Working Paper Series

Bank Runs and Investment Decisions Revisited

WP 04-03

Huberto M. Ennis
Federal Reserve Bank of Richmond

Todd Keister
Centro de Investigación Económica

THE FEDERAL RESERVE BANK OF RICHMOND
RICHMOND ▪ BALTIMORE ▪ CHARLOTTE

Bank Runs and Investment Decisions Revisited *

Huberto M. Ennis
Research Department, Federal Reserve Bank of Richmond

Todd Keister
Centro de Investigación Económica, ITAM

Federal Reserve Bank of Richmond Working Paper No. 04-03
March 2004

Abstract

We examine how the possibility of a bank run affects the deposit contract offered and the investment decisions made by a competitive bank. Cooper and Ross (1998) have shown that when the probability of a run is small, the bank will offer a contract that admits a bank-run equilibrium. We show that, in this case, the bank will chose to hold an amount of liquid reserves exactly equal to what withdrawal demand will be if a run does not occur. In other words, precautionary or "excess" liquidity will not be held. This result allows us to determine how the possibility of a bank run affects the level of illiquid investment chosen by a bank. We show that when the cost of liquidating investment early is high, the level of investment is decreasing in the probability of a run. However, when liquidation costs are moderate, the level of investment is actually *increasing* in the probability of a run.

JEL Classification Numbers**:** D84, E44, G21

Keywords**:** Bank runs; deposit contracts; bank reserves; sunspot equilibrium

---

* We thank seminar participants at the Cornell-Penn State Macroeconomics Workshop, the Federal Reserve Bank of Cleveland, ITAM, and the University of Guanajuato for useful comments. The views expressed herein are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

E-mail addresses: huberto.ennis@rich.frb.org (H. M. Ennis), keister@itam.mx (T. Keister)

1 Introduction

The direct consequences of a crisis in the banking system are fairly well understood and have been extensively documented.[1] It seems likely, however, that the mere possibility of a banking crisis will also have important macroeconomic consequences, even if a crisis does not occur. For example, suppose it is believed that with some small but positive probability, all depositors will suddenly rush to their banks and attempt to withdraw their money. This possibility will almost certainly influence the level of cash reserves that a bank will choose to hold. One of the main roles of the banking system is to perform maturity transformation, that is, to hold long-term assets while issuing short-term liabilities. If the possibility of a crisis leads banks to hold a more liquid portfolio, it will reduce the funds available for long-term investment and thereby have a substantial impact on real economic activity. In addition, the possibility of a run on the banking system will likely change the interest rate that a bank chooses to offer on deposits, thus changing the consumption profiles available to depositors. Such indirect effects of banking crises are less well understood, and are our focus in this paper. Specifically, we ask: How does the *possibility* of a bank run influence the deposit contract offered and the investment decisions made by a competitive bank?

We build on the work of Cooper and Ross [3], who first analyzed this issue in an extension of the classic Diamond and Dybvig [5] model. Diamond and Dybvig set up an environment with idiosyncratic liquidity-preference shocks and private information. They showed that by offering a contract that lets depositors choose when to withdraw their funds, a bank can implement the first best allocation in this environment as a Nash equilibrium. However, they also showed that there exists another equilibrium where all depositors attempt to withdraw their funds from the contractual arrangement early, fearing that other agents will do the same and that as a result there will be insufficient funds to cover all promised payments. In other words, the process of providing liquidity insurance makes a bank illiquid, and thereby makes it susceptible to a bank run. The Diamond-Dybvig analysis leaves open an important question: Why would an agent deposit in a bank if she expected a run to occur? Cooper and Ross [3] address this issue by introducing a sunspots-based equilibrium selection rule to the problem: If the banking contract is such that both equilibria exist, then the run equilibrium will be selected with a fixed probability q. They restrict the bank to offer *demand deposit* contracts, in which depositors are promised a fixed return and

[1] See, for example, Caprio and Klingebiel [2] and Boyd *et al*. [1].

are given this return as long as the bank has not run out of funds.[2] Banks operate in a perfectly competitive environment and hence choose the deposit contract that maximizes the expected utility of depositors. Cooper and Ross show that if the probability of a run is above some cutoff level, the bank will set the deposit contract such that waiting to withdraw is a dominant strategy for depositors who do not need liquidity. Such a contract guarantees that a bank run will not occur, and therefore is called *run proof*. However, they also show that if the probability of a bank run is below the cutoff, the bank will choose a contract that is not run proof. In this case, bank runs emerge as a truly equilibrium phenomenon; consumers are willing to deposit in the bank because the improved consumption profile that the bank offers (if a run does not occur) outweighs the possibility of losing one's deposit if a run does occur.

Our interest is in characterizing the properties of the optimal deposit contract, including the portfolio choice of the bank, and in examining how this contract changes when the probability of a bank run changes. The Cooper-Ross model has some attractive features for this purpose. First, the environment is well known and is a natural extension of that in Diamond and Dybvig [5]. We are adding nothing new in this dimension. In addition, the bank's maximization problem is conceptually easy to understand and allows one to gain a substantial amount of intuition about the various ways in which the possibility of a run changes a bank's incentives. Finally, the problem has a small number of choice variables, which allows us to give a relatively sharp characterization of the solution. In spite of this apparent simplicity, our analysis yields some surprising results. In fact, we show that some of the conclusions drawn in Cooper and Ross [3] are misleading, and our extension of their results helps clarify these issues. Cooper and Ross focus on the role of "excess liquidity," that is, liquid assets that a bank intends to hold over the long term, despite the fact that these assets yield a lower rate of return than illiquid investment. There are in principle two reasons why a bank might choose to hold excess liquidity. First, having a highly liquid portfolio would minimize liquidation costs if a run were to occur. Second, if the bank decides to instead offer a run-proof contract, holding a highly liquid portfolio might provide the best possible run-proof consumption profile. Assuming preferences are of the constant-relative-risk-aversion variety, we show that a bank will *never* hold excess liquidity for the first reason. That is, if a bank offers

[2] This restriction is designed to mimic a sequential service constraint; suspension of convertibility is not allowed. In the problem considered by Diamond and Dybvig [5], the restriction to deposit contracts was non-binding since the first-best allocation was being implemented. Here, however, the constraint is binding. Nevertheless, Peck and Shell [9] show that the same qualitative results can obtain when a broad class of possible contracts is considered.

a deposit contract that is not run proof, it will hold only enough liquid assets to exactly meet withdrawal demand in the event that a run does not occur. The sole reason, then, that a bank would hold excess liquidity is as a way of making the contract immune to runs. We provide an exact characterization of the situations when the best run-proof contract does indeed involve holding excess liquidity.

