PART III: DISTORTED CAPITAL FLOWS

Chapter 7: Bank Runs and Capital Flow Reversals

Introduction

Many emerging economies which had liberalized their capital markets generated massive capital inflows. Not much later, these inflows sharply reversed during the financial crises that started in East and South-East Asia in 1997 and then spread to Russia in 1998 and Brazil in 1999. In all cases, the liberalization of the capital markets was not accompanied by an appropriate reform of the banking sector which is preiminent in the financial system of these economies (although other financial intermediaries are growing in importance). Banks remained vulnerable to bank runs. Indeed, bank failures were intertwined with capital flow reversals, generating balance-of-payments crises.[1]

In this chapter we lay out the seminal model of bank runs which was developed by Diamond and Dybvig (1983). They were the first to formalize the maturity mismatch between the bank's assets and liabilities which is at the root of the vulnerability of a bank. This vulnerability may lend itself to a bank run. We then show how international capital flows interact with bank failures.

The Diamond-Dybvig Model of Bank Runs

Suppose that there is a single all-purpose commodity which serves as both a capital good and a consumption good. There are three periods: The present, the short-run, and the long-run. In period one (the present) there are two types of investment opportunities. One is a short-run opportunity which matures in the second-period yielding a return of r_S. This short-run investment technology is available also in period two. Thus, a unit of the good invested in period one in a short-run opportunity can be reinvested in the same type of opportunity in the second period, accumulating to $(1 + r_S)^2$ units of the good in period three. The second type of investment technology which is available in period one is of a long-run nature, that is, it has an incubation process lasting two periods. Thus, a unit of the good invested with this technology matures in period three, yielding an accumulated quantity of $(1 + r_L)^2$ units of the good. If this investment is terminated after one period only, that is, in period two, it accumulates to $1 + r_B$ units of the good. Naturally, it is assumed that such abrupt termination of a project which was a priori designed as a long-run project is costly in the sense that it yields upon such termination a return (r_B) which is significantly lower than the return (r_S) of a project which was originally designed for the short-run. Otherwise, no one will choose to invest in a short-term technology. On the other hand, a long-term investment which is held until its planned maturity date, that is, period three, yields a higher return than a short-term investment which is reinvested for an additional period. Thus, $r_L > r_S$. Summarizing:

$$r_L > r_S > r_B. \tag{7.1}$$

Consumers are ex-ante identical. Each consumer possesses an initial endowment in period one of one unit of the good; she possesses no endowment in any other subsequent period. There is a continuum of individuals whose number is normalized to one. Therefore, the aggregate initial endowment in the first period is one unit of the good. To simplify the analysis, we assume that the individual derives no utility from consuming in period one, so that she invents all of her first-period endowment. We further assume that the individual consumes either in period two or three (but not in both periods), depending on whether a "need" arises either in period two or in period three.

In period one, no individual knows in which subsequent period (two or three) such a consumption (liquidity) need will arise. For a proportion λ of the consumers the need will arise in period two; these consumers will be referred to as "early consumers". For the other $1-\lambda$ consumers, the need will arise in period three; these consumers will be referred to as "late consumers." In period one, all consumers are identical; each one of them faces a probability λ of turning to be an early consumer and a probability $1-\lambda$ of turning to become a late consumer. There is no aggregate (macro) uncertainty in this economy: Exactly λ consumers will be early consumers and exactly $1-\lambda$ consumers will be late consumers; also, the returns on the two types of assets (the short-term asset and the long-term asset) are safe. However, no one in period one knows which type of consumer she is going to be, so that there is a risk at the individual level but not at the aggregate level.

The expected utility of an individual in period one is

$$U(c_2, c_3) = \lambda u(c_2) + (1-\lambda)u(c_3), \tag{7.2}$$

where c_i is consumption in period $i = 2, 3$ and u exhibits risk-aversion (that is, u is concave).[2]

A crucial assumption in the model is that in the second period only the individual knows whether she became a late consumer or an early consumer. That is, the late consumption or early consumption state of the world is private information and no contract contingent on these states of the world can be enforced.

