Discussion Paper Series – CRC TR 224

Discussion Paper No. 540
Project C 03

Smart Banks

Alkis Georgiadis-Harris[1]
Maxi Guennewig[2]
Yuliyan Mitkov[3]

April 2024

[1]University of Warwick, Email: alkisgharris@gmail.com
[2]University of Bonn, Email: m.guennewig@uni-bonn.de
[3]University of Bonn, Email: ymitkov@uni-bonn.de

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Collaborative Research Center Transregio 224 - www.crctr224.de
Rheinische Friedrich-Wilhelms-Universität Bonn - Universität Mannheim

Smart Banks*

Alkis Georgiadis-Harris† Maxi Guennewig‡ Yuliyan Mitkov§

December 22, 2023

Abstract

Since Diamond and Dybvig (1983), banks have been viewed as inherently fragile. We challenge this view in a general mechanism design framework, where we allow for flexibility in the design of banking mechanisms while maintaining limited commitment of the intermediary to future mechanisms. We find that the unique equilibrium outcome is efficient. Consequently, runs cannot occur in equilibrium. Our analysis points to the ultimate source of fragility: banks are fragile if they cannot collect and optimally respond to useful information during a run and not because they engage in maturity transformation. We link our banking mechanisms to recent technological advances surrounding 'smart contracts,' which enrich the practical possibilities for banking arrangements.

Keywords: Bank runs, financial fragility, mechanism design, limited commitment, smart contracts.

JEL Codes: D82, G2

*We thank Fernando Alvarez, Bruno Biais, Xuesong Huang, Todd Keister, Cyril Monnet, Giorgia Piacentino, Bruno Sultanum, Elu von Thadden, seminar participants at the University of Mannheim and the University of Amsterdam, as well as conference participants at the Summer Workshop on Money, Banking, Payments, and Finance 2023. Maxi Guennewig and Yuliyan Mitkov gratefully acknowledge support from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through the CRC TR 224 (Project C03).

†Department of Economics, University of Warwick. Coventry CV4 7AL, United Kingdom. E-mail: *alkisgharris@gmail.com*.

‡Department of Economics, University of Bonn. Adenauerallee 24-42, 53113 Bonn, Germany. E-mail: *mguennewig@uni-bonn.de*.

§Department of Economics, University of Bonn. Adenauerallee 24-42, 53113 Bonn, Germany. E-mail: *ymitkov@uni-bonn.de*.

1 Introduction

Since the seminal contribution of Diamond and Dybvig (1983), maturity transformation has been viewed as an inherently fragile activity.[1] Specifically, a financial intermediary provides the efficient level of liquidity insurance by issuing deposits while investing in long-maturity, high-return assets. However, the intermediary becomes illiquid in the process: not all obligations can be honored at face value if all depositors request to be paid immediately. This illiquidity gives rise to self-fulfilling runs. This highlights an important trade-off between efficiency and stability: one can achieve stability by mandating banking arrangements with inefficiently low levels of maturity transformation, or must live with the possibility of bank runs.

Indeed, bank runs have been recurrent events.[2] Moreover, they often result in costly and distortionary government bailout interventions. It is then unsurprising that regulators have designed many policies to prevent bank runs, mainly in the form of deposit insurance and liquidity regulation. However, to design such policies optimally, one needs a thorough understanding of the conditions that give rise to the trade-off between efficiency and stability.

This paper sets out to understand the ultimate source of fragility in the Diamond-Dybvig model of bank runs. The model has two essential components. First, the depositors have *private information* about their liquidity needs, with some having an urgent need to consume, whereas others can wait and consume later. Second, those with an urgent need to consume must be paid on demand since forcing them to wait is inefficient, leading to a *sequential service* constraint.

Early work by Wallace (1988) identified sequential service as a *necessary* constraint in the design of intermediaries for fragility to emerge. Without sequential service, the bank can collect all withdrawal requests and then assign payments accordingly. In that case, depositors without an urgent need to consume would prefer to leave their funds in the bank to generate a positive return. As a consequence, runs cannot occur. Thus, a sequential service constraint is necessary for the Diamond-Dybvig model to be a theory of banking,

[1] Diamond and Dybvig (and to a lesser extent Bryant (1980)) were the first to formalize it, but the idea of inherent fragility goes further back. For example, J.P. Morgan said during the crisis of 1907: "If the people will keep their money in the banks, everything will be all right." See Bruner and Carr (2009), pp. 100–01.

[2] Many banks in various countries experienced runs during the Great Depression in the 1930s and the Great Financial Crisis of 2007-2009. The U.S. alone experienced systematic banking panics in 1857, 1873, 1893, 1907, 1931, and 1933. Other banking and financial crisis examples include the Nordic countries 1991-3, Mexico 1994-5, East Asia 1997-8, and Argentina 2001-2. More recently, Silicon Valley Bank in the United States and Credit Suisse in Switzerland experienced run-like events leading to their failure.

illiquidity, *and* fragility.

At the same time, sequential service is not *sufficient* for bank runs: whether or not banks are fragile depends on the details of the environment and, in particular, on how fast the bank infers that a run is underway by observing depositors' withdrawal behavior (Green and Lin, 2003; Peck and Shell, 2003; Ennis and Keister, 2009b). However, the run equilibria in this literature have an odd property: depositors know that a run is underway, but the bank does not. In such cases, it seems natural to allow for richer mechanisms that elicit additional information from the depositors. Several recent papers (Cavalcanti and Monteiro, 2016; Andolfatto et al., 2017) show how mechanisms designed so that a patient depositor who runs on the bank is willing to reveal her true type—thus alerting the bank to the run—can uniquely implement the efficient allocation. No bank runs can occur in equilibrium. The key idea is to impose a strict deposit freeze immediately after detecting a run.[3]

However, deposit freezes are highly inefficient during a run as some depositors with urgent needs to consume cannot access their funds within the bank. Deposit freezes thus require a high degree of commitment power. This is unreasonable in practice since, during crises, regulators routinely take control of banks to preserve depositor welfare.[4] Ennis and Keister (2009a, 2010) shows how the inability to commit to deposit freezes causes runs. Nevertheless, they study this problem in a restricted contracting environment, and consequently, it is hard to delineate the extent to which limited commitment drives fragility.

To fully understand the ultimate source of fragility, one needs to allow for flexibility in the design of mechanisms while relaxing the commitment power of the intermediary. To this end, we develop a 'mechanism selection game' among a financial intermediary ('Banker') and many agents ('depositors'), building on the canonical Diamond-Dybvig model. Each depositor is allocated either some immediate consumption out of current resources or a promise of future consumption out of future resources. Depositors interact with the Banker one-by-one, and we identify a *stage* in this game to be an interaction with a particular depositor.

The interaction between the Banker and each depositor is through a mechanism that

[3]These papers extend an idea from Diamond and Dybvig (1983) where the bank was assumed to freeze deposits once withdrawal demand reveals that a run is underway. This solution only works if there is little uncertainty about total fundamental withdrawal demand. In Cavalcanti and Monteiro (2016) and Andolfatto et al. (2017), the bank detects the run more quickly using an indirect mechanism, which can work even when there is significant aggregate uncertainty.

[4]This can vividly be seen in the 2001 Argentinian banking crisis, where banks reneged on their promised payment plans. See Ennis and Keister (2009a) for details. Wallace (1990) offers historical examples of partial suspensions, for example, in the 1907 financial crisis. Friedman and Schwartz (1963) describe how suspensions during the Great Depression came long after it was clear that banks were in trouble.

governs their communication and determines the allocation for that stage. Importantly, at the beginning of each stage, the Banker can unilaterally scrap the existing mechanism and replace it with a new one. That new mechanism then governs the current and all future interactions until it is replaced—if ever. Notice that implicit in the description of our game is some level of commitment power by the Banker: within a stage, the selected mechanism must govern at least the current interaction, and it can only be replaced after the current stage is completed. In that sense, the Banker has intra-stage but not inter-stage commitment.

We emphasize that our mechanism selection game captures the three key ingredients we seek to model: (i) sequential service, which is embedded in the structure of our extensive-form game; (ii) contractual flexibility, reflected in the lack of restrictions on the available mechanisms; and (iii) limited commitment, captured by the ability of the Banker to replace any mechanism ex-post, as the game unfolds.

We find that the *unique* equilibrium outcome of the mechanism selection game is *efficient*. Consequently, there cannot be runs in *any* equilibrium of this game. This means that with enough flexibility in the design of mechanisms, limited commitment poses no problems, and the trade-off between efficiency and fragility disappears. This result relies on two assumptions. First, the efficient outcome (subject to sequential service) *is* an equilibrium of the mechanism selection game. Second, the depositors follow the mechanisms' instructions when indifferent between messages—but are otherwise free to discard these instructions.

What underpins this result is the following. By 'separating' the mechanism from the Banker, we can construct mechanisms that collect useful information *even* in the event of a run. In particular, the mechanism not only services withdrawal requests but also learns whether withdrawing depositors have urgent needs to consume. This information is then used to set up ex-post efficient allocations. When such a mechanism is in place, only depositors with an urgent need to consume withdraw, *and* the Banker is deterred from replacing the mechanism. Consequently, bank runs are avoided, and efficiency is achieved in every equilibrium.

The mechanism we propose does not rely on government guarantees or banking regulation and can be implemented by private financial arrangements if there is sufficient contractual flexibility. Thus, the ultimate source of fragility is not maturity transformation, private information about liquidity needs, sequential service, or limited commitment. Instead, banks become fragile when they fail to collect *and* optimally respond to useful information.

A rich contracting space is not only a key theoretical aspect of this paper but also of increasing practical relevance. Specifically, recent technological advancements significantly enrich the contracts that can be conceivably implemented. Indeed, in the spirit of Brzustowski et al. (2023), our banking mechanisms can be interpreted as resting on two minimal properties of 'smart contracts': automatic execution and cryptographic encryption.[5]

Finally, we do not wish to proclaim the end of bank runs. Rather, our analysis implies that self-fulfilling bank runs do not emerge due to the inherent fragility of maturity transformation. Instead, distortions generated by government guarantees could be a cause of bank runs. Indeed, historical experience suggests that governments find it difficult to commit *not* to bail out financial institutions in times of crisis.[6] Even within a Diamond-Dybvig setting, a bank that anticipates being bailed out but does not fully internalize the cost of the public funds might optimally choose to operate with a fragile liability structure (Keister, 2016; Keister and Mitkov, 2023). Alternatively, bank runs may be (i) behavioral phenomena (e.g., manias, panics, and crashes as per Kindleberger et al. (2005)) (ii) caused by insolvency rather than illiquidity (Allen and Gale, 1998), or (iii) resulting from institutional restrictions on financial contracts.

Literature Review. The paper primarily speaks to the theoretical literature on financial fragility following the seminal contribution of Diamond and Dybvig (1983).

The Diamond-Dybvig model has two key features: private information about liquidity needs and a sequential service constraint. Sequential service, as envisioned by Diamond-Dybvig and formalized by Wallace (1988, 1990), implies that depositors urgently need to consume and must be paid on demand. At the same time, sequential service is a necessary but not sufficient condition for bank runs. Diamond and Dybvig (1983) show that when the proportion of impatient depositors is known, a simple deposit freeze eliminates runs at no cost in terms of efficiency.[7]

Diamond and Dybvig (1983) also recognized that the above contract is not optimal if the number of impatient depositors is uncertain. In an influential paper, Green and Lin (2003)

[5]A third desirable property of these contracts property is *immutability*, which amounts to the impossibility of ever altering or stopping the protocol and, in our case, corresponds to inter-stage commitment. Our mechanism does not require immutability, which is fortunate since, as already mentioned, this is not an appealing assumption in a banking context.

[6]Large theoretical literature shows how bailouts can distort banks' incentives. See, for example, Farhi and Tirole (2012), Chari and Kehoe (2016), Bianchi (2016), Nosal and Ordonez (2016), Dávila and Walther (2020), and Philippon and Wang (2023).

[7]Papers in the literature generate runs even in that case, by assuming that the bank must give a prespecified payment until it runs out of funds (Postlewaite and Vives, 1987; Cooper and Ross, 1998; Allen and Gale, 2004a; Goldstein and Pauzner, 2005). This 'simple contracts' approach is convenient but at odds with the observation that the liabilities of financial intermediaries are often altered in times of financial distress (Ennis and Keister, 2009a, 2010).

characterize the constraint efficient subject to sequential service and aggregate uncertainty. They show that the payment to the current depositor will depend on the number of past withdrawals - if this quantity is high, the current depositor is paid less, leading to a sort of partial suspension. Green and Lin (2003) demonstrate a striking result: a direct mechanism can uniquely implement the constraint efficient allocation. As a result, the bank in their setup is not inherently fragile.

