Charles M. Kahn

Econ 590

Demand-Deposit Contracts and the Probability of Bank Runs

By Itay Goldstein and Ady Pauzner

Journal of Finance, Vol LX, No 3, June 2005, pp 1293-1327

The Diamond-Dybvig model predicts two equilibria arise in a demandable debt contract, one with and one without bank runs. This is a useful theoretical observation but one which is unsatisfactory for empirical or policy purposes: the important thing is to predict when, or how likely bank runs are to occur. The problem is that each agent's actions in this situation depend on his expectations as to the actions of the other agents. When there are a large number of agents in this position, a useful toolbox for turning a multiple equilibrium model into a single equilibrium model is the set of techniques which come from the "global games" literature (Carlsson and van Damme (1993) or Morris and Shin (1998)). This paper applies these techniques to the Diamond-Dybvig model. As they note, the techniques have to be modified, because the basic idea from those papers assumes a monotonicity of responses of individuals to the likelihood of other individuals' actions; here for some regions of the space of expectations that is not the case.

The Model

The basic Diamond-Dybvig framework:

Three periods (0,1,2) one good, and a continuum of agents, each with an endowment of one unit in period 0. Consumption occurs in period 1 or 2; with probability λ the agent is impatient and has utility $u(c_1)$ and with the remaining probability $u(c_1 + c_2)$ where c_t denotes consumption at date t. Usual assumptions on u: twice continuously differentiable increasing and coefficient of relative risk aversion greater than 1, with $u(0)$ normalized to 0.

Technology generates 1 unit of output per input in period 0 if liquidated in period 1; if liquidated in period 2 it returns R units of output with probability $p(\theta)$ and 0 with remaining probability, where θ, parameterizing the state of the economy, is drawn from a uniform distribution on [0,1] and is only revealed to agents in period 2. The function p is increasing in θ and we assume the expectation over θ of $p(\theta)u(R)$ is greater than $u(1)$ so that for patient agents the project dominates storage.

In other words, the sole alteration of DD is that probability of second period payoff is random—but utility still linear in that probability. (For example, first best risk sharing sets

$$u'(c_1^{FB}) = R\, u'(R(1 - \lambda\, c_1^{FB})/(1 - \lambda))\, E[\, p(\theta)\,]$$

which, except for the final expectational component is the same as in DD.) As in DD, given preferences, risk sharing transfers wealth from patient to impatient agents.

Bank contracts:

These take the form of demandable debt in period 1 and equity shares in period 2. In other words, agents declare their type. The contract promises a fixed amount r_1 to agents declaring themselves impatient, provided the proportion n making such a declaration is less than $1/r_1$ so that the bank does not use up its resources in making the provision. If the number declaring exceeds this amount, than a fraction chosen randomly receive their promised amount and the remainder receive nothing; in other words, the payoff to an impatient agent in this case is r_1 with probability $(n\, r_1)^{-1}$ and 0 otherwise. Those declaring themselves impatient get equal shares of the remnant of the firm in period 2. If the number declaring impatience is less than $1/r_1$ then what is left is $(1 - n\, r_1)\, R$ with probability $p(\theta)$; otherwise it is zero.

We look for the contract maximizing customer expected utility and causing the bank to break even. This contract provides first best payoffs in one of its equilibria. But in the bank run equilibrium, all agents demand early withdrawal, and so payoff in period 1 is r_1 with probability r_1^{-1} and zero for sure in period 2. This equilibrium is inferior to autarky. The authors argue that this result is unsatisfactory in two ways: First, rather than having two separate equilibria, it is more satisfying to look for a single equilibrium in which the bank run occurs with some probability. Second we would expect that the likelihood of a bank run outcome should be linked in some way to the terms proposed in the contract.

Agents with Private Signals

So make the following modification: assume that at period 1 each agent obtains a noisy signal $\theta_i = \theta + \varepsilon_i$, where the error term ε_i is small and uniformly distributed over $[-\varepsilon, +\varepsilon]$. Now the decision to withdraw or not will depend on the agent's type and signal. The signal not only provides information about the period 2 chances for the bank, it also provides information about *others* beliefs about the period 2 chances. Impatient agents will always choose early withdrawal; patient agents' decisions will depend on the information and their expectations of others' actions.

The equilibrium in this model will be found by using some techniques based on the iterated elimination of dominated strategies. Key to the analysis is an understanding of what happens in extreme regions of the space of values of $p(\theta)$. At the bottom end, running could be a dominant strategy, regardless of what the other individuals do. Denote by $\underline{\theta}$ the value such that

$$u(r_1) = p(\underline{\theta})\, u(R(1 - \lambda\, r_1)/(1 - \lambda))$$

That is, the value such that a patient individual who knew this to be the true value would find it just as good to withdraw funds immediately as to take chances on the funds being available in the future. Note that in this case it is irrelevant what others think. Also note that this boundary increases as the value of r_1 in the bank's contract increases. Since the noise in the signal is never more than ε, we know that anyone who observes a signal $\theta_i < \underline{\theta}(r_1) - \varepsilon$ will withdraw early.