No Financial Intermediation

Suppose first that there are no banks or any other financial intermediaries. In Figure 7.1 the line AB describes the consumption possibility frontier of a representative individual. If she invests all of her unit endowment in a short-term investment technology, then she will be able to consume $1 + r_S$ units in period two, if she turns out to be an early consumer; she will consume $(1 + r_S)^2$ units in period three, if she turns out to be a late consumer. This consumption bundle is described by point B. At the other extreme, if she invests all of her unit endowment in a long-term investment technology, then if she turns out to be an early consumer in period two, she will be forced to prematurely terminate her investment in period two in which case she will get $1 + r_B$ units of consumption; if she turns out to be a late consumer, then she will hold her investment until its planned maturity in which case she will be able to consume $(1+r_L)^2$ units in period three. This consumption bundle is described by point A. Dividing her unit endowment between the two types of investment opportunities she can attain any consumption bundle along the line AB. Given this consumption possibility frontier, she chooses the consumption bundle D, where her indifference curve $U(c_2, c_3) = U_A$ is tangent to this frontier.

Financial Intermediation

One of the roles of a financial intermediary is the so-called "maturity transformation," that is, the matching between savers' needs of different maturities and investments of different maturities. Consider a perfectly competitive bank where all consumers deposit their unit endowment. Since it is certain that exactly λ consumers will turn out to be early consumers and $1 - \lambda$ consumers will turn out to be late consumers, then if the bank invests λ units in short-term investment technologies and $1 - \lambda$ units in long-term investment technologies, it will have $\lambda(1 + r_S)$ units of consumption in period two and $(1 - \lambda)(1 + r_L)^3$ units of consumption in period three, without having to terminate any long-term investment. Being competitive, the bank will offer all consumers a deposit contract in period one which promises a return of r_S, if the deposit is withdrawn in period two, and a return of r_L per period, if the deposit is withdrawn in period three. Note that the bank will not be able to distinguish between the two types of consumers in period two (because the type of a consumer is private information), and therefore it allows all (both late and early) consumers to withdraw their deposits in period two. However, it does not pay a late consumer to withdraw her deposit in period two, if she expects that only early consumers will withdraw their deposits in period two. Therefore, the deposit contract is incentive-compatible. Indeed, when only early consumers withdraw their deposits in period two, total withdrawals in this period are $\lambda(1 + r_S)$ which exactly matches the amount of short-term assets which matured and are at the bank's disposal in period two. This deposit contract offers the consumer the consumption bundle N in Figure 7.1.

The bank can also follow different investment strategies other than investing λ units in short-term technologies and $1 - \lambda$ units in long-term technologies. Accordingly, it can offer different deposit contracts other than a return of r_S for withdrawals in period two and a return of r_L per period for withdrawals in period three. For instance, the bank can invest β units in short-term technologies and $1 - \beta$ units in long-term technologies. Accordingly, the bank can choose any $0 \leq \beta \leq 1$ and offer any pair (d_S, d_L) of short and long-term per-period returns, respectively, which satisfy the following three constraints:

$$\lambda(1 + d_S) = \beta(1 + r_S) + \alpha(1 - \beta)(1 + r_B), \tag{7.3}$$

$$(1 - \lambda)(1 + d_L)^2 = (1 - \alpha)(1 - \beta)(1 + r_L)^2, \tag{7.4}$$

and

$$\alpha = 0, \tag{7.5}$$

where α is the fraction of long-term projects that the bank has to terminate in period two in order to meet its short-term liabilities. Since there are λ early consumers in period two who will all withdraw their deposits in that period, the bank will have to pay them an amount equalling $\lambda(1 + d_S)$. The bank has a total of β units invested in short-term technologies. These investments will hand to the bank an amount of $\beta(1 + r_S)$ units of consumption in period two. If total withdrawals at period two (namely, $\lambda(1 + d_S)$) exceed total available liquid funds (namely, $\beta(1 + r_S)$), the bank has to call in a fraction α of long-term investments with a penalty of earning a return of only r_B. The fraction α has to satisfy equation (7.3). In period three, the remaining fraction $(1-\alpha)$ of the long-term investments of the bank matures and provides an amount of $(1 - \alpha)(1 - \beta)(1 + r_L)^2$ units of consumption. This amount must suffice to pay total withdrawals which amount to $(1 - \lambda)(1 + d_L)^2$. This explains equation