These results enable us to characterize the relationship between the probability of a run and the level of illiquid investment undertaken by a bank. At first glance, it seems like the nature of this relationship should be straightforward: the more likely is a run, the more likely it is that the bank will have to liquidate all of its investment and therefore the less the bank should invest. However, we show that there is another, more subtle, effect at work. Since we know that a bank will use all liquid assets to meet short-term withdrawal demand, it follows that if the bank were to choose a more liquid portfolio it would also offer a higher return to agents who withdraw their deposits early. In other words, an increase in the liquidity of the bank's portfolio will necessarily be accompanied by an increase in short-term liabilities. In fact, we show that if the bank were to choose a more liquid portfolio, it would be able to serve *fewer* depositors in the event of a run. Decreasing the level of investment, then, leaves the bank with more resources if a run occurs, but it also implies that the resources will be shared by fewer depositors (and more depositors will receive nothing). We show that when liquidation costs are above some threshold, the first effect dominates in expected utility terms and investment is decreasing in the probability of a bank run. However, when liquidation costs are more moderate, the second effect dominates and an increase in the probability of a run leads to an *increase* in investment. This latter result is fairly counter-intuitive and has interesting implications. In this canonical model of bank runs, a more crisis-prone economy could have a *higher* level of investment and hence could be expected to experience faster growth, at least in periods where a run does not occur.

The remainder of the paper is organized as follows. In the next section, we briefly review the environment of Cooper and Ross [3] and then set up the maximization problem of a competitive bank. In Section 3 we describe the optimal deposit contract, examining the conditions under which it is run proof and the conditions under which excess liquidity will be held. In Section 4 we study how the level of investment responds to a change in the probability of a bank run, and in Section 5 we offer some concluding remarks.

2 The Bank's Problem

The environment we study is that of Cooper and Ross [3]; we follow their notation as much as possible. Consider an economy with a $[0, 1]$ continuum of ex-ante identical agents, each of whom lives for three periods. Each agent has an endowment (normalized to one) in period 0, and no endowment in periods 1 and 2. In period 0, an agent decides whether or not to deposit her endowment in the bank. With probability $\pi \in (0, 1)$, at the beginning of period 1 she will discover that she is impatient and only gets utility from consuming in period 1. With probability $(1 - \pi)$ she will discover that she is patient and derives utility only from consuming in period 2. Let $u(c)$ be the utility over consumption (in the appropriate period) and assume that u is strictly increasing, strictly concave, and satisfies $u(0) = 0$. There are two technologies for saving, which we call storage and investment. One unit of consumption placed in storage yields one unit of consumption in the following period. One unit of consumption placed in investment in period 0 yields R units of consumption in period 2, but only $(1 - \tau)$ units if the investment is liquidated in period 1, where $\tau \geq 0$ represents a liquidation cost.[3]

We study the problem of a bank that behaves competitively in the sense that it offers the contract that maximizes the expected utility of depositors. Following Cooper and Ross [3], we restrict the bank to offer deposit contracts of the following form. There is a fixed payment c_E that is promised to depositors who withdraw "early," that is, in period 1. The bank must pay this amount to each depositor arriving in period 1 unless it has completely run out of funds; no suspension of convertibility is allowed.[4] Whatever resources remain in the bank in period 2 are divided evenly among the remaining depositors. Let c_L denote the payment promised on these "late" withdrawals. In addition to choosing these two payments, the bank must also decide how to divide its portfolio between investment and storage. The bank will clearly place sufficient resources in storage to be able to pay c_E to all impatient depositors, since otherwise it would pay liquidation costs in period

[3] Diamond and Dybvig [5] assumed $\tau = 0$, in which case storage is a dominated technology and the bank's portfolio choice is trivial. Cooper and Ross [3] introduced the liquidation cost in order to address issues related to portfolio choice.

[4] As Diamond and Dybvig [5] show, when the fraction of impatient consumers is known with certainty a suspension-of-convertibility scheme can costlessly eliminate bank runs. However, issues of credibility and commitment might make such a suspension difficult to implement. In addition, when the fraction of impatient consumers is random the optimal contract from a broad class that includes total and partial suspensions of convertibility may permit runs to occur (see Peck and Shell [9]). For the issues discussed in this paper, the important thing is that bank runs can occur with positive probability in equilibrium, and the restrictions placed on the set of feasible contracts by Cooper and Ross [3] are a useful way of generating this feature.

1 even if a run did not occur. Let i denote the fraction of the bank's deposits placed into investment and $1 - i$ the fraction placed into storage; then we know that $c_E \leq 1 - i$ will hold. Define i_2 to be the difference between $1 - i$ and c_E, so that if a run does not occur the fraction $1 - i - i_2$ of deposits will be paid out in period 1. The variable i_2 represents *excess liquidity:* resources that the bank plans to keep in storage for two periods (if a run does not occur), even though investment offers a higher two-period return.

At the beginning of period 1, each depositor learns whether she is patient or impatient. She then decides whether to attempt to withdraw from the bank in period 1 or in period 2. If she is impatient, she will clearly choose to withdraw in period 1. If she is patient, however, her optimal action may depend on the choices of other patient depositors. We consider only symmetric outcomes. Suppose that a patient depositor believes that all other patient depositors will try to withdraw in period 1. The total amount of resources available to the bank in this case would be the funds held in storage plus the amount of goods that can be obtained by liquidating all investment. If the deposit contract is such that these resources are not enough to pay c_E to all depositors, the bank will run out of funds in period 1 (and have nothing left in period 2). The optimal strategy of a patient agent is then to "run" and attempt to withdraw in period 1. In other words, if $c_E > 1 - i\tau$ holds, then there exists a bank-run equilibrium of the game played by patient depositors. If, on the other hand, the bank has enough resources to meet all of its short-term obligations, a patient depositor knows that the bank will have sufficient funds left in period 2 to pay her the promised amount no matter how many early withdrawals take place. Waiting to withdraw is then a dominant strategy. That is, if $c_E \leq 1 - i\tau$ holds the contract is *run proof*; the only equilibrium of the game played by patient depositors has all of them withdrawing in period 2.

