Wrapped Stablecoins and Systemic Risk in DeFi

Pairie Koh* Kuntara Pukthuanthong†

Abstract

Wrapped stablecoins convert base stablecoins into receipt tokens that can be pledged repeatedly across DeFi lending protocols, implementing a form of shadow banking on-chain. We develop a model in which this recursive collateral reuse generates an endogenous liquidity multiplier that expands leverage in normal times but amplifies liquidation cascades when collateral constraints tighten. Using novel on-chain data from Aave deployments across eight blockchain networks, we document three findings. First, using regulatory announcements as instruments, we show that wrapped stablecoin creation causally increases recursive leverage: a one-standard-deviation increase raises the liquidity multiplier by 7% relative to its mean. Second, higher pre-event recursive leverage amplifies cryptocurrency volatility during stress episodes, with economically large effects concentrated in equally-weighted rather than value-weighted measures. Third, cross-chain bridge connectivity, not chain size, determines contagion speed: doubling bridge connections reduces time to stress arrival by 67% during the March 2023 USDC depeg. Wallet-level evidence confirms that agents execute the deposit-borrow-redeposit loops predicted by our model. Our results highlight a fundamental tradeoff: deeper collateral reuse improves funding efficiency but increases exposure to systemic liquidation cascades.

Keywords: Stablecoins, DeFi, Rehypothecation, Recursive Leverage, Shadow Banking, Systemic Risk, Collateral Chains, Liquidation Cascades

JEL Classification: G01, G23, G28, E44

*Independent Research. Email: hello.pairie@gmail.com

†Robert J. Trulaske, Sr. College of Business, University of Missouri. Email: pukthuanthongk@missouri.edu.

1 Introduction

DeFi is often described as bankless. Yet its funding markets are built on the same primitive that made shadow banking both scalable and fragile: collateral that can be pledged repeatedly. Wrapped stablecoins are the on-chain form of this collateral. They are private, money-like claims designed for collateral mobility rather than payments, and they sit at the center of leverage in DeFi. We examine collateral reuse through wrapped stablecoins, and we show theoretically and empirically how it expands leverage in normal times while generating abrupt unwind risk and liquidation cascades in times of stress.

Historically, collateral-based intermediation in shadow banking took two forms. First, collateral could be reused within dealer-based secured funding, such as repo markets in which broker-dealers financed inventories against high-quality collateral that was frequently rehypothecated along chains. This system supported large balance sheets in normal times, but its effective collateral supply depended on haircut policies and counterparties' willingness to accept and re-pledge collateral. Second, collateral-based intermediation could occur through pooled vehicles, such as money market funds and securitization conduits, where money-like claims were created against portfolios of collateral and liquidity was provided through short-term funding. These structures broadened the investor base and scale of funding but were vulnerable to runs and fire-sale dynamics when the quality or liquidity of collateral was questioned.

Wrapped stablecoins are the on-chain analogue of collateral that is engineered for reuse. Like rehypothecatable collateral in repo, they are designed to circulate across funding venues: a base stablecoin is packaged into a claim that can be posted, re-posted, and routed across protocols to support borrowing and leverage. Unlike a simple deposit-like claim held to redemption, wrapped stablecoins are valuable precisely because they can sit inside collateral chains. These two features—money-like acceptance as collateral and repeated pledgability—make wrapped stablecoins susceptible to the classic fragility of secured funding markets: when collateral constraints tighten, leverage can unwind abruptly through margin calls, forced liquidations, and spillovers across interconnected positions. Since collateral reuse is determined in equilibrium, when and why do recursive leverage loops build through wrapped stablecoins? And how do tightening collateral constraints translate into sudden liquidation cascades rather than smooth deleveraging?

We answer these questions by developing a framework for the market structure of wrapped stablecoins, showing how issuers, arbitrageurs, lending protocols, liquidators, and end users interact to determine the effective supply of money-like claims in DeFi and the likelihood of destabilizing deleveraging spirals. Our first contribution is to document the novel and surprising fact that wrapped stablecoins implement a form of shadow banking on chain. In particular, base stablecoins are transformed into receipt tokens that circulate as collateral and can be pledged repeatedly, so that the same underlying dollar claim supports multiple layers of borrowing across venues. This mechanism is surprising because it expands short-term funding capacity without any new outside deposits, and it implies that small changes in collateral terms can lead to large changes in aggregate credit.

Second, we show that while this on-chain credit intermediation improves market depth and funding capacity in normal times, it can increase fragility in stress in a way that closely parallels shadow banking runs. Stress events are possible despite transparent, overcollateralized lending, because the system relies on liquidity transformation through collateral values and liquidation technology rather than a stable deposit base. When the underlying stablecoin deviates from par, or when collateral haircuts and liquidation penalties tighten effective borrowing constraints, forced sales become more likely. If enough borrowers are pushed near liquidation thresholds, it is rational for liquidators to sell collateral into weakening markets, which lowers prices further, triggers additional liquidations, and produces a self-reinforcing liquidation cascade. We thus highlight a new tradeoff faced by wrapped stablecoin markets: deeper collateral reuse improves funding efficiency and compresses spreads, but it also increases exposure to sudden balance-sheet contraction when collateral values become state dependent.

Third, our framework has novel implications for policy and market design interventions that target stablecoins and DeFi leverage. Our results imply that improving day-to-day funding conditions and reducing the risk of systemic cascades are fundamentally different goals, and one may come at the expense of the other. For example, proposals that broaden collateral eligibility, raise collateral factors, or standardize wrapped receipt usability across protocols would expand effective liquidity and improve pricing, but they may have the unintended consequence of amplifying crash dynamics unless coupled with constraints that limit recursive reuse in stress. Similarly, proposals that focus only on issuer-level backing and redemption quality may address one source of fragility while leaving the shadow banking channel intact if rehypothecation chains remain unconstrained.

We introduce a novel dataset that links wrapped stablecoin creation to downstream collateral reuse and liquidation outcomes across major DeFi ecosystems. We focus on Ethereum and its major L2 chains plus Avalanche, where lending activity is deep and wrapped stablecoin dynamics are most prevalent, and we primarily focus on Aave due to data limitations that make comparable protocol policy variables difficult to measure consistently across platforms. Using on-chain events, we measure when stablecoins are deposited, when receipt tokens are minted, when those receipts are re-pledged, and when liquidations occur, and we link these outcomes to stablecoin price deviations in secondary markets. This allows us to document several stylized facts about on-chain shadow banking, including that the expansion and contraction of credit are concentrated in a small number of assets and venues, and that stress episodes feature abrupt increases in forced selling and rapid spillovers across connected markets.

We develop a tractable model of the two-layered market structure of base stablecoins and wrapped collateral claims, yielding analytical predictions for both effective liquidity and cascade risk. In the baseline model, agents hold stablecoins for transaction and portfolio motives, supply stablecoins as collateral to borrow, and can redeploy receipt tokens to expand positions, while liquidators enforce collateral constraints mechanically. The key friction is that collateral acceptance is state dependent, so a peg shock or a liquidity shock can tighten constraints abruptly, turning smooth selling pressure into forced sales. Importantly, what is unique about wrapped stablecoin market structure is that fragility is shaped by the degree of collateral reuse, because deeper reuse increases the sensitivity

of aggregate positions to small changes in effective collateral value.

There has been increased attention on stablecoin regulation and on the systemic risk of DeFi leverage across several jurisdictions (see, e.g., Allen, 2023). Regulators and market participants would like stablecoin-based markets to support efficient settlement and credit provision while avoiding destabilizing run-like dynamics. Our framework highlights that these desirable goals are distinct from each other and driven by different economic forces. In particular, we highlight the tradeoff between liquidity creation through on-chain shadow banking and financial stability through robust collateral constraints, and we show that some interventions may attain one goal at the expense of the other.

Our results provide important insights for the design of redemption rules, collateral policy, and risk controls in stablecoin-based DeFi. First, measures that increase collateral reuse can improve funding conditions, but they can also amplify liquidation-driven contagion unless accompanied by state-contingent constraints. Second, policies that improve the quality and liquidity of base stablecoin backing should be coordinated with policies governing rehypothecation and collateral eligibility, since these factors jointly determine both secondary market price stability and cascade risk. Third, we show that rules that limit recursive reuse in stress, for example by tightening collateral factors when peg deviations widen, can reduce systemic amplification while preserving many of the benefits of collateral reuse in normal times. Taken together, we highlight that while deep liquidity and low systemic risk are both desirable features of stablecoin markets, they are fundamentally distinct and respond differently to policy and market design interventions.

2 Related Literature

Our analysis of wrapped stablecoins as rehypothecatable money-like collateral builds on a large literature on shadow banking, collateral reuse, and secured funding fragility (e.g., (Adrian and Shin, 2010; Gorton and Metrick, 2012; Singh, 2011a; Pozsar et al., 2013)). A central lesson in this work is that balance-sheet capacity can be created endogenously through collateral chains in normal times, yet the same chains can unwind abruptly when collateral constraints tighten, producing forced deleveraging and fire-sale spillovers. From a legal and regulatory perspective, Allen (2023) argues that DeFi replicates the core vulnerabilities of traditional shadow banking—heightened complexity, leverage, rigidity, and susceptibility to runs—and advocates precautionary regulation to contain these risks before they become systemic. Our contribution is to formalize and quantify this channel: we show how the shadow-banking mechanism is implemented in a "bankless" setting through wrapped stablecoins that convert a base stablecoin into a claim optimized for repeated pledging across protocols, so collateral reuse becomes a primitive of on-chain leverage and its unwind.

A large literature on leveraged collateralized trading studies how margin requirements and haircuts tie leverage to private balance-sheet capacity (Gromb and Vayanos, 2002; Fostel and Geanakoplos, 2008; Brunnermeier and Pedersen, 2009; Geanakoplos, 2010). Much of this work focuses on shocks to the value of collateral that trigger deleveraging and price declines. We complement this perspective by showing that the ability to reuse collateral itself can be an independent and state-dependent source of instability: recursive

pledging generates an endogenous collateral (liquidity) multiplier in normal times, but it also makes leverage highly sensitive to tightening collateral constraints and liquidation technology in stress.

Our analysis also relates to the literature on runs, liquidity transformation, and fire-sale dynamics (e.g., (Diamond and Dybvig, 1983; Allen and Gale, 1998; Bernardo and Welch, 2004; Goldstein and Pauzner, 2005)). It has also been shown that runnable claims arise when investors can withdraw at favorable terms, as in money market funds with fixed-value redemptions (Kacperczyk and Schnabl, 2013; Sunderam, 2015; Parlatore, 2016; Schmidt, Timmermann, and Wermers, 2016), while claims traded at market prices are often viewed as less exposed to classic run incentives (Jacklin, 1987; Allen and Gale, 1998; Farhi, Golosov, and Tsyvinski, 2009; Koont et al., 2022). Recent work has begun to connect these traditional finance fragilities directly to the crypto ecosystem. Anadu et al. (2024) document that stablecoins exhibit run and flight-to-safety dynamics strikingly similar to money market funds, with investors fleeing riskier stablecoins for perceived safe havens during stress episodes. Their evidence establishes stablecoins as functional analogues of MMF shares in the digital asset ecosystem, reinforcing the relevance of the shadow-banking fragility channel we study. Our contribution extends this insight by showing that even after a stablecoin survives a run on its base layer, the wrapping and rehypothecation of that stablecoin creates an additional amplification channel through recursive collateral chains that operates independently of—and on top of—the base-layer run dynamics they document. When collateral capacity is created through rehypothecation and enforced through mechanical margin rules, stress can propagate through abrupt liquidation cascades rather than through withdrawal queues alone, even when positions are marked to market.

In contrast, a large literature on financial networks and systemic risk studies how the topology of interconnected balance sheets maps shocks into aggregate outcomes Eisenberg and Noe (2001); Elliott, Golub, and Jackson (2014). Our environment shares the emphasis on network structure, but differs in the source of link strength: interconnections are endogenously shaped by collateral schedules and recursive reuse, so systemic risk depends jointly on network topology and on the terms governing collateral creation, haircuts, and liquidation. This provides a tractable way to link cross-protocol heterogeneity to the severity of deleveraging during stress states.

More narrowly, our paper relates to the growing literature on stablecoins and their stability. Gorton et al. (2022) emphasize the role of stablecoins in facilitating leveraged crypto trading and the link between stablecoin demand and stability. Lyons and Viswanath-Natraj (2023) study creation and redemption activity and its response to secondary-market price deviations. Uhlig (2022) and Liu, Makarov, and Schoar (2023) analyze runs and depegging dynamics around the Terra/Luna episode, while Ahmed, Aldasoro, and Duley (2024) study how information about reserve asset quality affects run dynamics. Related theoretical work studies stablecoin design and stability under commitment and policy rules (d'Avernas, Maurin, and Vandeweyer, 2022; Routledge and Zetlin-Jones, 2022; Li and Mayer, 2021), and empirical work develops measures of stablecoin devaluation risk (Eichengreen, Nguyen, and Viswanath-Natraj, 2023). Our contribution to this literature is to shift the lens from issuance and redemption to secondary use: we study how wrapping

and rehypothecation turn stablecoins into rehypothecatable collateral that endogenously generates a liquidity multiplier in normal times and amplifies stress through mechanical collateral constraints and liquidation cascades.

Finally, our framework speaks to macroprudential policy with pecuniary externalities, where private choices ignore spillovers generated by forced deleveraging and fire sales Jeanne and Korinek (2010); Lorenzoni (2008); Bianchi and Mendoza (2018). Allen (2023) argues from a regulatory standpoint that precautionary intervention is warranted before DeFi leverage channels become interconnected with the broader financial system, a view our formal analysis supports. In our setting, the relevant policy levers are the on-chain analogues of collateral regulation: loan-to-value limits, haircuts (including stress-contingent haircuts), and constraints on recursive reuse. Because these parameters govern the sufficient statistics for amplification, they can improve resilience even when the underlying collateral is unchanged.

3 Model

We consider an infinite-horizon economy with a continuum of identical traders who can leverage their positions through a decentralized finance (DeFi) protocol. Time is discrete and indexed by $t = 0, 1, 2, \ldots$ The economy consists of three assets: a base stablecoin S that is fully backed by fiat and typically pegged to the dollar but can deviate during stress periods, a wrapped version W that represents claims on deposited stablecoins within a DeFi protocol, and a volatile crypto-asset V which pays dividends. A two-state Markov chain $s_t \in \{0, 1\}$ governs aggregate conditions, where $s_t = 0$ is normal and $s_t = 1$ is stress. During stress, the effective value of collateral becomes impaired: wrapped tokens face elevated failure risk $\lambda(s_t)$ and the base stablecoin may trade below par with an effective collateral haircut $\kappa(s_t) \leq 1$.

Our model's key innovation is the recursive leverage mechanism built by the wrapped stablecoin. When a trader deposits one unit of base stablecoin into the protocol, they receive one unit of wrapped token in return. In turn, this wrapped token can be posted as collateral for borrowing additional base stablecoins, subject to a loan-to-value ratio $\theta \in (0, 1)$. Subsequently, the borrowed stablecoins can be re-deposited to mint more wrapped tokens which creates a recursive process that theoretically allows leverage up to $1/(1-\theta)$. However, in practice protocol fees denoted as τ and operational frictions reduce the effective leverage below this.

We assume the DeFi protocol operates under perfect competition and offers two services which are wrapping stablecoins and providing collateralized loans. Unlike standard DSGE models that assume a fixed money supply, we model the base stablecoin supply S_t as demand-driven. The modeling choice aligns with the literature like Gorton et al. (2022) and Arner, Auer, and Frost (2020) which shows that the primary driver of stablecoin demand is in their role as collateral in lending markets and yield-generating DeFi protocols. The centralized issuer expands supply passively in response to user demands for redemption or deposits which aligns with real-world behavior of stablecoins like USDC where supply adjusts endogenously with user's need for transactional liquidity.

The stress regime s_t is the sole source of aggregate uncertainty in the baseline model.

All collateral frictions—the peg haircut κ_t, execution discount δ_t, settlement delay Δ_t, and failure hazard λ_t—are deterministic functions of s_t. Protocol fees τ_t are determined endogenously by competition (Appendix B) or by the constrained-efficient policy rule. Realized volatility is an endogenous outcome of liquidation dynamics rather than an exogenous process. This parsimonious structure ensures clean identification: every friction traces back to a single observable state.

We complement the literature by providing a tractable model and empirical validation that makes several key contributions to understanding systemic risk in DeFi markets. First, we construct a recursive equilibrium framework that captures how wrapped stablecoins create recursive leverage through rehypothecation, which generates a liquidity multiplier of $1/(1-\theta)$ that amplifies both credit creation and systemic fragility. Then, we show that wrapped tokens trade at a discount to base stablecoins during stress periods, with this basis widening as confidence deteriorates and protocol fees increase. Finally, our welfare analysis reveals that decentralized protocols systematically over-issue wrapped claims relative to the constrained-efficient allocation, and that state-contingent policy instruments—implementable through existing smart contract parameters—can correct this inefficiency at lower cost than static leverage caps.

3.1 Agents

3.1.1 Representative Trader

A continuum of identical traders indexed by $i \in [0,1]$ populate the economy. Each trader has standard CRRA preferences:

$$U(c_t) = \frac{c_t^{1-\sigma}}{1-\sigma}. \tag{1}$$

Traders form rational expectations over the transition of the stress state s_t (defined in Section 3.2.1) and over equilibrium asset prices. There is no heterogeneity in beliefs or information.

Each trader faces an intertemporal budget constraint incorporating both standard asset accumulation and collateralized borrowing through DeFi protocols:

$$c_t + B_{t+1} + S_{t+1} + P_t^W W_{t+1} + P_t^V V_{t+1} = R_t^f B_t + (1+r_t^S) S_t + (1+r_t^W) P_t^W W_t + (P_t^V + d_t) V_t - L_t, \tag{2}$$

where $S_t = S_t^{\text{wallet}} + S_t^{\text{vault}}$ represents total stablecoin holdings. The depeg factor $\kappa_t \in (0,1]$ equals one when the stablecoin maintains its peg and falls below one when the peg is impaired; its determination is governed by the stress regime defined in Section 3.2.1. The trader's optimal portfolio, leverage, and consumption decisions are characterized by the Euler equations and collateral constraints derived in the recursive competitive equilibrium (Section 3.5).

3.1.2 Non-Optimizing Institutional Agents

The model includes three institutional agents whose behavior is governed by market protocols and competitive behavior.

DeFi Lending Protocol. The DeFi lending protocol operates as a competitive intermediary that provides two integrated services.

- A **wrapping module** that converts one unit of the base stablecoin S into one unit of a non-redeemable wrapped receipt W.
- A **collateralized loan book** where traders can borrow base stablecoins against posted collateral including base stablecoins S, wrapped tokens W and volatile crypto assets V subject to a loan-to-value ratio $\theta \in (0, 1)$.

The protocol's zero profit condition requires that expected loan repayments equal the interest paid to wrapped token holders plus any liquidation costs:

$$\mathbb{E}_t\left[\text{loan repayments}\right] = (1 + r_t^W)W_{t,i} + \text{expected liquidation costs},$$

where r_t^W represents the endogenously determined interest rate on wrapped tokens.

Arbitrageurs and Liquidity Providers maintain price consistency and maximize expected discounted profits by exploiting price discrepancies between wrapped and base stablecoins.

Centralized Stablecoin Issuer serves as the bridge between the traditional finance and the DeFi ecosystem. It maintains full fiat reserves that back the base stablecoin in cash or short-term securities. It is important to note that the centralized stablecoin issuer acts passively and its sole function is maintaining the convertibility between on-chain stablecoins and off-chain fiat. The separation between base money issuance and credit creation mirrors traditional banking but without regulatory oversight of the credit multiplication process. The centralized issuer manages off-chain reserves backing the base stablecoin S, which is fully redeemable at par for fiat. It passively holds reserves in cash or short-term securities, does not lend or rehypothecate funds, and does not engage in on-chain DeFi activities.

