Universidad Nacional de La Plata

Economic Growth, Liquidity, and Bank Runs

Ennis, Huberto (Research Department, Federal Reserve Bank of
Richmond) and Keister, Todd (ITAM and University of Texas at
Austin)

Economic Growth, Liquidity, and Bank Runs[1]

Huberto Ennis

Research Department, Federal Reserve Bank of Richmond

Todd Keister

ITAM and University of Texas at Austin

March 16, 2002

Abstract

We examine the growth implications of bank runs. To do so, we construct an endogenous growth model in which bank runs occur with positive probability in equilibrium. In this setting, a bank run has a permanent effect on the capital stock and on the level of output. In addition, the possibility of a bank run changes the portfolio choice of banks and thereby affects the long-run growth rate. We consider two different equilibrium selection rules. In the first, a run is triggered by sunspots and occurs with a fixed probability. A higher probability of a run in this case leads banks to hold a more liquid portfolio, which decreases total investment and thereby reduces capital formation. Hence the economy grows slower, even when a run does not occur. Under the second selection rule, the probability of a run is influenced by the bank's portfolio choice. This leads banks to place *more* resources in long-term investment, and the economy both grows faster and experiences fewer runs.

---

1 Introduction

Bank runs and banking crises in general are an important economic phenomenon, both historically and in recent years. These crises are typically accompanied by costly economic disruptions, with estimates of the cost of recent crises ranging as high as 40-55% of GDP.[2] Much has been written about these crises, including analyses of their possible causes and their many effects on the economy. We focus on one particular facet: the effect of the possibility of bank runs on capital formation and thereby on economic growth. While there is now a fair amount of empirical evidence on the effects of banking crises on economic growth, not much theoretical work has been done on the subject. In this paper we present an endogenous growth model where bank runs occur with positive probability in equilibrium. This allows us to examine also how the *possibility* of a run changes the decisions made by agents in the economy and how these changes affect long-run economic growth.

Our model of the behavior of banks is in the tradition of Diamond and Dybvig [11], which highlights the role of the banking system in creating liquidity by taking in short-term deposits and making long-term investments. In particular, Diamond and Dybvig [11] shows how demand deposit contracts can easily lead to a situation where there are multiple equilibria of the game played by a bank's depositors (the "post-deposit" game), one where a bank run occurs and one where it does not. The optimal contract for the bank to offer in the "pre-deposit" phase therefore depends critically on how an equilibrium of the post-deposit game is selected. One approach is to assume that agents coordinate their actions on a sunspot variable, a publicly-observed random variable that is *extrinsic* in the sense that it has no effect on the fundamentals of the economy.[3] Peck and Shell [19] shows that if a sunspot-induced run is sufficiently unlikely, depositors can prefer a contract that permits runs, even when a broad set of possible deposit contracts is considered.[4] It is always feasible in this setting for the bank to choose a contract that prevents runs, but it may be too costly from an *ex ante* point of view.

We have kept our model of bank behavior as simple as possible, while retaining the spirit of the Peck-Shell [19] analysis. In particular, we restrict banks to offer simple demand deposit contracts in order to make the problem tractable even when there is a large number of depositors. There is a large literature on

---

[2] These are the estimates given in Caprio and Klingebiel [6] for the Chilean and Argentine crises, respectively, in the early 1980s. For a large number of other crises, they report costs in excess of 10% of GDP.

[3] In a recent comprehensive study of modern banking crises, Boyd *et al.* [4] conclude that the available evidence strongly points toward a sunspots-based explanation for the cause of these crises.

[4] This type of result is shown for a restricted set of deposit contracts in Cooper and Ross [10].

the roles of the banking system and its microeconomic structure,[5] and it is not our intention to contribute to these issues. Rather, we aim to highlight the basic growth implications of bank runs that are likely to follow from any model where the activities of the banking system matter for real allocations. Banks in our model take as given an equilibrium selection rule, and they offer the demand deposit contract that maximizes the expected utility of their depositors. The only critical aspect of our model of the banking system is that, as in Peck and Shell [19], banks may choose a contract that admits a run equilibrium to the post-deposit game.

In addition to choosing the contract that it offers depositors, the bank must also allocate its portfolio between storage and investment in new capital. Investment is illiquid in the sense that much of its value is lost if the project is terminated early.[6] We show how the possibility of a bank run influences the process of capital formation in two distinct ways. The first is obvious: when a run occurs the bank liquidates investment, which reduces the amount of new capital created in that period. The second is more subtle but no less important: since the bank is aware of the possibility of a run, it places a higher fraction of its reserves in storage. Because there is a large loss from liquidating investment, holding a more liquid portfolio allows the bank to serve a larger number of customers during a run and hence provides depositors with insurance against a bad realization of the sunspot signal. However, resources placed in storage do not produce new capital, and hence a more liquid bank portfolio implies a lower capital stock in the following period. Thus the mere possibility of a bank run reduces capital formation, even when a run does not occur.

The long-run impact of the possibility of bank runs depends critically on whether or not the actions of the bank affect the long-run growth rate of the economy. We embed our banking model in an Ak model of growth,[7] which generates the following results. Because the marginal product of capital is constant, the path of real output is history dependent and a bank run necessarily has a permanent effect. This implies that the true cost of a crisis is much larger than a short-run estimate (such as those given in the first paragraph above) would indicate.[8] In addition, even small changes in banks' portfolios will change the long-run growth rate and therefore have large effects on the level of real output over time.

We also examine the issue of equilibrium selection in the post-deposit game in more detail, using an approach that we have developed elsewhere (Ennis and Keister [12], [13]). The standard sunspots approach assumes that the probability of a run is a fixed constant (as long as both equilibria exist). In our approach, the probability of a run depends on the strength of the incentive for agents to run as measured by the *risk*

---

[5]   Freixas and Rochet [15] provide a detailed summary and list of references for this literature.

[6]   This modification of the Diamond-Dybvig approach was taken in Cooper and Ross [10], which showed that the possiblity of runs could lead the bank to hold excess liquidity.

[7]   In this we follow Bencivenga and Smith [3], which studies the role of financial intermediaries in promoting growth. Our model is in many ways similar to theirs, but they do not examine the possibility of bank runs. A key characteristic of the Ak model is that it has no transitional dynamics, which greatly simplifies our computations. The model can be viewed as approximating the balanced growth path of richer models. See Antinofi, Keister, and Shell [1]

[8]   Boyd, Kwak, and Smith [5] provide evidence that the effects of modern banking crises are indeed very long lived.

factor of the run equilibrium. (A low risk factor corresponds to a strongly risk-dominant equilibrium.) When the risk factor is very low, an agent would choose to run for a wide range of beliefs about the actions of other agents, and therefore we say that a run is relatively likely.[9] The risk factor is determined largely by the contract offered by the bank. We show through examples that when the bank recognizes that its choice of contract affects the probability of a run, it holds a less liquid portfolio. This may seem surprising at first. However, it is intuitively clear that this must be the case. In order to reduce the probability of a run, the bank must decrease the expected payoff of running (relative to the payoff of waiting). When the bank places more funds in investment and less in storage, the payoff of waiting when there is no run increases. The payoff of waiting when there is a run does not change (it is zero). Lower liquidity levels therefore imply a higher expected payoff of waiting and hence lead agents to be more likely to wait. We also show that the more sensitive the probability of run is to the portfolio of the bank, the more resources the bank puts into investment. As we mentioned above, less liquid bank portfolios lead to more capital formation and therefore higher growth rates. Hence there is no tradeoff between growth and stability is this model; less liquid portfolios bring higher growth with fewer bank runs. It should be born in mind, however, that less liquid portfolios provide less insurance for depositors against high early withdrawal demand. This is what prevents banks from eliminating runs altogether.

