Managing maturity transformation under aggregate uncertainty
Work In Progress

Adrian Penalver
Paris School of Economics

November 5, 2013

Abstract

When banks engage in maturity transformation, they repeatedly make long-term loan commitments funded by short-term deposits. With only idiosyncratic shocks to borrowers and depositors, a bank big enough to rely on the law of large numbers can satisfy this funding constraint every period with constant loan and deposit interest rates. With aggregate shocks as well, this is a much more challenging problem because how a bank resolves excess long-term loan demand in the current period affects the distribution it takes into the next which in turn affects the excess demand problem it will face then. This paper builds a model of this dynamic programming problem and solves it under various conditions. The paper shows how fixed rate loan contracts and aggregate savings behaviour can interact to affect the dynamics of aggregate output in a bank-intermediated economy.

1. Introduction

The ongoing economic difficulties in most of the world's advanced economies emphasises how important banks can be in the propagation of shocks. In response, numerous theoretical models have been developed to analyse how frictions in the financial system, and in particular banks, can affect the dynamics of economic output. The most common approach has been to extend the framework of Holmström and Tirole (1997) to a recursive setting so that the net worth of the banking system becomes the key state variable. Analysing frictions on the liability structure of banks' balance sheets has provided important theoretical and quantitative insights on the interaction between distress in the financial system and the real economy (Meh and Moran (2008), Gertler and Karadi (2009) and Gertler and Kiyotaki (2009)). But banks also face considerable (self-imposed) constraints on the asset side of their balance sheet through their willingness to make long term loan commitments. Long term loans have two effects: first, they leave the bank exposed to the performance of loans issued many periods ago and under potentially quite different conditions and second, they leave the bank with only limited room to manoeuvre in response to aggregate shocks. These two effects interact since how a bank reacts to balance sheet pressure in the current period becomes part of the historical legacy in the next. A bank's loan portfolio is a state variable under maturity transformation.

In most advanced economies, banks are the dominant financial intermediaries. As a result, the real economy is a counterpart to the balance sheet of the banking sector. How banks allocate long-term credit is equivalent, therefore, to how an economy allocates capital and the performance of this allocation maps to the credit risk of the banks' loan portfolios. Going in the other direction, aggregate constraints and the behavioural incentives of economic agents limit banks' ability to act. Understanding how banks solve this balance sheet management problem is crucial in analysing macroeconomic dynamics in bank-intermediated economies. The purpose of this paper is to examine this issue.

So what exactly is the bank balance sheet management problem? When banks engage in maturity transformation they repeatedly commit to long term loan contracts but with only a temporary supply of deposits. The potential fragility of this mismatch is probably one of the most widely studied phenomenon in economics (Diamond and Dybvig (1983)). But pure "sunspot" runs on banks are actually extremely rare even in countries without a deposit insurance scheme. The intuition for this stability is as old as credit banking itself. A big enough bank can rely on the law of large numbers to have a predictable inflow and outflow of deposits on one side of the balance sheet and predictable repayment of (and default on) loan principal on the other and thus have an accurate estimate of the amount of new long term loans it can extend. By judicious choice of the deposit and loan rate, a bank hopes to ensure that a steady state stock of long term loans is backed by a steady state stock of short-term deposits. If the only risks are the idiosyncratic ones faced by depositors and borrowers, then a big enough bank can regard the evolution of its balance sheet as completely deterministic. This is exactly what is assumed implicitly in standard two period models of banking with maturity transformation. Penalver (2013) develops the same idea in a recursive model using a dynamic entry and exit model (following Hopenhayn (1992)) where the result is a unique invariant distribution that meets the balance sheet constraint and which is governed by a constant deposit and loan interest rate. (For reasons which will become clear, that model will be referred to as the "steady-state model" in the remainder of this paper.)

This balance sheet management problem with maturity transformation becomes potentially a lot harder in the presence of aggregate exogenous shocks. A bank making a volume of new long term loan commitments in the current state can no longer by entirely confident that sufficient deposits will be available in future states. With stochastic shocks and long term loans, a bank is not only having to try to cope with the excess supply or demand for funds arising from its inherited balance sheet but is also trying to adjust in such a way that it minimises the potential complications for future periods. In this paper I extend the static framework of Penalver (2013) to a dynamic setting and solve this dynamic programming problem under 3 different conditions. In the first case, the bank has complete flexibility to adjust deposit and loan interest rates after observing the aggregate shock. In the second case, the bank commits to the loan interest rate at origination and can only adjust the deposit rate and the interest rate on new loans. In these first two cases, the aggregate savings rate does not respond to output. In the third case, aggregate savings do respond to the shock. These three cases yield starkly different results for aggregate output, default rates and aggregate productivity even though the underlying economic mechanisms are the same.

The structure of the paper is as follows. Section 2 sets out the general structure of the balance sheet management problem and summarises the solution of the steady-state model in which there is only idiosyncratic risk. Section 3 describes the case in which the bank varies the interest rate applicable to all loans - labelled the flexible rate case. Section 4 describes the case in which the bank varies the deposit rate and the interest rate on new loans - labelled the fixed rate case. Section 5 describes the same structure as Section 4 but with aggregate saving responding to aggregate output. Section 6 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 (2013), 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. Every agent is endowed with a unit of capital which they cannot add to or lose.[1] By carrying a single unit of capital through time whatever their circumstances substantially simplifies the analysis by making the equilibrium solely a function of the extensive margin. This assumption makes the distribution of *capital* across agents degenerate and 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). 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).[2]

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 unmodelled. 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 and 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. Since there are exit thresholds, this assumption implies that all projects will almost surely close at some future point.

[1] Agents might be able to store output or save by some other means but since the timing of consumption is irrelevant in the model, this restriction doesn't have any significant consequences.

