Banks and the credit cycle
(EARLY DRAFT)

Adrian Penalver
Paris School of Economics

March 25, 2014

Abstract

This paper presents a model of the credit cycle focusing on the time varying behaviour of banks. It delivers key features of the credit cycle including counter-cyclical credit standards and default rates as well as pro-cyclical loan and deposit interest rates. An interesting feature of the model is that all agents, including the bank, are making rational forward looking decisions using model-consistent expectations. There are behavioural frictions in the model but these do not include principal/agent problems. The credit cycle can therefore be explained by rational responses to incentives and without recourse to psychological or sociological factors or misalignment of incentives. The paper builds a model of credit risk management for a bank making long term loans to fund projects subject to idiosyncratic and aggregate shocks. The objective of the bank is to maximise profits, minimise variance in profits whilst trying to keep the endogenous volume of loans as close as possible to the endogenous supply of deposits. The solution to the dynamic programming problem faced by the bank is a set of policy functions for the loan interest rate, the deposit interest rate, a monitoring rate and a loan covenant threshold.

1. Introduction

Since the seminal papers of Bernanke and Gertler (1989) and Kiyotaki and Moore (1997), macroeconomists have modeled credit relationships as borrowing or collateral constraints to combat moral hazard. These models demonstrated the importance of asset prices and borrower net worth in propagating and amplifying shocks. Credit cycles, asset price cycles and business cycles all move together. Many recent papers have added a moral hazard friction between banks and their depositors to show that banking sector capital can amplify the effects on credit and output (Chen (2001), Meh and Moran (2008), Gertler and Karadi (2009) and Gertler and Kiyotaki (2009)).

These models are highly convincing on the adjustment mechanisms that they consider but they are deficient in several respects as complete descriptions of the credit cycle. Firstly, they cannot explain the counter-cyclical pattern of default rates. Debt repudiation is ruled out by construction and project failure (if considered) is a constant probability. Secondly, these models do not give any explicit role to financial institutions. They either don't exist or are simply a source of friction. Yet financial institutions, and more specifically, banks, behave quite differently during business cycles. Figure 1 illustrates the path of survey measures of bank credit standards in the United States and Europe. It is clear from the longer run of data in the United States that credit standards are tightened during recessions and relaxed during periods of growth. The European data show a similar pattern. During credit booms, loan maturities are longer and covenant restrictions are weaker. Thirdly, they do not explain how resources are allocated *between* firms over the credit cycle. Do credit crunches improve the allocation of resources through Schumperterian "waves of creative destruction" (Schumpeter (1939)) or do they "sully" the allocation mechanism (as in Barlevy (2002))?

The purpose of this paper is to develop a credit and business cycle model that explains these features of the credit cycle. Bank credit standards and default rates are counter-cyclical but interest rates are pro-cyclical. Moreover, some of those who borrow during booms would not receive credit during busts. What is perhaps most surprising about these results is that all agents are making fully rational and forward looking decisions. There is no myopia, greed or misperception of risk.

This paper builds on the bank-intermediated firm dynamics model of Penalver (2014). In that model, a large monopoly bank repeatedly makes long-term loan commitments to finance projects subject to persistent idiosyncratic stochastic shocks which are privately observable by the borrower. The bank pays a monitoring cost to observe the state of continuing borrowers and recalls loans found below a covenant threshold. The monitoring intensity and covenant threshold are fully anticipated by the borrowers and affect the loan interest rate they are willing

Figure 1: Credit standards in the US and Europe

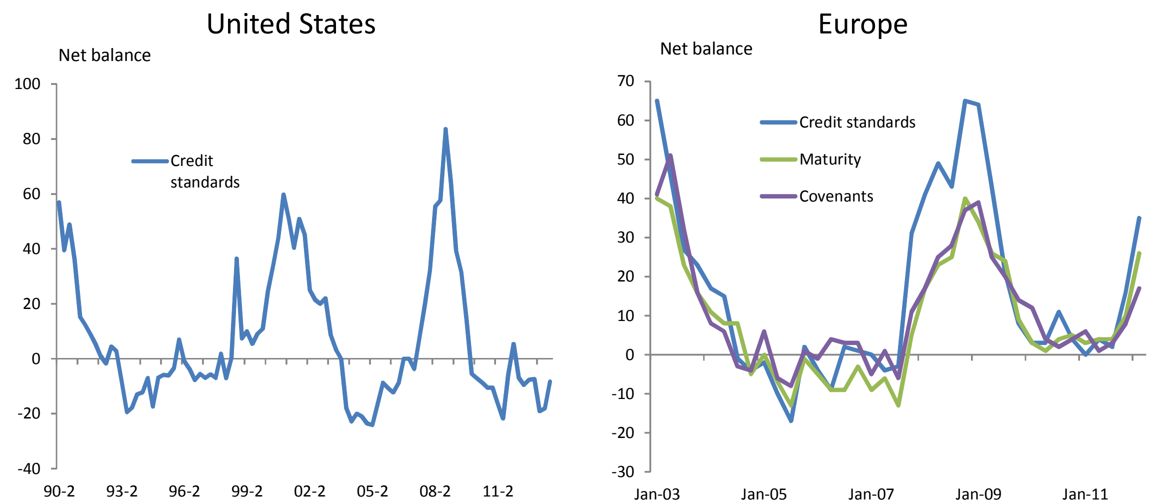

to pay. In picking the optimal monitoring intensity, covenant threshold and loan interest rate, the bank is trading off expected loss given default, the loan spread and monitoring costs, subject to the constraint that deposits must match loans. In that model, there is only idiosyncratic risk and since the bank can rely on the law of large numbers, constrained profit maximisation is a static problem. The optimal triplet of monitoring intensity, covenant threshold and loan interest rate is chosen over stationary distributions which satisfy the balance sheet constraint. For obvious reasons, the solution of that model will be refered to in this paper as the steady state. There are dynamics at the firm level but everything is constant for the bank.

The extension in this paper is to add an aggregate shock process through a common persistent shock to project profitability. This affects firm entry, exit and continuation decisions. As a result, the distribution of firms (over idiosyncratic profitability) is history dependent since it depends on entry and exit decisions made in earlier periods and thus the history of aggregate shocks. By construction, the distribution of firms is also the distribution of loans on the balance sheet of the the bank. In these circumstances, the bank faces a complex dynamic programming problem. How does it set deposit and loan interest rates and the loan terms and conditions in order to satisfy its budget constraint, knowing that the resulting distribution of firms (and loans) affects the decision problem it will face in subsequent periods? And if more than one set of policy rules can satisfy the balance sheet constraint (at least to a close approximation), which combination has the highest profits with the least variance? This paper answers those questions and shows that the resulting paths for the default rate, loan interest rates, loan covenants and loan monitoring intensity replicate key features of the credit cycle.

The structure of the paper is as follows. Section 2 sets out the general structure of the model and summarises the solution of the steady-state problem in which there is only idiosyncratic risk. Section 3 explains the solution when the aggregate shock is added. Section 4 analyses the credit cycle derived from an illustrative calibration of the model by describing the dynamics of loan interest rates, deposit interest rates, loan covenants, monitoring intensity and the default rate. Section 5 concludes.

2. Model set-up

This section starts by setting out the dynamic model in the most general terms and then summarises the equilibrium in steady-state.

The economy, following Penalver (2014), is populated by a measure 1 of infinitely small *ex ante* homogeneous agents. Agents are risk neutral, live forever and discount the future at rate β. Time is discrete.

Economic activity takes place through "projects" which can last indefinitely and deliver a gross stochastic idiosyncratic return per period $q(a)$ where $a \in A$ is a profitability index over the compact support $[0, 1]$. Some projects are more profitable than others at any point in time reflecting, for example, time varying differences in market power, productivity, managerial competence etc, the deeper sources of which are unmodeled. A project's profitability index evolves according to a time homogeneous Markov process represented by the cummulative distribution function $F(a', a)$. This process is assumed to have the following properties:

A (i) $F(a', a)$ is continuous in a and a'; (ii) profitability shocks are persistent, so $F(a', a)$ is strictly decreasing in a; (iii) but profitability shocks eventually die out and the monotone mixing condition is satisfied: $F^n(\epsilon, a) > 0\ \forall \epsilon$ for some n where $F^n(\epsilon, a)$ is the conditional probability distribution of profitability in n periods time given a. So from any given level of profitability, it is possible to transit to any other profitability level in a finite number of periods.[1] Since there are exit thresholds, this assumption implies that all projects will almost surely close at some future point.

At any point in time, some of the agents are "inventors". Similar to labour search models such as McCall (1970), inventors receive one idea per period with profitability a drawn from an independent and identical distribution, $G(a)$. If an inventor decides to commence a project, she pays a start up cost S and enters the following period with gross idiosyncratic return $q(a')$ according to $F(a', a)$.[2] Projects are assumed to require 2 units of capital. Entrepreneurs cannot get credit directly from inventors (perhaps for the reasons described by Diamond (1984)) but can borrow up to 1 unit from the bank (which is assumed to be binding).[3] If an inventor decides not to enter, she remains an inventor next period and in the absence of any better alternative, deposits her savings at the bank on which she receives a deposit rate τ.

The other group of agents are running projects and are called "entrepreneurs". Entrepreneurs make net profits per period $q(a) + z - \rho(.)$ where $\rho(.) \in \mathbb{R}^{+}$ is the loan interest rate and z is the aggregate shock process to be described below. Only entrepreneurs can costlessly observe the profitability state of the project and the bank will need to pay m per loan to get a perfectly accurate report. The gross idiosyncratic profitability function $q(a)$ is assumed to be continuous and strictly increasing in a.

