Run Equilibria in the Green-Lin Model of Financial Intermediation

Huberto M. Ennis
Department of Economics, Universidad Carlos III de Madrid
and
Research Department, Federal Reserve Bank of Richmond
hennis@eco.uc3m.es

Todd Keister
Research and Statistics Group, Federal Reserve Bank of New York
Todd.Keister@ny.frb.org

May 4, 2009

Abstract

We study the Green and Lin [J. Econ. Theory 109 (2003) 1-23] model of financial intermediation under a more general specification of the distribution of types across agents. We derive the efficient allocation in closed form. We show that, in some cases, the intermediary cannot uniquely implement the efficient allocation using a direct revelation mechanism. In these cases, the mechanism also admits an equilibrium in which some (but not all) agents "run" on the intermediary and withdraw their funds regardless of their true liquidity needs. In other words, self-fulfilling runs can arise in a generalized Green-Lin model and these runs are necessarily partial, with only some agents participating.

Forthcoming in the *Journal of Economic Theory*. We thank seminar participants at the Federal Reserve Banks of Chicago, New York, Philadelphia and Richmond; the Ohio State University; the Cornell - Penn State Macroeconomics Workshop; the Midwest Macroeconomics Meeting; North American Meeting of the Econometric Society; the annual meeting of the Society for Economic Dynamics; and the Conference on Money, Credit and Policy at Washington University for useful comments. We are especially grateful to Marco Bassetto, Ed Green, Ed Nosal, Jim Peck, Neil Wallace, Yaron Azrieli, the associate editor and two anonymous referees. Ennis acknowledges the financial support of the Ministry of Science and Technology in Spain (Projects 2008/00439/001 and 2009/00071/001). The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York, the Federal Reserve Bank of Richmond, or the Federal Reserve System.

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. Their key departure from the previous literature is to assume that agents have information about the order in which they will have an opportunity to withdraw. 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 that the efficient outcome will obtain in equilibrium, leaving no possibility of a self-fulfilling run.

We study the Green-Lin model under a more general specification of the distribution of preference types across agents. Whereas Green and Lin [6] assume that consumption needs are independent across agents, we allow for correlation. We show how the efficient allocation in this environment can be found by solving a finite dynamic-programming problem, and we derive this allocation in closed form. We then construct examples with the following properties. The efficient allocation is (Bayesian) incentive compatible and, hence, can be implemented by a direct revelation mechanism in which each agent reports his preference type to the intermediary. However, this mechanism also admits an equilibrium in which some, but not all, agents run on the intermediary and withdraw – claiming an immediate consumption need – regardless of their true type. In other words, we show that self-fulfilling runs *can* emerge in a generalized Green-Lin model, and that these runs are necessarily *partial*, with only some agents participating.

In the examples we construct, it is unlikely that all agents in the economy will face an immediate need to consume. Once a large number of withdrawals have taken place, therefore, the intermediary will infer that few of the remaining agents have immediate consumption needs. If some agents withdraw even though they do not need to consume right away, this inference will be incorrect. In other words, when some agents run, their actions tend to make the intermediary unduly optimistic about the consumption needs of the *remaining* agents. The intermediary will then conserve relatively few resources for future withdrawals. When the intermediary discovers that the consumption needs of the remaining agents are higher than anticipated, it will decrease all subsequent payments to agents, including the future payments to agents who have chosen not to withdraw.

Suppose, then, that an individual believes that the agents who have an opportunity to withdraw before she does will all run. She recognizes that if she does not withdraw, the payment she receives from the intermediary in the future will likely be small, which gives her an incentive to join the run and withdraw right away. Notice that this incentive applies even if she believes the agents who come after her will *not* participate in the run. The key point is that some of these agents may truly have immediate consumption needs and, given her beliefs about these agents' types, the intermediary has kept inadequate resources to deal with those needs. This incentive to run only applies if an individual's withdrawal opportunity is early enough, that is, if sufficiently many agents will contact the intermediary after her. As emphasized by Green and Lin [6, 7], an agent who knows he is the last to contact the intermediary never has an incentive to run. For this reason, the run equilibria we construct are necessarily partial; agents who are able to withdraw early do so, while those who act later only withdraw if they have an immediate consumption need.

Notice that the effects described above disappear when types are assumed to be independent, as in Green and Lin [6]. When an agent withdraws in that case, the action has no effect on the intermediary's perception of the types of the remaining agents since these types are simply independent draws from a given distribution. Correlation in types is necessary for these effects to be present and, hence, for our results to obtain.

Peck and Shell [9] also construct run equilibria in a model of financial intermediation 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 before knowing the order in which they will contact the intermediary. Relative to Green and Lin [6], this approach introduces an additional information friction into the environment: agents must act before knowing this particular (payoff-relevant) information. Whether such a friction is necessary to generate run equilibria has remained an open question. We show that this friction is not necessary; run equilibria can exist even when agents know the precise order in which they will contact the intermediary.

Our framework also allows us to clarify the precise nature of the Peck-Shell results. In addition to the informational assumptions described above, their examples rely on agents having different preferences than in the previous literature and it has been unclear whether this change in preferences is necessary for their results to obtain. In Section 4, we present examples in the spirit of Peck and Shell [9] in which types are independent and preferences – and all other aspects of the model except the informational assumptions – are exactly as in Green and Lin [6]. These examples demonstrate that the difference in information assumptions alone is sufficient to generate the Peck-Shell results.