Finally, we show that economies with more productive investment technologies will tend to grow faster for two reasons. First, a more productive investment technology induces the bank to invest more and generate more capital, which results in higher growth. This is completely standard. In addition, however, higher investment reduces the incentive for depositors to run and therefore also brings a lower likelihood of a run. As a result, the average growth rate over a long time period will be higher because runs will occur less often. In this way, the model provides an amplification mechanism for differences in the productivity of investment.

The outline of the remainder of the paper is as follows. The the next section we describe our model in detail. In section 3, we describe the equilibrium behavior of the economy and compute this behavior for some numerical examples. We also trace the implications of the possibility of bank runs for growth under the standard sunspots approach. In section 4, we present our alternate model of equilibrium selection and examine how this affects the decisions of banks and the long-run behavior of the economy. In section 5, we offer some concluding remarks.

---

[9]   Ennis and Keister [12] shows how an equilibrium selection rule of this form is the natural outcome of an adaptive learning process in a stochastic environment.

2 The Model

The economy consists of an infinite sequence of two-period lived, overlapping generations, plus an initial old generation. There is a single, perishable consumption good, which is produced using capital and labor. Agents in the initial old generation are each endowed with k_1 units of capital, and have preferences that are strictly increasing in consumption in the single period of their life. In each time period $t = 1, 2, \ldots$, a continuum of agents with unit mass is born. Each of these agents is endowed with one unit of labor when young and nothing when old, and each is either patient or impatient. Preferences are given by

$$u\left(c_{1,t}, c_{2,t}\right) = \left\{ \begin{array}{c} \frac{b_1}{\gamma}\left(c_{1,t}\right)^{\gamma} \\ \frac{b_2}{\gamma}\left(c_{1,t} + c_{2,t}\right)^{\gamma} \end{array} \right\} \text{ if the consumer is } \left\{ \begin{array}{c} \text{impatient} \\ \text{patient} \end{array} \right\}. \tag{1}$$

A fraction ϕ of consumers know at birth that they are patient. We refer to these as type I agents. The remaining (type II) agents will learn their type at the end of the first period of their lives, while there is still time to consume that period but after investment decisions are made.[10] Each type II agent in generation t is impatient with probability u_t and patient with probability $(1 - u_t)$. The realization of types is independent across agents, so that u_t is also the fraction of the population of type II agents in generation t that is impatient. The probability u_t is the realization of a random variable that gives the size of the aggregate liquidity shock in each period; high values of u_t correspond to high liquidity demand. We assume that u is independently and identically distributed over time, and that the distribution has a density function f.[11]

2.1 Production and investment

There is a large number of competitive firms who produce output using capital and labor as inputs according to the production function

$$Y_t = \overline{k}_t^{1-\theta} K_t^{\theta} L_t^{1-\theta},$$

where $\overline{k}_t$ is the average capital-labor ratio in the economy at time t, which is taken as given by each individual firm. Adding the capital externality is one way of preventing the marginal product of capital from falling too low as the economy grows and thereby generating endogenous growth.[12] There are many other, more interesting models with this property, including models of inventive activity. The externality-based approach allows us to keep the model simple and to abstract from transitional dynamics after a crisis, since

---

[10] We are thus collapsing the first two periods of the Diamond and Dybvig [11] setup into the first period of a young agent's life.

[11] Most of our analysis could be done under the assumption that u is a known constant over time, as in Cooper and Ross [10]. However, as Diamond and Dybvig [11] point out, in this case a simple suspension of convertibility is a costless way to eliminate the run equilibrium. When u is stochastic, however, the total suspension scheme is no longer optimal and the run equilibrium can exist under the optimal contract, as shown in Peck and Shell [19].

[12] The assumption that the exponent on the externality term is exactly equal to labor's share of income is of course special. Antinolfi, Keister and Shell [1] identify it as a bifurcation point in parameter space.

our economy will always be on a balanced growth path. Nevertheless, our banking model could easily be embedded in a richer model of growth. We discuss this extension briefly in section 5.

Capital is, of course, durable and therefore is one way for young agents to save. There are two other ways of saving, which we refer to as "storage" and "investment." One unit of consumption placed into storage at time t yields n units of consumption if liquidated later in period t and n^2 units of consumption if held until period $t + 1$. This storage can be thought of as a liquid asset such as money, with $n < 1$ representing an inflation tax.

One unit of consumption placed into investment in period t yields R units of capital in period $t+1$. This technology is the only way that new capital can be produced. If investment is liquidated early (at the end of period t), it yields $x < n$ units of consumption per unit invested. Hence investment is an illiquid asset, which yields a higher return than storage if held to maturity but a lower return if liquidated prematurely.

2.2 Timing of events

Period t begins with a stock of capital k_t owned by old agents. This capital is rented out to firms, who also employ young agents and thereby produce output. After production takes place, old agents sell the undepreciated capital. Letting q_t denote the price of capital, an old agent then has $(r_t + (1 - \delta) q_t)$ units of consumption for each unit of capital she began the period with. She consumes all of this and exits the economy.

Type I young agents know that they are patient, and therefore will save all of their income in whatever asset yields the highest return. The return to using a unit of consumption to purchase existing capital is

$$\frac{r_{t+1} + (1 - \delta) q_{t+1}}{q_t}.$$

The return from investing a unit of consumption in new capital formation is

$$R \left(r_{t+1} + (1 - \delta) q_{t+1} \right) \equiv \psi_{t+1}.$$

The decision rule of a type I agent is therefore the following:

$$\text{Invest in} \left\{ \begin{array}{c} \text{existing capital} \\ \text{either} \\ \text{new capital} \end{array} \right\} \text{ as } q_t \left\{ \begin{array}{c} > \\ = \\ < \end{array} \right\} \frac{1}{R}. \tag{2}$$

Since both strategies yield capital in period $t+1$, the agent simply chooses the option that yields more units of capital per unit of consumption invested.

