations below, and a brief description of it helps to illustrate the workings of the model.

The stationary state is one in which all the potential gains from trade are realized. Each person is matched with one employer and two stores. There are n shops, one trading in each of the n goods. To preserve symmetry across goods, we suppose that each good is the primary consumption good for exactly one shop owner. Each shop begins every week with actual, potential, and target input all equal to $n - 2$, which is the number of people endowed with each type of labor, and with actual and target sales equal to inventory holdings equal to actual output: $n - 2 - F$. Note that there is no bank lending in this stationary state. In particular, banks in our model are important mainly for facilitating entry and supporting the ability of firms to absorb shocks, not for sustaining a stationary state that the system might approximate. Further details of this stationary state are spelled out in appendix A.6.

5.2 Entry, exit, and systemic performance

As we shall see below, in our median simulation, the economy is able to achieve on average about 93 percent of capacity output. 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 consumed by customers of the old shop.

These events that reduce GDP are constantly being offset to some degree by the entry of new shops, which are able to satisfy customers in markets where there had previously been no viable shop, and by the exit of shops, which 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. Nevertheless, 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 due to layoffs or financial problems, entry can be harmful (to the level of real GDP) 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, whether because of financial difficulties, a surfeit of inventories, or a too high markup, exit can also be harmful in cases where the incumbent was previously doing well, because it can cascade across markets and thus cause a cumulative loss of output (Howitt, 2006).

Banks have two influences on this cumulative process of shop failures. One reflects the stabilizing role of banks in ameliorating this process. Specifically, 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 at the entry stage each week.

The other influence is the familiar “lending channel” of the financial accelerator. Namely, when a shop fails, this may cause its bank to get into trouble or at least get close to being in trouble, because some of its loans may 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 will fail for lack of finance.

As we shall see, the stabilizing influence of banks appears to play a more important role in our model than the financial accelerator, especially when the system is far below full-capacity utilization.

6 Calibration and empirical validation

This section first describes how we calibrate our model and then discusses its ability to match various stylized facts observed in real-world macroeconomic data.³⁴

Our model has 33 parameters, which we have calibrated to U.S. economic data. The parameters are categorized as personal, shop, bank, and government parameters. They are listed in Table 1, along with their assigned values and their six-letter codes referenced in the sensitivity analysis of appendix E.

Our calibration of these parameters took place at three different levels. At the first level, one subset of parameter values was chosen to match their empirical counterparts in the U.S. data and/or the values used in previous studies. At the second level, the values of other parameters were chosen so as to be internally consistent with median outcomes across simulations. At the third level, the values of the remaining parameters, for which we could find no convenient empirical counterparts, were chosen to make their median outcomes across simulations (loosely) match certain properties of the U.S. data.

6.1 First level of calibration

6.1.1 Personal 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 the elasticity of demand facing each shop would be $1 + \varepsilon/2 = 4.5$ in a no-shock stationary state with all shops charging

³⁴To the best of our knowledge, ours is the only agent-based macro model that picks all of its parameter values by calibration to actual data. Nevertheless, unlike Dosi et al. (2010, 2013) and Delli Gatti et al. (2011), we have made no attempt to empirically validate our model against stylized facts in microeconomic data. In particular, the fixed costs that make each market in our model a natural monopoly, combined with the ex ante symmetry we have imposed across markets, make it unlikely that the firm size distribution in our model has the fat tails observed in the real world.

TABLE 1: The calibrated parameters of the model

<u>Personal parameters</u>			
ρ	Rate of time preference (annual)	0.04	TMPRFY
ε	Demand parameter	7.0	DMDELS
δ	Quit rate (weekly)	0.00075	BRKFRC
λ_p	Permanent income adjustment speed (weekly)	0.4	INCSPD
σ	Job search probability	0.5	SRCHPP
θ	Frequency of innovation (weekly)	100	INFREQ
<u>Shop parameters</u>			
Δ	Length of the contract period (in weeks)	48	CNTPRD
F	Fixed cost (weekly)	3.5	FXDCST
λ_I	Inventory adjustment speed (weekly)	0.16	INVSPD
β	Wage adjustment coefficient (annual)	0.3	WAGSPD
$\bar{\mu}$	Average percentage markup over variable costs	0.138	MARKUP
ϕ	Failure rate of unprofitable shops (weekly)	0.011	FAILRT
S	Setup cost	15	SETUPC
IS	Critical inventory-to-sales ratio	3.0	TRIGGR
δ_p	Size of price cut (old price/new price)	1.017	PDELTA
<u>Bank parameters</u>			
h	Loan-to-value ratio	0.5	LVALRT
s	Loan spread (annual)	0.0175	SPREAD
C_b	Cost of foreclosure	0.1	BKRCST
l	Slope of loan approval schedule	9	PSLOPE
m	Number of banks	5	NBANKS
<u>Government parameters</u>			
<i>Fiscal policy</i>			
b^*	Target debt-to-GDP ratio	0.33	BNDINC
λ_τ	Fiscal adjustment speed (annual)	0.054	FSCSPD
<i>Monetary policy</i>			
$\widehat{\lambda}_{\pi 0}$	Inflation autocorrelation coefficient (initial estimate)	0.29	INFCOR
$\widehat{\lambda}_{y 0}$	Output autocorrelation coefficient (initial estimate)	0.66	GAPCOR
γ_π	Inflation coefficient in Taylor rule	1.5	INFTYR
γ_y	Output gap coefficient in Taylor rule	0.5	GAPTYR
π^*	Target inflation rate (annual)	0.03	PITRGY
r_0^*	Initial target real interest rate (annual)	0.032	RRTTRG
$\ln(Y^*) - \tilde{y}_0$	Initial estimate of output gap	0.074	GAPTRG
η_r	Adjustment speed of evolving real rate target	0.0075	ADJRRT
T_{cb}	Number of years before government's learning begins	10	LRNLAG
<i>Bank regulation</i>			
κ	Required capital adequacy ratio	0.08	CARATE
s_d	Premium on central bank discount rate (annual)	0.005	DISCPR

the same price. This lies in the range of values typically found in New Keynesian DSGE models.³⁵ Nonetheless, the elasticity of demand faced by a shop outside this stationary state, when it has rivals selling the identical good, will be larger than 4.5, because raising the price may induce a loss of all demand from any customer that finds another shop during the matching process.³⁶

6.1.2 Shop parameters

Smets and Wouters (2007) estimate the average duration of wage contracts in the U.S. 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 the wage contract period Δ to 48 weeks, which in our model is one year.

Estimates of the degree of returns to scale in the U.S. economy vary from 0 to about 30 percent. It 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

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

If the economy was operating with a 6 percent average unemployment rate, then the typical shop would have $x = 0.94 \cdot (n - 2) = 45.12$, so by setting the fixed cost F equal to 3.5 we get a typical degree of returns to scale equal to 8.4 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 $1 - (1 - 0.16)^4 \approx 0.5$.

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 the same relation at the firm level if we assume that the 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.

6.1.3 Bank parameters

The value of the normal annual loan spread s was set equal to 0.0175, which, according to the Survey of Business Lending Terms conducted by the Federal Reserve, is the average spread between lending and deposit rates for all commercial and industrial loans during the period 1986–2008. We set the cost of foreclosure equal to 0.1, which falls in the 2 to 20 percent range suggested by Bris et al. (2006).

³⁵Demand elasticity is set equal to 3 in Midrigan (2011), 4 in Nakamura and Steinsson (2010), and 7 in Golosov and Lucas (2007). The benchmark calibration in Burstein and Hellwig (2008) yields an elasticity of 4.4.

³⁶This is another consequence of our decision to model goods markets as “customer markets,” following Phelps and Winter (1970).

6.1.4 Government parameters

Fiscal policy. The target debt-to-GDP ratio b^* was set equal to 0.33, because this is the average ratio of federal marketable debt to GDP in the U.S. 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 $\hat{\lambda}_{y 0}$ were taken from estimates of univariate AR(1) processes on inflation and on (linearly detrended) log per-capita GDP using annual data for the U.S. over the 1984–2006 period. The coefficients γ_π and γ_y are Taylor’s original specification.³⁷ We took the inflation target π^* to equal 3 percent, which is the average in the U.S. over the 1984–2006 period.

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

6.2 Second level of calibration

6.2.1 Government targets – finessing Wicksell

The two government policy targets – the target real interest rate r^* and log potential output \tilde{y} – are chosen adaptively by the government. We chose the initial value r_0^* to be 0.032, which is equal to the median outcome across simulations in our baseline calibration. Likewise, we chose the initial value \tilde{y}_0 to be 7.6.³⁸ Since government adaptation was relatively quick and its learning period was quite long, our results were not sensitive to the choice of these initial values.

Note that our procedure of having the government 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 cumulative inflation.

6.2.2 Markups and the Lucas critique

In early trials of the model, we assumed that all shops applied the same markup, but we found that the results from 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

³⁷Parameters η_r and T_{cb} are introduced in appendix A.3, which describes all of the central bank’s learning and forecasting rules.

³⁸This corresponds to the logarithmic output gap (between capacity and actual GDP) of 0.074 reported in column 1 of Table 3.

TABLE 2: U.S. data vs. median outcomes in the model

Indicator	Data	Model
Inflation	3.0	2.9
Real interest rate	1.8	3.2
Unemployment rate	6.1	5.9
Unemployment duration	14	10
Volatility of output gap	2.0 to 3.2	2.8
Volatility of inflation	1.3	0.68
Autocorrelation of gap	20 to 76	43
Autocorrelation of inflation	16	18
Average markup	10 to 20	14
Exit rate	46	40
Job loss rate	0.69	0.59
Price-change frequency	4.0	4.1
Annual bank failure rate	0.51	0.50

are drawn in the same way that we decided on the initial values of the government targets – by internal consistency. Specifically, we set $\bar{\mu}$ equal to 0.138 to match the median markup across simulation runs in our baseline calibration.

6.3 Third level of calibration

This leaves 13 parameters still to be determined; namely, ϕ , S , IS, δ_p , δ , λ_p , σ , θ , l , η_r , T_{cb} , m , and h . These were chosen by searching (manually) for parameter values that would make the cross-run median values of thirteen different indicators in our model approximate their counterparts in the U.S. data. More specifically, we ran 10,000 simulations of 60 years. Each simulation run started near the no-shock stationary state and continued for 20 years before we started calculating the average value of each indicator variable across the remaining 40 years of that simulation. For each variable, we then computed the median across all simulation runs of these cross-year averages. The thirteen indicator variables are listed in Table 2, along with their actual values from U.S. data and their median values in our fully calibrated model.

All numbers are expressed in percentage points, except for unemployment duration, which is expressed in weeks, and price-change frequency, which is expressed in number of changes per year. The actual values from U.S. data for these thirteen indicators were determined as follows. The real annual interest rate is computed as the difference between the annual interest rate on 3-month T-bills (monthly data from the Federal Reserve) and the CPI inflation rate (monthly data from the U.S. Bureau of Labor Statistics), averaged over the 1984–2006 time period. Unemployment rate and duration are the averages over the period from 1969 to 2009. The volatility of the output gap is the standard deviation of linearly detrended and HP-filtered log per-capita GDP, and the volatility of inflation is the standard deviation of annual U.S. CPI inflation over the 1969–2009 period. The 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 et al. (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 percent of people working at a given date have been working continuously for the same employer for the past year.³⁹ Bils and Klenow (2004) find an average price-change frequency of 16 weeks, which in our model would imply an average annual price-change frequency of 3. Shops in our model, however, 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. We therefore aimed to match a price-change frequency of 4. Finally, according to the FDIC (Historical Statistics on Banking), the average commercial bank failure rate was about 0.51 percent per year over the period from 1984 to 2006.