If the deposit contract admits a bank run equilibrium, the bank and the depositors need to have expectations about which equilibrium will be played. Cooper and Ross [3] introduce a sunspots-based equilibrium selection rule: If the deposit contract is such that both equilibria exist, the run equilibrium will occur with a fixed probability q. If the deposit contract is run-proof, of course, the no-run equilibrium will obtain with certainty. The problem of the bank is therefore

$$\max_{\{c_E,c_L,i,i_2\}} \quad (1-qI_{\widehat{\pi}})\left[\pi u\left(c_E\right)+\left(1-\pi\right)u\left(c_L\right)\right]+qI_{\widehat{\pi}}\left[\widehat{\pi}u\left(c_E\right)\right] \tag{P}$$

s.t.

$$\begin{aligned}
\pi c_E &= 1-i-i_2 \\
(1-\pi)\,c_L &= iR+i_2 \\
i,i_2 &\geq 0 \\
\widehat{\pi} &= \frac{1-i\tau}{c_E} \\
\text{and } I_{\widehat{\pi}} &= \begin{cases} 1 & \text{if} \quad \widehat{\pi}<1 \\ 0 & \text{if} \quad \widehat{\pi}\geq 1 \end{cases} .
\end{aligned}$$

The variable $\widehat{\pi}$ gives the fraction of the bank's depositors who would be served in the event of a run; the run equilibrium exists if and only if this fraction is less than unity. The indicator function $I_{\widehat{\pi}}$ reflects the equilibrium selection rule, whereby the run equilibrium occurs with probability q if $\widehat{\pi}$ is less than unity, but otherwise occurs with probability zero.[5]

We begin by looking at the solution to this problem when q is zero, that is, when a bank run is ruled out by assumption. This solution generates the first-best allocation, which satisfies

$$u'\left(c_E\right)=Ru'\left(c_L\right).$$

In order for the issue of bank runs to be relevant, we need for the deposit contract generating this first-best allocation to satisfy $c_E>1-i\tau$, so that $\widehat{\pi}<1$ holds and the bank-run equilibrium exists. This condition, in turn, implies that the first-best allocation must satisfy

$$c_E>\frac{1-\tau}{1-\pi\tau} \quad \text{and} \quad c_L<\frac{R}{1-\pi\tau}.$$

In other words, the first-best allocation must be to the right of point A in Figure 1 (at point x). Hence a run equilibrium will exist when the bank tries to implement the first-best allocation if and only if

$$u'\left(\frac{1-\tau}{1-\pi\tau}\right)>Ru'\left(\frac{R}{1-\pi\tau}\right) \tag{1}$$

holds. We assume this condition throughout the analysis.

[5] See Ennis and Keister [7] for a more detailed discussion of equilibrium selection in this type of model. Also note that the presence of the indicator function implies that the problem (P) is fundamentally different from problem (6) in Cooper and Ross [3]. (P) is the bank's complete optimization problem, while their problem (6) is only intended to be an auxiliary problem useful for characterizing the solution to the complete problem.

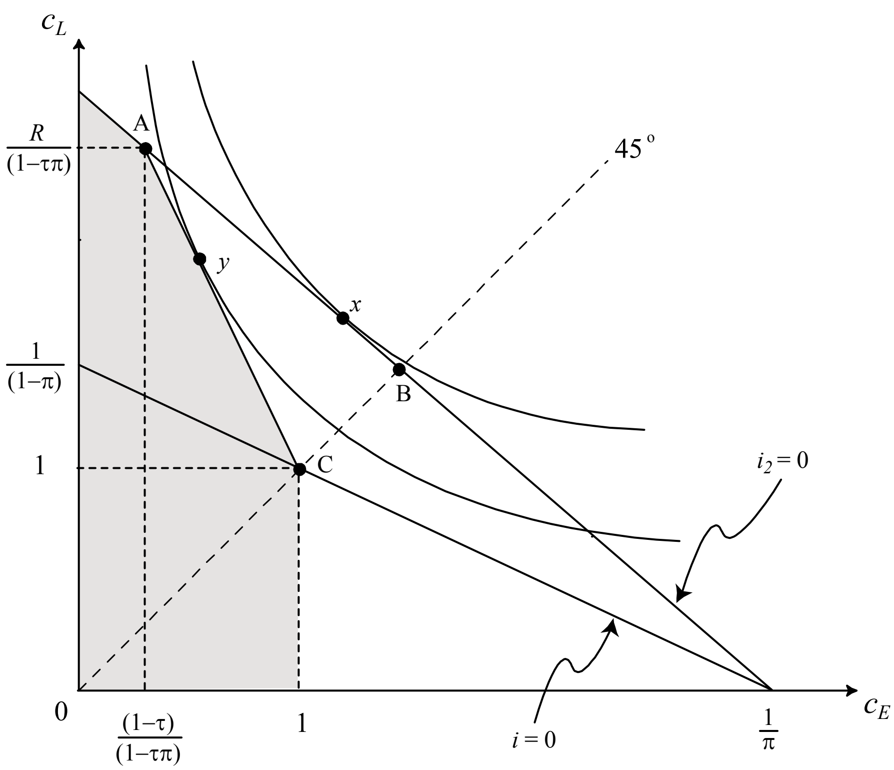

Figure 1: The set of feasible deposit contracts

Of course, the bank can always choose a contract such that $c_E \leq 1 - i\tau$ holds, in which case a run will not occur. If the bank were to choose such a contract, it would obviously choose the best contract in that category, which is called the *best run-proof contract*. In other words, if the solution to (P) is run proof, then it is also the solution to the following restricted problem

$$\max_{\{c_E, c_L, i, i_2\}} \pi u(c_E) + (1-\pi) u(c_L) \qquad (RP)$$

s.t.

$$\begin{aligned} \pi c_E &= 1 - i - i_2 \\ (1-\pi) c_L &= iR + i_2 \\ i, i_2 &\geq 0 \\ c_E &\leq 1 - i\tau. \end{aligned}$$

The only difference between problems (P) and (RP) is the constraint set; (RP) allows only run-proof contracts. On this set, the objective functions in the two problems are identical. As the statement of problem (RP) makes clear, this objective function has standard properties: it is continuous, strictly increasing, and strictly concave in the variables (c_E, c_L). In other words, the problem (RP) is a fairly standard maximization problem whose constraint set consists of the shaded region in Figure 1. Under condition (1), the slope of the indifference curve through point A is steeper than the slope of segment AB, and hence the solution to (RP) must lie on the line segment AC. This solution may either be at the "kink" point A or at a point of tangency like y.