3.2 The Leverage Mechanism

3.2.1 Stress Regime

The aggregate state of the DeFi ecosystem is governed by a two-state Markov chain $s_t \in \{0, 1\}$, where $s_t = 0$ denotes normal conditions and $s_t = 1$ denotes stress. Transition probabilities are:

$$\Pr(s_{t+1} = 1 \mid s_t = 0) = p_{01}, \qquad \Pr(s_{t+1} = 0 \mid s_t = 1) = p_{10}.$$

All collateral frictions are deterministic functions of this state:

$$\kappa_t = \kappa(s_t), \quad \delta_t = \delta(s_t), \quad \lambda_t = \lambda(s_t), \quad \Delta_t = \Delta(s_t),$$

with $\kappa(0) = 1$, $\delta(0) = 1$, $\lambda(0) \approx 0$, $\Delta(0) = 0$ in normal times, and $\kappa(1) < 1$, $\delta(1) < 1$, $\lambda(1) > 0$, $\Delta(1) > 0$ in stress. This parsimonious structure ensures that the model has a single identified source of aggregate uncertainty while preserving the rich micro-structure of DeFi stress channels described below.

3.2.2 Collateral Constraints

To formalize the claim that W acts as a near-perfect substitute for S in transactions but becomes weaker as collateral under stress, we introduce three frictions, each a deterministic function of the stress state s_t. First, $\Delta(s_t)$ represents settlement/withdrawal frictions (e.g., queues, pause switches, circuit breakers) that prevent immediate redemption or reuse. Second, $\delta(s_t) \in (0, 1]$ is an execution discount summarizing the quantitative effect of these lags on realized payoffs. Third, $\lambda(s_t) \in [0, 1)$ is a survival hazard associated with smart-contract failure or oracle manipulation, with $(1 - \lambda(s_t))$ representing the survival probability. In stress states, the effective payoff of W is scaled by $(1 - \lambda(s_t))\,\delta(s_t)$, while the base stablecoin S is marked to a peg haircut $\kappa(s_t) \in (0, 1]$. Together, these frictions imply that $W \approx S$ in tranquil times, but $W < S$ as collateral during stress which produces a stress-time basis that arises endogenously rather than being assumed.

Traders can borrow up to a fraction θ of their collateral value subject to the DeFi's protocol loan-to-value ratio. During normal times, the constraint is:

$$L_{t,i} \leq \theta \left(S_{t+1,i} + P_t^W W_{t+1,i} + P_t^V V_{t+1,i} \right) \tag{3}$$

In stress, collateral values are impaired by three forces: a peg haircut $\kappa_t \in (0, 1]$ on base stablecoins, execution lags summarized by $\delta_t \in (0, 1]$ (arising from Δ_t such as withdrawal queues, pause switches, and oracle/auction timing), and a survival hazard $\lambda_t \in [0, 1)$ for wrapped claims due to smart-contract or oracle risk. Incorporating these frictions yields the effective constraint:

$$L_{t,i} \leq \theta \left(\kappa_t S_{t+1,i} + (1 - \lambda_t)\,\delta_t\, P_t^W W_{t+1,i} + P_t^V V_{t+1,i} \right) \tag{4}$$

Conditional on the stress state s_t, the aggregate hazard $\lambda(s_t)$ reflects several operational risk channels that are first-order in DeFi. We decompose it into bridge risk, oracle manipulation, governance attack, and gas congestion:

$$\lambda(s_t) \equiv 1 - \prod_{r \in \{br, or, gov, gas\}} \big(1 - \lambda^r(s_t)\big), \tag{5}$$

where $\lambda^r(0) \approx 0$ for all r in normal times and $\lambda^r(1) > 0$ in stress, so that any component channel increases the aggregate hazard. These components map into (i) execution quality and (ii) fee spikes.

(i) Execution discount/lag. Oracle manipulation risk and gas congestion worsen liquidation execution and delay auction clearing:

$$\delta(s_t) = \delta(\lambda^{or}(s_t), \lambda^{gas}(s_t)), \qquad \Delta(s_t) = \Delta(\lambda^{gas}(s_t)), \tag{6}$$

with $\partial\delta/\partial\lambda^{or} \leq 0$, $\partial\delta/\partial\lambda^{gas} \leq 0$, and $\partial\Delta/\partial\lambda^{gas} \geq 0$.

(ii) Fee spikes. These risks also raise the effective wedge faced by borrowers/issuers through insurance premia, risk buffers, and congestion costs. We decompose the wedge as

$$\tau_t = \tau_t^{base} + \tau^{risk}(s_t), \qquad \tau^{risk}(s_t) = \tau^{risk}(\lambda^{br}(s_t), \lambda^{gov}(s_t), \lambda^{gas}(s_t)), \tag{7}$$

with $\partial\tau^{risk}/\partial\lambda^{br} \geq 0$, $\partial\tau^{risk}/\partial\lambda^{gov} \geq 0$, and $\partial\tau^{risk}/\partial\lambda^{gas} \geq 0$. The base fee τ_t^{base} is determined endogenously by protocol competition (Appendix B) or by the constrained-efficient LM-cap rule (Section 3.7).

This decomposition clarifies the empirical mapping: bridge/oracle/governance/gas shocks affect the model only through the wrapped failure probability (λ_t), execution quality (δ_t, Δ_t), and the wedge (τ_t).

Pricing convention. Throughout the paper we adopt an SDF-based pricing paradigm. The wrapped-collateral price q_t^W satisfies the Euler equation

$$q_t^W \, u'(c_t - G(h_t)) = \beta \, \mathbb{E}_t\left[u'(c_{t+1} - G(h_{t+1}))\left(1 + r_{t+1}^W + \theta_{t+1}\, m_{t+1}\right) q_{t+1}^W\right], \tag{8}$$

where the stochastic discount factor $\beta\, u'(c_{t+1})/u'(c_t)$ prices the wrapped claim inclusive of its yield r_{t+1}^W and collateral service flow $\theta_{t+1}\, m_{t+1}$. For the base stablecoin, we normalize $q_t^S = 1$ at all times and treat peg impairment $\kappa_t < 1$ as a collateral haircut only: when the peg weakens, the protocol marks down the collateral value of S to κ_t in the borrowing constraint (equation (14)) but the unit of account remains the stablecoin itself. This avoids conflating price deviations with collateral adequacy and ensures consistency across constraints, asset prices, and liquidation logic.

3.2.3 Recursive Leverage Multiplier

Our recursive leverage framework builds on the rehypothecation literature in traditional finance. Aitken and Singh (2010) illustrate how collateral is reused multiple times in the shadow banking system, creating what they term a 'collateral multiplier.' Singh (2011b) further quantifies this phenomenon by showing how the velocity of collateral (i.e., the rate at which it is reused) amplifies financial system leverage. Hence, we formalize an analogous mechanism for DeFi settings, where smart contracts enable transparent and automatic collateral reuse through wrapped tokens.

Definition (Recursive Leverage Process). Consider a DeFi lending protocol with loan-to-value ratio $\theta \in (0, 1)$ and protocol fee $\tau \in [0, 1-\theta)$. An agent with initial stablecoin holdings S_0 recursively: (i) deposits stablecoins as collateral, (ii) borrows fraction $(\theta - \tau)$ against the collateral, (iii) re-deposits borrowed funds as new collateral, (iv) repeats until economically negligible.

This recursive process generates geometric amplification of the initial capital. Therefore, the total amount of collateral deposited through this process is:

$$D = S_0 \cdot \sum_{k=0}^{\infty} (\theta - \tau)^k = \frac{S_0}{1 - \theta + \tau}$$

which shows how small sums of initial capital can be turned into significantly larger collateralized positions. The term $\frac{1}{1-\theta+\tau}$ represents a DeFi collateral multiplier, parallel to the traditional shadow banking collateral multiplier.

The infinite-series expression above is a useful benchmark, but real-world recursion is finite because the effective per-loop wedge increases with loop depth (gas/MEV costs, latency and liquidation risk, borrow caps, and rising price impact). We capture this by

allowing the effective wedge to depend on loop depth n:

$$\tau_{\text{eff}}(n) = \tau + \phi_1 n + \phi_2 n^2, \qquad \phi_1 \geq 0,\ \phi_2 > 0. \tag{9}$$

At depth n, the net reinvestment factor becomes $\theta - \tau_{\text{eff}}(n)$, which declines in n and eventually makes further looping unattractive or infeasible. The resulting deposited collateral is

$$D = S_0 \sum_{k=0}^{N_t} \prod_{j=1}^{k} \big(\theta - \tau_{\text{eff}}(j)\big), \tag{10}$$

where the endogenous maximum recursion depth N_t is the largest integer such that

$$\theta - \tau_{\text{eff}}(N_t) > 0 \quad \text{and} \quad \theta - \tau_{\text{eff}}(N_t + 1) \leq 0. \tag{11}$$

Thus, even when θ is high, convex depth frictions deliver a finite effective multiplier consistent with the empirically observed range (e.g., approximately 3–5× loops in typical environments).

Inadvertently, the convexity of the recursive leverage ratio in θ means that protocols competing to attract users will raise loan-to-value ratios significantly and create explosive leverage growth. Appendix B formalizes protocol competition over (θ, τ) and shows that a sector-wide LM-cap (minimum surcharge) continues to implement the constrained-efficient allocation under competition. It is important to note that unlike traditional rehypothecation that unwinds gradually through bilateral negotiations, DeFi liquidations execute automatically and simultaneously. This in turn, transforms individual failures into market-wide deleveraging spirals. The system's only endogenous stabilizer is protocol fees τ that act as the "sand in the gears" that limit maximum achievable leverage.

3.2.4 Liquidity Multiplier

We formalize the mechanism by which wrapped stablecoin expand the effective money supply in DeFi, extending the classical money multiplier framework to DeFi. Whereas Carlstrom and Fuerst (1997) illustrate the process where liquidity assets translate into credit through financial intermediaries, our analysis demonstrates how tokenized collateral receipts create synthetic liquidity that appears identical to base money from a transactional perspective. Unlike traditional banking where regulatory reserves constrain the money multiplier, DeFi protocols are only constrained by technological design and economic incentives. Hence, the parameter θ is endogenously shaped by protocol governance and market competition instead of regulatory fiat. This distinction is important because the absence of regulatory constraints allows leverage to reach extreme levels and the synthetic nature of wrapped token liquidity creates fragility that is absent in traditional systems. Furthermore, we will show that these features combine to generate the severe volatility spikes observed empirically (e.g., 23% spike in ETH volatility following the March 2023 USDC-depeg).

Definition (Liquidity Multiplier). Let S_t represent the total circulating supply of a base stable coin at t, and W_t denote the total outstanding wrapped versions (e.g., aUSDC,

cDAI, yvUSDC). The liquidity multiplier LM_t quantifies the ratio of total liquid claims to base money:

$$LM_t \equiv \frac{S_t + W_t}{S_t}$$

where wrapped tokens W_t are treated as perfect substitutes for base stablecoins in transaction services but not in collateral value.

Proposition (Equilibrium Liquidity Multiplier). In equilibrium with loan-to-value ratio $\theta \in (0, 1)$, full capital deployment, and binding collateral constraints, the liquidity multiplier equals:

$$LM^* = \frac{1}{1-\theta}$$

Proof. Let D_t represent the total deposits supporting wrapped issuance. Under full deployment (where all wrapped tokens serve as collateral for new loans), the recursive collateralization process yields:

$$D_t = S_t^{\text{dep}} + \theta S_t^{\text{dep}} + \theta^2 S_t^{\text{dep}} + \cdots = \frac{S_t^{\text{dep}}}{1-\theta}$$

Since $W_t = D_t$ in equilibrium and market clearing requires $S_t^{\text{dep}} = S_t^{\text{vault}} \leq S_t$:

$$LM = \frac{S_t + W_t}{S_t} = \frac{S_t + \frac{S_t^{\text{dep}}}{1-\theta}}{S_t}$$

Thus full deployment $S_t^{\text{dep}} = \frac{\theta S_t}{1-\theta}$ results in:

$$LM = 1 + \frac{\theta}{1-\theta} = \frac{1}{1-\theta}$$

Proposition (Bounds on Liquidity Creation). The liquidity multiplier satisfies:

$$1 \leq LM_t \leq \frac{1}{1-\theta}$$

with $LM_t = 1$ iff $W_t = 0$ (no wrapping) and $LM_t = \frac{1}{1-\theta}$ iff all efficient entrepreneurs achieve full collateralization.

Proof. The lower bound follows from $W_t \geq 0$. For the upper bound, note that maximum wrapped issuance happens when $S_t^{\text{dep}} = S_t$ and all tokens are recursively rehypothecated with parameter θ.

Corollary Formally, liquidity multiplier satisfies the following properties:

$$\frac{\partial LM}{\partial \theta} = \frac{1}{(1-\theta)^2} > 0 : \text{ Higher LTV ratios increase liquidity creation}$$

$$\frac{\partial^2 LM}{\partial \theta^2} = \frac{2}{(1-\theta)^3} > 0 : \text{ Liquidity creation is convex in leverage}$$

$$\lim_{\theta \to 1} LM = \infty : \text{ Perfect collateralization yields unbounded liquidity}$$

The liquidity multiplier $1/(1-\theta)$ is formally equivalent to the traditional money multiplier $1/rr$ where rr represents the reserve requirement and $1-\theta$ denotes the effective reserve of unwrapped collateral. This parallel extends the traditional monetary theory of banking Tobin (1963) to algorithmic finance and highlights how in both systems the ability to expand liquidity hinges on the fraction of base assets withheld from reuse (whether by regulation or protocol design).

During times of stress, the market applies haircuts to stablecoin collateral which violates previous assumptions of base stablecoins maintaining par value. For example, following the March 2023 USDC crisis the token was traded at 0.87 and lending protocols marked down collateral values accordingly. To measure effective liquidity under impaired collateral conditions and accounting for protocol fees, we extend our framework to the following.

Definition (Fee-Adjusted Liquidity Multiplier). For protocol m, let $\theta_m \in [0,1)$ denote the protocol-specific loan-to-value ratio, and let $\tau_{m,t} \in [0, 1-\theta_m)$ denote the borrow–supply rate spread (protocol fees and frictions). The baseline liquidity multiplier under peg ($\kappa_t = 1$) is:

$$LM_{m,t} = \frac{1}{1-\theta_m+\tau_{m,t}}. \tag{12}$$

This is the object we estimate in the empirical analysis.

Definition (Depeg-Adjusted Liquidity Multiplier). Let $\kappa_t \in (0,1]$ represent the market-implied value of the base stablecoin. The depeg-adjusted multiplier is:

$$LM_{m,t}^{\text{depeg}} = \frac{1}{\kappa_t(1-\theta_m+\tau_{m,t})} = \frac{LM_{m,t}}{\kappa_t}. \tag{13}$$

The depeg-adjusted multiplier captures how even modest depegging events can amplify into system-wide instability through three distinct channels:

1. A decline in κ_t reduces effective base collateral $\kappa_t S_t$ relative to outstanding claims, thereby increasing LM_t^{depeg}. Whereas a rise in τ_t enlarges the denominator $1-\theta+\tau_t$ and thus reduces LM_t^{depeg}.

2. This rise results in a fragile situation where synthetic liquidity appears high but in reality the reserves backing it are effectively impaired and creation costs are elevated.

3. In the limit, as $\kappa_t \to 0$ or $\tau_t \to (1-\theta)$, effective reserves vanish and the multiplier tends toward infinity, representing a complete collapse of peg credibility and a breakdown in trust.

The depeg-adjusted multiplier LM_t^{depeg} captures the procyclical dynamics that escalate confidence shocks into systemic crises. When confidence deteriorates—indicated by a drop in κ_t or rising stress-induced fees τ_t—the effective multiplier increases, magnifying the value of any existing leverage. This pushes leveraged positions beyond sustainable thresholds. In turn, the resulting feedback loop drives a downward spiral of collapsing collateral values and amplifies the initial shock.

3.3 The DeFi Financial Accelerator

While the liquidity multiplier LM_t^{depeg} provides a static measure of synthetic liquidity, it obscures the fragility embedded within the system. Hence, we adapt the financial accelerator framework of Bernanke and Gertler (1995) to DeFi to show how shocks to collateral values create endogenous feedback loops that amplify initial disturbances. The DeFi version is distinct from traditional financial accelerators because it is driven by automated smart contract liquidations and recursive rehypothecation. This means that timescales are compressed and responses are magnified, resulting in even more abrupt transitions from stability to collapse.

Definition (DeFi Financial Accelerator) refers to a recursive feedback mechanism whereby declines in stablecoin confidence κ_t trigger automatic liquidations. This in turn further erodes κ_t through market impact, amplifying initial shocks into systemic crises—especially when synthetic leverage denoted by LM_t^{depeg} is high.

Formally, the evolution of stablecoin confidence κ_t follows:

$$\kappa_{t+1} = \kappa_t - \phi(\ell_t, LM_t^{\text{depeg}})$$

where $\ell_t = 1 - G(\theta/\kappa_t)$ is the mass of liquidated positions and $\phi(\cdot)$ captures their endogenous price impact. When κ_t declines, more positions breach the liquidation threshold θ/κ_t, which intensifies deleveraging.

The system's fragility is governed by the feedback elasticity:

$$\varepsilon_t = \left(\frac{\partial \phi}{\partial \ell}\right) \cdot \left(\frac{\partial \ell}{\partial \kappa}\right) \cdot \kappa_t$$

where $\varepsilon_t > 1$ indicates explosive amplification. This is the DeFi version of a tipping point. In this regime, each dollar of liquidation contracts LM_t^{depeg} dollars of synthetic liquidity, creating nonlinear responses to modest shocks.

3.4 Network Structure and Contagion

We identify a previously uncharacterized form of systemic risk in decentralized finance: cross-protocol rehypothecation risk that comes from the recursive deployment of tokenized collateral across multiple DeFi platforms. In practice, wrapped stablecoins like aUSDC are routinely deposited as collateral in MakerDAO or Compound which generates layered leverage that the current "single-protocol" framework cannot capture. For example, the March 2023 USDC depeg event, triggered by Circle's $ 3.3 billion exposure to the failed Silicon Valley Bank caused USDC's decline to $0.87 and forced MakerDAO to implement emergency measures to address their $3.1 billion USDC collateral exposure. This in turn triggered widespread liquidations across Aave, Compound and other protocols because their USDC-backed positions became undercollateralized Chainalysis Team (2023). Unlike previous DeFi risks that operate within protocol boundaries we show that this operates at a network level and generates amplification effects that existing frameworks cannot detect or measure. Moreover, recent work Bappy, Cheon, and Islam (2025) has suggested that users' reliance on a few centralized platforms effectively concentrates network flows which supports our use of a dense rehypothecation matrix A and motivates modeling systemic fragility as a property of interconnections.

3.4.1 Rehypothecation Matrix

Our recursive-leverage framework builds directly on Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015) and Eisenberg and Noe (2001) by extending their network clearing models to incorporate endogenous velocity and multiplicative leverage effects. Hence, we propose a network adjacency matrix that provides the first systematic approach to quantify these contagion pathways. We will formalise cross-protocol rehypothecation by:

$$\mathbf{C}_t^{\text{eff}} = (I - \mathbf{A})^{-1}\mathbf{C}_t \tag{14}$$

where:

- $\mathbf{C}_t^{\text{eff}}$ is the effective collateral exposure across protocols at time t,
- I is the identity matrix,
- $\mathbf{A}$ is the rehypothecation matrix, where A_{ij} is the fraction of protocol i's collateral redeposited into protocol j,
- $\mathbf{C}_t$ is the vector of raw (visible) collateral holdings in each protocol at time t.

