Federal Reserve Bank of New York
Staff Reports

Run Equilibria in a Model of Financial Intermediation

Huberto M. Ennis
Todd Keister

Staff Report no. 312
January 2008

Run Equilibria in a Model of Financial Intermediation
Huberto M. Ennis and Todd Keister
Federal Reserve Bank of New York Staff Reports, no. 312
January 2008
JEL classification: D82, D84, G21

Abstract

We study the Green and Lin (2003) model of financial intermediation with two new features: traders may face a cost of contacting the intermediary, and consumption needs may be correlated across traders. We show that each feature is capable of generating an equilibrium in which some (but not all) traders "run" on the intermediary by withdrawing their funds at the first opportunity regardless of their true consumption needs. Our results also provide some insight into elements of the economic environment that are necessary for a run equilibrium to exist in general models of financial intermediation. In particular, our findings highlight the importance of information frictions that cause the intermediary and traders to have different beliefs, in equilibrium, about the consumption needs of traders who have yet to contact the intermediary.

Key words: bank runs, optimal contracts, private information, incentive feasibility, self-fulfilling expectations

---

Ennis: Federal Reserve Bank of Richmond (e-mail: huberto.ennis@rich.frb.org). Keister: Federal Reserve Bank of New York (e-mail: todd.keister@ny.frb.org). The authors thank participants at the Cornell–Penn State Macroeconomics Workshop for helpful comments.

1 Introduction

Bank runs and financial panics are often thought to be self-fulfilling phenomena, in the sense that individuals withdraw their funds in anticipation of a crisis and, together, these individual actions generate the crisis that everyone feared. A substantial literature has arisen asking whether or not, and under what circumstances, a self-fulfilling bank run can be the outcome of an economic model with optimizing agents and rational expectations. Early contributions to this literature assumed particular institutional arrangements, such as a bank offering a demand-deposit contract. In an influential recent paper, Green and Lin [6] study a model very much in the spirit of the classic work of Diamond and Dybvig [4] but with no restrictions on contracts other than those imposed by the physical environment. They derive a striking result: in their environment, the efficient allocation can be uniquely implemented. In other words, a financial intermediary can offer a contract that guarantees the efficient outcome will obtain in equilibrium, leaving no possibility of a self-fulfilling run.

The Green-Lin result opens the question of whether there exist any reasonable economic environments in which self-fulfilling runs can occur in the absence of arbitrary institutional restrictions. We identify two such environments, both of which are close variants of that in Green and Lin [6]. In one of our settings, consumption needs are correlated across agents, while in the other it is costly for agents to contact the financial intermediary. Despite their apparent differences, these two features have a common effect in equilibrium: they exacerbate the existing information frictions in the Green-Lin model and, as a result, allow for self-fulfilling runs to be consistent with equilibrium.

It is well known that some information frictions are necessary for the possibility of a bank run to arise. If agents' consumption needs were observable, for example, the intermediary could guarantee the efficient outcome by offering a simple insurance contract that makes payments to each depositor conditional on her realized preferences. In such a setting, agents would not even have the option to run. For this reason, Diamond and Dybvig [4] assumed that an agent's consumption needs are private information; implementing the efficient allocation then requires that agents be given a choice of when to withdraw their funds from the intermediary. This element of choice is clearly necessary for a run to occur.

Private information alone is not enough to generate a bank run in equilibrium, however. If the intermediary can condition the payments it offers to agents on the total demand for early withdrawals, it can again guarantee the efficient outcome, even though it does not know which or even how many agents truly need to consume right away. This is because, in such an environment, the efficient response of the intermediary to high withdrawal demand is to make a smaller payment to each withdrawing agent. In doing so, the intermediary makes withdrawing early less attractive, thereby decreasing the incentive for an agent to “panic” and withdraw when she has no immediate need to consume. The result is that, when the intermediary adjusts payments efficiently, each agent has a dominant strategy to withdraw only when she needs to consume.[1] No run can occur, even though agents have private information.

In order for a self-fulfilling run to be possible, then, some friction must prevent the intermediary from being able to condition payments to all agents on total withdrawal demand. To capture this idea, Diamond and Dybvig [4] included a first-come, first-served (or “sequential service”) constraint in their analysis. Wallace [10] formalized this notion by identifying features of the economic environment that would generate such a constraint.[2] In particular, Wallace assumed that agents are isolated from each other and visit the intermediary sequentially. Each agent must consume upon arrival at the intermediary and, therefore, the intermediary must make a payment to an agent before observing the actions of subsequent agents. Notice that, fundamentally, the sequential service constraint is another type of information friction; in an environment with sequential service, the intermediary must make payments to agents before observing total withdrawal demand.

Wallace [11] studied an environment with an explicit sequential service constraint and with aggregate uncertainty, where the number of agents who need to consume early is random. He showed that in states where many agents need to consume early, those who contact the intermediary first receive higher levels of consumption in the efficient allocation than those who contact the intermediary last. Such an event may be interpreted as a form of banking crisis, but it is clearly distinct from a self-fulfilling run on the intermediary, which leads to an inefficient allocation of resources. It remained an open question whether or not the combination of aggregate uncertainty and a sequential service constraint might generate self-fulfilling bank run equilibria.

Green and Lin [6] focused on a special case of aggregate uncertainty, where the realization of consumption needs is independent across agents, and recast the issue as a mechanism design problem. They departed from the previous literature by assuming that agents have information about the order in which they will have an opportunity to withdraw their funds. They showed that

[1] See Green and Lin [6, Theorem 1] for a formal statement and proof of this result.

[2] Such a constraint also makes banking essential, in the sense of being able to achieve outcomes not achievable through markets. See Jacklin [8] and Wallace [10] on this point.

the efficient allocation in this environment can be implemented using a direct revelation mechanism (i.e., this allocation is Bayesian incentive compatible). They also showed that, surprisingly, the direct revelation game always has a unique Bayesian Nash equilibrium allocation. In other words, in their version of the Diamond-Dybvig model, financial intermediaries are not inherently fragile; the efficient allocation can be implemented without raising the possibility of a bank run. This demonstrates that sequential service does not necessarily imply the possibility of a run. Andolfatto *et al.* [2] extended this result to a wider class of preferences. The result led Green and Lin [7] to ask "What's Missing" in the model. In other words, what feature(s) of the environment would permit bank runs to occur?

One answer to this question was provided by Peck and Shell [9], who also studied a model without any institutional or other restrictions on contracts. In the Peck-Shell model (as in the earlier work of Diamond and Dybvig [4] and others), agents must decide whether or not to withdraw their funds from the intermediary before knowing the order in which they would sequentially contact the intermediary. Relative to Green and Lin [6], this is an additional information friction: agents must act before knowing this payoff-relevant information. Peck and Shell present an example in which a bank run equilibrium exists. Their example, however, relies on agents having different preferences than in the previous literature, leaving open the question of what exactly is responsible for their result. In section 2 we present an example in the spirit of Peck and Shell, but in which preferences (and other aspects in the model) are exactly as in Green and Lin [6]. Our example demonstrates that it is the additional information friction, and not the difference in preferences, that accounts for the difference in results.

We then return to the Green-Lin framework, where agents know the order in which they are able to contact the intermediary, and investigate what other combinations of frictions can generate a run equilibrium. We present two variations on the Green-Lin environment. First, we allow for consumption needs to be correlated across agents. In this case, an agent's private information about her own type is also (private) information about the types of other agents. We construct an example where types are negatively correlated. In this example, an agent who withdraws early when she does not need to consume right away will make the intermediary unduly optimistic about the consumption needs of the remaining agents. The intermediary will then conserve relatively few resources for the next period and, as a result, agents who wait to withdraw will end up with low consumption levels. This fact, in turn, gives some agents an incentive to withdraw early if they expect others to do likewise. In this way, the asymmetry between the beliefs of the agent and the (unduly optimistic) intermediary allows a self-fulfilling run to be consistent with equilibrium. We also show that a bank run in this setting is necessarily partial, with only some of the agents participating.

In our second variant, agents face a cost of contacting the intermediary. In Green and Lin [6], all agents contact the intermediary in the early period regardless of their withdrawal intentions. The intermediary thus observes not only the decisions of agents who choose to withdraw in the current period, but also those of agents who decide instead to withdraw at a later date; current-period payments can then be conditioned on both types of information. We modify the environment by introducing a utility cost of contacting the intermediary. This cost can be thought of as the shoe-leather cost of physically visiting the intermediary, but it could also represent more broadly the cost of monitoring one's transaction needs (or lack thereof) and communicating them to the intermediary on a regular basis.

Agents have an opportunity to contact the intermediary in the same order as before, but now, because of the cost involved, they may not do so unless they want to withdraw their funds. If agents only contact the intermediary when they want to withdraw, the intermediary must act with less information than in the Green-Lin model.[3] We show that the introduction of costly communication can also generate an equilibrium in which some, but not all, agents run on the intermediary, even for the case where types are independent. The intuition is broadly similar to that described above for the case of correlated types. By restricting the flow of information to the intermediary, the additional friction can also create a wedge between the equilibrium beliefs of agents and the beliefs used to design the efficient allocation. This wedge can again generate incentives for some agents to withdraw early if they believe others are doing so.

In the next section, we present the general environment for our analysis and derive the efficient allocation. In Section 3, we discuss how this allocation can be implemented in the presence of private information. We also present the main result of Green and Lin [6] in the context of our model and an example in the spirit of Peck and Shell [9]. In Section 4, we discuss the case of correlated types and present examples of run equilibria for this case, while we do the same for the case of costly communication in Section 5. We offer some concluding remarks in Section 6.

[3] This analysis is, in some ways, closer to the approach taken in Diamond and Dybvig [4], where the intermediary was assumed to react only to withdrawal requests.

2 The Model

In this section we present a version of the Green-Lin model with two new features: traders may face a cost of contacting the intermediary and types may be correlated across traders. We then derive the efficient allocation in this environment under different assumptions about the size of the cost of contacting the intermediary.

2.1 The environment

There are two time periods, indexed by $t \in \{0,1\}$, and a finite number I of traders. Let $\mathbb{I} = \{1,2,\ldots,I\}$ denote the set of traders. There is a single good that can be consumed in each period. There is also an intermediary that acts as a benevolent planner and attempts to distribute resources to maximize traders' expected utility, subject to the constraints described below.

Technology. Traders are isolated from each other, but have an opportunity to contact the intermediary in each period in order to receive goods. Let $c_i^t \in \mathbb{R}_+$ denote the consumption of trader i in period t and let $d_i^t \in \{0,1\}$ be a binary variable that represents whether or not the trader contacts the intermediary in a given period. Specifically, let $d_i^t = 1$ if trader i contacts the intermediary in period t and $d_i^t = 0$ if she does not. Feasibility then requires that, for each trader i,

$$\left(1 - d_i^t\right) c_i^t = 0, \quad \text{for } t = 0,1. \tag{1}$$

In other words, if trader i does not contact the intermediary in period t, she cannot consume in that period. Goods are nonstorable and must be consumed immediately after contacting the intermediary in order to give utility.[4]

The intermediary has an aggregate endowment of I units of the good at date 0. Each unit of the good that is not consumed in the early period is transformed into R units of the good in period 1. Let a_i denote the individual allocation of trader i, that is, the specification of whether or not she contacts the intermediary and how much she consumes in each period,

$$a_i = \left(c_i^0, c_i^1, d_i^0, d_i^1\right).$$

Let $a = (a_1, \ldots, a_I)$ denote the complete vector of individual allocations. An (ex post) allocation

[4] This assumption implies that markets in which agents could trade after contacting the intermediary are infeasible. See Wallace [10].

in this environment is an assignment of an individual allocation a_i to each trader. We denote the set of *feasible (ex post) allocations* by

$$\mathbb{A} = \left\{ a : \mathbb{I} \to \mathbb{R}^2_+ \times \{0,1\}^2 : \sum_{i \in \mathbb{I}} \left(c_i^0 + \frac{c_i^1}{R} \right) \leq I \text{ and (1) holds} \right\}.$$

A *state-contingent allocation* is a mapping from states to (ex post) allocations; we denote such a mapping by $\mathbf{a}$. The set of feasible state-contingent allocations is then

$$\mathbb{F} = \left\{ \mathbf{a} : \Omega^I \to \mathbb{A} \right\}.$$

We use the bold-faced variables $\mathbf{c}$ and $\mathbf{d}$ to denote the consumption and contact components, respectively, of a state-contingent allocation $\mathbf{a}$.

Preferences. A trader's consumption preferences depend on her type $\omega_i \in \{0,1\}$. If $\omega_i = 0$, the trader is *impatient* and only cares about consumption in period 0. If $\omega_i = 1$, the trader is *patient* and cares about the sum of her consumption in the two periods. A trader's type is private information. Let $\omega = (\omega_1, \ldots, \omega_I)$ denote the vector of types for all traders. As discussed below, types will be revealed sequentially; we therefore refer to ω as the *history* of types. Let Ω denote the set $\{0,1\}$, so that we have $\omega_i \in \Omega$ and $\omega \in \Omega^I$.

