Bachelor Thesis

Unstable Stability: Measuring Peg Deviations and Run Risk in Stablecoins

Authors:

Laura Leimane and Kārlis Kauliņš

Supervisor:

Boriss Siliverstovs

JEL codes: E42, G23, G14

April 2026

Riga

COPYRIGHT DECLARATION AND LICENCE

Names of the authors in full: Laura Leimane, Kārlis Kauliņš

Title of the Thesis: Unstable Stability: *Measuring Peg Deviations and Run Risk in Stablecoins*

We hereby certify that the above-named thesis is entirely the work of the persons named below, and that all materials, sources and data used in the thesis have been duly referenced. This thesis – in its entirety or in any part thereof – has never been submitted to any other degree commission or published.

In accordance with Section 1 of the Copyright Law of Latvia, the persons named below are the authors of this thesis.

Pursuant to Article 40 of the Copyright Law the authors hereby agree and give an explicit licence to SSE Riga to deposit one digital copy of this thesis in the digital catalogue and data base at SSE Riga Library for an unlimited time and without royalty. The licence permits SSE Riga to grant access to the duly deposited thesis to all users of the catalogue and data base without royalty and limitations to downloading, copying and printing of the digital thesis in whole or in part provided we are indicated as the authors of the thesis according to Clause 4 Section 1 Article 14 of Copyright Law. We assert our right to be identified as the author(s) of this thesis whenever it is reproduced in full or in part.

Signed

_____ _____

Laura Leimane Kārlis Kauliņš

Date
31.03.2026

Table of contents

Abstract

This thesis examines how reliably major U.S.-dollar stablecoins maintain their peg and when stability breaks down. It compares three leading tokens – USDT, USDC and DAI – to link design features and market structure to peg behavior. Using hourly data from 2020–2025, the study combines prices and trading volumes with granular Ethereum on-chain indicators, including issuance activity, exchange flows, wallet participation and supply. The empirical strategy integrates fixed-effects panel regressions, panel quantile regressions and VAR-based spillover measures to relate peg deviations to on-chain and market conditions, including tail behavior and cross-coin linkages. Results show that USDC has the tightest and least persistent deviations, USDT is moderately stable, and DAI exhibits larger and longer-lasting mispricings, especially during and after the 2022 crypto crisis. Arbitrage effectiveness declines in volatile, illiquid periods, and large deviations adjust more slowly than small ones. The findings imply that reserve quality, governance and arbitrage capacity jointly shape stablecoin resilience, while limited spillovers indicate that peg risk remains mostly coin-specific. On-chain issuance, exchange flows and volatility emerge as candidate early-warning signals for regulators, issuers and DeFi protocol designers.

1. Introduction

Stablecoins have become a central component of the modern crypto-financial ecosystem, serving as the main transactional asset in decentralized finance (DeFi), the dominant trading pair on exchanges, and a key link between blockchain markets and traditional finance. They promise to maintain a stable value- most often one U.S. dollar per token- while preserving the programmability and global reach of digital assets. As adoption has grown, stablecoins now settle large volumes, back lending and collateral in DeFi, and increasingly hold portfolios of traditional safe assets such as U.S. Treasury bills, giving their stability broader financial-stability relevance.

Recent episodes show, however, that stablecoins are not inherently stable. The collapse of TerraUSD in 2022 demonstrated how design failures can trigger large peg breaks, liquidity spirals and contagion (Duan & Urquhart, 2023). Even seemingly robust designs have wobbled: in March 2023, USD Coin temporarily fell to 0.88 following reserve-related disclosures, illustrating how information shocks can spark run-like behavior (Ahmed et al., 2024). These events support the view of stablecoins as a form of private money subject to classic fragilities in confidence, redeemability and liquidity (Gorton et al., 2025).

Stablecoin systems rely on different stabilisation mechanisms – fiat reserves, over-collateralised smart contracts, algorithmic supply rules and arbitrage incentives, and their effectiveness varies. Arbitrage helps keep prices near the peg, but when it is costly, illiquid or operationally constrained, deviations can persist (Pernice, 2021). Empirical work shows that public information releases, market-wide volatility and the concentration of arbitrage capacity can amplify runs and de-pegs (Ahmed et al., 2024; Lee et al., 2025; Ma et al., 2025). Design choices regarding collateral, governance and redemption also matter for resilience (Ling et al., 2025; Jarno & Kołodziejczyk, 2021).

Despite a rapidly growing literature, important empirical gaps remain. Many studies focus on single stablecoins or short windows around specific events. Others model de-pegging risk or analyse particular drivers such as volatility or disclosures, but few integrate on-chain microstructure variables – issuance and redemption flows, exchange inflows, velocity – into a multi-year, multi-stablecoin framework. Event studies are widely used to characterise individual stress episodes, yet there is limited evidence on whether day-to-day blockchain behaviour systematically predicts future peg deviations. As a result, the link between micro-level on-chain activity and aggregate stability is still not well quantified (Ante et al., 2023; Naifar, 2025).

This thesis addresses these gaps using a dataset that combines blockchain transfer records, exchange wallet classifications and market prices for three major

stablecoins – USDT, USDC and DAI – over 2020–2025. The choice of this set is deliberate. USDT is the largest and most actively traded stablecoin, widely used as a quote asset on centralized exchanges. USDC is regarded as a relatively transparent and tightly regulated fiat-backed coin, making it a natural benchmark in policy debates (Berentsen & Schär, 2019; Jarno & Kołodziejczyk, 2021). DAI is the leading decentralized, crypto-collateralised stablecoin governed on-chain by MakerDAO, representing a contrasting protocol-based stabilisation model (De Sclavis et al., 2025; Huo et al., 2022). Together, these three dominate by market capitalisation, trading volumes and DeFi usage while spanning centralized versus decentralized architectures, which makes them a natural comparative set for studying how design and market structure shape peg dynamics.