Bearing in mind their current and expected future profits, entrepreneurs decide whether to continue in production next period or exit and switch to being an inventor. Entrepreneurs can exit in two ways:

- "Orderly" exit occurs if an entrepreneur absorbs current period payoffs losses, $q(a) + z - \rho(.)$, (in this situation, losses) and pays a liquidation cost L to close the project. These liquidation costs might be pecuniary such as termination pay, liquidating stock at below cost and administrative costs or non-pecuniary such as lost human capital and reputation.

[1] The numerical example uses an AR(1) process with Gaussian shocks approximated over a finite dimension grid using the method of Tauchen (86).

[2] Caballero and Hammour (1994) estimate entry costs at around half a year's operating costs.

[3] This is not a particularly satisfying assumption. But one could assume that there is a collateral constraint in one direction and in the other direction agents prefer to borrow rather than invest more of their own funds to diversify risk.

- "Default" occurs if an entrepreneur files for bankruptcy protection in which case current period losses are excused (including repayment of loan interest) but the agent pays an exogenous bankruptcy cost B.[4] Naturally, L and B are calibrated so that bankruptcy is prefered over orderly exit only in extreme circumstances.

It will be assumed that agents make very simple "wealth" management decisions by targeting an average level of capital of one unit over their infinite lives. There is no growth in this economy so a stationary average stock of capital is perfectly reasonable and making it one unit is just a convenient normalisation. Consistent with this, agents consume their income on average. But aggregate income (output) is going to vary due to the aggregate shock process and the endogenous allocation of resource, although this latter effect is very small. It will be assumed that aggregate savings absorb 90% of the variation in income and that all of these savings are intermediated through the bank. Precisely how this is allocated between individual agents is not modeled explicitly.[5]

The purpose of this somewhat inelegant assumption is to incorporate aggregate savings behaviour without having to track the intensive margin of adjustment. It thereby bypasses all the issues associated with precautionary savings in heterogeneous agent models such as those analysed by Aiyagari (1994) and Krusell and Smith (1998).[6] The dimension that is important in this model is the distribution of productive agents across profitability states. The model is closer in this regard to Lee and Mukoyama (2008).[7] Having included the assumption of procyclical saving, the effect is actually quite small for familiar reasons. Average output of the economy in the numerical simulation is around 17% of capital and the minimum and maximum levels of output over the simulation are 16% and 18% respectively. Thus the assumption of aggregate savings (and thus investment) absorbing 90%

[4] To simplfy subsequent calculations, it is assumed that entry, exit and bankruptcy costs are paid in the following period.

[5] But as an example, all agents could save 90% of the difference between their current income and the steady-state level.

[6] A future research project will, for completeness, try to endogenise savings choices across agents. As long as the entry and exit decisions are independent of the saving decisions, it is unlikely that this extension will offer any additional insights. It will likely merely validate the aggregate savings assumption.

[7] Khan and Thomas (2013) and Clementi and Palazzo (2013) track agents across both dimensions.

of the variation in output, means the capital stock varies by less than 1% either side of steady state.

Finally, it will be assumed that all agents have an exogenous endowment of income every period regardless of their circumstances sufficient to cover any expenses or losses. This exogenous endowment plays no role in the model beyond giving borrowers financial flexibility to continue with the project if they wish. In short, all exit decisions in the model are based on incentives not inability to pay.

Capital/savings are intermediated between inventors and entrepreneurs by a monopoly bank. As a consequence of the assumption of infinitely small agents and the law of large numbers, idiosyncratic risks have no aggregate stochastic effects. The bank offers the following deposit and loan contracts.

- Deposits earn the current deposit interest rate $\tau(.)$. Deposits can be withdrawn at the end of any period.

- The loan contract specifies the terms and conditions which apply in the current period. The terms and conditions are the bank policy functions for the loan interest rate $\rho(.)$, a monitoring intensity $\varphi(.)$ (where $0 < \varphi < 1$) and a covenant specifying a minimum net profit level $\xi(.)$. The loan is provided for the duration of the indefinite project but both parties have an option to terminate it each period. Each borrower has the option to repay the loan if he decides to exit production and the bank can demand full repayment if it discovers that the covenant condition has been breached. This demand is enforceable and by the endowment assumption, an entrepreneur has the ability to repay.

Since the bank uses the covenant to protect its interests, it follows that $\xi(.)$ must in all cases imply a trigger value of a at least as high as that at which entrepreneurs voluntarily exit or else the covenant would be redundant. Therefore the entrepreneur faces a utility loss from having his loan recalled and he cannot be expected to reveal the profitability state of the loan voluntarily. So to discover the state and enforce the covenant, the bank must monitor its continuing loan portfolio. The monitoring intensity $\varphi(.)$ gives the probability that any entrepreneur who chooses to remain in production is inspected.

The model in Penalver (2014) contains only idiosyncratic shocks and the problem for the bank is to find a loan interest rate, monitoring intensity and covenant threshold triplet (ρ, φ, ξ) which maximises profits subject to the constraint that the bank must have the same measure of deposits as loans. The equilibrium of that model is an invariant distribution with constant values for all the endogenous variables, which is why it is referred to as the steady state version of the more general model. So with constant values it does not matter in a sense whether the equilibrium loan terms $\left(\overline{\rho}, \bar{\varphi}, \bar{\xi}\right)$ are stipulated in the contract or are a rational expectation. In the model in this paper, there is an aggregate shock process so constant interest rates will no longer satisfy the balance sheet constraint. Therefore how the bank adjusts the loan terms and how expectations are formed are fundamentally important for the equilibrium.

There are 7 elements in the aggregate shock set, Z. They are symmetrically spaced with time homogeneous Markov transition probabilities $Z(z', z)$ which approximate an AR(1) process with normal errors using a Tauchen matrix (Tauchen (1986)). The aggregate shock is a common scalar shift in the gross return of all projects in that period.

To simplify notation, define $\Omega \equiv \{z, \theta_{-1}\}$ where θ_{-1} is an endogenous state variable at the end of the previous period. Exactly what is in θ is delayed until Section 3 when it becomes relevant. It is assumed that the endogenous state variable is accurately observed by the bank and the agents but only after they have taken the decisions that determine it. By contrast, the bank and the agents are assumed to be able to observe the aggregate shock before taking decisions. Ω is thus the information set on which everyone conditions their choices each period and gives the model a Markovian structure.

This is a recursive model and in each period the move order is the following:

1. Agents enter the period in their previously chosen situation (inventor or entrepreneur) with knowledge of θ_{-1}. The aggregate shock z is revealed. The inventors get an idea from $G(a)$ and entrepreneurs get a private update of their idiosyncratic profitability state according to $F(a', a)$.

2. The bank updates the loan interest rate, the monitoring intensity and the loan covenant threshold which in this recursive setting are functions of the information set. The loan interest rate is updated according to $\rho(\Omega)$, the monitoring intensity according to $\varphi(\Omega)$ and the covenant threshold according to $\xi(\Omega)$. In the interest of notational simplicity, summarise these terms as $\psi(\Omega)$. The bank also updates the deposit rate, $\tau(\Omega)$.

3. Entrepreneurs decide whether to continue with production next period or to exit either voluntarily or by defaulting. Payoffs are received and loan interest paid by non-defaulting entrepreneurs. Entrepreneurs who exit voluntarily inform the bank that they will repay their loan. Inventors receive deposit interest and decide whether to enter production next period based on their profitability draw.

4. The bank monitors ongoing loans at the stochastic rate $\varphi(\Omega)$ and recalls the loans of all entrepreneurs found below the covenant profitability threshold $\xi(\Omega)$.

5. The bank receives deposits from ongoing and new inventors and makes additional loans to entering entrepreneurs. θ is revealed.

Before diving into the details, which become quite complex in appearance very quickly, it is worth at this stage giving an overview of the solution.

The first step is determine what is is optimal for individual entrepreneurs and inventors to do given their profitability draws and what they know about the two aggregate states, z and θ_{-1}. Since these entry and exit decisions are forward looking, they will depend on the expectations entrepreneurs and inventors have about how their individual circumstances and the aggregate states will evolve. The exogenous aggregate state is discrete and the endogenous aggregate state is approximated at discrete points, so the aggregate state (and thus the information set, Ω) is a 2-dimensional grid. The economy will pass stochastically through these grid points. The agents, therefore, are working out the probabilities of passing from the current grid point to another.

The second step is to consider the problem faced by the bank. For any given shock to aggregate profitability, the effect on the loan portfolio will depend on the distribution of firms (and loans) over the idiosyncratic state. But this distribution depends on the history of entry and exit decisions and thus the history of aggregate shocks. These portfolio dynamics will affect the size of the loan portfolio and its profitability through time through default risk and loss given default. The objective of the bank is to maximise average profits, minimise the variance of profits whilst keeping the volume of loans and deposits as close to each other as possible.[8] The dynamic programming problem for the bank is to set the deposit interest rate, the loan interest rate, the monitoring intensity and the covenant threshold such that the demand for loans will bring the volume of loans into line with the volume of deposits but bearing in mind that these loan terms and the resulting distribution of entrepreneurs will be the state variable in subsequent periods. Since the model has a Markovian structure, the solution will be state-contingent policy rules. The initial conditions, the agents actions, the bank's policy functions and the aggregate shock process will determine the evolution of the economy through time. Since individual agents are assumed to be making decisions based on their state-contingent rational expectations, their beliefs will have to be fulfilled on average.