As described above, in order to consume in a particular period, a trader must contact the intermediary in that period. Contacting the intermediary may be costly. Specifically, we assume each trader loses $\delta_1 \geq 0$ units of utility if she contacts the intermediary once in the two periods and $\delta_2 \geq \delta_1$ units if she contacts the intermediary in both periods. Trader i's utility level is given by

$$v(a_i; \omega_i) = \frac{1}{1-\gamma} \left(c_i^0 + \omega_i c_i^1 \right)^{1-\gamma} - \delta_{d_i^0 + d_i^1}, \tag{2}$$

where $\gamma > 1$ is assumed to hold.[5]

Note that the cost of making a single trip to the intermediary, δ_1, does not depend on whether this trip occurs in period 0 or in period 1. Of course, the intermediary is only useful to a trader if she visits it at least once. The cost δ_1 thus represents a kind of fixed cost of intermediation and, without any loss of generality, we can set $\delta_1 = 0$.[6] The important cost in this model is δ_2,

[5] The assumption of a specific functional form for the utility function is not necessary here; Green and Lin [6] assume only that the coefficient of relative risk aversion in consumption is everywhere greater than unity. However, since this specific form simplifies the derivation of the efficient allocation substantially and is also used in our examples below, we make the assumption from the outset (as in Green and Lin [7]).

[6] Having $\delta_1 > 0$ would not change any of the analysis in this paper. It would, however, affect the attractiveness

which measures the cost to a trader of contacting the intermediary in *both* periods. As we will see below, the efficient allocation has each trader consuming in only one period (period 0 for impatient traders and period 1 for patient ones). If a trader contacts the intermediary twice, therefore, in one of the visits she will not receive any consumption. The purpose of this visit would be solely to communicate her type, which gives the intermediary useful information about the history ω. The cost δ_2, therefore, represents a form of information friction: how costly is it for the intermediary to learn the type of a trader even when that trader has no immediate need to consume? When $\delta_2 = \delta_1 = 0$, these preferences reduce to those used in Green and Lin [6] and elsewhere.

Uncertainty. Let $\mathcal{P}$ denote the probability measure on the set of all subsets of Ω^I. We assume that $\mathcal{P}$ treats all traders equally in the sense that each trader has the same *ex ante* probability of being patient. Specifically, we require that there exist a non-negative function p with

$$\sum_{\theta=0}^{I} p(\theta) = 1$$

such that

$$\mathcal{P}(\omega) = \frac{p(\theta(\omega))}{C(I, \theta(\omega))} \quad \text{for all } \omega, \tag{3}$$

where C is the standard combinatorial function

$$C(I, \theta) = \frac{I!}{\theta!\,(I-\theta)!}$$

and $\theta(\omega)$ is the number of patient traders in the state ω.

This approach is the same as that taken in Wallace [10] and can be thought of in the following way: nature first chooses θ according to the density function p, and then θ traders are chosen at random (with each trader equally likely to be chosen) and assigned $\omega_i = 1$. The remaining traders are assigned $\omega_i = 0$. The assumption of independent types used by Green and Lin [6] is a special case where the density p is given by the binomial distribution

$$p(\theta) = C(I, \theta)(1-\pi)^{\theta} \pi^{I-\theta},$$

with $\pi \geq 0$ being the probability with which each individual trader is impatient.

---

of intermediation relative to an outside option (such as autarky). To simplify the analysis, we abstract from such outside options here and, hence, there is no loss in normalizing δ_1 to zero.

Isolation and Sequential Service. Traders are isolated from each other and from the intermediary. They do not observe each others' actions. The only way information can be communicated in this environment is by traders contacting the intermediary. Each trader has an opportunity to contact the intermediary in each period. This opportunity arrives sequentially in a fixed order given by the index i, beginning with trader 1 and ending with trader I.[7]

This physical structure of the environment places two important restrictions on the allocation $\mathbf{a}$. First, whether or not trader i contacts the intermediary in period 0 can depend on her own type ω_i, but cannot depend on the type of any other trader since there is no way she could observe this information before her opportunity to contact the intermediary arrives. We can write this *isolation constraint* as

$$\mathbf{d}_i^0\left(\omega\right) = E\left[\mathbf{d}_i^0\left(\omega\right)\ \mid \omega_i\right]. \tag{4}$$

In other words, trader i's action can only depend on information that she has at the time the action is taken, and the only information she can possibly have before contacting the intermediary is her own type.

The second restriction is the *sequential service constraint*, which follows Wallace [10] and others. This constraint states that the period-0 consumption of trader i cannot depend on information the intermediary could not possibly have received from either trader i or the previous traders in the order. What information the intermediary could have received from these traders depends, in turn, on which of them have contacted the intermediary in the early period. In other words, the intermediary can only potentially receive information from trader j in period 0 if $d_j^0 = 1$. Let $\Lambda_i\left(\omega\right) \subset \mathbb{I}$ denote, for a given allocation, the set of traders up to trader i who contact the intermediary in period 0, that is,

$$\Lambda_i\left(\omega\right) = \left\{j \in \mathbb{I} : j \leq i \text{ and } d_j^0\left(\omega\right) = 1\right\}.$$

Then sequential service requires that the consumption of trader i depend only on information obtained from traders in the set Λ_i. This constraint can be written as

$$\mathbf{c}_i^0\left(\omega\right) = E\left[\mathbf{c}_i^0\left(\omega\right)\ \mid\ \left\{\omega_j\right\}_{j\in\Lambda_i(\omega)}\right]. \tag{5}$$

[7] We follow Green and Lin [7] and Andolfatto *et al.* [2] in assuming that traders contact the intermediary in a fixed order, rather than in a random order as in Green and Lin [6]. The two approaches lead to similar results, and adopting the fixed-order approach simplifies the notation considerably.

In other words, trader i must consume the same amount in any two states that the intermediary cannot possibly distinguish between given the information it could have potentially received so far. We denote the set of feasible state-contingent allocations that satisfy the isolation and sequential service constraints by

$$\mathbb{F}' = \{\mathbf{a} \in \mathbb{F} : (4) \text{ and } (5) \text{ hold}\}.$$

The simplicity of the expression in (5) belies the subtle complexities of sequential service in our environment. In particular, calculating the expectation on the right-hand side requires taking into account the circumstances under which each trader will and will not contact the intermediary in period 0. In other words, the expectation operator is itself a function of the allocation $\mathbf{a}$ through the contact component $\mathbf{d}$. In what follows, we simplify the analysis by focusing on two special cases. First we consider the case studied by Green and Lin [6], where $\delta_2 = 0$ holds and it is clearly efficient for all traders to contact the intermediary in period 0. In this case, the period-0 consumption of trader i can, in principle, depend on the entire partial history ω^i. In other words, in this case efficiency requires $d_i^0 = 1$ for all i and, given this fact, the sequential service constraint becomes

$$\mathbf{c}_i^0(\omega) = \mathbf{c}_i^0(\widehat{\omega}) \quad \text{for all } \omega,\ \widehat{\omega} \text{ such that } \omega^i = \widehat{\omega}^i,\ \text{ for all } i. \tag{6}$$

The second case we study is where δ_2 is large enough that, in the efficient allocation, each trader contacts the intermediary only once. In this case, impatient traders will contact the intermediary in period 0 and patient traders will contact the intermediary only in period 1. As a result, sequential service implies that the period-0 consumption of trader i can only depend on the number of impatient traders before her in the order; the intermediary has no information in period 0 about patient traders who might be before i in the order. In other words, when the first impatient trader arrives, the intermediary only knows that $\omega_i = 0$ for at least one trader i. This trader's arrival does not change the relative probabilities the intermediary assigns to any two states in which at least one trader is impatient. The consumption of the first impatient trader in the order must, therefore, be the same in all states.

More generally, let $\theta_i(\omega^i)$ denote the number of patient traders in the partial history ω^i. Then efficiency when δ_2 is large requires $d_i^0 = (1 - \omega_i)$ for all i and, given this fact, the sequential service constraint becomes

$$\mathbf{c}_i^0(\omega) = \mathbf{c}_j^0(\widehat{\omega}) \text{ for all } \omega, \widehat{\omega} \text{ with } \omega_i = \widehat{\omega}_j \text{ and } \theta_i\left(\omega^i\right) = \theta_j\left(\widehat{\omega}^j\right) \tag{7}$$

for all combinations of i and j. Note that the set of consumption allocations $\mathbf{c}$ satisfying (7) is a strict subset of those satisfying (6). In other words, sequential service is a stronger constraint when only impatient traders contact the intermediary in period 0 because it leads the intermediary to act with strictly less information.

Expected Utility. Once a trader learns her type, she seeks to maximize the expected value of the utility function v conditional on this type. We can write the information set of trader i as

$$\mathcal{E}_i = \left\{\varnothing, \Omega^I, \{\omega|\omega_i = 0\}, \{\omega|\omega_i = 1\}\right\}.$$

Given a state-contingent allocation $\mathbf{a}$ and a (true) state of nature ω^*, define

$$U_i(\mathbf{a}, \omega^*) = E\left[v(\mathbf{a}_i(\omega), \omega) \mid \mathcal{E}_i(\omega^*)\right].$$

Notice that the value taken by U_i depends only on the element $\mathbf{a}_i$ of the allocation $\mathbf{a}$; payments made to other traders do not directly affect trader i's utility. In addition, the function U_i is $\mathcal{E}_i$-measurable, implying that for a given allocation $\mathbf{a}$ it takes on at most two values, one for $\omega_i = 0$ and another for $\omega_i = 1$.

2.2 The efficient allocation when $\delta_2 = 0$

We now derive the efficient, symmetric state-contingent allocation, that is, the allocation the intermediary would assign if traders' types were observable.[8] We begin with the case where all traders contact the intermediary in period 0. While this solution has been partly characterized before for the case of independent types (see, for example, Green and Lin [7]), ours is the first complete solution of the efficient allocation in the Green-Lin model for an arbitrary number of traders, as well as the first to allow for correlation in types.

[8] Note that the efficient allocation here will typically be different from the full-information first-best allocation under no aggregate uncertainty as studied by Diamond and Dybvig [4]. When there is no aggregate uncertainty, the sequential service constraint is nonbinding and the first-best allocation is the same as in an environment without sequential service. In the presence of aggregate uncertainty, on the other hand, the sequential service constraint always binds in the efficient allocation.

The efficient allocation is the solution to

$$\max_{\mathbf{a}\in\mathbb{F}'} \sum_{i\in\mathbb{I}} E\left[U_i\left(\mathbf{a},\omega\right)\right]. \tag{8}$$

Let $\mathbf{a}^*$ denote this solution. We have argued above that when $\delta_2 = 0$, the efficient allocation has $d_i^0 = 1$ for all i and that the sequential service constraint reduces to (6). It is straightforward to show that, under the preferences in (2), efficiency requires that impatient traders only consume at date 0 and patient traders only consume at date 1. In other words, the efficient (state-contingent) allocation $\mathbf{a}^*$ must have

$$\mathbf{c}_i^0\left(\omega\right) = 0 \text{ if } \omega_i = 1 \qquad \text{and} \qquad \mathbf{c}_i^1\left(\omega\right) = 0 \text{ if } \omega_i = 0. \tag{9}$$

In addition, it is easy to see that the resources remaining at date 1 will be divided evenly among the patient traders in this allocation, that is,

$$\mathbf{c}_i^1\left(\omega\right) = \frac{R\left(I - \sum_{i=1}^{I} c_i^0\left(\omega\right)\right)}{\theta_I}. \tag{10}$$

All that remains, then, is to determine the payment that would be given to each trader i at date 0 if she is impatient, as a function of the partial history ω^i. In other words, we need to determine $\mathbf{c}_i^0\left(\omega\right)$ for histories with $\omega_i = 0$. These payments can be found by using the results above to reformulate (8) as a dynamic programming problem.

Our formulation of the problem makes use of some important implications of condition (3), which governs the correlation structure of types. First, the condition implies that any two histories ω and $\widehat{\omega}$ with $\theta\left(\omega\right) = \theta\left(\widehat{\omega}\right)$ are assigned the same probability by $\mathcal{P}$.[9] Second, consider the probability of some continuation history $\omega^{I-i} = \left(\omega_{i+1}, \ldots, \omega_I\right)$ conditional on the partial history $\omega^i = \left(\omega_1, \ldots .\omega_i\right)$. Condition (3) implies that this probability depends only on the number of patient traders in the partial history, denoted $\theta_i\left(\omega^i\right)$, and not on their positions within the history. Abusing notation slightly, let $\mathcal{P}\left(\omega^i\right)$ denote the probability of the partial history ω^i, that is, the probability of the set $\left\{\widetilde{\omega} \in \Omega^I : \widetilde{\omega}^i = \omega^i\right\}$. Then the following lemma establishes these two claims and, thus, shows how θ_i is a useful summary statistic for ω^i. A proof of this lemma is given in Appendix A.

9 This fact is easily seen in (3), where the expression on the right-hand side depends on $\theta\left(\omega\right)$ but not directly on ω.