Methodologically, the thesis combines event-study techniques with high-frequency panel econometrics. Panel regressions then relate hourly deviations to lagged deviations, on-chain flows, net issuance and market conditions, allowing the analysis of both routine dynamics and stressed periods within a unified framework.

The thesis aims to make three main contributions. First, it links high-frequency peg behavior to detailed on-chain microstructure across several major stablecoins and over multiple years, addressing the scarcity of empirical work that fully exploits blockchain data alongside prices (Ante et al., 2023; Naifar, 2025). Second, it explicitly combines event-study methods with dynamic panel models, analysing stress episodes and everyday conditions jointly rather than in isolation (MacKinlay, 1997; Duan & Urquhart, 2023). Third, it quantifies how much of major de-peg events can be explained by contemporaneous on-chain and market variables, offering insights for issuers, regulators and protocol designers on potential early-warning indicators and on where interventions could most effectively enhance resilience.

The research is guided by three questions:

RQ1: How frequent and how large are hourly peg deviations of USDT, USDC and DAI from their one-dollar target between 2020 and 2025?

A key dimension of peg stability that this thesis investigates is whether downward deviations correct at the same rate as upward deviations. When a stablecoin trades below its dollar peg, holders can redeem tokens directly with the issuer at par value, providing an immediate arbitrage channel. When a stablecoin trades above parity, correction instead requires the issuer to expand supply into secondary markets, a process that depends on operational capacity and governance decisions. If these two channels operate at different speeds, a systematic directional asymmetry should be detectable in the data. We test this directly, and the answer has

implications for how regulators and risk managers should distinguish between stablecoins trading at a discount and those trading at a premium.

RQ2: How do changes in net issuance, hourly transaction velocity and exchange inflows predict next-day peg deviations across these stablecoins?

RQ3: How do major de-peg episodes affect short-run peg dynamics, and to what extent do on-chain activity and market conditions explain the abnormal deviations observed during these events?

The goal is to provide empirical evidence on the stability of major stablecoins by linking peg deviations to granular, blockchain-based measures of market behaviour. By combining long-horizon on-chain and off-chain data with panel methods, the thesis seeks to clarify how stablecoin stability is maintained or lost and to draw implications for financial stability, regulatory design and the architecture of decentralized monetary systems.

Main contributions. This thesis makes the following contributions to the stablecoin literature:

- Quantifies the structural break in peg stability following the 2022–2023 crisis cascade, showing a 10× improvement in mean-reversion speed (half-life from 16.3h to 1.6h, pooled panel).
- Documents asymmetric arbitrage frictions across the distribution, with half-lives ranging from 2.3 hours at the 5th percentile to 53.3 hours at the 90th percentile (the 95th percentile is non-stationary, indicating permanent deviations in the most extreme tail).
- Finds a 7× asymmetry in mean-reversion speed: discount deviations (price below peg) resolve in approximately 2.3 hours versus approximately 16.2 hours for premium deviations, consistent with stronger arbitrage incentives on the downside.
- Provides the first comprehensive empirical assessment of systemic spillovers among major stablecoins using high-frequency on-chain data, finding that deviations remain predominantly coin-specific (2.00% total spillover index).
- Demonstrates that on-chain issuance and exchange-flow metrics have modest but statistically significant predictive power for peg deviations, supporting arguments for enhanced real-time disclosure in regulatory frameworks.

2. Literature review

2.1. Stablecoin Designs and Peg Mechanisms

The foundational literature categorizes stablecoins into three broad design types- fiat-backed, crypto-collateralized, and algorithmic. Each type embeds fundamentally different stabilization mechanisms. Fiat-backed stablecoins such as USDT and USDC maintain stability through off-chain reserves, typically consisting of cash, Treasuries, or short-term debt. Price stability is achieved through redeemability: users must believe that each token can be reliably exchanged for one U.S. dollar. The credibility of such claims depends on reserve quality and transparency (Berentsen & Schär, 2019; Jarno & Kołodziejczyk, 2021). Empirical evidence suggests that fiat-backed stablecoins with higher-quality collateral and more transparent attestations experience smaller deviations from parity.

Crypto-collateralized stablecoins such as Dai (DAI) rely on over-collateralized debt positions managed through decentralized smart contracts. The underlying mechanism adjusts supply based on collateral values, liquidation incentives, and governance decisions (De Sclavis et al., 2025; Huo et al., 2022). Although such systems avoid reliance on traditional custodians, they remain sensitive to crypto-market volatility and governance delays.

Algorithmic stablecoins rely on rules-based monetary policies rather than explicit collateral. Cerezo Sánchez (2019) provides early theoretical foundations for AI-governed stabilization rules, although empirical experience shows that algorithmic stabilization is highly fragile. TerraUSD's collapse demonstrated the reflexivity inherent in such systems: once confidence weakens, endogenous supply rules accelerate divergence (Duan & Urquhart, 2023). The collapse of TerraUSD has been seen as a defining event illustrating the limits of algorithmic supply adjustments unsupported by credible reserves.

The literature on arbitrage provides a theoretical backbone for understanding peg maintenance. Pernice (2021) demonstrates that stablecoin price processes are shaped by arbitrage bounds; as long as arbitrageurs can profitably redeem stablecoins at par, prices should not deviate significantly from the peg. However, transaction costs, delayed settlement, liquidity constraints, or insufficient collateral may prevent arbitrage from restoring parity quickly. Potter et al. (2021) builds on this idea by describing the "stablecoin trilemma": it is difficult for any stablecoin to achieve stability, decentralization, and scalability simultaneously. Each design sacrifices at least one dimension.