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

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 $F(a', a)$. Projects are assumed to require 2 units of capital, so the inventor needs to borrow 1 unit from a bank as direct lending is ruled out by assumption. If an inventor decides not to enter, she remains an inventor next period and in the absence of any better alternative, deposits her unit of capital at the bank on which she receives an *endogenously* determined deposit rate τ.

The other group of agents are running projects and are called "entrepreneurs". Entrepreneurs make net profits per period $q(a) + z - \rho(i, .)$ where $\rho(i, .) \in \mathbb{R}^+$ is the loan interest rate paid by cohort i entrepreneurs 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. A cohort is defined by the aggregate state vector in the period they decided to enter and signed a loan contract. This contract specifies whether or not the loan interest rate varies over time. There are a discrete number of cohorts k and they are ordered so that $\rho(i, .) < \rho(j, .)$ for $i < j$. 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 losses, $q(a)+z-\rho(i, .)$ 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.

- "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.[3] Naturally, L and B are calibrated so that bankruptcy is prefered over orderly exit only in extreme circumstances.

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

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. So all exit decisions in the model are based on incentives not inability to pay.

Capital in this model is 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, the *ex ante* probability distributions over the bank's loan portfolio are identical to the distributions of *ex post* outcomes along all dimensions in steady state. The bank offers the following deposit and loan contracts.

- Deposits earn the current deposit interest rate τ. Deposits can be withdrawn at the end of any period.

- The loan contract specifies a loan interest rate rule for the entering cohort i, $\rho(i,.)$, a monitoring intensity φ (where $0 < \varphi < 1$) and a covenant specifying a minimum net profit level ξ.[4] 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 ξ must 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 φ gives the probability that any entrepreneur who chooses to remain in production is inspected.

The model in Penalver (2013) contains only idiosyncratic shocks and the problem for the bank is to find a loan interest rate and monitoring pair (ρ, φ) which

[4]In Penalver (2013), $0 \leq \varphi \leq 1$ and is determined endogenously by a profit maximisation condition. In future work, φ and ζ will be the subject of policy rules. To focus only on the role of the fixed loan interest rate, φ is set at the profit maximising level of the steady-state obtained in the paper cited above.

maximises profits subject to the constraint that the bank must have the same measure of deposits and 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 $(\overline{\rho}, \bar{\varphi})$ are stipulated in the contract or are a rational expectation. In the model in this paper there is an aggregate shock process (and φ is fixed) so constant interest rates will no longer satisfy the balance sheet constraint. Therefore what type of loan contract is signed and how expectations are formed are fundamentally important for the equilibrium. There are two aggregate shock states, called "low" and "high" and denoted $z_t = \{-\epsilon, \epsilon\}$ with time homogeneous Markov transition probabilities $Z(z', z)$. The aggregate shock is assumed to be persistent and, for expositionary purposes, symmetric. The aggregate shock is a common scalar shift in the gross return of all projects in that period.

This paper will examine two options for the loan contract. One contract will be labelled "flexible" because

$$\rho(i, .) = \rho(j, .) \qquad \forall i, j$$

so that all borrowers pay the same loan interest rate in any period. Clearly this means that existing borrowers pay the same interest rate as that offered to new borrowers. An alternative contract, labelled "fixed", has $\rho(i)$ constant for each cohort.

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 the following sections. 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 and the idiosyncratic shocks are drawn. The inventors get an idea from $G(a)$ and entrepreneurs get an update of their profitability according to $F(a', a)$.

2. The bank updates its interest rates which in this recursive setting are a function of the information set. The deposit rate is updated for all depositors according to $\tau(\Omega)$ and the loan rate is set for new or all borrowers (depending on whether we are in the fixed or flexible case) according to $\rho(\Omega, i)$. The set of control variables is discrete in order for the cohort set to be discrete.

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 φ and recalls the loans of all entrepreneurs found below the covenant profitability threshold.

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

Given this economic structure and the timing convention, the behaviour of the agents is as follows.

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

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

where the value function of an inventor is denoted $V_I(a, \Omega, V_E)$ and $E\left[V_I(a', \Omega'; .) \mid \Omega\right] = \int\limits_{\Omega'}\int\limits_{A} V_I(a'\rho(d\Omega' \mid \Omega); .)G(da')$. The value of being an inventor is conditioned on the deposit rate and the interest rate available on new loans but given the assumed policy rule, these are functions of Ω.[5]

[5] 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 (2013).

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

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

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

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

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

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

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

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 $a < a_\delta(i,\Omega)$ where

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

The threshold for orderly exit, $a_X(i,\Omega)$, which also depends on the cohort and 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: $E[V_E(a',i,\Omega';V_I') \mid a,\Omega]$. Consider first an entrepreneur with profitability above the loan covenant threshold, $a \geq a_T(i,\Omega)$. In this case the entrepreneur faces no risk if the bank randomly chooses to monitor, so we can ignore the role of the bank and

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

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 $a_\delta(i,\Omega) < a < a_T(i,\Omega)$. In this case, if the entrepreneur decides to continue and escapes monitoring (with probability $(1-\varphi)$), 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 $a_\delta(i,\Omega) < a < a_T(i,\Omega)$

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

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

$$\int\limits_{\Omega'}\int\limits_{A} V_E(a\prime, i, \Omega'; .)J(d\Omega', \Omega)F(da\prime, a_X(i,\Omega)) = \int\limits_{\Omega'}\int\limits_{A} V_I(a', \rho(\Omega'); .)J(d\Omega', \Omega)G(da') - L \tag{6}$$

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 breaches the covenant. These four options partion the profitability space A into four regions based on three thresholds - the default threshold $a_\delta(i,\Omega)$, the voluntary exit threshold $a_X(i,\Omega)$ and the covenant threshold $a_T(i,\Omega)$. The thresholds are determined by the points at which a rational, forward-looking agent is indifferent between two naturally adjacent options. The thresholds will differ across cohorts because an entrepreneur with a lower interest rate always has higher net profitability than an entrepreneur with a higher interest rate for any given profitability state a. Entrepreneurs' decisions are also going to be a function of the state variables. This is direct in the case of the aggregate shock because it enters the payoff function of the entrepreneurs. The endogenous state variable enters indirectly because of its effect on bank interest rate setting behaviour and thus the value of the outside option of being an inventor. All choices are made based on the conditional expected pay-offs next period which are functions of the information set.