The interesting investment problem is that of Type II agents. Because these agents do not know their

preferences until after the opportunities to invest have passed, they will form coalitions that we call banks. The agents will choose to deposit their income in the bank, and the bank will place some of these resources in storage and the rest in investment in new capital.[13] We assume that banks offer simple demand deposit contracts, and that suspension of convertibility is not possible. Each depositor chooses to withdraw her funds from the bank in either period t or period $t + 1$. Agents who choose period t arrive at the bank in random order. The bank offers a fixed rate of return on deposits withdrawn in this period, and it must honor this contractual obligation unless it has completely run out of resources.[14] Whatever resources remain in period $t + 1$ are divided among the remaining depositors. The bank therefore can be viewed as choosing what fraction of deposits to place in storage (we denote this fraction η_t) and what return to offer to agents who withdraw their deposits in period t (assuming it has not run out of resources; we denote this $a_{1,t}$). These two numbers completely determine the payoffs received by a depositor under each possible contingency.

After the bank sets the contract and type II agents have deposited their income, we move to what Peck and Shell [19] call the "post-deposit game." Each agent learns whether she is impatient or patient, and then decides whether to go to the bank in period t or in period $t + 1$. Following the literature, we focus on symmetric, pure strategy equilibria of the game. There are two possible equilibria of this type: one where all agents go to the bank at the end of period t (a run) and one where only impatient consumers go (no run). In choosing the optimal deposit contract to offer, the bank needs to know how likely each of these outcomes is. The same is true of an agent deciding whether or not she wants to deposit her savings in the bank. The standard approach is to assume that agents coordinate their actions based on the realization of a sunspot signal, so that a run occurs with a fixed probability (assuming that the contract offered by the bank is such that both equilibria exist). Our interest is in both this approach and a modified version of it in which the probability of a run depends on the parameters of the deposit contract. We begin by defining an equilibrium selection mechanism.[15]

Definition: An *equilibrium selection mechanism* (ESM) is a function that assigns a probability π to the run equilibrium and $(1 - \pi)$ to the no-run equilibrium for each possible deposit contract. These probabilities must be feasible, meaning that $\pi = 0$ holds if the run equilibrium does not exist for a particular contract and $\pi = 1$ holds if the no-run equilibrium does not exist.

---

[13] We do not allows the bank to purchase existing capital. However, the extreme illiquidity of this asset will make it a dominated strategy in equilibrium.

[14] Our contract is not fully optimal. Analyzing the optimal contract with a continuum of agents is difficult, in part because increasing the consumption of an individual agent is costless. Peck and Shell [19] have a finite number of agents and show, in an environment similar to ours, that there are runs in equilibrium. Because we study a competitive economy, the number of agents in our model must be large, and having a continuum simplifies the analysis. Our simple contracts lead to equilibrium bank runs in the spirit of Peck and Shell [19], and we believe that the gain from allowing for more complex contracts would be small.

[15] See Ennis and Keister [13] for a detailed discussion of the equilibrium selection mechanism approach.

More generally, an ESM assigns a probability distribution over the set of equilibria of the post-deposit game to each possible deposit contract. This definition is more general than the standard sunspots approach in that the probability can vary continuously with the parameters of the contract, instead of being a fixed constant. The specification is simplified here by the fact that there are at most two (symmetric, pure strategy) equilibria, and hence the support of the distribution has at most two points. We should emphasize that while we speak of "selecting" between the two equilibria, the actual allocations in each equilibrium depend on the contract chosen by the bank. Hence the resulting stochastic equilibrium allocations are *not* (mere) randomizations over the allocations in the certainty equilibria. The possibility of a bank run will affect the contract chosen by the bank and will therefore affect the growth rate even if a run does not actually occur.

The feasibility constraints on the ESM are simply a way of including two natural constraints on the bank's problem. As is common in the literature, the bank will be able to choose a contract such that the run equilibrium does not exist. For such contracts, it recognizes that $\pi = 0$ holds and hence the no-run equilibrium will obtain. In addition, if the bank is not careful the contract might be such that patient agents always prefer to run; this would usually be referred to as a violation of the incentive compatibility constraint (which requires patient agents to prefer to wait if other patient agents are waiting). In this case the ESM would deliver $\pi = 1$, and the bank would recognize that such a contract could only lead to a run. Hence the incentive compatibility constraint is naturally embedded in the ESM approach. Beyond these two restrictions, the function π reflects the properties of the equilibrium selection process, whatever that may be. Note that the value of π cannot depend on u_t, since u_t is not known when agents choose their strategies in period t.[16]

2.3 The bank's problem

We assume that there is free entry into banking, so that competition will drive a bank to maximize the expected utility of depositors. The problem of the bank formed in period t can then be written as

$$
\max_{a_{1,t},\eta_t} \ \pi\left(a_{1,t},\eta_t\right) \int_0^1 \frac{\overline{\overline{u}}_t}{\gamma} \left(ub_1 + (1-u)\,b_2\right) a_{1,t}^\gamma f\left(u\right) du \ +
$$
$$
\left(1 - \pi\left(a_{1,t},\eta_t\right)\right) \left(\int_0^{\overline{\overline{u}}_t} \frac{b_1 u}{\gamma} a_{1,t}^\gamma f\left(u\right) du + \int_{\overline{\overline{u}}_t}^1 \frac{b_1 \overline{\overline{u}}_t}{\gamma} a_{1,t}^\gamma f\left(u\right) du + \int_0^{\overline{\overline{u}}_t} \frac{b_2\left(1-u\right)}{\gamma} a_{2,t}^\gamma f\left(u\right) du \right)
$$

subject to

$$
\overline{u}_t = \max\left(\frac{\eta_t n}{a_{1,t}}, 1 \right)
$$

---

[16] The probability could, however, depend on past realizations of u. We do not investigate this possibility here.

$$\overline{\overline{u}}_t = \max\left(\frac{\eta_t n + (1-\eta_t)\,x}{a_{1,t}}, 1\right)$$

$$a_{1,t} \geq 0$$

$$0 \leq \eta_t \leq 1$$

and

$$a_{2,t} = \left\{ \begin{array}{c} \frac{\eta_t n^2 + (1-\eta_t)\psi_{t+1} - nua_{1,t}}{1-u} \\ \frac{\eta_t n + (1-\eta_t)x - ua_{1,t}}{1-u}\psi_{t+1} \end{array} \right\} \text{ as } \left\{ \begin{array}{c} u \leq \overline{u}_t \\ \overline{u}_t \leq u \leq \overline{\overline{u}}_t \end{array} \right\}.$$

Notice that the form on the utility function is such that the size of the deposit made by the agent (which will equal the wage) cancels out of the objective and hence does not matter for the solution. This is important because, as the economy grows, wages and therefore deposits will grow. We see here that this growth does not affect the bank's decision problem.