As Table 2 shows, we were mostly successful in mimicking the actual data with our 13 parameters. Our model, however, underpredicts the duration of unemployment, the volatility of inflation, the job loss rate, and the exit rate.

6.4 Empirical validation

Before conducting experiments using our calibrated model, we follow what has now become standard procedure in the literature on agent-based macroeconomic modeling, that is, validating that the model is able to qualitatively match various stylized facts known to hold in real-world macroeconomic data (e.g., Dosi et al., 2010, 2013; Delli Gatti et al., 2011; Dawid et al., 2012, 2014). To this effect, we perform several tests, exploiting intertemporal variations in the pooled annual time series of different macroeconomic indicators generated by repeated simulations of our calibrated model. Specifically, for our empirical validation exercises, we employ a sample of 2,000 different 60-year simulations, and similar to the methodology in our third level of calibration, we restrict attention to the final 40 years' worth of data for each simulation.

In the first step of our empirical validation, we follow Delli Gatti et al. (2011) in documenting the existence of a core “triad” of real-world macroeconomic relationships – namely, the Phillips curve, Okun’s law, and the Beveridge curve – in our simulated data. Figure 1 depicts each of these relationships across annual observations in a sample of 25 simulations.⁴⁰ Consistently with the qualitative feature of each of these core macroeconomic patterns in nature, we find highly statistically significant negative relationships between (i) the unemployment rate and the inflation rate (panel a), as per the canonical Phillips curve; (ii) the unemployment rate and the change in the inflation rate, conditional on the lagged values of these variables (panel b), as per the variant of the Phillips curve estimated by Stock and Watson (1999); (iii) the unemployment rate and the

³⁹That is, $(1 - \alpha)^{48} = 0.718$ if $\alpha = 0.00688$.

⁴⁰Our decision to employ a sample of 25 simulations was made in order to maximize the visual clarity of the scatter plots shown in Figure 1. Employing information from a larger sample of 2,000 simulations yields similar but, obviously, even more statistically precise estimates of the depicted relationships.

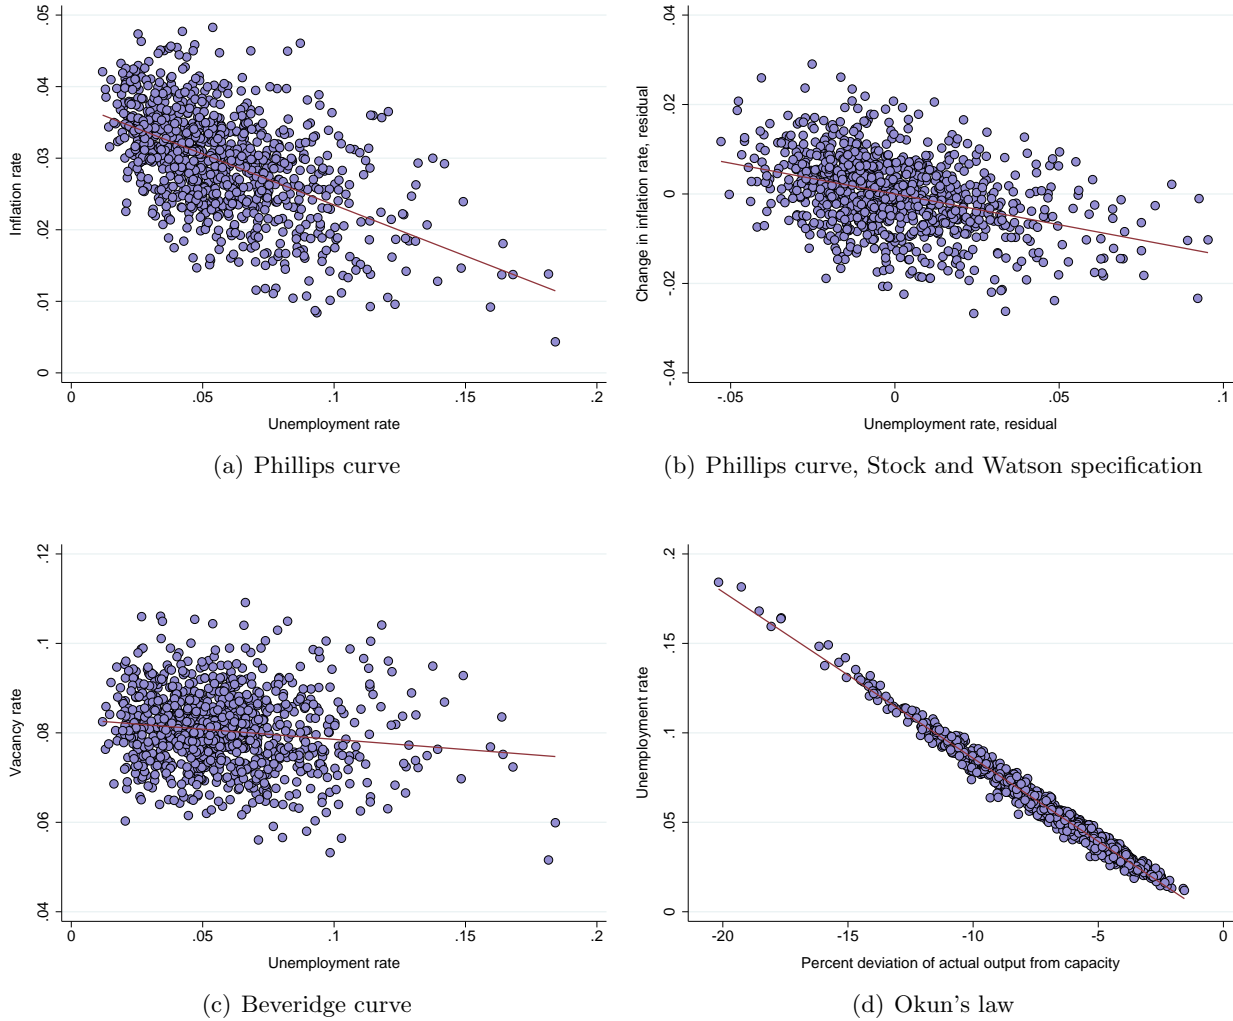


FIGURE 1: Macroeconomic relationships generated by the model

job vacancy rate (panel c), as per the Beveridge curve; and (iv) the percentage deviation of real GDP from capacity and the unemployment rate (panel d), as per Okun's law.⁴¹

For our second set of empirical validation results, we follow Dosi et al. (2010, 2013) and Dawid et al. (2012, 2014) to examine the extent to which our calibrated model is able to qualitatively replicate real-world observations of the co-movement of various macroeconomic indicators with real GDP over the business cycle, as documented for the postwar U.S. by Stock and Watson (1999). Exploiting pooled annual time series data from across our simulations, Table B.1 in appendix B presents the cross-correlation structure between output and each of 12 different macroeconomic variables (not counting output itself), including leads and lags of the latter of up to 4 years, along with the associated Monte Carlo standard errors in parentheses. As is evident from the table, in line with the results of Stock and Watson (1999), we find that on average in our simulations, consumption,

⁴¹We have additionally found, consistently with an alternative version of Okun's law, a statistically significant negative relationship between the change in the unemployment rate and the growth rate of real GDP.

investment, change in inventories, employment, the job vacancy rate, average labor productivity, the inflation rate, and the nominal interest rate are all procyclical, whereas the unemployment rate is countercyclical.⁴² Further, consistently with the empirical findings of Bils (1987) and Rotemberg and Woodford (1999), markups on average move countercyclically with output in our simulations. We also find that the government debt-to-GDP ratio is on average countercyclical in our simulations, in line with the pattern observed in the data for the U.S. economy during the 1966–2005 time horizon.⁴³

In the final step of our empirical validation, we follow Dosi et al. (2010) in confirming that the distribution of the growth rate of output from our simulations displays tails that are fatter than what one would expect from normally distributed shocks, in tune with the empirical evidence documented for OECD nations by Fagiolo et al. (2008). Specifically, employing pooled annual time series data on the growth rate of output from 500 simulations and exploiting the same maximum-likelihood procedure as in Fagiolo et al. (2008) and Dosi et al. (2010) to fit an exponential-power (EP) distribution to our simulated data, we recover a point estimate of 1.277 (with a standard error of 0.018) for the “shape” parameter of our fitted EP distribution, indicating the presence of tails that are significantly fatter than the Gaussian benchmark (and are, in fact, closer to those of the Laplace distribution). As in Dosi et al. (2010), we have also confirmed the presence of “super normal” tails in the growth-rate distributions of consumption and investment in the data generated by our simulations.⁴⁴

7 The main results

Having empirically validated our calibrated model, we next perform our main experiments. We conduct a total of 10,000 simulation runs, each allowed to continue for $T = 2880$ weeks (60 years). All runs are initiated near the hypothetical stationary state of the economy described in section 5.1 above. We do not start tabulating results from any run until 20 years have passed, in order to provide the system with sufficient time to possibly settle into a stochastic steady state. Figure 2 shows the time series of the median values across all 10,000 runs of the output gap (measured as the difference between log capacity GDP and log actual GDP on the left-hand-side vertical axis), the real interest rate, and the inflation rate (both measured on the right-hand-side axis). It appears that our 20-year adjustment period was indeed sufficiently long for the cross-run median values of

⁴²In addition, for most of these indicators, the lead/lag cross-correlation structure with output is, by and large, qualitatively similar to the results of Stock and Watson (1999).

⁴³Specifically, exploiting the annual time series of real GDP (in chained 2009 dollars) and total public debt as a percentage of GDP for the U.S. economy, 1966–2005, obtained from the FRED Economic Data website maintained by the Federal Reserve Bank of St. Louis, we find a contemporaneous correlation coefficient of -0.053 between HP-filtered log real GDP and the total public debt-to-GDP ratio.

⁴⁴In particular, in the case of consumption growth, we obtain an estimated “shape” parameter of 1.336 for our fitted EP distribution, whereas in the case of investment growth, this parameter estimate is 0.342. Thus, consistently with the model of Dosi et al. (2010), the distribution of investment growth in our model displays not just “super normal” but also “super Laplace” tails.




FIGURE 2: Time series of cross-run medians

these indicators to become more or less stable, except for the slight downward trend in the cross-run median output gap over the first half of the final 40-year period.

7.1 The corridor effect: normal and worst-case scenarios

Most of the simulation runs produce time paths that exhibit strong homeostatic tendencies. As mentioned in our introductory remarks, however, there are also a few “pathological” runs in which the market makers appear to have lost control of the system. This is an important nonlinear feature of the model, which seems to behave in a qualitatively different manner in “worst-case scenarios” as compared to “normal” ones.⁴⁵ To convey some idea of this qualitative difference, Figure 3 depicts the time paths of three major macroeconomic variables in one of the many normal runs of the system. A randomly chosen simulation would almost certainly exhibit similar characteristics. There are times when the output gap rises for a few years, but these times are soon followed by recovery, and the output gap rarely exceeds 0.12.

In contrast to these normal runs, Figure 4 illustrates what happens in some of the worst runs. Although there are times when the output gap seems to be falling back to a normal level, beyond some point, it spirals out of control. Indeed, there were 20 runs out of 10,000 in which the output gap

⁴⁵As we shall see, we define worst-case scenarios in terms of a 40-year average output gap. As such, these situations do not correspond to standard recessions, which typically last only a few years. We interpret them instead as periods of prolonged depression, recovery from which is slow and might even require structural and/or policy changes that we have not modeled.