2.2. Inventors

Each period, an inventor receives interest on her deposit $\tau(\Omega)$ and an idea with profitability index a. The bank will offer a loan contract $\rho(\Omega, i)$ in stage 2 and the inventor chooses between paying S and entering production next period as an agent of cohort i with profitability index a' given $F(da\prime, a)$ or waiting for another draw from $G(a)$ next period. The value function of an inventor is consequently

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

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

$$\int_{\Omega'}\int_{A} V_E(a\prime, i, \Omega'; .)J(d\Omega', \Omega)F(da\prime, a_E(\Omega)) - S = \int_{\Omega'}\int_{A} V_I(a', \rho(\Omega'); .)J(d\Omega', \Omega)G(da')$$

2.3. Equilibrium

Define $H([0, a), [1, j])$ as the measure of entrepreneurs at the *end* of each period with profitability index in the interval $[0, a)$ and in cohorts $[1, j]$. And to simplify notation define $FH(a'; i) = \int_A F(a\prime, a)H(da, i; .)$, which is the distribution of firms entering state a' from the distribution over a in the previous period for cohort i.

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:

$$H'([0, a'), [1, j]; \Omega, .) = \sum_{i=1}^{j} \left[\begin{array}{c} I(i(\Omega))D \int_{a_E(\Omega)}^{a'} G(a) \\ + \int_{a_X(i,\Omega)}^{1} FH(da'; i) - \varphi \int_{a_X(i,\Omega)}^{a_T(i,\Omega)} FH(da'; i) \end{array} \right] \quad (7)$$

The first term enumerates how many agents enter at profitability levels below a' given state Ω and thus in cohort $i(\Omega)$. ($I(i(\Omega))$ is an indicator function when state Ω returns cohort i). The middle term describes how many continuing entrepreneurs evolve into profitability subset $[0, a')$ from the measure of entrepreneurs above the voluntary exit threshold aggregated across each cohort i. The third term eliminates those entrepreneurs closed down by the bank because they are monitored and have their loan recalled, again aggregated across cohorts. Defaulting entrepreneurs are implicitly removed by the lower truncation of the distribution at $a_X(i,\Omega)$. Since each entrepreneur borrows one unit of capital, $H(A,[1,k];.)$ is the measure of the volume of loans outstanding at the end of each period. Intuitively, one can think of the bank as running separate portfolios for each cohort and in which each portfolio is evolving separately according to the optimal behavioural choices for that cohort. Every cohort portfolio is decaying through exit and one is refreshed each period according to Ω. As the system transits through Ω, each cohort is "visited" from time to time.

The balance sheet constraint for the bank is that it needs to have enough deposits to fund its loan portfolio. In the first two cases presented below in which there is no aggregate savings response, the bank has to solve

$$H(A,[1,k];\Omega,.) = D = \frac{1}{2}$$

every period (where everything depends on the conditioning assumptions denoted by "."). The following expression describes the bank's profits every period.

$$\begin{aligned}\Pi(\Omega) \quad &= \quad \sum_{i=1}^{k}\begin{bmatrix} \rho(\Omega,i)\int_{a_\delta(i,\Omega)}^{1} FH(da';i) \\ -\int_0^{a_\delta(i,\Omega)} \lambda(a')FH(da';i) \\ -\varphi m\int_{a_X(i,\Omega)}^{1} FH(da';i)\end{bmatrix} \\ &\quad -\tau(\Omega)\,H(A,[1,k];.) \end{aligned} \tag{8}$$

The three terms inside the square brackets relate to cohort-specific revenues and costs. The first term is the loan interest paid by non-defaulting borrowers. This depends on how many firms evolve into an individual state above the default threshold and the loan interest rate paid by that cohort $\rho(\Omega,i)$. The second term deducts loss given default assuming that the bank closing the project is inefficient (determined by the function $\lambda(a')$). The third term subtracts the costs incurred in monitoring the continuing firms of each cohort. The thresholds and firm distributions across each cohort are different and contingent on the information set. Profits each period are the sums across each cohort less the common deposit rate paid.

For future reference denote gross output as the weighted sum of individual gross profitability q:

$$Y = \sum_{i=1}^{k} \int_{A} q(a) FH(da; i)$$

The invariant model in Penalver (2013) is nested within the current framework by setting $\epsilon = 0$, assuming the deposit rate is fixed exogenously at $\bar{\tau}$ and noting that there is effectively only one cohort. Propositions 1 and 2 of that paper state that for a fixed monitoring intensity φ, there is a unique loan interest rate $\overline{\rho}$ that satisfies

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

and delivers an invariant distribution. 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 1. 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 illustrates the one period transition of the distribution in Figure 1 with the invariant distribution overlaid. (This illustrates the expression $FH(a, i)$.) 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 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

Figure 1: Invariant distribution of firms on the bank's balance sheet

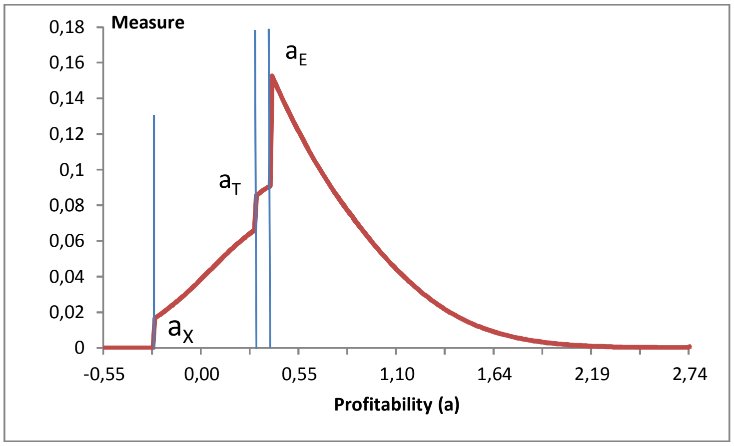

left, there are entrepreneurs that fall below a_X but above a_δ and exit voluntarily. Finally, there is a portion of the distribution that falls below a_δ and defaults.