Before describing the objective function, we introduce some additional notation and describe the constraints. Let $\alpha_t^s(u)$ denote the fraction of the stored goods that are paid out in period t in the no-run equilibrium. Similarly, let $\alpha_t^i(u)$ denote the fraction of investment that is liquidated early and paid out in period t in the no-run equilibrium. Both of these fractions depend on the contract $(\eta_t, a_{1,t})$. Because $x < n < R$ holds, the bank will never choose to pay agents withdrawing in period t with liquidated investment when stored goods are available. In other words, $\alpha_t^i > 0$ implies $\alpha_t^s = 1$. Hence the first constraint defines $\overline{u}_t$ as the value of u at which, for the given values of η_t and $a_{1,t}$, all stored goods have been given to withdrawing agents but no investment has been liquidated ($\alpha_t^s(\overline{u}_t) = 1$ and $\alpha_t^i(\overline{u}_t) = 0$). The second constraint defines $\overline{\overline{u}}_t$ as the point at which all investment has been liquidated and the bank has just run out of resources ($\alpha_t^s(\overline{\overline{u}}_t) = \alpha_t^i(\overline{\overline{u}}_t) = 1$). Hence if the realization of u is $\overline{\overline{u}}_t$ (or greater), agents arriving at the bank in period $t+1$ will receive nothing. The next two constraints are obvious bounds on the choice variables, and the final constraint simply says that the resources remaining in the bank at period $t+1$ are divided equally among the agents choosing to withdraw in that period.

Turning to the objective function, the first term gives the expected utility of an agent in the event of a bank run (and therefore is multiplied by the probability of a run π). Because the agent's place in line is random and the first $\overline{\overline{u}}_t$ depositors to arrive are served during a run, $\overline{\overline{u}}_t$ also gives the individual probability of being served. With probability u the agent will truly be impatient and hence have preference parameter b_1, while with probability $(1-u)$ she will be patient and will have parameter b_2. All agents who are served receive the rate of return $a_{1,t}$ on their deposits. The remaining term in the objective function gives the expected utility of a depositor when there is no run (and therefore is multiplied by $(1-\pi)$). If the fraction of patient agents is less than $\overline{\overline{u}}_t$, the bank will not run out of resources and the agent will receive return $a_{1,t}$

9

if she is impatient. If the fraction of impatient consumers is above $\overline{\overline{u}}_t$, however, she will receive $a_{1,t}$ with probability $\left(\overline{\overline{u}}_t/u\right) < 1$ and zero otherwise. Finally, with probability $(1 - u)$ she will be patient and receive the return $a_{2,t}$ as given in the final constraint. Note that if the fraction of impatient agents is above $\overline{\overline{u}}_t$, the bank runs out of resources in the first period and patient agents receive nothing.

This problem is difficult to address analytically, but an upper bound for $a_{1,t}$ can clearly be chosen large enough to not be binding. Therefore, as long as the function π is well-behaved, the problem is one of maximizing a continuous function over a compact set, and a solution must exist. We use $\left(a_{1,t}^*, \eta_t^*\right)$ to denote this solution.

Of course, type II agents always have the alternative of choosing not to participate in the banking system. In this case, the agent will store some goods and invest the rest in capital formation. The payoff for an agent taking this route is given by

$$\max_{\eta_t} \int_0^1 \frac{w_t^\gamma}{\gamma} \left(ub_1 \left[\eta_t n + (1 - \eta_t)x\right]^\gamma + (1 - u) b_2 \left[\psi_{t+1}(1 - \eta_t) + \eta_t n^2\right]^\gamma\right) f(u)\, du$$

subject to

$$0 \leq \eta_t \leq 1$$

If the probability of a bank run is very high, the agent will choose to remain in autarky. For a small enough probability of a run, however, it is not hard to see that the agent will prefer to deposit in the bank in order to insure against the preference shock. For some parameter values, the payoff from autarky is lower than the payoff that the bank can achieve designing a contract that does not allows for runs. In this case, the agent always chooses to participate in the banking system (this is the case for all of the examples in this paper).

3 Equilibrium

We now turn to the analysis of the equilibrium behavior of the economy. We first impose the market-clearing conditions and derive the equilibrium law of motion for the capital stock as a function of the contract offered by the bank. We then turn to some numerical examples in which we compute the solution to the bank's problem and simulate the equilibrium behavior of the economy.

3.1 Market clearing and aggregate investment

Firms are competitive and therefore factors are paid their marginal products

$$w_t^* = (1 - \theta)\overline{k}_t^{1-\theta} k_t^\theta$$

$$r_t^* = \theta \bar{k}_t^{1-\theta} k_t^{-(1-\theta)}.$$

In equilibrium all firms will choose the same capital-labor ratio, and therefore $\bar{k}_t = k_t$ must hold. The equilibrium wage and rental rate therefore reduce to

$$w_t^* = (1-\theta) k_t$$
$$r_t^* = \theta.$$

The marginal product of capital is constant in the Ak model, which is why the economy is always on a balanced growth path.

In the market for existing capital, supply is given by $(1-\delta) k_t$. If the price of capital q_t were greater than $\frac{1}{R}$, we have from (2) that demand for capital would be zero because new investment would yield a higher return. Therefore the market would not clear. Suppose instead that q_t were below $\frac{1}{R}$. Then from (2) we have that type I agents put all of their income into existing capital, so that total demand for existing capital would be equal to

$$\frac{\phi (1-\theta) k_t}{q_t}.$$

The market-clearing price would then be given by

$$q_t = \phi \frac{(1-\theta)}{1-\delta}.$$

We assume

$$\phi \geq \frac{1-\delta}{1-\theta} \left(\frac{1}{R} \right) \tag{3}$$

holds, which implies that the candidate price above is at least $\frac{1}{R}$, contradicting our original supposition. The role of type I agents in this model is to hold the stock of existing capital (which is completely illiquid) between periods. The assumption in (3) is simply that there are enough type I agents in the economy to prevent existing capital from trading at a discount. (Otherwise, if the discount were large enough, banks might want to invest in existing capital as well.) The only remaining possibility is then $q_t = \frac{1}{R}$. Under condition (3) this price clears the market because the demand for capital is perfectly elastic (and large enough). We state this result as a proposition.

Proposition 1 *Assume condition (3). Then the equilibrium price of capital is given by*

$$q_t = \frac{1}{R} \quad \text{for all } t.$$

Since we have shown both r_t and q_t to be constant over time, the bank's problem is exactly the same in

every period and the solution (η^*, a_1^*) will be independent of time. In other words, the bank will offer the same deposit contract in every period.

We use i_t^{I} to denote investment in new capital made by an individual type I agent. This must be equal to the income of the agent less her purchases of existing capital. Using the market-clearing condition for existing capital, we can write the total investment in new capital by type I agents as

$$
\begin{aligned}
\phi i_t^{\mathrm{I}} &= \phi w_t - q_t \left(1 - \delta\right) k_t \\
&= \left(\phi \left(1 - \theta\right) - \frac{1 - \delta}{R}\right) k_t.
\end{aligned}
$$

Finally, we need to calculate how much new investment is undertaken by banks. The amount of investment per type II agent is given by

$$
i_t^{\mathrm{II}} \left(u_t\right) = \left(1 - \eta\right) w_t \left(1 - \alpha^i \left(u_t\right)\right). \tag{4}
$$

In other words, new capital formation depends on the amount the bank invests and the fraction of that investment that is not liquidated prematurely. Note that this quantity depends on the realization of u_t. If the realization is high, so that a large fraction of the population is impatient, the bank will need to liquidate a large fraction of their investment in order to meet the demand for withdrawals. This will result in a lower capital stock in the following period. Hence the aggregate uncertainty about liquidity demand leads to equilibrium fluctuations in the growth rate of the capital stock and of output.