The objective function in the complete problem (P), on the other hand, is not standard, both because the indicator function leads to a discontinuity and because the variable $\widehat{\pi}$ is a decreasing function of c_E. It is important to bear in mind that when q is positive, the indifference curves drawn in Figure 1 do not apply in the non-run-proof region – the objective function may not be increasing or even concave in this region. We show below how these non-standard properties generate some surprising results. First, however, we address an issue that is ignored in the statement of the problem given above: We need for the solution to satisfy the incentive-compatibility constraint $c_E \leq c_L$. If this inequality did not hold, withdrawing in period 1 would be a dominant strategy for patient agents and hence a bank run would occur with certainty. Our first result shows that the solution to the problem (P) as stated above necessarily satisfies incentive compatibility, and hence (P) is indeed the correct problem for the bank to be solving.

Proposition 1 *The solution to* (P) *satisfies* $c_E < c_L$.

The proof of this proposition is contained in the appendix. One way of interpreting the result is as saying that depositors do not receive perfect insurance against the liquidity shock. Intuitively, perfect insurance is not optimal because insurance is costly in this setting; payments to impatient depositors must be made from relatively low-return assets. The proposition implies that the solution to (P) must lie in the triangle ABC in Figure 1, not including the boundary segment BC. The solution will lie on the segment AC if the optimal contract is run proof, and will lie either on the segment AB or in the interior of the triangle if it is not.

Notice that the statement of the proposition is actually stronger than the requirement for incentive compatibility to hold, since it involves a strict inequality rather than a weak one. It is easy to see from the figure that this stronger result directly implies that the solution to problem (P) cannot involve the bank holding only liquid assets.

Corollary 1 *The solution to* (P) *satisfies* $i > 0$.

3 The Optimal Deposit Contract

We now turn our attention to characterizing the solution to (P), asking under what conditions the optimal deposit contract is run proof and under what conditions it involves holding excess liquidity.

3.1 When is the optimal contract run proof?

Suppose that the likelihood of a run, q, and the cost of liquidation, τ, are both large enough that

$$q\tau > (1-q)(R-1) \tag{2}$$

holds. Our next result shows that this condition is sufficient to guarantee that the optimal deposit contract is run proof.

Proposition 2 *If (2) holds, the solution to (P) is a run-proof contract.*

The formal proof is contained in the appendix; the reasoning is as follows. Suppose the solution to (P) is not run proof, and consider the effect of decreasing i slightly while increasing i_2 to keep c_E unchanged. That is, suppose we shift some of the resources set aside for period 2 withdrawals from investment to storage. The benefit of such a change comes in the event of a run, since total liquidation costs will then be τ units lower for each unit less of investment. For each additional unit of resources the bank has during a run, it can give the fixed amount c_E to $(1/c_E)$ additional depositors. Hence the value of having an additional unit of resources is measured by the average utility of consumption $u(c_E)/c_E$. Since a run occurs with probability q, the gain in expected utility is therefore

$$q\tau\frac{u(c_E)}{c_E}.$$

However, this proposed change causes a decrease in utility in the event that a run does not occur. If there is no run, the bank will have earned a return of unity (instead of R) on the assets switched from investment to storage, and therefore total resources in period 2 will be $(R-1)$ units lower for each unit less of investment. The marginal value of a unit of resources in period 2 when a run did not occur in period 1 is measured by $u'(c_L)$, the marginal utility of a depositor withdrawing in period 2. Since the "no run" event occurs with probability $(1-q)$, the loss in expected utility is equal to

$$(1-q)(R-1)\,u'(c_L)\,.$$

From Proposition 1 we know that $c_E < c_L$ holds at the solution point. Strict concavity of the function u then implies that $u(c_E)/c_E$ is greater than $u'(c_L)$. Therefore, when (2) holds the gain in expected utility is necessarily greater than the loss, and hence the original contract under consideration cannot be optimal. Since this argument applies to any contract that is not run proof, the solution to (P) must be a run-proof contract.

Condition (2) was introduced by Cooper and Ross [3] for a different purpose, in an attempt to give a sufficient condition for excess liquidity to be held. It is important to bear in mind that when a run-proof contract is chosen, the objective function in (P) reduces to that in (RP). In other words, there is a discontinuity in the objective function along the line segment AC in Figure 1. Hence the above argument does *not* imply that the solution to (P) will have $i_2 > 0$ when condition (2) holds. It only shows that the solution to (P) is run-proof; problem (RP) can then be solved to find the properties of this optimal contract. We show below that condition (2) is neither necessary nor sufficient for $i_2 > 0$ to hold.[6]

Proposition 5 in Cooper and Ross [3] establishes a cut-off probability $q^* \in (0,1)$ such that for all $q > q^*$ the solution to problem (P) is run proof and for all $q < q^*$ it is not. Our Proposition 2 gives an upper bound for q^*, which we state as a corollary.

Corollary 2 *If q^* is the threshold value of q above which the solution to (P) is a run-proof contract, then*

$$q^* \leq \frac{R-1}{R-(1-\tau)}$$

must hold.

3.2 When is excess liquidity held?

For the remainder of the paper, we assume that the utility function of depositors is of the constant-

[6] In other words, the statement of Proposition 3 in Cooper and Ross [3] is misleading. That proposition states that if we consider the problem (P) without the indicator function (their problem (6)), then under condition (2) the solution to this problem has $i_2 > 0$. Our Proposition 2 shows that under condition (2) the solution to their problem (6) is *never* the solution to the bank's complete optimization problem (P), and hence is irrelevant for characterizing the optimal deposit contract.

relative-risk-aversion form

$$u(c) = \frac{c^{\alpha}}{\alpha} \quad \text{with} \quad 0 < \alpha < 1. \tag{3}$$

As discussed above, the benefit of an increase in the resources available to the bank in the event of a run is measured by the average utility of consumption $u(c_E)/c_E$, while this benefit is measured by the marginal utility of consumption when a run does not occur. Finding the optimal deposit contract when a run is possible, then, involves comparing average and marginal utilities of consumption. With the utility function in (3), marginal utility is always a fraction α of average utility, and hence sharp comparisons are possible. Most of our results below involve strict inequalities and hence are robust to small perturbations of the function u.