Lemma 1 *Under* (3), $\theta_i(\omega^i) = \theta_i(\widehat{\omega}^i)$ *implies both*

$$\mathcal{P}(\omega^i) = \mathcal{P}(\widehat{\omega}^i)$$

and

$$\mathcal{P}(\omega^i, \omega^{I-i}) = \mathcal{P}(\widehat{\omega}^i, \omega^{I-i}) \quad \textit{for all } \omega^{I-i}.$$

Now consider the problem faced by the intermediary when it encounters trader i. Let y_{i-1} denote the amount of resources it has remaining after the first $i-1$ encounters. If trader i is impatient, the intermediary must decide how much of y_{i-1} should be given to her and how much should be saved for future payments, including those to patient traders at date 1. The efficient payment to trader i will depend on both the types of all traders encountered so far and the probability distribution over types of the remaining traders. However, from Lemma 1 we know that the number of patient traders encountered so far, θ_{i-1}, is sufficient to determine this probability distribution. We can, therefore, determine this payment as a function of y_{i-1} and θ_{i-1} alone; let c_i^0 denote the payment.[10]

The proposition below presents the efficient payments c_i^0. The proof in the appendix consists of converting (8) into a dynamic programming problem and solving it backward. Presenting the solution requires one additional piece of notation: let $\pi_i(\theta)$ denote the probability of $\omega_i = 0$ conditional on θ of the first $i-1$ traders being patient.[11] We then have the following result.

Proposition 1 *The efficient allocation when all traders contact the intermediary in period* 0 *sets*

$$c_i^0 = \frac{y_{i-1}}{\psi_i(\theta_{i-1})^{\frac{1}{\gamma}} + 1} \quad \textit{for } i = 1, \ldots, I,$$

where $y_{i-1} = I - \sum_{j<i} c_j^0$ *and the functions* ψ_i *are defined recursively by* $\psi_I(x) = \left(xR^{\frac{1-\gamma}{\gamma}}\right)^{\gamma}$ *and*

$$\psi_i(x) = \pi_{i+1}(x)\left(\psi_{i+1}(x)^{\frac{1}{\gamma}} + 1\right)^{\gamma} + (1 - \pi_{i+1}(x))\,\psi_{i+1}(x+1) \tag{11}$$

for $i = 1, \ldots, I-1$.

A proof of the proposition is given in Appendix A. Note that equation (11) depends only on the conditional probabilities π_i and the parameters R and γ. This equation can, therefore, be used

[10] A comment on notation: The variable c_i^0 here denotes the payment given to depositor i at date 0 if she is impatient *conditional on* y_{i-1} and θ_{i-1}. Once we solve the full dynamic programming problem, we will be able to use this variable to calculate the payment as a function only of the partial history, denoted above by $\mathbf{c}_i^0(\omega^i)$.

[11] That the probability π_i depends only on θ follows from Lemma 1.

recursively to determine $\psi_i(\theta_{i-1})$ for any values of i and θ_{i-1}. The functions ψ_i then determine the payment c_i^0 to an impatient depositor following any partial history ω^i.

Example. Figure 1 depicts the efficient allocation for an example with 5 traders. Types are independent, with each trader having probability $1/2$ of being impatient; the other parameter values are given by $R = 1.1$ and $\gamma = 6$. The figure shows the possible period 0 consumption levels of each trader. The black dots correspond to partial histories in which trader 1 is impatient, while the red diamonds correspond to histories in which trader 1 is patient.

The level of consumption trader 1 receives if she is impatient is given by the first black dot in the figure. For trader 2, the consumption she receives in period 0 if she is impatient depends on the type of trader 1. If trader 1 was impatient, then the payment to trader 2 will be smaller (the black dot), while if trader 1 was patient the payment to trader 2 will be larger (the red diamond). For trader 3, there are four different possible consumption levels if she is impatient, depending on the types of the first two traders. The figure shows that trader 3's consumption is slightly higher following the partial history $\omega^2 = (0,1)$ than following $\omega^2 = (1,0)$. In general, trader i faces 2^{i-1} possible consumption levels, each corresponding to a particular realization of the types of the previous traders.

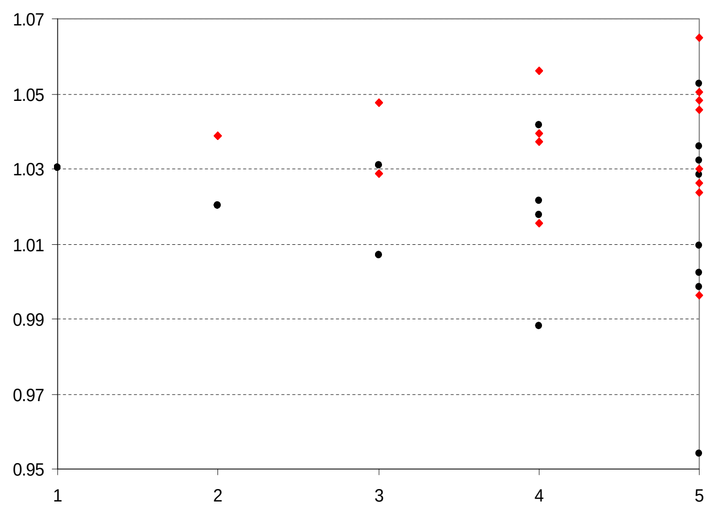

Figure 1: Efficient allocation when $\delta_2 = 0$

2.3 The efficient allocation when δ_2 is large

We now investigate how the efficient allocation changes when contacting the intermediary in both periods is costly. The efficient allocation is still the solution to the maximization problem in (8). However, as discussed above, a positive value of δ_2 may change the efficient pattern of traders contacting the intermediary. We assume δ_2 is large enough that it is inefficient for any trader to contact the intermediary twice.[12] It then follows immediately that each trader should contact the intermediary in period 0 if and only if she is impatient. As a result, the sequential service constraint reduces to (7).

As a first step in solving this problem, note that (7) implies that the efficient allocation can be summarized in a particularly simple way. First, we know that the efficient allocation will again satisfy (9) and (10), which state that only impatient traders consume in period 0 and that the resources remaining in period 1 will be divided evenly among the patient traders. Therefore, we only need to determine the efficient payment for each impatient trader to receive in period 0. Second, condition (7) implies that the first impatient trader must receive the same level of consumption regardless of her position i in the order; let x_1 denote this amount. Similarly, let x_n denote the amount of consumption received by the n^{th} impatient trader, which again must be the same regardless of her position in the order and, thus, the number of patient traders before her.[13] Then the period-0 payment to an impatient trader is

$$\mathbf{c}_i^0(\omega) = x_{i-\theta_i(\omega^i)} \quad \text{if } \omega_i = 0$$

and, thus, the sequence of numbers x_n completely summarizes the allocation $\mathbf{a}$.

To solve for this efficient schedule x, we formulate a dynamic programming problem similar to

[12] It is straightforward to show that such a level of δ_2 exists. It would also be interesting to study positive but small values of δ_2, in which case the efficient pattern of traders contacting the intermediary is potentially more complex. We leave this issue for future research.

[13] Notice that we assume the intermediary has no information about the number of patient traders who may have preceded the n^{th} impatient trader in the order. Adding some information along these lines would be an interesting extension. Suppose, for example, that the intermediary observes the "time" within period 0 at which each impatient trader arrives. If traders' decision opportunities were arranged deterministically in time (say, one trader per minute), then the intermediary could perfectly infer how many patient traders have passed their contact opportunity at any point in time; the analysis would then be isomorphic to the case of $\delta_2 = 0$. If, however, decision opportunities occur randomly in time, this inference would be imperfect and the type of informational friction we study here would arise. The present approach simplifies the analysis considerably and can be regarded as a useful benchmark for understanding more general information structures.

the one in Section 2.2. Define the following conditional probabilities:

$$q_n = \text{Prob}\left[I - \theta(\omega) \geq n \mid I - \theta(\omega) \geq n - 1\right].$$

After the intermediary has encountered $n-1$ impatient traders in period 0, q_n is the probability that it will meet at least one more. These conditional probabilities are easily computed for any distribution of types in the population using (3). The proposition below derives the efficient payment schedule x_n as a function of these probabilities.

Proposition 2 *The efficient payment schedule when only impatient traders contact the intermediary in period* 0 *sets*

$$x_n = \frac{z_{n-1}}{(\phi_n)^{\frac{1}{\gamma}} + 1} \quad \textit{for } n = 1, \ldots, I,$$

where $z_{n-1} = I - \sum_{j<n} x_j$ *and the constants* ϕ_n *are defined recursively by* $\phi_I = 0$ *and*

$$\phi_n = q_{n+1}\left({\phi_{n+1}}^{\frac{1}{\gamma}} + 1\right)^{\gamma} + (1 - q_{n+1})(I - n)^{\gamma} R^{1-\gamma}$$

for $n = 1, \ldots, I - 1$.

The variable z_{n-1} measures the amount of resources remaining when the intermediary encounters the n^{th} impatient depositor. The proposition shows that the fraction of the remaining resources this depositor will receive depends on the remaining conditional probabilities q_{n+1}, q_{n+2}, etc., as well as on the parameters R and γ. A proof of the proposition is given in Appendix A.

Example. Figure 2 plots the efficient payment schedule x when there are 20 traders and the parameter values are given by $R = 1.1$ and $\gamma = 6$, and types are independent with the probability of being impatient set to $\pi = 0.5$ for each trader. The lower curve in the figure presents, for each value of n, the consumption that the n^{th} impatient trader will receive in period 0. While this curve is strictly decreasing, it is initially close to being flat. In other words, the period 0 payment schedule resembles a *demand deposit contract* in which, initially, agents withdrawing funds from the intermediary receive (approximately) the same amount. Once the total number of early withdrawals exceeds a threshold, however, the intermediary starts to decrease the payment. This latter part of the curve resembles a "partial suspension of convertibility" (see Wallace [11]).

The upper curve in the figure represents the level of consumption that all patient traders will

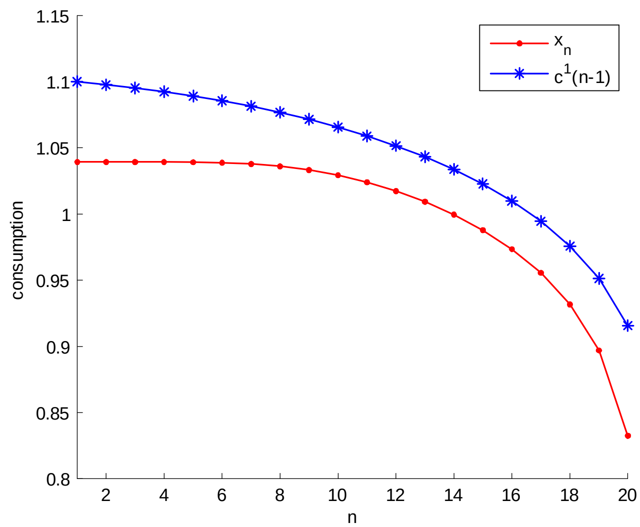

Figure 2: Efficient allocation under costly communication

receive in period 1 if there is a total of $n-1$ impatient traders. The fact that this latter curve lies everywhere above the former has the following interpretation. Consider the last trader in the order, trader I. Let the number of impatient traders before her be given by $n-1$. If she is impatient she will receive the consumption allocated for the n^{th} impatient trader, from the lower curve in the figure, while if she is patient she will receive the consumption allocated for patient traders when there are a total of $n-1$ impatient traders, which is the corresponding point on the upper curve. Thus the figure shows that the last trader always consumes more when she is patient than when she is impatient, regardless of the types of the other traders. Notice that this feature does not necessarily hold for other traders. Trader 1, for example, consumes more if she is patient when the total number of patient traders turns out to be small enough (fewer than 14 in the example). However, if sufficiently many of the other traders are impatient, trader 1 will end up consuming *less* if she is patient than if she is impatient.

Comparing Figures 1 and 2 shows that the efficient allocation can be expressed more simply when δ_2 is large. In this case, the intermediary collects less information in period 0 because receiving information from traders with no immediate need to consume is costly. As a result, the payments made to impatient traders are conditional on less information, leading the allocation to take a simpler form.

3 Implementation

Propositions 1 and 2 derive the efficient way to allocate resources as a function of traders' types. We now turn to the study of mechanisms designed to implement this efficient allocation in the presence of private information. We study direct revelation mechanisms, where traders are asked to report their own types. We ask whether the resulting game has an equilibrium where traders "run" on the intermediary by mis-reporting their types, and we review the answers to this question given by Green and Lin [6] and Peck and Shell [9] in the context of our model. In the following two sections, 4 and 5, we present our own answers to this question based on the effects of correlated types and costly communication, respectively.

3.1 Mechanisms and equilibrium

We study mechanisms in which each trader is asked to submit a message m_i from some set M. Let $m = (m_1, \ldots, m_I)$ denote a profile of messages. Trader i's *communication strategy* is an $\mathcal{E}_i$-measurable function $\mu_i : \Omega^I \rightarrow M$. A profile of communication strategies is $\mu(\omega) = (\mu_1(\omega), \ldots, \mu_I(\omega))$. We use μ_{-i} to denote the profile of strategies for all traders except i.

An *allocation rule* is a function α that assigns a feasible (ex post) allocation to any profile of messages m.[14] Let Γ denote the set of such rules, i.e.,

$$\Gamma = \left\{\alpha : M^I \rightarrow \mathbb{A}\right\}.$$

Given any allocation rule α and any profile of communication strategies μ, we can generate a state-contingent allocation by $\mathbf{a} = \alpha \circ \mu$, or, for each state ω,

$$\mathbf{a}(\omega) = \alpha(\mu(\omega)).$$

In other words, an allocation rule and a profile of communication strategies together create a mapping from states to feasible (ex post) allocations. We say that the allocation rule α *respects* the isolation and sequential service constraints if the corresponding state-contingent allocation $\mathbf{a}$ satisfies (4) and (5) for every profile of communication strategies μ. Let Γ' denote the set of feasible allocation rules that respect isolation and sequential service.