More recent research explores hybrid or next-generation stabilization mechanisms. Han et al. (2024) propose JANUS, a dual-token, multi-collateral model that incorporates AI-driven dynamic stabilization. Ling et al. (2025) develop a "Stablecoin LEGO" systemization framework to decompose stablecoin designs into modular components- collateral, governance, liquidation, oracle, and redemption modules, which interact to determine empirical price behavior. Automated market maker (AMM) models introduce yet another layer: AMM-based pegged assets can maintain quasi-stability through liquidity pool mechanics, but liquidity shortages or pool imbalance can trigger significant deviations (Bergault et al., 2024).

Theoretical and empirical work converge on three pillars of stablecoin price stability: credible redeemability, efficient arbitrage, and robust design. Redeemability at par anchors expectations but depends on reserve quality and disclosure, especially for fiat-backed coins (Berentsen & Schär, 2019; Jarno & Kołodziejczyk, 2021). Arbitrage bounds explain why, in normal conditions, prices stay close to the peg, yet also why frictions – transaction costs, settlement lags, and balance-sheet constraints – allow temporary deviations (Pernice, 2021; Potter et al., 2021). Design frameworks such as JANUS and the "Stablecoin LEGO" emphasise how collateral, governance, oracles, and redemption rules interact to determine which of these pillars fails first under stress (Han et al., 2024; Ling et al., 2025; Bergault et al., 2024).

2.2. Market Microstructure Theory, Arbitrage Limits, and Automated Market Makers

The arbitrage mechanism that Glosten and Milgrom (1985) and Kyle (1985) place at the centre of traditional market microstructure theory does not disappear in decentralised finance; it is instead encoded in protocol logic. Traditional market makers quote bid and ask prices continuously, adjusting their inventory in response to order flow and earning the spread as compensation for absorbing adverse selection. Prices return to fundamental values not because a central authority intervenes, but because arbitrageurs exploit any deviation between the market maker's quoted price and the price available elsewhere. This mechanism, in its essential form, is preserved in automated market makers.

The theoretical expectation that arbitrage ensures rapid peg restoration rests on the assumption that arbitrage is frictionless. In practice, Shleifer and Vishny (1997) demonstrate that arbitrage is constrained by capital requirements, funding risk, and horizon uncertainty; even when a mispricing is identified, the cost and risk of exploiting it may deter arbitrageurs. These limits to arbitrage are directly relevant to stablecoins. For fiat-collateralised stablecoins such as USDT and USDC, redemption at par requires

institutional access to the issuer's redemption facility, which is typically restricted to whitelisted counterparties with minimum transaction sizes. Retail arbitrageurs must instead trade through secondary market venues, incurring transaction costs and price impact. Observed half-lives therefore encode not just the speed of frictionless arbitrage but the combined effect of transaction costs, counterparty access constraints, and funding availability.

An automated market maker (AMM) is a smart contract that holds reserves of two or more assets and quotes exchange rates according to a deterministic formula. Uniswap v2, the canonical AMM architecture, governs prices through the constant product formula $x \cdot y = k$, where x and y are the reserve quantities of two assets and k is a constant (Adams et al., 2021). Any transaction that moves the ratio of reserves moves the quoted price, and arbitrageurs are incentivised to trade against any divergence between the AMM's quoted price and the price on external markets. The mechanism is identical in spirit to the inventory management of a traditional market maker: deviations are corrected by profit-seeking agents, and the protocol's design determines how quickly and at what cost (Angeris et al., 2021). Curve Finance introduces the stableswap invariant, a formula specifically engineered for assets expected to trade near parity (Egorov, 2019). By concentrating liquidity in the region where prices are close to one, the stableswap invariant reduces slippage for near-peg trades and thereby lowers the cost of arbitrage precisely where stablecoin deviations are most likely to occur. For the three stablecoins examined in this thesis, Curve pools constitute a primary venue for peg arbitrage, making the stableswap architecture directly relevant to the dynamics studied.

DAI's Peg Stability Module (PSM) represents a further institutional parallel. Introduced by MakerDAO in 2020, the PSM permits any user to exchange DAI for USDC at a fixed one-to-one rate, subject to a small fee (MakerDAO, 2020). This arrangement is structurally analogous to a currency board: the issuer commits to redeeming the liability at a fixed price, removing residual uncertainty about the peg. Unlike USDT and USDC, whose pegs depend on arbitrageurs trusting off-chain reserve claims and executing trades across venues, DAI's peg through the PSM is enforced on-chain and is unconditional within the module's capacity. This architectural difference generates a testable prediction: DAI's mean reversion dynamics should be governed by a different set of forces than those of the fiat-collateralised stablecoins.

The depth and activity of AMM liquidity pools constitute observable proxies for arbitrage capacity. A pool with deep reserves can absorb larger deviations without moving the quoted price substantially, while high swap volume reflects active arbitrage. These on-chain characteristics motivate the inclusion of liquidity and flow variables in

the panel regression framework of Research Question 2: they are not merely descriptive indicators but structural determinants of the speed at which peg deviations are corrected.

2.3. Run Dynamics and Systematic Fragility

Stablecoin instability closely resembles classic financial runs. Gorton et al. (2025) conceptualize stablecoins as a form of private money: their value depends on confidence in par redemption. During adverse events such as reserve uncertainty or market-wide stress users have incentives to redeem immediately or flee to safer alternatives. This behavior mirrors the dynamics of bank runs, where fear of insolvency becomes self-fulfilling. The literature shows that even fundamentally solvent stablecoins may experience significant runs if users doubt the reliability or speed of redemption.

Stablecoins are extremely sensitive to information about reserves and issuer operations. Ahmed et al. (2024) provide evidence that public information releases, including routine updates, can trigger substantial redemptions. The authors show that, paradoxically, transparency can amplify run risk when disclosures reduce, even marginally, the perceived safety of backing assets. This finding is significant because fiat-backed stablecoins have increasingly large positions in short-dated U.S. Treasuries (Ahmed et al., 2025), creating important spillovers into traditional money markets.