Equation (14) is well-defined only when the rehypothecation network does not "blow up" collateral through infinite recursion. We therefore restrict attention to the admissible set

$$\mathcal{A} \equiv \{A_t \in \mathbb{R}_+^{n\times n} : \rho(A_t) < 1\},$$

where $\rho(A_t)$ is the spectral radius. On $\mathcal{A}$, $(I - A_t)$ is invertible and admits the convergent Neumann series expansion

$$(I - A_t)^{-1} = \sum_{k=0}^{\infty} A_t^k,$$

so

$$C_t^{\text{eff}} = (I - A_t)^{-1}C_t = \sum_{k=0}^{\infty} A_t^k C_t.$$

This representation makes clear that C_t^{eff} aggregates all rehypothecation "paths" of length $k = 0, 1, 2, \ldots$, and $\rho(A_t) < 1$ guarantees that the total effective collateral exposure is finite.

To capture stress-time "isolation" (fragmentation of collateral chains), we allow the rehypothecation matrix to be state-dependent. Let $s_t \in \{0, 1\}$ be a stress indicator (e.g., $\kappa_t < 1$ or high liquidation mass). We model isolation as a reduction in network intensity or connectivity under stress:

$$A_t = \chi(s_t)\,\bar{A} \quad \text{with} \quad 0 < \chi(1) \le \chi(0) \le 1,$$

or, more generally, as sparsification $A_t = \bar{A} \odot M_t$ where M_t is a state-dependent adjacency mask that drops edges when bridges/oracles are congested or protocols pause withdrawals. In either case, $\rho(A_t)$ falls in stress, reflecting reduced cross-protocol reuse of collateral and more localized loss absorption. This mechanism maps directly into the phase-transition logic: the interaction between leverage θ and interconnection $\rho(A_t)$ governs whether shocks remain localized or become systemic.

3.4.2 Payment Equilibrium

We extend the notion of payment equilibrium introduced by Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015) to model the liquidation and repayment dynamics within an interconnected web of DeFi protocols. For the purposes of our setting, the nodes correspond to protocols or smart contracts and directed edges are the rehypothecated obligations.

Definition (DeFi Payment Equilibrium). For a given rehypothecation network $\{y_{ij}\}$ a vector of repayments $\{x_{ij}\}$ and liquidation decisions $\{\ell_i\}$ is the *payment equilibrium* if:

1. Each protocol's repayment and liquidation behavior clears the constraints:

$$\ell_i = \left[\min\left\{\frac{v_i + y_i - h_i}{\lambda_i},\, A_i\right\}\right]_+, \quad x_{ji} = \frac{y_{ji}}{y_i}\left[\min\left\{y_i, h_i + \lambda_i \ell_i - v_i\right\}\right]_+,$$

   where $h_i = c_i + z_i + \sum_{j \neq i} x_{ij}$ and $y_i = \sum_{j \neq i} y_{ji}$.

2. The realized repayments $\{x_{ij}\}$ are mutually consistent given endogenous liquidations $\{\ell_i\}$.

3. The system clears such that all protocols simultaneously optimize their obligations using available liquidity and liquidation proceeds.

We next establish existence and (under regularity) uniqueness of the payment equilibrium via a mapping to the classical clearing framework of Eisenberg and Noe (2001).

Existence and Uniqueness of the Payment Equilibrium. Let $y_{ji} \geq 0$ denote the nominal obligation of protocol i to creditor j, and define total nominal obligations

$$\bar{p}_i \equiv y_i := \sum_{j \neq i} y_{ji}.$$

Following Definition 3.4.2, proportional repayment implies that if protocol i makes total payment $p_i \in [0, \bar{p}_i]$, then creditor j receives

$$x_{ji} = \Pi_{ij}\, p_i, \qquad \Pi_{ij} := \begin{cases} y_{ji}/y_i, & y_i > 0, \\ 0, & y_i = 0. \end{cases}$$

Hence the total inflow to protocol i from the network equals

$$\sum_{j \neq i} x_{ij} = \sum_{j \neq i} \Pi_{ji}\, p_j = (\Pi^\top p)_i.$$

Let the pre-liquidation liquid resources of protocol i be

$$h_i(p) := c_i + z_i + (\Pi^\top p)_i,$$

and recall the liquidation rule in Definition 3.4.2:

$$\ell_i(p) = \left[\min\left\{\frac{v_i + \bar{p}_i - h_i(p)}{\lambda_i},\, A_i\right\}\right]_+.$$

Total funds available for repayment after liquidation are $h_i(p) + \lambda_i \ell_i(p) - v_i$. Using the identity

$$a + \min\{[b-a]_+,\, c\} = \min\{b,\; a+c\},$$

with $a = h_i(p) - v_i$, $b = \bar{p}_i$, and $c = \lambda_i A_i$, we obtain

$$h_i(p) + \lambda_i \ell_i(p) - v_i = \min\{\bar{p}_i,\; h_i(p) - v_i + \lambda_i A_i\}.$$

Therefore the induced total payment by protocol i can be written as

$$p_i = \Phi_i(p) := \min\{\bar{p}_i,\; e_i + (\Pi^\top p)_i\}, \qquad e_i := \left[c_i + z_i - v_i + \lambda_i A_i\right]_+,$$

and the vector of total payments p is a fixed point of the clearing map

$$\Phi(p) = (\Pi^\top p + e) \wedge \bar{p}.$$

This is exactly the Eisenberg and Noe (2001) clearing system. Their Theorem 1 implies that for any financial system $(\Pi, \bar{p}, e)$ there exists at least one clearing payment vector (indeed, greatest and least fixed points). Moreover, if the system is *regular* in the sense of Eisenberg and Noe (2001), then the clearing payment vector is unique. Hence a DeFi payment equilibrium exists, and it is unique under the corresponding regularity condition.

3.4.3 Phase Transition

Proposition (Phase Transition in Recursive Rehypothecation Networks)

Suppose that $\theta \in (0,1)$ is the common loan-to-value ratio imposed by the collateralized DeFi protocol, ϵ denotes the magnitude of an exogenous collateral shock and $\rho(A) \in [0,1)$ is the spectral radius of rehypothecation matrix A, where each element A_{ij} represents the fraction of protocol j's collateral that's recursively deposited into protocol i. Then there exists a critical threshold $\epsilon^*(\theta, \rho(A))$, such that:

(i) If $\epsilon > \epsilon^*(\theta, \rho(A))$ then the recursive rehypothecation structure is the least stable and least resilient DeFi networks.

(ii) If $\epsilon < \epsilon^*(\theta, \rho(A))$, then the recursive rehypothecation network is stable.

Thus, when the magnitude of the shock crosses the critical threshold, the network exhibits a form of phase transition: it flips from being the most to the least stable and resilient network similar to the instability of minimally diversified structures. The intuition is straightforward because each protocol recursively reuses collateral from all others, a single large shock to one node will propagate directly to all of its counterparties. Any subset of protocols where the localized collateral shortfall exceeds available liquidity within its downstream exposure set will cause the system to enter a deleveraging spiral and destabilize the payment equilibrium. In the limit as $\theta \to 1$ or $\rho(A) \to 1$, the system becomes fragile to arbitrarily small shocks.

However, as part (ii) shows not all rehypothecation networks are equally fragile in the presence of large shocks the presence of "weakly connected" substructures which guarantees that the losses rather than being transmitted across the protocol are borne primarily by the distressed protocol. The intuition behind the phase transition is that for sufficiently small shocks, recursive leverage functions as a stabilizing mechanism that spreads

the loss across the network. That in turn allows excess liquidity to absorb the idiosyncratic shortfalls. However, when shocks grow large and collateral chains become increasingly entangled then the same mechanism becomes a channel for amplification. The spectral radius $\rho(A)$ thus represents the extent of endogenous interconnection in rehypothecation chains and the loan-to-value ratio θ denotes the degree of leverage that's applied to each unit of collateral. The interaction between these two components will determine whether the contagion remains localized or turns systemic.

Section 3.4 defines the rehypothecation matrix A over protocols, where A_{ij} is the fraction of protocol i's collateral redeposited into protocol j, and shows that network amplification is governed by the spectral radius $\rho(A)$ via $C_t^{\text{eff}} = (I - A)^{-1} C_t$ in (20). The stress regime $s_t \in \{0, 1\}$ (as defined in Section 3.2.1) determines whether the network operates in the stable or fragile region of this phase transition. Hypothesis H3 is stated at the chain level: cross-chain bridges connect blockchains and determine the speed of contagion. This subsection shows that the same mathematical structure applies when the network nodes are chains and bridges form the adjacency structure.

Let $c \in \{1, \ldots, C\}$ index blockchains (chains). Let $C_{c,t}^{\text{ch}}$ denote the total stablecoin collateral on chain c at time t (aggregating across protocol deployments), and let $C_t^{\text{ch}} \in \mathbb{R}_+^C$ collect these quantities. Define a chain-level reallocation (bridge) matrix $A_t^{\text{ch}} \in \mathbb{R}_+^{C \times C}$ by

$$A_{cd,t}^{\text{ch}} \equiv \mathbb{E}_t \left[\frac{F_{c \to d,t}}{C_{c,t}^{\text{ch}}} \right], \qquad c \neq d,$$

where $F_{c \to d,t}$ is net bridge outflow from chain c to chain d over the relevant stress window. The bridge graph pins down the support of A_t^{ch}: $A_{cd,t}^{\text{ch}} > 0$ only if at least one active bridge connects c to d. (If one prefers an extensive-margin representation, write $A_t^{\text{ch}} = \bar{A}^{\text{ch}} \odot M_t^{\text{br}}$ where M_t^{br} is a bridge adjacency mask.)

The chain-level analogue of (20) is

$$C_t^{\text{ch,eff}} = (I - A_t^{\text{ch}})^{-1} C_t^{\text{ch}},$$

which is well-defined on the admissible set $\{A_t^{\text{ch}} : \rho(A_t^{\text{ch}}) < 1\}$. The Neumann expansion implies $(I - A_t^{\text{ch}})^{-1} = \sum_{k=0}^{\infty} (A_t^{\text{ch}})^k$, so $C_t^{\text{ch,eff}}$ aggregates multi-hop bridge paths of length $k = 0, 1, 2, \ldots$ exactly as in the protocol network. Hence bridge topology affects amplification through the same sufficient statistic, the spectral radius $\rho(A_t^{\text{ch}})$.

[Phase transition in bridge networks] Fix $\theta \in (0, 1)$ and let ε denote the magnitude of an exogenous collateral shock originating on a single chain. Under the chain-level representation above, there exists a critical threshold $\varepsilon_{\text{ch}}^*(\theta, \rho(A^{\text{ch}}))$ such that: (i) if $\varepsilon > \varepsilon_{\text{ch}}^*(\theta, \rho(A^{\text{ch}}))$, bridge-mediated reallocation makes deleveraging self-reinforcing and the shock propagates systemically across chains; and (ii) if $\varepsilon < \varepsilon_{\text{ch}}^*(\theta, \rho(A^{\text{ch}}))$, the shock remains localized. Moreover, $\varepsilon_{\text{ch}}^*$ is decreasing in $\rho(A^{\text{ch}})$, and in the limit $\theta \to 1$ or $\rho(A^{\text{ch}}) \to 1$ the system becomes fragile to arbitrarily small shocks.

For nonnegative matrices, the spectral radius is monotone in entries: if $0 \leq X \leq Y$ (entrywise), then $\rho(X) \leq \rho(Y)$ (and the inequality is strict under standard irreducibility conditions). Adding bridge connections expands the support of A^{ch} (or raises its intensity), which increases $\rho(A^{\text{ch}})$ and therefore lowers $\varepsilon_{\text{ch}}^*(\theta, \rho(A^{\text{ch}}))$. A linearized propagation of

stress across chains satisfies $x_{t+1} \approx \theta A^{\text{ch}} x_t$ in early stages, so the dominant growth rate is governed by $\theta\rho(A^{\text{ch}})$. This implies shorter time-to-arrival of stress on highly connected chains. Empirically, bridge count is a proxy for the extensive margin of connectivity in the bridge graph and therefore for $\rho(A^{\text{ch}})$, yielding the chain-level prediction tested in H3: chains with more bridges experience faster contagion arrival (lower delay).

3.5 Equilibrium

Three market clearing conditions must be satisfied for the equilibrium to clear. First, the stablecoin market clears when total holdings equal the fixed supply: $\int_0^1 S_{t,i}\, di + S_t^{\text{protocol}} = \bar{S}$. Second wrapped token issuance must balance where the protocol's wrapped token liabilities match exactly the sum of the user holdings:

$$\int_0^1 W_{t,i}\, di = W_t$$

. Any discrepancy (whether from smart contract exploits or token burn) immediately manifests as a change in W_t which provides on-chain transparency absent in traditional shadow banking. Third, individual borrowing decisions $L_{t,i}$ must aggregate to the system-wide leverage ratio L_t. This in turn affects equilibrium asset prices and risk dynamics. Together these conditions guarantee that the decentralized actions of traders, protocols, and arbitrageurs produce internally consistent market outcomes even without centralized coordination.

3.6 Recursive Competitive Equilibrium

Following the approach of Bianchi and Mendoza (2011) and Bianchi and Mendoza (2018), in this section we define the recursive competitive equilibrium where agents will make optimal decisions when given aggregate state variables and pricing functions. Our RCE extends previous standard frameworks in two key dimensions. First, we introduce a dual money multiplier that incorporates the wrapped stablecoin multiplier $\frac{W}{S} = \frac{1}{1-\theta}$ (multilayer financial intermediation) unlike the traditional models that focus on single-layer financial intermediation such as Brunnermeier and Sannikov (2014) and Gertler and Kiyotaki (2010). Second, the stress regime s_t introduces state-dependent haircuts: in stress states, wrapped token collateral is impaired through $\kappa(s_t)$, $\delta(s_t)$, and $\lambda(s_t)$, departing from previous models like Kiyotaki and Moore (1997) and Iacoviello (2005) that assumed fixed collateral ratios. We maintain the standard regularity conditions used in Bianchi and Mendoza (2018). In particular, the aggregate state space is compact (we restrict (B_t, W_t, s_t) to bounded sets), and preferences, constraints, and induced payoffs are continuous so the Bellman operator is well-behaved. In addition, the wrapped asset has a bounded, measurable payoff so its equilibrium price satisfies a standard Euler equation under the same stochastic discount factor as other traded assets (see Eq. (21)). Under these maintained assumptions, existence of a recursive competitive equilibrium is standard in this literature, so we proceed by characterizing the RCE via the Bellman problem, Euler equations, and market clearing.

Hence, we denote the state variables as the individual-level asset holdings and aggregate macro-financial conditions.

Individual states for each trader i are given by the triplet $(b_{t,i}, w_{t,i}, v_{t,i})$, where:

- $b_{t,i}$ denotes the trader's holdings of base stablecoins S,
- $w_{t,i}$ denotes the holdings of wrapped stablecoins W, and
- $v_{t,i}$ denotes the holdings of a volatile crypto asset V.

Aggregate states are defined as follows:

- $B_t = \int b_{t,i}\, di$: the total supply of base stablecoins in the economy, where the integration is over a continuum of agents on the unit interval,
- $W_t = \int w_{t,i}\, di$: the total supply of wrapped stablecoins in the economy, also integrated over the unit interval,
- $s_t \in \{0, 1\}$: the stress regime defined in Section 3.2.1, where $s_t = 1$ corresponds to stress with impaired collateral values and elevated failure risk.

Traders maximize their lifetime utility according to:

$$V(b_{t,i}, w_{t,i}, v_{t,i}, B_t, W_t, s_t) = \max_{c_{t,i}, b_{t+1,i}, w_{t+1,i}, v_{t+1,i}} \{u(c_{t,i}) + \beta \mathbb{E}_t [V(\cdot)]\}$$

Subject to:

1. Budget Constraint:

$$c_{t,i} + q_t^S b_{t+1,i} + q_t^W w_{t+1,i} + q_t^V v_{t+1,i} \leq \text{Income} + \text{Collateralized Loans}$$

2. Collateral Constraint:

$$\Lambda_{t,i} \leq \theta \left(q_t^S b_{t,i} + q_t^W w_{t,i} + q_t^V v_{t,i} \right)$$

Solving the intertemporal optimization problems yields the following first-order optimality conditions for risk-free bonds and wrapped stablecoins:

Risk-Free Bonds.

$$1 = \mathbb{E}_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} R_t^f \right]$$

Wrapped Stablecoins. With zero recovery in failure, the wrapped-token Euler is:

$$q_t^W = \mathbb{E}_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} \left(1 - \lambda_{t+1}\right) \delta_{t+1} \left(1 + r_{t+1}^W\right) \right]. \tag{15}$$

Base Stablecoin (S).

$$q_t^S = 1$$

Volatile Crypto Asset (V).

$$q_t^V = \mathbb{E}_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} \left(d_{t+1} + q_{t+1}^V \right) \right]$$

Market Clearing.

$$\int_0^1 b_{t,i}\, di + \text{Vault Reserves} = S, \quad \int_0^1 w_{t,i}\, di = W_t = \frac{S}{1-\theta}, \quad \int_0^1 v_{t,i}\, di = \bar{V}.$$

Definition (Recursive competitive equilibrium). A recursive competitive equilibrium is defined by:

- Value functions $V(b, w, v, B, W, s)$,
- Pricing functions $q^S(B, W, s)$, $q^W(B, W, s)$, $q^V(B, W, s)$,
- Policy functions $b'(b, w, v, B, W, s)$, $w'(\cdot)$, $v'(\cdot)$,
- Laws of motion for aggregates:

$$B' = \Gamma_B(B, W, s), \quad W' = \Gamma_W(B, W, s), \quad s' \sim \text{Markov}(s).$$

such that (i) traders optimize given prices and constraints, (ii) the DeFi protocol satisfies zero-profit and collateral rules, (iii) all markets clear each period, and (iv) perceived and actual laws of motion coincide.

3.7 Welfare Analysis

We analyze two benchmarks: (i) the decentralized equilibrium (DE) with collateral-based issuance limits, and (ii) the constrained-efficient allocation chosen by a time-consistent social planner. In the decentralized economy, issuance limits reduce borrowing by restricting balance-sheet growth and encouraging issuers to have precautionary buffers to mitigate the risk of binding constraints. We build on existing work including Bianchi and Mendoza (2018), Jeanne and Korinek (2010), Lorenzoni (2008) and Bianchi and Mendoza (2011) and demonstrate that when the collateral constraint is non-binding, the decentralized equilibrium with issuance limits still results in more issuance than the social planner would choose. The inefficiency comes from a pecuniary externality, where individual issuers do not internalize the effect of their issuance decisions on future collateral prices, redemption capacity or systemic risk.

To our best knowledge, this is the first application of the pecuniary externality framework from the macroprudential policy literature to the architecture of DeFi protocols.

In our setting, the externality stems from the recursive leverage embedded in wrapped-stablecoin issuance and from the state-dependent liquidity multiplier. These channels amplify shocks in ways not present in conventional financial systems and necessitate the use of policy instruments in the form of protocol-level issuance limits, dynamic fees, and collateral haircuts.

Hence, we formulate the planner's problem in a manner similar to the primal form of Bianchi and Mendoza (2018). The planner chooses allocations subject to the resource constraint, the implementability constraint, and the collateral constraint. The planner cannot set collateral prices directly, but recognizes that issuance decisions alter future prices. We isolate the prudential motive to reduce issuance in calm states from the ex-post motive to support prices in stress states, which provides a normative foundation for state-contingent protocol policies. We focus on Markov stationary policy rules expressed as functions of the payoff-relevant state variables (B, s) where B represents the outstanding stock of wrapped-stablecoin issuance s denotes the exogenous shocks to relevant states. In each period, the planner selects next-period issuance and recognizes both the contemporaneous impact on collateral values and balance-sheet constraints, and the fact that future planners will optimize under the conditions they inherit. The Markov Perfect Equilibrium is a set of policy functions such that: (i) given the policy functions of future planners, the current planner's choices solve its optimization problem, (ii) these policy functions are identical across planners such that no planner has an incentive to deviate from a rule others have followed. Denote the planner's policy for next-period issuance as $I(B, s)$; the function maps the current state into an issuance choice that satisfies the resource constraint, the implementability condition given by the private agent's Euler equation for collateral assets, and the collateral constraint. This embeds the general-equilibrium feedback from issuance to collateral prices and through to the tightness of system-wide borrowing limits.