The remainder of the paper is organized as follows. We present the environment and derive the efficient allocation in Sections 2 and 3, respectively. In Section 4, 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 examples in the spirit of Peck and Shell [9]. In Section 5, we study the effects of correlation in types and present our main results. We offer some concluding remarks in Section 6.

2 The Model

In this section we extend the Green-Lin model of financial intermediation to allow types to be correlated across agents. We largely follow the notation in Green and Lin [6], with some simplifications where possible.

2.1 The environment

There are two time periods, indexed by $t \in \{0, 1\}$, and a finite number I 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. Traders are isolated from each other and from the intermediary (as in Wallace [10]), but have an opportunity to contact the intermediary in each period in order to receive goods. Goods are nonstorable and must be consumed immediately after contacting the intermediary.[1] Let $a_i = (a_i^0, a_i^1)$ denote the consumption of trader i in each period and let $a = (a_1, \ldots, a_I)$ denote the complete vector of consumption bundles.

[1] This assumption implies that markets in which agents could trade after contacting the intermediary are infeasible. See Jacklin [8] and Wallace [10] on this point.

Preferences. Trader i's utility is given by

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

where $\omega_i \in \{0,1\}$ is her type and $\gamma > 1$ holds.[2] 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, traders will contact the intermediary 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 $\omega \in \Omega^I$.

Uncertainty. Let $\mathcal{P}$ denote the probability measure on the set of all subsets of Ω^I. We assume that there exists a non-negative function p with

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

such that

$$\mathcal{P}(\omega) = \frac{p(\theta(\omega))}{C(I, \theta(\omega))} \quad \text{for all } \omega, \tag{2}$$

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 $\omega_i = 0$. Note that, under this approach, each trader has the same *ex ante* probability of being patient.[3] The assumption of independent types used by Green and Lin [6] is a special case where the density p is given by the

[2] The assumption of a specific functional form for the utility function is not necessary here; Green and Lin [6] only place assumptions on the level of risk aversion. 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]).

[3] More specifically, traders are *exchangeable* in the sense that the distribution of $(\omega_1, \ldots, \omega_I)$ is invariant under permutations.

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.

Technology. The intermediary has an aggregate endowment of I units of the good in period 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. An (ex post) allocation in this environment is an assignment of an individual allocation a_i to each trader. Letting $\mathbb{I} = \{1, 2, \ldots, I\}$ denote the set of traders, the set of *feasible (ex post) allocations* is given by

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

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

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

Sequential Service. Traders contact the intermediary sequentially in period 0 in a fixed order given by the index i, beginning with trader 1 and ending with trader I.[4] As a result, the period-0 consumption of trader i can only depend on the partial history ω^i; the intermediary cannot possibly know the types of any of the remaining traders when this payment must be made. The constraint can be written as

$$\mathbf{a}_i^0(\omega) = E\left[\mathbf{a}_i^0(\omega) \mid \omega_1, \ldots, \omega_i\right]$$

or, alternatively, as

$$\mathbf{a}_i^0(\omega) = \mathbf{a}_i^0(\widehat{\omega}) \quad \text{for all } \omega, \widehat{\omega} \text{ such that } \omega^i = \widehat{\omega}^i. \tag{3}$$

In other words, trader i must consume the same amount in any two states that the intermediary cannot possibly distinguish given the information it could have potentially received so far. We

[4] We follow Green and Lin [7] and Andolfatto *et al.* [2] in assuming that traders know the exact order in which they will contact the intermediary, rather than having imperfect information about this order as in Green and Lin [6]. The two approaches lead to the same results and the former simplifies the notation considerably.

denote the set of feasible state-contingent allocations that satisfy the sequential service constraint by

$$\mathbb{F}' = \{\mathbf{a} \in \mathbb{F} : (3) \text{ holds}\}.$$

Expected Utility. Traders seek to maximize the expected value of the utility function v conditional on their information. We can write the information set of trader i as

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

that is, a trader knows only her own type. Given a state-contingent allocation $\mathbf{a}$ and a (true) state of nature ω^*, define

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

where the expectation over states ω is based on the probability distribution $\mathcal{P}$ updated to reflect trader i's private information. 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 $\omega_i = 1$.

3 The Efficient Allocation

We now derive the efficient, symmetric state-contingent allocation, that is, the allocation the intermediary would assign if traders' types were observable.[5] 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.

The efficient allocation is the solution to

$$\max_{\mathbf{a} \in \mathbb{F}'} \sum_{i \in \mathbb{I}} E\left[U_i(\mathbf{a}, \omega)\right]. \tag{4}$$

[5] 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.