(7.4). Constraints (7.3)-(7.5) can be consolidated into two constraints as follows:

$$(1-\lambda)(1+d_{\mathsf{L}})^2(1+r_{\mathsf{B}})+\lambda(1+d_{\mathsf{S}})(1+r_{\mathsf{L}})^2=[\beta(1+r_{\mathsf{S}})+(1-\beta)(1+r_{\mathsf{B}}]\,(1+r_{\mathsf{L}})^2, \tag{7.6}$$

and

$$\lambda(1+d_{\mathsf{S}})=\beta(1+r_{\mathsf{S}}). \tag{7.7}$$

Thus, for each investment strategy, that is, for each pair $(\beta, 1-\beta)$, there exists a continuum of deposit contracts, that is, pairs $(d_{\mathsf{S}}, d_{\mathsf{L}})$ that the bank can offer; these pairs have, of course to satisfy (7.6)-(7.7). Put differently, for each pair $(\beta, 1-\beta)$, there is a corresponding consumption possibility frontier in $(c_2, c_3)-$ space. The upper envelope of all of these frontiers, obtained by letting β vary from zero to one is termed the grand consumption possibility frontier, and is depicted by the curve HNK in Figure 7.1. Note that point N is indeed feasible: First, N is characterized by an investment strategy $(\beta, 1-\beta)=(\lambda, 1-\lambda)$ and a deposit contract $(d_{\mathsf{S}}, d_{\mathsf{L}})=(r_{\mathsf{S}}, r_{\mathsf{L}})$. It is straightforward to see that these values of $(\beta, 1-\beta)$ and $(d_{\mathsf{S}}, d_{\mathsf{L}})$ satisfy constraints (7.6)-(7.7). Next, note that N lies on the grand consumption possibility frontier: If we maximize c_3 subject to $c_2=1+r_{\mathsf{S}}$ and constraint (7.6)-(7.7), then the only solution is $(\beta, 1-\beta)=(\lambda, 1-\lambda)$ and $(d_{\mathsf{S}}, d_{\mathsf{L}})=(r_{\mathsf{S}}, r_{\mathsf{L}})$.[3]

Point M is a competitive equilibrium in which each consumer withdraws her deposit in period two, if and only if, she turns out to be an early consumer. This is a "good" financial equilibrium which clearly dominates the autarkic (with no financial intermediaries) equilibrium, D.

However, there is another, "bad" financial equilibrium of a bank run. As put by the Economist (1999, p.107): "Most of a Bank's liabilities have a shorter maturity than its assets. There is, therefore, a mismatch between the two. This leads to problems, if depositors become very worried about the quality of a bank's lending book that they demand their savings back. Although some overdrafts or credit lines can easily be called in, longer-term loans are much less liquid. This can cause a bank to fail."

As long as the bank does not invest all of its assets in short-term technologies, then everyone knows that the bank will not be able to pay a return d_S in period two, if all consumers run on the bank. This is because the competitive bank offered a short-term return d_S on the premise that only λ consumers (that is, the early consumers) will withdraw their deposits in period two; see equation 7.3. Thus, even a late consumer will find it optimal to join a run on the bank, because in such a case the bank will be forced to terminate abruptly all of its long-term projects ("liquidate all of its assets") in period two and divide whatever resources are available among panicked depositors on a first-come, first-served basis. If a late consumer does not join the run, she will be left with nothing in period three. Thus, a bank run is another financial equilibrium.

