SUERF
The European Money
and Finance Forum

SUERF Policy Brief
No 1409, April 2026

Stablecoins are run-optimised instruments


Jón Danielsson | SUERF Fellow and Systemic Risk Centre, London School of Economics
Robert Macrae | Systemic Risk Centre, London School of Economics

Keywords: Stablecoins, Artificial Intelligence (AI), crises, runs

JEL codes: G01, G23, H12

Abstract

Stablecoins are private digital money designed to trade at par with a fiat currency, usually the US dollar. This column argues that stablecoins are run-optimised in the sense that the promise of cheap and fast trading around the clock and across borders will lead to faster, synchronised exits that compress the timeline for stablecoin crises to minutes or hours. AI will exacerbate this.

Disclaimer: The views expressed are those of the authors and not necessarily those of the institutions the authors are affiliated with.

Introduction

Stablecoins are private money in token form, designed to trade at par with a reference fiat currency. Issuers mint and redeem tokens against reserve assets – cash and short-term government securities or bank deposits.

Today, stablecoins mainly serve as the cash leg of crypto markets. They are used for trading and settlement on exchanges, with growing use in decentralised finance. Cross-border transfers are on the rise, and in countries with weak currencies, stablecoins increasingly serve as a substitute store of value.

Stablecoins are overwhelmingly US dollar-denominated, with around $310 billion outstanding in February 2026. Roughly $185bn is in USD Tether (USDT) and $73bn in USD Coin (USDC). Currently, that is tiny relative to domestic deposits and payment systems, with M0 about $5.3 trillion and M1 $19 trillion, so stablecoins are not yet a significant crisis driver. Much of what follows is therefore conditional: what happens if and when stablecoins become large enough to matter?

Here, we focus on their behaviour under stress, when money-like assets are judged by a single criterion: can they be converted into fiat at par, quickly and at scale?

Parity is maintained, not inevitable

Stablecoins promise a one-to-one peg with fiat, but that peg is not self-executing; it is continuously produced through arbitrage and the issuers' interaction with the fiat financial system.

Typically, only primary authorised participants can redeem stablecoins into fiat directly with the issuer at, or close to, par. But even they are subject to terms and compliance requirements and are ultimately constrained by the availability of banking and market infrastructure.

Most holders, instead, enter and exit stablecoin positions through secondary markets such as centralised exchanges, decentralised exchanges, brokers, OTC desks and market makers. These venues typically operate 24/7, allowing near-instantaneous token-to-token swaps under normal market conditions. Arbitrageurs help maintain price parity across venues and, when necessary, transmit secondary market price signals into primary mint and redemption flows.

Under stress, liquidity providers and market makers may reduce risk or withdraw entirely, causing spreads to widen and, in extreme cases, allowing the peg to break. The consequence may be that the holders of stablecoins become unable to transact at meaningful scale without incurring severe discounts. If this happens, price ceases to represent mechanical par convertibility and instead reflects liquidity premia, counterparty risk, uncertainty regarding primary market redemption, and market assessments of reserve quality and accessibility.

Three failure modes matter:

First, when reserves are worth less than liabilities and the stablecoin is effectively bankrupt, an asset problem arises. This is not hypothetical: issuers earn seigniorage on reserves and face competitive pressure to reach for yield, especially if permitted to pay interest. At the same time, opacity of the composition of reserves weakens market discipline. In times of stress, the consequence may well be that liabilities exceed assets.

Second, if conversion fails even though reserves are adequate, a plumbing problem arises. This could happen because of internet outages, sanctions, banking hours, disrupted rails, operational outages, compliance freezes or collateral constraints by intermediaries.

Third, a belief problem can arise, à la Diamond and Dybvig (1983), when everyone runs because they expect everyone else to run. It is, after all, optimal to be first.

In practice, these modes interact. A plumbing failure – say banking rails are down over a weekend – might trigger a belief run that forces fire sales that could turn a solvent issuer into an insolvent one. Any single failure mode can activate the others, converting liquidity stress into solvency impairment.