Let $B(B_t, s_t)$ be the policy rule for next-period wrapped-stablecoin issuance that the social planner takes as given, and let $\{\, C(B_t, s_t),\ H(B_t, s_t),\ v(B_t, s_t),\ m(B_t, s_t),\ q^W(B_t, s_t) \,\}$ be the associated recursive functions that return the values of aggregate consumption, labor effort, protocol-level yield, the multiplier on the collateral constraint, and the wrapped-collateral price under that policy rule. Given these functions, the optimization problem of private agents yields a standard Euler equation for the wrapped collateral asset (see APPENDIX A.1). Incorporating this condition into the following social planner's asset-pricing implementability constraint yields:

$$q_t^W\, u'\big(c_t - G(h_t)\big) = \beta\, E_t\Big[u'\big(C_{t+1} - G(H_{t+1})\big)\,\big(q_{t+1}^W + r_{t+1}^W q_{t+1}^W + \theta_{t+1}(B_{t+1}, s_{t+1})\, m_{t+1}\, q_{t+1}^W\big)\Big] \tag{16}$$

This expression indicates that the choice of B_{t+1} by the planner affects wrapped-collateral prices through two mechanisms. First, it has a direct effect by altering the date-t marginal utility, which is seen in the denominator of the stochastic discount factor. Second, it has an indirect effect by modifying the issuance and balance-sheet positions that future planners will choose, in turn affecting their associated allocations and the resulting market-clearing prices.

As discussed above, the social planner (SP) chooses allocations to maximize the utility of the representative firm–household, subject to the resource constraint, the collateral constraint, and the implementability condition. The planner also considers the optimality

conditions for labor supply and the provision of wrapped-stablecoin collateral services, as well as the Kuhn–Tucker conditions associated with the collateral constraint in the decentralized equilibrium (DE). As shown in Section A.1 of the online appendix, however, these additional conditions are not binding in equilibrium and can therefore be omitted. Given the functions $\{B(B_t, s_t),\, C(B_t, s_t),\, H(B_t, s_t),\, v(B_t, s_t),\, m(B_t, s_t),\, Q^W(B_t, s_t)\}$, the SP's problem can be expressed in recursive form as:

$$
\begin{aligned}
V(B_t, s_t) = \max_{c_t,\, B_{t+1},\, q_t^W,\, h_t,\, v_t} \; & \left\{ u\left(c_t - G(h_t)\right) + \beta\, E_{s_{t+1}|s_t}\, V(B_{t+1}, s_{t+1}) \right\} \\
\text{s.t.} \quad & c_t + \frac{B_{t+1}}{R_t} = B_t + F(1, h_t, v_t) - p^v v_t, \\
& \frac{B_{t+1}}{R_t} - p^v v_t \geq -\theta_t q_t^W, \\
& q_t^W\, u'\left(c_t - G(h_t)\right) \\
& \quad = \beta\, E_{s_{t+1}|s_t} \Big[u'\left(C(B_{t+1}, s_{t+1}) - G(H(B_{t+1}, s_{t+1}))\right) \\
& \qquad \times \Big(Q^W(B_{t+1}, s_{t+1}) + z_{t+1} F_k\left(1, H(B_{t+1}, s_{t+1}), v(B_{t+1}, s_{t+1})\right) [1] \\
& \qquad\quad + \theta_{t+1} m(B_{t+1}, s_{t+1})\, Q^W(B_{t+1}, s_{t+1}) \Big) \Big],
\end{aligned}
\tag{17}
$$

where the final condition represents the implementability constraint implied by the Euler equation for the wrapped-collateral asset.

In the planner's formulation, the multiplier on the resource constraint is denoted $\mu \geq 0$. The collateral constraint has an associated multiplier $m^{\text{sp}} \geq 0$, which may differ from the private-sector multiplier m because the social and private gains from relaxing collateral requirements are not generally the same. The asset-pricing implementability constraint carries the multiplier $\gamma \geq 0$; this condition ensures that the planner's allocation choices are consistent with the price q^W implied by the wrapped-collateral Euler equation in the decentralized economy.

Definition. A recursive constrained-efficient equilibrium in this environment consists of:

- a policy rule for next-period issuance $B'(B, s)$,
- decision rules for consumption $c(B, s)$, labor $h(B, s)$, volatile-asset holdings $v(B, s)$, and the multiplier on the collateral constraint $m(B, s)$,
- the wrapped-collateral pricing function $q^W(B, s)$, and
- the value function $V(B, s)$,

together with:

- a conjectured policy rule of future planners $\mathcal{B}(B, s)$,

[1] Here z_t denotes the aggregate productivity factor following Bianchi and Mendoza (2011), distinct from the stress regime s_t defined in Section 3.2.1.

- the associated future decision rules $C(B,s)$, $H(B,s)$, $v^{\text{fut}}(B,s)$, $m^{\text{fut}}(B,s)$, and
- the future wrapped-collateral pricing function $Q^W(B,s)$,

such that:

1. **Planner's problem:** The value function $V(B,s)$ and the policy functions

$$\{\, B'(B,s),\ c(B,s),\ h(B,s),\ v(B,s),\ q^W(B,s)\,\}$$

solve the Bellman equation in problem (12), given the conjectured future rules

$$\{\, \mathcal{B}(B,s),\ C(B,s),\ H(B,s),\ v^{\text{fut}}(B,s),\ m^{\text{fut}}(B,s),\ Q^W(B,s)\,\},$$

and $m(B,s)$ satisfies the collateral constraint condition.

2. **Time-consistency (Markov stationarity):** The conjectured policy rules for future planners coincide with the optimal rules of the current planner:

$$\begin{aligned} B'(B,s) &= \mathcal{B}(B,s), & c(B,s) &= C(B,s),\\ h(B,s) &= H(B,s), & v(B,s) &= v^{\text{fut}}(B,s),\\ m(B,s) &= m^{\text{fut}}(B,s), & q^W(B,s) &= Q^W(B,s). \end{aligned}$$

The decentralized equilibrium and the time-consistent planner are distinct in 2 ways. First, individual issuers of wrapped stablecoins do not recognize how their borrowing and issuance decisions at date t influence collateral valuations at $t+1$ in states when the constraint on leverage becomes binding. Second, when the constraint binds at t, they fail to account for the fact that their current choices may raise asset prices and expand borrowing capacity (either directly or through their influence on future regulatory decisions). These issues are especially significant in the DeFi setting because recursive leverage and cross-protocol rehypothecation make collateral prices highly sensitive to issuance flows. Hence, we compare the planner's sequential optimality conditions for consumption, borrowing and collateral prices with those of the unregulated equilibrium.

The social planner's optimal conditions can be written as follows:

$$\ell_t = u'(c_t) - y_t\, u''(c_t)\, q_t \tag{18}$$

$$u'(c_t) = \beta R_t\, \mathbb{E}_t\Big[u'(c_{t+1}) - y_{t+1}u''(c_{t+1})Q_{t+1} + y_t Q_{t+1}\Big] \tag{19}$$

$$y_t = \kappa_t\, \frac{m_t^*}{u'(c_t)} \tag{20}$$

Q_{t+1} represents how a marginal increase in next-period borrowing b_{t+1} impacts the date-t collateral price q_t through the future planner's policies. In our context, this sensitivity has three DeFi-native channels: (i) a *confidence/marginal-utility channel*, where more issuance raises tomorrow's effective liquidity multiplier which makes constrained

states likelier and pushes up $u'(c_{t+1})$. This term enters negatively in the planner's Euler equation; (ii) a *price–liquidity channel*, where extra issuance temporarily lifts q_{t+1} by increasing collateral demand in tranquil regimes (typically positive); and (iii) a *constraint–network channel*, where higher b_{t+1} tightens liquidation thresholds system-wide via high loan-to-value ratios θ and cross-protocol rehypothecation (scaling with the spectral radius $\rho(A)$), which amplifies liquidations under smart-contract automation (negative and convex).

The term Q_{t+1} is highly state-dependent because peg haircuts, fee spikes, and dense collateral networks magnify the negative forces described above. In tranquil conditions, the positive price-liquidity channel may offset these effects, but outside of that Q_{t+1} is typically negative. This implies that the planner must raise the marginal cost of borrowing in states where the collateral constraint binds. This demonstrates our rationale for using state-contingent wedges. For example, capping the effective liquidity multiplier LM_t^{depeg} or imposing a surcharge that increases with $\rho(A)$ instead of relying on constant fees that can not adjust to risk conditions.

3.7.1 Comparison of Equilibria

Consider first the consumption condition. Under the planner, the shadow value of wealth is

$$\ell_t \;=\; \frac{\partial u(c_t)}{\partial c_t} \;-\; y_t \,\frac{\partial^2 u(c_t)}{\partial c_t^2}\, q_t \tag{21}$$

where y_t is the multiplier on the implementability constraint and q_t is the wrapped–collateral price. In the decentralized equilibrium (DE), the envelope condition yields $\ell_t = \partial u(c_t)/\partial c_t$. Hence the private shadow value equals the marginal utility of current consumption, whereas the social shadow value adds the term through which a marginal increase in c_t lowers marginal utility and relaxes the implementability constraint via its effect on q_t. Using the collateral–pricing link

$$y_t \;=\; \frac{\kappa_t\, m_t^*}{\dfrac{\partial u(c_t)}{\partial c_t}}, \tag{22}$$

the final term in (26) can be expressed as

$$-\,\frac{\partial^2 u(c_t)}{\partial c_t^2}\; q_t \left[\frac{\kappa_t\, m_t^*}{\dfrac{\partial u(c_t)}{\partial c_t}}\right].$$

If the collateral constraint is slack ($m_t^* = y_t = 0$), the social and private shadow values coincide. In this environment, binding occurs precisely in states with impaired pegs ($\kappa_t < 1$), elevated protocol fees (τ_t), or dense rehypothecation networks (large spectral radius $\rho(A)$), each of which compresses effective collateral values and elevates the payoff to relaxing the implementability constraint through q_t.

Turning to issuance, the planner's intertemporal condition is

$$\frac{\partial u(c_t)}{\partial c_t} = \beta R_t\, \mathbb{E}_t\!\left[\frac{\partial u(c_{t+1})}{\partial c_{t+1}} - \; y_{t+1}\,\frac{\partial^2 u(c_{t+1})}{\partial c_{t+1}^2}\, Q_{t+1} + \; y_t\, Q_{t+1}\right] + \; y_t\,\frac{\partial^2 u(c_t)}{\partial c_t^2}\, q_t + \; m_t^*, \tag{23}$$

where Q_{t+1} summarizes how today's issuance shifts tomorrow's collateral tightness and pricing. Compared with the DE Euler, $\frac{\partial u(c_t)}{\partial c_t} = \beta R_t \mathbb{E}_t[\partial u(c_{t+1})/\partial c_{t+1}]$ (when the DE constraint is slack), two forces emerge. First, a forward (ex-ante) channel operates even if $m_t = 0$: substituting (27) into (26) produces

$$-\frac{\partial^2 u(c_{t+1})}{\partial c_{t+1}^2}\, q_{t+1} \left[\frac{\kappa_{t+1}\, m_{t+1}^*}{\dfrac{\partial u(c_{t+1})}{\partial c_{t+1}}}\right],$$

which is strictly positive whenever $m_{t+1}^* > 0$. Economically, an extra unit of issuance in tranquil periods shifts probability mass toward states in which the collateral constraint binds tomorrow; the planner therefore borrows less ex ante. This effect scales with DeFi amplification: the depeg-adjusted leverage,

$$LM_t^{\text{depeg}} \;=\; \frac{1}{\kappa_t(1-\theta+\tau_t)},$$

and cross-protocol rehypothecation (large $\rho(A)$) both enlarge the set of states with $m_{t+1}^* > 0$ and increase the sensitivity of future tightness to today's issuance. Second, a contemporaneous (ex-post) channel operates only if $m_t > 0$: the term $y_t(\partial^2 u(c_t)/\partial c_t^2)q_t$ captures that a marginal rise in c_t props up q_t and loosens today's binding constraint, lowering the social return to saving in those states; the accompanying strategic term with Q_{t+1} reflects how current b_{t+1} alters the state faced by future planners and can have either sign. In wrapped-stablecoin markets, automated liquidations and oracle updates make q_t highly elastic exactly when $m_t > 0$, and peg-contingent haircuts κ_t further increase the benefit of supporting q_t in those states.

Taken together, the planner tilts away from issuance relative to the DE precisely in regions where DeFi's design primitives—large LM_t^{depeg}, high $\rho(A)$, and stressed pegs/fees (low κ_t, high τ_t)—make tomorrow's binding more likely and today's collateral price more elastic. The wedge between allocations is therefore greatest in those states, providing a direct mapping from protocol design into the quantitative size of the corrections.

3.7.2 Optimal Policy and Decentralization

We now show that the planner's allocation can be implemented in a regulated competitive equilibrium with a state–contingent *issuance wedge* applied at the protocol level. In our environment the wedge is delivered by a mint/borrow surcharge $\tau_t \geq 0$ and a state–contingent collateral schedule $\theta_t \in (0,1)$, which together alter the private Euler condition so that it coincides with the planner's Euler. As in the benchmark, lump–sum transfers rebate fee revenue so that only marginal incentives are affected, not resources. The decentralization argument is therefore identical in structure to the benchmark result with a debt tax, with the difference that the corrective instrument acts on wrapped–token issuance rather than on broad debt.

The optimal wedge inherits the two–part structure that follows from (14). The forward channel through q_{t+1} yields a prudential component that is nonnegative and is active only

when the collateral constraint may bind at $t+1$. Using the same steps that lead to the benchmark equation (17), the macroprudential term is

$$t_t^{MP} \;=\; -\,\frac{E_t\big[\,y_{t+1}\,u''(c_{t+1})\,Q_{t+1}\,\big]}{E_t\big[u'(c_{t+1})\big]} \;\geq 0, \tag{24}$$

with strict inequality whenever $P(m_{t+1}^* > 0) > 0$ since $u'(c_{t+1}) > 0$, $u''(c_{t+1}) < 0$, and $y \geq 0$. In our setting t_t^{MP} can be implemented either by choosing τ_t or by selecting an equivalent θ_t that induces the same right–hand side in the private Euler for the wrapped claim that appears in the implementability constraint. Because wrapped tokens expand transactional capacity one–for–one with the fee–adjusted, depeg–aware liquidity multiplier

$$LM_t' \;=\; \frac{1}{\kappa_t\,\big(1-\theta_t+\tau_t\big)}\,, \tag{25}$$

a simple and transparent way to deliver the prudential wedge is to cap the multiplier at a target LM^*. The protocol then selects

$$\tau_t \;=\; \max\!\left\{\,0,\;\frac{1}{\kappa_t\,LM^*}-\big(1-\theta_t\big)\,\right\}, \tag{26}$$

so that fees rise mechanically when the peg weakens (κ_t falls) or leverage is high (θ_t rises), precisely in the states in which $E_t[y_{t+1}]$ is large and Q_{t+1} is most negative. In this way (26) implements (24) with an instrument that is native to the wrapped–stablecoin architecture and preserves the nonnegativity of the prudential wedge in tranquil states.

Recursive rehypothecation and liquidity spirals enlarge the set of states with $m_{t+1}^* > 0$ by transmitting balance–sheet shocks across venues. Let $\rho(A_t)$ denote the spectral radius of the rehypothecation matrix. To reflect the increased force of the forward channel when network propagation is strong, the protocol adds a surcharge that is proportional to $[\rho(A_t)-\rho^*]_+$. When the DeFi financial accelerator elasticity ε_t exceeds unity, which marks the region where liquidations become self–reinforcing, the schedule tightens funding further by raising τ_t or by lowering θ_t:

$$\tau_t \;\leftarrow\; \tau_t \;+\; \phi\,\big[\rho(A_t)-\rho^*\big]_+ \;+\; \chi\,\mathbf{1}\{\varepsilon_t > 1\}. \tag{27}$$

Equations (26)–(27) therefore map the primitives that govern amplification in the wrapped–token system directly into the prudential component (24).

The contemporaneous channel through q_t enters the planner's Euler only when $m_t > 0$ and, as in the benchmark, has an ambiguous sign because one term lowers the social return to saving by propping up q_t when the constraint binds, while the strategic term depends on Q_{t+1}. In our implementation the prudential surcharge τ_t remains nonnegative; crisis management is delivered by a rule that is triggered only under $m_t > 0$ and that provides temporary support to q_t using a fee–funded stabilization module with rate–limited interventions or grace–period throttling of liquidations. This reproduces the sign of the crisis–time wedge without relying on negative fees in normal times and preserves the clean separation between good–times discipline and crisis–time relief that is implicit in (14).

Finally, constant wedges are dominated by the state–contingent schedule described above and can reduce welfare, since they fail to scale with binding risk and price impact. In the benchmark, fixed debt taxes have at best weak effects and can be welfare–reducing, while simple elastic rules help but remain inferior to the constrained–efficient policy. The same ranking applies here: a constant surcharge or a static LTV schedule is strictly inferior to a rule that tracks the peg factor κ_t, the collateral schedule θ_t, the forward price–impact embedded in Q_{t+1}, and the network amplification captured by $\rho(A_t)$.

4 Institutional Data

4.1 Wrapped Stablecoins

Wrapped Stablecoins arise in DeFi lending protocols when users deposit a base stablecoin and receive a transferable receipt token in return. For example, depositing USDC into Aave mints aUSDC, which represents a claim on the deposited USDC and can be pledged as collateral to borrow additional stablecoins. Because these receipt tokens are transferable and accepted as collateral, traders can execute deposit-borrow-redeposit loops in which the same underlying stablecoin supports multiple layers of borrowing.

In the empirical analysis, we define wrapped stablecoins as stablecoin receipt tokens issued by lending protocols against deposits of base stablecoins. We restrict this definition to protocol-issued receipt tokens, because they map directly into the recursive borrowing mechanism we study. We do not treat bridge-wrapped stablecoins or yield vault shares as wrapped stablecoins in the baseline, since their primary risks and value dynamics differ from receipt tokens and do not correspond to the same deposit-borrow-redeposit channel.

4.2 Data

Our empirical analysis focuses on Ethereum and major EVM chains where Aave is active: Ethereum, Arbitrum, Optimism, Base, Polygon, BNB Chain, Avalanche, and Gnosis. These networks host the deepest stablecoin lending markets in DeFi and are where Aave's deployments are most established, so stablecoin borrowing, leverage formation, and collateral reuse are most visible and consistently measured. We collect wrapped stablecoin net creation and cross-chain bridge flows from Flipside, which provides standardized labeled tables for token mints/burns and bridge activity across chains. We collect protocol borrowing and supply rates from Dune and The Graph and use these series to construct the stablecoin borrowing spread used in our recursive leverage and liquidity multiplier measures.

We focus on Aave for three reasons. First, Aave is a leading DeFi lending venue over our sample period, accounting for an average of 44.9% of total lending TVL.[2] Sec-

[2]Over January 2021–August 2025 (N = 1,674 daily observations), Aave's mean share of aggregate DeFi lending TVL is 44.9%. Using on-chain data indexed via The Graph, Aave V3 is also the primary destination for cross-protocol collateral flows, absorbing 96% of verified rehypothecation inflows ($4.7B of $4.9B), primarily liquid staking derivatives (wstETH, rETH). Appendix 6 describes variable construction and definitions.

ond, Aave's contract structure and risk parameter conventions are relatively standardized across chains, which allows comparable measurement of borrowing, utilization, and liquidation outcomes. Third, for many competing lending protocols, historical policy variables and borrowing spreads are harder to reconstruct consistently, making cross-protocol aggregation noisier. Additional evidence on cross-protocol collateral flow concentration in Aave is reported in the appendix.