In general, an allocation *mechanism* specifies both a message space and an allocation rule

[14] Green and Lin [6] allow α to depend on the true state ω as well as the message profile m. However, since the planner observes nothing about ω directly, there is no loss of generality in having α depend only on m.

(M, α). Following the literature, we consider direct mechanisms in which each trader is asked only to report her type, so that $M = \Omega = \{0, 1\}$. We can then refer to the allocation mechanism as simply being the rule α. We require $\alpha \in \Gamma'$.

After a mechanism α is chosen, traders play the resulting direct revelation game. A *Bayesian Nash Equilibrium* of this game is a communication-strategy profile μ^* such that, for all i and for all μ_i, we have[15]

$$U_i \left(\alpha \circ \left(\mu_{-i}^*, \mu_i \right), \omega \right) \leq U_i \left(\alpha \circ \left(\mu_{-i}^*, \mu_i^* \right), \omega \right) \quad \text{for all } \omega.$$

We say that an allocation is *implementable* if it is the outcome of a Bayesian Nash equilibrium of this game under some mechanism. In other words, $\mathbf{a}$ is implementable if there exists a mechanism α and an equilibrium strategy profile μ^* of the direct revelation game generated by α such that

$$\mathbf{a}(\omega) = \alpha \left(\mu^* (\omega) \right) \quad \text{for all } \omega. \tag{12}$$

An allocation is *truthfully implementable*, or (Bayesian) *incentive compatible*, if it can be implemented in an equilibrium where all traders report truthfully, that is, where $\mu_i^* = \omega_i$ for all i. The Revelation Principle tell us that an allocation is implementable if and only if it is incentive compatible.

Green and Lin [6] showed that when $\delta_2 = 0$ and types are independent, the efficient allocation is always incentive compatible. The same is true in our examples in the sections that follow. In other words, in all of these cases, the efficient allocation can be implemented by following a simple rule: treat all messages as truthful and assign allocations according to the general solution to (8) derived above. In what follows, we focus exclusively on this allocation rule, which we denote α^*.

Before moving on, we point out that some strategies in the direct revelation game generated by α^* are strictly dominated and, hence, cannot be part of any equilibrium. In particular, condition (9) states that any trader reporting to be patient will be given zero consumption at date 0. Furthermore, it is straightforward to show that all traders reporting to be impatient will receive positive consumption at date 0. Since impatient traders only care about consumption at date 0, lying when a trader is impatient is a strictly dominated strategy. For the analysis of equilibrium, therefore, we

15 A comment on notation: The requirement "for all ω" in this expression might seem strange, since a depositor does not know ω. Recall, however, that the function U_i takes on only two values, one for $\omega_i = 0$ and another for $\omega_i = 1$. Our notation follows Green and Lin [6].

only need to examine the action of a trader in the event that she is patient.

3.2 A unique implementation result (Green-Lin)

While incentive compatibility of the efficient allocation guarantees that it is *an* equilibrium of the direct revelation game under α^*, it may not be the *only* equilibrium. Our primary interest is in the possibility that there also exist "run" equilibria in which some traders mis-report their types in some states. The nature of the exercise we perform in this paper is the same as that in Diamond and Dybvig [4] and others. Suppose the intermediary tries to implement the efficient allocation using the rule α^*. Is there a run equilibrium of the resulting game?

When there are no reporting costs (*i.e.*, $\delta_2 = 0$) and types are independent, our model reduces to exactly that studied by Green and Lin [6]. They showed that, under the efficient allocation rule α^*, the direct revelation game has a *unique* equilibrium. In that equilibrium, all traders truthfully report their types; no one runs on the intermediary.

Proposition 3 *(Green and Lin* [6]*) If $\delta_2 = 0$ and types are independent, the direct revelation game associated with α^* has a unique Bayesian Nash equilibrium and the efficient allocation $\mathbf{a}^*$ obtains in that equilibrium.*

This remarkable result demonstrates that the basic elements of the Diamond-Dybvig framework – isolation, private information, and sequential service – do not necessarily open the door to a run equilibrium. In a particular environment that contains all of these features, an intermediary can, through the proper choice of contract, ensure that the efficient allocation obtains. The results of Diamond and Dybvig [4], and the sizable literature that has followed, thus depend crucially on some unmodelled restriction(s) that prevent an intermediary from following the efficient payment rule characterized in Proposition 1. Recent work by Andolfatto *et al.* [2] has extended Green and Lin's result to a broader class of preferences and has helped clarify the logic behind the arguments, particularly regarding the importance of the assumption that traders' types are independent.

Green and Lin conclude their study by asking "what's missing" from the model that prevents it from being able to generate self-fulfilling bank runs (see also Green and Lin [7]). In the remainder of this paper, we provide three possible answers to this question. The first of these answers, which follows Peck and Shell [9], is presented in the next subsection. In Sections 4 and 5 we present answers based on correlation in types and costly communication, respectively.

3.3 Run equilibria based on early decisions (Peck-Shell)

One way to modify the Green-Lin environment is to assume that traders must choose an action prior to learning the order in which they will contact the intermediary. Places in this order are then assigned at random, with each trader equally likely to occupy each place. This is the approach implicitly taken in the original work of Diamond and Dybvig [4] and in much of the subsequent literature. In this case, a trader's expected utility when choosing a strategy is an average of the utilities associated with each of the I places in the ordering[16]

$$\frac{1}{I}\sum_{i\in\mathbb{I}} E\left[U_i\left(\mathbf{a},\omega\right)\right]. \tag{13}$$

Note that this expression is equivalent to (8), the objective function of the intermediary.

We now show that a run equilibrium can exist in this modified environment. Our examples are very much in the spirit of Peck and Shell [9], who first showed that a run equilibrium can exist when no restrictions other than sequential service are placed on the intermediary's allocation rule. However, the preferences used in Peck and Shell [9] are not of the form in (2); rather, in their setting the marginal utility of consumption is higher for impatient traders than for patient traders. This approach simplifies the computations in their model by ensuring that an incentive compatibility constraint binds at the efficient allocation. Our example shows that differing marginal utilities are not necessary for this result to obtain. Everything in our examples below is exactly as in the Green-Lin model except the information that traders have when choosing an action. In particular, Proposition 1 still characterizes the efficient allocation in this setting.

The following proposition summarizes our results for this first modification of the Green-Lin model.

Proposition 4 *Suppose types are independent and* $\delta_2 = 0$. *When traders must choose a strategy before knowing their position in the order,* (i) *the efficient allocation* $\mathbf{a}^*$ *is incentive compatible, but* (ii) *for some parameter values the direct revelation game also has a run equilibrium.*

The proof of the first part of the proposition follows from Green and Lin [6], who showed that the

[16] A proper formulation of this case would introduce new notation to distinguish between a trader's index (or "name") and his eventual place in the ordering. (See Green and Lin [6] for such a formulation.) Doing so, however, complicates the presentation considerably. Since this issue only arises in the present subsection, we take the notational shortcut of having traders act before any "names" are assigned. This shortcut is purely a matter of notation; it does not change the underlying analysis in any way.

efficient allocation is incentive compatible when traders know their position in the order. In other words, once traders are assigned positions in the order, each of them will prefer to report truthfully if all others are doing so. It follows immediately that a trader who does not yet know her position in the order would make the same choice, since it will be a best response whatever position she is assigned. The proof of the second part of the proposition is by example.

Example. There are 15 traders. Types are independent, with each trader having probability 0.1 of being impatient; the other parameter values are given by $R = 1.1$ and $\gamma = 6$. We first calculate the efficient allocation $\mathbf{a}^*$ using Proposition 1. We then ask the following question. Suppose a trader believes that all others will run, that is, claim to be impatient regardless of their true types. Would this trader prefer to run as well or, if patient, would she prefer to wait and consume in period 1?

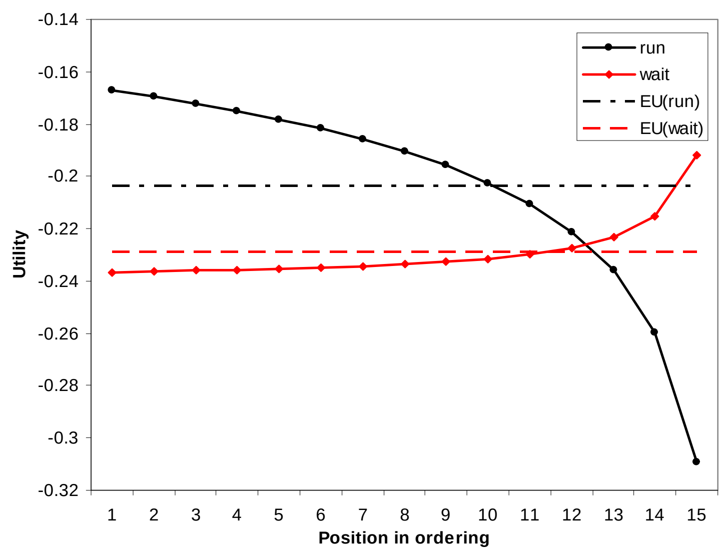

Figure 3: Expected utility if all other traders run

Figure 3 plots the utility associated with each of these actions for each possible position in the order, conditional on the trader in question being patient. The solid black line represents the utility from running, which is strictly decreasing in the trader's position in the order. The solid red line represents the utility of reporting truthfully and waiting until period 1 to consume. The figure shows that if the trader knew she would be among the first 12 traders to contact the intermediary, then, given the belief that all other traders will run, she would strictly prefer to run. However, if she knew she would be among the last three traders in the order, she would prefer to report truthfully and consume in period 1 if patient.

The dashed lines in the figure represent the expected value of each of the actions, given that a trader is equally likely to end up in each of the two positions. The figure demonstrates that, for this example, the trader strictly prefers to run. Therefore, an equilibrium exists in which all traders claim to be impatient and consume early; this outcome resembles a classic run on the intermediary.

The parameter values used in the above example are in no way special. It is easy to find other combinations that also generate a run equilibrium. This fact is demonstrated in Figure 4, which plots the gain in expected utility from following the run strategy (relative to reporting truthfully) for a trader who believes that all other traders will run. The run equilibrium exists if and only if this number is positive. The parameter values from the above example are represented by the solid red line. As the figure shows, the run equilibrium exists for these values whenever the number of traders is at least six. If the probability of impatience is increased to 0.5, the run equilibrium no longer exists when there are six traders. However, the dashed black line in the figure shows that it will exist if there are at least nine traders. These calculations demonstrate both that there is nothing special about the parameter values used for the example above and that increasing the number of traders makes it more likely that a run equilibrium will exist in this setting.

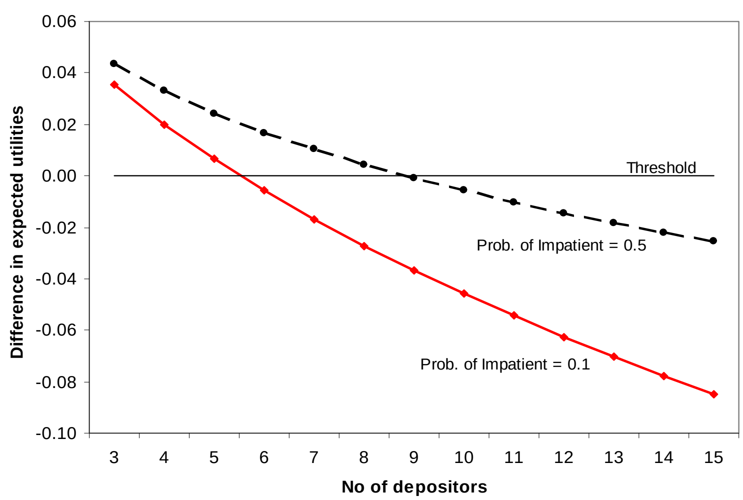

Figure 4: Incentive to run as I varies

Going back to Figure 3, notice an interesting feature in this graph: the last three traders to contact the intermediary in period 0 would actually be better off waiting until period 1, even though all other traders are claiming to be impatient. One can show that this reflects a general feature of the efficient allocation: once positions in the order are realized, traders I and $I-1$ are *always* made strictly better off by reporting truthfully.[17] Suppose, then, that a trader were somehow able to re-evaluate her decision once she arrives at the intermediary and discovers her position in the order. A patient trader who finds she is the last to arrive would prefer to report truthfully and wait until period 1 to consume.[18] Other traders should recognize this fact and adjust their forecasts of others' behavior accordingly.

It was precisely to capture these types of effects that Green and Lin [6] introduced the possibility that a trader's action could also depend on her position in the order. We focus on this case for the remainder of the paper. An immediate implication of the approach is that there cannot be an equilibrium in which all traders claim to be impatient. Any run on the intermediary would have to be *partial*, with only some traders participating. In the next two sections, we present modifications to the Green-Lin environment that generate such partial-run equilibria.

4 Correlated Types

In this section, we study the case where the realization of types is correlated across traders. All other features of the environment are exactly as in Green and Lin [6]; in particular, there are no costs of contacting the intermediary (*i.e.*, $\delta_2 = 0$) and traders always know the order in which this contact will occur. We show that a run equilibrium can exist in this setting; specifically, we establish the following result.