However, the uniqueness result in Green and Lin (2003) is very sensitive to the details of the environment. Ennis and Keister (2009b) study a version of the Green and Lin where depositors' liquidity shocks are correlated, showing that run equilibria exist under the optimal direct mechanism. The reason is that correlation in liquidity shocks can generate *belief divergence* between the depositors and the intermediary since the Banker might be too slow in inferring that a run is underway (see Section 5). Peck and Shell (2003) and Ennis and Keister (2009b) study modifications of Green and Lin where the depositors do not know their line position and show that run equilibrium exists under the optimal direct mechanism. The reason is that the unique equilibrium in Green and Lin relies on backward induction logic, which does not apply if the depositors do not have some information about their line position.[8]

The difference in run outcomes in the above paper can be traced to different informational frictions. However, the information of the different players can be viewed as an *outcome* of a mechanism design problem – a point recognized by Nosal and Wallace (2009a). They, in particular, advocate an approach where the depositors receive information about their preferences, and the intermediary can choose what information to reveal to the depositors. In the same spirit, we allow mechanisms to control the allocation *and* the flow of information to the depositors and the Banker. In this sense, we aim to capture the different approaches in the literature and provide a unifying framework to study fragility.

Cavalcanti and Monteiro (2016) and Andolfatto et al. (2017) study a general mechanism design problem with full commitment. Both papers show that an indirect mechanism that learns that a run is underway (by eliciting depositors' types) and then implements a strict deposit freeze can prevent runs. Full commitment is key to those mechanisms since, as already mentioned, deposit freezes are time-inconsistent: suspending withdrawals for all depositors, including those with an urgent need to consume, is inefficient ex-post. If the bank reneges on the deposit freeze, which is anticipated in equilibrium, then depositors

[8]Huang (2023) analyzes the environment of Ennis and Keister (2009b) but assumes that depositors observe all past withdrawals (as in Andolfatto et al. (2007a)). He uses forward induction logic to show that runs cannot happen under the optimal direct mechanism (see also Kinateder and Kiss (2014)).

without an urgent need to consume run on the bank in the first place Ennis and Keister (2009a, 2010). In contrast, the mechanism design problem in our case treats the Banker as a player and thus takes into account her incentives to replace the current mechanism. Nevertheless, we show that runs cannot occur under the optimal mechanism.

We also contribute to the literature on dynamic contracting without inter-temporal commitment, which goes back to Laffont and Tirole (1988, 1990).[9] The Principal (the Banker) posts a mechanism determining the 'rules of the game' for the remaining agents (the depositors). After the current interaction is over, the Principal can *unilaterally* scrap the existing mechanism and deploy a different mechanism that, from this moment on, would determine the 'rules of the game.' Importantly, the Principal has intra-temporal commitment since she cannot replace the mechanism before the current interaction ends.

Finally, we show how intra-temporal commitment can be attained in our setup through a 'smart contract' that has automatic execution. The promise of smart contracts from a mechanism design perspective is that they can be viewed as technologies that facilitate the separation of the mechanism from the mechanism designer.[10] This notion is not new to the Diamond-Dybvig tradition. Wallace (1988) describes the intermediary as a 'cash machine' that will dispense consumption on demand. Our mechanism selection game expands on this idea.

Outline. Section 2 sets up the model environment. Section 3 introduces the mechanism selection game. We present our main result in section 4. Section 5 contrasts our mechanism with a direct mechanism that admits a run equilibrium. We discuss our modeling assumptions in sections 6 and 7. Section 8 concludes.

2 The model

2.1 The environment

There are two time periods $t \in \{0, 1\}$, a finite number of depositors N, indexed by $i \in \{1, 2, ..., N\}$, and a Banker. In addition, there is a single good that can be consumed in each

[9]Papers in that tradition include Tirole (2016); Beccuti and Möller (2018); Doval and Skreta (2022); Brzustowski et al. (2023).

[10]Several recent papers have explored how smart contracts can expand the contracting space by, for example, facilitating enforcement and preventing opportunistic renegotiations. See e.g. Tinn (2018); Cong and He (2019); Bakos and Halaburda (2020); Holden and Malani (2021).

period. The purpose of a banking arrangement is to allocate consumption goods among the depositors over the two periods. In particular, let c_i denote depositor i's 'early' consumption in period-0, and C_i his 'late' consumption in period-1. Denote by $\mathbf{c}_i = (c_i, C_i)$ depositor i's consumption bundle.

Depositor preferences. A depositor's utility over his consumption bundle depends on his payoff type and is given by

$$U(\mathbf{c}_i, \omega_i) \;=\; \begin{cases} u(c_i) & \text{if } \omega_i = 0 \\ v(c_i + C_i) & \text{if } \omega_i = 1 \end{cases}$$

where $u : \mathbb{R}_+ \to \mathbb{R}$, $v : \mathbb{R}_+ \to \mathbb{R}$, and $\omega_i \in \{0, 1\}$ is the depositor's type. If $\omega_i = 0$, the depositor is *impatient* and only values early consumption. If $\omega_i = 1$, the depositor is *patient* and values early and late consumption equally.

Payoff-types. Let $\mathbb{P}$ denote a probability measure on the set of all subsets of Ω^N where $\mathbb{P}$ is constructed as follows. Nature first draws the total number of patient depositors $\phi \in \{0, 1, ..., N\}$ according to the probability mass function $p : \{0, 1, ..., N\} \to [0, 1]$. Then for each $i \in \{1, ..., N\}$ we have $\omega_i = 1$ with probability $\frac{\phi}{N}$ and $\omega_i = 0$ with probability $\frac{N-\phi}{N}$. Thus, each depositor has the same ex-ante probability of being impatient. This specification allows depositors' payoff types to be independent, as in Green and Lin (2003), or correlated, as in Ennis and Keister (2009b).

Technology. The bank has $Y > 0$ units of the good in period 0.[11] Each unit not consumed in period 0 is transformed (i.e., matures) into $R > 1$ units in period 1. A realized allocation is an assignment of a consumption bundle for each depositor $\mathbf{c} = (c_i, C_i)_{i=1}^N$, which is feasible if

$$\sum_{i=1}^{N} \left(c_i + \frac{C_i}{R} \right) \leq Y \tag{1}$$

Sequential service. A central element of our environment is the sequential service constraint as envisioned in Diamond and Dybvig (1983) and formalized in Wallace (1988).

[11] We abstract from the pre-deposit game in which depositors decide whether to pool their endowments in the bank (Peck and Shell, 2003; Peck and Setayesh, 2023). This is without loss of generality since all depositors are ex-ante identical, and our mechanism uniquely implements the efficient allocation.

Sequential service has two main features. First, the depositors must be serviced in the order in which they contact the bank in period 0 without the possibility of determining their consumption after interacting with all remaining depositors. That is, impatient depositors have an 'urgent need to consume' and cannot wait to be serviced. The second feature is isolation, which has two aspects: a physical and an informational. First, depositors cannot engage in any trade.[12] Second, depositors cannot communicate with each other or the Banker once depositor-specific information has been realized.

We assume the depositors arrive at the banker in the order given by their index i. One can think of i as the depositor's line position. As in Green and Lin (2000), we assume each depositor knows his line position. The probability of being patient or impatient is independent of the depositor's line position.[13]

Banker preferences. The Banker's preferences are given by

$$W(\mathbf{c}, \omega) = \sum_{i \leq N} U(\mathbf{c}_i, \omega_i) \mathbb{1}_{[i \in \mathcal{D}(\mathbf{c})]}$$

where $\omega = (\omega_1, ..., \omega_N) \in \Omega^N$ denotes the full profile of consumption types; and $\mathbb{1}_{[i \in \mathcal{D}(\mathbf{c})]} = 1$ if depositor i belongs to the *coalition of depositors* $\mathcal{D}(\mathbf{c})$, and $\mathbb{1}_{[i \in \mathcal{D}(\mathbf{c})]} = 0$ otherwise. Depositor i is part of the coalition if he has not contacted the bank or if he has contacted and was promised late consumption.

The Banker's utility is thus the equally weighted sum of utilities of all depositors in the banking coalition. In other words, the Banker is no longer concerned with the welfare of depositors who have 'closed their account.' The Banker's benevolence towards the coalition of depositors is an analogy for the joint behavior of the Banker and a regulator in times of crisis. The regulator essentially prevents any actions that do not maximize the ex-post welfare of the bank's depositors. This specification of Banker preferences serves as a useful benchmark, allowing us to study the inherent fragility of banks without the confounding effect of other frictions.[14]

[12]It is well-known that the possibility of trade undermines maturity transformation (see Jacklin (1987), Allen and Gale (2004b), and Farhi et al. (2009)). As was noted by Wallace (1988), the isolation assumption, which captures that depositors have limited access to financial and asset markets when consumption opportunities arise, also rules out trade among the depositors.

[13]Assuming that depositors know their line positions streamlines the analysis but is not important for our results. See Section 6.3 for discussion on an alternative specification in which depositors receive a signal about their line position, which can be fully precise (Green and Lin, 2000; Ennis and Keister, 2009b), partially precise (Green and Lin, 2003), or entirely uninformative (Peck and Shell, 2003).

[14]Our analysis will be the same if profit-maximizing banks were competing for deposits in period 0, but

2.2 Continuation efficient allocation

To derive the efficient allocation in our environment, we formulate the problem of a social planner who starts from a given partial history and observes the payoff type of each depositor as they arrive to withdraw.[15]

Specifically, a partial history h^i is constructed as follows. Let $Y^i = Y - \sum_{n=1}^{i-1} c_n$ denote the bank's remaining goods after the initial $i-1$ depositors have contacted the bank. Let $\omega^i = (\omega_n)_{n \le i-1} \in \Omega^{i-1}$ denote the payoff types of the first $i-1$ depositors to arrive at the bank. We write $\pi_n = 1$ ($\pi_n = 0$) if depositor n was (was not) promised late consumption, and use $\pi^i = (\pi_n)_{n \le i-1} \in \{0,1\}^{i-1}$ indicate those depositors among the first $i-1$ that were promised late payment. A partial history is then given by $h^i = (Y^i, \omega^i, \pi^i)$. Since we want to allow for arbitrary partial histories, do not impose the restriction $\omega_i = \pi_i$. That is, we allow for partial histories where some patient depositors were not promised late consumption.

Conditional on h^i, the planner selects a state-contingent allocation. Suppressing the dependence on h^i, we denote it by $\mathbf{c} : \Omega^N \to \times_N \mathbb{R}^2_+$. Let $\mathbf{c}_i(\omega) = (c_i(\omega), C_i(\omega))$ denote the state-contingent bundle for depositor i. The planner chooses $\mathbf{c}$ to solve the program

$$W_i^* \;=\; \max_{\mathbf{c}} \mathbb{E}\left[W(\mathbf{c},\omega) \,|\, h^i\right] \tag{2}$$

subject to the feasibility constraint for each $\omega \in \Omega^N$

$$\sum_{n=1}^{i-1} \pi_n \frac{C_i(\omega)}{R} + \sum_{i=n}^{N} \left(c_i(\omega) + \frac{C_i(\omega)}{R} \right) \le Y^i \tag{3}$$

and subject to the sequential service constraint

$$c_i(\omega) = c_i(\tilde{\omega}) \qquad \text{for all } \omega,\, \tilde{\omega} \text{ such that } \omega^i = \tilde{\omega}^i \tag{4}$$

The above implies that conditional on history h^i, depositor i must receive the same early consumption in two different states ω and $\tilde{\omega}$, which the planner cannot distinguish given the information generated in the first $i-1$ interactions.

Let $\mathbf{c}^*$ denote the solution of the program in (2) - (4). It is straightforward to show that

a (benevolent) regulator has the authority to reschedule payments once a run is detected. We stress that bank runs in the standard Diamond-Dybvig setting emerge even when the coalition of depositors is always operated in the best interest of its members.

[15] We adopt the convention of writing the 'efficient' allocation rather than the 'constrained efficient' since we treat the sequential service constraint as part of the environment.

$\mathbf{c}^*$ will satisfy the following: for any $\omega \in \Omega^N$, we have $c_i^*(\omega) = 0$ if $\omega_i = 1$ and $C_i^*(\omega) = 0$ if $\omega_i = 0$. An impatient depositor consumes only in period 0, and a patient depositor consumes only in period 1. Finally, we impose the following assumption.

