Finance Research Letters 87 (2026) 109110

Contents lists available at ScienceDirect

Finance Research Letters

journal homepage: www.elsevier.com/locate/frl

Who is still in line? How bank beliefs drive fragility under runs[☆]

Péter Csóka [a,b,*], Tamás Erb [c], Hubert János Kiss [b,d]

[a] Corvinus University of Budapest, Institute of Finance, 8 Fővám tér, 1093 Budapest, Hungary
[b] ELTE Centre for Economic and Regional Studies, Hungary
[c] Independent researcher
[d] Corvinus University of Budapest, Department of Economics, Hungary

ARTICLE INFO

JEL classification:
D04
G01
G21

Keywords:
Bank run
Belief formation
Artificial intelligence
Systemic risk

ABSTRACT

During a bank run, the bank observes massive withdrawals but cannot distinguish between genuine liquidity-needs and panic withdrawals. Optimal policy responses—e.g., rescheduling or suspending convertibility—depend on the bank's beliefs about the liquidity needs of those yet to be served. We revisit Ennis and Keister (2009), relaxing their assumption that the share of impatient depositors still in line mirrors the population share, allowing arbitrary feasible beliefs. Our analysis shows that greater bank pessimism—i.e., a higher perceived share of impatient depositors still in line—increases fragility, defined as the likelihood that a bank-run equilibrium coexists with the efficient no-run outcome, regardless of the policy tool employed. This suggests that managing banks' beliefs—by reducing undue pessimism—before a crisis may mitigate run incentives. In an era of real-time data and AI-driven monitoring, banks may strive to infer more precise information on remaining depositors—potentially pivotal for preventing panic-driven runs.

1. Introduction

Ennis and Keister (2009) show that *ex-post* efficient policy responses to bank runs can create *ex-ante* incentives for depositors to run. Their model assumes that, during a run, the proportion of depositors with urgent liquidity needs (*impatient* depositors) among those yet to decide (i.e., to withdraw or retain their deposits) matches the population share. However, there is no empirical support for this assumption.[1] Reliable evidence on the sequence of withdrawal decisions is lacking, making it difficult to identify early withdrawers and, consequently, to estimate the remaining liquidity needs. This empirical gap leaves room to form differing beliefs about the composition of remaining depositors. While traditional banking systems offer little insight into who withdraws and when, digital-finance ecosystems provide granular data on user behavior during liquidity crises. For instance, Saengchote and Samphantharak (2024) analyze the collapse of the IRON algorithmic stablecoin using blockchain transaction data and show that sophisticated users consistently exited earlier and more profitably than retail depositors. This underscores the importance of belief formation under uncertainty in banking.

☆ Péter Csóka was supported by the János Bolyai scholarship of the Hungarian Academy of Sciences and also by funding from the National Research, Development and Innovation Office – NKFIH, K-138826. Hubert János Kiss gratefully acknowledges financial support from the Hungarian Academy of Sciences (Momentum Grant No. LP2021-2).

* Corresponding author.
E-mail addresses: peter.csoka@uni-corvinus.hu (P. Csóka), tamas.erb@gmail.com (T. Erb), kiss.hubert@krtk.hun-ren.hu (H.J. Kiss).

[1] Studies document substantial heterogeneity in depositor behavior, but offer no clear evidence on who withdraws and when during a bank run; see Kelly and Gráda (2000), Gráda and White (2003), Iyer and Puri (2012), Iyer et al. (2016) and Kiss (2018).

https://doi.org/10.1016/j.frl.2025.109110
Received 1 August 2025; Received in revised form 17 November 2025; Accepted 24 November 2025
Available online 25 November 2025

Banks may hold varying beliefs about the share of impatient depositors among those still to be served. Suppose a bank has 100 depositors, 20 of whom are impatient. Assume that after 20 withdrawals the twenty-first depositor arrives with the same request. The bank recognizes that a run is underway and, following Ennis and Keister (2009), assumes that 20%—i.e., 16—of the remaining 80 depositors have urgent liquidity needs. Alternatively, a *pessimistic* bank might assume that a higher share (or number) of impatient depositors remains, implying that many early withdrawals were made by patient depositors (not in need of liquidity but acting out of panic), while an *optimistic* bank might assume the opposite.