Why stablecoins are run-optimised

While stablecoins are not intentionally "designed to run", in practice they are run-optimised because their architecture minimises exit friction. Transfers settle in near real time, markets operate 24/7, and holders can reallocate globally without geographic or institutional constraints. While the underlying vulnerabilities have always been present in banking systems, and were clearly identified by Bagehot in 1873, what is new is not merely faster settlement (this has also improved in fiat markets) but the removal of gating mechanisms, the programmability of liquidation, and the compression of coordination time. Ultimately, stablecoins are run-optimised because they create a heightened belief in liquidity, which intensifies the resulting panic when this belief collapses.

Six features matter.

1. Stablecoins trade and settle 24/7.
2. Primary mint and redemption ultimately settle through the fiat banking system and remain constrained by its operating hours and rails.
3. Stablecoins are cross-border instruments by default, so stress is not contained by a domestic perimeter.
4. The fastest run vector is stablecoin-to-stablecoin flight on secondary markets, which occurs near-instantaneously on-chain. Liquidity can migrate across issuers in seconds, transmitting stress before primary redemption is even activated.
5. Issuers possess private circuit breakers – freezes, blacklists, throttles – which may be necessary for compliance and risk control but create pre-emptive incentives to exit before constraints are imposed.
6. There is no explicit state backstop, but such a backstop can be implicit via the institutions that hold the stablecoins' reserves.

In combination, these features compress the time between concern and withdrawal, reducing the ability of issuers or authorities to slow down exits, and transforming confidence shocks into liquidity crises much faster than in crises past.

Traditional banking systems rely on public circuit breakers – deposit insurance, a lender of last resort, supervisory intervention and, historically, the ability to suspend convertibility – to contain runs. Stablecoins, by contrast, depend on private controls such as freezes, blacklists and redemption limits. These tools may manage operational or compliance risk, but they do not constitute a coordinated public backstop.

As a result, stablecoins intensify run dynamics, turning events that might once have been mediated by domestic institutions over days into episodes that are fast, global, market-based and increasingly automated.

How runs play out

The speed and form of exit depend on who is running. Retail holders are slowest because they are constrained by identity checks and banking rails. Institutional holders move faster through OTC desks and direct redemption. Fastest are on-chain algorithmic systems.

The first and fastest leg is stablecoin-to-stablecoin flight. When one issuer is perceived as weaker, liquidity migrates on-chain around the clock. This stage redistributes risk but does not yet shrink system-wide liquidity. Damage escalates only if secondary markets become illiquid and redemptions force reserve liquidation. If outside normal business hours, reserve liquidation is constrained by repo settlement, custodian operations and banking rails, threatening disorderly cascades when those markets reopen. The only defence may be redemption gates. But these are so deeply unpopular that stablecoin operators have, in normal times, a strong incentive to promise they will not be introduced. Competitive pressure is likely to make redemption gates impossible unless they are backed by a binding regulatory mandate for all operators.

If redemptions accelerate outside normal business hours, issuers may face timing mismatches between 24/7 token markets and business-hour reserve markets. This can amplify volatility when underlying markets reopen, echoing Black Monday in 1987.

Cross-currency runs introduce an additional layer of complexity. Large-scale migration from one currency-backed stablecoin to another may require issuers to rebalance reserves across jurisdictions with differing legal frameworks and crisis-management tools. Authorities in the affected jurisdictions may have neither equivalent powers nor the same incentives to respond, preventing coordinated cross-border action. That said, this channel is naturally constrained because a willing buyer must exist, and the redemption price of the local stablecoin might already reflect the expectations of the eventual crisis outcome.

Stablecoin runs into fiat follow a classic redemption dynamic. If holders lose confidence in an issuer, redemption demand rises as investors seek bank deposits or other fiat claims. Although these flows are slower than on-chain migrations because they are constrained by compliance procedures and banking settlement windows, sustained redemptions may require issuers to liquidate reserve assets. In stressed market conditions, this could create a fire-sale loop between falling reserve prices and rising redemption demand.

The reverse dynamic is also possible. In periods of banking stress, depositors may view short-duration Treasury-backed stablecoins as safer than uninsured deposits, triggering rapid inflows into stablecoins and accelerating deposit outflows from banks. However, in such a scenario, a more sensible course of action would be to buy the Treasuries directly.