Assumption 1. *The continuation efficient allocation* $\mathbf{c}^*$ *is incentive compatible. That is, for each* i

$$\mathbb{E}\left[v\big(C_i^*(\omega)\big) - v\big(c_i^*(\omega)\big) \,|\, \omega_i = 1\right] \; > \; 0$$

According to the above, any given patient depositor i strictly prefers late consumption whenever all patient depositors that contact the bank after him only receive a promise of late consumption. Assumption 1 implies that the efficient allocation $\mathbf{c}^*$ is (weakly) implementable. It follows that W_i^* is the best continuation equilibrium payoff for a Banker who has full commitment but does not observe depositors' payoff types. Whether or not this assumption holds depends on the primitives of the environment, such as the depositors' preferences and the probability distribution over payoff types. In Section 7, we show that this continuation incentive constraint holds for the canonical specifications in the literature.

3 The Mechanism Selection Game

This section introduces the *mechanism selection game*, which aims to unify existing approaches to the Diamond-Dybvig model. Specifically, this game is characterized by (i) *sequential service*, which is embedded in the structure of the extensive-form game, (ii) *limited commitment*, captured by the ability of the Banker to replace any mechanism ex-post as the game unfolds, and (iii) *contractual flexibility* reflected in the lack of restrictions on the available mechanisms.

Sequential service is reflected in two ways in our setup. First, the Banker's information sets at each stage i contain information only about previous interactions. Therefore, any allocation determined in stage i can only depend on information generated by the mechanisms deployed in previous stages. The depositors cannot trade and communicate with each other once private information has been revealed to the depositors. Our formulation of sequential service therefore follows Wallace (1988) and other papers building on this formulation.[16]

Limited commitment is reflected in the fact that during any stage i, the Banker cannot commit to any mechanism she will select in the following stages. However, the Banker

[16]See, among others, Green and Lin (2003), Peck and Shell (2003), Andolfatto et al. (2007b), Ennis and Keister (2009a), Ennis and Keister (2009b), Nosal and Wallace (2009b), Ennis and Keister (2010), Andolfatto et al. (2017) and Huang (2023).

can commit to the mechanism within a stage: after a mechanism has been selected, that mechanism must govern at least the current interaction and can only be replaced after the current stage ends. In that sense, the Banker has intra-stage but not inter-stage commitment. We explain how this commitment power can be practically achieved in Section 6.1.

Furthermore, in stage N, we force the consumption of each patient depositor who has been promised late consumption to be a pro-rata share of available resources. This specification embeds the idea that the (benevolent) Banker cannot promise late consumption levels, which leads to ex-post inefficiencies among those depositors that have been promised period 1 consumption.

3.1 Timing

Nature draws each depositor's payoff type ω, and period 0 begins. We split period 0 into N stages. At the beginning of stage 1, the Banker selects some initial mechanism M from a class $\mathcal{M}$, which we describe in detail in Section 3 below. The mechanism then manages the communication between the Banker and the first depositor and determines the allocation. Importantly, after one stage finishes and before the next starts, the Banker can unilaterally scrap the existing mechanism M and replace it with a new one, $M' \in \mathcal{M}$. The new mechanism would then govern the current and all future interactions—until it is replaced, if ever. Period 0 ends after stage N, i.e., when all N depositors have interacted with the selected mechanisms by the Banker.[17] Period 1 contains only one stage $N + 1$ during which the remaining resources are converted into $R \cdot Y_N$ units of goods allocated among all depositors who were promised late consumption.

3.2 Banking mechanisms

In this section, we specify the set of banking mechanisms, $\mathcal{M}$, the Banker chooses from in the mechanism selection game, starting with the formal definition of a typical mechanism. A banking mechanism M consists of a sequence of protocols governing the interaction between each depositor and the Banker in each stage. Each protocol determines the communication between the active depositor and the Banker via an extensive-form game whose terminal

[17] In our game, all depositors interact with the mechanism in period-0 as in Green and Lin (2003) and others. This assumption differs from Peck and Shell (2003), where depositors only contact the mechanism when they seek early consumption. See Section 6.2 for a discussion.

nodes determine the allocation implemented in that stage. A protocol's terminal nodes end the current stage.

Protocols. The interaction within a protocol is depicted in Figure 1 and unfolds as follows: (i) The Banker privately submits a message to the protocol. (ii) The protocol sends a private signal to the active depositor. (iii) The active depositor then submits a private message to the protocol. (iv) The protocol determines the allocation to the active depositor: either some immediate consumption or a promise of future consumption and a private signal to the Banker.[18] Formally, each banking mechanism $M \in \mathcal{M}$ takes the form $M = \{P_k\}_{k=1}^{N}$ with

$$P_k = (B_k, A_k, \mathcal{B}_k, \mathcal{T}_k, \boldsymbol{\tau}_k, \Gamma_k, \boldsymbol{c}_k, \boldsymbol{\pi}_k, \boldsymbol{\beta}_k) \tag{5}$$

where

- B_k, A_k: message space for the Banker and current depositor
- $\mathcal{B}_k, \mathcal{T}_k$: signal space for the Banker and the current depositor
- $\boldsymbol{\tau}_k : B_k \times (B_n, A_n, \mathcal{B}_n, \mathcal{T}_n)_{n=1}^{k-1} \to \Delta(\mathcal{T}_k)$: private signal distribution for the current depositor
- $\Gamma_k \subseteq \mathbb{R}^+$ is some set of possible early consumption levels, with $0 \in \Gamma_k$.
- $(\boldsymbol{c}_k, \boldsymbol{\pi}_k) : (B_k, A_k, \mathcal{T}_k) \times (B_n, A_n, \mathcal{B}_n, \mathcal{T}_n)_{n=1}^{k-1} \to \Delta(\Gamma_k \times \{0, 1\})$: distribution functions for early consumption $\boldsymbol{c}_k \in \Delta(\Gamma_k)$ and late consumption promise $\boldsymbol{\pi}_k \in \Delta(\{0, 1\})$ for the current depositor.
- $\boldsymbol{\beta}_k : (B_k, A_k, \mathcal{T}_k) \times \Gamma_k \times (B_n, A_n, \mathcal{B}_n, \mathcal{T}_n)_{n=1}^{k-1} \to \Delta(\mathcal{B}_k)$: private signal distribution for the Banker

We use bold symbols $\boldsymbol{\tau}_k$ (similarly $\boldsymbol{c}_k$, $\boldsymbol{\pi}_k$, and $\boldsymbol{\beta}_k$) to denote a probability distribution and τ (similarly c, π, and β) to denote a particular realization from that distribution. We also write $P_k(M)$ (similarly $B_k(M)$, $A_k(M)$, etc) when we wish to emphasize the mechanism, $M \in \mathcal{M}$, to which a given protocol P_k belongs. We assume, without loss of generality, that the Banker observes the allocation implemented at every stage. We also assume that all levels of early consumption of any mechanism are feasible along any path of play supp $\{\boldsymbol{c}_k(.|.)\} \subseteq [0, Y_k]$.

[18] One can think of a protocol as a trusted *mediator* between the Banker and the depositor who communicates with the parties and selects the allocation for that stage of the game.

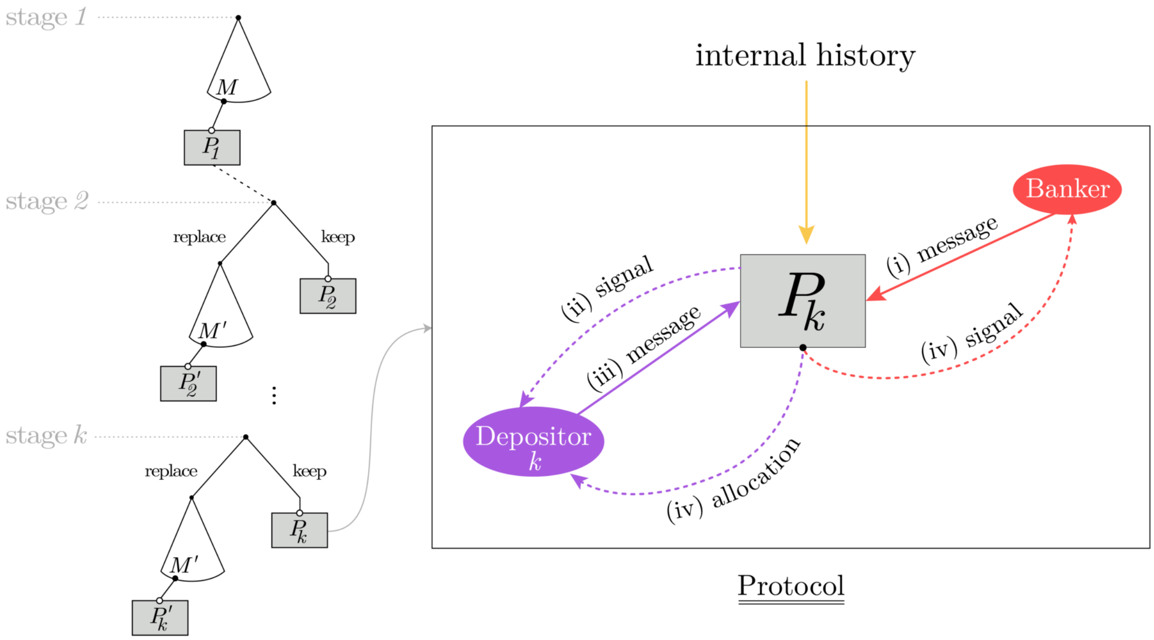

Figure 1: Illustration of the protocol

When initially deployed, each mechanism is provided with the exact value of the bank's remaining resources. However, the mechanism does not have any conception of how many depositors have already been serviced. That is, M does not know what is the true stage. Instead, it uses protocol 1 to interact with the first depositor it encounters, protocol 2 to interact with the second depositor, and so on. Any mechanism $M \in \mathcal{M}$ induces an internal history $h_k(M) \in \mathcal{H}_k(M)$ where $\mathcal{H}_k(M)$ denotes the set of all possible histories when mechanism M uses its k^{th} protocol. Each $P_k(M)$ can condition the information it generates and the implemented allocation on the history within that protocol and the entire history in preceding protocols of M, so long as it hasn't been replaced.[19]

Intra-stage interaction. We now describe the interaction between the current depositor and a protocol. Fix any stage $i \in \{1, ..., N\}$ and mechanism $M \in \mathcal{M}$ such that M has been continuously in place since stage $T \in \{1, ..., i-1\}$. Then M uses protocol P_k for the current depositor i, where $k = i - T + 1$. The protocol induces the following extensive-

[19] As in Brzustowski et al. (2023), that captures the idea that if a mechanism is replaced, it loses any stored information. We revisit this property in Section 6.1.

form game. First, the banker privately inputs a message $b_k \in B_k$ into the mechanism.[20] Second, the mechanism sends a private signal to the current depositor drawn from the probability distribution $\boldsymbol{\tau}_k(b_k|h_k) \in \Delta(\mathcal{T}_k)$. Third, the depositor inputs a private message $a_k \in A_k$ into the mechanism. Fourth, the mechanism generates an allocation for the current depositor, consisting of early consumption and a promise of late consumption drawn from the probability distributions $\boldsymbol{c}_k(b_k, a_k, \tau_k|h_k) \in \Delta([0, Y_k])$ and $\boldsymbol{\pi}_k(b_k, a_k, \tau_k|h_k) \in \Delta(\{0, 1\})$ respectively. Finally, the mechanism sends a signal to the banker drawn from the probability distribution $\boldsymbol{\beta}_k(b_k, a_k, \tau_k, c_k, \pi_k|h_k)$. The stage then ends. An illustration of a protocol is given in Figure 1.

Consider the following examples to see the breadth of mechanisms allowed in this environment. First, M can be designed to be transparent to the Banker, revealing full information about the interaction. Formally, $\boldsymbol{\beta}_k(b, \tau, a, c \mid h_k)$ outputs (τ, a, c) with probability 1, where (τ, a, c) is the triple of signal received by the depositor, message sent by the depositor, and realized allocation. Second, it can be designed to be relatively opaque for the depositors, that is, to reveal no information about the history, $\boldsymbol{\tau}_k(b_k \mid h_k)$ outputs $\bar{\tau}$ with probability 1, for some fixed signal $\bar{\tau} \in \mathcal{T}_k$, for all $b_k \in B_k$. Third, it can be designed to be opaque to the Banker, that is, to reveal nothing about the interaction except the realized allocation: $\boldsymbol{\beta}_k(b, a, \tau, c \mid h_k)$ outputs c with probability 1, for all inputs, $(b, a) \in B_k \times A_k$, of the Banker and the depositor; all signals $\tau \in \mathcal{T}_k$ received by the depositor; and all realized allocations $c \in [0, Y_k]$.