The law of motion for the capital stock is given by

$$
k_{t+1} = \left(1 - \delta\right) k_t + R \left(\phi i_t^{\mathrm{I}} + \left(1 - \phi\right) i_t^{\mathrm{II}} \left(u_t\right)\right).
$$

Since both i^{I} and $i^{\mathrm{II}} \left(u_t\right)$ are linear functions of k_t we have the following result.

Proposition 2 *For any period t and stock of capital k_t, the equilibrium growth rate of capital (k_{t+1}/k_t) is a random variable $g(u)$ independent of t. Furthermore, we have*

$$
g(u) \equiv R \left(1 - \theta\right) \left[\phi + \left(1 - \phi\right) \left(1 - \eta\right)(1 - \alpha^i \left(u\right))\right], \tag{5}
$$

where u is a random variable with distribution $f(u)$ when there is no run and $u = 1$ when there is a run. The function $\alpha^i \left(u\right)$ is given by

$$
\alpha^i \left(u\right) = \left\{
\begin{array}{ll}
0 & if \quad u < \overline{u} \\
\frac{u a_1 - \eta n}{(1 - \eta) x} & if \quad \overline{u} < u < \overline{\overline{u}} \\
1 & if \quad u > \overline{\overline{u}}
\end{array}
\right. .
$$

In summary, the difference equation describing the dynamic behavior of k_t is linear and stochastic. Notice that since η and the function α^i are the same in every period, the growth rate is a time-invariant, weakly

12

decreasing function of the realization of u. In equilibrium, aggregate output is given by $Y_t = k_t$ and hence the growth rate of output is the same as the growth rate of the capital stock in this economy.

Properties of this difference equation are difficult to derive analytically, because of the complexity of the bank's problem. We now look at some examples for which we can compute solutions and simulate the equilibrium behavior of the economy.

3.2 Implications for growth

In this subsection we investigate the growth implications of the possibility of bank runs by computing a representative example of the model presented in the previous section. In particular, consider the utility function (1) with the following parameter values: $\gamma = 0.4$, $b_1 = 2.5$ and $b_2 = 1$. We take the capital share of income θ to be equal to 0.4 and a 20% depreciation rate ($\delta = 0.2$). The return on storage n is set equal to 1 and the return on investment is given by the pair $(R = 3, \ x = 0.3)$.[17] We set the liquidation value of investment x relatively low because for values of x closer to n the bank finds liquidated investment a not-too-costly instrument for providing consumption to impatient agents. We want to make a clear distinction between storage (which yields consumption goods) and investment (which yields capital), and a low value of x is useful for this purpose. The total return on investment when not liquidated early is then given by

$$\psi = R\left(r + (1 - \delta)q\right) = 2.0$$

We assume that the value of u, the proportion of impatient agents in the population, is an independent and identically distributed random variable with a distribution function given by a beta distribution with parameters $(3, 9)$. The mean value of u is 0.25 and the standard deviation is around 0.12. If the variance of u is set very high, the effects of bank runs will not be substantially different from the effects of regularly occurring high liquidity demand shocks. Bank runs, however, are extreme events and the way to capture this in the present model is to assume a relatively low variance of u.

Given these parameter values and an ESM, we solve the bank's problem numerically.[18] We first consider the standard sunspots story, where the probability of a run is a fixed number for all deposit contracts under which both equilibria exist. It should be kept in mind that for a sufficiently high value of the probability of a run, the bank might want to choose a contract that rules out a run as an equilibrium outcome (i.e., has $\pi = 0$). In the present example, if the exogenous probability of a bank run is above $10\ \%$ the bank will choose a contract that does not admit a run equilibrium.[19] Furthermore, the payoff of autarky is smaller

---

[17] For values of $n < 1$ patient consumers would rather consume early in their life than use the liquid technology for saving.

[18] All calculations are done in Fortran. The source code is available from the authors upon request.

[19] The best contract that prevents runs has some interesting characteristics. For the parameter values being considered (and for a

than the payoff from this run-proof contract and hence type II agents always choose to participate in the banking sector.

In Table 1 we present the solution of the bank's problem for different levels of the probability of a bank-run. Note that the fraction of liquid assets in the bank's portfolio is increasing in the probability. As the probability of a run increases, it becomes more likely that the bank will have to liquidate investment early. But since the liquidation value is relatively low, the bank prefers to hold more of the liquid asset (storage) to deal with a run if it occurs.

Table 1

| Prob. of Run (π)	a_1	η	Prob$[u \geq \overline{u}]$	Prob$[u \geq \overline{\overline{u}}]$
0.00	1.124	0.444	0.126	0.017
0.04	1.120	0.456	0.110	0.014
0.06	1.116	0.462	0.101	0.013
0.08	1.112	0.468	0.093	0.012

The last two columns of the table show the probability that the bank will liquidate some investment early (Prob$[u \geq \overline{u}]$) and the probability that the bank will have to suspend payments in period t after running out of resources (Prob$[u \geq \overline{\overline{u}}]$), both conditional on the bank not suffering a run. We can see that as the probability of a run increases, the bank chooses a more liquid portfolio, thereby reducing the chances that it will have to liquidate investment early in the no-run situation (from over 0.12 to around 0.09). The higher level of liquidity also implies that the bank will have to suspend payments less often when there is no run. However, note that the reduction in the probability Prob$[u \geq \overline{\overline{u}}]$ is not enough to compensate the increase in π. In other words, the unconditional probability of suspension increases with the probability of a run.

As a benchmark case, we take the probability of a bank run to be equal to 0.06 (when both equilibria exist). Figure 1 shows the time series of the logarithm of the stock of capital for this economy in a representative 50 period simulation. It shows how a bank run causes an abrupt fall in the level of capital formation in the economy. However, not all of the major downturns in this figure are due to bank runs. In periods where the number of impatient agents u_t is very high, the bank liquidates some investment early to pay these agents, and this liquidation creates some of the observed fluctuations in the stock of capital. There are four bank runs during these fifty periods, at $t = 11, 29, 36$, and 38. There is no bank-run, for example, at $t = 17$ and 44.[20]

---

wide range around them) the "run-proof" contract has a relatively low level of storage in the bank's portfolio and a relatively low return on early withdrawals.

[20] During the 50 years of the National Banking Era (1863-1914) there were five major bank panics: 1873, 1884, 1890, 1893, and 1907. The Federal Reserve System was established after that, partly as a response to those regular periods of crisis.