3. The flexible loan rate case

The purpose of the current paper is to consider economic dynamics when there is an aggregate shock, $\epsilon > 0$. Clearly, constant loan interest and deposit rates will now no longer satisfy the balance sheet constraint. As a reference case, it is useful to very quickly describe the situation when the bank has the flexibility to change loan interest rates so that new and existing borrowers pay the same rate. The solution to this case is remarkably straightforward. Observe that if the bank sets

$$\rho = \overline{\rho} + z$$

and $\tau = \overline{\tau}$ then we have

$$q(a) - (\overline{\rho} + \epsilon) + \epsilon = q(a) - \overline{\rho} = q(a) - (\overline{\rho} - \epsilon) - \epsilon$$

Trivially, net profits are unaffected by the aggregate shock if the change in the loan interest rate exactly offsets it. Since net profits are the now the same regardless of the aggregate state, then the value of being an entrepreneur or an inventor is also invariant to the aggregate state and the entry and exit thresholds from the steady

Figure 2: One period transition of bank balance sheet

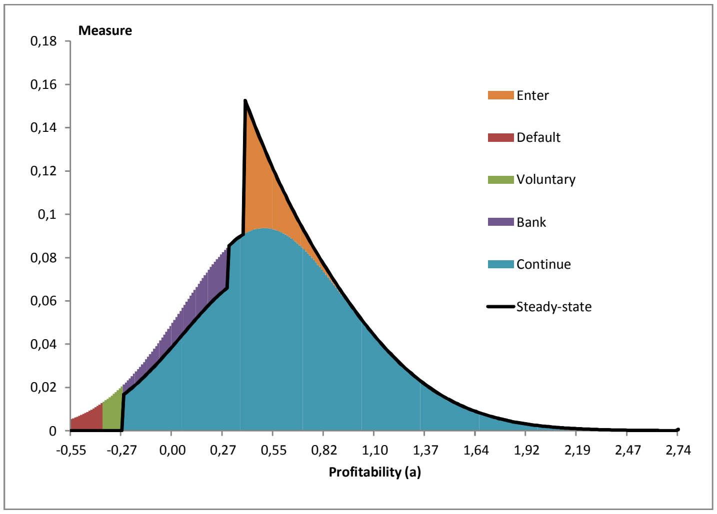

state model will always satisfy the indifference conditions in the aggregate shock case. With all entry and exit behaviour the same, the distribution of projects over A does not change and always satisfies the balance sheet constraint. Finally, since the distribution does not change, the default rate is constant over time. Aggregate output rises and falls by ϵ only when the aggregate shock strikes. As a result, the bank is providing full insurance against the aggregate shock even though the agents are risk neutral with the corrolary that it is the bank that fully absorbs movements in aggregate profits. The stark simplicity of this solution is, of course, due to the assumption that the effect of the common shock is constant across the distribution of the profitability state and a common interest rate rule. But this will help to make very clear how other contracts alter dynamics. In particular, it shows that maturity transformation *per se* does not make any difference to aggregate dynamics.

4. The fixed loan rate case

The solution to the fixed rate contract is considerably more difficult. In the fixed rate contract, the loan interest rate is constant throughout the life of the loan and the bank varies the interest rate on *new* loans as its control variable.

The bank retains the flexibility to change the deposit rate τ. In Section II, the loan interest rate rule for the bank was described generally as a function of the observed state variables, $\rho(\Omega, i)$. The fixed rate contract implies that a borrower who receives a loan at date s pays a loan interest rate $\rho(\Omega_s)$ at all dates $t > s$. If $\Omega_k = \Omega_l$ at dates k and l, then borrowers who took out loans at both dates pay the same interest rate at all subseqent dates and are thus in the same cohort. An immediate consequence of fixing the interest rate for the life of the loan is that the endogenous component of the contemporaneous state vector θ does not *directly* affect the value function of being an entrepreneur. There is, however, still an indirect channel because the expected value of being an inventor, which features in the exit options, does depend on all the elements of Ω. An inventor in the fixed rate case is now weighing up two risks: if she turns down entering production with her current profitability draw and her current loan rate offer, what is the prospect she will get a better combination of profitability draw and borrowing costs in the future and what will be the deposit rates on offer? The profitability risk (over a) is purely idiosyncratic and i.i.d and thus reflects the timeless distribution of $G(a)$. Expectations of future loan and deposit interest rates are, on the other hand, state contingent because of the persistence in the aggregate shock and potentially persistence in the endogenous state and the dependence of the interest rate rules on this state vector. Agents are assumed to use a forecasting function for the state vector as a function of current observables,

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

Moreover, if the bank rule and the forecast rules are linear, which will be the case for the numerical solution presented below, then

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

Similarly, if we assume a linear rule for the deposit rate $\tau = \tau(\Omega_t)$, then

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

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. Therefore, all agents will have a common view on the expected value of being an inventor. The key behavioural difference in the fixed rate case relative to the flexible rate case is that borrowers' reactions to changes in the value of being an entrepeneur will depend on their cohort. A cohort with high fixed loan interest rates is more likely to exit (less likely to continue) for a given expected value of switching to being an inventor than a cohort with a low fixed loan rate. Since different cohorts behave in different ways for the same information set, the distribution of cohorts across the bank's balance sheet will matter for the exit rate, including the default rate. The problem for the bank under a fixed rate contract is how to set the loan rate for new borrowers and the deposit rate for all lenders in order to meet the balance sheet constraint at all times, knowing that the loan terms persist for as long as the loans do. Of course in a stochastic setting with a discrete set of control variables, having exactly the same level of deposits as loans in all periods is an impossible task, so the more relaxed optimising criteria is to find a rule that minimises a loss function which is a negative function of bank profits and a positive function of the standard deviation of excess deposits.