FIGURE 3: A normal run




FIGURE 4: A collapse




FIGURE 5: Macroeconomic performance by decile

became infinite ($GDP = 0$) in finite time.⁴⁶ Such behavior could not be produced by a linear macro model. These collapses indicate the presence of something like Leijonhufvud’s (1973) “corridor,” that is, the economy is capable of absorbing shocks up to some limit but not beyond.

The notion that something qualitatively different happens when the economy substantially deteriorates becomes more evident when we group the runs into deciles from the distribution of the average value of the output gap across the final 40 years of a run. Figures C.1 and C.2 in appendix C show the time series of the cross-run average and variance of the output gap, respectively, within each decile. The time series for the worst (tenth) decile is a clear outlier in both cases. In comparison to the difference between the ninth and tenth deciles, the differences among the time series for all the other deciles appear to be quite small.

Figure 5 provides further evidence on how abruptly macroeconomic performance declines when going from the ninth to the tenth decile of runs. It shows the within-decile average behavior of six indicators, relative to the worst decile. All six measures show a rather smooth deterioration over the first nine deciles and then a much sharper decline when the worst decile is reached. For some measures, the decline from the ninth to the tenth decile is ten times the cumulative deterioration from the first to the ninth.

To some extent, this pattern is what one would expect from a more standard (but necessarily nonlinear) New Keynesian DSGE model that respected the zero lower bound on nominal interest

⁴⁶We excluded these runs from the computation of the cross-run median paths shown in Figure 2.

rates.⁴⁷ Indeed, the fraction of times in which the zero lower bound is hit rises gradually from 0.008 in the first decile to 0.034 in the ninth and then jumps drastically to 0.148 in the tenth decile.

There is, however, more going on. Figure D.1 in appendix D reports the results from rerunning the 10,000 simulations of our basic calibration, allowing the nominal interest rate to be determined by the Taylor rule even when this rule would make it negative. All other aspects of 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 nevertheless shows that there is more to the results in Figure 5 than just the zero lower bound. Specifically, Figure D.1 shows that although the deterioration in macro performance is less severe than before when going from the ninth to the tenth decile of runs, there is still a characteristic kink in each of these patterns (except for the job loss rate), much like before.

7.2 The role of banks: financial accelerator or stabilizer?

It is commonly believed that financial intermediaries in general, and banks in particular, exacerbate cyclical fluctuations through a financial accelerator mechanism. This section shows that banks can actually moderate disastrous scenarios, thus making the economy more stable than it would be in their absence. As Leijonhufvud has argued, a corridor effect is inherent in any control mechanism that does not have the capacity to counteract every possible deviation from a coordinated state of affairs. Indeed, as shown in the previous section, the presence of banks does not eliminate the corridor effect in our model. As we will see below, however, it does make the effect less pronounced: banks improve the economy’s performance, especially in the worst decile. In this sense, banks in our model provide a financial stabilizer whose effect dominates that of any financial accelerator that might be present.

In order to quantify the effect of banks on the economy’s performance, we conducted an experiment in which we shut down the banking sector and reran our simulations. To do this, we simply altered bank behavior in the financial market stage of each week so that the banks always imposed a credit limit of zero on all customers, thereby turning the banks into mere conduits for the private holding of government debt. In this case, a bank’s equity is equivalent 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 from this experiment are reported in Table 3, where we differentiate between normal and worst-case scenarios.⁴⁸ The former are represented by cross-run medians of the 40-year average values of various performance indicators, while the latter are represented by cross-run averages of the 40-year average values of those indicators in the worst (tenth) decile of runs (in terms of the output gap). As is evident from the first two columns of Table 3, with the exception of inflation, all indicators show a marked deterioration in median performance in the case where banks are shut

⁴⁷See, e.g., Eggertsson and Woodford (2003) or Levin et al. (2010).

⁴⁸The notion of “risky” banks is introduced in the next section.

TABLE 3: Macroeconomic outcomes with and without banks

Indicator	Median			Worst decile		
	(1)	(2)	(3)	(4)	(5)	(6)
	Banks	No banks	Risky banks	Banks	No banks	Risky banks
Inflation	2.9	2.5	3.0	2.6	1.3	2.9
Output gap	7.4	13	7.6	15	48	10
Unemp. rate	5.9	11	5.8	11	29	7.8
Unemp. duration	10	18	8.7	16	26	10
Job loss rate	0.59	0.68	0.69	0.77	2.2	0.78
Vol. of output gap	2.8	6.2	2.5	7.9	27	4.3
Vol. of inflation	0.68	1.1	0.60	1.1	1.8	0.73
Ann. bank fail. rate	0.5	0	2.5	0.84	0	3.1
Frac. of banks in troub.	4.0	0	1.0	16	0	1.7

down.⁴⁹ Columns 4 and 5 of Table 3 reveal that banks make an even bigger difference in the worst decile of runs. For almost all indicators, the difference in worst-decile averages between the two cases is considerably larger than the corresponding difference in median performance reported in the first two columns. For example, banks reduce the median unemployment rate from 11 to 5.9 percent, but they reduce the worst-decile average unemployment rate from 29 to 11 percent.

The reason why banks make an even bigger difference in worst-decile scenarios is that in those runs their help in sustaining shops facing financial difficulty and in facilitating the entry of new shops to replace the failed ones is most needed to avert a major economic downturn. 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 the shop’s former employees of their source of income and, thus, increasing the likelihood that other shops will fail for lack 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 workers to cover its fixed cost. The process continues until the former employees of failed shops find new employers, and this will often require new shops to enter. Sustaining shops at risk of failure and replacing shops that fail make it less likely that this cumulative process will cause the economy to spin out of control.

Figure 6 shows that banks do indeed play an important role in mitigating the cumulative process of shop failures. This figure records, for the cases with and without banks, the results from an experiment in which we forced a shop to fail. Specifically, we forced the failure at the beginning of year 20 in the simulation (i.e., at the end of the adjustment period) and repeated this simulation 10,000 times. For a counterfactual, we performed the same 10,000 simulations but without the shock that forces the shop failure. The left panel of Figure 6 shows the “trimmed” average difference in the output gap between the shocked and counterfactual simulations on a monthly basis for two years (24 months) following the shock, where the solid and dashed lines

⁴⁹Inflation is lower because the fact that firms always factor in the same target inflation rate π^* when adjusting wages implies that there is a long-run trade-off between output and inflation.




FIGURE 6: Average response of the output gap (left) and entry (right) to a shop failure

correspond to the cases with and without banks, respectively.⁵⁰ As the figure demonstrates, in either case there was a hump-shaped impulse response of the output gap, but the response is much stronger in the economy without banks. Moreover, the right panel of Figure 6 shows, as expected, that the entry of new shops in response to a shop failure rose by much more in the presence of banks than in their absence.

Another indication of the mechanism at work is given by panels (a) and (b) of Figure 7, which show by decile of simulation performance, the behavior of the number of shops and the fraction of entrepreneurs that fail the financial viability test and, thus, let the entry opportunity lapse. Without banks, moving from the best to the worst decile, the number of shops declines and the incidence of financially non-viable business plans increases, both of which clearly amplify the deterioration in macroeconomic conditions. With banks, however, these variables are roughly constant from the first through the ninth decile of runs, thus averting the amplification that occurs in the absence of a banking system.

The fact that the economy is less volatile and performs better with banks than without them shows that banks, while often depicted as creating a deviation-amplifying financial accelerator, actually provide a deviation-counteracting “financial stabilizer” in our model. Nevertheless, as can be seen in the top two panels of Figure 7, there is a limit to this stabilizer: in the worst decile of runs, banks can no longer avert the fall in the number of shops and the rise in the share of potential entrants failing the financial viability test. This apparent weakening of the financial stabilizer in the worst decile of runs could be interpreted as a kind of financial accelerator working through a bank-lending channel. That is, when shops fail, banks that have lent to these shops are more likely to become troubled, and troubled banks cut back on their lending, thus raising the incidence of shop failures due to this particular form of “credit rationing.” Consistently with this notion, in panels (a) and (b) of Figure 7, we observe very abrupt trend breaks in the case with banks, whereas the trajectories are much smoother in their absence. Moreover, in line with this mechanism, panel

⁵⁰That is, in order to avoid the noisiness of runs where the simulations spun out of control, we calculated, for each month, the average between the 25th and 75th percentile of all these differences.

