January 15, 2006

Federal Reserve Bank of Cleveland

Optimal Deposit Contracts: Do-It-Yourself Bank-Run Prevention for Banks

by Ed Nosal

Government provision of deposit insurance is often rationalized on the grounds that it stabilizes the banking system by removing the incentives for depositors to engage in bank runs. Even free market luminaries such as Milton Friedman have credited federal deposit insurance for the relative stability of the United States' banking system after World War II. The perception that deposit insurance improves the stability of the banking industry has been reinforced by a number of theoretical papers on bank runs, including the seminal work in this area by Douglas Diamond and Philip Dybvig.

However, as illustrated by the 1980s savings and loan debacle and more recent regional banking problems in the United States, deposit insurance is not a panacea. By effectively eliminating bank runs, deposit insurance has the unintended effect of reducing market discipline and increasing incentives for banks to take on risk. While deposit insurance may be effective in eliminating potentially destabilizing bank runs, it is far from clear that deposit insurance is the best way to promote the stability of banks and the banking system.

Despite the received wisdom that banks are prone to instability, a careful review of the economic literature on the topic suggests otherwise. Recent theoretical research, in particular, shows that banks themselves can prevent bank runs by modifying their demand deposit contracts. It turns out that in the absence of deposit insurance, banks will try to protect themselves from being the victim of bank runs by choosing an appropriate financial structure. In the end, it is difficult to tell a coherent story in which banks are inherently unstable.

Bank Runs

A bank run occurs when depositors believe that their bank will be unable to honor all of its deposit liabilities; depositors rush to the bank and attempt to withdraw as much of their money as they can, in order to avoid losing all or most of it. Given that a bank holds only a small fraction of its assets in the form of cash, a bank run can be disruptive because the bank will have to liquidate its noncash assets in order to meet the demands of its depositors.

In spite of these disruption costs, in some circumstances a bank run may not be such a bad thing. If a bank is insolvent, then, from society's point of view, it is optimal that the bank be dissolved, and the sooner the better. A bank run can send a very visible signal that something may be wrong; if the bank is indeed insolvent, a run may be an expedient way to bring about the optimal dissolution. But if the depositors' beliefs are incorrect and the bank is actually solvent prior to the bank run, then a bank run would be a very bad thing if it resulted in a solvent bank becoming insolvent.

How susceptible are banks in general to runs of the sort that turn solvent banks into insolvent ones? This is actually a hard question to answer. Although we have seen very few bank runs since the Great Depression—which might lead one to conclude that the banking sector is stable—federal deposit insurance, which came into existence during the Great Depression, may explain this. If depositors understand that their deposits are guaranteed to pay off up to some maximum amount (which today is $100,000), they will have little or no incentive to ever run on a bank.

The need for federal deposit insurance is often based on the claim that it prevents bank runs and makes the banking system more stable. But research shows that banks could prevent bank runs by constructing their deposit contracts appropriately, and, in the absence of deposit insurance, they would do so in their own self interest. Federal deposit insurance may be useful as insurance per se—protecting depositors against unforeseen accidents—but it should not be considered necessary for banking system stability.

Prior to federal deposit insurance, bank runs were a bit more commonplace. But it is difficult to determine the extent to which solvent banks became insolvent following a run. Gary Gorton provides some indirect evidence that suggests that for bank runs during the U.S. National Banking Era (1863–1914), depositors were responding to a worsening of their bank's financial net worth, meaning that depositors were probably running on insolvent—and not solvent—banks. But this evidence is inconclusive.

A Model of Bank Runs

Perhaps the most famous model of bank runs is the one articulated by Douglas Diamond and Philip Dybvig in 1983. Their model, and variations of it, are the closest thing that the profession has to a consensus model of bank runs. Here, I outline the basic intuition behind it.

An investor has $1 today. He can just hold onto the $1 if he wants or he can invest it in a long-term project and

ISSN 0428-1276

receive, say, $1.20 at a future date. The long-term investment project can always be liquidated at an earlier date. The model assumes for simplicity that if an investor liquidates his long-term investment project early, he will get back the dollar that he originally invested. The basic insights from the analysis do not depend on this assumption, and this assumption can be interpreted as saying that financial markets do a good job of converting relatively illiquid investments into liquid assets. At the time when the investor must decide what to do with his dollar—invest it or not—he does not know whether he will be impatient, meaning that he prefers to spend and consume at an early date, or patient, meaning that he prefers to spend and consume at a late date. Suppose that the investor believes there is a 50-50 chance that he will be impatient. What is his optimal investment and spending strategy? First of all, the individual investor will always invest in the long-term project. If he turns out to be impatient, and this happens with a 50 percent probability, he will liquidate the long-term investment and buy $1.00 worth of goods at the early date; if he turns out to be patient, he will liquidate his project at the late date and buy $1.20 worth of goods at that time.