This "bad" bank-run equilibrium is located somewhere on the curve QB in Figure 7.1. To see this, consider the resources available to the bank to divide among its panicked depositors in period two. If it were to invest all deposits in short-term technologies, then it will be able to pay $1 + r_S$ to all of its depositors in period two. Thus, an early consumer will then be able to consume $1 + r_S$ units in period two. If the consumer turns out to be a late consumer, then she will reinvest her withdrawn amount of $1 + r_S$ in period two in a short-term technology, to obtain $(1 + r_S)^2$ units of consumption in period three. This consumption bundle is described by point B in Figure 7.1. At the other extreme, if the bank were to invest all of its deposits in long-term technologies, it will have to liquidate all its assets in period two in the event of a bank run and get $1 + r_B$ units of consumption. Hence, a consumer that turned out to be an early consumer in period two can consume $1+r_B$ units of consumption in this period. A late consumer will receive $1 + r_B$ units in period two and invest it in a short-term technology, and will be able to consume $(1 + r_B)(1 + r_S)$ units of consumption in period three. This consumption-bundle is described by point Q in Figure 7.1. Other intermediate investment strategies, including the strategy that underlies the "good" equilibrium M, are located somewhere along the line QB.

Summing up

Financial intermediation can support a "good" equilibrium (point M) which unequivocally dominates the consumption possibility frontier in the absence of financial intermediation (the curve AB). But, financial intermediation could also turn into a bank run which is "bad." Such an equilibrium is located somewhere along the curve QB, and is clearly inferior to the equilibrium which prevails in the absence of financial intermediation (point D).

The "bad" equilibrium can be eliminated by "the lender of the last resort." For instance, a government deposit insurance will nullify the possibility of a bank run. Note that such an insurance scheme is indeed feasible as there is no aggregate risk in the economy. In fact, the insurance will never be exercised (that is, a bank run will never occur), if everyone indeed believes that the government has sufficient resources to bail out the bank on which the depositors may run.

Nevertheless, in a more realistic setup there is some aggregate risk. For instance, the return of long-term investment (namely, r_L) is typically risky. In this case, deposit insurance may be either infeasible (for instance, in a closed economy), or costly. It is also plugged by moral hazards, as the bank's investment strategy may be biased toward high-return, high-risk portfolio; in a "good" state of the world, it will reap the high return; in a "bad" state of the world, the government will bail it out.

The emergence of two different outcomes ("good" and "bad") from the same action (deposit contracts) by the banks and the consumers is a serious drawback or inconsistency of the Diamond-Dybvig model. The bank is actually offering depositors contracts that may end up in a bank run, without taking this possibility into account at all. Similarly, depositors accept such contracts that are not feasible in a state of a bank run, without paying any attention to such a possibility of which they are fully aware at the time they make the deposits. Loosely speaking, the equilibrium in this model is not fully rational if the bank-run equilibrium has some positive probability. If all depositors believe that there will be no bank run and act accordingly, then there will not be one; if all believe that there will be a bank run and act accordingly, then there will indeed be a bank run. And, as described by Morris and Shin (2000): "The shift in beliefs which underpins the switch from one equilibrium to another is left unexplained."

The next section extends the Diamond-Dybvig model in a way that results in a fully rational equilibrium.

A Rational Expectations Bank Run Equilibrium

We now assume that long-term investments are typically risky. Suppose then that r_L =

$r_L(\theta)$, where θ is the state of the world; and assume, for concreteness, that r_L is strictly increasing in θ and that no one has any private information about the state of the world before deciding whether to withdraw her deposit in period two. Instead, θ is revealed to all at the same time (that is, the state of the world becomes common knowledge), after they decide whether or not to withdraw their deposits in period two.

One can verify (see, for instance, Goldstein and Pauzner (1999)) that this setup still lends itself to two outcomes with self-fulfilling beliefs: a "good" one with no bank runs and a "bad" one with bank runs, without any explanation for either the formation of each set of beliefs, or for the switch from one set to another. As put by Morris and Shin (2000): "The apparent indeterminancy of beliefs can be seen as the consequence of two modelling assumptions.... First, the economic fundamentals are assumed to be common knowledge; Second, economic agents are assumed to be certain about others' behavior in equilibrium... Both assumptions allow agents' actions and beliefs to be perfectly coordinated in a way that invites a multiplicity of equilibria." They go on to explain what ingredient can be essential in rationalizing the model outcome. In their approach "agents have a small amount of idiosyncratic uncertainty about economic fundamentals. Even if this idiosyncratic uncertainty is small, agents will be uncertain about each other's behavior in equilibrium. This uncertainty allows us as modellers to pin down which set of self-fulfilling beliefs will prevail in equilibrium." They use this approach to tackle various issues such as bank runs (Morris and Shin (2000)) and currency crises (Morris and Shin (1998)).