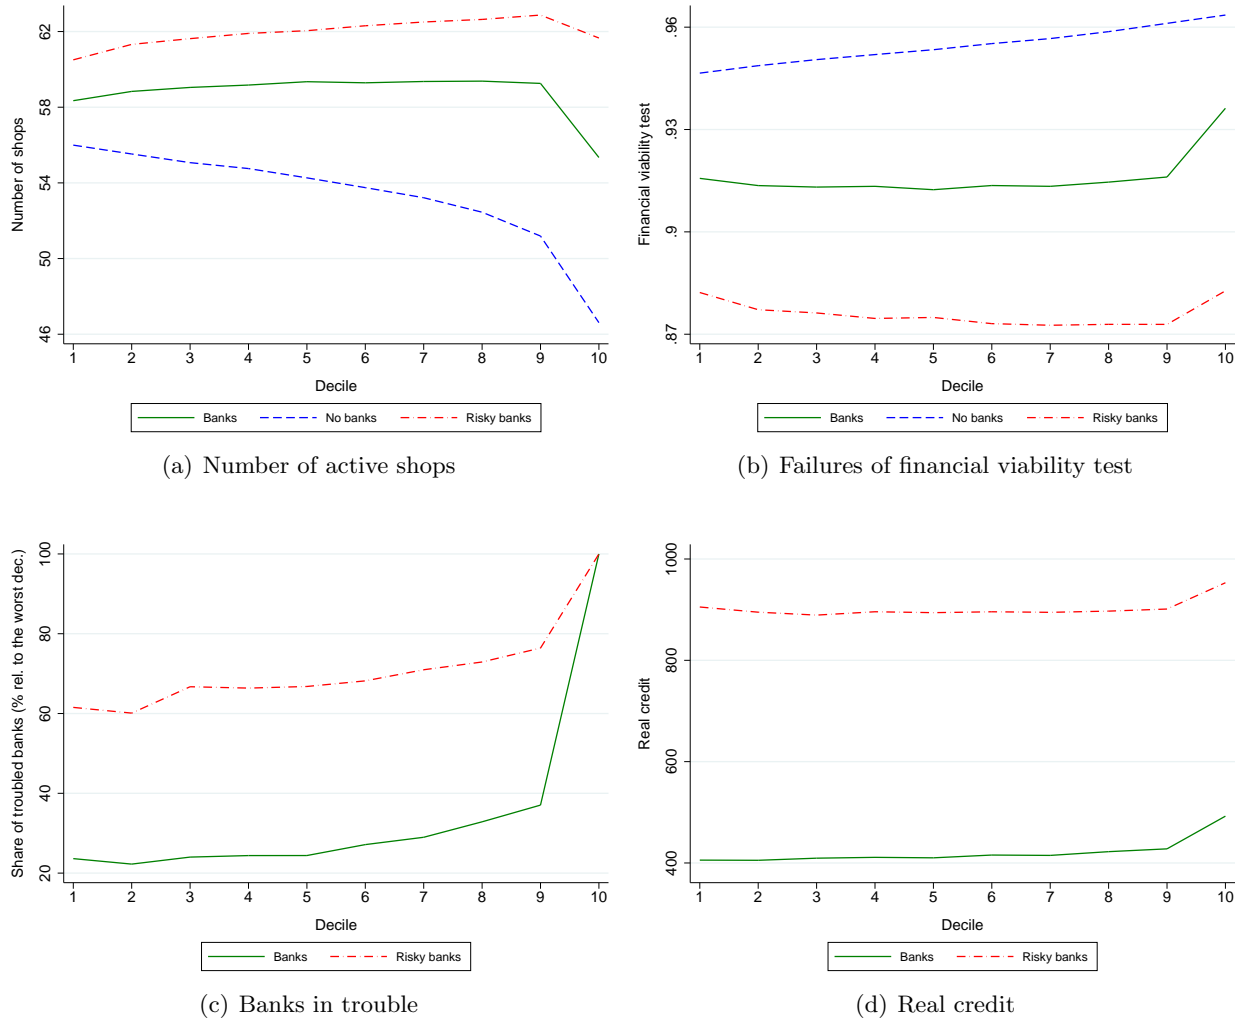


FIGURE 7: Bank lending and systemic performance

(c) shows that the fraction of troubled banks is indeed markedly higher in the worst decile of runs. Yet, the fact that the economy still performs better with banks than without them, especially in the worst decile, means that the financial accelerator is not as strong as the financial stabilizer. Indeed, despite the spike in the share of troubled banks in the worst decile of runs, panel (d) of Figure 7 shows a noticeable increase in the total volume of real credit in these runs, implying more intense lending activity on the part of non-troubled banks. In effect, our model shows that the financial accelerator, instead of amplifying the effects of shocks beyond what would be observed in the absence of banks, merely reduces the extent to which banks dampen the shocks relative to the case without banks.

7.3 Risky banks improve macroeconomic performance in worst-case scenarios

To further our understanding of how banks affect macroeconomic performance, we consider an alternative case in which banks are more “risky” by micro-prudential standards than the “safe” banks of our baseline case. Specifically, first, instead of constraining its customers with a loan-to-value ratio of 0.5, each bank in the risky case allows a higher loan-to-value ratio of 0.9, and second, the required capital adequacy ratio in the risky case is just 0.02, in contrast to the more credit-constraining value of 0.08 used in the safe case.

As shown in columns 1 and 3 of Table 3 above, median performance in the risky case is roughly comparable to that in our baseline safe case. Median unemployment rate is about the same, as is the median output gap, but volatility and unemployment duration are lower, while the job loss rate is higher. As expected, the incidence of bank failures is much higher in the risky case, whereas the fraction of banks in trouble is much lower, because being in trouble is defined more loosely for a risky bank. Columns 4 and 6 of Table 3, however, indicate that worst-decile performance is considerably better with risky rather than safe banks. This is because banks mitigate the deterioration in performance by sustaining shops at risk of failure and by financing the entry of new shops to replace those that have failed. Risky banks perform these roles better than safe banks, because their higher loan-to-value ratio makes them automatically willing to lend more when they are not in trouble. In addition, the lower capital requirement means that risky banks are less often in trouble and, hence, more able to lend.

These insights are corroborated by Figure 6 above. The left panel shows that the hump-shaped impulse response of the output gap to a shop failure is more moderate with risky banks than with safe ones, while the right panel demonstrates that the increase in entry following the shock is more pronounced with risky banks. On the other hand, Figure 8 shows that a troubled-bank shock has a larger effect on the output gap in the risky case, which is hardly surprising because the bank that falls into trouble cuts its lending by a larger amount on average when risky than when safe.⁵¹ But this amplification of the troubled-bank shock is mitigated by the fact that banks are less often in trouble in the risky case.

Accordingly, the financial accelerator, which weakens the financial stabilizer and, thus, contributes to the corridor effect in our model, should be less pronounced in the risky case, because there will be fewer shop failures (as a result of more freely available finance) and a lower incidence of banks becoming troubled and, hence, cutting off new borrowing. Indeed, the different panels of Figure 7 show that the accelerator was less at work in the worst decile in the risky case than in the safe case, as suggested by the more moderate trend breaks in the patterns under risky banks. Figure 9 further demonstrates that although there is still a sharp deterioration of macroeconomic performance from the ninth to the tenth decile of runs in the risky case, it is much less severe than in the safe case depicted earlier in Figure 5.

⁵¹For the troubled-bank shock, instead of forcing a shop to fail at the end of year 20, we imposed the same sanctions on a bank (for one year) as if it were troubled; that is, the bank could not initiate any new loans or pay its owner a dividend for one year.

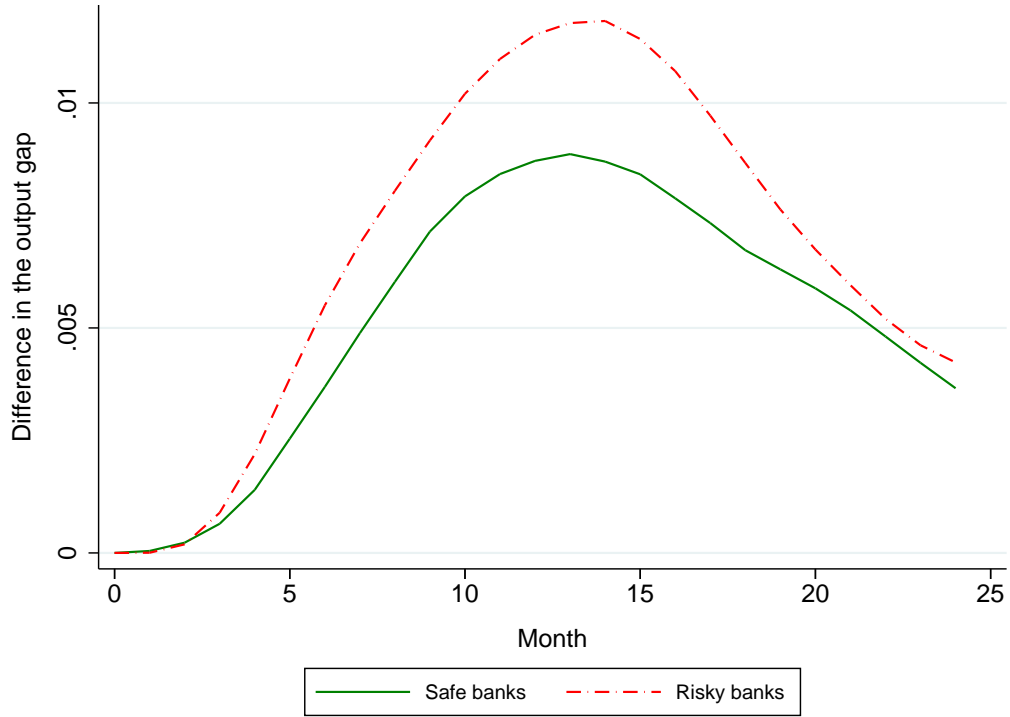


FIGURE 8: Average response to a troubled-bank shock: safe vs. risky banks




FIGURE 9: Macroeconomic performance by decile with risky banks




FIGURE 10: Effect of the loan-to-value ratio on the output gap

Another way to see the difference between safe and risky banks is to conduct the following experiment. We start with the safe case and increase the loan-to-value ratio gradually from 0.5 to the riskier value of 0.9 by increments of 0.1, while maintaining the required capital adequacy ratio of 0.08. For each of the five indicated values of the loan-to-value ratio, we simulated the model 10,000 times. The same set of simulations was then repeated for the risky case, that is, with the less credit-constraining required capital adequacy ratio of 0.02. Figure 10 shows that the increase in the loan-to-value ratio resulted in a slight deterioration of median performance (as measured by the output gap) but a significant improvement in worst-decile average performance, in both cases. Again, this suggests that willingness to lend makes relatively little difference in normal scenarios but helps substantially in the worst of runs.

In a related experiment, we gradually increase the required capital adequacy ratio from 0.02 to the more credit-constraining value of 0.16 by increments of 0.02, while maintaining the loan-to-value ratio at either 0.5 (for the safe case) or at 0.9 (for the risky case). Figure 11 shows that although a higher required capital adequacy ratio has little effect on median performance, worst-decile performance is significantly degraded by this tightening of capital requirements.⁵²

⁵²We also conducted an experiment in which the required capital adequacy ratio was allowed to change procyclically. In this experiment, a higher degree of procyclicality had almost no effect on macroeconomic performance in the risky case, but it did result in significantly better worst-decile performance in the safe case.




FIGURE 11: Effect of the required capital adequacy ratio on the output gap

8 Alternative assumptions regarding bank behavior

In appendix E, we show that our main results with respect to macroeconomic performance under safe and risky banks in normal and worst-case scenarios are robust to 25% variations in almost all of the model’s parameters. The present section demonstrates that our main results are also robust to alternative, and arguably more realistic, assumptions regarding bank behavior.

8.1 Making bank failure costly

The failure of a bank in our model is a salutary event. A bank that is about to fail is almost certainly in trouble, not making any new loans, and is subsequently replaced by a new bank that can resume making new loans. Thus, one reason for our finding that risky banks improve macroeconomic performance in the worst decile is that their higher failure rate does not impinge negatively on the real economy.

One argument for believing that bank failure does in fact impair macroeconomic performance is that it may result in the loss of important information that the failing bank had regarding its customers. In particular, a new bank may require some time to get to know the customer base, and in the meantime, it will not be as willing as the old bank would have been to extend credit to customers that look like poor risks on the basis of the hard information available to an outsider.

To explore this possibility, we altered the model along the following lines. In each market, we designated N^c of the $n - 2$ potential shop owners as “crooks.” We assumed that any bank that

has been in business for at least 48 weeks knows which of the agents in its sector are crooks and which are not, but that a new bank cannot distinguish crooks from other customers for the first 48 weeks. The new bank, however, does know the fraction of total loan applications that are coming from crooks. Thus, when a bank fails, the replacement bank will charge a higher loan interest rate in order to compensate for the predictable losses generated by fraudulent loan applications.

More specifically, during the entry phase, if a crook becomes an entrepreneur, he asks his bank for a line of credit. If the loan application is granted, he takes out a loan equal to the monetary cost S_N of his fixed capital minus his financial wealth, just like any other entrepreneur with the same financial wealth and no legacy capital. Unobserved by his bank, however, the crook will not open a shop (and we can assume that crooks do not even know how to open a shop). Instead, he deposits the loan in his bank account and never repays any amount of the loan. The following week, the bank will realize that he is a crook, but we assume that there is nothing the bank can do at this stage other than write the loan down to zero in its books.

We also assume that crooks can never become bank owners. Given that crooks and other non-bank-owning, non-shop-owning people have the same probability of becoming an entrepreneur, the fraction of loan applications coming from crooks will be $p^c = \frac{N^c n}{N - \Omega - 5}$, where Ω is the number of shops at the start of the week, assumed to be known by all banks. We further assume that a new bank (i.e., one that has been in business for less than 48 weeks) sets a loan rate i_L such that its expected return from a new loan is equal to $i_w + s/48$, the loan rate of an experienced bank. Since the bank loses interest plus principal from each loan to a crook, this means that

$$i_L = \frac{1 + i_w + s/48}{1 - p^c} - 1.$$

At the beginning of the goods and labor market trading stage of each week, we assume that for each legitimate shop owner that applies for a loan to finance a firesale purchase, there is a probability $\frac{p^c}{1 - p^c}$ that a crook will apply for a loan of the same size. Again, the new bank cannot distinguish these fraudulent loan applications from legitimate ones, although it does know that the fraction p^c of such loans are being granted to crooks. Once the 48-week initiation period is over, the bank simply rejects all loan applications coming from crooks.