We now investigate under what conditions the bank will actually choose to hold excess liquidity. There are two distinct reasons why the bank might choose to do this. First, having more assets in storage lowers liquidation costs and therefore makes depositors better off in the event of a run. Second, holding a very liquid portfolio might be the best way to eliminate the possibility of a run by satisfying the run-proof condition. Our next result shows that the first reason never applies. That is, a bank will never hold excess liquidity for the purpose of mitigating the effects of a potential run.

Proposition 3 *If a solution to the bank's problem is not run-proof, then the bank holds no excess liquidity (i.e., $i_2 = 0$).*

The proof of this proposition consists of showing that the solution to (P) does not lie in the interior of the triangle ABC in Figure 1. The boundary segment AC can be ruled out by the hypothesis of the proposition, while the boundary BC can be ruled out by Proposition 1. Hence the solution must lie either in the interior of the triangle, or on the boundary AB. In the appendix we show that the objective function cannot have an interior local maximum (although it may have a local minimum). This establishes that if the solution is not run proof, it must lie on the boundary AB where $i_2 = 0$ holds.

Proposition 3 shows that the bank will never hold excess liquidity for the purpose of providing funds to depositors in the event of a run. The only reason it might set i_2 positive, then, is as a way of eliminating the run equilibrium. We now provide conditions under which the level of excess liquidity in the best run-proof contract is positive. Under condition (1), the best run-proof contract must lie on the segment AC in Figure 1. If the contract is at point A then the bank will not be holding excess liquidity and the consumption allocation will be given by

$$c_E = \frac{1-\tau}{1-\pi\tau} \quad \text{and} \quad c_L = \frac{R}{1-\pi\tau}.$$

If the contract is to the right of point A, however, the bank will be holding excess liquidity. The condition that determines whether or not the best run proof contract is at point A involves the slope of the indifference curve of the objective function in problem (RP) at that point. Specifically, the best run proof contract is to the right of point A if

$$u'\left(\frac{1-\tau}{1-\tau\pi}\right) > \left[1 + \frac{R-1}{\pi\tau}\right] u'\left(\frac{R}{1-\tau\pi}\right) \tag{4}$$

holds. Note that

$$1 + \frac{R-1}{\pi\tau} > R$$

holds, so that (4) is a stronger requirement than (1). When (1) holds but (4) does not, the best run-proof deposit contract is at point A and has i_2 equal to zero. We summarize this discussion in the following proposition, whose formal proof is contained in the appendix.

Proposition 4 *The best run-proof contract satisfies $i_2 > 0$ if and only if (4) holds.*

Proposition 2 provides a sufficient condition for the best deposit contract to be run proof. This condition can be made to hold by choosing a sufficiently high value of q. At the same time, Proposition 4 gives the condition under which the best run-proof contract involves holding excess liquidity. Notice that this condition does not depend on q, so that it is clearly possible for both conditions to hold at the same time. Combining these two results therefore gives us a set of sufficient conditions under which the solution to the bank's problem (P) involves holding excess liquidity.[7]

Corollary 3 *If (4) and (2) hold, the solution to (P) has $i_2 > 0$.*

It is important to bear in mind that the reason the bank holds excess liquidity is *not* to be able to satisfy depositors' demand during a bank run, but instead to avoid the possibility of a run altogether. Under the assumption on preferences in (3), the Cooper-Ross model predicts that a bank will never hold excess liquidity for the purpose of providing funds to depositors in the event of a run.

[7] Note that Proposition 4 and Corollary 3 do not require that the utility function be of the form (3).

4 The Level of Investment

We now turn to the question of how the possibility of a run affects the amount of resources that the bank places into investment. In the present model the level of investment affects the consumption allocations of both types of depositors. In a dynamic setting where banks play an important role in the capital formation process, it becomes even more important. For example, Ennis and Keister [7] embed a similar banking environment into an endogenous growth model.[8] In such a setting, bank runs have two important effects on the growth rate of the economy. The first is obvious: a bank run leads to the liquidation of investment and therefore a lower stock of capital (and hence output) in the next period. The second effect is less direct but no less important: the possibility of a run affects the level of investment chosen by banks, which in turn affects the growth rate in periods when a run does not occur. Hence understanding how the investment decisions of banks are affected by the variable q in the present model is an important part of understanding the macroeconomic effects of bank runs.

At first glance, the answer to this question seems like it should be straightforward: when a run is more likely, there is a higher probability that the bank will have to liquidate all of its investment and therefore it seems optimal for the bank to invest less. In this section, we show that this intuition is not necessarily correct. An increase in the probability of a run can either decrease or increase the level of investment chosen by the bank, depending on parameter values. In particular, define

$$\tau^* \equiv \frac{(1-\alpha)\left(R^{\frac{1}{1-\alpha}} + \frac{\pi}{1-\pi}\right)}{(1-\alpha)\,R^{\frac{1}{1-\alpha}} + \frac{\pi}{1-\pi}}.$$

Note that $0 < \tau^* < 1$ holds. We then have the following result.[9]

Proposition 5 *If $\tau > \tau^*$ holds, i is strictly decreasing in q for $q \in (0, q^*)$. If $\tau < \tau^*$ holds, i is strictly increasing in q for $q \in (0, q^*)$.*

To see the intuition behind this result, suppose that q is equal to zero. In this case, the solution to (P) is to offer the first-best allocation, which is represented by point x in Figure 1. Now suppose that q increases slightly. The objective function in a neighborhood of the first-best allocation is

[8] See also Gaytan and Ranciere [8], which uses a different banking model to examine the relationship between bank runs and economic development.