Parametric restriction:

$$u(1)/u(R) > p(2\varepsilon)$$

If this restriction holds, then, even if r_1 goes as low as 1 in the bank contract (and it will never be lower than that), there are still values of θ low enough such that everybody, no matter how optimistic their individual signal, will still know that the true value of $\theta < \underline{\theta}$, and for all of them early withdrawal is a dominant strategy.

At the other end we want to make a similar argument, but this will not be possible with the technology used so far. Instead the authors modify the technology, to guarantee the existence of an upper dominance region parallel to this lower dominance region.

The assumption they use is as follows: There exists a value of $\bar{\theta}$ such that for all θ exceeding $\bar{\theta}$, $p(\theta) = 1$, and short term return is R (rather than 1, as it is for lower values of θ). In other words, in this region of very good news, not only is there no chance of second period failure, the amount available to the bank in the first period is so great that even if everybody decided to withdraw early, it would not exhaust first period funds, and so a patient individual is better off waiting until the second period regardless of what everyone else does.

[Note, we would need to verify that changing the technology in this way doesn't change the optimality of the demandable debt contract in the original DD framework]

Parametric restriction: analogous to previous restriction, assume $\bar{\theta} < 1 - 2\varepsilon$. Then there exists values θ high enough that everybody knows that it is optimal not to run when patient, regardless of others' beliefs.

Therefore we have pinned down everybody's behavior at these extreme signals. But that can be used to pin down behavior at slightly less extreme signals, and so forth, until, surprisingly, behavior is pinned down throughout the range of signals yielding a unique equilibrium:

Main theorem: There is a unique equilibrium in which patient agents run if they observe a signal below some threshold, and do not run if they observe a signal above.

Consequences: As a result for fundamentals far away from the threshold signal, all agents either run or not. For fundamentals close to the threshold signal, individuals with the most positive noise terms do not run and those with the most negative noise terms do, and the proportion running changes smoothly as the fundamental changes. Note the difference between this use of the signal based on fundamentals as a coordinating device and the effect of sunspot equilibria.

The authors emphasize that an innovation in their proof of this result over results in Carlsson and van Damme or Morris and Shin comes from the fact that their results rely on "global strategic complementarities"—an agent's incentive to take an action is higher when other agents take that action. In the case at hand, this is only true at low probabilities of early withdrawal: at low levels, as the probability increases that others are choosing early withdrawal, I want to do so as well. But if the bank is already in the bankruptcy region, this is no longer the case—the more other individuals there are withdrawing early, the lower my chances of getting to be first in line, and so the benefits decline (although they still exceed the benefit of not withdrawing). The authors show the weaker condition suffices for the result: the condition, which they call "one-sided strategic complementarity," is that benefits are increasing in number of agents so long as these benefits are negative. The proof is established in two parts; first that there is a unique threshold equilibria and the second, more technical result, that any equilibrium is a threshold equilibrium.

The intuition for the uniqueness of the threshold equilibrium is as follows. Start with the picture (figure 2 of the text) of the net incentive to withdraw in period 2.

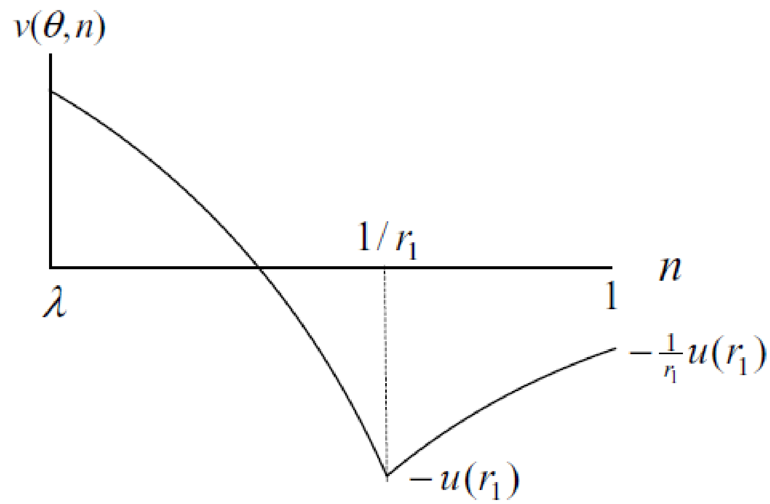

The x axis is the fraction withdrawing early—it always includes the fraction λ of truly impatient people. The y axis is the net benefit from late withdrawal—that is, the *negative* of net benefits of early withdrawal. In other words, the left side is the part showing the one-sided strategic complementarity: as number of early withdrawers increases the net benefit from staying in falls at first (and always falls in the region in which there is net positive benefit to staying).