In this variant of our model, the presence of crooks can make bank failure impair the real economy, because there will now be a very high interest rate to be paid by new entrants and by established firms in 20 percent of the economy for a year following the exit of a bank. This could potentially reverse our ranking of safe and risky banks, in terms of worst-decile macro performance, since risky banks fail more often than safe banks do.

Our results, however, indicate that it would require a lot of crooks in order for this ranking to be reversed. Table 4 shows the results for worst-decile average performance, under the safe and risky cases, from simulating this variant of our model with 4 crooks in each market. As is evident from a comparison with Table 3, under both the safe and risky cases, the economy's worst-decile average performance deteriorates further by almost all measures relative to the baseline case without crooks.

TABLE 4: Worst-decile average outcomes with “crooks” in the model

Indicator	Safe banks	Risky banks
Inflation	2.5	2.7
Output gap	16	14
Unemployment rate	12	11
Unemployment duration	18	14
Job loss rate	0.74	0.79
Volatility of output gap	8.3	7.4
Volatility of inflation	1.1	0.94
Annual bank failure rate	2.2	4.8
Fraction of banks in trouble	18	2.8

In these simulations, since the number of shops Ω at the start of each week is usually about 60, the fraction p^c of loans coming from crooks is more than 8.5 percent, meaning that the loan rate of a new bank, instead of being about 7 percent per annum, is over 9 percent per week! Nevertheless, it appears that macro performance is better under the risky case, which involves much more bank failure than does the safe case. Indeed, we found that the number of crooks had to be raised to 6 per market in order to make the economy’s worst-decile results better with safe than with risky banks.

8.2 Procyclical loan-to-value ratio

As emphasized earlier, we have assumed a highly stylized banking sector where, among other things, banks that are not in trouble do not endogenously alter their rules for approving collateralized loan applications, even if the economy is in recession and the customer’s financial condition is worsening. These simplifications do indeed produce some counterfactual results. In particular, in our baseline calibration, the correlation coefficient ρ_{YL} of nominal loans with nominal GDP (both differenced and logged) is negative in almost all runs, with a median value across runs of -0.341 , whereas in the data, $\rho_{YL} = 0.408$.⁵³ This section checks the sensitivity of our main results to an alteration of our model that brings the aforementioned aspects of bank behavior closer to reality. In particular, we now suppose that the loan-to-value ratio h is an increasing function of real GDP, with an adjustment speed calibrated to make ρ_{YL} approximately equal to 0.4.⁵⁴

⁵³This correlation was computed using U.S. annual seasonally-adjusted data on nominal GDP and commercial and industrial loans from all commercial banks (logged and first-differenced) over the 40-year period, 1961–2000. The series were downloaded from the FRED Economic Data website maintained by the Federal Reserve Banks of St. Louis.

⁵⁴Our formulation employs a “squasher” to keep h positive and to keep positive and negative adjustments symmetric. Specifically, we assume that

$$h = \bar{h} + \eta_h (y - \tilde{y}) \cdot f(y, h), \quad f(y, h) \equiv \frac{\bar{h}}{\sqrt{\eta_h^2 (y - \tilde{y})^2 + \bar{h}^2}},$$

where \bar{h} is the average loan-to-value ratio, still equal to 0.5 in the baseline (safe) case and 0.9 in the risky case, and η_h is the adjustment speed. We set η_h equal to 3, resulting in a median ρ_{YL} equal to 0.413 in the baseline case. Each bank makes this adjustment once per month.




FIGURE 12: Credit-GDP correlogram: model vs. actual data

As with our baseline calibration, prior to experimenting with this variant of our model, we validate its ability to qualitatively match certain real-world empirical patterns, paying particular attention this time around to the financial variables in our model while noting that all the model properties that were validated earlier for the baseline version remain qualitatively intact. Exploiting the pooled annual time series of real GDP and the supply of real credit across repeated simulations, Figure 12 depicts the cross-correlation structure between these two variables, documenting that the pattern in our simulated data accords well with that observed in the data for the U.S. economy over the 1961–2000 time period.⁵⁵ Furthermore, Table B.2 in appendix B reveals the cross-correlation structure in our simulated data between real GDP and each of four additional variables, including the real money supply, and the hazard rates of (i) firm exit due to bankruptcy, (ii) failing the financial viability test at market entry, and (iii) becoming a troubled bank. Consistently with the documented movement of the real money supply over the business cycle in the U.S. (Stock and Watson, 1999), Table B.2 shows that on average this variable is indeed procyclical in our simulations. Moreover, in line with both casual observation and priors based on the role of banks in our model, all three aforementioned hazard rates (particularly, the ones associated with market entry and exit) are on average countercyclical.

Generally speaking, allowing the loan-to-value ratio to behave procyclically in our model did little to change our main results from section 7. First, as Figure 13 shows, the system’s performance still exhibits a sharp deterioration in the worst decile of runs. In fact, for most macroeconomic

⁵⁵The cross-correlations in the U.S. data are produced using the annual time series of real GDP (in chained 2009 dollars) and the real value of commercial and industrial loans (obtained by dividing the corresponding nominal series by the GDP deflator) over the 40-year period, 1961–2000. The reported correlations are calculated for the HP-filtered logged series of real GDP and real credit.





FIGURE 13: Macroeconomic performance by decile with procyclical loan-to-value ratio

indicators, it is even sharper than in the baseline case with a constant loan-to-value ratio (see Figure 5). This is what we would have expected: with procyclical lending behavior, banks reduce lending precisely when loans are needed the most. This interpretation of our results is corroborated by Figure 14 which shows that the number of shops is lower, the fraction of entrepreneurs failing the financial viability test is higher, and the total volume of real credit is lower under a procyclical loan-to-value ratio than in the baseline case, and that these differences increase with the decile of runs, becoming most pronounced in the worst (tenth) decile. Interestingly, however, the share of troubled banks under a procyclical loan-to-value ratio is largely similar to the baseline case across all deciles.


Finally, as Table 5 indicates, under a procyclical loan-to-value ratio, the safe case is still much worse than the risky case in terms of worst-decile macroeconomic performance, and the two cases are still similar in terms of median performance. As mentioned earlier, we expected that making h procyclical would worsen economic performance. A comparison with Table 3 shows that this did indeed happen in terms of both median and worst-decile performance, and under both cases, although median performance changed only slightly in the risky banks case.

9 Conclusion


In this paper, we developed an agent-based macroeconomic model, focusing on the mechanism that normally makes a free-market economic system self-organizing and self-regulating, namely, its



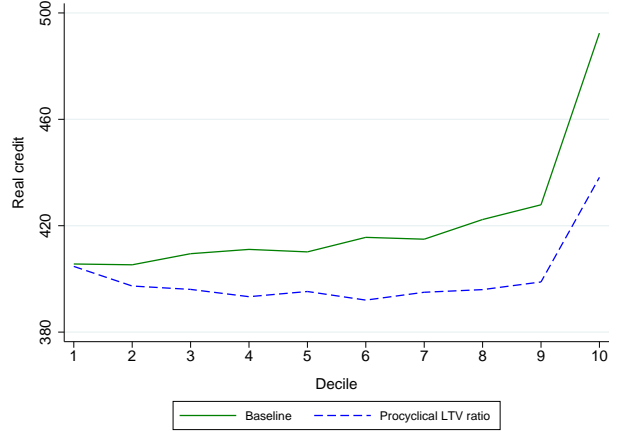
(a) Number of active shops



(b) Failures of financial viability test



(c) Banks in trouble



(d) Real credit

FIGURE 14: Bank lending and systemic performance with procyclical loan-to-value ratio

TABLE 5: Macroeconomic outcomes with procyclical loan-to-value ratio

Indicator	Median		Worst decile	
	Safe	Risky	Safe	Risky
Inflation	2.9	3.0	2.5	2.8
Output gap	7.7	7.6	20	13
Unemployment rate	6.2	5.8	13	8.7
Unemployment duration	11	8.8	18	11
Job loss rate	0.60	0.69	0.98	0.89
Volatility of output gap	3.4	2.6	12	7.2
Volatility of inflation	0.76	0.63	1.2	0.86
Annual bank failure rate	1.0	2.5	1.0	2.9
Fraction of banks in trouble	4.5	0.97	17	1.6

network of specialist trading enterprises. We applied our model to explore the role that the banking system plays in supporting this mechanism.

Our investigation generates a number of interesting results. First, it provides a framework for understanding rare disasters: states of the world in which the market coordination mechanism has lost control to such an extent that the evolving network of trade intermediaries is no longer able to counteract negative shocks, thus leading to a systemic collapse. Second, our findings indicate that by financing replacement firms and additionally supporting existing ones in the event of an adverse shock, banks confer a stabilizing influence on macroeconomic performance, and this influence is most noticeable in worst-case scenarios. In this respect, banks in our model provide a financial stabilizer that more than counteracts the familiar financial accelerator. Third, related to the stabilizing role of banks, our analysis highlights the conflict between micro-prudential bank regulation and macroeconomic stability: while tighter lending standards make little difference in normal scenarios, in worst-case scenarios, the economy exhibits significantly weaker performance under more stringent regulations (such as higher capital requirements or lower loan-to-value ratios), because these regulations suppress the lending activity that is especially needed during a rare disaster.

We conclude by reiterating the exploratory nature of our investigation but nevertheless emphasize its methodological contribution, in terms of yielding insights that might be elusive in more conventional frameworks for analyzing the impact of the banking system on macroeconomic performance.

References

- Aghion, P., Howitt, P., 2009. *The Economics of Growth*. Cambridge: MIT Press.
- Akerlof, G.A., Yellen, J.L., 1988. Fairness and unemployment. *American Economic Review: Papers and Proceedings* 78(2), 44–49.
- Altman, E.I., 1968. Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. *Journal of Finance* 23(4), 589–609.
- Amano, R., Moran, K., Murchison, S., Rennison, A., 2009. Trend inflation, wage and price rigidities, and productivity growth. *Journal of Monetary Economics* 56(3), 353–364.
- Ashraf, Q., Gershman, B., Howitt, P., 2016. How inflation affects macroeconomic performance: an agent-based computational investigation. *Macroeconomic Dynamics* 20(S2), 558–581.
- Assenza, T., Delli Gatti, D., Grazzini, J., 2015. Emergent dynamics of a macroeconomic agent based model with capital and credit. *Journal of Economic Dynamics and Control* 50, 5–28.
- Austin, J.S., Rosenbaum, D.I., 1990. The determinants of entry and exit rates into U.S. manufacturing industries. *Review of Industrial Organization* 5(2), 211–223.
- Bank for International Settlements, 1999. Capital requirements and bank behaviour: the impact of the Basle Accord. Basle Committee on Banking Supervision Working Papers No. 1.