The third step of the solution is to realise that agents' choices and expectations and the bank's policy rules are an enormous fixed point problem. Solving the model computationally involves a search algorithm that finds agents' choices and expectations that are mutually consistent with a given set of bank policy rules and then repeats this operation over different bank rules until the objective function is maximised.

The following subsections explain the optimal behaviour of the entrepreneurs and inventors and the resulting portfolio dynamics for any specified bank policy rules.

2.1. Entrepreneurs

Depending on the idiosyncratic profitability state, a, an entrepreneur chooses between default, orderly exit and continuation. If he decides to default to escape losses, he pays B next period and switches to being an inventor. The discounted value of defaulting given information set Ω is thus

$$V_B(\Omega) = \beta\{E[V_I(a', \Omega'; .) \mid \Omega] - B\} \tag{1}$$

where the value function of an inventor is denoted $V_I(a, \Omega, V_E)$.

[8] These three objectives cannot, of course, be separately maximised. They are jointly summarised in one continuous objective function with implicit trade-offs between the three elements. Nothing of any importance depends on the precise form of the objective function.

If he chooses orderly exit from production, he absorbs current losses, pays liquidation costs L next period and also enters next period as an inventor. The value of orderly exit in state a given Ω and the bank interest rate rules is

$$V_X(a,\Omega) = q(a) + z - \rho(\Omega) + \beta\{E[V_I(a',\Omega';.) \mid \Omega] - L\} \tag{2}$$

The remaining option is to continue in production next period. Naturally the conditional value of continuing in production is to receive current payoffs and so the discounted expected value of being an entrepreneur in the next period is

$$V_C(a,\Omega) = q(a) + z - \rho(\Omega) + \beta\{E[V_E(a',\Omega') \mid a,\Omega]\} \tag{3}$$

where the value function of an entrepreneur is denoted $V_E(a,\Omega;V_I)$.

Given the options available, the value of being an entrepreneur at the moment the shock is revealed is:

$$V_E(a,\Omega;V_I) = \max\{V_B(\Omega), V_X(a,\Omega), V_C(a,\Omega)\} \tag{4}$$

There is a natural ordering of the choices facing an entrepreneur. Bankruptcy costs will be assumed to be sufficiently large that entrepreneurs only choose this form of exit when facing a very bad profitability state. It is straightforward to see from equations (1) and (2) that entrepreneurs will default for all values of $a < a_\delta(\Omega)$ where

$$q(a_\delta(\Omega)) + z - \rho(\Omega) = \beta(L - B) \tag{5}$$

thus defines the state-contingent default threshold.

The threshold for orderly exit, $a_X(\Omega)$, which also depends on the information set, results from a comparison of V_X and V_C. The only tricky aspect of this problem is the conditional expected value of being an entrepreneur next period, $E[V_E(a',\Omega';V_I') \mid a,\Omega]$. Consider first an entrepreneur with profitability above the state-contingent loan covenant threshold, $a \geq a_T(\xi(\Omega))$. In this case the entrepreneur faces no risk if the bank randomly chooses to monitor his loan, so we can ignore the role of the bank and

$$E[V_E(a',\Omega';.) \mid a \geq a_T(\xi(\Omega)),\Omega] = \int_A \int_{\Omega'} V_E(a\prime,\Omega';.)J(d\Omega',\Omega)F(da\prime,a)$$

where $J(\Omega',\Omega)$ describes the evolution of the state variables and subsumes the stochastic process of z and the endogenous transition of θ. The calculation is more complex for an entrepreneur with $a_\delta(\Omega) < a < a_T(\xi(\Omega))$. In this case, if the entrepreneur decides to continue and escapes monitoring (with probability $(1-\varphi(\Omega))$), then he gets the conditional expected value of being an entrepreneur in the next period. If the entrepreneur tries to continue but is monitored then the loan is recalled by the bank, the project is shut down and he involuntarily reverts to being an inventor. Therefore for $a_\delta(\Omega) < a < a_T(\xi(\Omega))$

$$\begin{gathered} E\left[V_E(a', \Omega'; .) \mid a_\delta(\Omega) < a < a_T(\xi(\Omega))\right] = \\ (1-\varphi(\Omega)) \int_A \int_{\Omega'} V_E(a', \Omega'; .) J(d\Omega', \Omega) F(da', a) \\ + \varphi(\Omega)\left(E\left[V_I(a', \Omega'; .) \mid \Omega\right] - L\right) \end{gathered}$$

The voluntary exit threshold $a_X(\Omega)$ is the value of current period profitability at which an entrepreneur is indifferent between continuing or exiting voluntarily. Some simple cancelling defines $a_X(\Omega)$ as

$$\int_A \int_{\Omega'} V_E(a', \Omega'; .) J(d\Omega', \Omega) F(da', a_X(\Omega)) = \int_A \int_{\Omega'} V_I(a', \rho(\Omega'); .) J(d\Omega', \Omega) G(da') - L \tag{6}$$

These equations look highly complex but are actually intuitively quite simple. Each period the agent has three options - continue, exit in an orderly fashion or default - and the bank has the option of demanding repayment if the project is known to breach the covenant. These four options partion the profitability space A into four regions based on three thresholds - the default threshold $a_\delta(\Omega)$, the voluntary exit threshold $a_X(\Omega)$ and the covenant threshold $a_T(\Omega)$. The thresholds are determined by the points at which a rational, forward-looking agent is indifferent between two naturally adjacent options. The thresholds vary with the information set because this conditions the bank's policy rules, $\psi(\Omega)$. Note that the policy rules affect the value functions of inventors as well as entrepreneurs because each is a function of the other.[9] So changes in Ω have quite subtle effects on the various thresholds because the inside and outside options are both changing at once.

[9] Given that there is switching between being an entrepreneur and an inventor and vice versa, the value function of each is a function of the other. This interdependency is recorded in the initial definitions but left implicit in all subsequent notation. The proof of uniquely consistent value functions is in the appendix of Penalver (2014).

2.2. Inventors

Each period, an inventor receives interest on her deposit $\tau(\Omega)$ and an idea with profitability index a. The inventor chooses between paying S and entering production next period with profitability index a' given $F(da\prime, a)$ and loan terms $\psi(\Omega')$ or waiting for another draw from $G(a)$ next period. and $E\left[V_I(a', \Omega'; .) \mid \Omega\right] = \int\limits_A \int\limits_{\Omega'} V_I(a' \rho\left(d\Omega' \mid \Omega\right); .) G(da')$. The value of being an inventor is conditioned on the deposit rate and the expected loan terms available on subsequent entry, which given the assumed policy rule, are functions of Ω.The value function of an inventor is consequently

$$V_I(a, \Omega; V_E) = \tau(\Omega) + \beta \max \left\{ \begin{array}{c} \int\limits_A \int\limits_{\Omega'} V_E(a\prime, \Omega'; .) J(d\Omega', \Omega) F(da\prime, a) - S, \\ \int\limits_A \int\limits_{\Omega'} V_I(a', \rho\left(\Omega'\right); .) J(d\Omega', \Omega) G(da') \end{array} \right\}$$

giving an entry threshold $a_E\left(\Omega\right)$ defined by

$$\int\limits_A \int\limits_{\Omega'} V_E(a\prime, \Omega'; .) J(d\Omega', \Omega) F(da\prime, a_E\left(\Omega\right)) - S = \int\limits_A \int\limits_{\Omega'} V_I(a', \rho\left(\Omega'\right); .) J(d\Omega', \Omega) G(da')$$

2.3. Equilibrium

Define $H([0, a))$ as the measure of entrepreneurs at the *end* of each period with profitability index in the interval $[0, a)$. And to simplify notation define $FH(a') = \int_A F(a\prime, a) H(da; .)$, which is the measure of firms entering the interval $[0, a')$ given the distribution over a in the previous period.

With the behavioural assumptions of the model and denoting D as the measure of inventors with their funds on deposit at the start of each period, a law of motion for the distribution of entrepreneurs can be defined by:

$$H'([0, a'); \Omega, .) = D \int_{a_E(\Omega)}^{a'} G(a) + \int_{a_X(\Omega)}^{a'} FH(da') - \varphi(\Omega) \int_{a_X(\Omega)}^{a' < a_T(\Omega)} FH(da') \quad (7)$$

The first term measures how many agents enter at profitability levels below a' given state Ω. The middle term measures how many entrepreneurs evolve into a profitability state above the voluntary continuation threshold. The third term eliminates those entrepreneurs closed down by the bank because they are monitored and have their loan recalled. Defaulting entrepreneurs are implicitly removed because they are already below the lower truncation of the distribution at $a_X(\Omega)$. Since each entrepreneur borrows one unit of capital, $H(A;.)$ is the measure of the volume of loans outstanding at the end of each period.

The following expression describes the bank's profits each period:

$$\Pi(\Omega) = \rho(\Omega) \int_{a_\delta(\Omega)}^{1} FH(da') - \int_{0}^{a_\delta(\Omega)} \lambda(a', \Omega) FH(da') - \varphi(\Omega) m \int_{a_X(\Omega)}^{1} FH(da') - \tau\,(\Omega)\, D$$

The first term is the loan interest paid by non-defaulting borrowers. This depends on how many firms evolve into a profitability state above the default threshold and the loan interest rate $\rho(\Omega)$. The second term deducts state-contingent loss given default assuming that the bank closing the project is inefficient (determined by the function $\lambda(a', \Omega)$). The third term subtracts the costs incurred in monitoring the continuing firms at a rate $\varphi(\Omega)$. The final term deducts the deposit interest paid and thus depends on the deposit rate $\tau\,(\Omega)$.