The transmission channel to traditional finance runs through reserve portfolios. USD Tether and USD Coin alone hold over $250 billion in reserves – primarily short-dated US treasuries, reverse repo and money-market fund shares – which is small relative to the $28 trillion Treasury market. If stablecoins grow significantly, the end result may be pro-cyclical sales into already stressed markets, perhaps similar to money-market funds in 2008 and 2020, although Aldasoro et al. (2024) show that the transmission channels differ. Similarly, in March 2020, some bond ETFs were valued at ten percent below NAV, in part because of inability to arbitrage between liquid and illiquid assets.

The pattern is visible in recent episodes. In May 2022, the Terra Luna collapse destroyed over $40 billion in value, demonstrating how quickly confidence shocks can propagate through crypto markets. In March 2023, USD Coin de-pegged after the failure of Silicon Valley Bank (SVB) exposed reserve and custodian risk. The peg was restored only after US authorities guaranteed SVB deposits, reaffirming the convertibility of reserves.

AI and the rails

AI matters because it compresses and synchronises the critical decision in a crisis: whether to stay or to run.

Most of the time, financial institutions optimise for profit. Under stress, that flips – and survival becomes of the essence (Danielsson 2024). Then, if the institution decides to run, speed dominates: the first mover gets the best prices; the last mover faces large losses, even bankruptcy.

Runs are an example of strategic complementarity: if others run, the optimal response is to run too. When many institutions rely on similar data and model architectures, that coordination problem becomes sharper. AI sharpens it further. In normal conditions, it improves monitoring and risk management. Under stress, however, shared data sources, similar model architectures and automated execution pipelines can produce highly correlated responses. Signals are processed faster than human committees can convene, and exit decisions can be executed programmatically. This dynamic suppresses volatility in normal times and increases tail risk in times of stress (Danielsson 2025).

Stablecoins amplify these forces by providing continuous markets and near-frictionless on-chain exit. If AI-driven systems decide to reduce exposures, then price movements, on-chain outflows and liquidity withdrawal can reinforce one another in real time, turning strategic coordination into a mechanical cascade. At some point, liquidity providers are forced to step away, causing a self-fulfilling de-peg and a predictably rapid and aggressive response from the AIs that all receive the same shock at the same time.

The consequence could be a much faster and more vicious crisis than any we have ever seen before.

Conclusion

Stablecoins exist because of a demand for stable, accessible, liquid money. But the features that make this money so attractive – 24/7, borderless, low-friction – also make it a run-optimised instrument: in a crisis, the expectation of continuous liquidity can reverse abruptly under stress, turning confidence shocks into rapid, self-reinforcing exits. Ultimately, stablecoins might not be designed to run, but they are run-optimised.

AI adds a second accelerator: faster, more synchronised decisions about whether to stay or run. Stablecoins provide the execution rail; AI provides the decision engine. Together they compress the crisis timeline and raise the premium on preparedness.

The question is not whether stablecoins will face runs, but whether the response infrastructure can match the speed of the exit; and whether, as with SVB, fiat bailouts will again be required to contain a crypto crisis.

The traditional crisis-fighting infrastructure was, and remains, set up for the much slower crises of days past. That, by itself, makes a future crisis more likely.