Macro-financial variables also shape instability. Lee et al. (2025) show that de-pegging risk increases during periods of high volatility, deteriorating sentiment, or weak liquidity. Eichengreen et al. (2025) find structural patterns of devaluation driven by macroeconomic conditions, suggesting that stablecoin stability cannot be understood in isolation from broader markets. Abraham (2024) demonstrates that when stablecoins back lending or collateralized positions, feedback loops, such as liquidation cascades, can magnify deviations.

Ma et al. (2025) identify the centralization of arbitrage capacity as a key systemic risk: a small number of entities are responsible for restoring parity across multiple stablecoins. When these agents face funding shortages, constraints, or operational disruptions, arbitrage becomes ineffective, allowing deviations to persist.

Stablecoin instability often exhibits contagion patterns. Oxenhorn (2022) model de-pegging as a multivariate Hawkes process and find that de-peg events cluster strongly across time and across stablecoins. This supports Naifar's (2025) finding that stablecoins are embedded within a broader DeFi network characterized by systemic tail risk.

Strohmeyer, N., Vishwanath, S., & Fridovich-Keil, D. (2024) analyze dynamic games in DeFi environments and show that strategic user interactions can support or destabilize stablecoin systems depending on protocol design. Concentrated holdings further exacerbate tail risk (Naifar, 2025). Al-Afeef et al. (2024) suggest that high usage and transaction velocity may mitigate run pressure by anchoring demand for transactional purposes.

Together, these studies deepen understanding of run dynamics and directly motivate RQ2 and RQ3, which examine whether on-chain issuance, flows, velocity and adoption metrics can predict next-day peg deviations and how much of the major depeg events they can explain.

Despite broad agreement that confidence and liquidity are central, the literature disagrees on which mechanisms are most fundamental for stability. Some authors emphasise arbitrage capacity and its concentration, arguing that deviations persist primarily when key intermediaries are constrained (Pernice, 2021; Ma et al., 2025). Others highlight governance, collateral design and protocol rules – especially for decentralized and algorithmic stablecoins – as the main fault lines that determine whether runs can escalate (Cerezo Sánchez, 2019; Huo et al., 2022; Ling et al., 2025). A third strand stresses adoption and usage: high transactional velocity and broad user bases may either anchor demand or, if holdings are concentrated in “whales”, amplify systemic tail risk (Al-Afeef et al., 2024; Naifar, 2025; Strohmeyer et al., 2024). This thesis speaks to these debates by testing whether microstructure variables associated with arbitrage, governance-driven supply changes, and adoption dynamics actually predict peg deviations across different designs.

2.4. On-Chain Microstructure: Issuance, Flows and Adoption

Recent studies highlight the importance of blockchain microstructure as opposed to purely off-chain fundamentals in shaping peg stability. Stablecoin supply is dynamically adjusted through mint and burn transactions. Kwon et al. (2023) show that net issuance, particularly large redemption waves, predicts short-term de-pegs across major stablecoins. Huo et al. (2022) similarly emphasize that governance-driven supply mechanisms, especially in decentralized systems like DAI, significantly affect resilience under stress.

Exchange flows provide near-real-time indicators of run pressure. Joshi (2025) notes that spikes in exchange outflows precede negative peg deviations, as traders move stablecoins into self-custody or swap them for other assets. Somin et al. (2025) contribute a high-resolution Ethereum dataset enabling granular tracking of flows,

enabling researchers to distinguish retail flows from institutional ones using address clustering.

Adoption dynamics further influence stability. Greater wallet activity, particularly returning users, may reflect transactional demand that stabilizes the peg. Conversely, high concentration of holdings among “whales” can worsen run dynamics. Naifar (2025) finds that stablecoins exhibit significant tail risk exposure due to concentrated liquidity pockets, linking on-chain concentration to systemic fragility. Strohmeyer, N., Vishwanath, S., & Fridovich-Keil, D. (2024) further show that strategic behavior in DeFi environments, including dynamic game interactions, affects the stability of collateralized stablecoin ecosystems.

Velocity (the ratio of on-chain transfer volume to circulating supply) is another stabilizing factor cited in the literature. Higher velocity indicates that a stablecoin is used actively as a medium of exchange, potentially reducing the persistence of deviations. Al-Afeef et al. (2024) argue that stablecoins with high transactional use provide hedging benefits for crypto traders, indirectly supporting their peg during moderate stress.

Methodologically, the emphasis of on-chain microstructure naturally motivates the use of high-frequency panel regressions, which capture persistence and cross-sectional heterogeneity in deviations. What remains scarce, however, are applications that use granular blockchain data across multiple major stablecoins over several years. By doing so, this thesis extends panel approaches from traditional assets to the specific context of stablecoin microstructure.

2.5. Empirical Evidence on Peg Stability

The empirical literature on stablecoin stability uses event studies, panel regressions, de-pegging models, and systemic risk mapping. Event studies, based on MacKinlay (1997), help identify abnormal deviations during crisis episodes. Duan & Urquhart (2023) show that algorithmic stablecoins experience severe deviations, while fiat-backed coins deviate less but still significantly during stress.

Time-series models such as those developed by Djogbenou et al. (2023) reveal persistent but time-varying deviations in Tether’s price. Eichengreen et al. (2025) show that devaluation risk includes both idiosyncratic and systematic components. Cross-sectional literature highlights structural influences. Jarno & Kołodziejczyk (2021) show that collateral and transparency shape price stability. Ling et al. (2025) demonstrate how design modules combine to determine empirical outcomes. Ahmed et al. (2025) show that stablecoin reserve management influences Treasury markets, exposing systemic connections.