The invariant model in Penalver (2014) is nested within the current framework by setting the exogenous aggregate shock to zero and assuming the deposit rate is fixed exogenously at $\bar{\tau}$. An invariant, or steady-state, distribution is naturally one in which

$$H'(A, \overline{\rho}) = H(A, \overline{\rho}) = \bar{H}(A, \overline{\rho})$$

Propositions 1 and 2 of Penalver (2014) state that for a given monitoring intensity φ and covenant threshold ξ, there is a unique loan interest rate $\overline{\rho}$ that delivers an invariant distribution satisfying the balance sheet constraint

$$\bar{H}(A, \overline{\rho}) = D = \frac{1}{2}$$

Since there is a measure 1 of agents each with a unit of capital, the balance sheet constraint requires half the agents to be depositors and half the agents to be entrepreneurs at the end of each period. In the invariant version of the model, there is no need to condition on state variables and the entry and exit thresholds are also constant. An example of an invariant distribution of projects along the profitability index a is illustrated in Figure 2. It is easy to see the influence of the three behavioural thresholds on the distribution. Below a_X, there are no entrepreneurs in the distribution at the end of each period because they have either defaulted or exited voluntarily. Between a_X and a_T there are entrepreneurs that want to continue but are in breach of the loan covenant and thus at risk of having their loan recalled. Entrepreneurs in this region only survive if the bank does not monitored them. a_E marks the threshold at which it is just preferable to enter rather than wait another period. There is a concentration of entrepreneurs just above this level.

Figure 2: Invariant distribution of firms on the bank's balance sheet

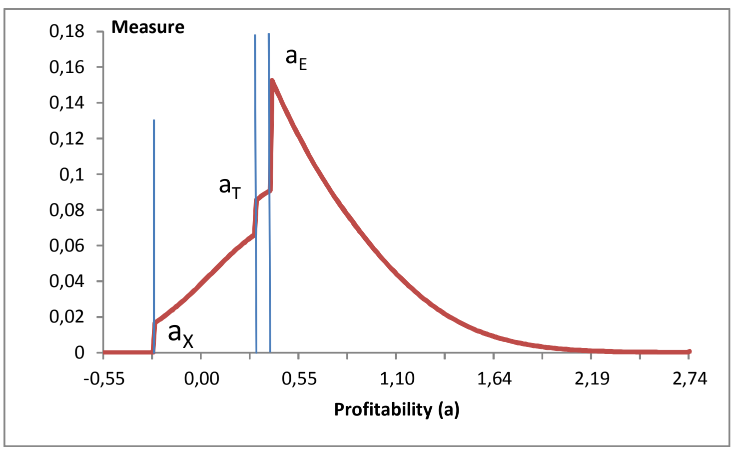

Figure 3 illustrates the one period transition of the distribution in Figure 2 with the invariant distribution overlaid. (This illustrates the expression $FH(a)$.) Looking from right to left, one can see that the upper tail of the distribution is entirely driven by the presence of a small number of existing entrepreneurs experiencing positive shocks. Since on average entrepreneurs with positive profitability experience a reversion towards the mean (of zero), there is a noticeable deterioration in the average quality of existing entrepreneurs - the distribution melts to the left. The distribution is refreshed by the entry of new entrepreneurs clustered above the entry threshold. Moving further to the left, a number of entrepreneurs fall below the threshold a_T but above a_X. These are the entrepreneurs that want

Figure 3: One period transition of bank balance sheet

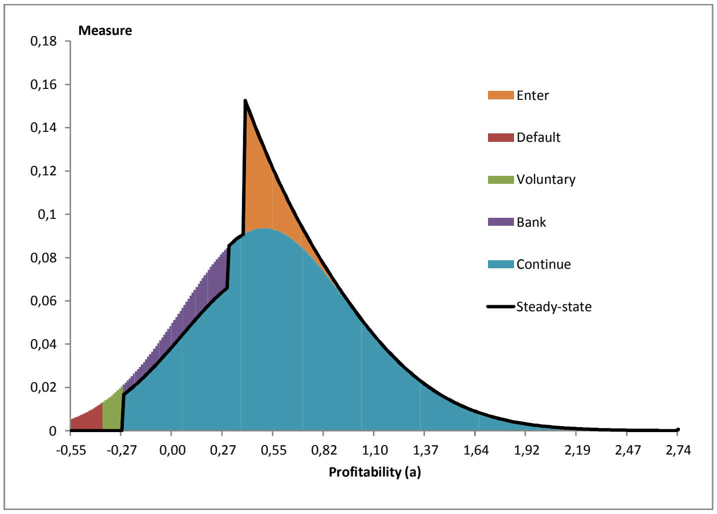

to continue but are at risk of having their loans recalled if the bank monitors them because they are in breach of the loan covenant. φ proportion of these entrepreneurs are monitored and exit and $1-\varphi$ are able to continue. Moving further to the left, there are entrepreneurs that fall below a_X but above a_δ and exit voluntarily. Finally, there is a portion of the distribution that falls below a_δ and defaults.

3. Solving the bank's dynamic problem

The solution to the bank's problem with only idiosyncratic risk was a fixed vector of loan terms. From the perspective of the bank, everything about that problem was deterministic and hence why it is the steady state of the model with aggregate shocks. With unchanging loan terms, the only thing that the agents had to forecast was their future individual state. Everything changes with an aggregate shock. The bank's loan portfolio is now stochastic and history dependent; the bank's behaviour is aggregate state contingent; and everyone will have to forecast what the aggregate states will be as well.

The exogenous shock process z is straightforward to forecast because it is a known AR(1) process with normal errors. The distribution of $z'|z$ in a discrete state space is just the row of a matrix.

The endogenous aggregate state variable θ is much harder. Equation (7) shows that the transition of the balance sheet is a complex function of the distribution over idiosyncratic profitability states. Part of this transition is under the control of the bank through its current choices, $\psi(\Omega)$, but the distribution is also affected by the history of past decisions. So in principle, θ should correspond to H. However, since A is a compact set of real numbers, H is an infinite dimensional object. As in Krusell and Smith (1998), to get traction with this problem, it is necessary to approximate H with a suitable summary statistic which in this case will simply be $\frac{H(A)}{\bar{H}(A)}$, ie the ratio of the volume of loans outstanding relative to the target. Henceforth $\theta \equiv \frac{H(A)}{\bar{H}(A)}$ will be referred to as "excess loans".

It was noted earlier that the information set Ω contains z and θ_{-1}. In other words, agents know the current exogenous aggregate state but only know the level of excess loans from the period before. In other words, the agents and the bank do not know the aggregate endogenous state that will arise out of the decisions that they are currently taking. This is not because agents are not fully rational. Indeed, they will be making accurate forecasts of what the endogenous state will be given the information that they have. In principle the bank ought to be able to rely on the monitoring process from the previous period and the same law of large numbers assumption plus knowledge of the entry process to estimate precisely the full distribution. But the spirit of the exercise is that agents and the bank do not have complete real time information. Monitoring is explicitly costly and it is not difficult to assume that information processing costs are high too.[10]

Agents are assumed to make forecasts for the aggregate state vector as a function of current observables,

$$E\left[\Omega' \mid \Omega\right] = \sigma\left(\Omega\right)$$

It will be further assumed that the forecasting function for the endogenous aggregate state and the bank policy rules are all linear. As a result, the expected loan

[10] In a sense, too, this is just a question of an arbitrary timing convention. If the distribution was defined at the start of the current period rather than the end, then decisions taken in the current period are based on current period information and forecasts of the consequences for next period's distribution. This looks much closer to the framework of Krusell and Smith.

interest rate next period, for example, will be a linear function of the information set:

$$E\left[\rho' \mid \Omega\right] = E\left[\rho\left(\Omega'\right) \mid \Omega\right] = \rho\left(\sigma\left(\Omega\right)\right)$$

with a similar structure for all the other bank policy rules. All agents have the same information set and all are rationally forward-looking, so all agents will have the same forecasts of future states and thus the same expectations about the path of loan interest rates and deposit rates. As a result, all agents will have a common view on the expected value of being an inventor and an entrepreneur.

For any given set of bank rules and forecasting function (whether accurate or not), we can solve for the behaviour of the agents which is completely summarised by the entry and exit thresholds. And for any given distribution of firms (and loans) from the previous period, we can solve the transition equation for the distribution. Thus for a given starting distribution and a sequence of aggregate shocks $\{z_t\}$, there is a unique path for the evolution of the bank's balance sheet and its profits. Clearly, for arbitrary bank policy functions one side or the other of the implied balance sheet could grow until it absorbed the full measure of capital and this would not be a feasible equilibrium. So the only acceptable rules are those that deliver a dynamically balanced measure of loans and deposits (at least to satisfactory precision). But there may be more than one (indeed infinite) set of rules that satisfy this dynamic balance sheet constraint. Different rules, though, imply different paths for profits with associated means and variances. It will be assumed that the bank seeks to maximise mean profits, minimise the variance of profits subject to minimising the deviation from the balance sheet constraint.[11] These multiple objectives are captured in a single objective function.