Here we follow Goldstein and Pauzner in supposing that in period two, before agent i decides on whether or not to withdraw her deposit in period two, she receives an imperfect signal θ_i about the true value of θ, which is private to her. Given the state of the world (that is, the fundamental, θ), the signal is uniformly distributed over $[\theta - \varepsilon, \theta + \varepsilon]$. Knowing the distribution of the signal, if agent i receives a signal θ_i, then she knows that no one received a signal below $\theta_i - 2\varepsilon$ or above $\theta_i + 2\varepsilon$.

Goldstein and Pauzner elegantly show how, for a given fundamental θ (which uniquely determines the long-term rate of return $r_L(\theta)$), the indeterminancy between the two outcomes regarding bank runs is eliminated. This is shown in Figure 7.2, where the fraction of consumers withdrawing their deposits in period two is described by the solid curve. Note that early consumers always withdraw their deposits in period two. Whether or not it pays a late consumer to withdraw her deposit in period two depends on d_S and on θ (because θ determines $d_L(\theta)$). Given d_S, they show that there is a unique threshold level of the fundamental, denoted by $\theta^*(d_S)$, such that if the true value of the fundamental is below $\theta^*(d_S) - \varepsilon$, there will be a bank run; if $\theta > \theta^*(d_S) + \varepsilon$, there will be no bank run; and if $\theta^*(d_S) - \varepsilon < \theta < \theta^*(d_S) + \varepsilon$, only a fraction of the late consumers withdraw their deposits in period two.[4] As the signals are positively correlated to the fundamental, this fraction (of consumers withdrawing their deposits in period two) is decreasing in the fundamental.

The rationale for this result is as follows. For a very large value of the fundamental θ, it is assumed that there will arise a lender of last resort that will guarantee the return $d_L(\theta)$. (Alternatively, for a very large value of the fundamental, the payoff of the long-term asset can arrive earlier in period 2 instead of period 3.) In this case every consumer will receive a "high" level of the signal. Then, no late consumer will withdraw early, no matter what she believes other late consumers may do. In this case, the fraction of early withdrawing consumers is therefore equal to λ, the fraction of early consumers. On the other hand, for a very low level of the fundamental, each agent will receive a "low" level of the signal and will early withdraw, even if she believes that no other late consumer will do so.

Consider now "intermediate" values of the fundamental θ. An agent learns two things from the signal: First, she learns about the fundamental θ (and the long-term return); Second, she learns about the signals of other agents (which form their beliefs). The higher her signal θ_i, the more she expects to earn if she is a late consumer and decides not to withdraw early. Also, the higher θ_i, the higher is the posterior expected value of θ and, consequently, the posterior expected value of other agents' signals. As a result, it becomes less likely that other agents will run on the bank. She then has less incentive for an early withdrawal. This consideration suggests how the indeterminancy between the two outcomes (of total bank run or no bank run) under common knowledge (more accurately, common ignorance) is eliminated. This is also why the fraction of early withdrawing consumers is declining in θ.

As expected and actually shown by Goldstein and Pauzner, the threshold $\theta^*(d_S)$ is equal to the level of the signal which equates the expected payoff of a late consumer from withdrawing her deposit in period two early, or waiting until its planned maturity in period three. (This expected payoff is conditioned on the level of the signal which also reveals some information about the signals received by all consumers, which, in turn, affects the number of early withdrawing consumers.) Evidently the threshold θ^* is increasing in the deposit-rate d_S.