Mechanism instructions. We allow the private signal sent by a mechanism to a depositor $\tau_k \in \mathcal{T}_k$ to include instructions. For example, the mechanism can instruct the current depositor to select a message $a_k \in A_k$ if impatient and another message $a'_k \in A_k$ if patient. The depositors can freely ignore these instructions whose use will become apparent when we define our equilibrium.

Finite mechanisms. Finally, we must impose a technical assumption to circumvent stubborn difficulties in belief consistency and sequential equilibrium required for our solution concept (see, for example, Myerson and Reny (2020)). Specifically, we assume that banking mechanisms are finite. They have finite messages, finite signals, and finite possible allocations. We call the set of all finite mechanisms $\overline{\mathcal{M}}$ and assume $\mathcal{M}$ is a finite subset of $\overline{\mathcal{M}}$.

[20]In the spirit of contractual richness, we allow the Banker to input messages to the protocol at the beginning of each stage since she may wish to condition the mechanism's allocations on her information.

The restriction to finite mechanisms will not hinder the ability to achieve efficiency. Since there are finitely many possible histories, a finite mechanism can implement the efficient allocation.

3.3 Strategies and payoffs

Depositors. Before interacting with the mechanism deployed by the Banker, depositor i knows his payoff-type $\omega_i \in \{0,1\}$ and line position $i \in \{1,...,N\}$. In addition, he observes the current mechanism M, protocol $P_k(M)$, and the signal sent by the protocol $\tau_k(M)$. The information set of this depositor is then $I_i^d = (\omega_i, M, P_k(M), \tau_k(M))$. Denote by $\mathcal{I}_i^d(M)$ the set of all possible information sets such that depositor i interacts with mechanism M. A pure strategy for depositor i, denoted σ_i^d, assigns to each information set, $I_i^d \in \cup_{M\in\mathcal{M}}\mathcal{I}_i^d(M) \equiv \mathcal{I}_i^d$, a message, $a_k \in A_k(M)$, to the corresponding protocol. That is,

$$\sigma_i^d = \left\{\sigma_i^d(M)\right\}_{M\in\mathcal{M}} \qquad \text{s.t.} \qquad \sigma_i^d(M) : \mathcal{I}_i^d(M) \rightarrow A_k(M).$$

Let Σ_i^d denote the set of all pure strategies for depositor i and $\sigma^d = (\sigma_1^d, ..., \sigma_N^d) \in \Sigma_1^d \times ... \times \Sigma_N^d$ denote a profile of pure strategies for all depositors. Finally, let $\Sigma^d = \Sigma_1^d \times ... \times \Sigma_N^d$ denote the set of all pure strategy profiles for the depositors.

Banker. The Banker's information before interacting with depositor i denoted I_i^b contains the collection of signals she received from all mechanisms deployed in the past and the realized allocation at each stage. Denote by $\mathcal{I}_i^b$ the set of all possible information sets for the Banker before interacting with depositor i. The Banker's pure strategy for stage i assigns to each information set $I_i^b \in \mathcal{I}_i^b$ a choice of mechanism $M \in \mathcal{M}$, and a message, $b_k \in B_k(M)$, to be inputted into the current protocol $P_k(M)$ of the chosen mechanism. That is, for each $I_i^b \in \mathcal{I}_i^b$

$$\sigma_i^b(I_i^b) = (M, b_k) \quad \text{s.t. } M \in \mathcal{M} \text{ and } b_k \in B_k(M)$$

Denote by Σ_i^b the set of all pure strategies for the Banker in stage i. A pure strategy in the normal form of the game is a pure strategy for each stage: $\sigma^b = \left(\sigma_1^b, ..., \sigma_N^b\right) \in \Sigma_1^b \times ... \times \Sigma_N^b$. We denote by $\Sigma^b = \Sigma_1^b \times ... \times \Sigma_N^b$, the set of all pure strategies for the Banker.

Payoffs. The banking mechanisms can be designed to randomize over signals and allocations, and consequently, the payoffs of the depositors and the Banker depend on the randomness of each mechanism. To characterize payoffs concisely, we associate to each mechanism $M \in \mathcal{M}$ a probability space with a random variable describing all of the uncertainty in M. We then combine all those probability spaces together with the payoff types ω to build the state of the world $\theta \in \Theta$. The state θ resolves all uncertainty in the mechanism selection game. Given a profile of strategies $\sigma \in \Sigma^b \times \Sigma^d$, and a realization of the state $\theta \in \Theta$, the consumption bundle for each depositor is then uniquely determined. We denote the consumption bundle as a function of σ and θ by $\mathbf{c}_i(\sigma, \theta)$. Depositor i's expected payoff given his information set $I_i^d \in \cup_{M \in \mathcal{M}} \mathcal{I}_i^d(M)$ is given by

$$\mathcal{U}_i^d(\sigma | I_i^d) = \mathbb{E}_\theta \left[U\left(\mathbf{c}_i(\sigma, \theta), \omega_i(\theta)\right) \mid I_i^d \right]$$

where $\omega_i(\theta)$ corresponds to the payoff-type of depositor i in state θ. The expectation is with respect to the distribution over the state of the world θ given the depositor's information (note that $\{\omega_i(\theta)\} \in I_i^d$). The Banker's expected payoff given her information set $I_i^b \in \mathcal{I}_i$ is given by

$$\mathcal{U}_i^b(\sigma | I_i^b) = \mathbb{E}_\theta \left[W(\mathbf{c}(\sigma, \theta), \omega(\theta)) \mid I_i^b \right]$$

3.4 Equilibrium concept

To capture self-fulfilling bank runs, we will study a type of correlated equilibria of the mechanism selection game. We further insist on sequential rationality and impose belief consistency as in Kreps and Wilson (1982) to rule out 'unreasonable beliefs' off equilibrium.

Equilibrium definition. *An assessment (γ^*, μ^*) where $\gamma^* \in \Delta(\Sigma^b \times \Sigma^d)$ and μ^* is a system of beliefs, is an equilibrium if it satisfies the following:*

- **Sequential Rationality**: *$\sigma^d = (\sigma_i^d)_{i \leq N}$ is such that*

$$\sigma_i^d(I_i) \in \underset{\tilde{\sigma}_i^d \in \Sigma_i^d}{\operatorname{argmax}} \ \mathbb{E}^{\mu^*}_{(\sigma^d_{-i}, \sigma^b)} \Big[\mathcal{U}_i^d\big(\tilde{\sigma}_i^d, \sigma_{-i}^d, \sigma^b \mid I_i^d\big) \,\Big|\, \sigma_i^d \Big] \qquad \textit{for all } i \leq N \textit{ and } I_i^d \in \mathcal{I}_i^d \quad (6)$$

  and $\sigma^b = (\sigma_i^b)_{i \leq N}$ is such that:

$$\sigma_i^b(I_i^b) \in \underset{\tilde{\sigma}_i^b \in \Sigma_i^b}{\operatorname{argmax}} \ \mathbb{E}^{\mu^*}_{\sigma^d} \Big[\mathcal{U}_i^b\big(\sigma^d, \tilde{\sigma}_i^b \mid I_i^b\big) \,\Big|\, \sigma^b \Big] \qquad \textit{for all } k \leq N \textit{ and } I_i^b \in \mathcal{I}_i^b \qquad (7)$$

where both expectations are taken with respect to μ^* *and each player uses* γ^* *to predict others strategies.*

- **Consistency**: $\mu^* = \lim_n \mu_n$ *where* $(\mu_n)_n$ *is a sequence of belief-systems derived by Bayes' rule given a sequence* $(\gamma_n)_n$ *with* $\lim_n \gamma_n = \gamma^*$.

- **Tie-Breaking**: *If a depositor is indifferent between messages at some information set, he follows the mechanism instructions.*

As is typical, one can interpret this as follows. Nature draws a pure strategy profile $(\sigma^b, \sigma^d) \in \Sigma^b \times \Sigma^d$ according to γ^*, and then 'informs' each player about their individual draw: σ_i^d for depositor i, and σ^b for the Banker. We can view these draws as strategy recommendations. Sequential rationality necessitates that players wish to follow the recommendation given their beliefs. Consistency ensures that equilibrium outcomes are not supported by 'unreasonable' beliefs off-path. Given our restriction to finite mechanisms, an equilibrium exists.

Our solution concept is general enough to encompass previously used solution concepts. For instance, sequential equilibria can be recovered by insisting that $\gamma^* \in \Delta(\Sigma^b \times \Sigma^d)$ is independent over $\Sigma^b \times \Sigma_1^d \times \cdots \times \Sigma_N^d$. Our solution concept is also rich enough to capture all sorts of sunspot equilibria. In particular, it is equivalent to the following. An underlying sunspot state is realized and not observed by anyone. Instead, each depositor and the Banker receive a signal of the sunspot state, where the joint distribution of the signals conditional on the sunspot state can be arbitrary. As applied to the banking literature, the sunspot approach assumes that all depositors perfectly observe the sunspot (see, for example, Peck and Shell (2003)).

The tie-breaking rule ensures that any multiplicity does not arise purely because patient depositors are indifferent between messages. In existing work, runs occur because patient depositors strictly prefer to seek early consumption over late consumption, not because they are indifferent between running and not running. Consequently, adding this refinement does not change the equilibrium outcome in these previous studies. We, therefore, view indifference as an uninteresting source of multiplicity in the context of bank runs. For this reason, we insist that each depositor follows the mechanisms' instructions when indifferent between messages.

The tie-breaking rule can be rationalized by a model where the Banker can commit an ε amount of funds to break indifferences, with $\varepsilon > 0$ arbitrarily small. For example, message a leads to some level of early consumption c, whereas message a' leads to early consumption

$c - \varepsilon$ plus $R\varepsilon$ of late consumption. Then, any impatient depositor strictly prefers message a to a', whereas a patient depositor strictly prefers message a' to a. Importantly, such a subsidy will *not* eliminate runs in setups where the Banker can only use mechanisms asking the depositors to request early or late consumption.

4 Main result

We now show that *any* equilibrium of the mechanism selection game is efficient. Recall that efficiency requires that all patient depositors consume late and only impatient depositors consume early. Runs—defined as outcomes in which at least one patient depositor consumes early—do not occur in any equilibrium.

We prove our result using a simple mechanism M^* explained below. We say the contract space is sufficiently rich if the Banker can select this mechanism.

Theorem. *Suppose $M^* \in \mathcal{M}$. Then, the unique equilibrium outcome of the mechanism selection game is efficient.*

We proceed as follows. First, we define a mechanism M^* that will be shown always to collect useful information from the depositors. The mechanism uses that information to implement the continuation efficient allocation after any history of play (on or off the equilibrium path). We show how our mechanism deters patient depositors from seeking early consumption and the banker from replacing it.

4.1 The mechanism M^*

Consider a mechanism M^* with the following protocol $P_k(M^*)$ in stage i:

- The Banker can only input a single message into M^*.[21]

- Depositor i only receives instructions from M^* on how to play (see below) and must input message $a_k \in \{0, 1, \alpha\}$.

- If $a_k \in \{0, \alpha\}$, he is allocated immediate consumption c_k, which equals the continuation efficient level according to the mechanism's internal history. If $a_k = 1$, the depositor is given a promise of late consumption.

[21]Hence, the Banker cannot shape the behavior of mechanism M^* using her message.

- The instruction for each depositor is the following. If you are impatient, choose message $a_k = 0$. If you are patient but want to consume early, choose $a_k = \alpha$.
- The mechanism does not reveal the depositors' messages to the Banker.

Suppose M^* has been in place since stage $T \leq i$, and the current stage is i. Hence, it will use protocol $P_{i-T+1}(M^*)$ at the current stage. The mechanism has collected past messages, allocated immediate consumption, and promised future consumption to those depositors who have interacted with it. To construct the history of past types $\omega^{i-1} = (\omega_n)_{n=T}^{i-1}$, M^* records depositor n as impatient if $a_n = 0$, and as patient if $a_n \in \{\alpha, 1\}$. The mechanism uses the history of past types to compute the continuation efficient early payment for the current stage.

Also, observe that the depositors are only asked to provide information that they actually *know*, such as whether they are impatient or patient and whether they want early or late payment.

4.2 Properties of M^*

We now prove a key property of our mechanism M^*.

Lemma 1. *Pick an equilibrium (γ^*, μ^*), and a path along which M^* has been continuously deployed. Then, at each stage i, there is a common certainty between the Banker and depositor-i that the history of past types constructed by M^* is correct.*

Proof. See Appendix A.