So the ultimate solution to the model is a fixed point in which:

- the bank follows linear policy rules which, given the behaviour of depositors and borrowers, and the characteristics of the common shock process, maximises the objective function;

- agents' expectations are model consistent; and

[11] Since balance sheet constraints are binding, one could assume that the bank holds a small buffer of liquid assets such as government bonds which pay the same as the deposit rate. The balance sheet constraint is met exactly by buying or selling this buffer stock. So the objective function of the bank is to keep this buffer stock as small as possible.

- depositors and borrowers act optimally given their heterogeneous circumstances, their expectations about the future paths of their private states and the aggregate state vector, and the policy rules of the bank.

The solution is found using numerical methods and the technical details are relegated to the appendix. The qualitative properties of the solution are general and are not parameter dependent. A sketch of the solution algorithm is as follows:

- Step 1: Guess parameters for the linear bank policy rules as functions of the information set, Ω. The rules are parameterised by the steady state loan rate and deviations in the state variables from their steady state levels. The rules for the loan interest rate, the monitoring rate and the covenant threshold are respectively:

$$\rho = \overline{\rho} + \rho_1 z + \rho_2(\theta_{-1} - \bar{\theta})$$

$$\tau = \overline{\tau} + \tau_1 z + \tau_2(\theta_{-1} - \bar{\theta})$$

$$\varphi = \overline{\varphi} + \varphi_1 z + \varphi_2(\theta_{-1} - \bar{\theta})$$

$$\xi = \overline{\xi} + \xi_1 z + \xi_2(\theta_{-1} - \bar{\theta})$$

- Step 2: Guess parameters for the linear forecasting rule for the agents for the endogenous variable θ as a function of Ω.

$$E[\theta] = \sigma_1 z + \sigma_2(\theta_{-1} - \bar{\theta})$$

- Step 3: Form a grid of a, z and θ. Calculate the value functions for inventors and entrepreneurs conditional on a and Ω and the forecast rule from Step 2 and the bank policy rules from Step 1. On the basis of these value functions, find the entry and exit thresholds for each grid point of Ω for each cohort.

- Step 4: Simulate the economy for 10,000 periods based on these rules, throw away the first 1,000 observations and estimate a new forecasting rule using ordinary least squares.

- Step 5: Repeat Steps 2 to 4 until the revision to the forecasting rule is reasonably small.

- Step 6: Calculate the loss function over the estimation range and repeat Steps 1 to 5 with a different bank policy rules using a search algorithm until the objective function is maximised.

Once this sequence of loops is completed, the solution is the profit-maximising fixed point of the problem. Given the number of steps in the algorithm, the two-dimensional aggregate state vector, the need to have a very granular idiosyncratic space so that changes in the rules result in changes in agents' behaviour (or else there are large flat regions over the minimisation surface) and the need to keep the run time to a reasonable length, many of the approximations are crude and the grid for the endogenous state variable is small (15 nodes). After all the number crunching, the optimal fixed interest rate rule for a simulation of the model is

$$\rho = 0.0675 + 0.5187z + 0.1698(\theta_{-1} - \bar{\theta}) \tag{8}$$

the deposit rule is

$$\tau = 0.035 + 0.7063z + 0.1063(\theta_{-1} - \bar{\theta}) \tag{9}$$

the monitoring intensity rule is

$$\varphi = 0.3067 - 0.4167z + 0.1642(\theta_{-1} - \bar{\theta}) \tag{10}$$

and the covenant threshold rule is

$$\xi = -0.0286 - 0.0833z + 0.0283(\theta_{-1} - \bar{\theta}) \tag{11}$$

In each case the first term is the value in the steady-state model.

The third term in each case captures the optimal response to excess loans the previous period. In each policy rule this is a positive value as might be expected. Since the loan distribution is slow moving because of its history dependence, excess loans the previous period will likely imply excess loans in the current period without any bank response. Since the objective is to try to keep excess loans to a minimum, it is sensible to lean against it with the instruments available. That all co-efficients are positive indicates that the optimal response is to use all instruments together. *Ceteris paribus*, raising the loan interest rate, raising the deposit rate, increasing the monitoring intensity and raising the covenant threshold, will all dampen the demand for loans which, since those not borrowing are lending, automatically increases the supply of deposits. But there is also a relationship between the instruments. Increasing the loan interest rate will also increase the default rate for any given z, so it is also prudent for the bank to monitor more intensely and tighten the covenant threshold as a precaution.

The second term captures the optimal response to the aggregate shock. Naturally the loan interest rate increases when the aggregate shock is positive in anticipation of increased loan demand. But notice that it does not increase one-for-one which would neutralise the effect of the aggregate shock on profits. In other words, the bank does not fully insure firms against the aggregate shock, which in turn implies that firm profits are pro-cyclical. As a direct consequence, default risk is also counter-cyclical, ceteris paribus, and thus the bank can afford to reduce the monitoring rate and loosen the covenant threshold. This relaxation in credit standards is optimal partly because it economises on costly monitoring when this is relatively unimportant but also because it allows the bank to increase the loan interest rate by more than it otherwise would be able to. Positive times are a good opportunity to trade off standards for credit spread. The reason why the bank increases loan interest rates less than one-for-one is because the deposit rate also moves. In fact, it moves more than the loan interest rate, so credit spreads narrow in equilibrium. This aggressive use of the deposit rate occurs entirely because of the objective of the bank to reduce the *variance* in profits. If this variance objective is turned off, so that the bank only cares about balance sheet constrained profit maximisation, then the optimal policy is to *reduce* the deposit rate during positive shocks and increase the loan interest rate by more than one-for-one. The logic behind this effect is that loss given default is procyclical, so a profit maximising bank would attempt to orchestrate a pro-cyclical default rate. Since this is clearly not a plausible outcome - deposit rates and default rates are procyclical - and profit smoothing is a reasonable objective, this is the equilibrium presented.

On a related point, if the aggregate savings rate was acyclical (ie constant) and the bank had no variance objective, then the optimal policy would indeed be to move the loan interest rate by exactly one-for one with all other co-efficients zero. To see why, notice that if the loan interest rate moved exactly one-for-one, then the aggregate shock would be completely neutralised from the point of view of the individual agents. The aggregate model would collapse to the idiosyncratic one for the agents, and the entry and exit thresholds would be constant. This would completely replicate the steady-state equilibrium and thus satisfy the balance sheet constraint at all times. Since the steady state itself is already the result of a static maximisation problem (Penalver (2014)), it will also be the one that maximises profits on average. The bank, however, fully absorbs aggregate profit shocks and thus has maximal variance. It is the presence of procyclical aggregate savings or the variance objective which disturbs this as a possible equilibrium.

The expectations function is

$$E\theta_t = \bar{\theta} - 0.0552z + 0.8094(\theta_{-1} - \bar{\theta}) \tag{12}$$

The third term indicates that excess loans are expected to be smaller but the same sign as in the previous period, so they stabilise gradually. The second term indicates that excess loans are *negatively* related to the aggregate shock. This term is the net effect of two opposing forces. On the one hand, a positive aggregate shock is on average positive news for enterpreneurs and increases loan demand. On the other hand, savings are assumed to be pro-cyclical. The negative co-efficient indicates that the second effect dominates the former.

The objective of the bank in this model is to maximise profits whilst minimising deviations between the volume of loans and the volume of deposits. An important question, therefore, is how successful is the bank in achieving its balance sheet objective. It is useful to consider as a benchmark what happens if the bank does *not* vary the loan interest rate, the deposit interest rate, the monitoring intensity or the loan covenant threshold over time but keeps them at their steady state values. In this case, the relative size of the two sides of the balance sheet moves around over time, importantly, it does *not* exhibit any systematic drift. In other words, the balance sheet condition is met on average even with completely unresponsive rules. The variance, however, is quite wide at around 4.7%. So if it is optimal to have active policy rules, it must be because it reduces this variance. And this is indeed the case, with state-contingent rules lowering the variance to around 0.2%.

Before turning to a discussion of the dynamics, it is important to consider briefly the accuracy of the model. As is well known, a draw-back of heterogeneous agent models solved using numerical methods is that it is impossible to know (or bound) the true solution and therefore impossible to gauge the accuracy of the approximation. Den Haan (2010) argues that using the R^2 of the forecast equation, one of the options proposed by Krusell and Smith (1998), is neither a necessary nor sufficient statistic for measuring accuracy, primarily because the true value is used as the updated lagged dependent variable each period. A high R^2 therefore does not give any indication whether long range forecasts are accurate. Den Haan

Figure 4: Forecast accuracy plot

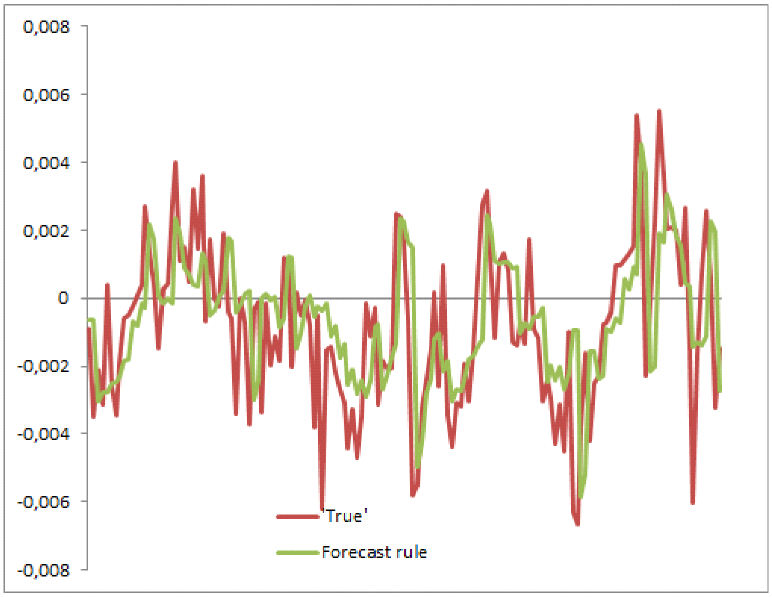

(2010) instead suggests comparing iterations of the original forecast equation with the model output using a new draw of aggregate shocks. A plot of this iteration against the true value can not only tell whether the forecast is accurate on average but also whether there are systematic errors.