It is now time to be more precise about the endogenous state variable θ. Equation (7) shows that the transition of the balance sheet is a complex function of its distribution over idiosyncratic profitability states. Differences in the shape of this distribution will alter the rate at which loans are repaid or defaulted on and thus the amount of spare lending capacity available. The exit rates also depend on the cohort: *ceteris paribus*, a cohort with a higher loan interest rate will quit in a higher profitability state than a cohort with a low interest rate. 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 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.[6] Henceforth $\theta \equiv \frac{H(A)}{\bar{H}(A)}$ will be referred to as "excess loans".

A solution to the model outlined in Section II for the fixed rate loan case is a fixed point in which:

- the bank follows loan and deposit interest rate setting policy rules which, given the behaviour of depositors and borrowers, and the characteristics of the common shock process, minimises the loss function;
- agents' expectations are model consistent; and

[6] In early research for this paper, the portfolio average interest rate was also added to the list of endogenous variables. This substantially complicated the solution algorithm without adding any accuracy to the solution and was dropped.

- 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: Posit linear rules for the deposit interest rate and the interest rate for new borrowers as a function 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.

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

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

- Step 2: Posit a 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 in each cohort conditional on a and Ω and the forecast rule from Step 2 and the bank rule 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 2,600 periods based on these rules, throw away the first 500 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. Calculate the loss function over the estimation range.

- Step 6: Repeat Steps 1 to 5 with a different bank interest rate rules using a search algorithm until the loss function is minimised.

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 (9 nodes). After all the number crunching, the optimal fixed interest rate rule for a simulation of the model is

$$\rho = 0.0613 + 0.0002z + 0(\theta_{-1} - \bar{\theta}) \tag{9}$$

$$\tau = 0.035 + 0.0026z + 0.896(\theta_{-1} - \bar{\theta}) \tag{10}$$

with an expectations function

$$E\theta_t = 0.0035z + 0.861(\theta_{-1} - \bar{\theta}) \tag{11}$$

The solution parameters are non-negative and this accords with common sense. It is clear from comparison of the parameters of the two rules that the deposit rate is doing most of the work to stabilise the balance sheet. The deposit rate responds more to the change in the aggregate shock and is the only variable to change in response to the balance sheet measure. This division of labour between the two control variables is completely intuitive: changing the loan rate has persistent effects on the balance sheet and thus introduces endogenous variation whereas changing the deposit rate does not. Nevertheless, there is still some benefit in moving the loan rate on new borrowers in response to the aggregate shock. The average distance from the balance sheet condition is less than 1% which is reasonable given the approximations made.

The important question, though, is how does the assumption of fixed rate loans alter the dynamic response to aggregate shocks? The charts that follow use the NBER convention of shading for recessions which in this case corresponds to periods with the low shock state. By construction, there are on average as many low as high periods over the simulation and the average expected duration of any shock is 5 periods.

Figure 3 illustrates the evolution of productivity over 2 sequences of high states and an intervening sequence of low states. In this economy, the only productive input, capital, is in fixed supply so the evolution of productivity is driven by

Figure 3: Productivity - fixed rate case

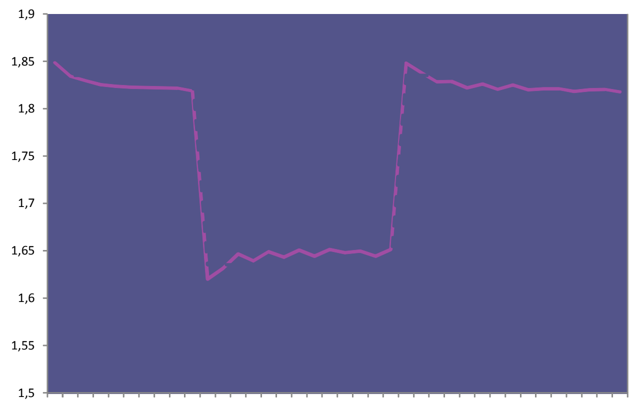

only two effects - the exogenous shock process and the endogenous effect from the allocation of capital between cohorts. Figure 3 clearly shows that the former effect dominates. This is a result of the general equilibrium setting rather than the calibration parameters (although that probably matters too). In order to satisfy the balance sheet constraint in this closed economy, prices need to move to shift enough capital between cohorts and even the distortion induced by fixed interest rates does not prevent the price mechanism and rational expectations inducing the least productive to depart and the most productive to enter. In other words, allocative distortions only work at the margin. It is interesting to note that one reason why the effect is relatively small is because the outside option of being an inventor is state-contingent and therefore helps to regulate exit. Models of dynamic entry and exit which assume that the outside option is fixed (for example at zero) completely miss this stabilising force. Moreover, it is actually the relative stability of the cross-sectional distribution that allows the Krusell and Smith first-moment approximation technique to work. If the reallocation effects were very large so that the cross-sectional distribution varied significantly over time, then other moments would become important. The dashed line in the figure shows the evolution of productivity in the flexible rate case. The fixed rate case deviates from this case in only small ways but these are nevertheless interesting as will be discussed below. Before describing the evolution of endogenous productivity, it is useful to show first one difference between the two cases which is dramatic.