Although our analysis centers on a single protocol, the mechanism we study is not Aave-specific: deposit–borrow–redeposit loops arise generically in collateralized lending with transferable receipt tokens. We view Aave as a representative and clean setting where these dynamics operate at scale, and focusing on one dominant protocol reduces measurement error that would arise from pooling across platforms with heterogeneous liquidation mechanics, oracle designs, and collateral conventions.

5 Empirical Results

We have three main sets of results. First, we show that wrapped stablecoin creation causally increases leverage in DeFi lending markets. Second, we find that higher aggregate recursive leverage amplifies cryptocurrency volatility during market stress events. Third, we document rapid cross-chain contagion: a stablecoin shock on one chain propagates to nearly half of all connected chains in under an hour, with the fastest transmission occurring in just 8 minutes.

Some clarifications are in order when taking the model predictions to the data. In particular, the recursive leverage channel we identify operates through the mechanical relationship between wrapped stablecoin deposits and collateralized borrowing. While our instrumental variable approach isolates exogenous variation in wrapped stablecoin supply, we acknowledge that leverage dynamics may differ across lending protocols depending on collateral parameters, liquidation thresholds, and utilization rates. At the same time, the aggregate effect we estimate is represented as a weighted average across the DeFi ecosystem.

For our volatility amplification results, we study twelve stress events spanning October 2021 to August 2024 where the recursive leverage channel is the most salient. These episodes—including protocol exploits (Cream Finance, Wormhole, Euler Finance, Curve Vyper), market crashes (December 2021, January 2022, August 2024 carry unwind), stablecoin stress events (Terra/UST, USDC depeg, USDT Curve 3Pool), and CeFi contagion (Celsius/3AC, FTX collapse)—represent periods when deleveraging pressures were acute and the feedback loop between collateral values and forced liquidations was most likely to bind.

Finally, our cross-chain contagion analysis isolates the speed of shock transmission through bridge infrastructure. We define contagion as the first minute where net bridge outflows exceed two standard deviations below the pre-event mean for each destination chain. The striking speed of propagation 43% of chains affected within 60-minutes underscores how interconnected the multi-chain stablecoin ecosystem has become, with implications for systemic risk that extend beyond any single blockchain.

5.1 Wrapped Stablecoins and Leverage

We verify the equilibrium liquidity multiplier proposition by showing that wrapped stablecoin creation causally increases leverage in DeFi lending markets—the recursive leverage channel described in Section 2. The key empirical challenge is that leverage demand and stablecoin creation likely respond to common shocks. To address this endogeneity, we instrument for wrapped stablecoin creation using regulatory announcements that affect stablecoin issuance but do not directly impact lending market leverage (we describe the selection process in Appendix C3). Throughout the empirics, we define the liquidity multiplier as

$$LM_{c,t} \;=\; \frac{1}{1-\theta_c+\tau_{c,t}},$$

where θ_c is the value-weighted average loan-to-value ratio across stablecoin markets on chain c, and $\tau_{c,t} = r^b_{c,t} - r^s_{c,t}$ is the borrow–supply rate spread (the effective wedge). This corresponds to the baseline (peg-on) liquidity multiplier from Equation (19), setting $\kappa_t = 1$.

We estimate the following specification:

First stage:

$$\text{WrappedCreation}_{c,t} = \pi_0 + \pi_1 Z_{\text{regulatory},t} + \gamma X_{c,t} + \alpha_c + \delta_t + u_{c,t}.$$

Second stage:

$$\text{LM}_{c,t} = \beta_0 + \beta_1 \widehat{\text{WrappedCreation}}_{c,t} + \gamma X_{c,t} + \alpha_c + \delta_t + \varepsilon_{c,t}.$$

where $\text{LM}_{c,t}$ is the liquidity multiplier for chain c on day t, and $\widehat{\text{WrappedCreation}}_{c,t}$ is instrumented wrapped stablecoin net creation (millions USD). $X_{c,t}$ includes controls for major token returns, log TVL, DEX volume, gas fees, and supply growth. α_c are chain fixed effects, and δ_t are month fixed effects. We cluster standard errors at the chain–month level (257 clusters).

Table 1 presents our two-stage least squares estimates. The first-stage partial F-statistic is 19.97, well above the conventional threshold of 10, indicating that regulatory announcements are a strong predictor of wrapped stablecoin creation. In the second stage, we estimate that a one-million-dollar increase in wrapped stablecoin creation raises the liquidity multiplier by 0.00073 units ($t = 5.12$, $p < 0.001$). A one standard deviation increase in wrapped stablecoin creation—approximately \$111 million—raises the liquidity multiplier by 7% relative to its mean.

One concern is that our sample includes only eight chains, raising questions about the reliability of clustered standard errors. We address this in several ways. First, we cluster at the chain–month level (257 clusters) rather than the chain level, which provides sufficient clusters for asymptotic inference under the assumption that within-chain shocks dissipate within a month. Second, we implement two-way clustering by chain and month following Cameron, Gelbach, and Miller (2011), which yields similar standard errors. Third, we employ a pair cluster bootstrap, which resamples entire chains with replacement; the 95%

Table 1: Wrapped Stablecoin Creation and DeFi Leverage: 2SLS Estimates

	(1) OLS	(2) 2SLS
Panel A: Second Stage — Dependent Variable: Liquidity Multiplier		
Wrapped Creation ($M)	0.00019***	0.00073***
	(0.00003)	(0.00014)
Major Token Return	−0.251***	−0.250***
	(0.076)	(0.076)
Log TVL	0.018***	0.018***
	(0.007)	(0.007)
DEX Volume	−0.000	−0.000
	(0.000)	(0.000)
Gas Fee	0.000	0.000
	(0.000)	(0.000)
Supply Growth	0.092	0.092
	(0.081)	(0.081)
Panel B: First Stage — Dependent Variable: Wrapped Creation ($M)		
$Z_{regulatory}$		32.70***
		(3.20)
Partial F-statistic		19.97
Chain FE	Yes	Yes
Month FE	Yes	Yes
Observations	7,578	7,578
Chains	8	8
Clusters (chain-month)	257	257

Notes: Dependent variable is the liquidity multiplier. Wrapped Creation is net minting minus redemptions ($M). Instrument $Z_{regulatory}$ indicates regulatory announcement days. Standard errors clustered at chain-month level in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

confidence interval excludes zero. We acknowledge that with few chains, inference should be interpreted cautiously, though the coefficient is consistently positive and economically meaningful across all specifications.

Another concern is that regulatory announcements may correlate with broader market conditions that independently affect leverage. In Table 6, Panel B, we control for major token returns, total value locked, DEX volume, gas fees, and supply growth—variables that capture market-wide conditions. The coefficient on wrapped stablecoin creation remains stable after including these controls, suggesting that our instrument operates through the wrapped stablecoin channel rather than correlated market factors.

Finally, we compare our 2SLS estimates to naive OLS. The OLS coefficient is 0.00019, nearly four times smaller than the 2SLS estimate. This pattern is consistent with omitted variables that simultaneously reduce both stablecoin creation and leverage demand—for instance, periods of market stress when issuers restrict minting and traders deleverage. The 2SLS approach recovers the structural effect of wrapped stablecoin supply on leverage, which OLS attenuates.

We conduct several additional tests to validate our identification strategy (Appendix D.1). Table 4 confirms that the instrument operates specifically through the regulatory channel: replacing it with bridge exploits or operational events yields weak first stages, while lagged and cumulative transformations of the regulatory shock produce consistent results. A placebo test shifting announcement dates by 30 days eliminates significance (Table 6), and the exclusion restriction holds for non-stablecoin assets (Table 3). Subsample stability tests confirm positive, significant effects wherever the instrument is strong (Table 5).

5.2 Recursive Leverage and Crypto Volatility

We test whether recursive leverage in DeFi lending protocols amplifies cryptocurrency volatility during stress events. Our identification strategy exploits within-event variation in time-varying recursive leverage across twelve major crypto stress episodes spanning October 2021 to August 2024, observed at the protocol–chain level across seven Aave deployments (V1 Ethereum, V2 Ethereum, V2 Avalanche, V3 Arbitrum, V3 Avalanche, V3 Ethereum, V3 Polygon). These episodes represent plausibly exogenous shocks to the crypto ecosystem, allowing us to study how leverage conditions affect volatility transmission.

For each event, we construct a symmetric ±48-hour window centered on the event timestamp, yielding up to 1,152 five-minute intervals per event per protocol–chain pair. We measure recursive leverage $\mathrm{RL}_{i,t}$ at the protocol–chain level at each five-minute interval, preserving both cross-sectional and time-series variation.

We estimate:

$$\mathrm{Vol}_{e,t} = \alpha + \beta_1 \mathrm{Post}_{e,t} + \beta_2 \mathrm{RL}_{i,t} + \beta_3 \left(\mathrm{Post}_{e,t} \times \mathrm{RL}_{i,t}\right) + \varepsilon_{e,i,t}, \tag{28}$$

where $\mathrm{Vol}_{e,t}$ is the absolute five-minute BTC return for event e at time t, $\mathrm{Post}_{e,t}$ indicates observations after the event, and $\mathrm{RL}_{i,t}$ is protocol–chain-level recursive leverage at time t. Standard errors are clustered at the event × protocol–chain level. The coefficient

Table 2: Recursive Leverage and Volatility Amplification During Stress Events

	(1) RL Interaction	(2) With Controls	(3) Severity Interaction
Post	−0.001577***	−0.001577***	
	(0.000506)	(0.000506)	
RL_t	−0.000037	−0.000037	
	(0.000097)	(0.000097)	
Post × RL_t	0.000690***	0.000690**	
	(0.000166)		
		[0.024]	
Post × RL^{pre}			−0.000118
			(0.000280)
Post × RL^{pre} × Severity			0.000223***
			(0.000049)
Major token return		Yes	
Observations	50,622	50,622	50,622
Events	12	12	12
Protocol–chain clusters	46	46	46
R^2	0.058	0.058	0.100
Wild cluster bootstrap p		0.024	
Bootstrap 95% CI		[0.000091, 0.001279]	

Notes: Dependent variable is absolute 5-minute BTC return (volatility proxy). Sample includes symmetric ±48-hour windows around twelve stress events spanning October 2021 to August 2024: Cream Finance exploit, December 2021 flash crash, January 2022 correction, Wormhole bridge exploit, Terra/UST collapse, Celsius/3AC contagion, FTX collapse, Euler Finance exploit, USDC depeg, USDT Curve stress, Curve Vyper exploit, and August 2024 carry trade unwind. RL_t is protocol–chain-level recursive leverage measured at each 5-minute interval. Panel is at the event × protocol × chain × timestamp level across seven Aave deployments (V1 Ethereum, V2 Ethereum, V2 Avalanche, V3 Arbitrum, V3 Avalanche, V3 Ethereum, V3 Polygon). Column (1) reports clustered standard errors (event × protocol–chain) in parentheses. Column (2) adds major token return as a control and reports the wild cluster bootstrap p-value in brackets (9,999 Rademacher replications, clustered at the event level with 12 clusters). Column (3) interacts post × pre-event RL with standardized event severity (BTC price drop). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$ (based on wild cluster bootstrap in Column 2; clustered SEs elsewhere).

of interest is β_3, which captures whether higher recursive leverage amplifies post-event volatility.

Table 2 reports the results. Column (1) presents our main specification. The interaction term $\text{Post} \times \text{RL}_t$ is positive and highly significant ($\hat{\beta}_3 = 0.000690$, $p < 0.001$), indicating that protocol–chain pairs with higher recursive leverage experience amplified volatility in the post-event window. The Post indicator is negative and significant (-0.00158, $p < 0.01$), reflecting that absent elevated leverage, five-minute absolute BTC returns tend to decline after the initial shock as markets stabilize.

Column (2) adds the major token return as a control and reports wild cluster bootstrap inference to address the concern that standard clustered standard errors may be unreliable with a moderate number of clusters. The point estimate is unchanged at 0.000690, and the bootstrap p-value is 0.024 (9,999 Rademacher replications clustered at the event level), confirming statistical significance at the 5% level. The bootstrap 95% confidence interval is $[0.000091, 0.001279]$, excluding zero.

Column (3) explores whether the amplification effect varies with event severity by interacting $\text{Post} \times \text{RL}^{\text{pre}}$ with standardized BTC price drop. The triple interaction is positive and highly significant (0.000223, $p < 0.001$), while the uninteracted $\text{Post} \times \text{RL}^{\text{pre}}$ term is near zero and insignificant (-0.000118, $p = 0.67$). This indicates that the leverage–volatility amplification channel activates primarily during severe shocks; in mild events, pre-existing leverage has little effect on post-event volatility.

Several caveats apply. First, our volatility measure captures market-wide BTC volatility rather than protocol-specific outcomes, and therefore does not identify the precise transmission mechanism (e.g., forced liquidations, margin calls, or sentiment spillovers). As a channel validation, Appendix Table 8 replaces BTC with ETH as the outcome variable: the $\text{Post} \times \text{RL}^{\text{pre}}$ coefficient is approximately 1.8 times larger for ETH than for BTC and marginally significant ($p = 0.091$), consistent with the DeFi channel operating more directly through ETH-denominated collateral markets. Second, although we mitigate reverse causality by using the pre-determined component of leverage, we cannot fully rule out anticipatory leverage adjustments in expectation of increased volatility.

Despite these limitations, the findings support the hypothesis that recursive leverage in DeFi lending protocols amplifies cryptocurrency volatility during stress events. The mechanism is intuitive: higher leverage implies that a given price decline triggers larger forced selling, which further depresses prices and induces additional liquidations. The severity interaction in Column (3) further suggests that this amplification is nonlinear, binding primarily when shocks are large enough to trigger cascading liquidations.

5.3 Cross Chain Contagion and Bridge Connectivity

We analyze the propagation of the March 2023 USDC depeg shock across 21 blockchain networks to identify determinants of cross-chain contagion speed. The shock originated on March 10, 2023 at 22:00 UTC, when USDC first traded below \$0.995 following news of Circle's \$3.3 billion exposure to Silicon Valley Bank.

Table 9 reports the time elapsed between the initial shock and the first detection of stress on each destination chain. We define stress as the first minute where net bridge

outflows exceed two standard deviations below the pre-event (7-day) mean for that chain.

Contagion was rapid but heterogeneous. Arbitrum experienced stress within 8 minutes of the initial shock, followed by Fantom (17 minutes), Ethereum (20 minutes), and BSC (26 minutes). Nine of 21 chains (43%) experienced stress within one hour. However, several chains exhibited substantially delayed contagion: Solana (396 minutes), zkSync Lite (1,526 minutes), and Injective (1,640 minutes). This variation—ranging from 8 minutes to over 27 hours—motivates our investigation into the determinants of contagion speed.

We estimate the relationship between chain characteristics and contagion speed using OLS regression with robust standard errors:

$$\log(\text{Delay}_i) = \alpha + \beta_1 \log(\text{TVL}_i) + \beta_2 \log(\text{Bridges}_i) + \varepsilon_i \tag{29}$$

where Delay_i is minutes to first stress, TVL_i is total value locked in billions USD, and Bridges_i is the number of active bridge protocols connecting chain i to other networks.

Table 10 presents the results. In univariate specifications, both TVL (Column 1) and bridge count (Column 2) significantly predict contagion speed. However, in the combined specification (Column 3), bridge connectivity remains highly significant ($\beta = -1.578$, $p < 0.01$) while TVL becomes statistically indistinguishable from zero ($\beta = 0.065$, $p = 0.75$). The coefficient on log(Bridges) implies that a doubling of bridge connections is associated with a $(1 - 2^{-1.578}) \times 100 = 67\%$ reduction in time to contagion.

The economic magnitude is substantial. Chains with more than three bridge connections experienced stress in 45 minutes on average, compared to 550 minutes for chains with three or fewer bridges—a difference of over 12-fold (Mann-Whitney $p < 0.001$).

We conduct two sets of robustness tests to assess the sensitivity of our findings.

Our baseline defines stress arrival as net outflows exceeding two standard deviations below the pre-event mean. Panel D of Table 16 reports results using alternative thresholds. Varying the relative threshold from 1.5σ to 2.5σ yields coefficients ranging from -1.460 to -1.498, all significant at $p < 0.001$. Using absolute dollar thresholds ($10k to $100k net outflows) produces coefficients between -1.263 and -1.867, with all seven specifications significant at the 5% level (six at 1%). The bridge connectivity result is not sensitive to how we operationalize stress arrival.

Panel E of Table 16 replicates the analysis across six additional major stress events: the Terra/Luna collapse (May 2022), FTX collapse (November 2022), Celsius/3AC crisis (June 2022), January 2022 market correction, USDT Curve pool stress (June 2023), and the Curve Vyper exploit (July 2023). All seven events—including the main USDC depeg—exhibit negative coefficients on bridge connectivity, demonstrating 100% sign consistency. Four of seven events reach statistical significance at the 10% level, with the Terra/Luna collapse ($\beta = -1.413$, $p = 0.013$) and FTX collapse ($\beta = -1.255$, $p = 0.055$) showing effects comparable in magnitude to the main specification. The average coefficient across all events is -0.947. The finding that bridge connectivity predicts faster contagion is not specific to the USDC depeg but generalizes across diverse stress events.

Our findings suggest that bridge infrastructure, rather than chain size per se, determines how quickly cross-chain contagion propagates. This is consistent with bridges serving as transmission channels for capital flight during stress events: more bridges provide more pathways for rapid outflows.

However, several caveats apply. First, TVL and bridge count are highly correlated ($r = 0.86$), reflecting the fact that larger chains tend to attract more bridge deployments. While bridge connectivity dominates in our horse-race specification and the coefficient remains negative in 99% of bootstrap resamples, we cannot definitively rule out that TVL contributes to contagion speed through channels correlated with bridge count. The variance inflation factor of 3.8 is within acceptable bounds (below 5) but indicates elevated multicollinearity.

Second, while our out-of-sample tests demonstrate that the bridge connectivity effect replicates across multiple stress events, weaker effects in smaller events (e.g., the Celsius/3AC crisis and Curve exploit) suggest that statistical power may be limited when contagion is less pronounced.

Third, our bridge count estimates are based on ecosystem knowledge rather than precise on-chain measurement. Future work could leverage bridge transaction data to construct more granular measures of cross-chain connectivity.

Despite these limitations, our results have implications for systemic risk monitoring in decentralized finance. If bridge connectivity determines contagion speed, then network topology—not just chain-level metrics like TVL—should inform risk assessment. Highly connected chains may benefit from deeper liquidity and arbitrage efficiency in normal times, but face faster capital flight during stress events. Regulators and protocol designers seeking to contain cross-chain contagion should consider the trade-offs inherent in bridge connectivity: isolation slows contagion but may also reduce market efficiency and increase fragmentation.

6 Conclusion

This paper documents how wrapped stablecoins implement a form of shadow banking on-chain, creating recursive collateral chains that amplify both liquidity and fragility in decentralized finance. We develop a tractable model in which the reuse of stablecoin collateral through lending protocols generates an endogenous liquidity multiplier, and we show that the same mechanism that expands funding capacity in normal times produces abrupt liquidation cascades when collateral constraints tighten.

Our empirical analysis, based on novel on-chain data from Aave deployments across eight blockchain networks, provides three main findings. First, wrapped stablecoin creation causally increases recursive leverage: using regulatory announcements as instruments, we estimate that a one-standard-deviation increase in wrapped creation raises the liquidity multiplier by 7% relative to its mean, with a first-stage F-statistic of 19.97. Second, higher pre-event recursive leverage amplifies cryptocurrency volatility during stress episodes, with the amplification effect concentrated in equally-weighted rather than value-weighted measures, consistent with smaller positions being more vulnerable to mechanical liquidation. Third, cross-chain bridge connectivity, not chain size, determines the speed at which stress propagates across networks: doubling bridge connections reduces time to stress arrival by 67% during the March 2023 USDC depeg. Wallet-level evidence confirms that agents execute the deposit-borrow-redeposit loops predicted by the model, achieving effective leverage ratios between 1.6× and 2.5× through recursive collateral reuse.