Let $\mathbf{a}^*$ denote this solution. It is straightforward to show that, under the preferences in (1), efficiency requires that impatient traders only consume in period 0 and patient traders only consume in period 1. In other words, $\mathbf{a}^*$ must satisfy

$$\mathbf{a}_i^0(\omega) = 0 \ \text{ if } \ \omega_i = 1 \qquad \text{and} \qquad \mathbf{a}_i^1(\omega) = 0 \ \text{ if } \ \omega_i = 0. \tag{5}$$

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

$$\mathbf{a}_i^1(\omega) = \frac{R\left(I - \sum_{i=1}^{I} \mathbf{a}_i^0(\omega)\right)}{\theta(\omega)}. \tag{6}$$

All that remains, then, is to determine the payment that would be given to each trader i in period 0 if she is impatient, as a function of the partial history ω^i. In other words, we need to determine $\mathbf{a}_i^0(\omega)$ for histories with $\omega_i = 0$. These payments can be found by using the results above to reformulate (4) as a finite dynamic programming problem.

Our formulation of the problem makes use of some important implications of condition (2), which governs the correlation structure of types. First, the condition implies that any two histories ω and $\widehat{\omega}$ with $\theta(\omega) = \theta(\widehat{\omega})$ are assigned the same probability by $\mathcal{P}$.[6] Second, consider the probability of some continuation history $\omega^{I-i} = (\omega_{i+1}, \ldots, \omega_I)$ conditional on the partial history $\omega^i = (\omega_1, \ldots .\omega_i)$. Condition (2) implies that this probability depends only on the number of patient traders in the partial history, denoted $\theta_i(\omega^i)$, and not on their positions within the history. Abusing notation slightly, let $\mathcal{P}(\omega^i)$ 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 the appendix.

Lemma 1 *Under* (2), $\theta_i(\omega^i) = \theta_i\left(\widehat{\omega}^i\right)$ *implies both*

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

and

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

[6] This fact is easily seen in (2), where the expression on the right-hand side depends on $\theta(\omega)$ but not directly on ω.

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 in period 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 a_i^0 denote the payment.[7]

The proposition below presents the efficient payments a_i^0. The proof consists of converting (4) 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.[8] We then have the following result.

Proposition 1 *The efficient allocation sets*

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

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

$$\psi_i(x) = \pi_{i+1}(x)\left(\psi_{i+1}(x)^{\frac{1}{\gamma}} + 1\right)^{\gamma} + (1 - \pi_{i+1}(x))\,\psi_{i+1}(x+1) \tag{7}$$

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

A proof of the proposition is given in the appendix. Note that (7) depends only on the conditional probabilities π_i and the parameters R and γ. This equation can, therefore, be used recursively to determine $\psi_i(\theta_{i-1})$ for any values of i and θ_{i-1}. These functions ψ_i can then be used recursively, starting with trader 1, to determine the efficient payment to an impatient trader 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

[7] A comment on notation: The variable a_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 $\mathbf{a}_i^0(\omega^i)$.

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

given by $R = 1.1$ and $\gamma = 6$. The figure shows the possible period 0 consumption levels of each trader. The circles correspond to partial histories in which trader 1 is impatient, while the triangles correspond to histories in which trader 1 is patient.

The level of consumption trader 1 receives if she is impatient is given by the left-most circle 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 circle), while if trader 1 was patient the payment to trader 2 will be larger (the triangle). 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.

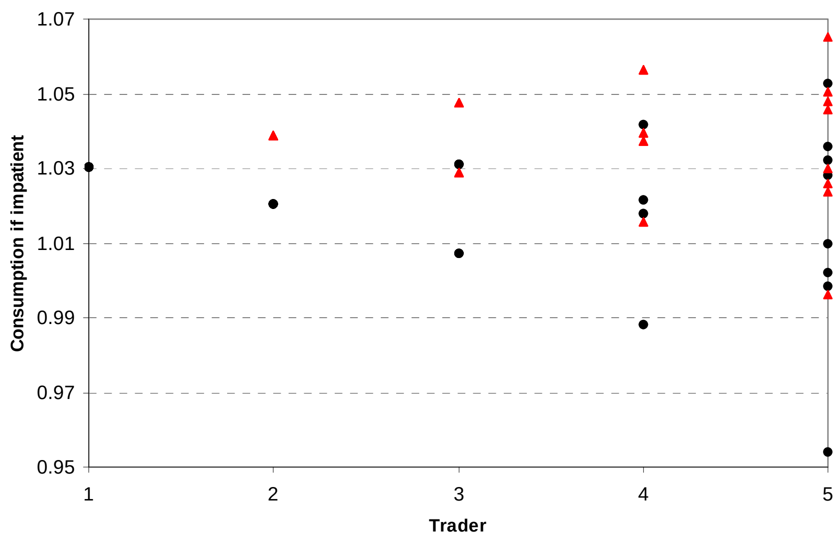

Figure 1: Efficient allocation with independent types

4 Implementation

Proposition 1 derives 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. Our question of interest is whether or not the resulting game has an equilibrium where traders "run" on the intermediary by mis-reporting their types.

4.1 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 $\mu_i : \Omega^I \to M$. 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.[9] Let Γ denote the set of such rules, i.e.,

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

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 sequential service* if the corresponding state-contingent allocation $\mathbf{a}$ satisfies (3) for every profile of communication strategies μ. Let Γ' denote the set of feasible allocation rules that respect sequential service.

In general, an allocation *mechanism* specifies both a message space and an allocation rule (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[10]

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

[9] 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.

[10] 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].