[9] This proposition is similar in spirit to Proposition 4 in the working paper Cooper and Ross [4]. However, our result characterizes the solution to the bank's complete problem (P), whereas their result applied to an artificial problem where i_2 was required to be zero. In addition, our assumption on preferences (3) allows us to provide a more complete characterization.

equivalent to

$$(1-q)\left[\pi u\left(c_E\right)+(1-\pi)\,u\left(c_L\right)\right]+q\widehat{\pi}u\left(c_E\right). \tag{5}$$

Increasing q generates a continuous change in this (local) objective and hence the solution to the problem will move continuously, as long as q stays below q^*. Proposition 3 tells us that while $q < q^*$ holds, the solution has $i_2 = 0$. In other words, as q increases the solution to the bank's problem must move along the line segment AB in Figure 1. The first term in (5) is maximized at the point x, and hence there is no first-order loss in the value of this term from deviating in either direction along the boundary. Therefore, whether the level of investment should be increased or decreased when q is raised above zero is determined entirely by the last term in (5), which is the expected utility of depositors in the event of a run.

If a run occurs, there is an obvious benefit to having less investment: total liquidation costs will be lower and the bank will have more resources to give to depositors. However, there is a cost to having less investment as well. Because we know $i_2 = 0$ holds, the first constraint in problem (P) sets πc_E equal to $1-i$. Decreasing investment therefore implies raising the amount promised to depositors who withdraw in the first period. In fact, it is straightforward to show that $\widehat{\pi}$ is strictly increasing in i, and hence decreasing i implies that *fewer* depositors will be served during a run. In other words, lowering the level of investment implies that during a run a depositor will be given a larger amount if she is served, but that she will be served with a lower probability. Which of these two effects dominates in terms of expected utility depends on the size of the liquidation cost τ. If τ is very large, then lowering investment leads to a large increase in the total resources available to the bank and the first effect dominates. In this case, the answer is what one would (naively) expect: The possibility of a run leads the bank to choose a more liquid portfolio. However, when τ is relatively small, lowering investment leads to a modest increase in the total resources available and the second effect dominates. In this case the answer is surprising: The possibility of a run leads the bank to choose a *less* liquid portfolio. Applied to a growth setting, then, the Cooper-Ross model predicts that when liquidation costs are small, a more crisis-prone economy would have higher levels of investment and growth, at least in periods where a crisis does not occur.

It is important to bear in mind that the bank *could* decrease liquidation costs without the negative side effect of decreasing the number of depositors served in the event of a run. By setting i_2 positive, the bank could decrease i and leave c_E unchanged, for example, guaranteeing that more depositors would be served in the event of a run. However, Proposition 3 tells us that this response would not be optimal for any $q < q^*$. The reason is that if a run does not occur, holding excess liquidity is very costly. That is, while moving slightly away from the point x along the segment AB in Figure 1 does not cause a first-order loss in expected utility in the no-run outcome, moving strictly inside the constraint set does. Hence the optimal response of the bank to a small probability of a run will always be to adjust only the level of investment, and the direction of this adjustment will depend on the size of τ. The intuitive arguments we have given here are valid for q close to zero, but the proof of the proposition in the appendix shows that they extend to all q less than q^*.

5 Concluding Remarks

We have studied what is in many ways a fairly simple model where a bank run is possible in equilibrium. We have maintained the simplicity of the environment for several reasons. First, the bank runs that we study here are *exactly* the kind that were first heuristically described by Diamond and Dybvig [5] in their seminal paper. The analysis of Cooper and Ross [3] can be seen as a formalization of this argument and a fleshing out of its implications. We have been able to deepen and clarify the results of Cooper and Ross [3] in important dimensions. Additionally, and perhaps most importantly, the simplicity of the environment allows us to obtain a better understanding of the various ways in which the possibility of a bank run affects the incentives of a bank offering demand deposit contracts. In fact, some of our results are far from obvious, and hence this simple environment seems to be a very useful step in understanding more general issues. There are many ways in which extending the model presented here could give interesting new results. Just to mention one, modeling bank runs explicitly as an equilibrium phenomenon changes the willingness of agents to participate in the banking system. In particular, when runs are possible, agents may not want to deposit all of their resources in the bank (as is assumed in the Cooper-Ross framework). Even though the main incentive effects that we study in this paper will not change much by permitting agents to keep some resources outside the banking sector, the macroeconomic implications of such disintermediation could be significant (see Ennis-Keister [7]). In this sense, extending the analysis presented here to richer environments seems to be a promising agenda. The analysis in this paper should prove helpful in such a program, especially to the extent that it makes starkly clear that the possibility of a run can change the incentives of a bank in relatively intricate ways.

Appendix A. Proofs

Proposition 1: *The solution to* (P) *satisfies* $c_E < c_L$.

Proof: We break the proof into two parts: the case where the solution to (P) is run proof and the case where it is not. If the solution to (P) has $c_E \leq 1 - i\tau$ (which implies $c_E \leq 1$) then this solution must also solve (RP). Since the objective function in (RP) is strictly increasing in c_E and c_L, we know that the solution must lie on the outer boundary of the constraint set, either at the point A or along the segment AC in Figure 1. Hence $c_E \leq c_L$ necessarily holds. To show that this inequality is strict, we only need to show that the point C cannot be the solution. The slope of the indifference curve evaluated at any point on the 45-degree line is equal to $-\pi/\left(1-\pi\right)$. The slope of the line segment AC is equal to

$$-\frac{R-1+\tau\pi}{\tau\left(1-\pi\right)} < -\frac{\pi}{1-\pi}.$$

That is, the indifference curve through the point C is flatter than the boundary segment AC, and hence C cannot be the solution.