Our results carry implications for the design of DeFi protocols and for macroprudential policy in digital asset markets. The fundamental tradeoff we identify—that deeper collateral reuse improves funding efficiency but increases exposure to systemic liquidation cascades—suggests that policies aimed at improving day-to-day funding conditions and those aimed at reducing crisis severity are fundamentally different goals, and that achieving one may come at the expense of the other. Our welfare analysis shows that state-contingent collateral schedules, in which loan-to-value ratios and protocol fees adjust with peg conditions and network stress, can reduce the wedge between the decentralized equilibrium and the constrained-efficient allocation. In contrast, static collateral requirements or constant fee structures are strictly dominated.

Several limitations of our analysis point to directions for future work. Our event study relies on a relatively small number of major stress episodes, and while our wild cluster bootstrap and out-of-sample replication strengthen the inference, a longer sample with more events would allow for finer decomposition of the amplification channel. Our cross-chain contagion analysis uses bridge counts rather than transaction-level bridge flow data, which future work could leverage as bridge analytics improve. The model abstracts from governance dynamics—in practice, protocol parameters are adjusted through governance proposals with meaningful delays, which may affect the feasibility of implementing the state-contingent policies our welfare analysis recommends. Finally, as DeFi lending markets evolve to incorporate new collateral types such as liquid staking derivatives and real-world assets, the recursive leverage mechanisms we document may manifest in new and potentially more complex forms, warranting continued attention from both researchers and market participants.

Appendix A: Theoretical Analysis and Proofs

This section presents the theoretical foundations and proofs for the analysis of our paper. Throughout the appendix, z_t denotes the aggregate productivity factor as in Bianchi and Mendoza (2011); the stress regime is denoted s_t (see Section 3.2.1).

A.1: Recursive Constrained-Efficient Planner Formulation

We lay out the analytical framework for the Markov-perfect equilibrium and show that the recursive planner's problem is given by (21).

Given prices $\{q_t^W, p_v, R_t\}$ and transfers consistent with the resource constraint, the representative private agent chooses $\{c_t, h_t, v_t, w_{t+1}\}_{t\geq 0}$ to maximize

$$\mathbb{E}_0 \sum_{t\geq 0} \beta^t \, u(c_t - G(h_t)),$$

subject to the period-t budget constraint and the next-period collateral constraint

$$\frac{B_{t+1}}{R_t} - p_v v_t \;\geq\; -\theta_t q_t^W.$$

First-order conditions imply:

$$z_t F_h(1, h_t, v_t) = G'(h_t), \tag{A.1}$$

$$z_t F_v(1, h_t, v_t) = p_v(1 + \theta_t \, \mu_t / u'(c_t - G(h_t))), \tag{A.2}$$

$$q_t^W \, u'(c_t - G(h_t)) = \beta \, \mathbb{E}_t\Big[u'(c_{t+1} - G(h_{t+1}))\big(q_{t+1}^W + r_{t+1}^W q_{t+1}^W + \theta_{t+1} m_{t+1} q_{t+1}^W\big)\Big], \tag{A.3}$$

with $\mu_t \geq 0$ and $\mu_t\Big(\frac{B_{t+1}}{R_t} - p_v v_t + \theta_t q_t^W\Big) = 0$. Equations (A.1)–(A.3) and the Kuhn–Tucker conditions are the implementability constraints used below in the planner's problem.

Problem 1. Formulation of the planner's problem We evaluate a time-consistent (Markov-perfect) planner who chooses policies and takes the policy rules of future planners as given. Let $B(B_t, s_t)$ denote the conjectured policy for next-period issuance and let $\{C(B,s), H(B,s), v(B,s), m(B,s), Q^W(B,s)\}$ be the associated recursive allocations and wrapped-collateral pricing function. The decentralized private problem implies a wrapped-asset Euler equation. Imposing it yields the SP's asset-pricing implementability constraint.

In recursive form, the planner's problem is stated with all constraints as follows:

$$c_t + \tfrac{B_{t+1}}{R_t} = B_t + F(1, h_t, v_t) - p_v v_t \tag{SP1}$$

$$\tfrac{B_{t+1}}{R_t} - p_v v_t \geq -\theta_t q_t^W \tag{SP2}$$

$$q_t^W u'(c_t - G(h_t)) = \beta \mathbb{E}_t \Big[\cdots \Big] \tag{SP3}$$

$$z_t F_h(1, h_t, v_t) = G'(h_t) \tag{SP4}$$

$$z_t F_v(1, h_t, v_t) = p_v \Big(1 + \theta_t \mu_t / u'(c_t - G(h_t)) \Big) \tag{SP5}$$

$$\mu_t \geq 0, \quad \mu_t \Big(\tfrac{B_{t+1}}{R_t} - p_v v_t + \theta_t q_t^W \Big) = 0 \tag{SP6–7}$$

Proposition: Relaxed planner problem In the time-consistent planner's problem A.1, we will show that constraints (SP4)-(SP7) do not bind. Thus, the planner can equivalently solve the relaxed recursive problem consisting of the objective and constraints (SP1)-(SP3) only.

Consider the relaxed problem that includes (SP1)–(SP3) only and denote the multipliers on (SP1), (SP2) and (SP3) by $\mu_t \geq 0$, $m_t^{sp} \geq 0$, and $\gamma_t \geq 0$, respectively. The Lagrangian is

$$\begin{aligned}\mathcal{L} &= u(c_t - G(h_t)) + \beta\, \mathbb{E}_t V(B_{t+1}, s_{t+1}) \\ &\quad + \mu_t \Big[B_t + F(1, h_t, v_t) - p_v v_t - \tfrac{B_{t+1}}{R_t} - c_t \Big] + m_t^{sp} \Big[\tfrac{B_{t+1}}{R_t} - p_v v_t + \theta_t q_t^W \Big] + \gamma_t \left[\text{(SP3) residual} \right],\end{aligned}$$

First-order conditions with respect to c_t, h_t and v_t yield, respectively,

$$u'(c_t - G(h_t)) = \mu_t, \qquad \mu_t F_h(1, h_t, v_t) = \mu_t G'(h_t), \qquad \mu_t F_v(1, h_t, v_t) = \mu_t p_v + \big(p_v\, \theta_t\big) m_t^{sp} / u'(c_t - G(h_t)).$$

The FOCs of the relaxed problem exactly reproduce the labor and input optimality conditions, i.e. (SP4)–(SP5), once we set $\mu_t = u'(c_t - G(h_t))$. Since (SP2) is an inequality, the associated Kuhn-Tucker conditions imply $m_t^{sp} \geq 0$ and $m_t^{sp}(\tfrac{B_{t+1}}{R_t} - p_v v_t + \theta_t q_t^W) = 0$, which are precisely (SP6)–(SP7). Hence, the solution to the relaxed problem already satisfies (SP4)–(SP7). Conversely, any solution to the full problem satisfies the relaxed one. This completes the proof that the two problems are equivalent, and it is sufficient to study the relaxed version with (SP1)–(SP3).

A.2 Proof of Proposition 1

The constrained–efficient (time–consistent) planner's allocation can be implemented as a regulated decentralized equilibrium by a state–contingent pair of protocol instruments (τ_t, θ_t).

The constrained–efficient equilibrium is defined by sequences $\{c_t, h_t, v_t, B_{t+1}, q_t^W, m_t\}_{t \geq 0}$ that satisfy the resource constraint, the collateral constraint, and the asset–pricing implementability condition implied by the private Euler for the wrapped claim together with the recursive implementability condition and complementary slackness.

In the regulated decentralized setting, let the protocol post an issuance fee $\tau_t \geq 0$ and an LTV schedule $\theta_t \in (0,1)$. Private agents take prices as exogenous and obey their Euler condition for the wrapped asset. We express this post-wedge Euler condition by multiplying the laissez-faire right-hand side by a policy wedge represented by $\tau_t(\tau_t, \theta_t)$. We define the issuance fee τ_t such that the decentralized Euler condition matches the planner's implementability condition:

$$1+\tau_t \;\equiv\; \frac{E_t\big[u'(C_{t+1}-G(H_{t+1}))\left\{Q^W_{t+1}+z_{t+1}F_k(1,H_{t+1},v_{t+1})+\theta_{t+1}m_{t+1}Q^W_{t+1}\right\}\big]}{E_t\big[u'(c_{t+1}-G(h_{t+1}))\left\{Q^W_{t+1}+r^W_{t+1}q^W_{t+1}\right\}\big]}. \quad \text{(A.4)}$$

With $1+\tau_t$ so defined, substituting (A.4) into the decentralized Euler delivers exactly (SP3)/(22); together with (SP1)–(SP2) and complementary slackness, this yields the same system that characterizes the planner's allocation, hence implementation.

When the collateral constraint is not binding at t, the wedge reduces to the macroprudential component:

$$t^{MP}_t \;=\; -\frac{E_t[y_{t+1}\,u''(c_{t+1})\,Q_{t+1}]}{E_t[u'(c_{t+1})]} \;\geq 0, \quad \text{(A.5)}$$

which can be implemented either via τ_t or via an equivalent θ_t that induces the same right–hand side in the private Euler.

Because wrapped issuance expands transactional capacity one–for–one with the fee– and peg–adjusted liquidity multiplier

$$LM^{\text{depeg}}_t \;=\; \frac{1}{\kappa_t\left(1-\theta_t+\tau_t\right)}, \quad \text{(A.6)}$$

a constructive decentralization rule is to cap LM^{depeg}_t at a target LM^* and set

$$\tau_t \;=\; \max\left\{0,\; \frac{1}{\kappa_t\,LM^*}-\left(1-\theta_t\right)\right\}, \quad \text{(A.7)}$$

so fees rise mechanically when the peg weakens or leverage is high, precisely in the states where the forward externality is strongest. Fee revenues are rebated lump–sum so resources are unaffected and only marginal incentives change.

A.3 Optimal Fee/LTV Problem: Equivalence Result

Consider the planner's problem to determine an optimal issuance fee on wrapped–token borrowing, given that future fees $T(B,s)$ are chosen by future planners, which are associated with policies $\{B(b,s), C(b,s), H(b,s), v(b,s), m(b,s), Q^W(b,s)\}$. The planner chooses the fee to maximize utility subject to all competitive–equilibrium conditions including the household Euler with fee:

$$u'(c-G(h)) \;=\; \beta R(1+\tau)\,E[u'(C(b',s')-G(H(b',s')))]+m. \quad \text{(A.8)}$$

together with the resource constraint, the collateral constraint, and the asset–pricing implementability conditions implied by the wrapped–asset Euler (cf. (21)–(22)).

Formally, the problem is

$$V(b,s) = \max_{c,b',q^W,h,v,m,\tau} \; u(c - G(h)) + \beta\, E_{s'|s}\, V(b',s') \tag{A.9}$$

$$\text{s.t.} \quad c + \frac{b'}{R} = b + F(1,h,v) - p^v v,$$

$$\frac{b'}{R} - \theta\, p^v v \;\geq\; -\kappa_t\, q_t^W,$$

$$q_t^W\, u'(c - G(h)) = \beta E_{s'|s}\Big[u'(C(b',s') - G(H(b',s')))\big(Q^W(b',s') + z' F_k(1, H(b',s'), v(b',s')) + \theta' m(b',s'$$

$$u'(c - G(h)) = \beta R(1+\tau)\, E_{s'|s}\big[u'(C(b',s') - G(H(b',s')))\big] + m,$$

with the complementary–slackness conditions associated with the collateral constraint.

Relative to problem (SP), problem (A.9) contains one additional restriction (A.8) and one additional policy instrument τ. Let τ^* be the optimal fee that solves (A.9); the Markov–perfect equilibrium (MPE) condition implies the fee rule satisfies $\tau(B,s) = T(B,s)$.

Proposition A.3. *A sequence of allocations and prices constitutes a constrained–efficient equilibrium if and only if it is the outcome of a Markov–perfect equilibrium in which the protocol chooses sequentially a fee on issuance.*

Consider solving the relaxed problem obtained from (A.9) by dropping the additional restriction (A.8). The resulting problem is the same as the primal planner's problem (cf. (12) together with Appendix A.1), so allocations and prices satisfy (12) if and only if they satisfy the optimal–fee problem. In addition, notice that τ only appears in (A.8). Hence, setting

$$1 + \tau \;=\; \frac{u'(c - G(h))}{\beta R\, E[u'(C(b',s') - G(H(b',s')))]} \;-\; \frac{m}{\beta R\, E[u'(C(b',s') - G(H(b',s')))]} \tag{A.10}$$

implies that condition (A.8) is satisfied.

A.4 Non–Negative Fee on Issuance

We solve for the optimal fee on wrapped–token issuance, as referenced in problem (A.9), but consider fees that cannot be negative. This restriction is natural because a negative fee would be equivalent to a subsidy on issuance financed by lump–sum taxation. The unconstrained optimal fee is non–negative whenever the collateral constraint is slack, but it may turn negative when the constraint binds. Consequently, the non–negativity constraint on τ can become binding.

The planner's problem is the same as in (A.9), with the additional implementability restriction:

$$1 + \tau_t \;=\; \frac{u'(c_t - G(h_t)) - m_t}{\beta R_t\, E_t[u'(c_{t+1} - G(h_{t+1}))]} \;\geq\; 1, \tag{A.11}$$

which enforces $\tau_t \geq 0$. Following the same steps as in Appendix A.1 and the imple-

mentability system (21)–(22), the regulator's problem can be written as

$$V(b,s) = \max_{c,b',q^W,h,v,m} \; u(c-G(h)) + \beta E_{s'|s} V(b',s') \tag{A.12}$$

$$\text{s.t.} \quad c + \frac{b'}{R_t} = b + F(1,h,v) - p^v v,$$

$$\frac{b'}{R_t} - \theta\, p^v v \;\geq\; -\kappa_t\, q_t^W,$$

$$q_t^W\, u'(c-G(h)) = \beta E_{s'|s}\Big[u'(C(b',s') - G(H(b',s')))\big(Q^W(b',s') + z' F_k(1,H(b',s'),v(b',s')) + \theta' m(b',s'$$

$$\frac{u'(c_t - G(h_t)) - m_t}{\beta R_t\, E_t[u'(c_{t+1} - G(h_{t+1}))]} \;\geq\; 1,$$

together with the complementary–slackness conditions for the collateral constraint and the fee floor.

Let $\gamma_t^p \geq 0$ denote the Lagrange multiplier on the non–negativity constraint (A.11). When the collateral constraint is not binding at t and the fee floor does not bind at t, the optimal fee simplifies. Since either constraint may bind at $t+1$ with positive probability, the optimal fee reduces to the macroprudential component plus a forward–looking correction that accounts for the potential binding of the fee floor in the next period:

$$1+\tau_t^{\text{fee}\geq 0} \;=\; \frac{1}{\beta R_t\, E_t[u'(c_{t+1})]} \left\{ \beta R_t\, E_t[u'(c_{t+1}) - y_{t+1}\, u''(c_{t+1})\, Q_{t+1}] \;+\; \gamma_{t+1}^p\, \frac{u''(c_{t+1})}{\beta R_t\, E_t[u'(c_{t+1})]} \right\}. \tag{A.13}$$

This means in tranquil states the fee equals the usual prudential surcharge. If there is a positive probability that the fee floor binds tomorrow, then the additional forward term proportional to γ_{t+1}^p raises today's fee to compensate for the future loss of policy flexibility.

Finally, when the unconstrained fee implied by (A.11) would be negative, the non–negativity restriction binds, $\gamma_t^p > 0$, and the optimal fee is $\tau_t = 0$. The prudential logic in tranquil states is unaffected, and crisis–time support can instead be provided by temporary measures affecting q_t^W without violating $\tau_t \geq 0$.

A.5 Derivation of the Collateral Constraint

We derive the collateral constraint as an incentive–compatibility condition under limited enforcement. Lending occurs in decentralized, non–exclusive protocols, where borrowers may pledge assets and issue wrapped tokens. At the beginning of period t, the household can borrow and then divert resources within the period, avoiding any repayment costs next period when debt is due. However, protocol rules guarantee that undercollateralization is immediately detected: liquidators seize a fraction κ_t of posted collateral at the oracle price q_t^W. After liquidation, the wallet regains market access and can repurchase assets at q_t^W.

Let $\tilde{d}_{t+1}$ represent funds raised at t that are diverted and not repaid at $t+1$. The value of default is

$$V^d(\tilde{d}_{t+1}, b_t, k_t, X_t) \;=\; \max_{b_{t+1},\, k_{t+1},\, c_t,\, h_t,\, v_t} \; u(c_t - G(h_t)) + \beta E_t\, V(b_{t+1}, k_{t+1}, X_{t+1}) \tag{A.14}$$

subject to

$$q_t^W k_{t+1} + c_t + \frac{b_{t+1}}{R_t} = \underbrace{\tilde{d}_{t+1}}_{\text{diverted funds}} + \underbrace{q_t^W(1-\kappa_t)k_t}_{\text{post–seizure collateral}} + \underbrace{b_t + z_t F(1, h_t, v_t) - p_t^v v_t}_{\text{other resources}},$$

$$\frac{b_{t+1}}{R_t} - \theta_t p_t^v v_t \geq -\kappa_t\, q_t^W \qquad \text{(new credit limit under reentry)}.$$

If the household does not default, it gets the current utility plus the continuation value of next period's state with debt b_{t+1}:

$$V^r(b_t, k_t, X_t) \;=\; \max_{b_{t+1},\, k_{t+1},\, c_t,\, h_t,\, v_t} u(c_t - G(h_t)) + \beta E_t\, V(b_{t+1}, k_{t+1}, X_{t+1}) \tag{A.15}$$

subject to

$$q_t^W k_{t+1} + c_t + \frac{b_{t+1}}{R_t} = q_t^W k_t + b_t + z_t F(1, h_t, v_t) - p_t^v v_t,$$

$$\frac{b_{t+1}}{R_t} - \theta_t p_t^v v_t \geq -\kappa_t\, q_t^W.$$

Inspecting the two budget constraints shows that a deviating borrower who diverts $\tilde{d}_{t+1}$ gains $\tilde{d}_{t+1}$ today but loses $\kappa_t q_t^W k_t$ through collateral seizure. Hence, repayment is optimal if and only if

$$\tilde{d}_{t+1} \;\leq\; \kappa_t\, q_t^W\, k_t. \tag{A.16}$$

Therefore, condition (A.16) is equivalent to the borrowing limit

$$\frac{b_{t+1}}{R_t} - \theta_t p_t^v v_t \;\geq\; -\kappa_t\, q_t^W, \tag{A.17}$$

which is the collateral constraint used in the decentralized equilibrium.

The same collateral constraint applies under both decentralized allocations and the planner's allocation because enforcement is embedded in protocol rules. Smart contracts allow liquidators to seize only a fraction κ_t of collateral at the oracle price q_t^W, after which wallets can immediately repurchase assets in open markets. Thus, the feasible credit set is identical across the two environments, and (A.17) constrains the planner's problem as well.

Appendix B: Protocol Competition

Appendix B will formalize protocol competition. Platforms choose leverage terms (θ, τ) to maximize fee income subject to break-even risk, which can induce a race to high LTV. A sector-wide LM cap (equivalently, a minimum surcharge schedule) implements the constrained-efficient allocation under competition by replicating the planner wedge from Proposition 1.

B.1 Competition.

There are $N \geq 2$ platforms. In each t, platform i chooses $(\theta_{i,t}, \tau_{i,t})$ with $\theta_{i,t} \in (0,1)$ and $\tau_{i,t} \geq 0$. Loan/issuance volume is $L_{i,t} = \sigma_{i,t}(v(\theta_{i,t}) - \tau_{i,t})\,\mathcal{L}_t$ where $\mathcal{L}_t > 0$ and $\sigma_{i,t}(\cdot)$ is increasing (more attractive terms attract flow). Break-even risk requires

$$\tau_{i,t} \geq c_t(\theta_{i,t}), \qquad c_t'(\theta) > 0. \tag{30}$$

Platforms maximize fee income $\Pi_{i,t} = \tau_{i,t} L_{i,t}$.