In practice, banks are unlikely to assume that the share of impatient depositors still in line exactly matches the population average π. Instead, they may update beliefs π^{br} based on observable cues during a run. For example, if early withdrawals come primarily from small retail accounts, the bank may infer that many impatient depositors remain, forming a pessimistic posterior. Banks may also infer from account types, recent communication signals, or the geographic/digital footprint of early activity. Even informal heuristics—e.g., patterns in past stress events—may shape belief formation. These considerations make belief heterogeneity feasible and empirically plausible.

Beliefs about the liquidity types of remaining depositors may shape how much the bank pays early withdrawers, which affects depositors' incentives to withdraw. Whether pessimistic or optimistic beliefs increase fragility is not immediately clear. A pessimistic bank expects relatively more depositors to withdraw and may therefore liquidate a larger share of its long-term investments to ensure that early withdrawers receive sufficiently high payments. However, this strategy reduces the funds available for late consumption, thereby increasing the incentives for patient depositors to run. Alternatively, a pessimistic bank could respond by offering lower payments to early withdrawers to conserve resources for those remaining. Yet if the bank aims to maximize total depositor welfare, such a strict payout policy may be suboptimal. An optimistic bank, by contrast, expects fewer withdrawals and may thus be willing to offer more generous early payments to accommodate impatient depositors with urgent liquidity needs. But this generosity can deplete liquid resources and force costly liquidation if more depositors than expected withdraw, amplifying fragility. Which belief weakens *ex-ante* run incentives is an open question we address.

We revisit the framework of Ennis and Keister (2009) with two modifications: (1) we use a discrete number of depositors instead of a continuum, and (2) we allow the bank to hold any feasible beliefs about the share of patient depositors among those yet to make a decision (withdraw or keep their funds deposited).[2] These beliefs matter because Davis and Reilly (2016) show experimentally that a credible tough stance (lower payouts) can reduce run likelihood. Their experiment captures the difference between tough and lenient responses through payoff structures but lacks a theoretical foundation for these choices. We provide this foundation, showing that beliefs about the remaining depositors' types can rationalize different payment schedules as optimal responses.

We analyze the *ex-post* optimization problem under payment rescheduling and the timing of convertibility suspension. Our results show that greater pessimism—interpreted as a larger number of remaining impatient depositors—consistently increases fragility, irrespective of the policy instrument. This highlights the importance of accurately estimating the liquidity needs of remaining depositors—increasingly feasible with AI and real-time data analytics.

The remainder of the paper is organized as follows. Section 2 presents the model. Section 3 analyzes payment rescheduling under heterogeneous beliefs and provides a numerical illustration. Section 4 concludes. The analysis of suspension of convertibility is relegated to the Appendix.

2. The model

2.1. Environment

Consider three dates $t \in \{0, 1, 2\}$ and $N > 2$ *ex-ante* identical depositors, each endowed with one unit of the consumption good at $t = 0$ and with preferences

$$v(c_1, c_2; \theta) = u(c_1 + \theta c_2),$$

where u is strictly increasing, strictly concave, $u(0) = 0$, and $u'(0) = \infty$. The privately observed type $\theta \in \{0, 1\}$ is revealed at $t = 1$: $\theta = 0$ denotes an *impatient* depositor who values only period-1 consumption, while $\theta = 1$ denotes a *patient* depositor who also values period-2 consumption. Let $0 < \kappa < N$ be the number of impatient depositors; the remaining $N - \kappa$ are patient, so aggregate liquidity demand is known.

The bank allocates deposits to a storage technology yielding one unit in either period, and a long-term investment yielding $R > 1$ in period 2 but only $1 - \tau$ in period 1 if liquidated early, where $\tau \in (0, 1)$ is the liquidation loss.

At $t = 1$ depositors learn their type and decide whether to withdraw immediately or wait until $t = 2$. Withdrawals follow a sequential-service constraint: payments depend only on withdrawals already processed. Impatient depositors withdraw in period 1 because they do not value consumption in period 2. Patient depositors may either wait or withdraw early.