The first thing that can be shown with the Diamond and Dybvig model is that banks are useful in that they can achieve better outcomes for investors than investors can get by acting in isolation of one another. For example, investors don't like risk; specifically, they prefer outcomes that produce less variation in their income or consumption. Banks can provide a smoother spending profile for an investor—in this case, something smoother than a choice between $1.00 at the early date or $1.20 at the late date.

Let a bank enter the scene and assume that there are many investors with a dollar to invest. A bank will take deposits, promising its depositors withdrawal possibilities—either for the early or late date. Mirroring what happens at the individual level, suppose in this economy that half of the investors turn out to be impatient and the other half patient; there is no economywide risk in the sense that there is no uncertainty regarding the proportion of investors that will be impatient or patient.

It is possible for the bank to improve upon the investor's optimal strategy. Suppose that the bank receives $1.00 from each investor and invests the total proceeds in the long-term investment project. In return, the bank promises to give $1.05 to anyone who wants to make a withdrawal at the early date or $1.14 to anyone who wants to make a withdrawal at the late date. Note that half of the investors will be impatient and will want to have money at the early date. In order to pay these people off, the bank must liquidate 52.5 percent of the long-term investment, i.e., 0.525/0.5 = 1.05. Hence, 48.5 percent of the original long-term investment remains intact, and the proceeds of the investment can be shared by the remaining 50 percent of the investors at the late date, i.e., (0.475 x 1.20)/0.5 = 1.14. The bank's deposit contract of promising $1.05 early or $1.14 late—which I will compactly represent by (1.05, 1.14)—is "smoother" than an investor's payoff if he does not use a bank. If investors really don't like risk, then they will prefer the deposit contract (1.05, 1.14) to the payoff (1, 1.20) that they can get for themselves.

It appears that investors are better off in a world where there is a bank. As well, it appears as if the bank will be solvent in the sense that it is able to pay off all of its obligations: The deposit contract is designed to be able to pay off half of the depositors at the early date and the remaining half at the late date.

Bank Runs and Suspensions

Unfortunately, impatient investors choosing to withdraw $1.05 at the early date and patient players waiting to withdraw $1.14 at the late date is but only one possible outcome in this banking scenario. To see this, suppose that after investors give their $1.00 to the bank in exchange for the deposit contract (1.05, 1.14), each investor comes to believe (for some reason) that all other investors are going to make an early withdrawal. And this is independent of whether investors are patient or impatient. When an investor believes that all other investors are going to make an early withdrawal from the bank, then his best response is also to make an early withdrawal even if he is patient. The reasoning is straightforward: If all other investors attempt to make an early withdrawal, the bank will be unable to pay off everyone because the total liquidation value of the long-term investment can only guarantee a payoff of $1.00 for each investor. But $1.05 is promised. If investors queue up at the bank and the bank pays off investors on a first-come-first-served basis, then those investors who are lined up near the end of the line will receive nothing. Obviously, in this scenario the bank will be unable to make any kind of payoff at the later date. Hence, a patient investor who understands all this will want to make an early withdrawal because a 95 percent chance of getting $1.05 at the early date is better than a 100 percent chance of receiving nothing at the later date.

This scenario makes the bank look as if it is unstable: At the time investments are undertaken and deposit contracts are negotiated, the bank is solvent—that is, the bank is expecting to pay off half of its depositors early and the other half later. But when the unanticipated bank run occurs, the bank becomes insolvent.

It turns out that in this simple example, a bank run can be prevented if the bank offers a better deposit contract to investors. By offering the deposit contract (1.05, 1.14) the bank effectively invites a run to happen if the depositors (for some reason) come to believe that all other depositors are going to make an early withdrawal.

Suppose that in addition to the payoff (1.05, 1.14), the bank's deposit contract also stipulates that the bank will "suspend payments" at the early date after it pays out 52.5 percent of the liquidated value of its original investment. With this added stipulation, can a bank run occur? Once again, suppose that investors come to believe that all other investors will attempt to make an early withdrawal. What is the optimal response for a patient investor? Unlike the first scenario, the patient investor will not make a run for the bank. The reason is clear. Because the bank suspends payments after it liquidates a certain fraction of the original investment, it essentially guarantees that any (patient) investor who waits until the late date to make a withdrawal will receive $1.14 for sure. Knowing this, the patient investor has no incentive to make a run on the bank at the early date because the patient investor prefers receiving $1.14 at the late date as opposed to $1.05 at the early date. So, at least in this simple example, if banks design their deposits contracts in an appropriate way, they can prevent bank runs from occurring.

In practice, banks have suspended payments to their depositors. In the National Banking Era, banks would at times suspend conversion of their banknotes into specie (gold) if they were expecting a destabilizing run to occur. After a while, these suspensions lifted, and banks would once again convert their banknotes into specie. In Scotland, bankers had option clauses embedded in their banknotes, which gave bankers the option of not converting their banknotes into specie over a specific period of time. This option was viewed as a protective measure against bank runs.