Ante et al.(2023) provide a comprehensive literature review and highlight fragmentation: empirical research remains narrow in scope, often focusing on single stablecoins, isolated de-pegging events, or price-based predictors while overlooking microstructure variables. In empirical applications, arbitrage activity is not directly observable and is typically proxied through market and on-chain variables such as issuance and redemption flows, exchange inflows, and trading volume. These variables reflect the intensity and constraints of arbitrage, allowing to assess how effectively deviations from the peg are corrected in practice.

Although the empirical literature has grown rapidly, several important gaps remain. First, existing studies largely focus on single stablecoins or isolated episodes, leaving limited understanding of how multiple major stablecoins behave relative to one another across longer time horizons. Second, despite the theoretical emphasis on microstructure variables such as issuance, exchange flows, and adoption metrics, empirical work seldom incorporates these on-chain measures into systematic analyses of peg stability. Third, predictive modelling remains relatively underdeveloped: only a few studies examine whether micro-level on-chain behavior anticipates next-day peg deviations, even though real-time blockchain data offers unique opportunities for forward-looking analysis. Fourth, event studies and panel regression approaches tend to be used in isolation rather than in combination, resulting in fragmented insights into both short-term stress dynamics and more persistent deviations. Fifth, although velocity is frequently invoked as a theoretical stabilizing force linked to transactional demand, empirical evidence on its relationship to stability is scarce. Finally, while several papers document clustering of de-pegging events, systematic analysis of cross-stablecoin contagion and co-instability remains limited, especially in studies that span multiple designs and multiple years.

Taken together, existing studies suggest three main explanatory blocks for peg stability: (i) arbitrage and balance-sheet frictions, (ii) design and governance features, and (iii) on-chain usage and adoption patterns. RQ1 focuses on descriptive stability metrics across designs, RQ2 operationalizes the arbitrage and microstructure channels using issuance, flows, and velocity, and RQ3 combines event-study methods with these variables to study how shocks propagate across coins. This mapping provides a simple conceptual framework that links the main strands of the literature to the empirical design developed in Sections 2 and 3.

3. Methodology

3.1. Overview

The empirical analysis models hourly peg dynamics for USDT, USDC and DAI using a reduced-form autoregressive panel with on-chain and market covariates. The goal is to estimate how lagged deviations, blockchain activity and exchange conditions relate to short-run deviations from the one-dollar peg, and how these dynamics change across regimes and tails of the distribution. The econometric strategy has four stages: (i) panel fixed-effects regression with Driscoll–Kraay standard errors, (ii) structural break analysis around the 2022 crypto crisis, (iii) panel quantile regression to model tail behaviour, and (iv) VAR-based spillover analysis with Granger causality, FEVD and Diebold–Yilmaz spillover indices.

3.2. Data Sources

Hourly prices and trading volumes are obtained from CryptoCompare, which provides OHLC data aggregated across major exchanges and is widely used in high-frequency stablecoin research. On-chain data come from Ethereum BigQuery public tables: all ERC-20 transfer, mint and burn events for the USDT, USDC and DAI contract addresses are extracted following Somin et al. (2025). Only verified contract addresses are used. Matching by hourly timestamps produces a balanced panel of three stablecoins observed at 51,634 hourly points over 2020–2025.

3.3. Variable Construction

3.3.1. Dependent Variable: Peg Deviation

The main dependent variable is the absolute hourly deviation from the one-dollar peg:

$$DEV_{i,t} = | 1 - p_{i,t} |,$$

where $p_{i,t}$ is the hourly closing price of stablecoin i. This follows the arbitrage bounds view in which deviations above and below one are symmetric mispricings.

3.3.2. On-Chain and Market Covariates

$$\boldsymbol{X}_{it} = [\text{Returns}_{it}, \ln\text{TradingVol}_{it}, \sigma_{i,t}, \text{Velocity}_{it}, \text{NetIssuance}_{it}, \text{WalletRatio}_{it}, \ln\text{TransferVol}_{it}]'$$

The covariate vector $X_{i,t}$ includes on-chain measures capturing activity, supply and participation:

- Transfer counts and volumes.

- Mint and burn volumes.
- Unique sending and receiving wallets.
- Circulating supply.

These variables reflect settlement intensity, net issuance and breadth of usage, which prior work links to stability, run risk and resilience.

Market-side variables include:

- Trading volume.
- Simple returns.
- Log returns,

$$r_{i,t}^{\log} = \ln(p_{i,t}) - \ln(p_{i,t-1}).$$

To capture volatility clustering and "limits to arbitrage", a 24-hour rolling volatility measure $\sigma_{i,t}$ is constructed from hourly returns, summarising local uncertainty that affects arbitrage costs and risk-bearing capacity. All variables are aggregated to hourly frequency.

Variable Definitions and Theoretical Justification

Variable	Construction	Unit	Theoretical Role
$\lvert DEV_{i,t}\rvert$	$\lvert P_{i,t} - 1\rvert$	$	Outcome: absolute peg deviation
ρ (lag)	$\lvert DEV_{i,t-1}\rvert$	$	AR(1) persistence (mean-reversion)
Net issuance	$\ln(\text{mint}_{i,t}) - \ln(\text{burn}_{i,t})$	log-ratio	Arbitrage-driven supply response
Velocity	transfer vol/supply$_{i,t}$	ratio	Transactional demand
Exchange inflow	$\Delta\ln(\text{exchange wallet balance}_{i,t})$	log-diff	Redemption / selling pressure

Variable	Construction	Unit	Theoretical Role
Volatility	Realized standard deviation of hourly returns (24h rolling)	%	Market stress
Volume	$\ln(\text{trading volume}_{i,t})$	log-USD	Liquidity

Note. All on-chain variables are sourced from Google BigQuery ERC-20 public tables (Somin et al., 2025). Prices from CryptoCompare.