Figure 4 illustrates a representative segment of the iteration of the forecast equation and the model output ('True').[12] The True value is highly volatile so it is unsurprising that the forecast rule cannot completely track it. Indeed, there is evidently a lag between the forecast rule and the true value. But overall, this is a highly encouraging fit.

4. Aggregate dynamics

[12]The forecast equation actually uses deviations from the steady-state value as the dependent variable.

This section describes the behaviour of the key credit cycle variables over a (common) representative sample of the simulation.[13] The paths of the control variables are, of course, simply a function of the aggregate state variables from the policy functions above. In the actual simulation, the endogenous variable is only represented at grid points and so there are only a fixed number of values for each of the control variables. But the simulation always keeps track of the actual value of θ. The Figures below record the path of the control variables implied by the policy rules and the actual value of θ.

Figure 5 illustrates the path of the two aggregate state variables. This clearly shows that the endogenous state - excess loans - is counter-cyclical. At first glance, this is a somewhat surprising result since the (implicit) demand curve for loans is procyclical. Two forces are at work. Firstly, by assumption, aggregate savings are procyclical so both sides of the balance sheet are naturally expanding during positive shocks. Secondly, although the aggregate shock process is persistent, it is also mean reverting. An inventor observing a positive aggregate shock in the current period expects a slightly less positive aggregate shock state in the following period. Since the inventor knows that the bank's interest rate setting rule moves less than one for one, this implies a positive and gradually decaying profit path. If there were no entry and exit costs, so that firms opportunistically entered and exited to take advantage of current period shocks, then expectations beyond the next period are not relevant. Positive entry and exit costs, however, dampen the enthusiasm for an inventor to enter for a given positive shock. Note also that agent behaviour is conditioned on the bank policy rules which are parameterised (implicitly) on the basis of the same decaying expectations. Thus both agents and the bank are positively surprised when a positive aggregate state becomes more positive, as for example occurs towards the right hand side of the Figure. A sequence of positive surprises, thus leads to a surprise fall in excess loans because the savings response dominates loan. However, this dominance is really very small. The range of the right hand axis is between $\pm 0.4\%$. It is also worth observing at this point that the endogenous state variable exhibits very low frequency variation as well depending on whether the recent history of aggregate shocks has been on average above or below the mean. It can be seen from the Figure that excess loans is on average negative over this particular segment of the simulation.

[13] All data in this section have been smoothed using Henderson trend weights to remove high frequency noise in the time series due to various approximations in the numerical simulation. Nothing of substance is hidden by this smoothing.

Figure 5: Simulated aggregate state variables

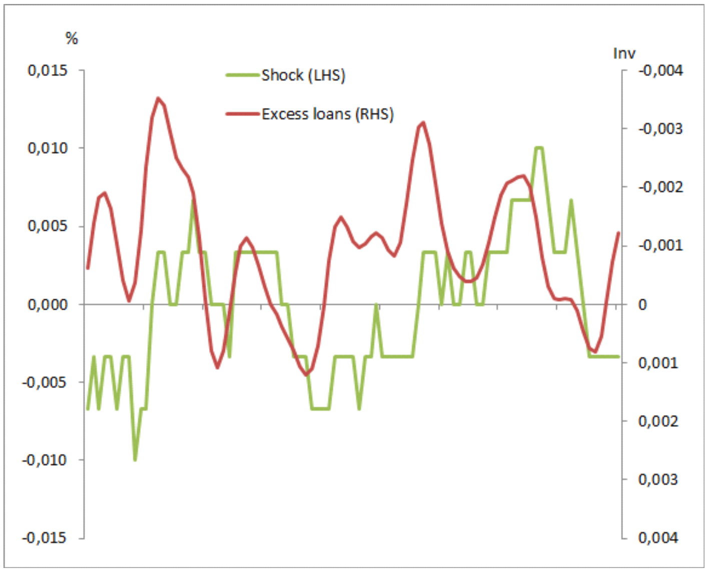

Having established the paths of the aggregate state variables, we can now analyse the behaviour of the control variables. Figure 6 plots the path of the loan interest rate over the cycle which is clearly procyclical.

Figure 7 decomposes the variation of loan interest rates from their steady-state value into the contribution of the two aggregate state variables. Given the positive co-efficients on both variables in the bank loan interest rate setting rule, equation () and that the two aggregate state variables move in opposite directions, then it is straightforward to see why the contribution of the endogenous state offsets the contribution of the aggregate shock. But the latter effect clearly dominates the former, a pattern that occurs with all the other control variables.

Figure 8 plots all of the control variables over the simulation range. Loan interest rates and deposit interest rates are procyclical but monitoring intensity and the covenant threshold are counter-cyclical. Why does the bank use the monitoring intensity and covenant threshold in tandem? Of course it is intuitive that should do so, rather than, for example, move in opposite directions. But nor is immediately obvious that it wouldn't be better to keep one instrument fixed and just move the other. For example, why not keep the monitoring rate

Figure 6: Loan interest rates and the aggregate shock

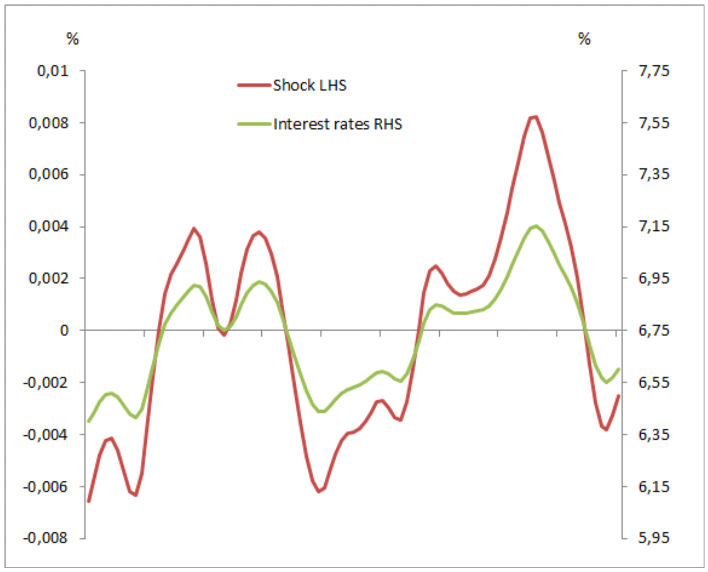

Figure 7: Decomposition of the loan interest rate path

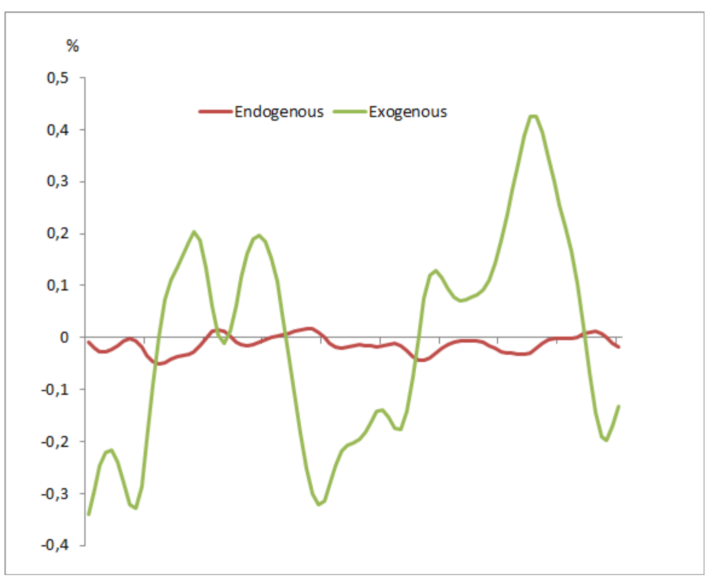

fixed and just move the loan covenant? Consider a negative shock that prompts the bank to wish to tighten credit standards. Raising the profit threshold gives the bank the right to eject more firms at the lower end of the distribution and thus will reduce default risk. But the higher it sets the threshold, the further the marginally ejected firm is from default, so it is a decreasingly effective instrument. And with a fixed monitoring rate, the bank is still not capturing all the high risk firms below that are actually below the threshold.

Consider instead that the bank keeps the covenant threshold constant but shifts the monitoring intensity. A higher monitoring intensity identifies more firms below the covenant threshold and this will reduce default risk. But at the same time, much of this additional monitoring effort is wasted since it will find firms that are above the threshold and thus allowed to continue. Since monitoring is expensive, this is on its own an inefficient way to reduce default risk. The best strategy is to use both instruments together.

Two further series of interest are the default rate and the entry and exit rate illustrated in Figures 9 and 10 respectively. The default rate is clearly counter-cyclical as one would expect. Since the loan interest rate moves by less than one-for-one with the aggregate shock, net profits are procyclical. Since the bankruptcy cost is not assumed to change over time, a larger fraction of entrepreneurs find themselves in the region in which it is better to walk away than liquidate in an orderly manner during negative shocks.