More formally, let the multipliers on the first four constraints listed in problem (RP) be given by λ_E, λ_L, γ, and β, respectively, and let δ denote the multiplier on the run-proof condition $c_E \leq 1 - i\tau$. Then the necessary first-order conditions characterizing the best run-proof contract include:

$$c_E \quad : \quad \pi u'(c_E) - \lambda_E \pi - \delta = 0 \tag{6a}$$

$$c_L \quad : \quad u'(c_L) - \lambda_L = 0 \tag{6b}$$

$$i \quad : \quad -\lambda_E + \lambda_L R - \delta\tau + \gamma = 0 \tag{6c}$$

$$i_2 \quad : \quad -\lambda_E + \lambda_L + \beta = 0. \tag{6d}$$

where all of the multipliers must be non-negative. We also have the complementary slackness conditions

$$\gamma i = 0, \qquad \beta i_2 = 0, \qquad \text{and} \qquad \delta(1 - i\tau - c_E) = 0. \tag{7}$$

The first-order conditions can be combined to yield

$$u'\left(c_E\right) - u'\left(c_L\right) = \beta + \frac{\delta}{\pi}.$$

Therefore, we need to show that any solution to problem (RP) satisfies $\beta + \delta/\pi > 0$, since the strict concavity of u then implies that $c_E < c_L$ then must hold. Assume to the contrary that both β and δ are zero (note that by definition both must be non-negative). Then, from the first order conditions we would have $\lambda_E = \lambda_L = u'(c_L)$. Using condition (6c) would then yield

$$\gamma = (1 - R)u'(c_L) < 0$$

which is a contradiction because γ must be non-negative. Therefore it cannot be the case that both β and δ are zero in the solution to problem (RP), and hence we have $\beta + \delta/\pi > 0$. This establishes that the conclusion of the proposition must hold if the solution is a run-proof contract.

If the solution to problem (P) is not run proof (i.e., has $c_E > 1 - i\tau$), then in a neighborhood of the solution point, the objective function is equivalent to

$$(1 - q)\left[\pi u\left(c_E\right) + (1 - \pi)\, u\left(c_L\right)\right] + q\widehat{\pi} u\left(c_E\right). \tag{8}$$

The first-order conditions characterizing a solution in this region are given by

$$\begin{aligned}
c_E &: \quad (1-q)\pi u'(c_E) - q\widehat{\pi}\left[\frac{u(c_E)}{c_E} - u'(c_E)\right] - \lambda_E \pi = 0 && \text{(9a)}\\
c_L &: \quad (1-q)u'(c_L) - \lambda_L = 0 && \text{(9b)}\\
i &: \quad -q\tau\frac{u(c_E)}{c_E} - \lambda_E + \lambda_L R + \gamma = 0 && \text{(9c)}\\
i_2 &: \quad -\lambda_E + \lambda_L + \beta = 0, && \text{(9d)}
\end{aligned}$$

where all of the multipliers must be non-negative. Substituting (9a) and (9b) into (9d) yields

$$\beta = (1 - q)\left[u'(c_E) - u'(c_L)\right] - q\frac{\widehat{\pi}}{\pi}\left[\frac{u(c_E)}{c_E} - u'(c_E)\right].$$

Because β is non-negative, $u(0) = 0$, and u is strictly concave, this equation implies that $u'(c_E) - u'(c_L) > 0$ holds and hence that we must have $c_E < c_L$. This establishes that the conclusion of the proposition also holds if the solution to (P) is not a run-proof contract, and therefore completes the proof. ■

Proposition 2: *If (2) holds, the solution to (P) is a run-proof contract.*

Proof: Suppose the solution is not run proof. Then we have $c_E > 1 - i\tau$ and $\widehat{\pi} < 1$, which implies that there is an open ball around the solution point where the objective function is equivalent to

(8). As in the proof of Proposition 1, the first-order conditions characterizing a maximum of this new objective subject to the feasibility constraints are given by (9a) - (9d). We also have the complementary slackness conditions (7). By Corollary 1 we only need to look for a solution with $i > 0$, so that $\gamma = 0$ holds. In this case, substituting (9b) and (9d) into (9c) yields

$$q\tau \frac{u(c_E)}{c_E} = (1-q)(R-1)u'(c_L) - \beta.$$

However, using the concavity of u and the result of Proposition 1 (that $c_E < c_L$ holds), this condition contradicts (2). ■

Proposition 3: *If a solution to the bank's problem (P) is not run-proof, then the bank holds no excess liquidity (i.e., $i_2 = 0$).*

Proof: Suppose the solution to (P) is not run proof. Then the first-order conditions for problem (P) are again (9a) - (9d). By corollary 1 we know that $i > 0$ and hence that $\gamma = 0$ holds. If $i_2 > 0$ holds in the solution, then $\beta = 0$ must hold, which in turn implies that $\lambda_E = \lambda_L = (1-q)c_L^{\alpha-1}$ holds. Combining this result with expression (9c) we obtain

$$(R-1)(1-q)c_L^{\alpha-1} = q\tau \frac{c_E^{\alpha-1}}{\alpha}.$$

This expression tell us that when the solution to problem (P) is not run-proof, if the bank is holding excess liquidity (that is, if $i_2 > 0$), then the optimal values of c_E and c_L lie on a ray from the origin in the (c_E, c_L) plane. In other words, we have $c_L = \phi c_E$ where ϕ is a constant given by

$$\phi \equiv \left(\frac{\alpha(1-q)(R-1)}{q\tau}\right)^{\frac{1}{1-\alpha}}.$$

We now show that these conditions cannot characterize a maximum of problem (P). Using the first two constraints in problem (P) we can rewrite the level of investment as

$$i = \frac{\pi c_E + (1-\pi)c_L - 1}{R-1}$$

and hence we have

$$\widehat{\pi} = \frac{1-i\tau}{c_E} = \frac{(R-1) + \tau\left[1 - \pi c_E - (1-\pi)c_L\right]}{(R-1)\,c_E}.$$

Using this expression for $\widehat{\pi}$ and the relationship $c_L = \phi c_E$, we can rewrite the objective as a function of the single variable c_E. Calling this function $f(c_E)$ we have

$$f(c_E) \equiv A\frac{c_E^{\alpha}}{\alpha} + B\frac{c_E^{\alpha-1}}{\alpha},$$

where A and B are constants given by

$$A \equiv (1-q)\left[\pi + (1-\pi)\phi^{\alpha}\right] - \frac{q\tau\left[\pi + (1-\pi)\phi\right]}{R-1}$$

and

$$B = q\frac{R-1+\tau}{R-1} > 0.$$

Since the solution being considered is interior, the following two necessary conditions must hold at the optimum: $f'(c_E) = 0$ and $f''(c_E) \leq 0$. The first-order condition implies

$$c_E = (1-\alpha)B/\alpha A. \tag{10}$$

The second-order condition is given by

$$f''(c_E) = -(1-\alpha)Ac_E^{\alpha-2} + \frac{(1-\alpha)(2-\alpha)}{\alpha}Bc_E^{\alpha-3} \leq 0,$$

which can only hold if $c_E \geq (2-\alpha)B/\alpha A$. This inequality contradicts (10) and therefore the proposed solution cannot satisfy the necessary conditions for a maximum. Hence we must have $i_2 = 0$ at the solution. In other words, evaluated along the ray where $c_L = \phi c_E$, the objective function cannot have a local maximum (although it may have a local minimum). If the solution does not lie on the run-proof boundary, then it must lie on the other boundary, where $i_2 = 0$ holds. ■