Proposition 5 *Suppose $\delta_2 = 0$. When types are correlated, for some parameter values it is the case that (i) the efficient allocation $\mathbf{a}^*$ is incentive compatible and (ii) the direct revelation game also has a run equilibrium.*

The proof is by example. We begin by presenting the simplest example in which a run equilibrium can arise. We describe the intuition behind this example in some detail; this discussion will also be useful for understanding the case of costly communication presented in Section 5. We then show that the result does not depend on the particular features of the simple example by constructing a richer example where a run equilibrium exists and is driven by the same underlying intuition.

[17] This result is stated as Lemma 2 in Appendix A and a proof is offered there.

[18] Recall that in this section δ_2 is assumed to be zero, so that there is no cost to the trader of contacting the intermediary again in period 1.

4.1 A basic example

To keep the presentation as simple as possible, we use the minimal number of traders needed to generate a run equilibrium. As mentioned above (in subsection 3.3), truth-telling is a dominant strategy for the last two traders in the order. Hence, a run equilibrium clearly cannot exist when $I = 2$. It is fairly easy to see that one cannot exist when $I = 3$, either, as long as the efficient allocation is strictly incentive compatible; the first trader knows that the last two will report truthfully and, therefore, incentive compatibility implies that she will prefer to do the same. Hence, the minimum number of traders needed to construct a run equilibrium is four, and we use $I = 4$ in our example.

Parameter Values. The number of patient traders is very likely to take on one particular value in this example; in this sense, there is little aggregate uncertainty. Specifically, we set

$$p(2) = 1 - \varepsilon, \quad \text{and} \tag{14}$$

$$p(\theta) = \frac{\varepsilon}{4} \qquad \text{for } n = 0, 1, 3, 4,$$

where $p(\theta)$ is, as defined in (3), the probability of the set of states in which exactly θ traders are patient. We choose ε to be small (we use $\varepsilon = 0.4\%$). We set the other parameter values to $R = 2$ and $\gamma = 6$.

The Efficient Allocation. The efficient allocation is calculated using Proposition 1 and is depicted in Figure 5. While the allocation has the same general structure as in the case of independent types (see Figure 1), the nature of the correlation in this example simplifies the pattern of payments and makes developing intuition fairly easy. Suppose for a moment that ε were zero, so that there is no aggregate uncertainty; two traders would be impatient and two patient with certainty. The efficient allocation would then be as in the textbook version of the Diamond-Dybvig model: a common payment $\bar{c}^0 > 1$ will be given to the impatient traders and a larger payment $\bar{c}^1$ to the patient traders, regardless of their order of arrival at the intermediary.

When ε is positive, the efficient allocation is more complex than the simple Diamond-Dybvig allocation, but as long as ε is small the allocations will, in broad terms, be similar. In particular, Figure 5 shows that the intermediary will give relatively large payments (greater than 1) to the first two impatient traders it encounters. If there are exactly two impatient traders, both will receive consumption very close to the Diamond-Dybvig level of $\bar{c}^0$ (around 1.3 in the figure). If the in-

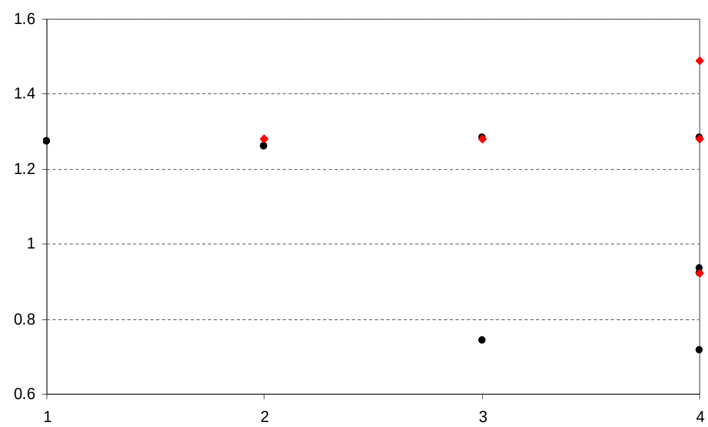

Figure 5: Efficient allocation with correlated types

termediary encounters a third impatient trader, however, it realizes that one of the low-probability states has occurred and will adjust payments accordingly. In such cases, the payment to the third impatient trader – and the fourth, if there is one – will be much lower, as will the payment to any patient trader. This decrease in payment size reflects the fact that the first two payments made (both close to $\bar{c}^0$) were based on a belief about the state ω that turned out to be very "optimistic" relative to the realization. Wallace [11] refers to this pattern where impatient traders who contact the intermediary late in the order receive less than those who arrived earlier as a *partial suspension of convertibility*.[19]

Incentive Compatibility. Now suppose the intermediary attempts to implement this efficient allocation using a direct revelation mechanism. We first check whether this can be done; in other words, is the efficient allocation incentive compatible? We know that a trader always strictly prefers to report truthfully when impatient. Therefore, we only need to compare the expected utility of reporting truthfully with that of following the run strategy, which sets $\mu_i = 0$ regardless of the trader's true type. The comparison is presented in Figure 6. The dashed black line in the figure depicts the gain in expected utility from choosing the run strategy (relative to reporting truthfully) for each trader under the assumption that all others are reporting truthfully. The fact that the line is negative everywhere indicates that all traders derive higher utility from reporting truthfully; hence, there exists a truth-telling equilibrium that implements the efficient allocation in this example.

[19] If the first three traders are *patient*, the intermediary will again realize that a low-probability state has occurred. In this case, if trader 4 is impatient she will receive a higher-than-usual level of consumption (about 1.5 in the figure).

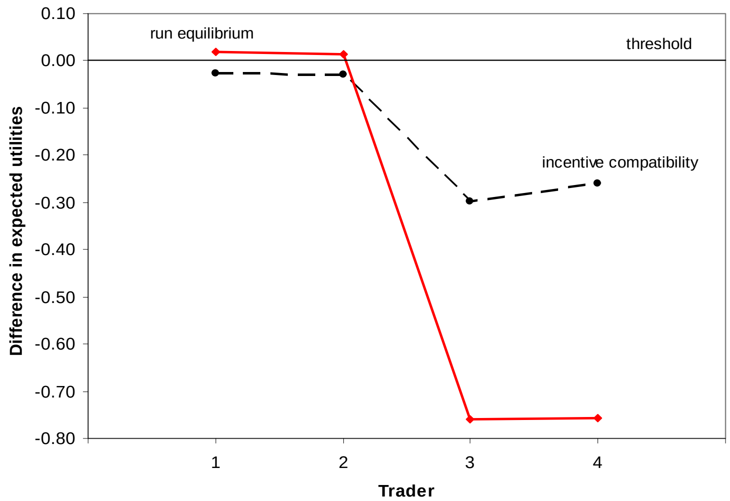

Figure 6: Individual incentive to run with correlated types

A Run Equilibrium. Next, we construct another equilibrium of the direct revelation game. In this equilibrium, the first two traders follow the run strategy while the last two traders report truthfully. The equilibrium communication strategies are, therefore, given by

$$\mu_i(\omega) = \left\{ \begin{array}{ccc} 0 & \text{for} & i = 1,2 \\ \omega_i & & i = 3,4 \end{array} \right\}. \tag{15}$$

Lemma 2 (in Appendix A) tells us that the strategies in (15) are optimal for traders 3 and 4. To verify that this strategy profile is indeed an equilibrium, therefore, we only need to show that, taking the strategies of others as given, both traders 1 and 2 will prefer to mis-report when they are patient.

Consider first the decision of trader 2. If she is patient, she knows it is very likely that exactly two of the other traders are impatient. The important question, from her point of view, is whether or not trader 1 is one of them. It is possible that trader 1 was indeed impatient, so that his report of $m_1 = 0$ was actually truthful. In this case, only one of the remaining traders (3 and 4) is likely to be impatient. If trader 2 reports truthfully, there are likely to be only two payments made at date 0 and, therefore, her payment at date 1 will be relatively large (close to the $\bar{c}^1$ of the Diamond-Dybvig model).[20] Reporting truthfully would then be the best choice.

If, on the other hand, trader 1 is patient (and, hence, his report was untruthful), then it is very

[20] Of course, it is possible that more than two traders will be impatient, but this risk is of order ε and thus is relatively unimportant in this example.

likely that both traders 3 and 4 will be impatient. In this case, if trader 2 reports truthfully, there will likely be three payments made at date 0 and the amount left for her at date 1 will be substantially smaller. If she lies, on the other hand, her report of impatient will be only the second one received by the intermediary and she will receive a larger payment at date 0 (similar to the $\bar{c}^0$ of Diamond-Dybvig). In this case, lying would be the best response.

Given her beliefs about the likelihood of each of these two cases (which are based on the probability distribution $\mathcal{P}$ updated to include her own private information), trader 2 must decide how to report. The solid red line in Figure 6 presents the gain in expected utility from choosing the run strategy under the assumption that other traders are following the strategy profile in (15). The fact that this line is positive at trader 2 shows that, in this example, behaving in accordance with (15) – mis-reporting her type when patient – is the optimal choice for trader 2.

The decision problem faced by trader 1 is similar. From (15) he knows that trader 2 will report impatient, but he does not know whether or not this report will be truthful. If trader 2 is actually patient, then it is likely that both traders 3 and 4 are impatient. In this case, if he reports truthfully there would likely be three payments made at date 0 and, therefore, a relatively small amount left for him at date 1. If trader 2 is actually impatient, however, trader 1 would be better off by reporting truthfully, as we would likely receive a large payment at date 1 (similar to the $\bar{c}^1$ of Diamond-Dybvig). The fact that the solid red line in Figure 6 is positive at trader 1 shows that he will also strictly prefer to follow (15) if all other traders are doing so.

Finally, note that the solid red line is negative for traders 3 and 4, which confirms that they prefer to report truthfully even when traders 1 and 2 follow the run strategy. The figure thus shows that the strategy profile in (15) is indeed an equilibrium for the chosen parameter values. This demonstrates that a run equilibrium can exist in the Green-Lin model when types are correlated across traders.[21]

4.2 Intuition

It is interesting to examine the behavior of trader 2 in this example. She chooses to run even though she believes that both of the traders after her will, following the strategy profile in (15),

[21] Of course, if the intermediary anticipates that traders will follow the strategy profile in (15), it might not want to use the allocation rule α^*. Peck and Shell [9] show, in a related model, how the existence of a run equilibrium in the direct revelation game (as shown here) implies the existence of a run equilibrium in an "overall" game where the intermediary anticipates traders' actions. It would be straightforward to apply the same approach here. See also Ennis and Keister [5].

report truthfully. This behavior is somewhat surprising in light of the results in Green and Lin [6, Lemma 5], which showed that it cannot arise in the model with independent types. In particular, Green and Lin demonstrated that, in their model, when a trader believes everyone after her will report truthfully, she strictly prefers to report truthfully regardless of the actions of the traders who contacted the intermediary before her. The behavior of trader 2 in this example is, therefore, critical to understanding why the Green-Lin unique implementation result does not extend to the case of correlated types. The key to understanding this behavior, in turn, is to compare the equilibrium beliefs of trader 2 with the beliefs used to calculate the payoffs for the direct revelation game.

The direct revelation game is designed to implement the efficient allocation $\mathbf{a}^*$ if all traders report truthfully. When the first two traders both report to be impatient, the payment offered to trader 2 is based on the belief, derived from the probabilities in (14), that traders 3 and 4 are very likely to both be patient. As a result of this belief, trader 2 is offered a relatively large payment – close to the $\bar{c}^0$ of the Diamond-Dybvig model. In a sense, when both traders 1 and 2 report to be impatient, the intermediary is "optimistic" that the early withdrawals will end there, and the payment offered to trader 2 reflects this optimism.

In the run equilibrium, trader 2's belief about the types of traders 3 and 4 is significantly different from that described above. Suppose trader 2 is patient. She recognizes that trader 1 will report to be impatient regardless of his true type. She thus recognizes that, following a withdrawal by trader 1, there is a significant chance that both traders 3 and 4 will be impatient. Relative to the belief used to design the allocation mechanism, trader 2 is more "pessimistic" about the number of additional early withdrawals.

This pessimistic belief makes waiting until period 1 less attractive for trader 2. She knows that, if she reports truthfully and waits to consume, there will almost certainly be at least one more early withdrawal. Moreover, she believes there is a significant chance that traders 3 and 4 will both be impatient, in which case the intermediary will face a third early withdrawal. Because the intermediary considered three early withdrawals to be unlikely *ex ante*, this event is associated with a substantial decrease in the payments to all traders who have yet to consume. Trader 2's pessimistic belief about the number of early withdrawals thus makes her less willing to report truthfully and, hence, makes running – and claiming the period 0 payment based on optimistic belief – more attractive. Trader 1 faces similar incentives when he believes that trader 2 will run.