By observing the protocol number, each depositor interacting with any mechanism $M \in \mathcal{M}$ can verify whether this mechanism has been deployed in all previous stages. Using this information, depositors make inferences about the mechanism's internal history and form beliefs about the actions of future depositors and the Banker. Lemma 1 states that in any equilibrium where M^* has been deployed in all past stages, all players agree that the mechanism's internal history corresponds to the game's history. As a result, all players agree that the level of early consumption offered at the current stage is continuation efficient. Specifically, by our tie-breaking rule, each depositor interacting with M^* reveals his type. That is, impatient depositors only select message $a_k = 0$, and patient depositors select messages $a_k \in \{\alpha, 1\}$. As a result, if M^* has been continuously deployed up to the current stage, it correctly sorts all past depositors into patient and impatient (regardless of the history of play) and uses that information to compute and then offer the continuation efficient early payment to those depositors inputting message 0 or α.

4.3 Proof of Theorem

We prove our theorem by induction. First, we show from first principles that M^* always attains the continuation efficient outcome at the final stage N when it has been deployed in all previous stages $i < N$. Based on this, we formulate our induction hypothesis: if M^* has been in place in all stages $n \leq i$, then continuation efficiency is achieved at stage $i + 1$. We then show that by selecting M^*, continuation efficiency is already achieved at stage n.

First, suppose M^* has been continuously deployed up to and including stage $n = N - 1$. By Lemma 1, it is common certainty that M^* has constructed its internal history correctly. Suppose M^* is also used at stage N. The early consumption set by M^* solves the program

$$W_N^* = \max_{c_N \in [0, Y_N]} u(c_N) + \sum_{i<N} \pi_i \cdot v(C)$$

where

$$C = \frac{R(Y^N - c_N)}{\sum_{i<N} \pi_i}$$

The first-order condition is $u'(c_N) = Rv'(C)$. By Assumption 1, we have that $c_N < C$ for each history. Hence, depositor N withdraws early if and only if he is impatient: $\sigma^*(I_N^d) = 0$ if $\omega_N = 0$ and $\sigma^*(I_N^d) = 1$ if $\omega_N = 1$. The Banker thus achieves the efficient continuation welfare W_N^*, which is an upper bound on the Banker's payoff after any history. The Banker can thus guarantee this upper bound by deploying M^*.

Based on this result, we formulate the following induction hypothesis: Suppose M^* has been deployed in all stages $n \leq i$. Then, continuation efficiency is achieved at stage $i + 1$.

We now prove that selecting M^* at stage i ensures that continuation efficiency is already achieved at stage i. Since M^* has been deployed in all stages 1 to $i - 1$, it is common certainty that M^*'s internal history is correct (Lemma 1). Hence, it is common certainty that M^* sets immediate consumption to the continuation efficient level. Note that continuation efficiency from stage i+1 onwards requires that patient depositors consume late and impatient depositors consume early. Further note that by Assumption 1, the continuation efficient allocation is incentive compatible.

Consider the strategy of depositor i when interacting with M^*. If impatient, he is indifferent between messages $\{0, \alpha\}$. Following M^*'s tie-breaking rule, depositor i selects 0. On the other hand, a patient depositor prefers to select message 1 over messages $\{0, \alpha\}$ given incentive compatibility. It follows that the Banker can guarantee the efficient continuation welfare W_i^* by selecting M^*, which is the upper bound. Hence, equilibrium welfare at stage i equals W_i^*.

Pushing the induction argument to stage 1 implies that welfare equals W_1^*. Hence, the unique equilibrium outcome is efficient, which completes the proof.

5 Ultimate source of fragility

Our mechanism M^* correctly identifies depositors as patient or impatient—even in the event of a run—and uses this information to re-optimize. Doing so deters patient depositors from withdrawing early and the Banker from replacing the mechanism. This result allows us to identify the ultimate source of fragility in Diamond-Dybvig models: banks are fragile if they cannot collect *and* respond to useful information during a run.

To illustrate this point, suppose the Banker is not using a mechanism that elicits depositors' types. Instead, we follow the standard exercise in the literature of deriving the Banker's best equilibrium payoff for a direct mechanism and then checking if there are multiple equilibria.

In particular, we will construct a run equilibrium in the spirit of Ennis and Keister (2009b). Suppose there are $N = 4$ depositors. There is no aggregate uncertainty for simplicity: 2 depositors are impatient, and the other 2 are patient. As before, depositor 1 is the first to contact the bank, depositor 2 is the second, and so on. The depositors know their line position, which is independent of their payoff type. There is a sunspot that is *bad* with a probability close to zero and *good* with the complement probability. All depositors observe the realized sunspot before contacting the bank.

We want to specify a run strategy profile for the depositors and show that it is consistent with equilibrium. Direct computation establishes that the last two depositors do not run in any equilibrium (Green and Lin, 2003; Ennis and Keister, 2009b). Thus, if there is a run, it will necessarily be *partial* and involve a critical 'last depositor to run' who anticipates that all patient depositors after him will not run but nevertheless prefer to run. For that to happen, the early consumption offered to that depositors cannot be continuation efficient but instead must be too high (otherwise, Assumption 1 implies this depositor will not run).

That is, the Banker must be unduly *optimistic* about the remaining number of impatient depositors. The example we present in this section has exactly this property.

Specifically, consider the following strategy profiles. If the sunspot is *good*, all depositors report their type truthfully. If the sunspot is *bad*, depositors 1 and 2 always report impatient. In contrast, depositors 3 and 4 continue to report truthfully. The Banker does not observe the sunspot state but makes inferences based on withdrawals. Since the probability of the bad sunspot is close to zero and at least two depositors will always request early payment, the early payment offered to depositors 1 and 2 will be almost the same.

The Banker infers a run is underway and adjusts consumption levels as soon as more than two depositors report impatient in period 0. However, this is not enough to avert a run. Suppose the sunspot state is bad and depositor 2 is patient. Hence, depositor 2 knows that depositor 1 withdrew even when patient. In contrast, the Banker assigns a high probability that depositor 1 withdrew only when impatient, given the low ex-ante probability of a bad sunspot. The depositors' and the Banker's *beliefs diverge*: depositor 2 assigns a higher probability than the Banker that depositors 3 and 4 are impatient. Consequently, the level of early consumption offered is too high, and the level of late consumption is reduced even further, given the number of impatient depositors waiting to be served. Depositor 2 thus runs. The same logic applies to depositor 1. Thus, depositors follow the suggested run strategy profile, leading to a run equilibrium.

Figure 2 compares the level of early consumption made by a direct mechanism and the mechanism M^* assuming that the first two depositors are patient and the sunspot is *bad*.[22] In particular, the first two depositors are patient but report 'impatient' under a direct mechanism, whereas they report 'patient but want early payment' under the mechanism M^*.

The direct mechanism sets the level of early consumption to about 1.07 for each of the first two depositors. Then, when depositor 3 reports impatient, the Banker infers that a run is underway and cuts the early consumption to the remaining two depositors to about 0.97. However, this is not enough to deter a run. *Limited commitment* is important here: a strict deposit freeze after two early withdrawals would make it a strictly dominant strategy for each patient depositor to withdraw late. However, freezing deposits is time-inconsistent during a run since there will be impatient depositors left, and giving them zero consumption is not ex-post efficient.

In contrast, the mechanism M^* learns that a run is underway after depositor 1. The mechanism has the correct beliefs about future payoff types and sets early consumption to

[22] Depositor's preferences are $U(\mathbf{c}_i, \omega_i) = \frac{(c_i+\omega_i C_i)^{1-\zeta}}{1-\zeta}$. We set $R = 1.2$, $Y = 4$, and $\zeta = 4$.

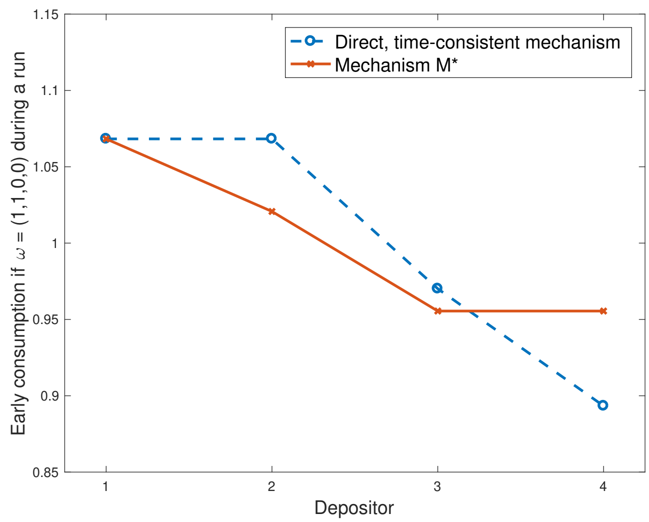

Figure 2: Illustration: Early consumption for two different mechanisms.

the continuation efficient level. Depositor 2 then prefers not to run. In turn, depositor 1 also prefers not to run. Thus, depositors do not follow the suggested run strategy profile for M^*.

Tie-breaking rule. This example also highlights the role of our tie-breaking assumption. If depositors do not follow the mechanisms' instructions and patient but running depositors select message 0 rather than α, then M^* would fail to elicit types. Beliefs between depositors and the mechanism diverge. Runs can then also occur under M^*, and the equilibrium outcome of the mechanism selection game is not necessarily efficient. However, any run equilibrium is *unstable*: whenever the Banker can commit to an infinitesimal subsidy paid in period 1 to each depositor who selected message α in period 0, then running but patient depositors strictly prefer to select message α over message 0. Notice that an infinitesimal subsidy would *not* eliminate runs under direct mechanisms since depositors 1 and 2 *strictly* prefer to run.

Revisiting Green and Lin (2003). Connecting our results to those in Green and Lin will be useful to highlight the importance of Lemma 1. Recall that the payoff type of each depositor in Green and Lin (2003) is independently and identically distributed: each depositor is impatient with probability $\lambda \in [0, 1]$ (see Section 7 for more details). They show that the optimal direct mechanism with commitment strongly implements the efficient allocation. That is, there are no bank-run equilibria under full commitment.

But what if, as in our case, the Banker cannot commit not to replace the current mechanism in the future? As it turns out, the Green and Lin set-up is the one case where Lemma 1 applies under direct mechanisms, which only ask each depositor to report 'impatient' or 'patient.' The reason is that the probability distribution over future types is independent of the past reports. The mechanism must only know the current stage to compute the continuation efficient allocation after any history (including those with a run). Stated differently, the direct mechanism of Green and Lin (2003) *is* time-consistent. However, as soon as depositors' payoff types are dependent (as in the example we presented), mechanisms only asking them to report their type would fail to correctly sort past depositors into patient and impatient in the event of a run and, as a result, would fail to assign the continuation efficient allocation.

6 Discussion

6.1 Relation to smart contracts

The mechanism selection game is characterized by limited commitment: the Banker has intra-stage but not inter-stage commitment. We now argue that such limited commitment power is attained by recent technological advancements, particularly smart contracts.

Smart Contracts are protocols that specify the contracting parties' actions and determine allocations given these actions. Several properties commonly identify them. The first is *automatic execution*. With automatic execution, allocations can be implemented without needing final approval by the transacting parties. In our case, this means that once the allocation within a stage has been determined, the Banker cannot renege. The second property is *cryptographic security*: information communicated to the protocol is encrypted. This prevents parties from accessing it even if they can access the contract's code. Further, if the mechanism is discarded, any information stored within it is lost.

In our context, the combination of automatic execution and cryptographic security means that the mechanisms can implement allocations contingent on depositors' messages while not revealing these messages to the Banker. Recall that the mechanism M^* *does not* reveal depositors' messages to the Banker. This property was irrelevant when the Banker was assumed to have intra-stage commitment. Let us now weaken the Banker's intra-stage commitment power by allowing her to 'restart' the current interaction before the allocation for that stage has been determined.

In particular, we modify the interaction within a protocol in Figure 1 by allowing the Banker to intervene before the current interaction has concluded. (i) The Banker privately submits a message to the protocol. (ii) The protocol sends a private signal to the current depositor. (iii) The current depositor then submits a private message to the protocol. (iv) The Banker can 'restart' the current interaction by discarding the current mechanism and deploying a new mechanism. (v) If the current interaction is not aborted, the protocol determines the allocation to the current depositor: either some immediate consumption or a promise of future consumption and a private signal to the Banker.