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

$$\mathbf{a}(\omega) = \alpha(\mu^*(\omega)) \quad \text{for all } \omega. \tag{8}$$

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 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 (4) derived in Section 3. In what follows, we focus exclusively on this allocation rule, which we denote α^*.

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?

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 (5) states that any trader reporting to be patient will be given zero consumption in period 0. Furthermore, it is straightforward to show that all traders reporting to be impatient will receive positive consumption in period 0. Since impatient traders only care about consumption in period 0, lying when a trader is impatient is a strictly dominated strategy. For the analysis of equilibrium, therefore, we only need to examine the action of a trader in the event that she is patient.

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

Suppose, for a moment, 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. 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

$$\frac{1}{I}\sum_{i\in\mathbb{I}} E\left[U_i\left(\mathbf{a},\omega\right)\right]. \tag{9}$$

Note that this expression is equivalent to (4), the objective function of the intermediary.[11]

This approach was implicitly taken in the original work of Diamond and Dybvig [4] and in much of the subsequent literature, including the recent work of Peck and Shell [9]. Relative to the Green-Lin model described above, it contains an additional information friction: a trader's place in the order is not known to her when she must choose an action. We show that, under this approach, an equilibrium can exist in which all traders run on the intermediary. 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 (1); rather, in their setting the marginal utility of consumption is higher for impatient traders than for patient traders. This assumption simplifies the computations in their model by ensuring that an incentive compatibility constraint binds at the efficient allocation. We show that differing marginal utilities are not necessary for their result to obtain. Everything in our examples below is exactly as in the Green-Lin model except for the information that traders have when choosing an action. In particular, Proposition 1 still characterizes the efficient allocation in this setting, and types are assumed to be independent.

Proposition 2 *Suppose types are independent. 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 first part of the proposition follows directly from Green and Lin [6], who showed that the efficient allocation is incentive compatible when types are independent and 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, therefore, that a trader who does

[11] A proper formulation of this case would distinguish between a trader's index (or "name") and her eventual place in the ordering. (See Green and Lin [6] for such a formulation.) Doing so, however, adds an extra layer of notation without introducing any new insights. Since the issue only arises in the present subsection, we take the notational shortcut of having traders act before any "names" are assigned. This shortcut does not change the essential analysis in any way.

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?

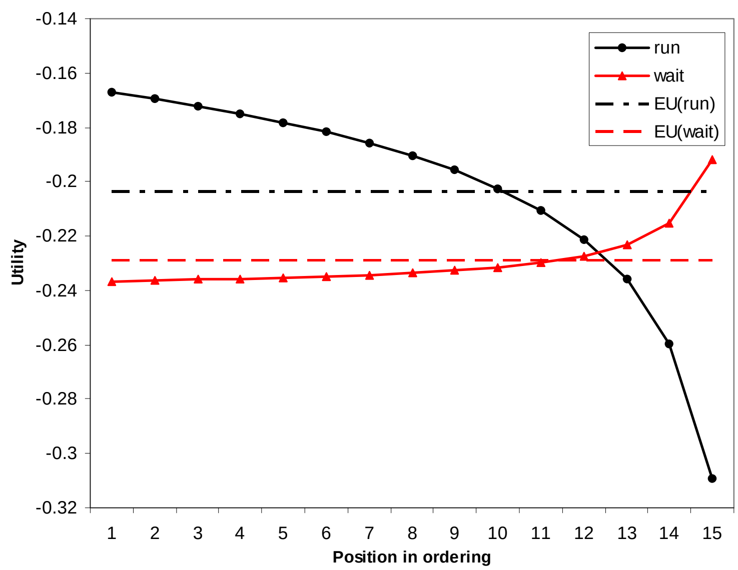

Figure 2: Expected utility if all other traders run

Figure 2 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 circles represents the utility from running, which is strictly decreasing in the trader's position in the order. The triangles 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 action, given that a trader is equally likely to end up in each of the 15 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 3, 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 upper curve. 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 lower curve 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 tends to make it more likely that a run equilibrium will exist in this setting.

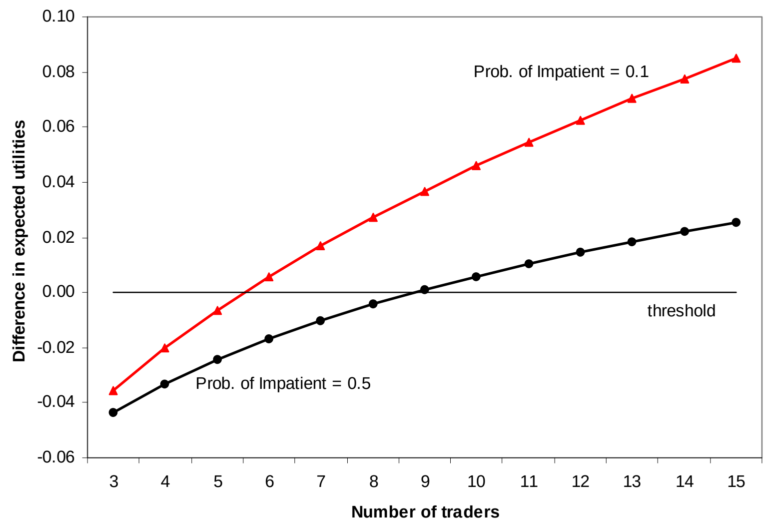

Figure 3: Incentive to run as I varies

4.3 A unique implementation result (Green-Lin)

Going back to Figure 2, 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.[12] Suppose, then, that a trader had some information about where she is likely to be in the order. A patient trader who believes she is the last to arrive would prefer to report truthfully and wait until period 1 to consume. All 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. They showed that in this case, assuming independent types, the direct revelation game has a *unique* equilibrium under the efficient allocation rule α^*. In that equilibrium, all traders truthfully report their types; no one runs on the intermediary.

