Notes on Diamond and Dybvig (JPE, 1983)

David Andolfatto

February 2008

1 Basics

There are $N > 1$ agents (a finite integer). Each agent has an endowment $y > 0$. There are two periods, $t = 1, 2$ and agents have preferences for 'current' and 'future' consumption (c, f) given by:

$$U(c\ , f; \omega) = \begin{cases} u\,(c + \eta f) & \text{if} \quad \omega = i; \\ u\,(c + f) & \text{if} \quad \omega = p; \end{cases}$$

where u is increasing and concave, and $0 \leq \eta < 1$. Hence, one can interpret $\omega = i$ as someone who turns out to be 'impatient' and $\omega = p$ as someone who turns out to be 'patient.' Let $0 < I < N$ denote the number of agents who will turn out to be impatient. However, *ex ante*, no agent knows his type. Hence, there is no aggregate uncertainty, but agents do face idiosyncratic risk.[1] Define $\pi \equiv I/N$.

All agents have access to a storage technology. In particular, any output not consumed at date 1 can be invested for consumption at date 2; let $R > 1$ denote the gross return on investment.

The autarkic allocation is denoted $\{c^a(\omega), f^a(\omega)\}$. Assume that $\eta R < 1$. Then clearly, $\{c^a(i), f^a(i)\} = \{y, 0\}$ and $\{c^a(p), f^a(p)\} = \{0, Ry\}$. This will generate an *ex ante* utility payoff:

$$V^a = \pi u\,(y) + (1 - \pi) u\,(Ry)\,.$$

The first-best allocation is denoted $\{c^*(\omega), f^*(\omega)\}$. Clearly, $c^*(p) = 0$ and $f^*(i) = 0$. Hence, the decision concerns $\{c^*(i), f^*(p)\}$. A planner (who can exploit the fact that there is no aggregate uncertainty) faces the following resource constraint:

$$\pi N c(i) + R^{-1}(1 - \pi) N f(p) \leq N y.$$

Naturally, this constraint will bind; hence, write:

$$f(p) = (1 - \pi)^{-1} R\,[y - \pi c(i)]$$

and insert into the objective:

$$\max_{c(i)} \left\{ \pi u(c(i)) + (1 - \pi) u\left((1 - \pi)^{-1} R\,[y - \pi c(i)]\right) \right\}.$$

[1] Diamond and Dybvig consider a continuum of agents with *i.i.d.* preference shocks.

The solution must satisfy:

$$u'\left(c^*(i)\right) = Ru'\left(f^*(p)\right); \tag{1}$$

with $f^*(p) = (1-\pi)^{-1}R\left[y - \pi c^*(i)\right]$. Notice that as $R > 1$, we have:

$$y < c^*(i) < f^*(p) < Ry.$$

It should be evident that:

$$V^* = \pi u\left(c^*(i)\right) + (1-\pi)u\left(f^*(p)\right) > V^a.$$

One interpretation here is that the 'planner' is a bank that accepts deposits y and promises a set of returns R_S^*, R_L^*. If you keep your money in the bank, you earn a return $R_L^* = f^*(p)/y$. If you decide to make an early withdrawal, you receive a lower return $R_S^* = c^*(i)/y$. Your desire to insure against risk means that you prefer the smoothed return to the return pattern you could achieve on your own; i.e., (R_S^*, R_L^*) is preferred to $(1, R)$.

2 Private Information

Imagine that the realization of each person's preference shock $\omega \in \{i, p\}$ is private information. Is the first-best allocation implementable? To answer this question, we can construct a mechanism that asks agents to play a simple revelation game (where they simultaneously make a report of their type to the mechanism). By the revelation principle, any achievable allocation in this environment is the allocation corresponding to a truth-telling equilibrium of a revelation game.