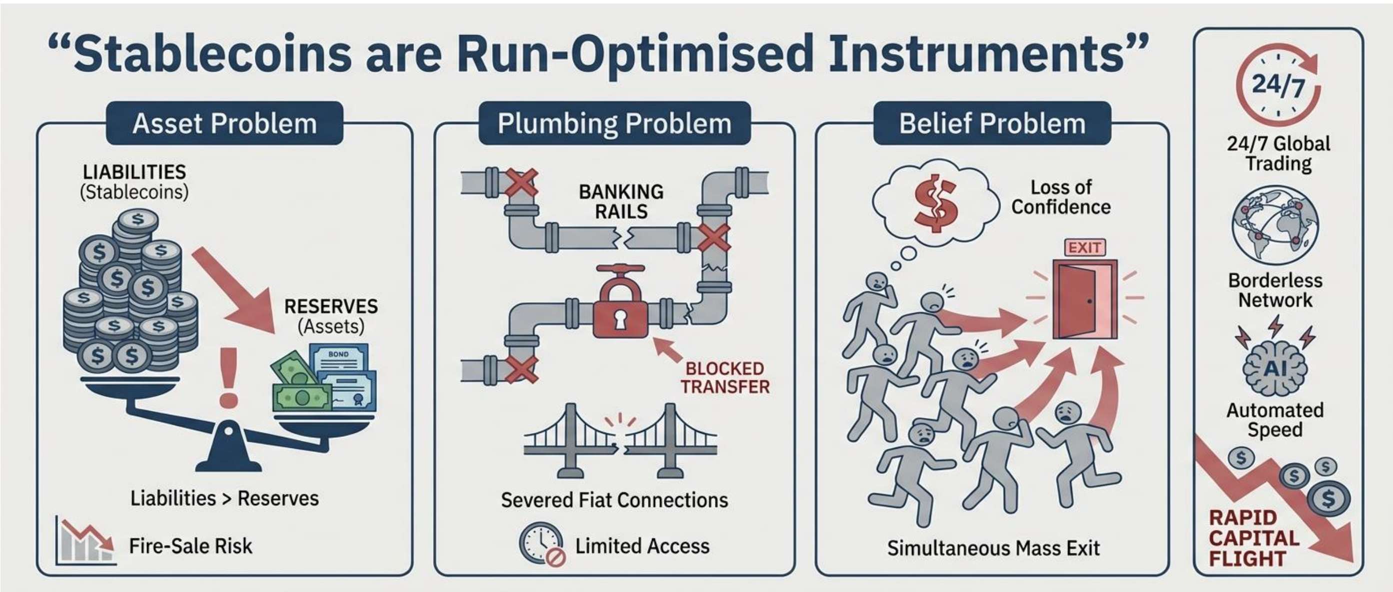

Note: Image generated with Google Gemini

References

Aldasoro, I, G Cornelli, L Gambacorta, M Ferrari Minesso, and M M Habib (2024), "Stablecoins and Money Market Funds: Less Similar than You Think", VoxEU.org, 22 November. https://cepr.org/voxeu/columns/stablecoins-and-money-market-funds-less-similar-you-think.

Danielsson, J (2024), "The One-in-a-Thousand-Day Problem", VoxEU.org. https://cepr.org/voxeu/columns/one-thousand-day-problem.

Danielsson, J (2025), "Artificial Intelligence and Stability", VoxEU.org, 6 February. https://cepr.org/voxeu/columns/artificial-intelligence-and-stability.

Diamond, D W and P H Dybvig (1983), "Bank Runs, Deposit Insurance, and Liquidity", Journal of Political Economy 91(3): 401–19.

About the author(s)

Jón Daníelsson is a SUERF Fellow, the Director of Systemic Risk Centre and Professor of Finance at the London School of Economics. Since receiving his PhD, Jón's work has focused on how economic policy can lead to prosperity or disaster. He is an authority on both the technical aspects of risk forecasting and the optimal policies that governments and regulators should pursue in this area. Jón has written three highly regarded books: The Illusion of Control (Yale University Press, 2022), which was included on the Financial Times “Best books of 2022” list; Financial Risk Forecasting (Wiley, 2011); and Global Financial Systems: Stability and Risk (Pearson, 2013). He has also contributed numerous academic papers on systemic risk, artificial intelligence, financial risk forecasting, financial regulation and related topics.

Robert Macrae is a Research Associate at LSE's Systemic Risk Centre. Robert has spent the last 30 years applying an engineering approach to problems in value investment and risk. Robert studied electrical engineering at Imperial College, London and moved into finance in 1990, eventually managing several value-based equity hedge funds. He is now a research associate at the LSE Financial Markets Group with areas of interest spanning value, hedge funds, risk control, tail risk, financial crises, financial regulation and machine learning.

Designed by the Information Management and Services Division of the Oesterreichische Nationalbank (OeNB)

SUERF Secretariat
c/o OeNB, Otto-Wagner-Platz 3A-1090 Vienna, Austria
Phone: +43 1 40 420 7206
E-Mail: suerf@oenb.at
Website: https://www.suerf.org/