Figure 4: Default rate - fixed rate case

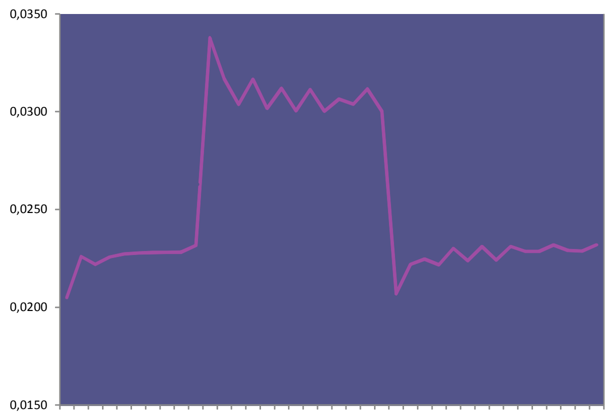

Figure 4 illustrates the default rate in the fixed rate case with the constant value in the flexible rate case (again the dashed line). The default rate is now strongly procyclical. When a negative aggregate shock hits, there is a brief spike up and then convergence to an above average default rate. (There is a symmetric effect after a positive aggregate shock.) The jump up in default after an adverse shock is entirely intuitive. A mass of firms that have been willing to continue whilst times are good suddenly find themselves in the default region. However, the subsequent convergence to above and below average default rates is, in a sense, an ex post illusion. At any date, firms make continuation decisions based in part on the expected aggregate shock state in the next period. For the Markov process, this lies somewhere between the two actual states. Thus when a negative shock follows another it is an adverse shock relative to expectations. Some firms that chose to continue in the hope that a positive shock might follow are disappointed and default. (The arguments are symmetrical when a positive shock strikes.)

Figure 5 illustrates the evolution of endogenous productivity over a long time horizon of the simulation. Although this is a tough chart to digest, the long run perspective offers some illuminating insights into the way the fixed rate contracts influence the efficiency of the allocation of capital over time. It is evident that there are sharp 'Schumpeterian' jumps in productivity when a negative shock strikes downturns and correspondingly a drop in productivity when a positive shock strikes. The negative shock pushes a lot of less productive firms into default

Figure 5: Endogenous productivity - fixed rate case

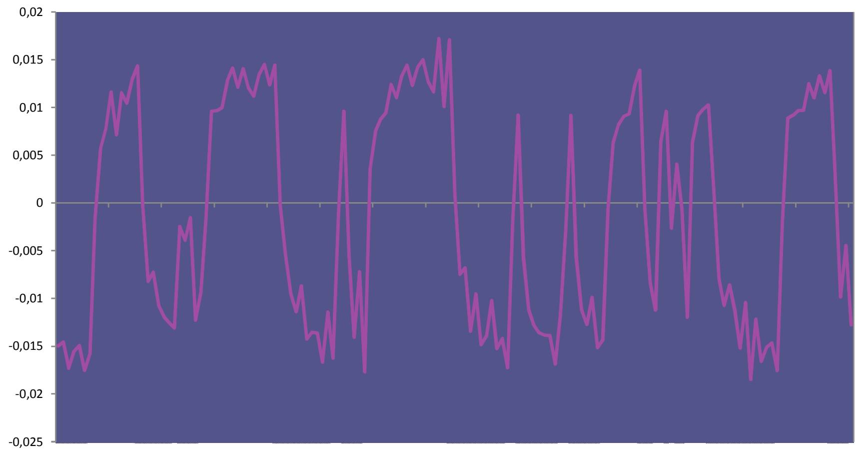

and others into the region in which they voluntarily exit. Yet another group find themselves below the covenant threshold. Even during a downturn, these resources can be allocated to entering firms with higher productivity. This process works even though there is an endogenous reduction in the value of being an inventor through a reduction in the deposit rate. A second striking feature of this chart is that long duration positive shocks can lead to ongoing reductions in productivity. Interesting, this effect does not appear to be symmetric. After a negative shock strikes, there is a jump and then not much further improvement. This asymmetry occurs because a negative shock chops away at the bottom end of the distribution, whereas a positive shock 'opens' up regions of the productivity space in which the distribution can subsequently move. It is interesting also to consider note that during positive shocks, productivity declines despite that fact that the average entrant is of increasingly higher productivity. This happens because positive shocks induce more firms to stay in production and with less exit, there is less capacity to fund new entrants (even with the endogenous response of the deposit and loan interest rate). The bank, of course, rations out these available funds to the highest productivity entrants and for this reason the average entrant quality is increasing. But, and this is the crucial point, the infra-marginal entrant lost is of higher productivity than the marginal incumbent who remains.

5. The financial accelerator case

The analysis presented in the previous section - which delivered a Schumpeterian adjustment process - was predicated on a fixed supply of capital. There was, by construction, no additional resources in the economy so the only effect came through reallocation between agents. This section explores a relaxation of this assumption by introducing an aggregate savings reponse to positive and negative aggregate shocks. This is done through the, admittedly ad hoc, assumption that aggregate savings increase by a fraction of the difference between total output and steady state output and is labelled the 'financial accelerator' case. This has no micro foundations in the model which completely abstracts from consumption timing decisions by agents. It is also in conflict with the assumption that agents have only a fixed unit of capital which closed down analysis of the intensive margin. On the other hand, the assumption is simple, transparent and delivers interesting results without having to develop a much more complicated extension of the model.

As in the previous case, the bank is assumed to make fixed rate loan contracts with borrowers for a single unit of capital for projects needing two units. The solution approach mirrors that for the fixed rate loan case in Section 4 by searching for a fixed point of linear bank interest rate setting rules, a linear expectations functions and individual optimisation decisions by agents. The following rules solve the problem:

$$\rho = 0.0613 + 0.0002z + 0(\theta_{-1} - \bar{\theta})$$

$$\tau = 0.035 + 0.0018z + 0.95(\theta_{-1} - \bar{\theta})$$

The deposit rate is moved less aggressively in response to the aggregate shock but is more responsive to tha aggregate state. Figure 6 shows that the evolution of productivity is still dominated by the aggregate shock rather than the allocation effect. (In the financial accelerator case, output is much more cyclical because of the cyclical change in capital availability. So to keep the charts comparable, this resource effect has been scaled away.) The endogenous effect, however, is now pro-cyclical rather than counter cyclical.