Notice that when types are independent, this disparity in beliefs cannot arise. In that case, trader

2's knowledge of her own type and the equilibrium strategy of Trader 1 do not provide her with any additional information about the likely types of traders 3 and 4. Her belief about these types remains identical to the belief used to design the efficient payment schedule – each of these traders has an independent probability π of being impatient. The payments offered by the intermediary are, therefore, "appropriate" given trader 2's belief and, as shown by Green and Lin [6], lead trader 2 to strictly prefer truthful reporting.[22]

4.3 Another example

The example presented above is, in some ways, rather special: there are only four traders and there is almost no aggregate uncertainty about the number of impatient traders. While these features were useful for generating intuition about why a run equilibrium can exist when types are correlated, they are by no means necessary for the result to obtain. We demonstrate this fact by presenting an example in which the same results obtain with 10 traders and a significant amount of aggregate uncertainty. Many other examples with similar features are possible, of course.

Parameter Values. For this example, we use $I = 10$. The parameter values $R = 2$ and $\gamma = 6$ are unchanged from the simple example above. We set the density function p for the number of patient traders as follows

$$p(\theta) = \frac{1-\varepsilon}{5} \text{ for } \theta = 3, \ldots 7, \text{ and} \tag{16}$$

$$p(\theta) = \frac{\varepsilon}{6} \text{ for } \theta = 0, 1, 2, 8, 9, 10.$$

We again choose ε to be very small (we set $\varepsilon = 0.006\%$ for this example). In other words, this example is designed so that the number of patient traders is very likely to fall somewhere between 3 and 7 out of the 10 total traders. We have made each of these possibilities equally likely just for simplicity. The important feature of the specification here is that it is very unlikely that almost all of the traders will be impatient; in particular, such events are substantially less likely than in the case of independent types.

Incentive Compatibility. The efficient allocation is again calculated as in Proposition 1. With the larger number of traders in this example, the structure of the efficient allocation is much more complex than before; however, using the proposition its calculation remains straightforward. We

[22] See also the discussion in Andolfatto *et al.* [2] regarding the importance of the independence assumption in deriving Green and Lin's main result.

first ask if this allocation is incentive compatible. To do so, we again compare the gain in expected utility from choosing the run strategy for each trader (relative to reporting truthfully) under the assumption that all other traders report truthfully. This gain is plotted as the dashed black line in Figure 7. The fact that the line is negative for all traders indicates that the efficient allocation is indeed incentive compatible in this example.

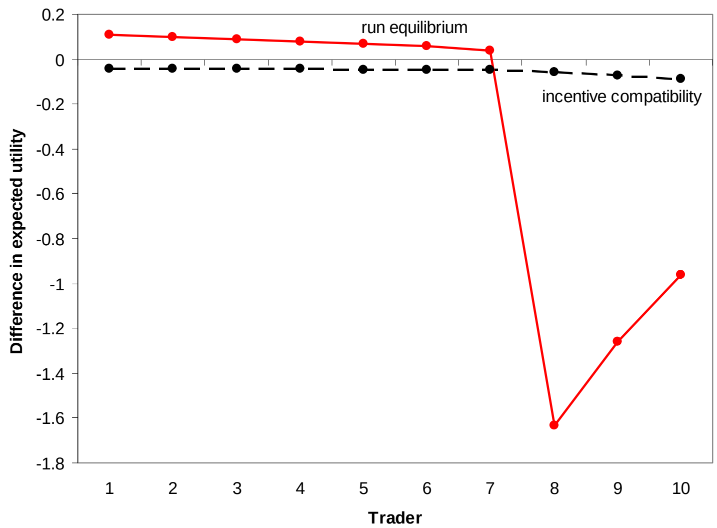

Figure 7: Incentive to run in the richer example

A Run Equilibrium. Next, we construct a partial run equilibrium for this example. The basic form of this equilibrium is the same as in the simple example above: traders who are early in the order choose to run, while those who are later in the order report truthfully. Specifically, we propose the following strategy profile as a potential equilibrium

$$\mu_i(\omega) = \left\{ \begin{array}{ccc} 0 & \text{for} & i = 1, \ldots, 7 \\ \omega_i & & i = 8,\ 9,\ 10 \end{array} \right\}. \tag{17}$$

The solid red line in Figure 7 plots the expected gain from following the run strategy (relative to reporting truthfully) when all other traders are expected to follow the strategies in (17). The figure shows that this gain is positive for the first seven traders and negative for the last three. In other words, if each trader believes that all others will follow the strategies in (17), she strictly prefers to do so as well. As a result, the strategy profile (17) is an equilibrium of the direct revelation game.

Intuition. As before, the key to understanding why a run equilibrium exists is to look at the behavior of the last trader to follow the run strategy. In the present example, trader 7 plays this critical role. Why does she run even though she anticipates that all traders after her will report truthfully?

The intuition behind trader 7's behavior in this example is the same as that described above for trader 2 in the basic example. When the first 7 traders report to be impatient, the payment given to the 7th trader is based on the belief that the remaining three traders are very likely to all be patient. This belief is generated by the probabilities in (16), which state that the number of impatient traders is very unlikely to be greater than 7. As a result, the payment offered to trader 7 in this situation is relatively large, reflecting the "optimistic" belief that additional early withdrawals are unlikely.

In the run equilibrium, however, trader 7 has a very different belief about the types of the remaining traders. She recognizes that the first six traders have reported to be impatient regardless of their true types and, therefore, she believes it is quite likely that two or even all three of the remaining traders will be impatient. In such a case, the intermediary would face an unexpectedly high level of early withdrawals and traders who have reported to be patient would receive relatively low levels of consumption. Under this "pessimistic" belief, therefore, reporting truthfully and waiting until period 1 to consume is substantially less attractive for trader 7. Running – and claiming the period 0 payments based on a more optimistic belief – is more attractive. In the example, her belief leads her to follow the run strategy.

This reasoning demonstrates that the lack of aggregate uncertainty in our simple example with 4 traders is not important for the result; rather, the important feature is that a very high number of impatient traders is very unlikely to be realized. In such a situation, the efficient allocation has the property that an unexpectedly high level of early withdrawal demand will lead to a "crisis" in which any remaining impatient traders, as well as *all* patient traders, receive relatively low levels of consumption. The possibility of such an event gives a trader who is pessimistic about the types of the traders after her in the order an incentive to withdraw early. Together, our examples above show not only that introducing correlation in traders' types can generate a wedge between the equilibrium beliefs of the traders and the beliefs used to design the efficient allocation, but also that this wedge can generate a run equilibrium in the direct revelation game.

5 Costly Communication

We now return to the case where types are independent across traders and study the effects of costly communication between traders and the intermediary. In particular, we now assume that δ_2 is large, so that traders face a significant cost of contacting the intermediary in both periods and the efficient allocation is as given in Proposition 2. While the details of this case are very different from the correlated types case studied in the previous section, the equilibrium effects are strikingly similar. In particular, we show that by slowing down the flow of information to the intermediary, costly communication can also drive a wedge between the equilibrium beliefs of traders and those used to derive the efficient allocation. This wedge can then generate an equilibrium in which some traders run on the intermediary in much the same way it did in the examples above. We should emphasize that, in this section, all aspects of the environment other than δ_2 are exactly as in Green and Lin [6]. Our results for this case are summarized in the following proposition.

Proposition 6 *When $\delta_2 > 0,$ for some parameter values it is the case that (i) the efficient allocation $\mathbf{a}^*$ is incentive compatible and (ii) the direct revelation game also has a run equilibrium. This is true even when types are independent across traders.*

The proof is again by example. We begin the analysis by discussing strategy profiles and the conditions under which a run equilibrium exists. As in the previous section, a run equilibrium will necessarily be partial, with only some traders participating. We then present an example in which such an equilibrium arises. The intuition behind this result is remarkably similar to that for the result in Section 4 above, despite the apparent differences in the form of the efficient allocation and the nature of uncertainty in the two sections. In particular, the costly-communication friction here allows the equilibrium beliefs of traders about the number of additional early withdrawals to deviate from the beliefs used to design the allocation rule. As in the previous section, this disparity in beliefs plays a crucial role in generating an incentive for some traders to run on the intermediary. We discuss this intuition in detail at the end of this section.

5.1 Partial-run equilibria

The efficient allocation for the case where δ_2 is large was presented in Proposition 2. It is straightforward to show that reporting truthfully is a strictly dominant strategy for trader I under the efficient allocation rule, for the same reasons as in Green and Lin [6] and in Section 4 above. If a run equilibrium exists, therefore, it must once again be partial, with only some traders participating in the run. Suppose we look for an equilibrium of the form used in the previous section, where traders who are early in the order choose to run while traders late in the order report truthfully. More precisely, suppose all traders up to some "critical" trader i^* run, while all traders after i^* report truthfully. Define the strategy profile $\widehat{\mu}(i^*)$ by

$$\widehat{\mu}_i(i^*) = \left\{ \begin{array}{c} 0 \\ \omega_i \end{array} \right\} \text{ for } \left\{ \begin{array}{c} i \leq i^* \\ i > i^* \end{array} \right\} \quad \text{for some } 0 < i^* < I. \tag{18}$$

To see if such a profile is consistent with equilibrium for some value of i^*, we define an auxiliary function z. This function measures the expected utility for a patient trader of reporting truthfully if she believes that all traders before her will run, but all traders after her will report truthfully. Specifically, define

$$z(i) = U_i(\alpha^* \circ \widehat{\mu}(i); 1). \tag{19}$$

In other words, $z(i)$ captures the expected utility of the critical trader i^* in (18) if she is patient and she deviates from that strategy profile by reporting truthfully. It is straightforward to show that z is a strictly decreasing function; higher values of i correspond to larger (partial) runs and, hence, they also correspond to lower consumption levels in period 1.

To check if the strategy profile (18) is an equilibrium, we must check if each trader is choosing a best response to the actions of others. Start with the critical trader i^*. she believes that all traders before her will report impatient, so that if she does the same she will receive the payment x_{i^*} in period 0. As described above, $z(i^*)$ measures her expected utility from reporting truthfully if she is patient. she will, therefore, choose to follow (18) and run if

$$v(x_{i^*}) > z(i^*). \tag{20}$$

This condition is also sufficient to guarantee that all traders before i^* will strictly prefer to run; this result follow from the fact that x_n is a decreasing sequence and that all traders reporting to be patient receive the same consumption in period 1.

Next, consider trader $i^* + 1$. The profile in (18) calls for her to report truthfully, which requires

$$v(x_{i^*+1}) < z(i^* + 1). \tag{21}$$

The left-hand side of this expression is the utility this trader would get if she runs. In this case, her report of impatient would follow those of the first i^* traders in the order. If, instead, she reports truthfully, she will receive the expected value of waiting when i^* traders participate in the run; by definition this value is given by $z(i^* + 1)$. Condition (21) thus says that trader $i^* + 1$ prefers to report truthfully if she is patient. Again using the monotonicity of x_n and the fact that all traders reporting to be patient receive the same consumption, the condition is also sufficient to guarantee that all traders after $i^* + 1$ will also prefer to report truthfully. Therefore, the strategy profile in (18) is an equilibrium of the direct revelation game under costly reporting if both (20) and (21) hold for some i^*. These concise conditions simplify the process of finding run equilibria, as we demonstrate in the example below.

5.2 An example

Our example is based on the same parameter values that were used to illustrate the efficient allocation in Section 2.3: $I = 20$, $R = 1.1$, $\gamma = 6$, and independent types with $\pi = 0.5$. Recall that the efficient allocation for this example was presented graphically in Figure 2. Figure 8 plots three curves, two of which are simply the expected utility associated with the consumption levels from the earlier figure. The third curve is the function z defined in (19).

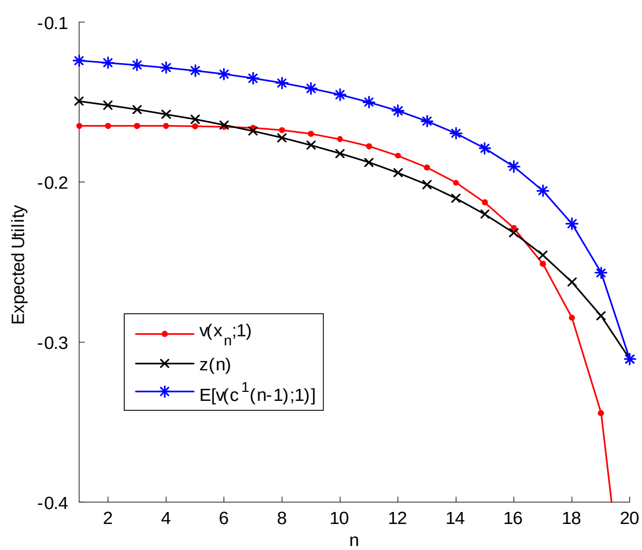

Figure 8: A partial-run equilibrium under costly communication

To understand these curves, first consider whether or not (18) could be an equilibrium strategy profile with $i^* = 1$. For this to be the case, condition (20) would require that the utility from consuming x_1 be less than the expected value of waiting $z(1)$. The figures shows that this is not the case and, hence, that the proposed strategy profile is not an equilibrium. If trader 1 believes that all traders after her will report truthfully, she strictly prefers to report truthfully as well. The figures shows that the same is true of traders 2 through 6, so that none of these could be the critical trader i^* in an equilibrium strategy profile.