Finally, Figure 10 illustrates the entry and exit rates. The entry rate is pro-cyclical and the exit rate is counter-cyclical. The exit rate combines voluntary exit, loan recall and default. Loan recall is counter-cyclical through the change in monitoring intensity and the covenant threshold. With positive shocks, fewer firms are monitored and the threshold is less severe so the survival rate is higher and vice versa during negative shocks. Voluntary exit is also counter-cyclical. Since the loan interest rate moves by less than one-for-one, profits are counter-cyclical. And the aggregate shock is persistent so profit expectations are also counter-cyclical as well. As a result, more firms fall into the region in which it is better to liquidate than continue. Likewise, the entry rate is pro-cyclical because a greater fraction of inventors will expect to find it beneficial to enter. Note that the process for generating ideas for projects, $G(a)$, is constant so this procyclicality of entry has nothing to do with innovation.

Figure 8: Control variables over the simulation

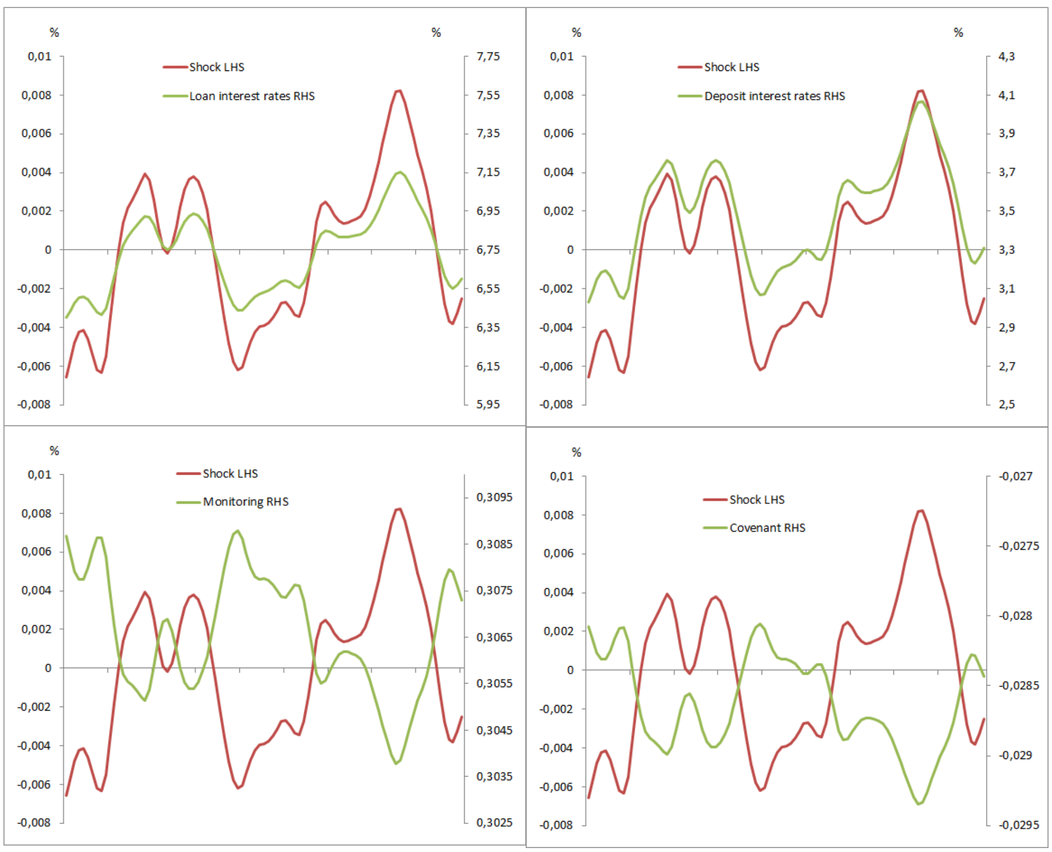

Figure 9: Simulated default rate

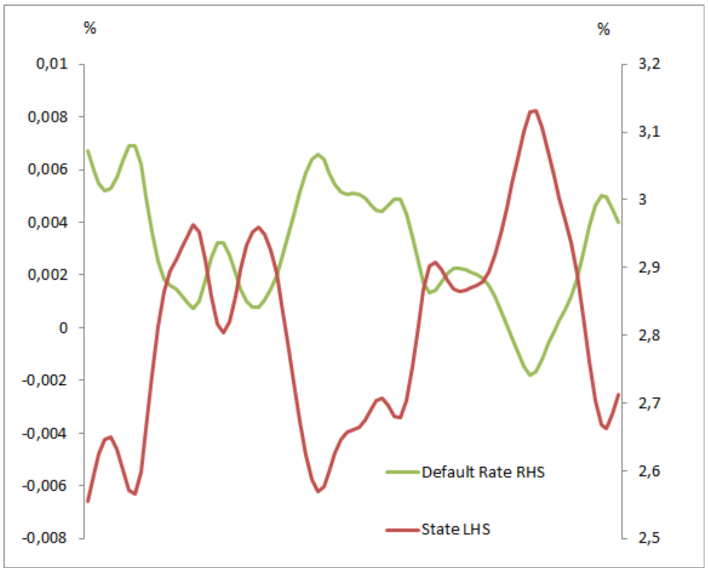

Figure 10: Entry and exit rates over the simulation

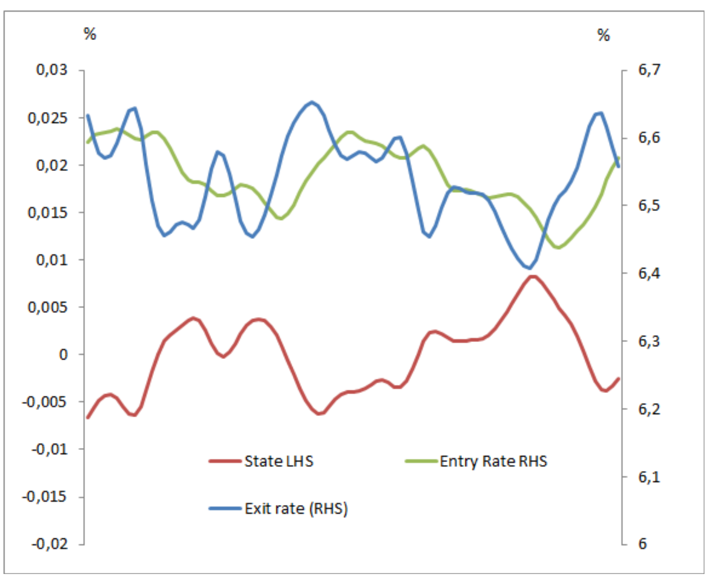

In a sense it is unsurprising that the entry rate is procyclical and the exit rate is counter-cyclical since by construction the measure of firms has to be procyclical because of the assumption of procyclical aggregate savings. With a higher level of savings during positive shocks, the volume of loans has to expand if the bank is to meet its balance sheet objective. But by assumption, all loans are of a fixed size, so all the adjustment has to take place at the extensive margin. In this model, there are two endogenous margins to work with and it is unsurprising that they share the burden.

A counterpart to the procyclical entry and exit rate is procyclical *endogenous* productivity. Figure 11 shows that endogenous productivity - the reallocation of resources between firms - is small relative to exogenous productivity (driven by the aggregate shock). As savings grow during periods of positive shocks, the entry rate increases. This occurs despite the fact that the marginal entrant during positive shocks has lower profitability (productivity) than the marginal entrant during negative shocks and that more marginal firms persist during positive shocks. Indeed, the marginal entrant has less than average productivity! To understand why endogenous productivity is procyclical, it is necessary to appreciate the subtle dynamics of the invariant distribution of the steady state. Recall that idiosyncratic shocks are mean reverting. If the initial distribution were left to replicate with no entry or exit, the distribution would eventually be degenerate to a mass point at the mean. But with exit, there is truncation from below of the distribution. So absent any entry, the distribution would both shrink in size and the average value would shrink. It is entry at a threshold above the mean that refreshes the distribution. The invariant distribution is one in which the entry process exactly compensates for the decay process. A temporary higher entry rate gives an additional positive nudge to the distribution.

5. Conclusion

Much of the popular debate on the credit cycle explains bank and borrower behaviour as a psychological cycle (see also Kindleberger (1978) or Minksy (1992)). As the business cycle picks up, investors select high quality projects and entrepreneurs borrow prudently. Asset prices start to rise and soon investors become greedy for returns. Memories of previous losses begin to fade and analysts assert that good times are set to continue. Investors start to take more and more risks and trade-off credit standards for yield. Greed becomes recklessness and then

Figure 11: Productivity over the simulation

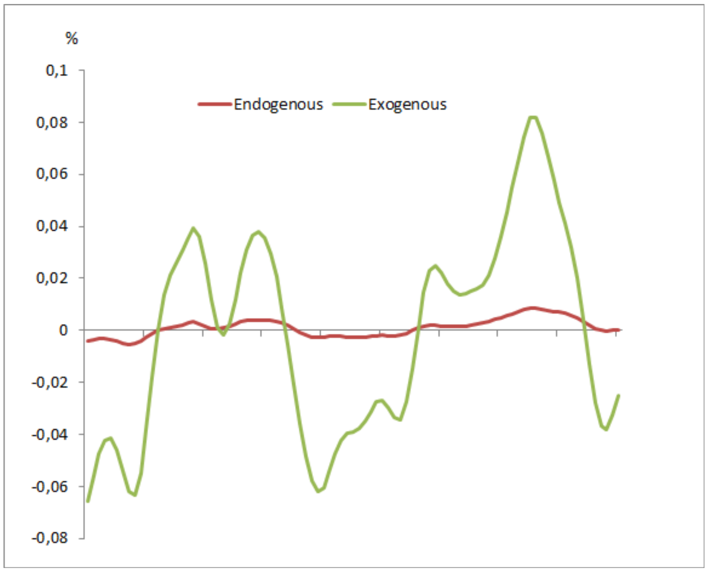

suddenly there is a panic and crash. Investors shut the stabe door after the horse has bolted by pulling back sharply on risky investments and drastically tighted their credit standards. Feast turns to famine and even the best quality firms have difficulty finding funding. Only gradually does confidence return and the cycle begin again. This is a compelling narrative. It is also a form of morality play which accords with the cultural perception of indebtedness as a form of sin.