Observe that the Banker's beliefs are the same before and after the depositor inputs a message (but before the allocation has been determined, at which point it is too late to interfere in the current stage). This is the case since the Banker does not observe the actual message and does not learn anything she did not already know. Thus, if the Banker wanted to change the current mechanism at point (iv), she might as well have done that *before* the interaction started.[23] However, we already showed that the Banker is deterred from replacing M^* after one interaction concludes and before the next begins. The analysis is then unchanged from before, and the unique equilibrium outcome of the mechanism selection game continues to be efficient.

In contrast, suppose the depositors' messages are not encrypted, as will necessarily be the case if the depositor has to interact directly with the Banker (as when there is no automatic execution). Then, after observing message α (patient but wants early payment), the Banker wants to renege and promise future payment. Patient depositors recognize that and select message 0 when they want early payment, which precludes the mechanism from always collecting useful information (i.e., Lemma 1 fails) and making runs possible.

In that sense, we interpret the option of using automatic execution and cryptographic security as expanding the contracting space by allowing the Banker to pre-commit to follow the 'rules of the game' at least for the current interaction. The above logic potentially applies to any dynamic contracting where the principal posts a contract with automatic execution before the start of each stage. What is special in our case is that automatic execution is sufficient for the design of run-proof financial arrangements that still provide the efficient level of maturity transformation (as long as Assumption 1 on the primitives of the environment holds).

[23]Note that she would strictly prefer to change the mechanism before the interaction begins if there is a (small) cost to changing the mechanism once the interaction has started.

6.2 *Forcing to contact the mechanism*

The baseline setup required each depositor to contact the mechanism deployed by the Banker in period 0 as in Green and Lin (2003). This differs from Peck and Shell (2003), where only those who seek early consumption contact the bank. What if forcing to contact the mechanism in period 0 is a *design choice* rather than a feature of the environment? Notice that the Banker reduces the informational frictions about aggregate liquidity needs when forcing all depositors to contact in period 0 since a depositor reporting patient carries information about the types of the remaining depositors (when types are correlated). The Banker thus wants to generate information during a run *and* in normal times. Our analysis follows through as before.

One could instead envisage a scenario in which it is inefficient or infeasible to force patient depositors to contact the bank, possibly due to the cost of interacting with the bank as in Ennis and Keister (2016).[24] First, Lemma 1 applies in that case as well: each depositor interacting with M^* reveals his type. Moreover, running necessarily requires contacting the bank in period 0. What changes is the efficient allocation in Section 2 since now the Banker *only* observes the payoff types of those depositors choosing to withdraw. In particular, the k-th depositor who arrives to withdraw is not necessarily the k-th depositor with an opportunity to withdraw since some patient depositors with an earlier opportunity may have decided to wait (whereas, in the baseline case, the k-th depositor to arrive is always the k-th depositor with an opportunity to withdraw).

Nevertheless, as long as the resulting continuation efficient allocation after any history is incentive-compatible also in this informational environment (i.e., satisfies Assumption 1), our results apply as in the baseline case.

6.3 *Information about the line*

We assumed that depositors learn their line position before withdrawals begin (as in Green and Lin (2000) and others). Since our mechanism relies on backward induction logic, it might seem that it will not work if we relax this assumption. However, we now show that our mechanism still holds even if depositors have no information (as in Green and Lin (2003)) or only partial information (as in Peck and Shell (2003)) about their position in the withdrawal order. The idea is that, through a proper design, depositors can infer their line

[24] Depositors who do not contact the bank in period 0 remain in the coalition of depositors.

position by observing the protocol number of the mechanism they interact with.

Specifically, consider a depositor interacting with protocol k in some mechanism M. This depositor only knows that the current stage is greater than or equal to k and thus cannot infer that M has been in place for all past interactions or that he is the k-th depositor in the withdrawal order. Hence, we cannot apply Lemma 1 and use the backward induction argument from the previous section. So, consider the following modification of the environment: when the coalition of depositors is initially formed (before any depositor-specific information has been realized), the Banker deploys a mechanism that each depositor gets to observe. Call this the *default mechanism.* Moreover, when a depositor subsequently arrives to withdraw, he can verify (i) whether the current is the default mechanism and (ii) whether the default mechanism has been deployed for all past depositors.[25]

Now, suppose the Banker selects M^* as the default mechanism. Lemma 1 applies, and our analysis follows through verbatim.

7 Incentive constraints

When is Assumption 1 on the continuation incentive constraints satisfied? To answer this question, we adopt the canonical Diamond and Dybvig (1983) preferences, summarize existing results in the literature and also establish a new result in Proposition 1.[26] Specifically, the utility of depositor i is given by

$$U(\mathbf{c}_i, \omega_i) \;=\; \begin{cases} u(c_i) & \text{if } \;\omega_i = 0 \\ \rho u(c_i + C_i) & \text{if } \;\omega_i = 1 \end{cases}$$

where $u : \mathbb{R}_+ \to \mathbb{R}$, $\rho \in (0, 1]$, and $R\rho > 1$. The function u is strictly increasing, twice continuously differentiable, strictly concave, and satisfies the Inada conditions where $\lim_{x\to 0} u'(x) = \infty$ and $\lim_{x\to\infty} u'(x) = 0$. Further, u features a degree of relative risk aversion weakly greater than 1 for any $x \in \mathbb{R}_+$. That is, $-\frac{xu''(x)}{u'(x)} \geq 1$. Also, recall that $p : \{0, 1, ..., N\} \to [0, 1]$ is an exogenous probability mass function such that $p(\phi)$ is the probability that the number of patient depositors is $\phi \in \{0, 1, ..., N\}$. Then, Assumption 1 is satisfied if one of the following is true.

[25] That is, the Banker cannot cheat by discarding the default mechanism - in favor of another mechanism - for some of the previous stages and then bringing it back for the current stage. The depositors would detect such a deviation (such verification can be achieved through modern cryptography).

[26] The importance of continuation incentive constraints was first recognized by Green and Lin (2003) and further analyzed by Andolfatto et al. (2007a), Ennis and Keister (2009b) and Huang (2023).

(i) Aggregate certainty: p is degenerate at a point $\phi \in \{0, 1, ..., N\}$. Equivalently, the Banker observes the realized value of ϕ before withdrawals begin.

(ii) Independent types: p is given by the binomial distribution, where $\lambda \in [0, 1]$ is the probability that each depositor is impatient:

$$p(\phi) \;=\; \binom{N}{\phi} \lambda^{\phi} (1-\lambda)^{1-\phi},$$

(iii) Positively correlated types: p satisfies the monotonicity condition

$$\frac{\mathbb{P}(\omega \,|\, \phi(\omega) = \kappa)}{\mathbb{P}(\omega \,|\, \phi(\omega) = \kappa + 1)} \;\geq\; \frac{\mathbb{P}(\omega \,|\, \phi(\omega) = \kappa + 1)}{\mathbb{P}(\omega \,|\, \phi(\omega) = \kappa + 2)}$$

where $\phi(\omega)$ denotes the number of patient depositors for a given realization of the payoff types.

Case (i) follows from direct computation. Consider a history for which the bank has $Y > 0$ goods left. The number of depositors remaining in the coalition of depositors is N, of which η depositors are impatient. The efficient allocation (c^*, C^*) solves the program

$$\max_{(c,C)} \;\eta \cdot u(c) \;+\; (N - \eta) \cdot \rho u(C)$$

$$\text{s.t.} \quad \eta \cdot c \;+\; (N - \eta) \cdot \frac{C}{R} \;\leq\; Y$$

The first-order condition is given by $u'(c^*) = \rho R u'(C^*)$. Since $\rho R > 1$, it implies $c^* < C^*$ for all histories.

Cases (ii) and (iii), in addition, assume u has absolute risk aversion, which is non-increasing everywhere: $\frac{d}{dc}\frac{u''(c)}{u'(c)} \geq 0$ for any $c \in \mathbb{R}_+$. Case (ii) was shown by Green and Lin (2003) whereas case (iii) was established by Huang (2023). Next, we establish the following result, allowing for an arbitrary correlation in types.

Proposition 1. *Suppose* $U(\mathbf{c}_i, \omega_i) = \frac{1}{1-\zeta}\big(c_i + \omega_i C_i\big)^{1-\zeta}$ *with* $\zeta > 1$. *Then, Assumption 1 is satisfied.*

Proof. See Appendix B.

This proposition is of independent interest since it answers an open question in the literature by providing conditions for the incentive constraints to hold under arbitrary correlation in payoff types.

Finally, we should remark it is possible to construct environments in which the continuation incentive constraints are not satisfied. This is most easily illustrated for case (i), with $R > 1$ and $\rho < 1$ such that $\rho R < 1$. The efficient allocation then features $c^* > C^*$, which is not incentive-compatible. To the best of our knowledge, the only examples in the literature in which the incentive constraints are not satisfied feature $\rho R < 1$ and are given by Peck and Shell (2003) and Sultanum (2022).

8 Conclusion

We revisit the Diamond-Dybvig model of bank runs in a general mechanism design framework with limited commitment. We present a simple mechanism resting on minimal properties of smart contracts, which achieves efficiency in all equilibria—as long as some weak assumptions on the primitives of the environment are satisfied.

Our results highlight the inability to collect and respond to useful information during a run as the ultimate source of financial fragility. Banks are not inherently fragile because they engage in maturity transformation or because they cannot commit to a plan of action during crises.

If banks are not inherently fragile, then what causes runs? Sources of fragility outside of the model may include depositors that are not fully rational or government guarantees that are distortionary and push intermediaries towards liability structures that are not run-proof. Runs may also be a phenomenon of insolvency rather than illiquidity or an outcome of institutional or political factors preventing financial intermediaries from adopting a run-proof liability structure.

References

Allen, F. and D. Gale (1998): "Optimal financial crises," *The journal of finance*, 53, 1245–1284.

——— (2004a): "Competition and financial stability," *Journal of money, credit and banking*, 453–480.

——— (2004b): "Financial Intermediaries and Markets," *Econometrica*, 72, 1023–1061.

Andolfatto, D., E. Nosal, and B. Sultanum (2017): "Preventing Bank Runs," *Theoretical Economics*, 12, 1003–1028.

Andolfatto, D., E. Nosal, and N. Wallace (2007a): "The role of independence in the Green–Lin Diamond–Dybvig model," *Journal of Economic Theory*, 137, 709–715.

——— (2007b): "The Role of Independence in the Green–Lin Diamond–Dybvig model," *Journal of Economic Theory*, 137, 709–715.

Bakos, Y. and H. Halaburda (2020): "Platform competition with multihoming on both sides: Subsidize or not?" *Management Science*, 66, 5599–5607.

Beccuti, J. and M. Möller (2018): "Dynamic adverse selection with a patient seller," *Journal of economic theory*, 173, 95–117.

Bianchi, J. (2016): "Efficient bailouts?" *American Economic Review*, 106, 3607–3659.

Bruner, R. F. and S. D. Carr (2009): *The panic of 1907: Lessons learned from the market's perfect storm*, John Wiley & Sons.

Bryant, J. (1980): "A model of reserves, bank runs, and deposit insurance," *Journal of Banking & Finance*, 4, 335–344.

Brzustowski, T., A. Georgiadis-Harris, and B. Szentes (2023): "Smart Contracts and the Coase Conjecture," *American Economic Review*, 113, 1334–1359.

Cavalcanti, R. d. O. and P. K. Monteiro (2016): "Enriching information to prevent bank runs," *Economic Theory*, 62, 477–494.

Chari, V. V. and P. J. Kehoe (2016): "Bailouts, time inconsistency, and optimal regulation: A macroeconomic view," *American Economic Review*, 106, 2458–2493.

Cong, L. W. and Z. He (2019): "Blockchain disruption and smart contracts," *The Review of Financial Studies*, 32, 1754–1797.

Cooper, R. and T. W. Ross (1998): "Bank runs: Liquidity costs and investment distortions," *Journal of Monetary Economics*, 41, 27–38.

Dávila, E. and A. Walther (2020): "Does size matter? Bailouts with large and small banks," *Journal of Financial Economics*, 136, 1–22.

Diamond, D. W. and P. H. Dybvig (1983): "Bank Runs, Deposit Insurance, and Liquidity," *Journal of Political Economy*, 91, 401–419.

Doval, L. and V. Skreta (2022): "Mechanism Design with Limited Commitment," *Econometrica*, 90, 1463–1500.

Ennis, H. M. and T. Keister (2009a): "Bank Runs and Institutions: The Perils of Intervention," *American Economic Review*, 99, 1588–1607.