This evolution of endogenous productivity is shown more clearly over a long period of the simulation in Figure 7. The underlying pattern of exogenous shocks is identical between this simulation and that of the fixed rate case and all the structural parameters of the model are identical. So the strikingly different results are only due to the introduction of the aggregate savings assumption. Endogenous

Figure 6: Productivity - financial accelerator case

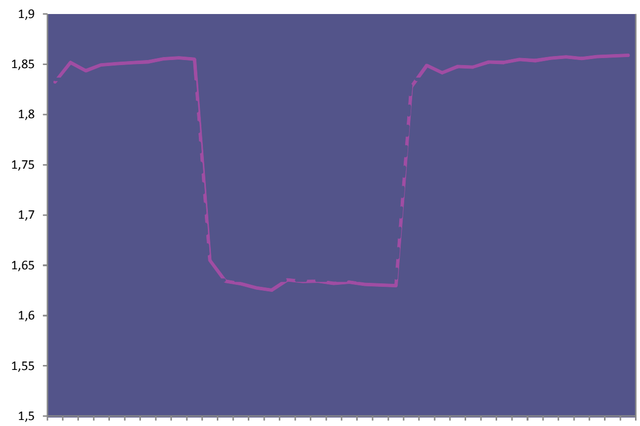

productivity is pro-cyclical because when savings rise with output (of course less than one-for-one), then more resources are available to fund entry. So even though there is still a decline in the average profitability of the marginal entrant, there is a rise in overall productivity. This can occur even if the marginal entrant has below average productivity. To understand this slightly counterintuitive result, it is useful to think of entry as leaning against the natural tendency of the distribution to deteriorate. The stationary equilibrium occurs when these two forces cancel out. Adding entrants gives an extra positive push. But this does not last indefinitely. One can see from the figure that when there is a long sequence of positive shocks, the bulge of higher productivity firms eventually begins to deteriorate at a rate faster than the higher entry rate can offset it and endogenous productivity starts to revert towards the average.

6. Conclusion

We are now 7 years into an economic downturn triggered by a collapse in the residential real estate market in some countries followed by a global banking crisis. Bank balance sheets remain laden with weak and non-performing loans and they still have very limited capacity to extend new credit. At the same time, many of the countries at the centre of the crisis are stuck with physical capital allocated

Figure 7: Endogenous productivity - financial accelerator case

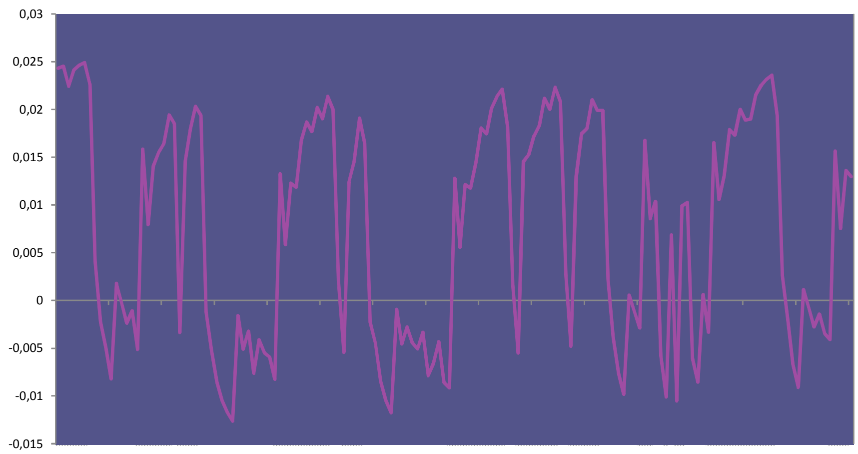

to low value projects (viz a surplus of residential housing). These facts are not independent phenomena and result directly from two functions of the banking system: maturity transformation and credit risk management. Banks determine how many of an economy's potential risky long-term projects receive credit and on what contractual terms. If physical capital were infinitely fungible, then all projects could use putty rented in spot markets and the adjective "long-term" would have no meaning. There would also be no possibility that physical capital could be "misallocated" over time. Therefore, implicit in both notions is the idea that projects are not homogeneous and it is costly to convert the capital used in one project into another. So how capital is allocated matters for future economic performance and these allocation incentives change over time.

This paper suggests why it is difficult to discern clear patterns in the response of productivity over the business cycle and between countries. Firstly, the magnitude will depend on the relative proportion of fixed and variable rate loan contracts. The more variable rate loans, even with long maturity lending, the more stable productivity can be expected to be over the cycle. A similar comment applies if entry and exit costs are relatively low. For countries with relatively high entry and exit costs (for example those with highly specialised industries) and fixed rate loan contracts, the direction of response depends to a considerable extent on the response of bank funding and in particular the aggregate savings rate.

If negative aggregate shocks are compounded by a fall in bank deposits (for example through the withdrawal of foreign funding), then aggregate productivity can fall. By contrast, if funding holds up after a negative shock, the Schumpeterian cleansing can improve aggregate productivity.