B.2 Race to high θ.

For any $\theta_{i,t}$, raising $\tau_{i,t}$ above $c_t(\theta_{i,t})$ weakly reduces $v(\theta_{i,t}) - \tau_{i,t}$ and hence weakly reduces $L_{i,t}$, so an optimal fee satisfies $\tau_{i,t} = c_t(\theta_{i,t})$. Thus the platform problem reduces to choosing $\theta_{i,t}$ to maximize $\Pi_{i,t}(\theta_{i,t}) = c_t(\theta_{i,t})\,\sigma_{i,t}(v(\theta_{i,t}) - c_t(\theta_{i,t}))\,\mathcal{L}_t$.

We assume that the net-attractiveness term $v(\theta) - c_t(\theta)$ is strictly increasing:

$$v'(\theta) > c_t'(\theta) \quad \text{for all } \theta \in (0,1). \tag{31}$$

Proposition Under the previous assumption, competition induces a race to high LTV: in any equilibrium, platforms choose $\theta_{i,t}$ at the maximal feasible value.

Proof. Since $\sigma_{i,t}$ is increasing, Assumption X.1 implies higher $\theta_{i,t}$ strictly raises $v(\theta_{i,t}) - c_t(\theta_{i,t})$ and hence strictly raises $L_{i,t}$, while also raising the fee rate $c_t(\theta_{i,t})$. Therefore $\Pi_{i,t}$ is strictly increasing in $\theta_{i,t}$, so the best response sets $\theta_{i,t}$ at the upper bound. □

B.3: Sector-wide LM cap and implementation.

The stress multiplier is $LM_t^{\text{depeg}} = 1/[\kappa_t(1-\theta_t+\tau_t)]$. A cap $LM_t^{\text{depeg}} \leq LM^*$ is equivalent to

$$\tau_t \geq \tau_t^{cap}(\theta_t) \equiv \max\left\{0,\ \frac{1}{\kappa_t LM^*} - (1-\theta_t)\right\}. \tag{32}$$

Let τ_t^{SP} denote the planner-implementing wedge from Proposition 1. Impose the sector-wide schedule

$$\tau_{i,t} = \max\left\{c_t(\theta_{i,t}),\ \tau_t^{cap}(\theta_{i,t}),\ \tau_t^{SP}\right\}. \tag{33}$$

Choose LM^* so that $\tau_t^{cap}(\theta_t^{SP}) = \tau_t^{SP}$ along the planner path. Then in the regulated platform game, θ_t^{SP} is optimal and the induced wedge equals τ_t^{SP}.

Proposition Under (33) and $\tau_t^{cap}(\theta_t^{SP}) = \tau_t^{SP}$, the regulated competitive equilibrium implements the constrained-efficient allocation.

Proof. At θ_t^{SP}, (33) implies $\tau_{i,t} = \tau_t^{SP}$ for all i. By Proposition 1, substituting τ_t^{SP} into the decentralized Euler equation for the wrapped claim yields the planner implementability condition, and together with feasibility/collateral constraints implies the planner allocation. □

Appendix C: Data Collection

C.1 Cross-Chain Wrapped Stablecoin Data

We construct a daily panel for stablecoin minting and redemption for Ethereum, BNB Smart Chain (BSC), Avalanche, Optimism, Polygon, Base, Gnosis, and BSC. ERC-20 lets anyone deploy a token that declares USDC and arbitrary decimals. In practice, there are hundreds of look-alikes that mint a cartoon-large supply in a single tx (e.g., 1e40) or report the wrong decimals, then never trade. Hence, we are simply not able to use criteria like ranking by issuance as that would pick them as the "largest USDC". Even with criteria like name/symbol whitelists, these are still spoofable and repeatedly surfaced fake contracts that dwarf real flows. Hence, given that spoof resistance cannot be solved we anchor inclusion externally: (1) issuer-canonical contracts from issuer documentation and (2) issuer-recognized wrapped/bridge contracts from official bridge registries that are USD-pegged only. We freeze this master list ex-ante and tag assets as base (issuer-canonical on that chain) or wrapped (issuer-canonical lives elsewhere). We include wrapped tokens only when (i) listed in the bridge's official registry (e.g., Wormhole Portal token list) or (ii) independently validated on-chain via bridge-router mints/withdrawals and contract-deployer provenance (in this case we disclose the absence of a live registry page and will only include such assets in robustness).

C.2: Spread

We measure the cost of recursive borrowing using a protocol–chain spread τ. For each protocol ℓ on chain c at time t, we define

$$\tau_{\ell,c,t} = \max\Big(r^{\text{borrow}}_{\ell,c,t} - r^{\text{supply}}_{\ell,c,t},\, 0\Big), \tag{34}$$

where $r^{\text{borrow}}_{\ell,c,t}$ and $r^{\text{supply}}_{\ell,c,t}$ are the annualized stablecoin borrowing and supply rates. We interpret $\tau_{\ell,c,t}$ as a wedge in the recursive leverage formula: higher spreads increase the carrying cost of levered positions and reduce effective leverage capacity.

When multiple stablecoins trade within a protocol on a given chain (e.g., USDC, USDT, DAI), we collapse to a single protocol-level borrowing rate using variable debt outstanding as weights:

$$\bar{r}^{\text{borrow}}_{\ell,c,t} = \frac{\sum_{m\in M_c} D_{m,t}\, r^{\text{borrow}}_{m,t}}{\sum_{m\in M_c} D_{m,t}}, \tag{35}$$

where $D_{m,t}$ denotes variable debt outstanding in market m at time t and M_c is the set of stablecoin markets on chain c. We use variable debt rather than total debt because variable-rate positions capture marginal borrowing activity and are most informative about current leverage conditions. We aggregate supply rates analogously, using supply balances as weights.

All rate observations are aligned to a common timestamp. For daily data, we anchor rates to end-of-day (23:59:59 UTC). For five-minute data, we floor timestamps to the nearest five-minute boundary and collapse duplicates by taking the mean, ensuring consistent alignment across protocols with different block times and indexing cadences.

We exclude AMPL (Ampleforth) from all aggregations. AMPL's rebasing mechanism mechanically expands and contracts supply, which can distort reported APRs in ways unrelated to borrowing costs and introduce spurious variation in τ driven by rebase events rather than market conditions.

Protocol coverage is constructed as follows. For Aave V2 and V3 deployments, we auto-discover active reserve markets by querying the protocol's on-chain registry contracts, which enumerate supported assets and their configuration; this provides complete coverage as markets are added. For Compound V2, auto-discovery is not possible due to the absence of a comparable registry, so we maintain a manually curated list of cToken addresses for each deployment and update it periodically to reflect governance-approved market additions. Compound V3 (Comet) uses an isolated-market architecture that we query directly via the Comet contract interface.

Finally, we address rare data artifacts that imply mechanically impossible negative supplies. We reconstruct protocol-level APRs as USD-value-weighted averages of per-reserve rates using balances derived from ERC-20 transfer logs. In rare cases, incomplete early transfer histories can produce negative implied supplies and slightly negative protocol APRs. We screen for such artifacts and clamp implied balances at zero prior to aggregation. These cases occur only for Aave V2 on Avalanche, where closer inspection indicates negative implied aToken supplies; for all other markets, no negative protocol APRs are observed, so this adjustment has no effect.

C.3: Regulatory Announcements

We pre-register a regulator universe (NYDFS, NY Attorney General, CFTC, SEC) and include only primary, official releases from these authorities. An item enters the event set if it is the first public release of a binding action—license approval/denial, court order/injunction, settlement/consent order, or enforcement action—and it explicitly targets a named stablecoin or named stablecoin issuer in a way that can be mapped to on-chain exposures. Each event must have a clearly identifiable announcement date on the regulator page/document (or on an official regulator index that lists the release date). We exclude speeches, consultations, non-binding commentary, general crypto actions that do not explicitly mention stablecoins, protocol-level DeFi parameter changes, and duplicative restatements of prior actions; if multiple releases occur for the same matter, we either retain only the earliest release or group later stages under a pre-specified event-family identifier for robustness.

C.4: Control Variables

We include time-varying controls to absorb confounding variation in chain-level activity and market conditions. The daily return on Bitcoin (ret_{maj}) captures market-wide sentiment and price movements that influence leverage demand. Log total value locked (log(TVL)) proxies for chain-level liquidity depth. Decentralized exchange trading volume (dex_vol) measures on-chain trading activity, which is mechanically related to stablecoin demand. Gas fees (gas_fee) capture network congestion that may affect arbitrage and

liquidation dynamics. Stablecoin supply growth (supply_growth) controls for aggregate minting and redemption activity, isolating the effect of wrapped stablecoin creation from broader stablecoin dynamics.

For the panel regression (H1), we include all five controls—ret_{maj}, log(TVL), dex_vol, gas_fee, and supply_growth—along with chain and month fixed effects. For the event-study analysis (H2), we include the same controls measured at five-minute frequency, with event and protocol–chain fixed effects. For the cross-chain contagion analysis, we use log(TVL) and bridge_count (the number of active bridge protocols on each chain) as the primary explanatory variables in a cross-sectional regression; we additionally report specifications controlling for pre-event platform concentration (pre_HHI), measured as the Herfindahl–Hirschman Index of bridge flows in the seven days prior to the shock. All continuous variables are winsorized at the 1st and 99th percentiles.

Appendix D: Result Tables

D.1: Wrapped Stablecoin and Leverage

Table 3: Exclusion Restriction Tests: Regulatory Announcements and Borrowing Rates

	Panel A: Non-Stablecoin (Placebo)		Panel B: Stablecoin (Positive Control)	
	(1) Pooled	(2) Asset FE	(3) Pooled	(4) Asset FE
Regulatory Event	−0.152	−0.200	2.046∗	2.059∗
	(0.394)	(0.412)	(1.104)	(1.111)
Asset FE	No	Yes	No	Yes
Observations	27,254	27,254	16,400	16,400
R^2	0.000	0.000	0.000	0.000

	Panel C: Heterogeneity by Targeted Stablecoin
	(5)
Regulatory Event	2.768
	(1.893)
Directly Targeted	−1.431
	(0.997)
Reg Event × Targeted	−0.427
	(2.251)
Observations	10,571
R^2	0.003

Notes: This table tests the exclusion restriction for our IV strategy for our first hypothesis. The dependent variable is the variable borrowing rate (APY, percentage points). Panel A (columns 1–2) presents placebo tests using non-stablecoin assets (ETH, BTC, LINK, etc.); Panel B (columns 3–4) shows positive control results for stablecoin assets. Panel C tests whether regulatory events have larger effects on directly targeted stablecoins (USDT, BUSD, TUSD) compared to untargeted stablecoins (USDC, DAI, GUSD, SUSD). If regulatory announcements affect leverage only through the stablecoin channel, coefficients in Panel A should be zero while Panel B shows significant effects. Standard errors clustered by asset are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 4: Instrument Channel Specificity and Timing Robustness

	(1) Main	(2) Bridge	(3) Operational	(4) Lag 7d	(5) Lag 30d	(6) Cumulative
Panel A: Second Stage — Dependent Variable: Liquidity Multiplier						
Wrapped Creation ($M)	0.00073***	−0.00318	−0.00073	0.00101***	0.00020	0.00730***
	(0.00014)	(0.00542)	(0.00114)	(0.00016)	(0.00015)	(0.00180)
Panel B: First Stage Diagnostics						
Partial F-statistic	19.97	3.16	2.74	10.19	75.08	72.32
First-stage coefficient	32.70	20.96	−30.98	23.38	63.19	−21.45
Strong instrument (F > 10)	✓			✓	✓	✓
Instrument	Z_t^{reg}	Z_t^{bridge}	Z_t^{oper}	Z_{t-7}^{reg}	Z_{t-30}^{reg}	$\sum_{s\leq t} Z_s^{reg}$
Instrument events	11	91	33	11	11	—
Chain FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	7,578	7,578	7,578	7,578	7,578	7,578
Clusters	257	257	257	257	257	257

Notes: This table evaluates instrument validity for Hypothesis 1. Column (1) reports the main specification using the regulatory announcement instrument. Columns (2)–(3) test alternative DeFi shock channels: bridge exploits and protocol operational events. The weak first-stage F-statistics ($F = 3.16$ and $F = 2.74$) confirm that the instrument operates specifically through the regulatory channel, not through generic DeFi shocks. Columns (4)–(6) examine robustness to instrument timing using lagged and cumulative transformations of the regulatory shock. All specifications include chain and month fixed effects, and control for major crypto returns, log TVL, DEX volume, gas fees, and stablecoin supply growth. Standard errors clustered at the chain-month level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 5: Subsample Stability Tests: Wrapped Stablecoin Creation and Leverage

Subsample	Coefficient	Std. Error	p-value	F-stat	N
Full Sample	0.0007***	(0.0001)	0.000	14.63	7,578
Post Terra-Luna	0.0007***	(0.0001)	0.000	10.90	7,343
High TVL Chains	0.0005**	(0.0002)	0.017	16.52	2,826
Low TVL Chains	-0.0200	(0.0245)	0.415	0.27†	4,752
Ethereum Only	0.0001	(0.0001)	0.287	18.84	1,309
L2s Only	-0.0020*	(0.0011)	0.074	0.97†	4,806
Panel A: Time Period					
Post Terra-Luna: May 9, 2022 onwards					
Panel B: Chain Characteristics					
High TVL: Ethereum, BNB, Arbitrum					
Low TVL: Avalanche, Base, Gnosis, Optimism, Polygon					
L2s: Arbitrum, Base, Gnosis, Optimism, Polygon					

Notes: This table presents 2SLS estimates of the effect of wrapped stablecoin creation (in millions USD) on the liquidity multiplier (leverage) across different subsamples. The instrument is regulatory announcements affecting stablecoin issuers. All specifications include chain and month fixed effects (where applicable) and control for major crypto returns, log TVL, DEX volume, gas fees, and stablecoin supply growth. Standard errors clustered by chain-month are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. †Weak instrument (F < 10); coefficient estimates unreliable.

Table 6: Placebo Test: Shifted Regulatory Announcement Dates

	Main (Actual Dates)	**Placebo** (+30 Days)
Panel A: Second Stage (2SLS)		
Wrapped Creation (mm USD)	0.000734***	0.000195
	(0.000143)	(0.000151)
t-statistic	5.121	1.294
p-value	0.000	0.196
Panel B: First Stage		
Partial F-statistic	19.97	75.07
Panel C: Sample		
Observations	7,578	7,578
Instrument Events	11	11
Clusters (chain-month)	257	257
Controls	Yes	Yes
Chain FE	Yes	Yes
Month FE	Yes	Yes

Notes: This table reports placebo test results for Hypothesis 1. The main specification uses actual regulatory announcement dates as the instrument. The placebo specification shifts all announcement dates forward by 30 days. Standard errors clustered at the chain-month level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

D.2: Recursive Leverage and Volatility

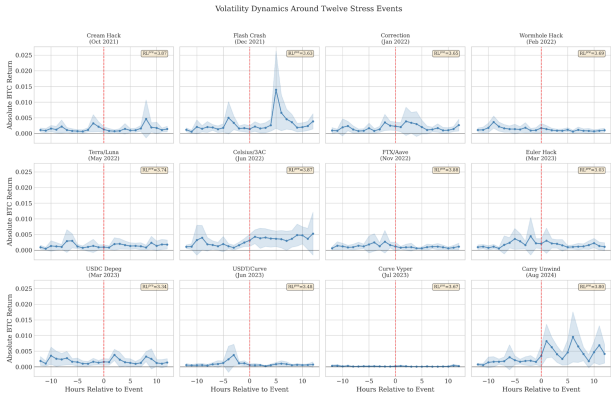

Figure 1: Average absolute 5-minute BTC returns around twelve stress events. Dashed line marks event time ($t = 0$). Pre-event aggregate recursive leverage (RL_{pre}) displayed for each event. Higher RL is associated with larger post-event volatility increases.

Table 7: Placebo Outcomes: Recursive Leverage and Non-DeFi Asset Volatility

	(1) BTC	(2) S&P 500	(3) Gold	(4) EUR/USD	(5) Crude Oil
Post	−0.0274	−0.0104	−0.0082	0.0172	0.0291
	(0.0975)	(0.0360)	(0.0408)	(0.0161)	(0.0469)
RL^{pre}	−0.0116	0.0014	−0.0065	0.0045	0.0005
	(0.0170)	(0.0084)	(0.0086)	(0.0030)	(0.0058)
Post × RL^{pre}	0.0107	0.0024	0.0022	−0.0044	−0.0054
	(0.0267)	(0.0099)	(0.0110)	(0.0044)	(0.0129)
Observations	132	80	80	83	80
Events	12	12	12	12	12
R^2	0.040	0.013	0.034	0.039	0.083

Notes: Dependent variable is daily absolute return (volatility proxy). Column (1) uses Bitcoin as a positive control; Columns (2)–(5) use non-DeFi placebo assets that are not used as collateral in Aave lending protocols. RL^{pre} is average recursive leverage across all protocol–chain pairs in the 24 hours before the event. Sample includes ±5 trading days around twelve stress events (October 2021 to August 2024). HC3 heteroskedasticity-robust standard errors in parentheses. None of the placebo interaction terms are statistically significant at conventional levels ($p > 0.30$ for all), confirming that the RL amplification channel is specific to crypto markets. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 8: H2 Event Study: BTC vs. ETH Volatility as Outcome

	Aggregated		Disaggregated	
	(1) BTC	(2) ETH	(3) BTC	(4) ETH
Post	−0.002780	−0.005618	−0.000047	−0.000678
	(0.003649)	(0.003754)	(0.000653)	(0.000663)
RL^{pre}	0.000490***	0.000612***	−0.000121	−0.000214
	(0.000068)	(0.000078)	(0.000117)	(0.000135)
Post × RL^{pre}	0.001009	0.001822*	0.000244	0.000431
	(0.001031)	(0.001078)	(0.000235)	(0.000271)
Observations	13,836	13,836	48,714	48,714
Events	12	12	12	12
R^2	0.133	0.154	0.145	0.160
Event FE	Yes	Yes	Yes	Yes
Clustering	Event	Event	Event	Event

Notes: Dependent variable is absolute 5-minute return (volatility proxy). Columns (1)–(2) aggregate to the event-time level (one observation per event × timestamp); Columns (3)–(4) use protocol–chain-level observations. RL^{pre} is pre-event recursive leverage measured at $k = -1$. The ETH coefficient on Post × RL^{pre} is approximately 1.8 times larger than the BTC coefficient in both specifications, consistent with the DeFi amplification channel operating more directly through ETH-denominated collateral. Standard errors clustered by event in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

D.3: Cross Chain Contagion and Bridge Connectivity

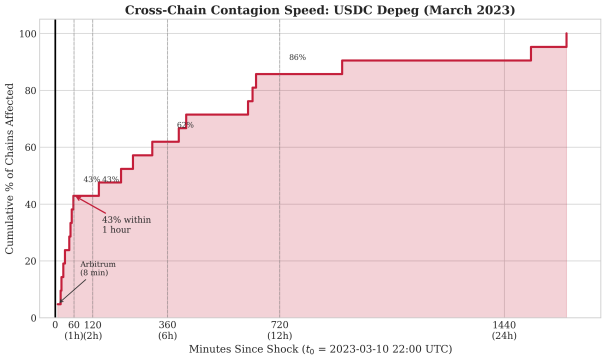

Figure 2: Cumulative Cross-Chain Contagion During USDC Depeg March 2023

Table 9: Cross-Chain Contagion Timing: USDC Depeg (March 2023)

Chain	Delay (min)	First Stress (UTC)	Outflow ($)	Bridges
Arbitrum	8	2023-03-10 22:08	−154K	8
Fantom	17	2023-03-10 22:17	−10K	5
Ethereum	20	2023-03-10 22:20	−5,132K	15
BSC	26	2023-03-10 22:26	−939K	8
Gnosis	31	2023-03-10 22:31	−176K	2
Avalanche	45	2023-03-10 22:45	−944K	7
Near	49	2023-03-10 22:49	−461K	2
Moonbeam	53	2023-03-10 22:53	−19K	3
Polygon	58	2023-03-10 22:58	−2,702K	10
Optimism	140	2023-03-11 00:20	−252K	6
Celo	211	2023-03-11 01:31	−99K	2
Tron	311	2023-03-11 03:11	−15K	2
Solana	396	2023-03-11 04:36	−333K	3
Astar	618	2023-03-11 08:18	−7K	1
Aurora	633	2023-03-11 08:33	−98K	2
Aptos	644	2023-03-11 08:44	−55K	2
zkSync Lite	1526	2023-03-11 23:26	−29K	1
Injective	1640	2023-03-12 01:20	−12K	1

Notes: Shock timestamp t_0 = 2023-03-10 22:00 UTC (first hour USDC < \$0.995). Delay is minutes from t_0 to first minute with net outflows exceeding two standard deviations below pre-event mean. Bridges = estimated number of active bridge protocols on each chain in March 2023. N=21 chains total; 3 minor chains omitted for space.