Suppose we are contemplating a possible equilibrium with threshold θ', and consider the incentives of an agent who obtained signal θ_i. Let $\Delta(\theta_i, \theta')$ represent the expectation of the utility differential; that is simply the expectation of v above given the agent's subjective probability distribution of θ, uniform around θ_i. When integrating over v to determine this expectation it is also necessary to factor in the consideration that for different possible values of θ the number of early withdrawers that the individual believes will show up also changes. In any threshold equilibrium (if one exists) then the agent will prefer to run if her signal θ_i lies below the threshold and not to run if it lies above, so Δ(θ', θ') will have to be zero. But this holds at exactly one point: we know that Δ(θ', θ') is negative at very low values and positive at very high values (the dominance regions) and increasing in θ' (if both the threshold signal and the individual's signal increase by the same amount, the subjective probability that others are withdrawing is unchanged, but waiting is more valuable). Thus there is a unique θ* such that Δ(θ*, θ*).

Thus there is only one candidate for a threshold equilibrium. For it actually to be an equilibrium we need to check that given that threshold for everybody else the single agent's best response respects that threshold. To show this, the authors rely on the single crossing property of v(θ, n(θ, θ')) (Note this is not quite the function in the graph above; rather than holding the number of individuals running constant, we are letting it vary as the state changes. But it turns out this function also exhibits single crossing, and is in fact a weaker condition than the one-sided strategic complementarity). $\Delta(\theta_i, \theta')$ is the integral of this function over the interval $[\theta_i - \varepsilon, \theta_i + \varepsilon]$. Δ(θ*, θ*) = 0.

As in the following picture, depicts, we see that starting from an interval in which the integral of v is zero (this is the interval around θ*, by definition and it perforce contains the point where v equals zero since the function is single-crossing), movement of the interval to the left forces the integral negative (the left portion integrates solely over a negative amount; the portion taken away on the right must be positive, since the whole of the interval is zero) similarly for a move to the right. In other words, single crossing in n induces single crossing of $\Delta(\theta_i, \theta^*)$ in θ_i and this is indeed a threshold equilibrium.

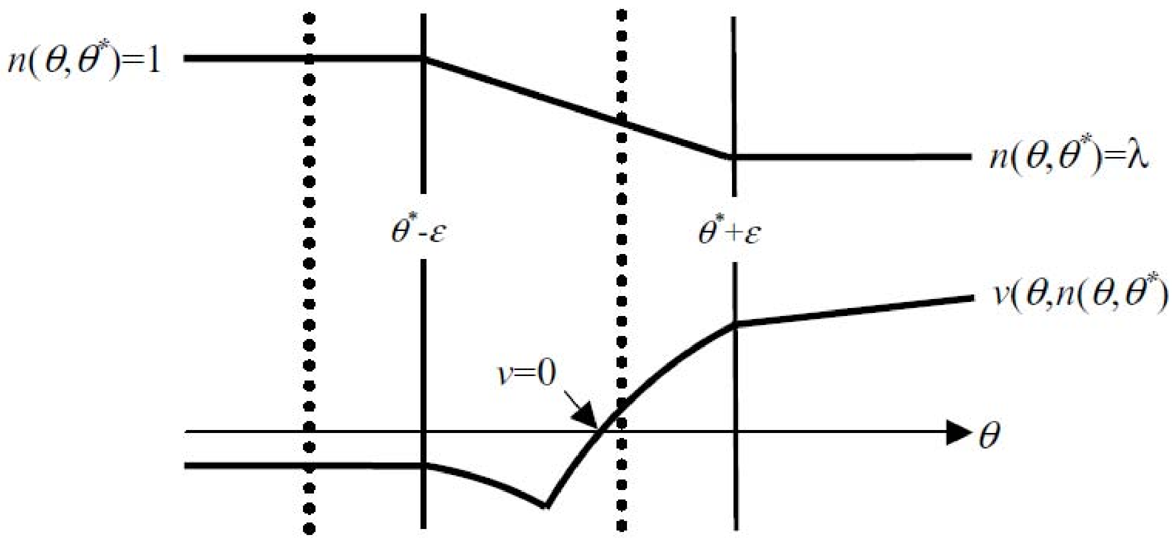

Implications

In this discussion we have omitted the dependence on the choice of the particular demandable debt contract r_1.   The threshold signal can be calculated fairly simply when ε approaches 0 and $\bar{\theta}$ approaches 1. The threshold is increasing in r_1; in other words, increases in promised short term payoff make runs more likely. The optimal r_1 is lower than first best, because lowering the promise reduces the chance of a run. However, provided the lower dominance region is not too large, it is still useful for banks to engage in risk sharing—that is, $r_1 > 1$.