Consider any given agent then and assume that all other agents report truthfully. Then the agent under consideration will weakly prefer to report truthfully iff the following incentive-compatibility (IC) conditions are satisfied:

$$\begin{aligned} u(c^*(i)) &\geq u\left(\eta f^*(p)\right); \\ u(f^*(p)) &\geq u(c^*(i)). \end{aligned}$$

The second IC constraint is clearly satisfied. The first IC constraint will be satisfied for a sufficiently small η. I have already assumed that $\eta R < 1$; is this sufficient? This does appear to be the case (you can check).

Hence, every agent has an incentive to report truthfully, conditional on all other agents reporting truthfully. Since this is true for all agents, truth-telling is an equilibrium strategy profile. Moreover, this revelation game implements the first-best allocation.

The key observation here is that when depositor type is private information, an efficient allocation can be implemented by an institution that resembles a bank that accepts deposits and, in exchange, issues **demand-deposit** liabilities. That is, the debt instrument gives the holder to redeem it, for a pre-specified payment, on demand (depending solely on his own private inclination).

The demandable debt instrument here plays a different role from the demandable debt instrument that arises in Calomiris and Kahn (1991); in particular, the Diamond-Dybvig model abstracts from bank agency problems (indeed, the bank is like a benevolent planner). In this model, the bank issues liabilities that can be redeemed on demand for risk-sharing purposes (allowing individuals to smooth consumption across different contingencies). Absent private information, a regular insurance contract would do the same thing. But if it is costly (or impossible) to demonstrate verifiable evidence of your type, then it makes sense to allow for self-reports (make the debt object demandable). You might note that every time you visit an ATM to make a withdrawal, the bank is essentially allowing you to make a self-report of your type (someone who needs cash now).

3 Multiple Equilibria and Bank Runs

It is important to understand that the Revelation Principle simply asserts that a given allocation might be implementable as an equilibrium to a revelation game in which all agents tell the truth. Whether this equilibrium is unique is an entirely different matter. In general, any given revelation game may have multiple equilibria; in particular, equilibria in which agents do not report their types truthfully.

In the banking literature, a "bank-run" would correspond to a scenario where patient agents misreport their type *en masse*. That is, even agents who do not have a pressing need for "liquidity" may be compelled to "run" to the bank to make a withdrawal if they–for whatever reason–believe that other patient agents will behave in a similar manner.

The basic intuition for this possibility is as follows. Consider the first-best allocation. This allocation promises anyone who wants to withdraw funds early an amount $c^*(i) > y$. But if all agents choose–for whatever reason–to withdraw early, the bank would be obliged to payout the amount $Nc^*(i)$; which clearly violates its resource constraint (it does not have the deposits to make good on it obligations). If depositors are served on a first-come first-served basis (a sequential service constraint), then a patient depositor will want to "run" the bank if he believes that others will do likewise. By running and potentially collecting the amount $c^*(i)$, he does better than by waiting and receiving nothing at all.

Despite the intuitive appeal of the argument above, the notion of an equilibrium bank run has proven difficult to formalize in a persuasive manner. Generating a bank-run equilibrium is easy to do; but only if one relies on suboptimal (*ad hoc*) contracts (as is frequently done in the literature). If the threat of a bank-run is a real possibility, one might reasonably imagine that an institution would alter the contractual form to mitigate the possibility. Let me show you what I mean by this.

3.1 A Bank Run Example

For simplicity, assume that $N = 3$ and that $I = 1$ (one agent is impatient and two are patient). Of course, agents do not know their types *ex ante*; all they know is that with probability 1/3 they will be impatient, and with probability 2/3 they will be patient. The actual "names" of the agents do not matter; i.e., we can treat $\{i, p, p\}$, $\{p, i, p\}$, and $\{p, p, i\}$ as equivalent outcomes.