Figure 1
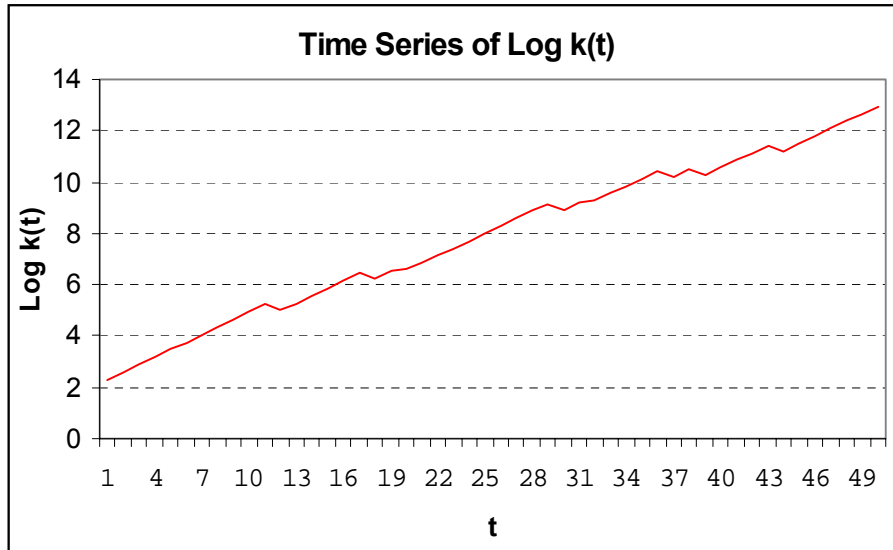
In Table 2 we present the growth rate of capital (which in this economy is the same as the growth rate of output) conditional on not having a run as well as the unconditional growth rate. This numbers are the result of 20 simulations of 50 periods each. We consider the periods in this overlapping generations economy to represent approximately 5 years and we report the annual growth rates.[21]

Table 2

Prob. of Run	Growth Conditional on No Run	Unconditional Growth
0.00	5.77%	5.77%
0.04	5.65%	5.37%
0.06	5.61%	5.15%
0.08	5.55%	4.93%

Note that the unconditional growth rate of the economy tends to be lower for economies with a higher probability of a bank-run. Two effects combine to generate this fact. First, the average growth rate of the economy when there is no run (the conditional-on-no-run growth rate) is lower because the bank places more resources in storage and fewer in investment. This is the effect isolated in the middle column of the table. Second, bank runs induce early liquidation when they occur, and that also reduces capital formation

---

[21] Miron ([18]) studies banking panics in the US during the period 1890-1908. He estimates that the probability of a financial panic in any given year was around 0.30. We are considering much lower probabilities of sunspot-driven runs in our computations. However, two factors make our numbers reasonable. First, the banking system in our model may experience distress due to unusually high levels of the proportion of impatient agents u. In fact, conditional on no run, there is around a ten percent probability of early liquidation of investment in our calibration (see Table 1). Some of these events would be included in Miron's definition of a panic. Second, the period studied by Miron seems to be a period with unusually high frequency of runs. Miron also reports the growth rate of output during these 18 years. The growth rate conditional on no run was 6.82% and the unconditional growth rate was 3.75% (see our Table 2 for comparison).

on average. Both of these effects are detrimental for long run growth.

So far we have taken the standard sunspots approach to equilibrium selection, where the bank takes the probability of a run as exogenously given and believes that it cannot influence the probability by changing the composition of its portfolio. However, it seems intuitively plausible that the portfolio chosen by the bank might actually be a useful predictor of the likelihood of a run. In other words, if the bank chooses a contract that is "closer" to a contract that would eliminate the run equilibrium, it seems reasonable to think that the probability of having a run might go down. In the next section we describe an alternative approach that yields an equilibrium selection mechanism with this property, and we investigate the growth implications of the new mechanism.

4   Risk Factor Based Equilibrium Selection

In this section we examine a more general ESM where the probability of a run can vary continuously with the deposit contract chosen by the bank. It seems entirely reasonable to think that the outcome of the post-deposit game can depend on the relative payoffs obtained by agent in the alternative scenarios, even when both outcomes are Nash equilibria. This is, in fact, precisely the idea that motivates the use of risk dominance as an equilibrium selection mechanism (Harsanyi and Selten [16]).[22] However, risk dominance selects a single equilibrium for each contract that the bank could choose. In other words, whereas the sunspots approach assigns a fixed probability to the run equilibrium (whenever it exists), risk dominance assigns a probability of either zero or one, depending on the contract. We find this unappealing because it implies that the bank can rule out runs entirely by choosing a contract that makes the no-run equilibrium barely risk dominant (while a very similar contract would lead to a run with certainty). We find a probabilistic approach much more realistic. In our approach, a run is more likely to occur when the equilibrium is risk dominant, but still can occur when it is not. In other words, we keep the idea that under certain circumstances a bank run is, to some extent, a chance event. However, the likelihood of this chance event now depends on the actions of the bank. Starting from any contract that permits a run equilibrium, slightly reducing the relative payoff of running will slightly reduce the probability of a run. Of course, there are still contracts for which only a bank run can happen, or for which a run cannot happen; these are the situations where only one of the equilibria of the post deposit game exists. However, we believe that when both equilibria exist, each one can obtain and therefore should be assigned positive probability. Hence our approach retains the probabilistic property of the sunspots approach while allowing the portfolio decision of the bank

---

[22] See Temzelides [20] for an evolutionary justification of using risk dominance as the equilibrium selection mechanism in a model with bank runs. Catalan [9] uses the risk dominance approach to study the relationships between the bank's balance sheet, deposit insurance, and the likelihood of runs.

to influence the probability of a run.

Specifically, we assume that the probability of a run π is a decreasing function of the risk factor of the run equilibrium. We begin by defining the risk factor.[23]

Definition: The *risk factor of the run equilibrium* is the smallest probability ρ such that if a patient agent believes that all other agents will run with probability strictly greater than ρ, then running is her unique optimal action.

The risk factor is therefore determined by the following expression,

$$\rho \bar{\bar{u}} a_1^\gamma + (1 - \rho) \left(\int_0^{\bar{\bar{u}}} a_1^\gamma f(u) du + \int_{\bar{\bar{u}}}^1 \frac{\bar{\bar{u}}}{u} a_1^\gamma f(u) du \right) = (1 - \rho) \int_0^{\bar{\bar{u}}} a_2^\gamma(u) f(u) du.$$

The left-hand side of this expression is the expected value for a patient agent of running to the bank when she believes that with probability ρ everybody else will run. The right-hand side is the expected value of not running given the same belief. The expression says that if a patient agent assigns probability ρ to the event of a run on the bank, she is indifferent between running and not running. If she assigns a higher probability to a run, she would strictly prefer to run. We assume that the higher the risk factor of the run equilibrium, the lower the probability of a run. Whatever determines the individual agent's prior belief about the possibility of a run on the banking system, the higher the risk factor ρ, the lower the likelihood that this belief will be greater than ρ, and hence the lower the likelihood that the agent would decide to run.