——— (2009b): "Run Equilibria in the Green–Lin Model of Financial Intermediation," *Journal of Economic Theory*, 144, 1996–2020.

——— (2010): "Banking Panics and Policy Responses," *Journal of Monetary Economics*, 57, 404–419.

——— (2016): "Optimal banking contracts and financial fragility," *Economic Theory*, 61, 335–363.

Farhi, E., M. Golosov, and A. Tsyvinski (2009): "A Theory of Liquidity and Regulation of Financial Intermediation," *The Review of Economic Studies*, 76, 973–992.

Farhi, E. and J. Tirole (2012): "Collective moral hazard, maturity mismatch, and systemic bailouts," *American Economic Review*, 102, 60–93.

Goldstein, I. and A. Pauzner (2005): "Demand–deposit contracts and the probability of bank runs," *the Journal of Finance*, 60, 1293–1327.

Green, E. J. and P. Lin (2000): "Diamond and Dybvig's classic theory of financial intermediation: what's missing?" *Federal Reserve Bank of Minneapolis Quarterly Review*, 24, 3–13.

——— (2003): "Implementing Efficient Allocations in a Model of Financial Intermediation," *Journal of Economic Theory*, 109, 1–23.

Holden, R. and A. Malani (2021): *Can blockchain solve the hold-up problem in contracts?*, Cambridge University Press.

Huang, X. (2023): "Sophisticated Banking Contracts and Fragility when Withdrawal Information is Public," *Theoretical Economics (forthcoming).*

Jacklin, C. J. (1987): "Demand Deposits, Trading Restrictions, and Risk Sharing," *Contractual Arrangements for Intertemporal Trade (Minneapolis: University of Minnesota Press)*, 26–47.

Keister, T. (2016): "Bailouts and Financial Fragility," *The Review of Economic Studies*, 83, 704–736.

Keister, T. and Y. Mitkov (2023): "Allocating losses: Bail-ins, Bailouts and Bank Regulation," *Journal of Economic Theory (forthcoming).*

Kinateder, M. and H. J. Kiss (2014): "Sequential decisions in the Diamond–Dybvig banking model," *Journal of Financial Stability*, 15, 149–160.

Kindleberger, C. P., R. Z. Aliber, and R. M. Solow (2005): *Manias, panics, and crashes: A history of financial crises*, vol. 7, Palgrave Macmillan London.

Kreps, D. M. and R. Wilson (1982): "Sequential equilibria," *Econometrica: Journal of the Econometric Society*, 863–894.

Laffont, J.-J. and J. Tirole (1988): "The dynamics of incentive contracts," *Econometrica: Journal of the Econometric Society*, 1153–1175.

——— (1990): "Adverse selection and renegotiation in procurement," *The Review of Economic Studies*, 57, 597–625.

Myerson, R. B. and P. J. Reny (2020): "Perfect Conditional ε-Equilibria of Multi-Stage Games With Infinite Sets of Signals and Actions," *Econometrica*, 88, 495–531.

Nosal, E. and N. Wallace (2009a): "Information revelation in the Diamond-Dybvig banking model," *Federal Reserve Bank of Chicago Policy Discussion Paper Series (December).*

——— (2009b): "Information Revelation in the Diamond-Dybvig Banking Model," *Federal Reserve Bank of Chicago Policy Discussion Paper Series (December).*

Nosal, J. B. and G. Ordonez (2016): "Uncertainty as commitment," *Journal of Monetary Economics*, 80, 124–140.

Peck, J. and A. Setayesh (2023): "Bank runs and the optimality of limited banking," *Review of Economic Dynamics*, 47, 100–110.

Peck, J. and K. Shell (2003): "Equilibrium Bank Runs," *Journal of Political Economy*, 111, 103–123.

Philippon, T. and O. Wang (2023): "Let the worst one fail: a credible solution to the too-big-to-fail conundrum," *The Quarterly Journal of Economics*, 138, 1233–1271.

Postlewaite, A. and X. Vives (1987): "Bank runs as an equilibrium phenomenon," *Journal of political Economy*, 95, 485–491.

Sultanum, B. (2022): "Discussion of "currency stability using blockchain technology"," *Journal of Economic Dynamics and Control*, 142, 104156.

Tinn, K. (2018): "'Smart'Contracts and External Financing," *Available at SSRN 3072854*.

Tirole, J. (2016): "From bottom of the barrel to cream of the crop: Sequential screening with positive selection," *Econometrica*, 84, 1291–1343.

Wallace, N. (1988): "Another Attempt to Explain an Illiquid Banking System: With Sequential Service Taken Seriously," *Federal Reserve Bank Minneapolis Quarterly Review 12*, 3–16.

——— (1990): "A banking model in which partial suspension is best," *Federal Reserve Bank of Minneapolis. Quarterly Review-Federal Reserve Bank of Minneapolis*, 14, 11.

Appendices

A Proof of Lemma 1

Proof. Pick an arbitrary equilibrium (γ^*, μ^*) and a path along which only M^* has been in place. Let $A_k = [\omega_k = 0, a_k \in \{1, \alpha\}] \cup [\omega_k = 1, a_k = 0]$ be the event that M^* has the wrong information about Depositor-k's payoff-type.

Now consider the information set I_n^b of the Banker at stage n along such a path. Since she does not observe anything about the history (apart from the fact that M^* has been in place), this information set's nodes correspond exactly to the possible sequences of types and messages within M^*.[27] Now let $A = \cup_{k<n} A_k$ be the set of nodes where the mechanism has some wrong information. We will show that $\mu^*(A) = 0$.

To this end, by belief consistency there exists some sequence $(\gamma^m)_{m\in\mathbb{N}}$ of completely mixed strategies such that $\lim_m \gamma^m = \gamma^*$, and $\mu^* = \lim_m \mu^m$, where $\mu^m(d) = \frac{\mathbb{P}^{\gamma^m}(d)}{\mathbb{P}^{\gamma^m}(I(d))}$, for each decision node d, and each information set $I(d)$.

Recall that by our tie-breaking assumption, $\gamma^*(A) = 0$.[28] Since $(\gamma^m)_{m\in\mathbb{N}}$ converges to γ^*, we must have that for any $\varepsilon > 0$, there exists $N_\varepsilon \in \mathbb{N}$ such that $\gamma^m(A) < \varepsilon$ for all $m \geq N_\varepsilon$. Thus, by Bayes' rule

$$\frac{\mu^m(A)}{\mu^m(A^c)} = \frac{\gamma^m(A)}{\gamma^m(A^c)} < \frac{\varepsilon}{1-\varepsilon}$$

Consequently, $\lim_m \frac{\mu^m(A)}{\mu^m(A^c)} = 0$ which implies $\mu^*(A) = \lim_m \mu^m(A) = 0$. That is, the Banker is certain that M^* has the correct information.

Note that Depositor-n learns nothing more than the Banker about the history at his succeeding information set I_n, even after M^* is kept in place. Moreover, whether or not the mechanism has the correct information is independent of Depositor-n's private information. Hence, $\mu^*(A) = 0$ is sufficient to conclude that Depositor-n is also certain that the mechanism has the correct information. Finally, the assertion about common certainty follows from the fact that the assessment itself (γ^*, μ^*) is common certainty. □

B Proof of Proposition 1

Suppose that $U(\mathbf{c}_i, \omega_i) = \frac{(c_i+\omega_i C_i)^{1-\zeta}}{1-\zeta}$ for all i, with $\zeta > 1$. Consider some history $h^i = (Y^i, \omega^i, \pi^i)$. Let $\Phi^i = \sum_{n<i} \omega_n$ and $\Pi^i = \sum_{n<i} \pi_n$. Lemma 2 below characterizes the continuation efficient allocation and is a restatement of Proposition 1 in Ennis and Keister (2009b):

Lemma 2. *The continuation efficient allocation sets*

$$c_i^* = \frac{Y^i}{\psi_i(\Phi^i, \Pi^i)^{\frac{1}{\zeta}} + 1} \qquad \text{(IB.1)}$$

[27]Allocations can be derived with certainty from those.

[28]$\gamma^*(\cdot)$ is the probability measure induced by the equilibrium strategy over the set of decision nodes.

where

$$\psi_i(x_1,x_2) \;=\; \mathbb{P}\left[\omega_{i+1}=0|x_1\right]\cdot\left(\psi_{i+1}(x_1,x_2)^{\frac{1}{\zeta}}+1\right)^{\zeta}+\mathbb{P}\left[\omega_{i+1}=1|x_1\right]\cdot\psi_{i+1}(x_1+1,x_2+1) \tag{IB.2}$$

for $i=\{1,...,N-1\}$ *and* $\psi_N(x_1,x_2)=\left(x_2\cdot R^{\frac{1-\zeta}{\zeta}}\right)^{\zeta}$ *for* $i=N$.

For simplicity, we drop the explicit dependence of $\mathbf{c}^*$ on the payoff state ω. We now prove Proposition 1 using four Lemmas. First, we prove a useful property of ψ_i (Lemma 3). We use this property to show that the expected late consumption, conditional on any history, is below R times the early consumption offered to the current depositor after that history (Lemma 4). This is intuitive as it expresses the desire of the planner to perform some maturity transformation. Third, we show that early consumption increases after a patient report (Lemma 5). This is also intuitive as the Banker will be keen to do more maturity transformation when there are more patient depositors.

Using Lemma 4, we show that a virtual continuation incentive constraint holds (Lemma 6). That is the continuation 'incentive constraint' from the perspective of an *impatient* depositor is satisfied.

Finally, we complete the proof of Proposition 1 using an induction argument. To this end, we show that the continuation incentive constraint holds for the last depositor N. Then, we assume it holds for arbitrary depositor-$(n+1)$ and prove it holds for depositor-n. In this step, we break up the evaluation of the expected utility of late consumption of a patient depositor-n into a convex combination of expected utility of consumption of a future *impatient* depositor-$(n+1)$, and a future *patient* depositor-$(n+1)$. Then, Lemma 6 and the induction hypothesis imply that the continuation IC for patient depositor-n holds if early consumption increases after a patient report, which is given in Lemma 5.

Lemma 3. *For all i, we find that:*

$$\psi_i(\Phi+1,\Pi+1)^{\frac{1}{\zeta}}-\psi_i(\Phi,\Pi)^{\frac{1}{\zeta}}\;\leq\;1$$

Proof. The claim is true for $i=N$ by direct computation. Now, suppose that it is true for $i=m+1$. We will prove it for $i=m$. The induction hypothesis implies:

$$\left(\psi_{m+1}(\Phi,\Pi)^{\frac{1}{\zeta}}+1\right)^{\zeta}\;\geq\;\psi_{m+1}(\Phi+1,\Pi+1) \tag{IB.3}$$

and hence from the definition of $\psi_m(\Phi, \Pi)$:

$$\psi_{m+1}(\Phi+1, \Pi+1) \ \leq \ \psi_m(\Phi, \Pi) \ \leq \ \Big(\psi_{m+1}(\Phi, \Pi)^{\frac{1}{\zeta}} + 1\Big)^{\zeta} \tag{IB.4}$$

Therefore:

$$\psi_m(\Phi+1, \Pi+1)^{\frac{1}{\zeta}} - \psi_m(\Phi, \Pi)^{\frac{1}{\zeta}} \ \leq \ \Big(\psi_{m+1}(\Phi+1, \Pi+1)^{\frac{1}{\zeta}} + 1\Big) - \psi_{m+1}(\Phi+1, \Pi+1)^{\frac{1}{\zeta}} \ \leq \ 1$$

which proves the claim. □

Lemma 4. *For all i, and for all histories h^i,*

$$\mathbb{E}\left[C^*|h^i\right] \ \leq \ Rc_i^* \tag{IB.5}$$

Proof. For each i, for each $\Pi^i = \Pi$ and $\Phi^i = \Phi$, consider the difference equation:

$$\overline{\psi}_i(\Phi, \Pi)^{\frac{1}{\zeta}} - \overline{\zeta}_{i+1}(\Phi, \Pi)^{\frac{1}{\zeta}} \ = \ 1, \quad i = n, \ldots, N-1, \quad \text{and} \quad \overline{\psi}_N(\Phi, \Pi) \ = \ \Big(\Pi \cdot R^{\frac{1-\zeta}{\zeta}}\Big)^{\zeta} \tag{IB.6}$$