Next, consider trader 7, the first trader for whom the z curve lies below the utility value of reporting to be impatient and consuming x_n. Given that the first 6 traders are running, trader 7 would also choose to run. However, the strategy profile in (18) with $i^* = 7$ is not an equilibrium because, as the figure shows, trader 8 would also choose to run given that the first 7 are running. The unique equilibrium of the form given in (18), therefore, occurs where the z curve crosses $v(x_n; 1)$ from below, with $i^* = 16$. Given that the first 15 traders are running, trader 16 will choose to do so as well. Given that the first 16 traders are running, however, trader 17 will choose to report truthfully. Hence, the strategy profile (18) with $i^* = 16$ comprises an equilibrium of the direct revelation game. It is worth emphasizing that there is nothing special about the parameter values used in this example; it is easy to constrict similar examples using a wide range of parameter values.

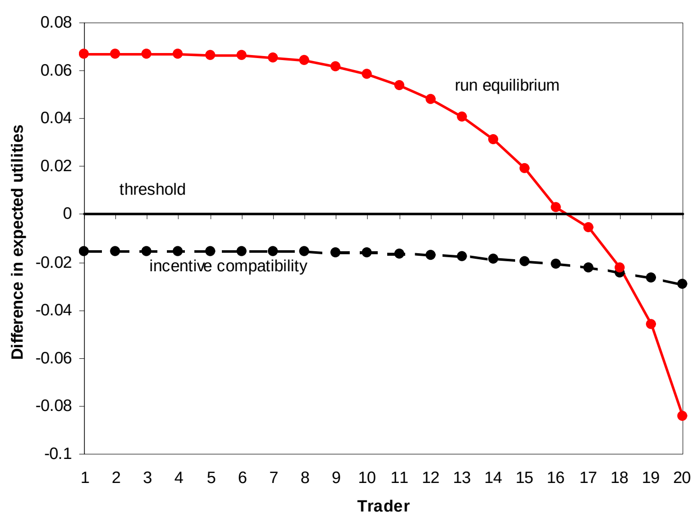

Figure 9: Individual incentive to run under costly communication

Figure 9 confirms that this strategy profile is an equilibrium by plotting the gain in expected utility from following the run strategy (relative to reporting truthfully) for each trader, under the assumption that all other traders follow the strategy profile (18) with $i^* = 16$. The solid red line shows that, given these beliefs, the first 16 traders will indeed prefer to follow the run strategy while the last four will prefer to report truthfully. The figure also verifies that the efficient allocation is incentive compatible in this example. The dashed black line plots the gain in expected utility from following the run strategy for each trader under the assumption that all other traders report truthfully. The fact that this line is negative everywhere demonstrates that truthful reporting is also an equilibrium of the direct revelation game and, hence, the efficient allocation can be implemented.

5.3 Intuition

To gain intuition for why the partial-run strategy profile in (18) is an equilibrium, it is useful to examine the behavior of the "critical" trader in that profile, i^*. This trader follows the run strategy even though she believes everyone after her will report truthfully. As discussed above, this type of behavior is inconsistent with equilibrium in the model of Green and Lin [6]. How does the introduction of costly communication allow this behavior to arise?

The key to understanding the behavior of the critical trader is again to compare this trader's equilibrium belief about the number of early withdrawals with the beliefs used to design the efficient allocation rule. In the example with 20 traders given above, the efficient payment to the 16th impatient trader depends on the probability that the intermediary will encounter a 17th impatient trader (and an 18th, and so on). Given that types are independent and the probability of being impatient is set to one-half, all of these conditional probabilities are close to zero. In other words, the consumption of the 16th impatient trader in the efficient allocation is based on the belief that she is very likely to be the *last* impatient trader and that the four traders who have not contacted the intermediary are very likely all patient. In this sense, when a 16th impatient trader arrives in period 0, the intermediary is "optimistic" that this will be the last early withdrawal, and the payment given to the trader is based on this optimism.

In the partial run equilibrium, however, trader 16 recognizes that the number of early withdrawals will likely be much larger than the number of impatient traders. Importantly, she knows that when she contacts the intermediary in period 0, four traders have not yet had the *opportunity* to contact the intermediary. She expects that, on average, two of these traders will be impatient. In other words, trader 16 realizes that the early withdrawals are unlikely to end with her. Any additional withdrawals will further deplete the intermediary's resources, lowering the consumption she would receive if she were to wait and contact the intermediary in period 1. Her more "pessimistic" belief about the number of additional early withdrawals thus makes running – and accepting the payment based on the more optimistic belief – an attractive strategy.[23]

Notice how the friction of costly communication makes this divergence in beliefs possible. When trader 16 arrives, the intermediary does not know that she is 16th in the order. In fact, the efficient payments are based, implicitly, on the belief that if a 16th impatient trader arrives, she is likely to be the *last* trader in the order. This "misunderstanding" arises because the intermediary does not observe the reports of patient depositors in period 0. In the Green-Lin model, the intermediary would always know that this trader is 16th in the order because it can simply count up all of the reports (both 'patient' and 'impatient') that it has received so far. Both the intermediary and trader 16 thus know that there are four traders who still have an opportunity to contact the intermediary in period 0. Because types are independent, if these traders report truthfully then the intermediary and trader 16 must have the same belief about the number of additional early withdrawals, regardless of what strategies the earlier traders have followed.

Because of this agreement in beliefs, the payment offered to trader 16 in the Green-Lin model will always appear "appropriate" given her beliefs and, as a result, she will choose to report truthfully. This reasoning is central to the unique implementation result in Green and Lin [6]. Under costly communication, in contrast, the divergence in beliefs discussed above arises naturally whenever one or more traders follows a non-truthful strategy. The example presented here shows how this divergence in beliefs can be strong enough to generate a run equilibrium in the direct revelation game.

Notice that the intuition given here is remarkably similar to that discussed in Section 4.2 above. In both cases, there is a critical trader who chooses to run even though she expects everyone after her to report truthfully. The payment this trader receives was designed, in both cases, under the belief that her withdrawal would likely be the last one in period 0 and, as a result, this payment is relatively generous. In the partial-run equilibrium, however, the trader is more pessimistic about

[23] It is easy to see that, given the strategy profile in (18), the incentive for all traders before trader 16 to join in the run are even stronger. The payment trader 1 receives if he withdraws early, for example, is based on the belief that the number of early withdrawals will be, on average, around 10. In equilibrium, however, he anticipates that there will be at least 16 and on average 18 early withdrawals. This large gap in beliefs makes running very attractive from trader 1's point of view. The incentive for traders 2 through 15 to participate in the run lie somewhere in between those of the first and the 16th trader.

the number of additional early withdrawals. In Section 4 this pessimism was generated by the correlation structure in types, while in the present section it resulted from the fact that the intermediary does not observe whether or not some traders early in the order have reported ‘patient’. In both cases, however, the result is the same: the disparity in beliefs gives the trader an incentive to mis-report her type and take the early payment based on the more optimistic belief.

6 Conclusion

Green and Lin [6] derived a surprising result regarding the ability of a financial intermediary to generate an efficient allocation of resources without introducing the type of financial fragility that appeared in the earlier literature. This result led them to ask "what’s missing" from the model that prevents it from being able to explain observed instances of bank runs and related financial crises (see also Green and Lin [7]).[24] We have examined this question systematically and shown how stronger information frictions are needed in order for a run to be possible. We presented two environments, both small deviations from that in Green and Lin [6], in which runs can occur. In one environment agents’ types are correlated, while in the other communicating with the intermediary is costly.

Our results show that, in each of these two cases, the direct revelation game associated with the efficient allocation rule has an equilibrium in which some traders run on the intermediary. A natural next step in the analysis would be to make the intermediary a player in the game and ask how it would react to the possibility that traders might run. The intermediary might, for example, prefer to use another allocation rule under which there is a unique equilibrium. Fortunately this issue has been addressed in related models in the existing literature and the results of such an exercise are now well known: if the probability the intermediary assigns to a run is small enough, it will choose to follow a rule close to the efficient allocation rule and, as a result, a run can occur with positive probability in an equilibrium of this extended game. (See, for example, Cooper and Ross [3], Peck and Shell [9], and Ennis and Keister [5].)

[24] Andolfatto and Nosal [1] introduce moral hazard on the part of the intermediary into the Green-Lin framework, but conclude it is not a potential source of run equilibria.

Appendix A. Proofs

Lemma 1: Under (3), $\theta_i(\omega^i) = \theta_i(\widehat{\omega}^i)$ implies both

$$\mathcal{P}(\omega^i) = \mathcal{P}(\widehat{\omega}^i)$$

and

$$\mathcal{P}(\omega^i, \omega^{I-i}) = \mathcal{P}(\widehat{\omega}^i, \omega^{I-i}) \quad \text{for all } \omega^{I-i}.$$

Proof. We begin with the second part. Let θ_{I-i} denote the number of patient traders in the continuation history ω^{I-1}. Then, using (3), we can write

$$\begin{aligned}
\mathcal{P}(\omega^i, \omega^{I-i}) &= \frac{p(\theta_i(\omega^i) + \theta_{I-i}(\omega^{I-i}))}{C(I, \theta_i(\omega^i) + \theta_{I-i}(\omega^{I-i}))} \\
&= \frac{p(\theta_i(\widehat{\omega}^i) + \theta_{I-i}(\omega^{I-i}))}{C(I, \theta_i(\widehat{\omega}^i) + \theta_{I-i}(\omega^{I-i}))} \\
&= \mathcal{P}(\widehat{\omega}^i, \omega^{I-i}),
\end{aligned}$$

where the second equality follows from the hypothesis of the lemma and the final equality from (3). This result allows us to establish the first as follows:

$$\begin{aligned}
\mathcal{P}(\omega^i) &= \sum_{\widetilde{\omega}^{I-1} \in \Omega^{I-i}} \mathcal{P}(\omega^i, \widetilde{\omega}^{I-i}) \\
&= \sum_{\widetilde{\omega}^{I-1} \in \Omega^{I-i}} \mathcal{P}(\widehat{\omega}^i, \widetilde{\omega}^{I-i}) \\
&= \mathcal{P}(\widehat{\omega}^i).
\end{aligned}$$

■

Proposition 1: *The efficient allocation when all traders contact the intermediary in period* 0 *sets*

$$c_i^0 = \frac{y_{i-1}}{\psi_i(\theta_{i-1})^{\frac{1}{\gamma}} + 1} \quad \text{for } i = 1, \ldots, I,$$

where $y_{i-1} = I - \sum_{j<i} c_j^0$ *and the functions* ψ_i *are defined recursively by* $\psi_I(x) = \left(xR^{\frac{1-\gamma}{\gamma}}\right)^{\gamma}$ *and*

$$\psi_i(x) = \pi_{i+1}(x)\left(\psi_{i+1}(x)^{\frac{1}{\gamma}} + 1\right)^{\gamma} + (1 - \pi_{i+1}(x))\,\psi_{i+1}(x+1)$$

for $i = 1, \ldots, I-1$.

Proof. Let V_i^0 denote the sum of the expected utilities of all traders who have not yet consumed when the intermediary encounters trader i, conditional on trader i being impatient and the intermediary dividing the available resources y_{i-1} efficiently among these traders. Specifically, this sum includes the utility levels of trader i, all traders after i in the sequence, and all traders before i who are patient and thus will consume at date 1. Let V_i^1 denote this same sum of expected utilities conditional instead on trader i being patient. These two value functions must satisfy the following recursive equations:

$$V_i^0\left(y_{i-1}, \theta_{i-1}\right) = \max_{\{c_i^0\}} \left\{ \begin{array}{c} \frac{\left(c_i^0\right)^{1-\gamma}}{1-\gamma} + \pi_{i+1}\left(\theta_{i-1}\right) V_{i+1}^0\left(y_{i-1} - c_i^0, \theta_{i-1}\right) + \\ \left(1 - \pi_{i+1}\left(\theta_{i-1}\right)\right) V_{i+1}^1\left(y_{i-1} - c_i^0, \theta_{i-1}\right) \end{array} \right\} \tag{22}$$

and

$$V_i^1\left(y_{i-1}, \theta_{i-1}\right) = \left\{ \begin{array}{c} \pi_{i+1}\left(\theta_{i-1} + 1\right) V_{i+1}^0\left(y_{i-1}, \theta_{i-1} + 1\right) + \\ \left(1 - \pi_{i+1}\left(\theta_{i-1} + 1\right)\right) V_{i+1}^1\left(y_{i-1}, \theta_{i-1} + 1\right) \end{array} \right\} \tag{23}$$

for $i = 1, \ldots I$. The function $\pi_i\left(\theta\right)$ in these equations represents the probability that $\omega_i = 0$ conditional on θ of the first $i-1$ traders being patient.

After the intermediary has encountered all I traders at date 0 and given consumption to the impatient ones, it will divide the remaining resources y_I, augmented by the return R, evenly among the θ_I patient traders at date 1. We therefore have the following terminal condition

$$V_{I+1}^0\left(y_I, \theta_I\right) = V_{I+1}^1\left(y_I, \theta_I\right) = \frac{\theta_I}{1-\gamma}\left(\frac{Ry_I}{\theta_I}\right)^{1-\gamma}.$$

The combination of this equation, the initial conditions $y_0 = I$ and $\theta_0 = 0$, and equations (22) and (23) constitutes the dynamic programming problem whose solution gives the efficient payment schedule.