Partial Suspensions

Perhaps the above example is too simple because it was constructed so that there is no economywide risk. I can introduce economywide risk by assuming that the probability that one investor is impatient, which is equal to 0.5, is independent of the probability that any other investor will be patient. So, for example, if there are 10 investors in the world, then there is a positive probability that there can be 1, 2, and all the way up to 10 investors who are impatient. In this example, the probability that there is exactly one impatient investor is equal to 0.5^{10}. Unlike the previous example, when there was no economywide risk, the bank does not know at the time that it collects money from investors and issues deposit contracts exactly how many investors will be patient and how many will be impatient. Therefore, it may not be a good idea for the bank to offer the deposit contract (1.05, 1.14) with a suspension clause, as above, since the suspension might kick in when there still are impatient investors who need funds for spending. But without a suspension clause, a deposit contract, like the one described in the above example, is always subject to a bank run. Does this imply that banks are unstable when there is economywide risk?

Just as in the case when there was no economywide risk, when a bank efficiently uses all information at its disposal, it can design a deposit contract that prevents bank runs from occurring. Although the deposit contract when there is economywide risk is a bit complicated, the intuition behind it is straightforward. In previous examples I assumed that there was a 50-50 chance that an investor would be impatient. If I assume instead that there is only a 10 percent chance that an investor will be impatient and there is no economywide risk, then it is feasible for a bank to offer the deposit contract (1.15, 1.18), along with a suspension rule that only 10 percent of the original investment will be liquidated at the early date. The point here is that if only 10 percent of the investors turn out to be impatient, as opposed to 50 percent in previous example, the payoff to an impatient investor will increase. Recall that when there was a 50-50 chance of an investor being impatient, the best deposit contract was (1.05, 1.14): Hence, as the proportion of impatient investors increases, their payoffs will decrease.

Using the insight that the payoffs to investors will decrease as the proportion of impatient investors increases, when there is economywide risk—meaning that we do not know how many impatient investors there will be—the optimal deposit contract will be characterized by the first impatient investor who wants to withdraw early receiving more than the second impatient investor, who receives more than the third impatient investor, and so on. That is, as the proportion of impatient investors increases, their individual payoffs will decrease and, hence, the optimal deposit contract is characterized by a declining payout to impatient investors.

This deposit contract will also have two rather important features. First, the deposit contract must be feasible: Since it may be the case that all investors turn out to be impatient—although this event may be unlikely—the deposit contract must ensure that if everyone shows up at the bank at the early date to make a withdrawal, everyone will receive the appropriate payoff when all of the investment is liquidated. And second, the deposit contract must be incentive compatible: The levels of the deposit payoffs for early withdrawals will be such that patient investors will prefer to wait to withdraw at the later date rather than line up and make an early withdrawal. For the kind of economywide risk that we have described, this deposit contract will not invite a bank run. The bank will always be able to offer a payoff to investors who withdraw late because the bank contract is feasible; the patient investor does not have to worry about the bank not having any resources at the late date. Even if a patient investor arrives early—perhaps anticipating a bank run—he will always prefer not to make a withdrawal at that time, but rather wait until later, because the bank contract is incentive compatible. This kind of deposit contract has been called a "partial suspension" contract because the payoff to a withdrawer at the early date is suspended in favor of a lower payoff for any subsequent early-date withdrawers.

Deposit Insurance

It appears that the banking sector will be stable if banks offer deposit contracts that permit them to either totally or partially suspend payments to depositors. The reader, however, should not take the contracts described above too literally. What the reader should take away is that if banks understand that depositors may, for some reason, attempt to make a run on a solvent bank, then they (the banks) can offer contracts with features that prevent bank runs occurring in the first place.

The fact that banks are able to protect themselves against bank runs does not imply that deposit insurance is without value. The pure insurance aspect of deposit insurance clearly has social value. Just as people find it valuable to purchase home insurance to protect themselves against an accident, deposit insurance provides the same benefit to depositors. That is, a depositor's bank may suffer an adverse shock that would diminish the value of existing deposit accounts; deposit insurance protects deposit holders against such an accident.

Recommended Reading

Douglas Diamond, and Philip Dybvig. (1983). "Bank Runs, Deposit Insurance, and Liquidity," *Journal of Political Economy*, 91, 401–19.

Federal Reserve Bank of Cleveland
Research Department
P.O. Box 6387
Cleveland, OH 44101

Return Service Requested:
Please send corrected mailing label to the above address.

PRSRT STD
U.S. Postage Paid
Cleveland, OH
Permit No. 385

Ed Nosal is a senior economic advisor at the Federal Reserve Bank of Cleveland.

Economic Commentary *is published by the Research Department of the Federal Reserve Bank of Cleveland. To receive copies or to be placed on the mailing list, e-mail your request to* 4d.subscriptions@clev.frb.org *or fax it to* 216-579-3050. Economic Commentary *is also available at the Cleveland Fed's site on the World Wide Web:* www.clevelandfed.org/research.

We invite comments, questions, and suggestions. E-mail us at editor@clev.frb.org.