This paper shows that such behaviour can instead be completely consistent with rational behaviour on the part of borrowers and banks. As long as there is some persistence in aggregate demand, it makes sense for banks and borrowers to change behaviour when the business cycle turns. This is true even when loans are made for long duration. If there were a fixed pool of savings, then the best option would be to adjust loan interest rates in line with the aggregate shock. Borrowers are insured against the aggregate shock and this neutralises its effect on the balance sheet of the bank. But if aggregate savings are procyclical, then unless the distribution of potential entrants shifts too (which of course it might), the relatively profitability of the marginal entrant must fall. Inducing the marginal entrant to borrow requires a relaxation in credit terms. This could be done solely by moving loan interest rates less than one-for-one with the shock. This shfts some of the aggregate shock risk onto borrowers and implies that default risk will be counter-cyclical. It thus makes sense for the bank to adjust its credit standards - lowering credit standards when default risk is expected to remain low - in return for yield.

The analysis in this paper suggests that the task of macroprudential policy-makers is tougher than it already appears. It is not enough to look at measures of credit standards to determine whether the banking system is acting imprudently because the benchmark of rationally forward-looking behaviour is also cyclical. The analysis also suggests that comparing standards with measures of growth will also not be entirely reliable. What matter are measures of demand and output relative to trend. In the model, everybody knows that there is no drift in aggregate productivity so the deviation is clear for all to see whereas in reality trend productivity is hard to gauge in real time. Nevertheless, macroprudential policy-makers are going to have to take a stand on these structural issues and then make a judgement on the appropriateness of credit standards relative to this assessment. This is clearly an immense challenge.

6. Bibliography

References

[1] Aiyagari, R. (1994), 'Uninsured idiosyncratic risk and aggregate saving', *Quarterly Journal of Economics*, vol. 109, pp.659-684.

[2] Barlevy, G. (2002), 'The sullying effect of recessions' Review of Economic Studies, vol. 69(1), pp. 65-96.

[3] Caballero, R. and Hammour , M. (1994), 'The cleansing effect of recessions' American Economic Review, vol. 84(5), pp. 1350-1368.

[4] Den Haan, W. (2010), 'Assessing the Accuracy of the Aggregate Law of Motion in Models with Heterogeneous Agents', *Journal of Economic Dynamics and Control* vol. 34(1), pp. 79-99.

[5] Gertler, M. and P. Karadi (2011), 'A model of unconventional monetary policy', *Journal of Monetary Economics*, vol. 58, pp. 17-34.

[6] Gertler, M. and N. Kiyotaki (2010) 'Financial intermediation and credit policy in business cycle analysis', in **Handbook of Monetary Economics** edited by B. M. Friedman & M. Woodford, edition 1, volume 3, number 3, Elsevier, October 2010.

[7] Hopenhayn, H. (1992), 'Entry, exit and firm dynamics in long run equilibrium', *Econometrica*, vol. 60, 5, pp. 1127-1150.

[8] Khan, A. and J. Thomas (2013), 'Credit shocks and aggregate fluctuations in an economy with production heterogeneity' *Journal of Political Economy* forthcoming.

[9] Kindleberger, C. (1978), **Manias, panics and crashes**, 6th edition, Palgrave Macmillan,

[10] Krusell, P. and A. Smith (1998), 'Income and wealth heterogeneity in the macroeconomy' *Journal of Political Economy*, vol. 106, pp.867-896.

[11] Lee, Y. and T. Mukoyama (2008), 'Entry, exit and plant-level dynamics over the business cycle' *Federal Reserve Bank of Cleveland Working Paper* 07-18R.

[12] McCall, J. (1970), 'Economics of information and job search' *Quarterly Journal of Economics*, vol 84, pp. 113-126.

[13] Meh, C. and K. Moran (2008), 'The role of bank capital in the propagation of shocks' *Bank of Canada working paper* 2008-36.

[14] Minsky, H. (1992), 'The financial instability hypothesis' Working paper no. 74, Levy Economics Institute, Bard College.

[15] Penalver, A. (2014), 'Bank credit standards and firm dynamics' Paris School of Economics mimeo.

[16] Schumpeter, J. (1939), 'Business cycles: A theoretical, historical and statistical analysis of the capitalist process', McGraw-Hill, New York.

[17] Tauchen, G. (1986) 'Finite state Markov chain approximations to univariate and vector autoregressions', *Economic Letters*, vol. 20, 2, pp. 177-181.

7. Appendix

The model takes place in $A \times Z \times \theta$ space, respectively the idiosyncratic state, the aggregate exogenous shock and the endogenous aggregate state. This space is approximated by a 3-dimensional grid $5001 \times 7 \times 15$. The idiosyncratic grid has to be highly granular for the effect of different behavioural rules of the bank to have observable effects on the entry and exit threshold. A coarse grid would severely limit the precision with which the rules could be calculated. The solution to the model relies on a simulation for a 10,000 periods based on a random sequence of aggregate exogenous shocks consistent with the transition probability $Z(z', z)$. $Z(z', z)$ is a Tauchen matrix approximation for an AR(1) with zero drift, an autoregessive co-efficient of 0.9 and standard deviation of 0.25. The shocks are evenly space between ± 1.74 standard deviations. The aggregate economy passes stochastically from gridpoint to gridpoint in $Z \times \theta$ space. Since the actually evolution of θ hits a gridpoint with probability zero, the economy is randomly assigned to one of the two nearest gridpoints with the probability based on the relative distance between the true value and the gridpoints.

Entrepreneurs and inventor receive draws from A so it is necessary to solve for the value functions V_E and V_I at each point on the $A \times Z \times \theta$ grid. This is a complicated system of value functions since each is dependent on the other and the transition probabilities. The idiosyncratic state also follows an AR(1) process and approximated by a Tauchen matrix, $F(a', a)$. $F(a', a)$ has zero drift, autogressive co-efficient 0.75, standard deviation 0.3 and the outer limits are ± 4 standard deviations.

The solution to the model is a fixed point in which:

- the bank follows linear policy rules which, given the behaviour of depositors and borrowers, and the characteristics of the common shock process, maximises the objective function;
- agents' expectations are model consistent; and
- depositors and borrowers act optimally given their heterogeneous circumstances, their expectations about the future paths of their private states and the aggregate state vector, and the policy rules of the bank.

The solution is found using numerical methods using the following steps:

- Step 1: Guess parameters for the linear bank policy rules as functions of the information set, Ω.
- Step 2: Guess parameters for the linear forecasting rule for the agents for the endogenous variable θ as a function of Ω.

$$E[\theta] = \sigma_1 z + \sigma_2(\theta_{-1} - \bar{\theta})$$

- Step 3: Form a grid of a, z and θ. Use iteration to solve the system of the value functions for inventors and entrepreneurs conditional on a and Ω and the forecast rule from Step 2 and the bank policy rules from Step 1. On the basis of these value functions, find the entry and exit thresholds for each grid point of Ω for each cohort.
- Step 4: Simulate the economy for 10,000 periods based on these rules, throw away the first 1,000 observations and estimate a new forecasting rule using ordinary least squares.
- Step 5: Repeat Steps 2 to 4 until the revision to the co-efficients on the forecasting rule is less than 1%. This is not very precise but since it is within a larger loop, I believe it to be a reasonable trade-off between accuracy and time.
- Step 6: Calculate the loss function over the estimation range and repeat Steps 1 to 5 with a different bank policy rules using the Nelder-Mead search algorithm until the objective function is maximised. The objective function is specified as

$$OBJ = \overline{\Pi_t} - \gamma * var(\Pi_t) - var(e_t)^{0.7}$$

  where the first term is average bank profits over the truncated simulation range, the second term is variance in profits and the third term is the average absolute error in the balance sheet constraint. The first term is almost completely irrelevant for the solution because the shocks are symmetric so the profit maximising co-efficients of the steady-state (taken from Penalver (2014)) are extremely close to those which maximise the dynamic model. The exponential on the third term is designed to reduce the importance of the balance sheet objective when it is very close to being satisfied. This is intended to capture the idea in the main text that a bank might have a small buffer of government bonds paying close to the deposit rate which it can use to cover temporary imbalances between deposits and loans. It is unlikely that a bank will reject strategies which increase average profits or reduce variance in profits just to reduce variance in this buffer.

In Step 6, the algorithm is searching over 8 co-efficients so the outer loop needs to be repeated many times. With this many co-efficients, the algorithm takes several days to solve even with judicious choice of starting values. How close the co-efficients are to their true values for some of the less significant parameters is of course difficult to say. [Starting from different initial co-efficients yields very similar results.]