[2] In Ennis and Keister (2009), the continuum assumption justifies a constant share of impatient depositors via the law of large numbers.

2.2. First-best allocation

A benevolent planner observes each depositor's type and assigns per-capita consumption c_E to impatient depositors in period 1 and c_L to patient depositors in period 2. Let I be the funds invested in the long-term investment, with the remainder $N - I$ stored. The planner solves

$$\max_{c_E, c_L, I} \kappa u(c_E) + (N - \kappa) u(c_L) \tag{1}$$

subject to

$$\kappa c_E = N - I, \tag{2}$$

$$(N - \kappa) c_L = RI, \tag{3}$$

$$c_E, c_L \geq 0, \qquad 0 \leq I \leq N. \tag{4}$$

Constraints (2)–(4) state that aggregate early consumption equals the liquid funds available in period 1, aggregate late consumption is financed by the matured long-asset return, and all variables remain non-negative and within the initial endowment.

To resemble (Ennis and Keister, 2009), we set the proportion of impatient depositors $\pi = \kappa/N$ and the investment share $i = I/N$. The optimization problem becomes

$$\max_{c_E, c_L, i} \pi u(c_E) + (1 - \pi) u(c_L) \tag{5}$$

subject to

$$\pi c_E = 1 - i, \tag{6}$$

$$(1 - \pi) c_L = Ri, \tag{7}$$

$$c_E, c_L \geq 0, \qquad 0 \leq i \leq 1. \tag{8}$$

Following Ennis and Keister (2009), we illustrate our findings with the constant relative risk-averse utility function. With $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$, $\gamma > 0$, the first-best allocation is

$$c_E^* = \frac{1}{\pi + (1 - \pi) R^{(1-\gamma)/\gamma}}, \tag{9}$$

$$c_L^* = \frac{R^{1/\gamma}}{\pi + (1 - \pi) R^{(1-\gamma)/\gamma}}, \tag{10}$$

implying $i^* = 1 - \pi c_E^*$. To ensure runs are feasible, we assume

$$c_E^* > 1 - \tau i^*. \tag{11}$$

This condition means the bank lacks sufficient liquid assets to pay c_E^* to all depositors in period 1 due to limited liquid funds, making runs feasible when patient depositors panic and withdraw early.

Diamond and Dybvig (1983) show that a demand–deposit contract can implement the first-best allocation despite private depositor types. Depositors choose whether to withdraw in period 1 or 2, and the bank pays c_E^* to early withdrawers until funds are exhausted. Those who keep their funds in the bank share the matured assets pro rata. If the bank invests $I^* = N - \kappa c_E^*$ and offers c_E^* in period 1, an equilibrium exists in which impatient depositors withdraw early and receive c_E^*, while patient depositors wait and receive c_L^* in period 2. However, a bank run equilibrium also exists: if patient depositors expect everyone to withdraw early, they do so too to avoid receiving nothing, since the bank cannot pay c_E^* to all in period 1.

Fragility arises when the run equilibrium coexists with the efficient, no-run outcome. Ennis and Keister (2009) study two policy tools—payment rescheduling and suspension of convertibility (deposit freeze). We focus on rescheduling; results for suspension of convertibility are relegated to the Appendix.

3. Rescheduling payments

Once the first $\kappa = \pi N$ depositors have withdrawn, the bank reoptimizes with $N - \kappa$ depositors and $N - \kappa c_E^*$ funds remaining. Keeping the original investment share $i^* = 1 - \pi c_E^*$ fixed, the bank chooses new early and late payments (c_E^{br}, c_L^{br}) to

$$\max_{c_E^{br}, c_L^{br}} (1 - \pi) \left[\pi^{br} u(c_E^{br}) + (1 - \pi^{br}) u(c_L^{br}) \right] \tag{12}$$

subject to

$$(1 - \pi) \pi^{br} c_E^{br} \leq (1 - \tau) i^*, \tag{13}$$

$$(1 - \pi)(1 - \pi^{br}) \frac{c_L^{br}}{R} = i^* - \frac{(1 - \pi) \pi^{br} c_E^{br}}{1 - \tau}, \tag{14}$$