Each on-chain variable is expected to relate to peg deviations through a specific economic mechanism. Net issuance (positive when tokens are minted, negative when redeemed) increases circulating supply and is expected to exert below-peg pressure when positive ($\beta < 0$ for the issuance coefficient): excess supply creates downward price pressure that arbitrageurs must correct. Token velocity, measured as transaction volume normalised by circulating supply, proxies for the intensity of on-chain transfer activity; higher velocity may reflect arbitrage trades correcting deviations, suggesting a negative coefficient (faster correction when activity is high). Exchange inflows measure the volume of tokens transferred into exchange wallets; large exchange inflows may precede redemption or sale pressure and are expected to be associated with negative deviations ($\beta < 0$). Wallet concentration (Gini coefficient of the top-decile wallet share) captures counterparty concentration risk; a more concentrated token supply increases the potential for coordinated redemption runs and is expected to be positively associated with deviation magnitudes.

These directional predictions follow directly from the limits-to-arbitrage and market microstructure frameworks outlined in Section 2.2. In Section 4.2, we assess whether the data confirm these hypotheses and discuss cases where sign reversals suggest mechanisms beyond those predicted by the baseline theory.

3.4. Data Cleaning

Data cleaning follows Ahmed et al. (2025) and standard high-frequency crypto practices. Zero-value candles, duplicated rows, stale prices and malformed entries are removed. Hours with valid prices but zero on-chain activity are retained so that low-activity periods are not systematically excluded. Short gaps are forward-filled; longer

gaps are removed, yielding a clean, balanced panel suitable for time-series and panel analysis.

3.5. Econometric Specification

Section 3.5 describes six econometric methods applied in sequence, each addressing a distinct aspect of the three research questions. To address Research Question 1 - whether peg deviations exhibit persistent mean reversion and how persistence varies across coins and over time - we apply an AR(1) panel fixed-effects model (Section 3.5.1) and structural break analysis around the May 2022 Terra/Luna collapse (Section 3.5.3). Directional asymmetry in mean reversion, a subsidiary question within RQ1, is addressed through a signed deviation analysis (Section 3.5.4a). To address Research Question 2 - whether on-chain microstructure variables predict deviations beyond autoregressive dynamics - on-chain covariates are incorporated into the baseline panel specification and their marginal contribution is assessed at different points of the deviation distribution using quantile regression (Section 3.5.4). To address Research Question 3 - how peg deviations evolve around major stress episodes and whether shocks propagate across coins - we employ a VAR model with Granger causality tests and forecast error variance decomposition (Section 3.5.6 and 3.5.7), the Diebold-Yilmaz spillover index (Section 3.5.8).

3.5.1. Panel Structure and Fixed Effects

To test whether peg deviations are persistent mean-reverting processes driven by arbitrage constraints that differ across stablecoin designs, we apply an AR(1) panel fixed-effects model with coin-specific intercepts and Driscoll–Kraay standard errors that accommodate cross-sectional and temporal dependence.

The baseline specification is an AR(1) panel model with stablecoin fixed effects:

$$DEV_{i,t} = \alpha_i + \rho\, DEV_{i,t-1} + \beta' X_{i,t} + u_{i,t}.$$

Keller (2024) employs bank × time panels with multiple layers of fixed effects to study covered interest parity (CIP) deviations and their impact on lending behavior, demonstrating that panel structures allow identification of both time-invariant entity characteristics and time-varying arbitrage dynamics. Here, the panel (stablecoin × time) increases statistical power compared with separate time-series ($N = 3$ coins and $T \approx 51{,}634$ hours, the panel provides roughly 155,000 observations compared to ~51,600 per individual time series), which is important for precisely estimating persistence and half-lives over a six-year window. It also allows common responses to volatility and on-chain activity to be estimated while controlling for coin-specific design differences.

Coin fixed effects α_i are included because DAI, USDC and USDT differ fundamentally in stabilisation mechanisms: USDC and USDT are centrally issued, fiat-backed with bank-based redemption, whereas DAI is a decentralised, over-collateralised DeFi system with smart-contract and governance risk. These differences show up as persistent level shifts in typical deviations, volatility and arbitrage speed are exactly what entity fixed effects are designed to absorb; otherwise they would bias ρ and β.

Time fixed effects λ_t are deliberately omitted. Time dummies would absorb all common time-series variation, including the volatility and volume dynamics that are key regressors, and FEVD results show that over 96% of short-run variance in each coin's deviations is idiosyncratic. They would also soak up the regime change around the 2022 crisis, making structural-break tests on persistence uninformative. Banti & Phylaktis (2015) and Krohn & Sushko (2022) face similar trade-offs in FX liquidity panels and explicitly choose not to include time FE when the goal is to estimate how time-varying aggregate risk (funding stress, volatility) affects spreads.

3.5.2. Driscoll–Kraay Standard Errors

Hourly stablecoin data exhibit (i) heteroskedasticity (volatility clustering), (ii) serial correlation and (iii) cross-sectional dependence due to shared exchanges, common arbitrageurs and joint DeFi liquidity pools such as Curve's 3pool. One-way entity clustering addresses within-coin dependence but not cross-coin correlation; two-way clustering requires large N and T; Newey–West HAC is designed for a single series. The Driscoll–Kraay (1998) estimator allows for heteroskedasticity, autocorrelation up to lag L (set to $\lfloor 4(T/100)^{\frac{2}{9}} \rfloor$ following the rule-of-thumb for HAC estimators, yielding 15 lags) and arbitrary cross-sectional dependence when T is large, a condition satisfied here. Banti & Phylaktis (2015) use Driscoll–Kraay SEs in panel regressions of FX bid-ask spreads on funding constraints, noting that currency pairs are linked through dealer balance sheets and arbitrage flows.