Table 10: Determinants of Cross-Chain Contagion Speed

	(1) TVL Only	(2) Bridges Only	(3) Combined
log(TVL)	−0.460***		0.065
	(0.101)		(0.117)
log(Bridge Count)		−1.435***	−1.578***
		(0.210)	(0.381)
Constant	4.231***	6.417***	6.666***
	(0.346)	(0.300)	(0.544)
Observations	21	21	21
R-squared	0.426	0.621	0.623

Notes: Dependent variable is log(minutes to first stress). Robust standard errors (HC1) in parentheses. TVL measured in billions USD as of March 2023. Bridge count is the estimated number of active bridge protocols connecting each chain. *** p<0.01, ** p<0.05, * p<0.1.

Table 11: Summary Statistics: Cross-Chain Contagion Sample

Variable	Mean	Std. Dev.	Min	Max
Panel A: Contagion Timing (N=21 chains)				
Delay (minutes)	382	461	8	1,640
First stress outflow (thousands)	553	1,187	7	5,132
Panel B: Chain Characteristics				
TVL (billions)	2.34	6.21	0.005	28.5
Bridge count	3.9	3.7	1	15
Panel C: Pre-Event Bridge Flows (7-day window)				
Mean daily net inflow (thousands)	142	389	-52	1,840
Std. dev. daily net inflow (thousands)	287	412	18	1,650
Platform HHI	0.993	0.009	0.975	1.000

Notes: Panel A reports contagion timing outcomes. Delay is minutes from shock timestamp (t_0 = 2023-03-10 22:00 UTC) to first minute with net outflows exceeding two standard deviations below pre-event mean. Panel B reports chain characteristics as of March 2023. TVL from DeFiLlama; bridge count estimated from ecosystem data. Panel C reports pre-event bridge flow statistics computed over the 7 days prior to t_0. Platform HHI is the Herfindahl-Hirschman Index of bridge platform concentration.

Table 12: Wallet-Level Recursive Leverage Loop During USDC Depeg

Step	Timestamp (UTC)	Action	Amount ($)	Cumulative
Wallet: `0xc79c30ef1941002c54293a028cf252dfb0ddd2aa`				
1	2023-03-11 11:29	Deposit USDC	12,296,427	Collateral: $12.3M
2	2023-03-11 11:29	Borrow USDC	1,435,332	Debt: $1.4M
3	2023-03-11 14:16	Deposit USDC	3,095,167	Collateral: $15.4M
4	2023-03-11 14:19	Borrow USDC	405,675	Debt: $1.8M
5	2023-03-11 20:34	Deposit USDC	7,865,641	Collateral: $23.3M
6	2023-03-11 20:38	Borrow USDC	530,248	Debt: $2.4M
7	2023-03-11 20:55	Deposit USDC	7,664,299	Collateral: $30.9M
8	2023-03-11 21:05	Borrow USDC	4,749,627	Debt: $7.1M
Result:				
Initial Capital				$12,296,427
Total Position (Deposited)				$30,921,535
Total Borrowed				$7,120,882
Effective Leverage				**2.51x**
Duration				9.6 hours
Liquidated?				No

Notes: This table documents a real recursive leverage loop observed on Aave V2 during the USDC depeg event (March 10–14, 2023). The wallet executes a classic "deposit → borrow → re-deposit" strategy, achieving 2.51x effective leverage from an initial $12.3M position. Theoretical maximum leverage at 75% LTV is 4x. The wallet survived the depeg event without liquidation, consistent with same-asset recursive leverage being protected during price dislocations (both collateral and debt lose value proportionally).

Table 13: Robustness: Alternative Stress Arrival Definitions

Stress Definition	Coefficient	Std. Error	N
Panel A: Relative Thresholds (σ below pre-event mean)			
1.5σ threshold	$-1.460 * **$	(0.252)	18
2.0σ threshold (baseline)	$-1.498 * **$	(0.246)	18
2.5σ threshold	$-1.466 * **$	(0.224)	18
Panel B: Absolute Dollar Thresholds			
\$10k outflow	$-1.867 * **$	(0.385)	17
\$25k outflow	$-1.614 * **$	(0.419)	15
\$50k outflow	$-1.263 * **$	(0.408)	14
\$100k outflow	$-1.470 * *$	(0.628)	13

Notes: Dependent variable is log(minutes to first stress arrival). All specifications regress log(delay) on log(bridges) with robust standard errors. Panel A varies the threshold for stress arrival relative to pre-event volatility: $k\sigma$ means net outflows exceeding k standard deviations below the 7-day pre-event mean. Panel B uses absolute dollar thresholds regardless of pre-event volatility. *** p<0.01, ** p<0.05, * p<0.1.

D.4: Recursive Leverage Example

Table 14: Bridge Connectivity and Contagion Speed: Evidence Across Seven Stress Events

Event	Date	N	log(Bridges)	p-value	R^2
Panel A: 2022 Events					
Terra/Luna collapse	May 2022	11	−1.413**	0.013	0.294
			(0.569)		
Celsius/3AC crisis	Jun 2022	13	−0.432	0.368	0.042
			(0.480)		
FTX collapse	Nov 2022	15	−1.255*	0.055	0.274
			(0.654)		
Jan 2022 correction	Jan 2022	12	−0.326	0.201	0.082
			(0.255)		
Panel B: 2023 Events					
USDC depeg (SVB)	Mar 2023	24	−1.679***	<0.001	0.612
			(0.190)		
USDT Curve stress	Jun 2023	19	−0.963*	0.057	0.171
			(0.505)		
Curve Vyper exploit	Jul 2023	21	−0.562	0.307	0.048
			(0.549)		
Panel C: Meta-Analysis					
Events with negative coefficient			7/7 (100%)		
Events significant at 10%			4/7 (57%)		
Events significant at 5%			2/7 (29%)		
Average coefficient			−0.947		
Median coefficient			−0.963		

Notes: This table reports the effect of bridge connectivity on cross-chain contagion speed across seven major crypto stress events. The dependent variable is log(minutes to first stress arrival) for each destination chain. log(Bridges) is the log of the number of active bridge protocols connecting to each chain. All regressions use robust standard errors (in parentheses). The sample excludes the shock-origin chain where applicable (e.g., Terra for the Terra/Luna event). *** p<0.01, ** p<0.05, * p<0.1.

Table 15: Reverse Causality Check: Controlling for User Sophistication

	(1) Baseline	(2) +Wallets	(3) +Avg Tx	(4) +DEX Vol
log(Bridges)	$-1.498 * **$	$-1.192 * **$	$-1.137 * *$	-0.364
	(0.246)	(0.455)	(0.520)	(1.295)
log(Unique Wallets)		-0.149	-0.146	-0.346
		(0.190)	(0.203)	(0.212)
log(Avg Tx Size)			-0.058	-0.103
			(0.192)	(0.241)
log(DEX Volume)				-0.023
				(0.228)
Constant	$6.561 * **$	$7.044 * **$	$7.443 * **$	$8.271 * *$
	(0.397)	(0.704)	(1.104)	(3.753)
Observations	18	18	18	16
R-squared	0.605	0.618	0.621	0.582

Notes: Dependent variable is log(minutes to first stress arrival). User sophistication proxies computed from 7-day pre-event period: Unique Wallets is total unique sender/receiver addresses in bridge transactions; Avg Tx Size is mean USD per bridge transaction; DEX Volume is average daily decentralized exchange volume. Column 4 should be interpreted with caution: VIF for log(Bridges) = 10.2 indicates severe multicollinearity with DEX Volume ($\rho = 0.88$), inflating standard errors and preventing identification of individual effects. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 16: Robustness Checks: Bridge Connectivity and Contagion Speed

Specification	Coefficient	p-value
Panel A: Alternative Specifications		
Rank regression: rank(Bridges)	−0.701	0.008
Excluding Tron: log(Bridges)	−1.473	<0.001
Panel B: Binary Treatment Test		
	Mean Delay	N
Many bridges (>3)	45 min	7
Few bridges (≤3)	550 min	14
Difference	505 min	
Mann-Whitney U		<0.001
Panel C: Diagnostics		
Breusch-Pagan test		p = 0.760
Shapiro-Wilk test		p = 0.074
VIF	3.80	
Correlation(TVL, Bridges)	0.86	
Bootstrap sign stability	99% negative	

Notes: Panel A reports coefficient on bridge connectivity under alternative specifications. Panel B compares mean delay for chains above vs. below median bridge count. Panel C reports diagnostic tests. Bootstrap based on 1,000 resamples; sign stability indicates percentage of resamples where log(Bridges) coefficient is negative.

Table 17: Wallet-Level Evidence of Recursive Leverage: Example 1

Step	Timestamp (UTC)	Action	Amount (USDC)	Cumulative Position
Wallet: `0xc79c30ef1941002c54293a028cf252dfb0ddd2aa` *Date: March 11, 2023 (USDC Depeg Event)*				
1	11:29:11	Deposit	$12,296,427	Collateral: $12.3M
2	11:29:11	Borrow	$1,435,332	Debt: $1.4M
3	14:15:59	Deposit	$3,095,167	Collateral: $15.4M
4	14:18:59	Borrow	$405,675	Debt: $1.8M
5	20:33:47	Deposit	$7,865,641	Collateral: $23.3M
6	20:37:59	Borrow	$530,248	Debt: $2.4M
7	20:55:23	Deposit	$7,664,299	Collateral: $30.9M
8	21:04:59	Borrow	$4,749,627	Debt: $7.1M
Summary:				
Initial Capital				$12,296,427
Final Position (Total Deposited)				$30,921,535
Effective Leverage				**2.51**×

Table 18: Wallet-Level Evidence of Recursive Leverage: Example 2

Step	Timestamp (UTC)	Action	Amount (USDT)	Cumulative Position
Wallet: `0x777777c9898d384f785ee44acfe945efdff5f3e0` *Date: March 2023 (USDC Depeg Event)*				
1	–	Deposit	$5,638,740	Collateral: $5.6M
2	–	Borrow	$478,530	Debt: $0.5M
3	–	Deposit	$289,122	Collateral: $5.9M
4	–	Borrow	$478,530	Debt: $1.0M
		... continues for 12 iterations ...		
Summary:				
Initial Capital				$5,638,740
Final Position (Total Deposited)				$9,106,201
Effective Leverage				**1.61**×
Iterations				12

Table 19: Wallet-Level Evidence of Recursive Leverage: Example 3

Step	Timestamp (UTC)	Action	Amount (DAI)	Cumulative Position
Wallet: `0x5cef0f6e78b0703aedcc0f6a1f4a9d76ca8c4d7e` *Date: November 2022 (FTX Collapse)*				
1	–	Deposit	$91,840	Collateral: $91.8K
2	–	Borrow	$61,500	Debt: $61.5K
3	–	Deposit	$61,500	Collateral: $153.3K
4	–	Borrow	$61,000	Debt: $122.5K
5	–	Deposit	$32,000	Collateral: $185.3K
6	–	Borrow	$61,000	Debt: $183.5K
Summary:				
Initial Capital				$91,840
Final Position (Total Deposited)				$185,340
Effective Leverage				**2.02**×
Iterations				3

Acknowledgments

We thank seminar participants at the University of Missouri for helpful comments and suggestions. We are grateful to Dune Analytics for data access. All errors are our own.

References

Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi. 2015. Systemic risk and stability in financial networks. *American Economic Review* 105:564–608.

Adrian, T., and H. S. Shin. 2010. Liquidity and leverage. *Journal of Financial Intermediation* 19:418–37.

Ahmed, R., I. Aldasoro, and C. Duley. 2024. Public information and stablecoin runs. BIS Working Paper 1164, Bank for International Settlements. Revised 2025.

Aitken, J., and M. M. Singh. 2010. *The (sizable) role of rehypothecation in the shadow banking system.* International Monetary Fund.

Allen, F., and D. Gale. 1998. Optimal financial crises. *Journal of Finance* 53.

Allen, H. J. 2023. DeFi: Shadow banking 2.0? *William & Mary Law Review* 64:919–81.

Anadu, K., J. Azar, M. Cipriani, T. M. Eisenbach, C. Huang, M. Landoni, G. La Spada, M. Macchiavelli, A. Malfroy-Camine, and J. C. Wang. 2024. Runs and flights to safety: Are stablecoins the new money market funds? Staff Report 1073, Federal Reserve Bank of New York.

Arner, D. W., R. Auer, and J. Frost. 2020. Stablecoins: risks, potential and regulation .

Bappy, F. H., E. Cheon, and T. Islam. 2025. Centralized trust in decentralized systems: Unveiling hidden contradictions in blockchain and cryptocurrency. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, 1960–71.

Bernanke, B. S., and M. Gertler. 1995. Inside the black box: the credit channel of monetary policy transmission. *Journal of Economic perspectives* 9:27–48.

Bernardo, A. E., and I. Welch. 2004. Liquidity and financial market runs. *Quarterly Journal of Economics* 119:135–58.

Bianchi, J., and E. G. Mendoza. 2018. Optimal time-consistent macroprudential policy. *Journal of Political Economy* 126:588–634.

Bianchi, J., and M. E. G. Mendoza. 2011. *Overborrowing, financial crises and 'macroprudential'policy.* International Monetary Fund.

Brunnermeier, M. K., and L. H. Pedersen. 2009. Market liquidity and funding liquidity. *Review of Financial Studies* 22:2201–38.

Brunnermeier, M. K., and Y. Sannikov. 2014. A macroeconomic model with a financial sector. *American Economic Review* 104.

Cameron, A. C., J. B. Gelbach, and D. L. Miller. 2011. Robust inference with multiway clustering. *Journal of Business & Economic Statistics* 29:238–49.

Carlstrom, C. T., and T. S. Fuerst. 1997. Agency costs, net worth, and business fluctuations: A computable general equilibrium analysis. *The American Economic Review* 893–910.

Chainalysis Team. 2023. Here's What On-Chain Data Tells Us About Crypto's Reaction to the Demise of Silicon Valley Bank And Its Impact on USDC. https://www.chainalysis.com/blog/crypto-market-usdc-silicon-valley-bank/. Accessed: 2025-07-05.

d'Avernas, A., V. Maurin, and Q. Vandeweyer. 2022. Can stablecoins be stable? BFI Working Paper 2022-131, Becker Friedman Institute.

Diamond, D. W., and P. H. Dybvig. 1983. Bank runs, deposit insurance, and liquidity. *Journal of Political Economy* 91:401–19.

Eichengreen, B., M. T. Nguyen, and G. Viswanath-Natraj. 2023. Stablecoin devaluation risk. Working Paper, SSRN. Working paper.

Eisenberg, L., and T. H. Noe. 2001. Systemic risk in financial systems. *Management Science* 47:236–49.

Elliott, M., B. Golub, and M. O. Jackson. 2014. Financial networks and contagion. *American Economic Review* 104:3115–53.

Farhi, E., M. Golosov, and A. Tsyvinski. 2009. A theory of liquidity and regulation of financial intermediation. *Review of Economic Studies* 76:973–92.

Fostel, A., and J. Geanakoplos. 2008. Leverage cycles and the anxious economy. *American Economic Review* 98:1211–44.

Geanakoplos, J. 2010. The leverage cycle. Cowles Foundation Discussion Paper 1751, Cowles Foundation for Research in Economics, Yale University.

Gertler, M., and N. Kiyotaki. 2010. Financial intermediation and credit policy in business cycle analysis. In *Handbook of monetary economics*, vol. 3, 547–99. Elsevier.

Goldstein, I., and A. Pauzner. 2005. Demand-deposit contracts and the probability of bank runs. *Journal of Finance* 60.

Gorton, G., and A. Metrick. 2012. Securitized banking and the run on repo. *Journal of Financial Economics* 104:425–51.

Gorton, G. B., E. C. Klee, C. P. Ross, S. Y. Ross, and A. P. Vardoulakis. 2022. Leverage and stablecoin pegs. Working Paper, National Bureau of Economic Research.

Gromb, D., and D. Vayanos. 2002. Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics* 66:361–407.

Iacoviello, M. 2005. House prices, borrowing constraints, and monetary policy in the business cycle. *American economic review* 95:739–64.

Jacklin, C. J. 1987. Demand deposits, trading restrictions, and risk sharing. In *Contractual Arrangements for Intertemporal Trade*. University of Minnesota Press.

Jeanne, O., and A. Korinek. 2010. Excessive volatility in capital flows: A pigouvian taxation approach. *American Economic Review* 100:403–7.

Kacperczyk, M., and P. Schnabl. 2013. How safe are money market funds? *Quarterly Journal of Economics* 128:1073–122.

Kiyotaki, N., and J. Moore. 1997. Credit cycles. *Journal of political economy* 105:211–48.

Koont, N., Y. Ma, L. Pastor, and Y. Zeng. 2022. Steering a ship in illiquid waters: Active management of passive funds. NBER Working Paper 30039, National Bureau of Economic Research.

Li, Y., and S. Mayer. 2021. Money creation in decentralized finance: A dynamic model of stablecoin and crypto shadow banking. Working Paper, SSRN. Working paper.

Liu, J., I. Makarov, and A. Schoar. 2023. Anatomy of a run: The terra luna crash. NBER Working Paper 31160, National Bureau of Economic Research.

Lorenzoni, G. 2008. Inefficient credit booms. *The Review of Economic Studies* 75:809–33.

Lyons, R. K., and G. Viswanath-Natraj. 2023. What keeps stablecoins stable? *Journal of International Money and Finance* 131. Article 102777.

Parlatore, C. 2016. Fragility in money market funds: Sponsor support and regulation. *Journal of Financial Economics* 121:595–623.

Pozsar, Z., T. Adrian, A. Ashcraft, and H. Boesky. 2013. Shadow banking. *Federal Reserve Bank of New York Economic Policy Review* 19.

Routledge, B., and A. Zetlin-Jones. 2022. Currency stability using blockchain technology. *Journal of Economic Dynamics and Control* 142. Article 104155.

Schmidt, L., A. Timmermann, and R. Wermers. 2016. Runs on money market mutual funds. *American Economic Review* 106:2625–57.

Singh, M. 2011a. Velocity of pledged collateral: Analysis and implications. IMF Working Paper WP/11/256, International Monetary Fund.

Singh, M. M. 2011b. *Velocity of pledged collateral: analysis and implications*. International Monetary Fund.

Sunderam, A. 2015. Money creation and the shadow banking system. *Review of Financial Studies* 28:939–77.

Tobin, J. 1963. Commercial banks as creators of 'money .

Uhlig, H. 2022. A luna-tic stablecoin crash. NBER Working Paper 30256, National Bureau of Economic Research.