Proposition 3 *(Green and Lin* [6]*) If 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 – illiquidity, private information, and sequential service – do not necessarily open the door to self-fulfilling runs. In a particular environment that contains these features, an intermediary can ensure that the efficient allocation obtains through the proper choice of contract. 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.

The proof of the Green-Lin result uses iterated deletion of strictly dominated strategies in a backward-induction fashion. We stated above that reporting truthfully is a strictly dominant strategy for the last two traders in the order (see Lemma 2 in the appendix). As a first step, then, all other strategies for these traders can be deleted. Once this is done, reporting truthfully becomes a strictly dominant strategy for the third-to-last trader, and the process continues. The strategy profile in which all traders report truthfully is the unique profile that survives this process and, hence, is the unique Bayesian Nash equilibrium of the game.

One might be tempted to conclude from this argument that the unique implementation result is "simply" a matter of backward induction and will thus obtain in *any* environment where traders

[12] This result is stated as Lemma 2 in the appendix and a proof is given there.

know the order in which they will contact the intermediary. In the next section, we show this is not true. In other words, the arguments behind the results in Green and Lin [6] are more subtle than they might appear at first and, as a result, introducing correlation in traders' types can generate equilibria in which some traders choose to run.

5 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, all traders know the order in which they will contact the intermediary. 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. We show that a run equilibrium can exist in this setting; specifically, we establish the following result.

Proposition 4 *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 and we describe the intuition behind this example in some detail. 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.

5.1 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 in the previous section, 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

$$p(2) = 1 - \varepsilon, \quad \text{and} \tag{10}$$

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

where $p(\theta)$ is, as defined in (2), 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 4. 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.

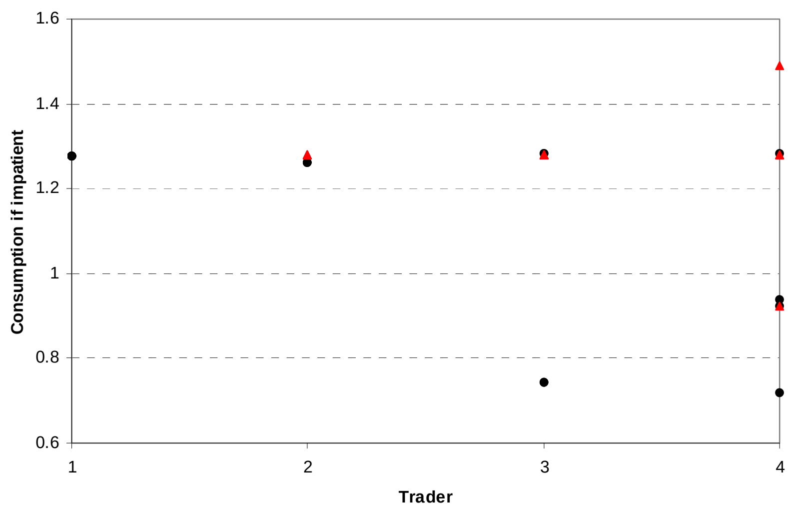

Figure 4: Efficient allocation with correlated types

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 4 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 intermediary 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 it made (both close to $\bar{c}^0$) were based on a belief about the state ω that turned out to be "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*.[13]

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 5. The dashed 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 below zero 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.

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

$$\mu_i(\omega) = \left\{ \begin{array}{ccc} 0 & \text{for} & i = 1, 2 \\ \omega_i & & i = 3, 4 \end{array} \right\}. \tag{11}$$

Lemma 2 (see the appendix) tells us that the strategies in (11) 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

[13] 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).