3.5.3. Structural Break Analysis Around May 2022

To test whether the 2022 crypto crisis materially altered the speed of arbitrage-driven mean reversion - and whether on-chain variables gained or lost predictive importance - we split the sample at 12 May 2022 and estimate the panel model separately on pre- and post-crisis subsamples with a formal test for coefficient shifts.

To examine whether arbitrage dynamics changed after the 2022 crypto crisis, the sample is split at 12 May 2022 (Terra/Luna collapse) and the panel model is

estimated separately on pre- and post-crisis subsamples. This mirrors CIP studies that split around major crises (2008 GFC, Eurozone crisis, Covid-19) to test whether deviations became more or less persistent as regulation and balance sheet constraints evolved. This reasoning follows the approach in Keller (2024), who similarly omits time FE when testing how CIP deviations change over regimes defined by crisis episodes.

Terra's failure, subsequent stress at Celsius, Voyager and Three Arrows Capital, and the FTX bankruptcy significantly reshaped stablecoin markets by eliminating major algorithmic designs, increasing regulatory scrutiny and transparency, and accelerating the development of sophisticated arbitrage infrastructure. Comparing pre- and post-2022 estimates of ρ and key β's shows whether arbitrage became faster and whether the influence of volatility and on-chain activity changed. As an alternative, a pooled model with a crisis dummy POST and interaction term is estimated:

$$DEV_{i,t} = \alpha_i + \rho_1 DEV_{i,t-1} + \rho_2\left(POST \times DEV_{i,t-1}\right) + \beta' X_{i,t} + u_{i,t},$$

with a Wald test on $H_0: \rho_2 = 0$ providing a formal test for a shift in persistence.

3.5.4. Quantile Regression: Tail Behaviour and Limits to Arbitrage

To test whether peg stability varies asymmetrically across the distribution of deviations - for example, whether small deviations are highly transient while extreme deviations exhibit breakdown in arbitrage mechanisms - we re-estimate the model at five quantiles of the conditional deviation distribution using panel quantile regression with block bootstrap inference.

Crypto returns and stablecoin deviations are heavy-tailed and non-normal, so conditional mean models can miss important tail behaviour. Standard least squares estimate conditional means, but in the presence of heavy tails and asymmetric arbitrage costs, quantile regression provides a more complete picture of the full conditional distribution. To capture asymmetric arbitrage dynamics across the distribution of deviations, the baseline model is re-estimated using panel quantile regression at five quantiles:

$$\tau \in \{0.05,\ 0.25,\ 0.50,\ 0.75,\ 0.95\},$$

following Baruník & Křehlík (2018).

The 5th percentile corresponds to small, efficiently arbitraged mispricings; the median reflects typical states; the 95th percentile represents severe de-pegs where funding constraints, capital limits and risk aversion are likely to bind, slowing arbitrage. The focus is on how the persistence coefficient $\rho(\tau)$ and key slopes $\beta(\tau)$ evolve across quantiles. A rising $\rho(\tau)$ toward the upper tail would be direct evidence that large

deviations are more persistent than small ones, consistent with “limits to arbitrage” theories and results showing that systemic risk factors concentrate in the tails.

Because the quantile regression objective is non-differentiable, inference is based on 200 bootstrap replications using a block bootstrap with 24-hour blocks, which preserves the serial correlation structure of the hourly data within each day, following Koenker (2005) and financial applications of quantile regression.

3.5.5. Directional Asymmetry Analysis

To assess directional asymmetry in mean reversion, the baseline panel specification is re-estimated using signed deviations that retain the direction of peg deviations. The signed deviation is defined as $\text{SDEV}_{i,t} = p_{i,t} - 1$, where positive values indicate premium deviations (price above $1.00) and negative values indicate discount deviations (price below $1.00). This approach is motivated by the asymmetric arbitrage mechanics inherent to stablecoin design: when prices fall below par, arbitrageurs can immediately purchase discounted tokens and redeem them at par with the issuer, creating a mechanical price floor and strong downside correction incentives. Conversely, when prices rise above $1.00, parity restoration requires the issuer to mint and distribute new supply into secondary markets, a process constrained by operational capacity and governance arrangements that limits upside correction speed. Following limits-to-arbitrage theory (Shleifer & Vishny, 1997), we expect these structural differences to manifest as faster mean reversion for discount deviations relative to premium deviations.

The panel AR(1) model is estimated separately on the subset of observations where $\text{SDEV}_{i,t-1} < 0$ (discount regime) and where $\text{SDEV}_{i,t-1} > 0$ (premium regime). The persistence coefficient $\hat{\rho}_d$ from the discount subsample and $\hat{\rho}_p$ from the premium subsample are compared using implied half-lives computed as $\text{HL} = -\ln(2)/\ln(\hat{\rho})$. A finding of $\text{HL}_d < \text{HL}_p$ confirms that downward deviations correct faster, consistent with the structural advantage of direct issuer redemption over secondary-market supply injection.

3.5.6. Cointegration and VAR for Long- and Short-Run Dynamics

All three stablecoins target a 1:1 USD peg, suggesting a common long-run equilibrium enforced by arbitrage, analogous to covered interest parity or exchange rate pegs. Johansen (1991) cointegration tests are applied to the vector of deviations to determine the rank of the cointegration space, allowing for multiple cointegration vectors and endogenous feedback. A full cointegration rank would indicate that each deviation is stationary yet shares common long-run dynamics linked by no-arbitrage.

Banti & Phylaktis (2015) use VAR-based models to study how funding liquidity shocks in one currency pair spill over to others.

Given stationarity and cointegration, a VAR in levels is estimated:

$$Y_t = C + \sum_{k=1}^{p} A_k Y_{t-k} + U_t,$$

where Y_t contains deviations of DAI, USDC and USDT. Each equation includes lagged values of all three coins, allowing for feedback and lead–lag effects in peg adjustment, such as USDT shocks propagating via shared liquidity pools or USDC leading price discovery due to higher liquidity. The lag length p is selected using the Hannan–Quinn Information Criterion (HQIC), which applies a stronger parsimony penalty than the Akaike Information Criterion, reducing the risk of overfitting. HQIC reaches its minimum at lag 24 (see Appendix Table A.2), corresponding roughly to observed half-life ranges.