Consider a direct revelation game in which agents are required to report their types simultaneously (there is no sequential service constraint). Note that since there is no aggregate uncertainty over true types, the mechanism will have some capacity for detecting lies. For example, the mechanism will know that at least one agent lied if it receives a set of reports containing only one p. On the other hand, the mechanism will not be able to distinguish (say) between the reports $\{i, p, p\}$ and $\{p, p, i\}$. We can abstract from this latter consideration (I am not even sure it is important) by assuming $\eta = 0$. In this case, it is a dominant strategy for the impatient agent to report truthfully. Now the mechanism can detect for sure whether someone has lied (although, not necessarily the identity of who lied).

An agent's true type ω lies in the set $\{i, p\} \equiv \Omega$. Likewise, an agent's report of his type m lies in the set $M = \Omega$. In the revelation game, each agent formulates a (pure) strategy:

$$s : \Omega \times M \times M \to M.$$

Let Σ denote the set of (pure) strategies $s(\omega, m, m)$. A **truth-telling strategy** is defined by $s(\omega, m, m) \equiv \omega$.

By construction, the impatient agent will always tell the truth; i.e., $s(i, m, m) = i$ for all $(m, m) \in M^2$. I am going to define a **bank run** as a scenario in which both patient agents misreport their types; i.e., $s(p, m, m) = i$ for some $(m, m) \in M^2$ (for both p agents). The question is whether a strategy profile of this form exists as an equilibrium.

The Revelation Principle allowed us to restrict attention to allocations in which agents report truthfully. However, as we are now interested in the possibility of non-truth-telling behavior, we have to specify the allocation for all potential strategy profiles. In particular, we have to be concerned about feasibility. For example, we know that assigning $c^*(i)$ to each agent is infeasible under a bank run strategy profile.

To this end, define the functions:

$$\begin{aligned} c &: M^3 \to \mathbb{R}_+; \\ f &: M^3 \to \mathbb{R}_+. \end{aligned}$$

Interpret $c(m_j, m_{-j})$ as current consumption for agent j given his own report m_j and the reports submitted by the other agents m_{-j}; and likewise for $f(m_j, m_{-j})$. An **allocation** is defined to be $\{c_j, f_j : j = 1, 2, 3\}$. A **feasible allocation** satisfies:

$$R\left[Ny - \sum_{j=1}^{3} c(m_j, m_{-j})\right] \geq \sum_{j=1}^{3} f(m_j, m_{-j}) \text{ for all } (m_j, m_{-j}) \in M^3.$$

Now, consider the following allocation $\{\hat{c}, \hat{f} : j = 1, 2, 3\}$:

$$\hat{c}(i, m, m) = \begin{cases} c^*(i) & \text{if} \quad (m, m) = (p, p); \\ 0 & \quad\quad \text{otherwise;} \end{cases}$$

$$\hat{c}(p, m, m) = 0 \text{ for all } (m, m) \in M^2;$$

$$\hat{f}(p, m, m) = \begin{cases} f^*(p) & \text{if} \quad (m, m) = (p, i) \text{ or } (m, m) = (i, p); \\ 0 & \quad\quad \text{otherwise;} \end{cases}$$

$$\hat{f}(i, m, m) = 0 \text{ for all } (m, m) \in M^2.$$

In short, the allocation assigns the first-best allocation to the agents if the mechanism collects a truthful set of reports; and otherwise assigns zero consumption to all agents. This allocation is clearly feasible.

The idea I am developing here reminds me very much of Holmstrom (1982). That is, the "team production" problem here does not feature unobservable effort; but rather, unobservable types (and observable reports). If all agents "play nice" (report truthfully), they achieve the first-best allocation; if even one agent deviates from truthtelling, all agents get massacred. Note that I am assuming here (as is standard) that the mechanism can **commit** to deliver such an allocation. I am also allowing for **weak budget balance** (the resource constraint need not be satisfied with equality).