$$c_E^{br}, c_L^{br} \geq 0, \tag{15}$$

$$\frac{1}{N(1-\pi)} \leq \pi^{br} \leq \frac{\pi}{1-\pi}. \tag{16}$$

Constraint (13) limits total early payouts to the cash obtained by liquidating the long investment. Constraint (14) allocates the remaining investment proceeds (net of liquidation losses) to those the bank believes are patient. Constraints (15) impose non-negativity on both early and late consumption. Constraint (16) captures the feasible *posterior* beliefs about the share of impatient depositors still in line: after κ early withdrawals, at least one and at most κ impatient depositors may remain, so π^{br} must lie between $1/[N(1-\pi)]$ and $\pi/(1-\pi)$.[3]

The first-order condition yields that a solution to this problem should satisfy

$$u'(c_E^{br}) = \frac{R}{1-\tau} u'(c_L^{br}).$$

A bank run, as in Ennis and Keister (2009), occurs only if the following two conditions are met simultaneously:

(i) *Run-feasibility:* $c_E^* > 1 - \tau i^*$, i.e. the bank would default if all depositors rushed to the window in period 1.
(ii) *Run-profitability:* $c_E^* > c_L^{br*}$, i.e. a patient depositor strictly prefers to withdraw early anticipating the bank's rescheduling policy.

We now turn to the CRRA utility function

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}, \qquad \gamma > 0,\ \gamma \neq 1.$$

The first-order conditions yield the following interior solution

$$c_E^{br*} = \frac{(1-\tau)\, i^*}{1-\pi} \; \frac{1}{\pi^{br} + (1-\pi^{br})\, R^{\frac{1}{\gamma}-1} (1-\tau)^{1-\frac{1}{\gamma}}},$$

$$c_L^{br*} = \frac{(1-\tau)\, i^*}{1-\pi} \; \frac{\left(R/(1-\tau)\right)^{1/\gamma}}{\pi^{br} + (1-\pi^{br})\, R^{\frac{1}{\gamma}-1} (1-\tau)^{1-\frac{1}{\gamma}}}.$$

A run can occur only if a patient depositor prefers the early payment, c_E^*, to the rescheduled late payment, c_L^{br*}. After substitutions, the inequality $c_E^* > c_L^{br*}$ reduces to

$$\pi^{br} > \frac{R^{\frac{1}{\gamma}-1} (1-\tau)^{1-\frac{1}{\gamma}} \left(R^{1/\gamma} - 1\right)}{1 - R^{\frac{1}{\gamma}-1} (1-\tau)^{1-\frac{1}{\gamma}}} \tag{17}$$

Condition (17) states that more pessimistic beliefs (higher π^{br}) make the condition more likely to hold, increasing fragility. Condition (17) is monotonic in π^{br}: for any given (τ, R, γ), a more pessimistic belief ($\pi^{br} \uparrow$) increases the left-hand side, while the right-hand side remains fixed, thus strictly increasing the likelihood of fragility.

Intuitively, greater pessimism leads the bank to liquidate a larger share of its long-term investment, boosting the early payment but reducing the late one. Anticipating a lower c_L^{br*}, patient depositors choose to run. Importantly, condition (17) does not depend on the population share of impatient depositors (π), implying that a bank with sufficiently optimistic beliefs can remain non-fragile, whereas a bank whose beliefs reflect π may still be fragile. We illustrate our findings with Fig. 1.

The x-axis represents the bank's belief π^{br}, defined as the perceived share of impatient depositors among those still awaiting service after the initial wave of withdrawals. The y-axis shows the liquidation cost τ. Each panel corresponds to a different level of risk aversion, with $\gamma \in \{2,\ 6\}$, and overlays fragility regions for three values of the long-term return $R \in \{1.05,\ 1.15,\ 2.0\}$. For each (γ, R) combination, shaded areas indicate parameter pairs (π^{br}, τ) where the bank is fragile—that is, where both a run is feasible and patient depositors have an incentive to join it.