For a solution $\overline{\psi}$ to Equation (IB.6), we have:

$$\psi_i(\Phi, \Pi) \ \leq \ \overline{\psi}_i(\Phi, \Pi)$$

Indeed, from (IB.4) we have for each $i \leq N-1$,

$$\psi_i(\Phi, \Pi)^{\frac{1}{\zeta}} - \overline{\psi}_i(\Phi, \Pi)^{\frac{1}{\zeta}} \ \leq \ \psi_{i+1}(\Phi, \Pi)^{\frac{1}{\zeta}} - \overline{\psi}_{i+1}(\Phi, \Pi)^{\frac{1}{\zeta}}$$

and hence:

$$\psi_i(\Phi, \Pi)^{\frac{1}{\zeta}} - \overline{\psi}_i(\Phi, \Pi)^{\frac{1}{\zeta}} \ \leq \ \psi_N(\Phi, \Pi)^{\frac{1}{\zeta}} - \overline{\psi}_N(\Phi, \Pi)^{\frac{1}{\zeta}} \ = \ 0$$

To solve Equation (IB.6) we define the variable $f(n) = \overline{\psi}_n(\Phi, \Pi)^{\frac{1}{\zeta}}$. Then the equation writes:

$$f(n+1) - f(n) \ = \ -1, \qquad f(N) \ = \ \Big(\Pi \cdot R^{\frac{1-\psi}{\psi}}\Big)$$

Summing from i to $N-1$ we have:

$$f(N) - f(i) \;=\; \sum_{n=i}^{N-1} \Big(f(n+1) - f(n) \Big) \;=\; -(N-i)$$

So,

$$\overline{\psi}_i(\Phi, \Pi)^{\frac{1}{\zeta}} = N - i + \left(\Pi \cdot R^{\frac{1-\zeta}{\zeta}} \right)$$

From this we get:

$$c_i^* \;\geq\; \frac{Y^i}{N - i + 1 + \left(\Pi^i \cdot R^{\frac{1-\zeta}{\zeta}} \right)} \;>\; \frac{Y^i}{N - i + 1 + \Pi^i} \;=:\; \hat{c}_i \tag{IB.7}$$

where $\hat{c}_i$ is also a stochastic process adapted to h^i. The last inequality follows from $R^{\frac{1-\zeta}{\zeta}} < 1$. Next, we compute conditional expectations of late consumption. We have,

$$\mathbb{E}\left[C^* \mid h^i \right] \;=\; \mathbb{E}\left[R \cdot \frac{Y^i - \sum_{n=i}^{N} (1-\omega_n) \cdot c_n^*}{\Pi^i + \sum_{n=i}^{N} \omega_n} \;\Big|\; h^i \right]$$

From the bound $c_n^* > \hat{c}_n$ we have:

$$\mathbb{E}\left[R \cdot \frac{Y^i - \sum_{n=k}^{N} (1-\omega_n) \cdot c_n^*}{\Pi^i + \sum_{n=k}^{N} \omega_n} \;\Big|\; h^i \right] \;<\; \mathbb{E}\left[R \cdot \frac{Y^i - \sum_{n=k}^{N} (1-\omega_n) \cdot \hat{c}_n}{\Pi^i + \sum_{n=k}^{N} \omega_n} \;\Big|\; h^i \right]$$

We now prove that $\mathbb{P}\Big(\hat{c}_n = \hat{c}_i \mid h^i \Big) = 1$ for all $n \geq i$. Each time $\omega_n = 0$, we have:

$$\hat{c}_{n+1} \;=\; \frac{Y^n - \hat{c}_n}{N - (n+1) + 1 + \Pi^n} \;=\; \frac{Y^n - \frac{Y^n}{N-n+1+\Pi^n}}{N - n + \Pi^n} \;=\; \frac{Y^n \cdot \frac{N-n+\Pi^n}{N-n+1+\Pi^n}}{N - n + \Pi^n} \;=\; \hat{c}_n$$

Similarly, each time $\omega_n = 1$, we have:

$$\hat{c}_{n+1} \;=\; \frac{Y^n}{N - (n+1) + 1 + (\Pi^n + 1)} \;=\; \frac{Y^n}{N - n + 1 + \Pi^n} \;=\; \hat{c}_n$$

That is, $\hat{c}_n$ is constant with probability 1 along any path. Hence:

$$\begin{aligned}
\mathbb{E}\left[R \cdot \frac{Y^i - \sum_{n=i}^N (1-\omega_n)\cdot \hat{c}_n}{\Pi^i + \sum_{n=i}^N \omega_n} \;\Big|\; h^i\right] &= \mathbb{E}\left[R \cdot \frac{Y^i - \hat{c}_i \cdot \sum_{n=i}^N (1-\omega_n)}{\Pi^i + \sum_{n=i}^N \omega_n} \;\Big|\; h^i\right] \\
&= \mathbb{E}\left[R \cdot \frac{Y^i}{N-i+1+\Pi^i} \;\Big|\; h^i\right] \\
&= R \cdot \hat{c}_i \\
&\leq R \cdot c_i^*
\end{aligned}$$

This completes the proof for this lemma. □

Lemma 5. *Consider two histories* $h^i = (Y^i, \omega^i, \pi^i)$ *and* $h^{i+1} = (Y^{i+1}, \omega^{i+1}, \pi^{i+1})$*, where* $Y^{i+1} = Y^i = Y$*,* $\Phi^{i+1} = \Phi^i + 1$ *and* $\Pi^{i+1} = \Pi^i + 1$*. Then*

$$c_{i+1}^* \;\geq\; c_i^*$$

Proof. First, note that the claim is true if and only if:

$$\psi_i(\Phi, \Pi) \;\geq\; \psi_{i+1}(\Phi+1, \Pi+1), \quad \text{for all } i \text{ and } \Pi$$

We write:

$$\psi_i(\Phi,\Pi) - \psi_{i+1}(\Phi+1,\Pi+1) \;=\; \mathbb{P}\left[\omega_{i+1} = 0 | h^i\right] \cdot \left[\left(\psi_{k+1}(\Phi,\Pi)^{\frac{1}{\zeta}} + 1\right)^{\zeta} - \psi_{k+1}(\Phi+1, \Pi+1)\right]$$

From Lemma 3 we have for any i and Π,

$$\left(\psi_{i+1}(\Phi,\Pi)^{\frac{1}{\zeta}} + 1\right)^{\psi} - \psi_{i+1}(\Phi+1,\Pi+1) \geq 0$$

which proves the claim. □

Lemma 6 (Virtual IC)**.** *For each* i*, and each history* h^i*,*

$$u(c_i^*) \;<\; \mathbb{E}\left[u(C^*) \mid h^i, \omega_i = 0\right]$$

Proof. Let $\varphi(c) = \frac{u(c)}{u'(c)}$, and note that CRRA with $\zeta > 1$ implies that $\varphi(c) = \frac{1}{1-\zeta} \cdot c$. That is, $\varphi(c)$ is strictly decreasing and (weakly) convex, with $\varphi(0) = 0$. Recall that the FOC determining c_i^* can be written as[29]

$$u'(c_i^*) \;=\; R \cdot \mathbb{E}\left[u'(C^*) \;\middle|\; h^i, \omega_i = 0\right] \;=:\; R \cdot \mathbb{E}_i^0\left[u'(C^*)\right]$$

where for notational simplicity in what follows, we suppress the conditioning arguments. We have:

$$\begin{aligned}
\mathbb{E}_i^0\left[u(C^*)\right] \;&=\; \mathbb{E}_i^0\Big[u'(C^*)\varphi(C^*)\Big] \\
&=\; \mathbb{E}_i^0\Big[u'(C^*)\Big] \cdot \mathbb{E}_i^0\Big[\varphi(C^*)\Big] + \mathrm{Cov}_i^0\Big[u'(C^*), \varphi(C^*)\Big] \\
&>\; \mathbb{E}_i^0\Big[u'(C^*)\Big] \cdot \mathbb{E}_i^0\Big[\varphi(C^*)\Big] \\
&=\; \frac{1}{R} u'(c_i^*) \cdot \mathbb{E}_i^0\Big[\varphi(C^*)\Big]
\end{aligned}$$

where the first inequality follows from the fact that u' and φ are both strictly decreasing functions making their covariance positive; and the last equality follows from the FOC.

Now observe that by convexity of φ,

$$\mathbb{E}_i^0\big[\varphi(C^*)\big] \;\geq\; \varphi\left(\mathbb{E}_i^0\big[C^*\big]\right)$$

From Lemma 4, we have $\mathbb{E}_i^0\big[C^*\big] \leq R \cdot c_i^*$, and thus because φ is decreasing,

$$\varphi\Big(\mathbb{E}_i^0\big[C^*\big]\Big) \;\geq\; \varphi(R \cdot c_i^*)$$

Using again the fact that φ is convex with $\varphi(0) = 0$,

$$\varphi(R \cdot c_i^*) \;\geq\; R \cdot \varphi(c_i^*)$$

[29] This follows from the Bellman equation. First, the FOC sets the marginal utility of early consumption equal to the expectation of the derivative of the value function with respect to resources, conditional on the current trader being impatient. One then establishes by the Envelope Theorem that the derivative of the value function is a martingale and hence equal to the expectation of its terminal value, conditional on the current trader being impatient. The terminal value is simply the derivative of the utility of terminal consumption with respect to resources—hence the formula.

since $R > 1$. Finally, putting all these together we conclude that:

$$\mathbb{E}_i^0\left[u(C^*)\right] \;>\; \frac{1}{R}u'(c_i^*) \cdot \mathbb{E}_i^0\big[\varphi(C^*)\big] \;\geq\; u'(c_i^*) \cdot \varphi(c_i^*) \;=\; u(c_i^*)$$

which proves the claim. □

We are now ready to complete the proof. To economize on notation, we write:

$$\mathbb{E}_i^1\left[u(C^*)\right] \;:=\; \mathbb{E}\left[u(C^*) \;\middle|\; h^i, \omega_i = 1\right]$$

We prove the claim by induction. It is easy to show that for depositor-N:

$$\mathbb{E}_N^1\big[u(C^*) - u(c_N^*)\big] \;>\; 0$$

We now assume that the continuation IC holds for arbitrary $n+1$:

$$\mathbb{E}_{n+1}^1\left[u(C^*) - u(c_{n+1}^*)\right] \;>\; 0 \quad \text{for all } h^{n+1}$$

We have:

$$\mathbb{E}_n^1\big[u(C^*) - u(c_n^*)\big] \;=\; \mathbb{E}_n^1\left[u(C^*) - u(c_{n+1}^*) + \left\{u(c_{n+1}^*) - u(c_n^*)\right\}\right]$$

where, as above, c_n^* and c_{n+1}^* are adapted to histories h^n and h^{n+1}, respectively, where $Y^{n+1} = Y^n = Y$, $\Phi^{n+1} = \Phi^n + 1$ and $\Pi^{n+1} = \Pi^n + 1$. By Lemma 5, we have that:

$$\mathbb{E}_n^1\big[u(C^*) - u(c_n^*)\big] \;\geq\; \mathbb{E}_n^1\left[u(C^*) - u(c_{n+1}^*)\right]$$

Hence, it is sufficient to prove that $\mathbb{E}_n^1\left[u(C^*) - u(c_{n+1}^*)\right] \geq 0$. We decompose this using iterated expectations as follows:

$$\begin{aligned}\mathbb{E}_n^1\left[u(C^*) - u(c_{n+1}^*)\right] \;=\;& \mathbb{P}\left[\omega_{n+1} = 0 \;\middle|\; h^n, \omega_n = 1\right] \cdot \mathbb{E}_n^1\left[u(C^*) - u(c_{n+1}^*) \;\middle|\; \omega_{n+1} = 0\right] \\ &+\; \mathbb{P}\left[\omega_{n+1} = 1 \;\middle|\; h^n, \omega_n = 1\right] \cdot \mathbb{E}_n^1\left[u(C) - u(c_{n+1}^*) \;\middle|\; \omega_{n+1} = 1\right]\end{aligned}$$

Notice that conditioning on the current trader-n being patient determines the history h^{n+1}

as above, and we have:

$$\begin{aligned} \mathbb{E}_n^1\left[u(C^*) - u(c_{n+1}^*)\right] \;&=\; \mathbb{P}\left[\omega_{n+1} = 0 \;\middle|\; h^{n+1}\right] \cdot \mathbb{E}_{n+1}^0\left[u(C^*) - u(c_{n+1}^*)\right] \\ &\quad +\; \mathbb{P}\left[\omega_{n+1} = 1 \;\middle|\; h^{n+1}\right] \cdot \mathbb{E}_{n+1}^1\left[u(C^*) - u(c_{n+1}^*)\right] \end{aligned}$$

The first term on the RHS is strictly positive by Lemma 6. The second term in the sum of the RHS is strictly positive by the induction hypothesis. This completes the proof.