As is common in finite-horizon dynamic programming problems, we start by solving the last decision problem the intermediary faces. Suppose trader I is impatient. Then, given θ_{I-1} and y_{I-1}, the maximization problem in (22) reduces to

$$\max_{\{c_I^0\}} \frac{\left(c_I^0\right)^{1-\gamma}}{1-\gamma} + \frac{\theta_{I-1}}{1-\gamma}\left(\frac{R\left(y_{I-1} - c_I^0\right)}{\theta_{I-1}}\right)^{1-\gamma}.$$

The solution to this problem sets

$$c_I^0 \left(y_{I-1}, \theta_{I-1} \right) = \frac{y_{I-1}}{\psi_I \left(\theta_{I-1} \right)^{\frac{1}{\gamma}} + 1},$$

where

$$\psi_I \left(x \right) \equiv \left(x R^{\frac{1-\gamma}{\gamma}} \right)^{\gamma}. \tag{24}$$

Substituting the solution back into the objective function and doing some straightforward algebra yields the value function

$$V_I^0 \left(y_{I-1}, \theta_{I-1} \right) = \frac{\left(y_{I-1} \right)^{1-\gamma}}{1-\gamma} \left(\psi_I \left(\theta_{I-1} \right)^{\frac{1}{\gamma}} + 1 \right)^{\gamma}.$$

If, on the other hand, trader I is patient, all of y_{I-1} is carried into date 1 and the value function is given by

$$V_I^1 \left(y_{I-1}, \theta_{I-1} \right) = \left(\theta_{I-1} + 1 \right) \frac{1}{1-\gamma} \left(\frac{R y_{I-1}}{\theta_{I-1} + 1} \right)^{1-\gamma},$$

which can also be written as

$$V_I^1 \left(y_{I-1}, \theta_{I-1} \right) = \frac{\left(y_{I-1} \right)^{1-\gamma}}{1-\gamma} \psi_I \left(\theta_{I-1} + 1 \right).$$

It is straightforward to use this same procedure to show that, for any trader $i < I$, the solution to the maximization problem in (22) sets

$$c_i^0 = \frac{y_{I-1}}{\psi_i \left(\theta_{i-1} \right)^{\frac{1}{\gamma}} + 1},$$

where

$$\psi_i \left(x \right) = \pi_{i+1} \left(x \right) \left(\psi_{I+1} \left(x \right)^{\frac{1}{\gamma}} + 1 \right)^{\gamma} + \left(1 - \pi_{i+1} \left(x \right) \right) \psi_{I+1} \left(x + 1 \right).$$

Note that, together with the "terminal" value ψ_I in (24), equation (11) can be used recursively to determine $\psi_i \left(\theta_{i-1} \right)$ for any values of i and θ_{i-1}. The associated value functions are

$$V_i^0 \left(y_{i-1}, \theta_{i-1} \right) = \frac{\left(y_{i-1} \right)^{1-\gamma}}{1-\gamma} \left(\psi_i \left(\theta_{i-1} \right)^{\frac{1}{\gamma}} + 1 \right)^{\gamma}$$

and

$$V_i^1 \left(y_{i-1}, \theta_{i-1} \right) = \frac{\left(y_{i-1} \right)^{1-\gamma}}{1-\gamma} \psi_i \left(\theta_{i-1} + 1 \right). \qquad \blacksquare$$

Lemma 2 *Under the mechanism α^*, reporting truthfully (that is, the strategy $\mu_i = \omega_i$) is a strictly dominant strategy for traders I and $I-1$.*

Proof. We already know that reporting truthfully is strictly preferred if a trader is impatient, so we only need to consider the case where each trader is patient. Consider first trader I. For any level of remaining resources y_{I-1}, the efficient allocation gives her the following payments depending on her report:

$$\text{lie:} \quad \frac{y_{I-1}}{\psi_I(\theta_{I-1})^{\frac{1}{\gamma}}+1} \quad \text{where } \psi_I(\theta_{I-1}) = \left(\theta_{I-1}R^{\frac{1-\gamma}{\gamma}}\right)^{\gamma}.$$

$$\text{truth:} \quad \frac{Ry_{I-1}}{\theta_{I-1}+1}$$

Truth-telling is strictly preferred if

$$\frac{R}{\theta_{I-1}+1} > \frac{1}{\theta_{I-1}R^{\frac{1-\gamma}{\gamma}}+1} = \frac{R}{\theta_{I-1}R^{\frac{1}{\gamma}}+R}$$

or if

$$\theta_{I-1}R^{\frac{1}{\gamma}} + R > \theta_{I-1}+1.$$

Since $R > 1$ and $\gamma > 0$, this condition holds for all $\theta_{I-1} \geq 0$. In other words, trader I strictly prefers to report truthfully regardless of the reports of other traders.

Next, consider the decision problem of trader $I-1$ in the event that he is patient. Let ϕ denote the probability he places on trader I reporting impatient. Then for a given level y_{I-2} of remaining resources, the expected utility of trader $I-1$ under α^* is

$$\text{lie:} \quad \tfrac{1}{1-\gamma}\left(c_{I-1}^0(y_{I-2},\theta_{I-2})\right)^{1-\gamma}$$

$$\text{truth:} \quad \phi\tfrac{1}{1-\gamma}\left(\tfrac{R\left(y_{I-2}-c_I^0(y_{I-2},\theta_{I-2}+1)\right)}{\theta_{I-2}+1}\right)^{1-\gamma} + (1-\phi)\tfrac{1}{1-\gamma}\left(\tfrac{Ry_{I-2}}{\theta_{I-2}+2}\right)^{1-\gamma}$$

where c_{I-1}^0 and c_I^0 are as derived in Section 2.2. It is straightforward to show that

$$\frac{Ry_{I-2}}{\theta_{I-2}+2} > \frac{R\left(y_{I-2}-c_I^0(y_{I-2},\theta_{I-2}+1)\right)}{\theta_{I-2}+1}$$

holds for all y_{I-2} and all θ_{I-2} (substitute in for c_I^0 and simplify). In other words, if trader $I-1$ reports patient, her consumption will be higher if trader I also reports patient than if the latter reports impatient. The claim will be proven, therefore, for any value of ϕ if we can show

$$\frac{R\left(y_{I-2}-c_{I}^{0}\left(y_{I-2},\theta_{I-2}+1\right)\right)}{\theta_{I-2}+1}>c_{I-1}^{0}\left(y_{I-2},\theta_{I-2}\right),$$

which can be reduced to

$$\frac{R^{\frac{1}{\gamma}}}{\left(\theta_{I-2}+1\right)R^{\frac{1-\gamma}{\gamma}}+1}>\frac{1}{\left(p_{I}\left(\theta_{I-2}\right)\left(\theta_{I-2}R^{\frac{1-\gamma}{\gamma}}+1\right)^{\gamma}+\left(1-p_{I}\left(\theta_{I-2}\right)\right)\left(\left(\theta_{I-2}+1\right)R^{\frac{1-\gamma}{\gamma}}\right)^{\gamma}\right)^{\frac{1}{\gamma}}+1}$$

Since $\gamma>0$ and $R>1$, we have

$$\theta_{I-2}R^{\frac{1-\gamma}{\gamma}}+1>\left(\theta_{I-2}+1\right)R^{\frac{1-\gamma}{\gamma}}$$

which implies that the denominator on the right-hand side is larger than that on the left-hand side. Since $R^{\frac{1}{\gamma}}>1$, the numerator on the left-hand side is larger, and hence the condition must hold. These calculations show that the consumption trader $I-1$ receives at date 1 if he reports patient is greater than the consumption he receives at date 0 if he reports impatient, even if trader I is certain to report impatient and independent of the reports of all previous traders. Therefore, reporting truthfully is also a strictly dominant strategy for trader $I-1$. ■

Proposition 2: *The efficient payment schedule when only impatient traders contact the intermediary in period* 0 *sets*

$$x_{n}=\frac{z_{n-1}}{\left(\phi_{n}\right)^{\frac{1}{\gamma}}+1}\quad\text{for } n=1,\ldots,I,$$

where $z_{n-1}=I-\sum_{j<n}x_{j}$ *and the constants* ϕ_{n} *are defined recursively by* $\phi_{I}=0$ *and*

$$\phi_{n}=q_{n+1}\left(\phi_{n+1}{}^{\frac{1}{\gamma}}+1\right)^{\gamma}+\left(1-q_{n+1}\right)\left(I-n\right)^{\gamma}R^{1-\gamma}$$

for $n=1,\ldots,I-1$.

Proof. Let V_n denote the sum of the expected utilities of all traders who have not yet consumed when the intermediary encounters the n^{th} impatient trader, conditional on the intermediary dividing the available resources z_{n-1} efficiently among these traders. These values must satisfy the following recursive equation:

$$V_n(z_{n-1}) = \max_{\{x_n\}} \left\{ \begin{array}{c} \frac{(x_n)^{1-\gamma}}{1-\gamma} + q_{n+1} V_{n+1}(z_{n-1} - x_n) + \\ (1 - q_{n+1})(I - n) \frac{1}{1-\gamma} \left(\frac{R(z_{n-1} - x_n)}{I-n} \right)^{1-\gamma} \end{array} \right\}, \tag{25}$$

for $n = 1, \ldots I$.

If all I traders are impatient, the intermediary will give all of the remaining resources to the last trader when she reports. We therefore have the following terminal condition

$$V_I(z_{I-1}) = \frac{1}{1-\gamma}(z_{I-1})^{1-\gamma}.$$

The combination of this equation, the initial condition $z_0 = I$, and equation (25) constitutes the dynamic programming problem whose solution gives the efficient payment schedule.

Consider the decision problem faced by the intermediary if it faces an $(I-1)^{th}$ impatient trader. Given z_{I-2}, the maximization problem in (25) reduces to

$$\max_{\{x_{I-1}\}} \frac{(x_{I-1})^{1-\gamma}}{1-\gamma} + q_I \frac{(z_{I-2} - x_{I-1})^{1-\gamma}}{1-\gamma} + (1 - q_I) \frac{(R(z_{I-2} - x_{I-1}))^{1-\gamma}}{1-\gamma}.$$

The solution to this problem sets

$$x_{I-1} = \frac{z_{I-2}}{(\phi_{I-1})^{\frac{1}{\gamma}} + 1},$$

where

$$\phi_{I-1} \equiv q_I + (1 - q_I) R^{1-\gamma}. \tag{26}$$

Substituting the solution back into the objective function and doing some straightforward algebra yields the value function

$$V_{I-1}(z_{I-2}) = \frac{(z_{I-2})^{1-\gamma}}{1-\gamma} \left((\phi_{I-1})^{\frac{1}{\gamma}} + 1 \right)^{\gamma}.$$

The function V_{I-1} captures the utility of the last two traders to report to the intermediary in the event that at least $I-1$ traders are impatient. In this case, the $(I-1)^{th}$ trader to report is necessarily impatient. The I^{th} may also be impatient, reporting in period 0, or patient, in which case she will report in period 1. The probabilities of these events (given by q_I) are contained in the constant ϕ_{I-1}.

It is straightforward to use this same procedure to show that, for any $n < I$, the solution to the maximization problem in (25) sets

$$x_n = \frac{z_{n-1}}{(\phi_n)^{\frac{1}{\gamma}} + 1},$$

where

$$\phi_n = q_{n+1} \left(\phi_{n+1}{}^{\frac{1}{\gamma}} + 1 \right)^{\gamma} + (1 - q_{n+1}) \, (I - n)^{\gamma} \, R^{1-\gamma}. \tag{27}$$

Note that condition (26) emerges naturally from (27) using the “terminal” value $\phi_I = 0$. ■

References

[1] Andolfatto, David and Ed Nosal, "Bank incentives, contract design, and bank runs," *Journal of Economic Theory* (forthcoming).

[2] Andolfatto, David, Ed Nosal, and Neil Wallace, "The role of independence in the Green-Lin Diamond-Dybvig Model," *Journal of Economic Theory* 137 (2007), 709-715.

[3] Cooper, Russell and Thomas W. Ross, "Bank runs: liquidity costs and investment distortions," *Journal of Monetary Economics* 41 (1998), 27-38.

[4] Diamond, Douglas W. and Phillip H. Dybvig, "Bank runs, deposit insurance, and liquidity," *Journal of Political Economy* 91 (1983), 401-419.

[5] Ennis, Huberto M. and Todd Keister, "Commitment and equilibrium bank runs," Staff Report 274, Federal Reserve Bank of New York, January 2007.

[6] Green, Edward J. and Ping Lin, "Implementing efficient allocations in a model of financial intermediation," *Journal of Economic Theory* 109 (2003), 1-23.

[7] Green, Edward J. and Ping Lin, "Diamond and Dybvig's classic theory of financial intermediation: What's missing?" Federal Reserve Bank of Minneapolis *Quarterly Review* 24 (2000), 3-13.

[8] Jacklin, Charles, 1987, Demand deposits, trading restrictions, and risk sharing, in: E. Prescott and N. Wallace, eds., Contractual arrangements for intertemporal trade (University of Minnesota Press, Minneapolis, MN) 26-47.

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

[10] Wallace, Neil. "Another attempt to explain an illiquid banking system: the Diamond and Dybvig model with sequential service taken seriously," *Federal Reserve Bank of Minneapolis Quarterly Review* 12 (Fall 1988), 3-16.

[11] Wallace, Neil. "A banking model in which partial suspension is best," *Federal Reserve Bank of Minneapolis Quarterly Review* 14 (Fall 1990), 11-23.