Holding other things constant, the risk factor of the run equilibrium is decreasing in the return offered on period t withdrawals a_1. Higher values of a_1 increase the incentive for agents to withdraw their funds from the bank early. They also make not running less attractive because even if there is no run, the bank will have fewer resources in period $t + 1$ and hence a_2 will be lower. The relationship between the risk factor and η, the fraction of the bank's portfolio that is in storage, is not monotonic. If η is very low, then with high probability the realization of u_t will be such that investment is liquidated. In such a situation, increasing η decreases the amount of liquidation and therefore increases the amount of resources available in the second period. This makes waiting a more attractive strategy, and as a result the risk factor is increasing in η. If, on the other hand, η is very high, then with high probability the realization of u_t will be such that no investment is liquidated. In this case decreasing η would increase the resources available in the second period, and hence the risk factor is decreasing in η. For moderate values of η, the risk factor is fairly flat and the effects of a change in a_1 will typically dominate the effects of a change in η.

In earlier work, we have shown how an adaptive learning process in a stochastic environment naturally

---

[23] See Young [22] for an extended discussion of risk factors and risk dominance.

generates an equilibrium selection mechanism in which the probability of an equilibrium is strictly decreasing in its risk factor (see Ennis and Keister [12]). Because our goal here is to examine the basic implications that follow from this approach, the exact relationship between the risk factor of the run equilibrium and the probability of a run is not very important. Rather than specifying a learning model, we posit a simple linear relationship of the form

$$\pi(\rho) = m - h \cdot \rho,$$

where m and h are constants that allow us to calibrate the equilibrium probability π to a reasonable number.[24] We are interested in studying the growth implications of the fact that the bank recognizes that its choice of contract influences the likelihood of a run. For this we compare the outcomes in two different economies. In one, banks are sophisticated decision makers and realize the endogeneity of the likelihood of runs. In the other, banks are naive and take the probability as given even though their decisions do influence it. For this second case we concentrate on the rational expectations equilibrium (i.e., the probability taken as given by the bank is confirmed by their decision, given the equilibrium selection function $\pi(\rho)$). We continue our analysis using the example introduced in the previous section. We assume that $m = 0.1$ and $h = 0.05$.

Table 3 shows the solution to the problem for both the sophisticated and the naive bank.

Table 3

	a_1^*	η^*	π^*	Growth Conditional on No Run	Unconditional Growth
Sophisticated Bank	0.995	0.415	0.085	6.47%	5.77%
Naive Bank	1.110	0.473	0.091	5.51%	4.78%

As expected, the sophisticated bank chooses a contract that implies a higher risk factor for the run equilibrium and thereby lowers the probability of a run π^*. The bank does this by lowering a_1 and η. These adjustments in the contract may seem not very intuitive at first but they are a direct consequence of the equilibrium selection process we are considering. By lowering a_1 and η, the bank lowers the contingent payoff from running to the bank and increases the payoff of waiting to withdraw in the second period of life. The lower value of a_1 allows the bank to put more resources into investment without having to liquidate more often. This leads to a larger return for agents who arrive in the second period. Notice that this is the opposite of the narrow-banking proposal of Friedman [14]. Friedman argued that demand deposits should

---

[24]   If the slope of the function π is very high the bank will have an incentive to make ρ high enough to eliminate the runs. We are interested in studying situations where runs are possible and we hence calibrate the value of h to be relatively small. Furthermore, it is quite possible that for large changes in ρ the linearity assumption about the equilibrium selection function $\pi(\rho)$ is not really appropriate. Finally, note that when $h = 0$ holds, this equilibrium selection mechanism reduces to the standard sunspots approach studied earlier.

be backed entirely by safe, short term assets (such as storage in our model). This would enable the bank to meet all of their obligations during a run, which would in turn prevent a run from happening. The problem with this approach is that preventing the bank from undertaking investment is costly because investment offers a much higher return than storage. (See Wallace [21].) Our analysis shows that a better approach for a bank facing the possibility of a run is to structure the contract to reward agents heavily for waiting. This involves holding *fewer* liquid assets and putting more resources into (illiquid) investment. In addition, the narrow-banking proposal is aimed at eliminating runs entirely. We see here that, at least from the standpoint of the welfare of generation t depositors, it is not optimal for the bank to eliminate the possibility of a run. The bank chooses to reduce the probability, but not all the way to zero.

The last two columns give the implications of these differences for the growth rate of the economy. The economy with sophisticated banks has both a higher level of investment and a lower return on early withdrawals. These two facts tend to increase the growth rate of the economy in periods with no runs. In such periods, the economy with sophisticated banks grows at an average rate that is 89 basis point higher than the economy with naive banks. Furthermore, the economy with sophisticated banks has a lower equilibrium probability of bank runs. Bank runs are detrimental to the process of capital formation and hence reduce the average growth rate of the economy. When we take the lower frequency of runs into account, the economy with sophisticated banks grows on average around 92 basis points faster than the economy with naive banks.

In tables 4 and 5 we further study the economy with sophisticated banks. Table 4 shows that when h (the sensitivity of the equilibrium selection function to the risk factor of the run equilibrium) is higher, the bank chooses less liquidity and a lower return on early withdrawals. These two changes have the direct effect of increasing capital formation and hence increasing the growth rate of the economy. In addition, the change in the contract reduces the probability of bank runs, and this further increases the long-run average growth rate. In other words, the more influence the bank's portfolio has on the likelihood of runs, the faster the economy will grow.

Table 4

(m, h)	a_1^*	η^*	π^*	Growth
$(0.1, 0.05)$	0.995	0.415	0.085	5.77%
$(0.1, 0.06)$	0.980	0.405	0.082	5.93%

In Table 5 we present the equilibrium outcome for different values of the return on the investment technology and the liquidation cost.

Table 5

(R, x)	a_1^*	η^*	π^*	Growth
$(3.00, 0.3)$	0.995	0.415	0.085	5.77%
$(3.25, 0.3)$	0.971	0.398	0.084	7.64%
$(3.00, 0.4)$	1.015	0.378	0.086	6.12%

The growth rate is higher for higher values of R for several reasons. The first is apparent from (5). When a fixed amount of investment yields more capital, the economy will grow faster. The second reason is also standard: when investment offers a higher return, banks will choose to invest more. However, in our model there is also a third effect. An increase in investment implies that banks are giving a relatively higher payoff to agents in the second period, and hence the risk factor of the run equilibrium is higher. Thus the change in the bank's portfolio also decreases the probability of a bank run, which increases the long run average growth rate even more.

In economies where the liquidation cost of investment is lower (higher values of x), the bank would again choose to place more resources in investment. However, the bank would increase the return for early withdrawals because early liquidation is less costly. This second effects tends to decrease the risk factor of running and hence in this case the probability of a run increases. Here the two effects are pointing in opposite directions with respect to long-run growth: more investment tends to raise the average growth rate, while a higher frequency of bank runs lowers it. In our example the first effect dominates, because the movement in η is much larger than the change in π. As a result, the long-run average growth rate goes up.

5 Concluding Remarks

We conclude by offering some remarks on possible extensions and modifications of the model that could be used to address some additional issues regarding the relationship between growth and banking crises.