Fig. 1 reveals three key patterns. First, fragility decreases with higher returns R: greater long-term yields enhance the value of waiting, shrinking the region where a run is profitable. Second, fragility is more pronounced under higher γ—that is, when depositors are more risk-averse. Third, and most important for this study, fragility increases with more pessimistic beliefs: as π^{br} rises, the bank anticipates a greater number of impatient depositors still in line and responds by liquidating more assets, which lowers late consumption and incentivizes patient depositors to withdraw.

As a concrete example, consider the case where $\gamma = 2$, $R = 1.15$, $\tau = 0.5$, and $\pi = 0.2$, which implies first-best payments upon withdrawal of $c_E^* \approx 1.06$. The re-optimized late payment, c_L^{br*}, depends on the bank's belief. A bank with a very optimistic belief ($\pi^{br} = 0.1$) would offer $c_L^{br*} \approx 1.08$, resulting in no fragility. However, even a moderately optimistic bank with $\pi^{br} = 0.15$ would offer $c_L^{br*} \approx 1.05$, which already falls in the fragile region. The same applies to a bank whose belief reflects the population share of impatient depositors ($\pi^{br} = 0.2$), as well as to any bank holding more pessimistic beliefs ($\pi^{br} > \pi$).

[3] After the first $\kappa = \pi N$ withdrawals, $N - \kappa = N(1-\pi)$ depositors remain in line. Let m denote the (unknown) number of impatient depositors among them.

$$1 \leq m \leq \kappa \quad \Longrightarrow \quad \frac{1}{N(1-\pi)} \leq \pi^{br} = \frac{m}{N(1-\pi)} \leq \frac{\kappa}{N(1-\pi)} = \frac{\pi}{1-\pi}.$$

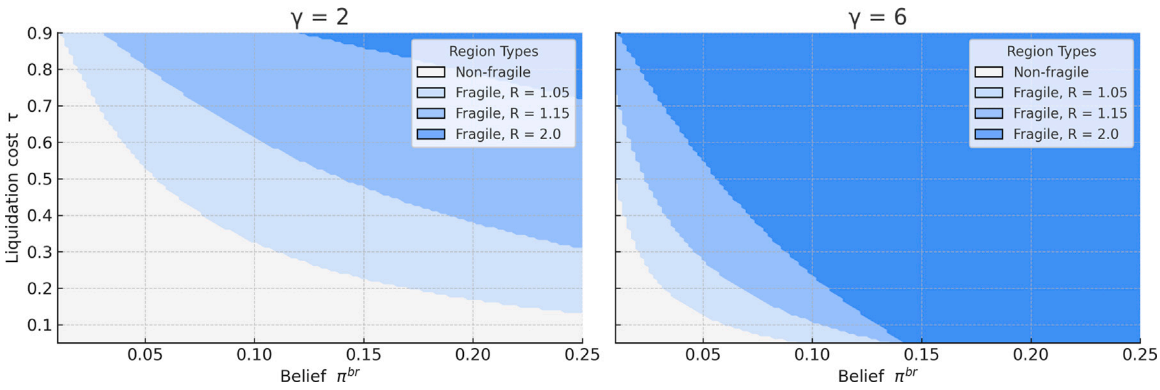

Fig. 1. Contour plots of fragility as a function of risk aversion (γ), return on the long-term asset (R), liquidation cost (τ), and the bank's belief (π^{br}) about the share of impatient depositors.

Relating back to Davis and Reilly (2016), the observed reduction in bank runs under a tougher stance may be micro-founded by the presence of optimistic beliefs among banks.

The Appendix presents the same exercise for the suspension of convertibility. There, we find qualitatively similar results: more pessimistic beliefs lead to increased fragility. This reinforces the main finding in this section.

4. Conclusion

Once a run begins, banks must rely on beliefs—not facts—about who remains in line. By allowing beliefs to vary within their feasible range, we show that the two main policy instruments—payment rescheduling and suspension of convertibility—share an Achilles' heel: more pessimistic beliefs increase fragility.

We model the bank as holding a fixed belief about the share of remaining impatient depositors; however, it is natural to assume that beliefs may be updated in response to additional withdrawals. Exploring such dynamic belief-updating processes represents a promising avenue for future research. Moreover, while our model isolates belief-driven fragility within a single institution, systemic vulnerability often arises from financial interconnections. Recent empirical evidence by Nguyen and Le (2025) demonstrates that network spillovers among Vietnamese banks played a significant role during episodes of market stress, underscoring the relevance of fragility in interconnected systems.