Having eliminated the indeterminancy of the outcome of each deposit contract, the bank and the depositors can now assess the probability of a bank run, based on the underlined, exogenously given probability of the fundamental (namely, θ). Therefore, the bank will take this probability into account in designing its deposit contracts and depositors will do so in choosing their most preferred deposit contracts. In contrast to the Diamond-Dybvig model, the bank now realizes that if it offers a higher short-term deposit rate (d_S), then it increases the probability of a bank run, that is, the solid curve in Figure 7.2 shifts to the right. The outcome of this extended model is a rational expectations equilibrium of financial intermediation.

To sum up, if there is common knowledge about the fundamentals, there is a possibility of multiple equilibria. This means that at each realization of the fundamental, agents may coordinate on any one of these multiple equilibria. In contrast, with noisy signals, the withdrawal action on which agents coordinate for each realization of the fundamental must be consistent with withdrawal actions they may take at adjacent values of the fundamentals. This places restrictions on their withdrawal behavior which in this case brings about unique equilibrium.

Bank Failures and Balance-of-Payments Crises

Formally speaking, the aforementioned analysis applies to economies that could be either isolated from or integrated into the world capital markets. However, a globalized economy is more vulnerable to bank runs. Furthermore, such episodes may be accompanied by (or interacting with) balance-of-payments crises, especially in emerging economies or in economies that recently underwent a capital market liberalization.[5] In this section we highlight three possible mechanisms through which this interrelation takes place.

As already mentioned, a bank run can be averted by a credible lender of last resort, or a credible government-sponsored deposit insurance. In either case, the government is essentially making a credible promise to bail out the creditors of the bank (that is, short and long-term depositors, other banks extending credit to the said bank, etc.). However, while domestic politics may be sympathetic totwards bailing out domestic creditors, it may be reluctant to bail out foreign creditors. Hypothetically, a government may distinguish between domestic and foreign creditors and provide insurance to the former only. In such a case, a bank with a high fraction of foreign creditors will be more vulnerable to a bank run. Such a bank run will generate a capital flight which may even reverse an otherwise inflow of capital, thereby creating a balance of payments crisis.

Such a distinction between foreign and domestic creditors may revoke international retaliation. Practically, it is also infeasible as foreign creditors may sell their claims to domestic residents. Nonetheless, domestic opposition to bailing out foreign creditors does not lessen. This makes it harder for the government to altogether implement a bail-out program. The credibility of a government-sponsored insurance scheme is eroded as the fraction of foreign credits out of all credits rises, thereby increasing the probability of a bank run. With a high fraction of foreign credits, bank failures may coincide with capital flow reversals and balance-of-payment crises.[6]

A second mechanism intertwining bank failures with capita flow reversals in emerging and developing countries may work as follows. Foreign creditors in such economies may have a better access to the world capital markets than domestic creditors. (For example, domestic residents may still face legal restrictions on investing abroad.) In this case, when a late foreign depositor withdraws her deposit early, then she may face better reinvestment opportunities than her domestic counterpart. Therefore, foreign creditors may be more "trigger happy" in running on the bank than domestic creditors. Again, as the fraction of foreign credits out of all credits rises, the probability of a bank run, accompanied by a capital flight, increases.

A third mechanism which connects bank crises and currency crashes is rooted in currency mismatch of financial intermediaries' balance sheet (see Ratelet and Sachs (1999) and Goldstein (2000)), in addition to maturity mismatch. With short-term liabilities of the intermediaries denominated in foreign currency (the foreign deposits) and long-term assets (the financial intermediary loans) denominated in domestic currency, a withdrawal of foreign deposits may lead to a bank crisis. Foreign currency outflows put pressures on the currency peg. To the extent that the probability that the monetary authorities would be willing to defend the peg (when it becomes more costly to do so) rises, the likelihood of currency depreciation increases. But, when the probability that the currency will crash rises, the likelihood of early withdrawal of deposits from banks also increases, thereby raising the likelihood of bank runs. Therefore, the currency mismatch in the financial intermediaries' balance sheet generates a circular interaction between currency crashes and bank crises.