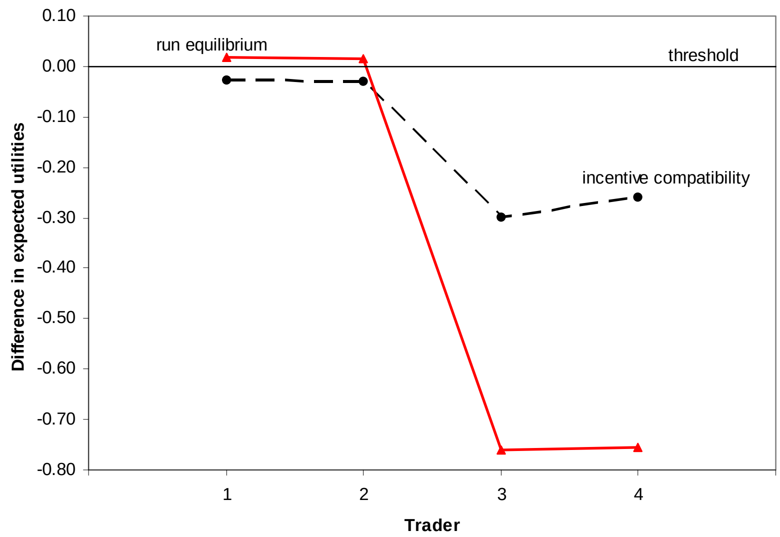

Figure 5: Individual incentive to run with correlated types

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. She expects trader 1 to report $m_1 = 0$ regardless of his true type. It is possible that trader 1 is indeed impatient, which would likely imply that only one of the remaining traders (3 or 4) is impatient. If trader 2 reports truthfully, in this case only two payments are likely to be made in period 0 and, therefore, her payment in period 1 will be relatively large (close to the $\bar{c}^1$ of the Diamond-Dybvig model).[14] Reporting truthfully would then be the best choice.

If, on the other hand, trader 1 is patient (and, hence, his report is untruthful), then it is very 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 in period 0 and the amount left for her in period 1 will be substantially smaller. If she lies, on the other hand, her report of impatient will be only the second received by the intermediary and she will receive a larger payment in period 0 (close 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 line in Figure 5 presents the gain in expected utility from choosing the run strategy under the assumption that other traders are following the strategy profile in (11). The fact

[14] 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.

that this line is positive at trader 2 shows that, in this example, behaving in accordance with (11) – mis-reporting her type when patient – is the optimal choice for trader 2.

The decision problem faced by trader 1 is similar. The fact that the solid line in Figure 5 lies above zero for him shows that he will also strictly prefer to follow (11) if all other traders are doing so. Finally, note that the solid 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 (11) is indeed an equilibrium for the chosen parameter values.

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 (11), report truthfully. This behavior is somewhat surprising in light of the results in Green and Lin [6, Lemma 5], which show that it cannot arise in the model with independent types. In particular, Green and Lin demonstrate that, under the efficient allocation rule, a trader who believes everyone after her will report truthfully strictly prefers to do the same, regardless of the actions of the traders who contact the intermediary before her. We call this property *continuation incentive compatibility* or *continuation IC*. The key to understanding why the Green-Lin unique implementation result does not extend to the case of correlated types, therefore, is understanding why the continuation IC property does not hold for trader 2 under the efficient allocation rule in this example.

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 assessment, derived from the probabilities in (10), that traders 3 and 4 are very likely to both be patient. As a result of this assessment, 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 the assessment 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 assessment used to design the allocation rule, 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 running – and claiming the period 0 payment based on an unduly optimistic assessment – more attractive. In this way, the correlation in traders' types breaks the continuation IC property of the efficient allocation.

Notice that these effects cannot arise when types are independent. 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 assessment 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.

5.2 Another example

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

Parameter Values. Let $I = 10$. The parameter values $R = 2$ and $\gamma = 6$ are unchanged from the simple example above. Set the density function p for the number of patient traders as follows

$$p(\theta) = \frac{1-\varepsilon}{5} \text{ for } \theta = 3, \ldots 7, \text{ and} \tag{12}$$

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

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 more complex than before; however, using the proposition its calculation remains straightforward. We 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 line in Figure 6. The fact that the line is below zero for all traders indicates that the efficient allocation is indeed incentive compatible in this example.

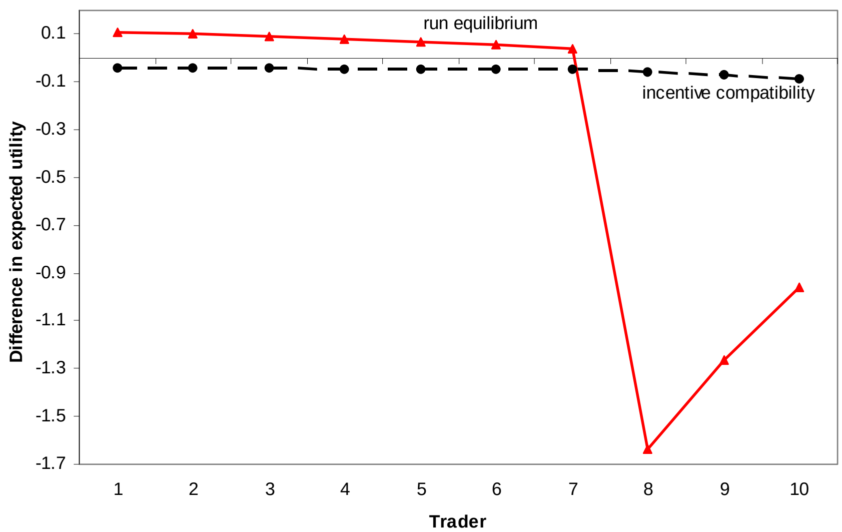

Figure 6: Incentive to run in the richer example

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

$$\mu_i(\omega) = \left\{ \begin{array}{ccc} 0 & \text{for} & i = 1, \ldots, 7 \\ \omega_i & & i = 8,\ 9,\ 10 \end{array} \right\}. \tag{13}$$

The solid line in Figure 6 plots the expected gain from following the run strategy (relative to reporting truthfully) when all other traders are expected to follow the strategies in (13). 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 (13), she strictly prefers to do so as well. The strategy profile (13) is, therefore, an equilibrium of the direct revelation game.