- Barro, R.J., Grossman, H.I., 1976. *Money, Employment, and Inflation*. New York: Cambridge University Press.
- Basu, N., Pryor, R., Quint, T., 1998. ASPEN: a microsimulation model of the economy. *Computational Economics* 12(3), 223–241.
- Baumol, W.J., 2010. *The Microtheory of Innovative Entrepreneurship*. Princeton: Princeton University Press.
- Beaver, W.H., McNichols, M.F., Rhie, J.-W., 2005. Have financial statements become less informative? Evidence from the ability of financial ratios to predict bankruptcy. *Review of Accounting Studies* 10(1), 93–122.
- Bennett, R.L., Bergmann, B.R., 1986. *A Microsimulated Transactions Model of the United States Economy*. Baltimore: Johns Hopkins University Press.
- Bernanke, B., Gertler, M., 1989. Agency costs, net worth, and business fluctuations. *American Economic Review* 79(1), 14–31.
- Bils, M., 1987. The cyclical behavior of marginal cost and price. *American Economic Review* 77(5), 838–855.
- Bils, M., Klenow, P.J., 2004. Some evidence on the importance of sticky prices. *Journal of Political Economy* 112(5), 947–985.
- Bohn, H., 1998. The behavior of U.S. public debt and deficits. *Quarterly Journal of Economics* 113(3), 949–963.
- Braun, M., Larrain, B., 2005. Finance and the business cycle: international, inter-industry evidence. *Journal of Finance* 60(3), 1097–1128.
- Bris, A., Welch, I., Zhu, N., 2006. The costs of bankruptcy: Chapter 7 liquidation versus Chapter 11 reorganization. *Journal of Finance* 61(3), 1253–1303.
- Burstein, A., Hellwig, C., 2008. Welfare costs of inflation in a menu cost model. *American Economic Review: Papers and Proceedings* 98(2), 438–443.
- Chinazzi, M., Fagiolo, G., 2015. Systemic risk, contagion, and financial networks: a survey. Institute of Economics, Scuola Superiore Sant’Anna, Laboratory of Economics and Management (LEM) Working Paper Series No. 2013/08.
- Cincotti, S., Raberto, M., Teglio, A., 2010. Credit money and macroeconomic instability in the agent-based model and simulator Eurace. *Economics: The Open-Access, Open-Assessment E-Journal* 4, 1–32.
- Clower, R.W., 1965. The Keynesian counter-revolution: a theoretical appraisal. In: Hahn, F.H., Brechling, F.P.R. (Eds.), *The Theory of Interest Rates*. London: Macmillan, pp. 103–125.
- Coleman, N., Feler, L., 2015. Bank ownership, lending, and local economic performance during the 2008–2009 financial crisis. *Journal of Monetary Economics* 71, 50–66.

- Dawid, H., Gemkow, S., Harting, P., van der Hoog, S., Neugart, M., 2012. The Eurace@Unibi model: an agent-based macroeconomic model for economic policy analysis. Bielefeld University, Faculty of Business Administration and Economics, Working Papers in Economics and Management No. 05-2012.
- Dawid, H., Harting, P., Neugart, M., 2014. Economic convergence: policy implications from a heterogeneous agent model. *Journal of Economic Dynamics and Control* 44, 54–80.
- Deissenberg, C., van der Hoog, S., Dawid, H., 2008. EURACE: a massively parallel agent-based model of the European economy. *Applied Mathematics and Computation* 204(2), 541–552.
- Delli Gatti, D., Desiderio, S., Gaffeo, E., Cirillo, P., Gallegati, M., 2011. *Macroeconomics from the Bottom-up*. Milan: Springer.
- Delli Gatti, D., Di Guilmi, C., Gaffeo, E., Giulioni, G., Gallegati, M., Palestrini, A., 2005. A new approach to business fluctuations: heterogeneous interacting agents, scaling laws and financial fragility. *Journal of Economic Behavior and Organization* 56(4), 489–512.
- Delli Gatti, D., Gallegati, M., Greenwald, B., Russo, A., Stiglitz, J.E., 2010. The financial accelerator in an evolving credit network. *Journal of Economic Dynamics and Control* 34(9), 1627–1650.
- Dosi, G., Fagiolo, G., Napoletano, M., Roventini, A., 2013. Income distribution, credit and fiscal policies in an agent-based Keynesian model. *Journal of Economic Dynamics and Control* 37(8), 1598–1625.
- Dosi, G., Fagiolo, G., Napoletano, M., Roventini, A., Treibich, T., 2015. Fiscal and monetary policies in complex evolving economies. *Journal of Economic Dynamics and Control* 52, 166–189.
- Dosi, G., Fagiolo, G., Roventini, A., 2010. Schumpeter meeting Keynes: a policy-friendly model of endogenous growth and business cycles. *Journal of Economic Dynamics and Control* 34(9), 1748–1767.
- Dunne, T., Roberts, M.J., Samuelson, L., 1988. Patterns of firm entry and exit in U.S. manufacturing industries. *RAND Journal of Economics* 19(4), 495–515.
- Durlauf, S.N., Maccini, L.J., 1995. Measuring noise in inventory models. *Journal of Monetary Economics* 36(1), 65–89.
- Easterly, W., Islam, R., Stiglitz, J.E., 2001. Shaken and stirred: explaining growth volatility. In: Pleskovic, B., Stern, N. (Eds.), *Annual World Bank Conference on Development Economics 2000*. Washington: The World Bank, pp. 191–211.
- Eggertsson, G.B., Woodford, M., 2003. The zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity* 34(1), 139–233.
- Eliasson, G., 1977. Competition and market processes in a simulation model of the Swedish economy. *American Economic Review: Papers and Proceedings* 67(1), 277–281.

- Fabiani, S., Druant, M., Hernando, I., Kwapil, C., Landau, B., Loupias, C., Martins, F., Mathä, T., Sabbatini, R., Stahl, H., Stokman, A., 2006. What firms' surveys tell us about price-setting behavior in the Euro area. *International Journal of Central Banking* 2(3), 3–47.
- Fagiolo, G., Napoletano, M., Roventini, A., 2008. Are output growth-rate distributions fat-tailed? Some evidence from OECD countries. *Journal of Applied Econometrics* 23(5), 639–669.
- Fonseca, R., Lopez-Garcia, P., Pissarides, C.A., 2001. Entrepreneurship, start-up costs and employment. *European Economic Review* 45(4–6), 692–705.
- Gertler, M., Kiyotaki, N., 2010. Financial intermediation and credit policy in business cycle analysis. In: Friedman, B.M., Woodford, M. (Eds.), *Handbook of Monetary Economics*, Vol. 3. Amsterdam: North-Holland, pp. 547–599.
- Golosov, M., Lucas, Jr., R.E., 2007. Menu costs and Phillips curves. *Journal of Political Economy* 115(2), 171–199.
- Greenwald, B.C., Stiglitz, J.E., 1993. Financial market imperfections and business cycles. *Quarterly Journal of Economics* 108(1), 77–114.
- Hall, R.E., 1995. Lost jobs. *Brookings Papers on Economic Activity* 26(1), 221–273.
- Holmstrom, B., Tirole, J., 1997. Financial intermediation, loanable funds, and the real sector. *Quarterly Journal of Economics* 112(3), 663–691.
- Howitt, P., 2006. The microfoundations of the Keynesian multiplier process. *Journal of Economic Interaction and Coordination* 1(1), 33–44.
- Howitt, P., 2008. Macroeconomics with intelligent autonomous agents. In: Farmer, R.E.A. (Ed.), *Macroeconomics in the Small and the Large: Essays on Microfoundations, Macroeconomic Applications and Economic History in Honor of Axel Leijonhufvud*. Cheltenham: Edward Elgar, pp. 157–177.
- Howitt, P., Clower, R., 2000. The emergence of economic organization. *Journal of Economic Behavior and Organization* 41(1), 55–84.
- Khemani, R.S., Shapiro, D.M., 1986. The determinants of new plant entry in Canada. *Applied Economics* 18(11), 1243–1257.
- Kiyotaki, N., Moore, J., 1997. Credit cycles. *Journal of Political Economy* 105(2), 211–248.
- Leijonhufvud, A., 1973. Effective demand failures. *Swedish Journal of Economics* 75(1), 27–48.
- Leijonhufvud, A., 1993. Towards a not-too-rational macroeconomics. *Southern Economic Journal* 60(1), 1–13.
- Leijonhufvud, A., 2006. Agent-based macro. In: Tesfatsion, L., Judd, K.L. (Eds.), *Handbook of Computational Economics*, Vol. 2. Amsterdam: North-Holland, pp. 1625–1637.

- Levin, A., López-Salido, D., Nelson, E., Yun, T., 2010. Limitations on the effectiveness of forward guidance at the zero lower bound. *International Journal of Central Banking* 6(1), 143–189.
- Levine, R., 2005. Finance and growth: theory and evidence. In: Aghion, P., Durlauf, S.N. (Eds.), *Handbook of Economic Growth*, Vol. 1A. Amsterdam: North-Holland, pp. 865–934.
- Loayza, N.V., Raddatz, C., 2007. The structural determinants of external vulnerability. *World Bank Economic Review* 21(3), 359–387.
- Midrigan, V., 2011. Menu costs, multi-product firms, and aggregate fluctuations. *Econometrica* 79(4), 1139–1180.
- Myers, S.C., Majluf, N.S., 1984. Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics* 13(2), 187–221.
- Nakamura, E., Steinsson, J., 2010. Monetary non-neutrality in a multisector menu cost model. *Quarterly Journal of Economics* 125(3), 961–1013.
- Orcutt, G., Caldwell, S., Wertheimer, II, R., 1976. *Policy Exploration through Microanalytic Simulation*. Washington: The Urban Institute.
- Patinkin, D., 1956. *Money, Interest, and Prices: An Integration of Monetary and Value Theory*. 1st ed. Evanston: Row, Peterson and Co.
- Phelps, E.S., Winter, Jr., S.G., 1970. Optimal price policy under atomistic competition. In: Phelps, E.S., et al. (Eds.), *Microeconomic Foundations of Employment and Inflation Theory*. New York: Norton, pp. 309–337.
- Pissarides, C.A., 2000. *Equilibrium Unemployment Theory*. 2nd ed. Cambridge: MIT Press.
- Popoyan, L., Napoletano, M., Roventini, A., 2016. Taming macroeconomic instability: monetary and macro prudential policy interactions in an agent-based model. Institute of Economics, Scuola Superiore Sant’Anna, Laboratory of Economics and Management (LEM) Working Paper Series No. 2015/33.
- Poterba, J.M., 2000. Stock market wealth and consumption. *Journal of Economic Perspectives* 14(2), 99–118.
- Raddatz, C., 2006. Liquidity needs and vulnerability to financial underdevelopment. *Journal of Financial Economics* 80(3), 677–722.
- Ramey, V.A., West, K.D., 1999. Inventories. In: Taylor, J.B., Woodford, M. (Eds.), *Handbook of Macroeconomics*, Vol. 1B. Amsterdam: North-Holland, pp. 863–923.
- Roberts, J.M., 1995. New Keynesian economics and the Phillips curve. *Journal of Money, Credit and Banking* 27(4), 975–984.
- Rotemberg, J.J., Woodford, M., 1999. The cyclical behavior of prices and costs. In: Taylor, J.B., Woodford, M. (Eds.), *Handbook of Macroeconomics*, Vol. 1B. Amsterdam: North-Holland, pp. 1051–1135.