Richer Models of Growth

In order to address the issue of how the economy responds in the short run to a crisis, our banking model could be embedded in a richer growth model where transitional dynamics play a role. The model in Jones and Manuelli [17] and the model of inventive activity in Antinolfi, Keister, and Shell [1] may be promising candidates. In both cases, the marginal product of capital would increase in the period after the run because of the decrease in the stock of capital. This would lead banks to shift their portfolios toward investment, which would increase the growth rate during the transition. In addition, our analysis above shows that this change in portfolios would decrease the probability of a run (this is similar to the middle line of Table 4), which would on average lead to a faster recovery. As the economy approaches a balanced growth path,

the marginal product of capital would gradually fall and banks would gradually shift their portfolios back toward more storage. Hence the probability of a run would increase back to its original level. The main obstacle to performing this analysis is computational. Because the return on capital would constantly be changing, the bank's problem would need to be solved separately in each period.

The Welfare Cost of Bank Runs

Another potentially interesting extension would be an analysis of the welfare cost of bank runs. In our model, as in Peck and Shell [19], allowing bank runs to occur with positive probability can be part of the (constrained) optimal risk-sharing arrangement among depositors in generation t. That is, taking the equilibrium selection mechanism as given, the utility of these depositors can be higher under a contract that allows runs to occur than under any contract that does not allow them. However, much of the cost of a bank run is external to this group. As is evident from Figure 1, a bank run has a permanent effect and the capital stock and real output never recover to the original trend. Therefore all future generations suffer as a consequence of a run today. The question then arises whether bank runs can occur with positive probability in a (constrained) Pareto optimal allocation in a dynamic model.[25] The overlapping generations framework seems well suited to address this issue because it is, as Cass and Shell [8] argue, "the only genuinely dynamic, basically disaggregative framework available" for macroeconomic analysis (p.260). In particular, the external effect arises because the model recognizes that a new set of agents will populate the economy in the future, and today's depositors do not take the welfare of these agents into account. It seems possible that the long-run damage caused by a bank run is so large that allowing runs with positive probability cannot be optimal in the dynamic setting.

The Correlation Between Bank Runs and Liquidity Demand

The seasonality of liquidity demand and its implications for the timing of bank runs has received a fair amount of attention in the banking literature. Miron [18] estimates that bank panics in the U.S. prior to the establishment of the Federal Reserve System were substantially more likely to occur in seasons with high loan demand or low deposit demand. One way of investigating this issue in our model would be to have the distribution of the number of impatient agents (represented by the function f) cycle over time between, say, one distribution with a relatively low mean and another with a higher mean. In this way liquidity demand would have a clear seasonal pattern over time but remain random. The bank would choose different portfolios under the different distributions, and the incentive for an agent to run (as measured by the risk factor of the run equilibrium) would reflect both the change in f and the portfolio change. It would

---

[25] One difficulty with performing this analysis is finding the proper concept of Pareto optimality to use in this setting. See Balasko and Shell [2] for a discussion of different concepts of optimality in overlapping generations models.

be interesting to see if such a setup could deliver the stylized fact that runs are more common in periods where liquidity demand is expected to be high.

An alternate approach would be to change the timing of events in the model so that agents gain some information about the number of impatient depositors u_t before they decide whether or not to go to the bank in period t. The bank's portfolio choice would be made before this information is revealed and therefore would remain constant over time. The information about u_t would certainly be useful for an agent in deciding whether or not to run – a high value of u_t implies that the bank will need to liquidate investment to pay depositors, which implies a lower return for agents who wait to withdraw in period $t+1$. In this way the effect of high liquidity demand on the incentive for an agent to run could be isolated, and it seems likely that the correlation between liquidity demand and the probability of a run would be positive. Of course, studying seasonality requires interpreting a time period as being fairly short (less than a year). Hence some modification of our overlapping generations setup might be desirable for such an analysis.

References

[1] G. Antinolfi, T. Keister, and K. Shell, Growth dynamics and returns to scale: Bifurcation analysis, *J. Econ. Theory* **96** (2001), 70-96.

[2] Y. Balasko and K. Shell, The overlapping generations model, I: The case of pure exchange without money. *J. Econ. Theory* **23** (1980), 281-306.

[3] V. Bencivenga, and B. D. Smith, Financial intermediation and endogenous growth, *Rev. Econ. Stud.* **58** (1991), 195-209.

[4] J. H. Boyd, P. Gomis, S. Kwak, and B. D. Smith, A user's guide to banking crises, mimeo., January 2001.

[5] J. H. Boyd, S. Kwak, and B. D. Smith, The real output losses associated with modern banking crises, or "The good, the bad, and the ugly," mimeo., February 2002.

[6] G. Caprio Jr. and D. Klingebiel, Bank insolvencies: Cross country experience, World Bank Policy Research Working Paper 1620, July 1996.

[7] D. Cass and K. Shell, Do sunspots matter? *J. Pol. Econ.* **91** (1983), 193-227.

[8] D. Cass and K. Shell, In defense of a basic approach, *in* "Models of Monetary Economies" (J. Kareken and N. Wallace, Eds.), pp. 251-260, Federal Reserve Bank of Minneapolis, 1980.

[9]   M. Catalan, Endogenous bank runs, mimeo., 2000.

[10] R. Cooper and T. W. Ross, Bank runs: Liquidity costs and investment decisions, *J. Monet. Econ.* **41** (1998), 27-38.

[11] D. W. Diamond and P. H. Dybvig, Bank runs, deposit insurance, and liquidity, *J. Pol. Econ.* **91** (1983), 401-419.

[12] H. M. Ennis and T. Keister, Optimal policy with probabilistic equilibrium selection, FRB of Richmond WP # 01-03, 2001.

[13] H. M. Ennis and T. Keister, Aggregate demand management and equilibrium selection, mimeo., 2002.

[14] M. Friedman, "A program for monetary stability," Fordham University Press, New York, 1959.

[15] X. Freixas and J.-C. Rochet, "Microeconomics of banking," MIT Press, Cambridge, MA, 1997.

[16] J. Harsanyi, and R. Selten, "A General Theory of Equilibrium Selection in Games," MIT Press, Cambridge, MA, 1988.

[17] L. Jones and R. Manuelli, A convex model of equilibrium growth: Theory and policy implications, *J. Pol. Econ.* **98** (1990), 1008-1038.

[18] J. A. Miron, Financial panics, the seasonality of the nominal interest rate, and the founding of the Fed, *Amer. Econ. Rev.* **76** (1986), 125-140.

[19] J. Peck and K. Shell, Equilibrium bank runs, CAE Working Paper # 01-10R, Cornell University, 2001.

[20] T. Temzelides, Evolution, coordination, and banking panics, *J. Monet. Econ.* **40** (1997), 163-183.

[21] N. Wallace, Narrow banking meets the Diamond-Dybvig model, *Fed. Res. Bank Minn. Quart. Rev.* **20** (1996), 3-13.

[22] H. P. Young, "Individual Strategy and Social Structure: An Evolutionary Theory of Institutions," Princeton University Press, Princeton, NJ, 1998.