These results imply a policy takeaway: regulators should discourage unduly pessimistic bank beliefs. One avenue is to leverage real-time data and AI-driven tools to better infer depositors' liquidity-needs. For example, banks can monitor withdrawal timing, account characteristics, and historical behavior to estimate the likely share of remaining impatient depositors. Although not directly comparable, evidence from a different bank-run setting—where delayed policy responses increase welfare losses yet reduce fragility by making banks more cautious (Izumi and Li, 2025)—points to divergent mechanisms relative to our belief-driven result, underscoring the need for further research.

Machine-learning models could aggregate such signals—combined with metadata like device use or geolocation—to produce dynamic estimates of the remaining type distribution π^{br}. Natural language processing (NLP) may support inference by analyzing customer communications for urgency cues. These tools could inform automated or semi-automated adjustments to payout policies or suspension timing, reducing unnecessary pessimism during runs.

Anchoring beliefs in richer real-time information may directly mitigate fragility. A natural extension would endogenize belief formation via noisy signals or incorporate costly information acquisition, linking fragility to emerging research on information design in financial regulation.

While we focus on the effects of varying bank beliefs within their feasible range, developments during a run may reveal a sharp divergence from actual liquidity needs. For example, court-ordered access to funds, as in Argentina's crisis (Ennis and Keister, 2009), may expose a different composition of types than what the bank had assumed. Such divergence may undermine the bank's response. How institutions respond—by reoptimizing payments, freezing withdrawals, or coordinating with regulators—remains open.

CRediT authorship contribution statement

Péter Csóka: Writing – review & editing, Validation, Funding acquisition, Conceptualization. **Tamás Erb:** Writing – review & editing, Writing – original draft, Visualization, Investigation, Formal analysis, Conceptualization. **Hubert János Kiss:** Writing – review & editing, Writing – original draft, Visualization, Supervision, Investigation, Formal analysis, Conceptualization.

Declaration of Generative AI and AI-assisted technologies in the writing process

During preparation, the authors used ChatGPT-4o to improve English and readability. The authors reviewed and edited the content and take full responsibility for the publication.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Suspension of convertibility

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.frl.2025.109110.

Data availability

No data were used for the research described in the article.

References

Davis, D.D., Reilly, R.J., 2016. On freezing depositor funds at financially distressed banks: An experimental analysis. J. Money Credit. Bank. 48 (5), 989–1017.
Diamond, D.W., Dybvig, P.H., 1983. Bank runs, deposit insurance and liquidity. J. Political Econ. 91 (3), 401–419.
Ennis, H.M., Keister, T., 2009. Bank runs and institutions: The perils of interventions. Am. Econ. Rev. 99 (4), 1588–1607.
Gráda, C.Ó., White, E.N., 2003. The panics of 1854 and 1857: A view from the emigrant industrial savings bank. J. Econ. Hist. 63 (1), 213–240.
Iyer, R., Puri, M., 2012. Understanding bank runs: The importance of depositor-bank relationships and networks. Am. Econ. Rev. 102 (4), 1414–1445.
Iyer, R., Puri, M., Ryan, N., 2016. A tale of two runs: Depositor responses to bank solvency risk. J. Finance.
Izumi, R., Li, Y., 2025. Rapid bank runs and delayed policy responses. J. Financial Stab. 101422.
Kelly, M., Gráda, C.Ó., 2000. Market contagion: Evidence from the panics of 1854 and 1857. Am. Econ. Rev. 90 (5), 1110–1124.
Kiss, H.J., 2018. Depositors' behaviour in times of mass deposit withdrawals. Financial Econ. Rev. 17 (4), 95–111.
Nguyen, A.T., Le, T.T., 2025. In bank runs and market stress, it matters how networks impact: Exploring the financial connectedness in Vietnam. Finance Res. Lett. 72, 106489.
Saengchote, K., Samphantharak, K., 2024. Digital money creation and algorithmic stablecoin run. Finance Res. Lett. 64, 105435.