Intuition. As before, the key to understanding why a run equilibrium exists is to examine why the continuation IC property fails to hold under the efficient allocation rule. In this example, the behavior of trader 7 is critical. Why does she run even though she anticipates that everyone after her will report truthfully? The intuition for this behavior is very similar to that for trader 2's behavior in the previous example.

When the first seven traders all report to be impatient, the payment given to trader 7 is based on the assessment – generated using the probabilities in (12) – that the remaining three traders are very likely to all be patient. As a result, this payment is relatively large, reflecting the "optimistic" view that further withdrawals in period 0 are unlikely. In the run equilibrium, however, Trader 7 has a more pessimistic belief about the types of the remaining traders and, hence, about the number of additional withdrawals. 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 belief, running – and claiming the period 0 payment based on the optimistic assessment – is more attractive than reporting truthfully and waiting until period 1.

5.3 Discussion

In the Green-Lin model, a trader correctly anticipates the equilibrium strategies of all other traders, but she does not directly observe their actions. Other approaches are possible, of course, and it may be interesting to study the extent to which the insights we present here can be applied in related environments. Andolfatto et al. [2], for example, modify the Green-Lin model by allowing a trader to observe the reports of the previous traders prior to announcing her type. The set of feasible allocations is unchanged in this modified environment. The direct revelation game is different, however, because a trader has more information when she chooses her report. Despite this fact, the

Green-Lin results continue to hold in the modified environment; the efficient allocation is again incentive compatible and can be uniquely implemented when types are independent.[15]

Andolfatto et al. [2] discuss the possibility of constructing run equilibria based on correlated types in their modified environment, but they do not offer any examples. Constructing equilibria in their environment is more involved because it requires specifying traders' beliefs at decision nodes that lie off the equilibrium path of play, which then raises the issue of the "reasonableness" of those beliefs according to various equilibrium refinement criteria. They argue that for the special case of $I = 3$, any run equilibrium would fail to satisfy a natural refinement like the intuitive criterion. The logic they describe does not directly apply to examples with a larger number of traders, however.

We study the environment as specified by Green and Lin [6], where traders do not observe each other's actions and these issues do not arise. This setting has the advantage of capturing the essential features of financial intermediation, including the timing of agents' decisions, in a parsimonious way. In fact, one of the important contributions of Green and Lin [6] is to show how a banking model with dynamic elements – and which allows the use of backward-induction reasoning – can be formulated as an essentially static (simultaneous-move) game. Whether our techniques, as illustrated in the examples above, can be applied to generate robust partial-run equilibria in environments like that in Andolfatto et al. [2] is an interesting question for future research.

6 Concluding Remarks

Observers frequently claim that runs on banks and other financial intermediaries are driven largely by self-fulfilling beliefs. Because these events typically occur during periods of financial turmoil, determining their underlying causes empirically has proven difficult. It is, therefore, useful to ask whether self-fulfilling runs are theoretically plausible, in the sense of being equilibrium outcomes of a reasonable economic model. Green and Lin [6] derive a remarkable result in this respect: in their version of the classic Diamond-Dybvig model, an intermediary can generate the efficient allocation of resources without introducing the type of self-fulfilling runs that appeared in the earlier literature.

[15] In fact, Andolfatto et al. [2] prove that *any* allocation satisfying incentive compatibility in the modified environment can be uniquely implemented when types are independent. They also show that this result holds for a broader set of preferences than studied in Green and Lin [6].

The Green-Lin model incorporates the basic elements that are commonly believed to open to the door to self-fulfilling runs, including illiquidity, private information, and sequential service. Their key departure from the existing literature is to assume that an individual has some information about where things stand when her opportunity to withdraw arrives. They show how this information can have a large effect in equilibrium. For an individual whose opportunity to withdraw comes very late, this information eliminates her incentive to run on the intermediary. Using a backward induction argument, they then show how the incentives for the remaining individuals to run unravel.

One might be tempted to infer from this logic that self-fulfilling run equilibria are necessarily an artifact of arbitrary modeling restrictions. In particular, one might think that run equilibria simply cannot occur in an environment in which (i) individuals are able to condition their decisions on their order of arrival at the intermediary and (ii) the payment schedule takes full advantage of the information available to the intermediary. We show that this is not the case. To do so, we change the Green-Lin model in only one respect: we allow for liquidity needs to be correlated across agents. We then construct equilibria in which some agents run on the intermediary. Our results show that it is difficult to dismiss the possibility of self-fulfilling runs on purely theoretical grounds.[16]

In the model we have presented here, two distinct types of crises can occur in equilibrium. In addition to a self-fulfilling run, the economy can experience a "fundamentals" crisis in which a large fraction of the population has immediate liquidity needs. Interestingly, an outside observer would have great difficulty distinguishing between these two types of crises. In both cases, individuals who withdraw late receive much lower consumption than those who withdrew earlier. Determining the underlying cause of a particular crisis is made particularly difficult by a new feature of our equilibrium construction: self-fulfilling runs here are partial, with only some agents participating. Despite this difficulty, the model has a key empirical prediction: the incidence of crises can be higher than would be warranted by the behavior of economic fundamentals alone.