7. Appendix

The model presented in this paper is dynamic, has heterogeneous agents and individual and aggregate stochastic shocks. Despite sharing many similarities with other heterogeneous agent models with aggregate uncertainty (such as described in Allais *et al* (2013) and including the seminal model of Krusell and Smith (1998)), there are quite a few differences too and these hinder direct comparison. Standard models with heterogeneous agent and aggregate uncertainty have agents choosing a continuous control variable (usual a level of savings) when subject to very simple shocks (employed or unemployed, boom or bust) and a unique market clearing solution. The margin of interest is the distribution of agents over the endogenous state variable. In this model the aggregate shock process is the same - a first order Markov process in two states - but there the similarities end. The heterogeneous agents in this model are making very simple binary decisions (stay or switch) subject to a continuous idiosyncratic shock process. There are no Euler conditions, only threshold points. By shutting down the intensive margin, all the adjustment dynamics come from the threshold conditions shifting around and therefore rates of inflow and outflow. The margin of interest is the distribution of agents over the idiosyncratic state. Market clearing occurs when the bank the same measure of loans outstanding as deposits. In the flexible and fixed rate cases, this is known *a priori* since the measure 1 of agents must be divided in two. But in this model, there is an actor (namely the bank) who cares about the shape of the distribution as well as needing to satisfy this balance sheet constraint. Moreover by being able to set deposit and loan rates, the bank has sufficient instruments to influence both. In effect, the bank is picking possible paths for the economy through its choice of deposit and loan rate rules. Finding the optimal choice is the key decision in the model.

The dynamic model has a steady-state if the aggregate shock is switched off. With a symmetric shock process it is natural to presume that the economy fluctuates around this steady state and steady state values are used as starting points for the simulations and for the constants in the interest rate setting rules.

There are 6 steps in the solution algorithm

- Step 1: Posit parameters for the linear rules $\hat{\rho}$ and $\hat{\omega}$ for the bank's interest rate setting rules for depositors and new borrowers (respectively) as a function of the observable state variables Ω.

- Step 2: Posit parameters for a linear forecasting rule $\hat{\sigma}$ for the agents for the endogenous variable θ as a function of Ω.

- Step 3: Calculate the value functions for inventors and entrepreneurs in each cohort conditional on a and Ω and the forecast rule from Step 2 and the bank rules from Step 1. On the basis of these value functions, calculate the observable state contingent entry and exit rates.

- Step 4: Simulate the economy based on these rules for 2600 periods, throw away the first 500 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. Calculate the average absolute net deposits over the estimation range.

- Step 6: Repeat Steps 1 to 5 with a different bank interest rate rule until a rule is found which most closely satisfies the balance sheet constraint over time.

The two outer loops are relatively straightforward but highly time consuming. The parameters in Step 6 are solved using the Nelder-Mead technique.

Step 3 is the heart of the algorithm. There are two states in the aggregate shock process and the endogenous states, θ, is further simplified to 15 grid points giving a 2×15 grid of observation variables, Ω. For each point on the observation grid, there is an associated control variable determined by the posited linear rule $\hat{\rho}\left(\sigma\left(\Omega_t\right)\right)$. In the case of the fixed interest rate rule, the control variable is the interest rate charged on new loans. There are thus 30 possible cohorts of entrepreneurs, one for each of the 30 aggregate state pairs. For each cohort, there is a value function determined by the maximum of the 3 options across a profitability grid of 5001 points. Up to the this point, the entrepreneurs' choices are based on things they know for certain - their interest rate, their profitability state and the aggregate shock - and on which they can form an expectation for the value of being an entrepreneur next period based on the stochastic processes F and Z. But the exit options depend on an expectation of the value of being an entrepreneur which in turn depends on the expected interest rate available on new loans. Given the bank's loan interest rate setting rule, this implies forecasting the observation variables for next period. The state conditional forecast is again easy. The endogenous observation variable, however, is much more difficult. As is explained in the text, the agents are assumed to use a linear forecasting rule, $\hat{\sigma}$, based on the previous period's endogenous state and the current period's exogenous state. Of course, the point forecast will almost never coincide with one of the endogenous grid points so probabilities are assigned to the two nearest points according to their relative distance to the point forecast. This creates a system of mutually inter-related value functions for the entrepreneurs and inventors - conditional on the posited rules $\hat{\rho}$ and $\hat{\sigma}$ - which is iterated until convergence.

In Step 4, the model is simulated for 2600 periods. The endogenous state is again a problem since the size of the balance sheet at the end of each period will almost never lie on one of the grid points. In this case, the endogenous state is randomly assigned to one of the two nearest points where the probability is determined by the relative distance to the actual value. Thus the model moves from grid point to grid point along the simulation path. In conjunction with the random switching of aggregate shock states, the model transits through the 2 dimensional state space approximated by the 2×15 grid. The survival distribution for each of the 30 cohorts is tracked and one of them is refreshed by new entrants. At the end of each simulation, the average absolute distance from the balance sheet constraint is calculated using only the final 2100 observations. (These low number of simulations is due to memory constraints - the state space in each period is already $5001 \times 30 \times 2 \times 15$.)

Each time the outer loops are repeated, the same sequences of random numbers are used. Several sequences were checked to confirm the main findings of the numerical simulations.

8. Bibliography

References

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

[2] Algan, Y., O. Allais, W.J. Den Haan, and P. Rendahl, forthcoming, Solving and simulating models with heterogeneous agents and aggregate uncertainty, in **Handbook of Computational Economics** vol. 3, K. Schmedders and K.L. Judd (eds.).

[3] Clementi, G. and D. Palazzo (2013) 'Entry, exit, firm dynamics and aggregate fluctuations' NBER Working Paper no 19217 (July 4, 2013).

[4] Diamond, D. and P. Dybvig (1983), 'Bank runs, deposit insurance, and liquidity', *Journal of Political Economy* vol. 91, pp. 401–419.

[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] Holmström, B. and J. Tirole (1997), 'Financial intermediation, loanable funds and the real sector', the *Quarterly Journal of Economics*, vol. 112, no. 3, pp. 663-691.

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

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

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

[11] Laffont, J.-J. and J. Tirole (1988), 'The dynamics of incentive contracts', *Econometrica*, vol. 56, 5, pp. 1153-1175.

[12] 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.

[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] Penalver, A. (2013), 'Pre-crisis credit standards: monetary policy or the savings glut?' Paris School of Economics mimeo.