- Shyam-Sunder, L., Myers, S.C., 1999. Testing static tradeoff against pecking order models of capital structure. *Journal of Financial Economics* 51(2), 219–244.
- Smets, F., Wouters, R., 2007. Shocks and frictions in U.S. business cycles: a Bayesian DSGE approach. *American Economic Review* 97(3), 586–606.
- Stigler, G.J., 1961. The economics of information. *Journal of Political Economy* 69(3), 213–225.
- Stock, J.H., Watson, M.W., 1999. Business cycle fluctuations in U.S. macroeconomic time series. In: Taylor, J.B., Woodford, M. (Eds.), *Handbook of Macroeconomics*, Vol. 1A. Amsterdam: North-Holland, pp. 3–64.
- Taylor, J.B., 1993. Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy* 39, 195–214.
- Tesfatsion, L., Judd, K.L. (Eds.), 2006. *Handbook of Computational Economics*, Vol. 2. Amsterdam: North-Holland.
- Wallis, J.J., North, D.C., 1986. Measuring the transaction sector in the American economy, 1870–1970. In: Engerman, S.L., Gallman, R.E. (Eds.), *Long Term Factors in American Economic Growth*. Chicago: University of Chicago Press, pp. 95–162.
- Wiklund, J., Baker, T., Shepherd, D., 2010. The age-effect of financial indicators as buffers against the liability of newness. *Journal of Business Venturing* 25(4), 423–437.
- Williamson, S.D., 1987. Financial intermediation, business failures, and real business cycles. *Journal of Political Economy* 95(6), 1196–1216.
- Woodford, M., 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton: Princeton University Press.

Appendices

A Details of the model

A.1 Calculation of profit

A shop owner's profit Π is defined as:

$$\Pi = R - (1 + i_D) \text{WB} - (i_L - i_D) L, \quad (\text{A1})$$

where i_D and i_L are, respectively, the weekly nominal interest rates on deposits and loans set by the shop owner's bank last week, R is the shop's sales revenue, WB is its wage bill, and L is its bank loan. This is consistent with our earlier definition of the shop owner's financial wealth as

$$A = M + D - L, \quad (\text{A2})$$

where M and D are, respectively, his holdings of cash and bank deposits. That is, according to the Haig-Simons definition, a shop owner's income is the change in his wealth that is not attributable to his consumption expenditures. Profit is income minus the "normal" return to his wealth ($i_D A$). Thus,

$$\Pi = \Delta A - i_D A + (1 + i_D) PCE, \quad (\text{A3})$$

where PCE is personal consumption expenditure. We multiply PCE by $(1 + i_D)$, because as we have seen in section 4.3 above, the agent holds cash in advance to finance all planned expenditures, so that one dollar in expenditure reduces next week's wealth by $(1 + i_D)$ dollars.

To see that definition (A1) is implied by (A2) and (A3), we just need to substitute for ΔA the difference between inflows of cash and outflows of cash. Consider first the case in which the shop owner has no outstanding bank loan ($L = 0$). In this case, the inflows are the revenues R and the interest income on deposits $i_D D$, whereas the outflows are WB and PCE . Therefore,

$$\Delta A = R - \text{WB} - PCE + i_D D.$$

But, $D = A - M$, and by the cash-in-advance constraint, $M = \text{WB} + PCE$. Hence,

$$\Delta A = R - \text{WB} - PCE + i_D (A - \text{WB} - PCE).$$

The formula (A1) follows from this and (A3).

Next, consider the case where $L > 0$. As we have seen in section 4.3, the shop owner will exhaust his deposit holdings before taking out a loan. So in this case, $D = 0$, $A = M - L$, and the change in wealth is the same as before but, instead of an inflow of interest income, there is an

outflow of loan interest $i_L L$. Therefore,

$$\begin{aligned}\Delta A &= R - \text{WB} - PCE - i_L L \\ &= R - \text{WB} - PCE - i_D (M - A) - (i_L - i_D) L \\ &= R - (1 + i_D) (\text{WB} + PCE) + i_D A - (i_L - i_D) L,\end{aligned}$$

where the last line is derived using the cash-in-advance constraint: $M = \text{WB} + PCE$. Again, the formula (A1) follows from this and (A3).

In all cases, profit is defined using the most recent actual values of R , WB , and L , except in the entry stage (section 4.1), where Π refers to the profit implied by the prospective entrant's business plan. In the latter case, R is the revenue that would be raised if the sales target y^{trg} were achieved at the price implied by the business plan, WB is the wage cost of hiring enough workers ($y^{\text{trg}} + F - 1$) to cover the expected weekly variable and fixed cost at the planned wage w , and L is the size of loan needed to finance the setup cost and the first period's planned expenditure.

A.2 A shop owner's budget planning

The shop owner's constraints in his financial transactions are:

$$\begin{aligned}M &\geq 0, D \geq 0, L \geq 0, \\ M + D/(1 + i_D) + \bar{L} &= \bar{M} + \bar{D} + L/(1 + i_L), \\ L &\leq \text{CL} \cdot (1 + i_L).\end{aligned}$$

It is possible for a shop owner to satisfy all of his constraints if and only if $\bar{M} + \bar{D} + \text{CL} \geq \bar{L}$. Whenever this inequality fails to hold, the shop is declared bankrupt.

Recall that the shop owner's financial wealth is $A = \bar{M} + \bar{D} - \bar{L}$. The following 5 cases are possible:

1. If $A + \text{CL} < 0$, the shop goes bankrupt.
2. If $0 \leq A + \text{CL} < \text{WB}^{\text{trg}}$, the shop owner sets $E = 0$, withdraws all of his deposits, and borrows as much as he can from the bank (CL). His cash holdings after repaying the loan are thus $M = A + \text{CL}$.
3. If $\text{WB}^{\text{trg}} \leq A + \text{CL} < \text{WB}^{\text{trg}} + E$, the shop owner withdraws all deposits, borrows as much as he can from the bank, and sets $E = A + \text{CL} - \text{WB}^{\text{trg}}$; that is, his priority is to have enough cash to cover the target payroll. As before, $M = A + \text{CL}$.
4. If $A < \text{WB}^{\text{trg}} + E \leq A + \text{CL}$, the shop owner can afford to finance the entire wage bill and desired consumption expenditure, but cannot pay off the whole outstanding loan. He then pays off as much of the loan as he can and sets $D = 0$ and $M = \text{WB}^{\text{trg}} + E$.

5. If $WB^{\text{trg}} + E \leq A$, the shop owner can afford the entire wage bill plus desired consumption expenditure, and he can also repay the whole outstanding loan. In this case, the shop owner pays off the loan, sets $M = WB^{\text{trg}} + E$, and puts his remaining wealth into his deposit account.

A.3 Estimates and forecasts used in monetary policy

On each fixed action date, the government adjusts its estimate of the real interest rate as follows:

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

where η_r is a fixed target interest rate adjustment coefficient, r_0^* is the initial real interest rate target, and $f(\pi, r^*)$ is a ‘‘squasher.’’ As such, the interest rate target is increased or decreased as the current inflation rate exceeds or falls short of the government’s target. The squasher makes sure that this change is symmetric (S-shaped) around the point $\pi = \pi^*$ and never causes r^* to become negative. Around the point $(\pi, r^*) = (\pi^*, r_0^*)$, $\Delta r^* \approx \eta_r (\pi - \pi^*)$, that is, the adjustment of the target interest rate is roughly proportional to the deviation of actual inflation from the government’s target inflation rate.

The government models the year-to-year behavior of weekly log GDP as the following AR(1) process: $y_t = \alpha_{yt} + \lambda_{yt} y_{t-1} + \xi_{yt}$, where $\tilde{y} \equiv \alpha_y / (1 - \lambda_y)$ is defined as potential weekly log GDP and ξ_{yt} is an i.i.d. disturbance term. We assume that the government begins to adjust all the estimates and targets after T_{cb} years. Given the initial estimates $(\hat{\alpha}_{yt-1}, \hat{\lambda}_{yt-1})$, the government re-estimates these parameters once per year using the following recursive OLS scheme on annual data:⁵⁶

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

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 log potential GDP, equal to $\hat{\alpha}_{yt} / (1 - \hat{\lambda}_{yt})$, is the one subsequently used in the Taylor rule.

Similarly, the government models the annual evolution of inflation as the following AR(1) process: $z_t = \lambda_{\pi t} z_{t-1} + \xi_{\pi t}$, 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 estimate $\hat{\lambda}_{\pi 0}$, the government recursively re-estimates this parameter on an annual basis after the learning period is over, that is,

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

⁵⁶Actually, we suppose that the government starts directly with an estimate \tilde{y}_0 of potential output, which it uses in the Taylor rule for the first T_{cb} years. It then uses the recursive OLS scheme starting with $\hat{\alpha}_{T_{\text{cb}}-1} = (1 - \hat{\lambda}_{y0}) \tilde{y}_0$.

A.4 The capitalization factor

A person who could count on the government's projections of inflation and interest rates, as described in section 4.5 and appendix A.3, who would never be rationed in purchasing consumption goods, who faced consumption prices that rose each period at the economy-wide rate of inflation, who held only enough cash to satisfy the cash-in-advance constraint each period, who invested the rest of his financial wealth in deposits yielding i_{wt} , and who had a constant perpetual stream of income, starting this week, equal to Y^P in today's ($t = 1$) dollars, would be able to afford any consumption stream $\{c_t\}_1^\infty$ satisfying⁵⁷

$$\sum_{t=1}^{\infty} R_t c_t = \frac{A_1 + VY^P}{P_1}, \quad (\text{A4})$$

where V is the capitalization factor from section 4.5 and each R_t is the present value of a real dollar in week t , defined recursively as

$$R_1 = 1, \quad R_{t+1} = \frac{1 + \pi_{wt}}{1 + i_{wt}} R_t \quad \forall t \geq 1.$$

The division of each term in equation (2) by $1 + \pi_{wt}$, as well as the fact that income at t is discounted by R_{t+1} in equation (A4) while consumption is only discounted by R_t , reflects the cash-in-advance constraint; the dollar earned in t cannot be spent until $t + 1$.

Note that the expenditure function (1) would apply if the person knew for certain what future incomes and interest rates would be and if he were choosing E so as to maximize a standard intertemporal additive logarithmic utility function with a weekly rate of time preference $\rho_w = v/(1 - v)$, subject to the lifetime budget constraint (A4). We use this interpretation of the expenditure function when calibrating the model, which we did in terms of the annual rate of time preference ρ , defined by $(1 + \rho) = (1 + \rho_w)^{48}$.

A.5 Wage and price setting

The wage adjustment equation used by each firm is:

$$w = \bar{w} \cdot \left[\left(1 + \beta \cdot \left(\frac{\overline{x^{\text{trg}}}}{\overline{x^{\text{pot}}}} - 1 \right) \right) \cdot (1 + \pi^*) \right]^{\Delta/48},$$

where \bar{w} is the pre-existing wage, $\overline{x^{\text{trg}}}$ is the average input target over the past Δ weeks, and $\overline{x^{\text{pot}}}$ is the average potential input over the same period (i.e., the number of people having an employment relationship with the shop, even if they were laid off or if they refused to work because they were not paid).⁵⁸ The parameter β thus indexes the degree of wage and price flexibility in the economy.

⁵⁷To see this, note that the flow budget constraint facing the individual each period would be $A_{t+1} = (A_t - P_t c_t)(1 + i_{wt}) + P_t Y^P / P_1$, which, together with the no-Ponzi-game condition, is equivalent to equation (A4).

⁵⁸In computing this expression, we use the maximum of $\overline{x^{\text{pot}}}$ and the shop's fixed cost F to avoid division by zero when potential employment falls to zero.

This wage adjustment anticipates inflation over the coming contract period at an annual rate equal to the government's target π^* .