We did not study how the intermediary would react to this higher probability of crises, but doing so would be fairly straightforward. To address this issue, the analysis would need to be expanded

[16] Modifying the Green-Lin model in other ways is likely to yield additional insights, both about the potential for self-fulfilling runs and about the structure of financial intermediation in general. Andolfatto and Nosal [1], for example, introduce a self-interested intermediary into the Green-Lin framework and study how this modification affects the optimal allocation rule and the potential for self-fulfilling runs, concluding it is not a potential source of the latter.

to include the intermediary as a player in the game. This step has been taken in related models in the existing literature and the results of such an exercise are now well known. (See, for example, Cooper and Ross [3], Peck and Shell [9], and Ennis and Keister [5].) A run typically cannot occur with certainty in an equilibrium of the expanded game, because the intermediary would then choose a rule under which truth-telling is the unique equilibrium. If the ex ante probability of a run is small enough, however, the intermediary will choose to follow an allocation rule close to the efficient rule studied here. A partial run can then be consistent with equilibrium, as we have shown. Following the steps in the existing literature, one could then construct a correlated equilibrium of the overall game in which a partial run – based on the strategies identified in this paper – occurs with positive probability.

Appendix A. Proofs

Lemma 1: Under (2), $\theta_i(\omega^i) = \theta_i\left(\widehat{\omega}^i\right)$ implies both

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

and

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

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 (2), we can write

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

where the second equality follows from the hypothesis of the lemma and the final equality from

(2). This result allows us to establish the first as follows:

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

■

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

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

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

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

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 in period 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:

$$V_i^0\left(y_{i-1},\theta_{i-1}\right) = \max_{\left\{a_i^0\right\}} \left\{ \begin{array}{c} \frac{\left(a_i^0\right)^{1-\gamma}}{1-\gamma} + \pi_{i+1}\left(\theta_{i-1}\right) V_{i+1}^0\left(y_{i-1}-a_i^0,\theta_{i-1}\right) + \\ \left(1-\pi_{i+1}\left(\theta_{i-1}\right)\right) V_{i+1}^1\left(y_{i-1}-a_i^0,\theta_{i-1}\right) \end{array} \right\} \tag{14}$$

and

$$V_i^1\left(y_{i-1},\theta_{i-1}\right) = \left\{ \begin{array}{c} \pi_{i+1}\left(\theta_{i-1}+1\right) V_{i+1}^0\left(y_{i-1},\theta_{i-1}+1\right) + \\ \left(1-\pi_{i+1}\left(\theta_{i-1}+1\right)\right) V_{i+1}^1\left(y_{i-1},\theta_{i-1}+1\right) \end{array} \right\} \tag{15}$$

for $i = 1,\ldots I$. The function $\pi_i(\theta)$ 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 in period 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 in period 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 (14) and (15) constitutes the dynamic programming problem whose solution gives the efficient payment schedule.

As is typical in finite 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 (14) reduces to

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

The solution to this problem sets

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

where

$$\psi_I\left(x\right)\equiv\left(xR^{\frac{1-\gamma}{\gamma}}\right)^{\gamma}. \tag{16}$$

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 period 1 and the value function is given by

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

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 (14) sets

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

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 (16), Eq. (7) 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

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

■

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:

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

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

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

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

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

where a^0_{I-1} and a^0_I are as derived in Section 3. It is straightforward to show that

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

holds for all y_{I-2} and all θ_{I-2} (substitute in for a^0_I 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

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

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 in period 1 if he reports patient is greater than the consumption he receives in period 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 $I-1$. ■

References

[1] D. Andolfatto, E. Nosal, Bank incentives, contract design and bank runs, J. Econ. Theory 142 (2008) 28-47.

[2] D. Andolfatto, E. Nosal, N. Wallace, The role of independence in the Green-Lin Diamond-Dybvig Model, J. Econ. Theory 137 (2007) 709-715.

[3] R. Cooper, T.W. Ross, Bank runs: liquidity costs and investment distortions, J. Monet. Econ. 41 (1998) 27-38.

[4] D.W. Diamond, P.H. Dybvig, Bank runs, deposit insurance, and liquidity, J. Polit. Economy 91 (1983) 401-419.

[5] H.M. Ennis, T. Keister, Commitment and equilibrium bank runs, Staff Report 274, Fed. Reserve Bank of New York, January 2007.

[6] E.J. Green, P. Lin, Implementing efficient allocations in a model of financial intermediation, J. Econ. Theory 109 (2003) 1-23.

[7] E.J. Green, P. Lin, Diamond and Dybvig's classic theory of financial intermediation: what's missing? Fed. Reserve Bank of Minneapolis Quart. Review 24 (2000) 3-13.

[8] C. Jacklin, Demand deposits, trading restrictions, and risk sharing, in: E. Prescott, N. Wallace (Eds.), Contractual Arrangements for Intertemporal Trade, University of Minnesota Press, Minneapolis, MN, 1987, pp. 26-47.

[9] J. Peck, K. Shell, Equilibrium bank runs, J. Polit. Economy 111 (2003) 103-123.

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

[11] N. Wallace, A banking model in which partial suspension is best, Fed. Reserve Bank of Minneapolis Quart. Review 14 (1990) 11-23.