Proposition 4: *The best-run proof contract satisfies $i_2 > 0$ if and only if (4) holds.*

Proof: Condition (4) is stronger than condition (1). Hence $c_E = 1 - i\tau$ must hold at the solution to (RP) and we can rewrite the constraints as

$$c_E = \frac{\tau + R - 1}{\pi\tau + R - 1} - \frac{(1-\pi)\tau}{\pi\tau + R - 1}c_L,$$

and

$$\frac{1-\tau}{1-\tau\pi} \leq c_E \leq 1.$$

We have $i_2 > 0$ whenever

$$c_E > \frac{1-\tau}{1-\tau\pi} \tag{11}$$

holds (see Figure 1). Suppose we look at the first-order condition of the problem ignoring the inequality constraints; this condition is given by

$$\left[1 + \frac{R-1}{\pi\tau}\right] u'(c_L) = u'(c_E).$$

If (4) holds, the solution to the first-order condition satisfies (11). In this case, the solution to the first-order condition is necessarily the solution to (RP), and hence that solution has $i_2 > 0$. Conversely, if (4) does not hold, then the solution to (RP) is the boundary point

$$c_E = \frac{1-\tau}{1-\tau\pi}$$

where $i_2 = 0$ holds. ■

Proposition 5: *If $\tau > \tau^*$ holds, i is strictly decreasing in q for $q \in (0, q^*)$. If $\tau < \tau^*$ holds, i is strictly increasing in q for $q \in (0, q^*)$.*

Proof: Fix some $\widehat{q} < q^*$. Then for values of q in a neighborhood of $\widehat{q}$, we know that the solution to (P) is not run proof and, by Proposition 3, therefore has $i_2^* = 0$. In other words, the problem (P) is locally equivalent to the problem of choosing the level of investment i to maximize (5), where the consumption allocations are given by

$$\pi c_E = 1 - i \quad \text{and} \quad (1-\pi)\, c_L = Ri.$$

These constraints can be substituted into the objective function to yield a maximization problem in a single variable. Because the solution is not run proof, we know that it satisfies $c_E > (1 - i\tau)$, which is equivalent to

$$i < \frac{1-\pi}{1-\pi\tau}.$$

On the interval $\left(0, \frac{1-\pi}{1-\pi\tau}\right)$, the objective function is smooth. It may not be strictly concave on the entire interval, but it can be shown to have a unique local maximum (which must be the global maximum because the solution is interior) and to be strictly concave at this maximum point. Therefore, the first-order condition

$$(1-q)\left[R\left(\frac{Ri}{1-\pi}\right)^{\alpha-1}-\left(\frac{1-i}{\pi}\right)^{\alpha-1}\right]+\frac{q}{\alpha}\left[(1-\alpha)\frac{1-i\tau}{1-i}-\tau\right]\left(\frac{1-i}{\pi}\right)^{\alpha-1}=0 \quad (12)$$

implicitly defines a differentiable function $i(q)$ in a neighborhood of $\widehat{q}$. Letting $G(i,q)$ denote the left-hand side of (12), the derivative of this function is given by

$$\frac{di}{dq}=-\frac{G_2(i,q)}{G_1(i,q)},$$

where subscripts on the function G represent partial derivatives. Strict concavity of the objective at the solution point implies that the denominator of this expression is negative, and hence that di/dq has the same sign as G_2.

The derivative G_2 can be written as

$$G_2(i,q)\equiv\frac{1}{\alpha}\left[(1-\alpha)\frac{1-i\tau}{1-i}-\tau\right](c_E)^{\alpha-1}-\left[R(c_L)^{\alpha-1}-(c_E)^{\alpha-1}\right] \quad (13)$$

Notice that this expression is independent of q. Suppose that $\widehat{q}$ is very close to zero. Then the solution to (P) will be very close to the first-best allocation, which implies that the second term in (13) is very close to zero and the sign of G_2 is equal to the sign of the first term. The value of i will also be close to the first-best value

$$i^*=\frac{R^{\frac{\alpha}{1-\alpha}}}{R^{\frac{\alpha}{1-\alpha}}+\frac{\pi}{1-\pi}}.$$

If $\tau<\tau^*$ holds, it is straightforward to show that the first term in (13) is negative. Therefore, for values of $\widehat{q}$ close to zero we have that G_2 is negative and i will be strictly decreasing in q at $\widehat{q}$. Furthermore, it is straightforward to show that the cross-partial derivative G_{21} is always strictly positive. Thus for larger values of $\widehat{q}$, the optimal value of i is lower than i^* and hence G_2 is again negative, implying that i is still strictly decreasing in q. This analysis holds as long as the first-order condition (12) characterizes the solution to (P), which is for all $q<q^*$, and hence establishes the first part of the proposition. The case of $\tau<\tau^*$ is completely symmetric. ■

REFERENCES

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

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

[3] Cooper, Russell and Thomas W. Ross (1998), Bank Runs: Liquidity Costs and Investment Distortions, *Journal of Monetary Economics* 41, pp. 27-38.

[4] Cooper, Russell and Thomas W. Ross (1991), Bank Runs: Liquidity and Incentives, NBER Working Paper #3921 (November).

[5] Diamond, Douglas W. and Philip H. Dybvig (1983), Bank Runs, Deposit Insurance, and Liquidity, *Journal of Political Economy* 91, pp. 401-419.

[6] Ennis, Huberto M. and Todd Keister (forthcoming), Government Policy and the Probability of Coordination Failures, *European Economic Review*.

[7] Ennis, Huberto M. and Todd Keister (2003), Economic Growth, Liquidity, and Bank Runs, *Journal of Economic Theory* 109, pp. 220-245.

[8] Gaytan, Alejandro and Romain Ranciere (2003), Banks, Liquidity Crises and Economic Growth, mimeo.

[9] Peck, James and Karl Shell (2003) Equilibrium bank runs. *Journal of Political Economy* 111, pp. 103-123.