It should be clear enough that truth-telling is an equilibrium. But it also appears that the bank run strategy profile is an equilibrium. That is, imagine that one of the patient agents expects that the other patient agent is going to lie. Then no matter what this agent reports, he will expect to receive zero consumption. Since he is indifferent between lying and telling the truth, assume that he lies (intuitively, one might expect that he would instead tell the truth; but we could replace zero current consumption with a strictly positive "epsilon"). As the other patient agent faces a symmetric situation, he too is indifferent between lying and telling the truth. Hence, lying for both patient agents is an equilibrium.

3.2 Preventing Bank Runs

A bank run equilibrium exists in the example above because of poor institutional design (a poorly designed allocation). It is, in fact, a simple matter to modify the allocation in a manner so that truth-telling is the unique equilibrium (this may not be true for more general environments). For example, imagine that instead of assigning zero consumption when the mechanism detects a lie, that it instead allocates $\hat{c}(i) = y$ to all agents reporting $m = i$ and $\hat{f}(p) = Ry$ to all agents reporting $m = p$. One might interpret such an action as **partial suspension** (the bank allows agents to withdraw y instead of $c^*(i) > y$) in the event of a bank run (such behavior has been observed in reality).

Now, imagine that a patient agent expects the other to misreport his type. Then the bank will be left with y units of output ($2y$ will be withdrawn) and will invest the proceeds, generating a future return Ry for the patient agent. If this agent waits, he receives $u(Ry)$; if he runs, he receives $u(y)$. Clearly, he will prefer not to run; and similarly for the other patient agent.

In fact, under this institutional structure, all agents find truth-telling a **dominant strategy** (it is individually rational to tell the truth regardless of how others report their types). It follows directly that truth-telling is the unique equilibrium.

There is an issue as to how general this result is. Green and Lin (2003) extend the environment here to allow for aggregate uncertainty over types together with an explicit **sequential service constraint**. They find that the first-best allocation is uniquely implementable (although, the optimal allocation looks quite complicated and rather unrealistic). Andolfatto, Nosal and Wallace (2007) generalize their result. On the other hand, Peck and Shell (2003) argue for the existence of an equilibrium bank run; their result, however, appears to hinge on an *ad hoc* (but intuitively plausible) restriction placed on the mechanism.

In a recent paper, Ennis and Keister (2007a) claim to find bank run equilibria in a Green-Lin style model (I have not read the paper carefully). Their environment extends Green and Lin (2003) by assuming that agents face a cost of sending messages (contacting the bank).

In another paper, Ennis and Keister (2007b) consider an environment without aggregate uncertainty, but where the mechanism lacks commitment. They demonstrate that the lack of commitment can generate equilibrium bank runs. Unfortunately, the Revelation Principle generally does not hold when there is a lack of commitment; see Bester and Strausz (2000).

4 References

1. Andolfatto, David, Ed Nosal and Neil Wallace (2007). "The Role of Independence in the Green-Lin Diamond-Dybvig Model," *Journal of Economic Theory*, 137: 709–715.

2. Bester, Helmut and Roland Strausz (2000). "Imperfect Commitment and the Revelation Principle: The Multi-Agent Case," *Economic Letters*, 69(2): 165–171.

3. Calomiris, Charles and Charles Kahn (1991). “The Role of Demandable Debt in Structuring Optimal Banking Arrangements,” *American Economic Review*, 81(3): 497–513.

4. Diamond, Doug and Philip Dybvig (1983). “Bank Runs, Deposit Insurance, and Liquidity,” *Journal of Political Economy*, 91: 401–419.

5. Ennis, Huberto and Todd Keister (2007a). “Run Equilibria in a Model of Financial Intermediation,” Manuscript.

6. Ennis, Huberto and Todd Keister (2007b). “Commitment and Equilibrium Bank Runs,” Manuscript.

7. Green, Ed and Ping Lin (2003). “Implementing Efficient Mechanisms in a Model of Financial Intermediation,” *Journal of Economic Theory*, 109: 1–23.

8. Peck, James and Karl Shell (2003). “Equilibrium Bank Runs,” *Journal of Political Economy*, 111: 103–123.