The shop's "normal" price is $p^{\text{nor}} = (1 + \mu)w / (1 - \tau)$, which would equate its after-tax price to its wage times its desired markup, corresponding to the rule discussed earlier in section 4.1. The shop will choose this normal price unless its inventories are too far from the desired level, namely, its target sales. Specifically, it will set

$$p = \begin{cases} p^{\text{nor}} \cdot \delta_p^{-1}, & \text{if } I > y^{\text{trg}} \cdot \text{IS}; \\ p^{\text{nor}} \cdot \delta_p, & \text{if } I < y^{\text{trg}} \cdot \text{IS}^{-1}; \\ p^{\text{nor}}, & \text{otherwise.} \end{cases}$$

The frequency of price changes will therefore be endogenous. A shop will change its posted price 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 the price when its inventory-to-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 threshold, it raises its price by that same factor.

A.6 The no-shock flexible-price stationary state

This appendix section provides further details regarding the no-shock stationary state described in section 5.1. Each shop begins every week with a common wage rate equal to $W = (1 + \pi_w^*)W_0$, where W_0 was the common wage rate last week, and with a price equal to $P = (1 + \bar{\mu})W / (1 - \tau)$, where the tax rate τ equals

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

As mentioned in section 4.7, this is the tax rate that leaves the government's real debt-to-GDP ratio undisturbed. We are assuming that all markups in this stationary state are equal to $\bar{\mu}$.

There are no bank loans outstanding, so banks are just conduits, converting deposits into government bonds. The initial outstanding stock of bonds is $B = b^* (1 + i_w) \cdot 48Y^* \cdot P_0$, where $P_0 = P / (1 + \pi_w^*)$ is last week's price level. These bonds are held equally by all banks, and each non-bank-owning person holds the amount B/N of bank deposits. The weekly interest rate i_w is given by $1 + i_w = (1 + \rho_w)(1 + \pi_w^*)$. The money supply at the start of the week is $M = W_0(N - n) + (1 - \tau)P_0Y^*$, which is the sum of all wage receipts of people not owning a shop and all sales receipts (ex taxes) of shop owners from last period.

Each person starts the period with an effective wage equal to W_0 and with effective prices for both consumption goods equal to P_0 . Each shop owner begins with a permanent income equal to last week's profit $\Pi = [(\bar{\mu} - i_w)(n - 2 - F) - (1 + i_w)(F - 1)]W_0$ and with money holdings equal to last week's revenue $(1 - \tau)P_0(n - 2 - F)$. Each person not owning a shop starts with money holdings equal to permanent income, which is equal to last period's wage income W_0 .

No one holds any legacy capital, no banks hold seized capital, and the firesale queues are all empty. Banks hold no reserves, so each bank's equity is its share of the government debt minus its customers' deposits, which amounts to B/N .

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. The government's real interest target is $r^* = \rho$, and its estimate of log potential GDP is $\tilde{y}_0 = \ln(Y^*)$. Its latest published capitalization factor is

$$V = \frac{1}{1 + \pi_w^*} \cdot \frac{1}{\rho_w}.$$

It is straightforward to verify that this configuration will repeat itself indefinitely, with all nominal magnitudes – money and bond holdings, actual and effective wages and prices, and permanent incomes – rising each week at the constant rate π_w^* , provided that the fixed cost F is small enough that shops always pass the profitability test during the exit stage. Indeed, in our baseline calibration, $F = 3.5$, whereas the critical value F^* at which the profitability test fails is 63.

B Empirical validation tables

TABLE B.1: Cross-correlations of selected series with output (constant LTV ratio)

Series (lag/lead)	-4	-3	-2	-1	0	1	2	3	4
Output	-0.046 (0.004)	-0.026 (0.004)	0.068 (0.004)	0.402 (0.003)	1.000 (0.000)	0.402 (0.003)	0.068 (0.004)	-0.026 (0.004)	-0.046 (0.004)
Consumption	-0.045 (0.004)	-0.024 (0.004)	0.077 (0.004)	0.434 (0.003)	0.991 (0.000)	0.358 (0.004)	0.060 (0.004)	-0.026 (0.004)	-0.045 (0.004)
Investment	-0.006 (0.003)	-0.019 (0.003)	-0.079 (0.003)	-0.248 (0.003)	0.037 (0.003)	0.320 (0.003)	0.062 (0.003)	0.001 (0.003)	-0.002 (0.003)
Ch. in invent.	-0.005 (0.003)	-0.015 (0.003)	-0.066 (0.003)	-0.191 (0.003)	0.066 (0.003)	0.289 (0.003)	0.035 (0.003)	-0.015 (0.003)	-0.010 (0.003)
Employment	-0.047 (0.004)	-0.026 (0.004)	0.071 (0.004)	0.413 (0.003)	0.992 (0.000)	0.370 (0.003)	0.049 (0.004)	-0.034 (0.004)	-0.048 (0.004)
Unemp. rate	0.047 (0.004)	0.026 (0.004)	-0.071 (0.004)	-0.413 (0.003)	-0.992 (0.000)	-0.370 (0.003)	-0.049 (0.004)	0.034 (0.004)	0.048 (0.004)
Vacancy rate	0.020 (0.004)	0.030 (0.004)	0.053 (0.004)	0.086 (0.005)	0.091 (0.005)	-0.103 (0.005)	-0.183 (0.004)	-0.137 (0.004)	-0.092 (0.004)
Productivity	-0.025 (0.004)	-0.019 (0.004)	0.006 (0.004)	0.136 (0.004)	0.577 (0.003)	0.414 (0.003)	0.152 (0.004)	0.030 (0.004)	-0.016 (0.004)
Inflation rate	-0.009 (0.004)	0.008 (0.004)	0.069 (0.004)	0.287 (0.004)	0.516 (0.003)	0.202 (0.004)	-0.069 (0.004)	-0.091 (0.004)	-0.084 (0.004)
Markup	0.077 (0.004)	0.087 (0.005)	0.089 (0.005)	0.065 (0.005)	-0.056 (0.005)	-0.167 (0.005)	-0.152 (0.005)	-0.118 (0.005)	-0.093 (0.005)
Nom. int. rate	-0.034 (0.004)	-0.020 (0.004)	0.049 (0.004)	0.295 (0.004)	0.883 (0.001)	0.529 (0.003)	0.070 (0.004)	-0.062 (0.004)	-0.082 (0.004)
Debt/GDP	-0.105 (0.004)	-0.133 (0.004)	-0.199 (0.004)	-0.384 (0.004)	-0.472 (0.005)	0.067 (0.004)	0.278 (0.003)	0.299 (0.003)	0.270 (0.004)

TABLE B.2: Cross-correlations of selected series with output (procyclical LTV ratio)

Series (lag/lead)	-4	-3	-2	-1	0	1	2	3	4
Real money	-0.012 (0.004)	0.043 (0.004)	0.211 (0.004)	0.671 (0.003)	0.591 (0.003)	0.072 (0.004)	-0.083 (0.004)	-0.106 (0.004)	-0.098 (0.004)
Real credit	0.000 (0.004)	0.016 (0.004)	0.050 (0.004)	0.150 (0.005)	0.480 (0.004)	0.469 (0.004)	0.193 (0.004)	0.023 (0.004)	-0.058 (0.004)
Bankrupt. haz.	0.019 (0.004)	0.006 (0.004)	-0.041 (0.004)	-0.200 (0.004)	-0.386 (0.004)	-0.149 (0.004)	-0.064 (0.004)	-0.008 (0.004)	0.013 (0.004)
Fin. const. haz.	0.036 (0.005)	0.028 (0.005)	0.011 (0.005)	-0.016 (0.005)	-0.116 (0.005)	-0.131 (0.005)	-0.081 (0.005)	-0.064 (0.005)	-0.051 (0.004)
Tr. bank haz.	0.019 (0.004)	0.012 (0.005)	-0.004 (0.005)	-0.029 (0.005)	-0.054 (0.005)	-0.089 (0.005)	-0.083 (0.005)	-0.063 (0.005)	-0.044 (0.005)

C Time series of cross-run statistics for output gap by decile




FIGURE C.1: Time series of cross-run mean output gap: deciles 1–9 (left) and 10 (right)




FIGURE C.2: Time series of cross-run variance of output gap: deciles 1–9 (left) and 10 (right)

D Macroeconomic performance without the zero lower bound




FIGURE D.1: Macroeconomic performance by decile without the zero lower bound

E Sensitivity analysis

In this section, we subject our main findings to a sensitivity analysis. In particular, we perturbed each of the model’s 33 parameters in turn, first setting it higher than the baseline value indicated in Table 1 and then setting it lower, and recorded the median value of the 40-year average output gap across 10,000 simulation runs. For all but two parameters, the perturbation was either plus or minus 25 percent of the baseline value. One exception was the size of price change δ_p , in which case we perturbed $\delta_p - 1$ by plus or minus 25 percent. The other exception was the number of banks m , which needs to be a divisor of the number of goods 50 to preserve the symmetry of the model’s setup across sectors; in this case, we set the perturbed values of m at the nearest divisors of 50, namely $m = 10$ and $m = 2$.

Figure E.1 shows the effects of varying each parameter on the median output gap.⁵⁹ There are four parameters that can raise the output gap (i.e., reduce output) by more than 1 percentage point, namely the fixed cost F , the length of the wage contract period Δ , the job search propensity σ , and the inventory adjustment speed λ_I . Of these, the fixed cost is the most important. When

⁵⁹The results from sensitivity analyses conducted on the median values of the other eight indicators listed in Table 3 are qualitatively similar, and are available from the authors upon request.




FIGURE E.1: Effect on median output gap from perturbing each parameter

it rises, firms are less likely to enter and more likely to fail, resulting in a lower number of shops on average and, thus, a greater degree of market disorganization.

More than half the parameters appear to have little effect on median performance. As shown in Figure E.1, 15 of the 33 parameters have an effect of less than or equal to one-tenth of a percentage point on median output when varied ± 25 percent. These parameters include all the initialization values used in the government’s forecasts, namely its initial estimates for the output gap $\ln(Y^*) - \tilde{y}_0$, its initial estimate for the autoregressive coefficient $\lambda_{\pi 0}$ in its forecast equation for inflation, and its initial real interest rate target r_0^* . They also include technical parameters such as the adjustment speed η_r in revising the estimate of the natural interest rate and the length T_{cb} of its learning period. There are, however, another 14 parameters whose values we thought would matter but that seem to have little effect on median performance. These parameters include the inflation target π^* , the other two Taylor-rule parameters γ_y and γ_π , and other “policy” parameters such as the loan-to-value ratio h , the required capital adequacy ratio κ , and the Fed’s discount rate premium s_d . Several private sector parameters are also relatively unimportant by the same criterion, namely the speed of adjustment λ_p of permanent income, the setup cost S , the foreclosure cost parameter C_b , the annual loan spread s , and the slope l of each bank’s loan approval schedule.

Despite the fact that several parameters have a quantitatively significant effect on median performance, Figure E.2 shows that our main result with respect to macroeconomic performance with safe and risky banks in normal and worst-decile scenarios is robust to 25 percent perturbations in almost all parameters. Specifically, the light blue line in this figure shows the safe-risky difference




FIGURE E.2: Effect on safe-risky output gap difference from perturbing each parameter

in the worst-decile average output gap in our baseline calibration, which is 0.046. The lightly shaded bars indicate the same difference under parameter perturbations, orange denoting +25 percent and green denoting -25 percent. All but 2 of these 66 bars indicate a positive difference, as in our baseline calibration, and neither of the 2 exceptions yielded a difference larger than 0.01 in absolute value. The parameter whose perturbations had the most effect on this difference was again the fixed cost F , which yielded a safe-risky difference of 0.262 when it was increased by 25 percent.