3.5.7. Granger Causality and FEVD

To test whether shocks to one stablecoin's peg deviation lead or lag shocks to others - identifying which coin drives price discovery - and whether such causal links translate into economically meaningful spillovers or remain statistical artifacts, we estimate a VAR(24) model in deviation levels and apply block-wise Granger causality tests and forecast error variance decomposition.

Within the VAR, Granger causality tests (block Wald tests on lag coefficients) are used to ask whether past values of coin A improve forecasts of coin B beyond B's own history. This is purely predictive, not structural causality, but it identifies which stablecoin tends to move first in terms of peg deviations.

To assess economic, rather than just statistical, importance, forecast error variance decomposition (FEVD) is applied. FEVD breaks each coin's H-step-ahead forecast error variance into components attributable to its own shocks and to shocks in other coins. Results showing that 96–99% of variance is own-driven imply that, despite statistically significant Granger causality, spillovers are small in magnitude and arbitrage remains fragmented across coins rather than fully integrated. Krohn & Sushko (2022) uses FEVD to show that while FX spot and swap markets are statistically linked (Granger-causal), the quantitative spillovers are modest (own-variance dominates).

3.5.8. Diebold–Yilmaz Spillover Index

To test whether peg risk concentrates within individual coins or propagates across the stablecoin complex as contagion, we compute the Diebold–Yilmaz spillover

index aggregating directional information flows and assess the magnitude of shock transmission relative to idiosyncratic variation.

The Diebold and Yilmaz (2012) spillover index aggregates FEVDs into a compact set of contagion measures. It reports:

- Total spillover: the share of variance due to cross-coin shocks (off-diagonal FEVD elements).
- Directional spillovers “To” and “From” each coin.
- Net spillovers (To – From), classifying coins as net transmitters or receivers of shocks.

  It is written as:

$$S(H) = \frac{\sum_{i=1}^{N}\sum_{j=1,\, j\neq i}^{N} \tilde{\theta}_{ij}(H)}{\sum_{i=1}^{N}\sum_{j=1}^{N} \tilde{\theta}_{ij}(H)} \times 100$$

where $\tilde{\theta}_{ij}(H)$ represents the contribution of variable j to the forecast error variance of variable i at horizon H.

Here, the index is applied to absolute peg deviations rather than returns or volatility, so spillovers are interpreted as contagion in inefficiency: a de-peg in one coin raising deviations in others indicates that confidence shocks or arbitrage frictions propagate across the stablecoin complex. Recent evidence on stablecoin runs and co-instability supports this interpretation, and the finding that USDT behaves as a net receiver rather than a transmitter challenges common narratives about Tether as the primary systemic risk source.

3.6. Interpretation

In the panel fixed-effects model, ρ summarises the average persistence of peg deviations across DAI, USDC and USDT; the implied half-life is

$$\text{Half-life} = \frac{\ln(0.5)}{\ln(\rho)}.$$

Lower ρ corresponds to faster arbitrage and shorter-lived deviations. Where $\rho_1 < 1$ ensures stationarity and the half-life measures the time required for deviations to decay by 50% The slope vector β shows how on-chain activity (transfers, issuance, wallet counts) and market conditions (volume, returns, volatility) correlate with stability, averaged across coins after controlling for fixed design features.

Comparing ρ and β across pre- and post-2022 subsamples reveals whether the crisis strengthened or weakened arbitrage and re-weighted the importance of volatility and flows. Quantile regression estimates $\rho(\tau)$ indicate whether large deviations are

more persistent than small ones; a rising $\rho(\tau)$ toward the upper tail is interpreted as evidence that limits to arbitrage bind most strongly during extreme de-pegs.

In the VAR, significant Granger causality identifies leading coins in the adjustment process, while FEVD and the Diebold–Yilmaz indices indicate whether those links translate into economically meaningful spillovers or remain mostly idiosyncratic.

3.7. Robustness Checks

Robustness checks include re-estimating the panel with percentage deviations instead of absolute deviations,

$$devpct_{i,t} = | 1 - p_{i,t} | \times 100,$$

testing alternative clipping thresholds for prices, and estimating reduced models without supply or without market variables to check coefficient stability. AR(2) and AR(3) specifications are also estimated, and entity- and two-way-clustered standard errors are compared against Driscoll–Kraay to assess the sensitivity of inference. Finally, the crisis interaction specification with POST × $DEV_{i,t-1}$ is used as an alternative to the sample split for structural-break testing.

3.8. Limitations

Five limitations constrain the generalizability of the results:

Endogeneity. Endogeneity exists between on-chain variables (issuance flows, token velocity, exchange activity) and peg deviations. Large issuance or redemption flows may both respond to peg deviations - if market participants react to instability - and cause deviations, if issuance disrupts liquidity. This thesis employs lagged values of on-chain covariates to mitigate reverse causality, reducing but not eliminating the endogeneity concern. Full resolution would require an instrumental variable approach identifying exogenous shifters of issuance or exchange flows independent of peg pressures. Such an approach is beyond the scope of this thesis and left for future work. Readers should interpret on-chain variable coefficients as associations conditional on the reduced-form specification rather than causal effects.

Symmetry assumption. The baseline panel model treats positive and negative deviations symmetrically by using absolute values. The quantile regression results in Section and the directional asymmetry analysis in Section 3.5.5 suggest that this masks important asymmetries; however, a fully asymmetric structural model lies outside the scope of this thesis.

Look-ahead bias in event selection. The analysis focuses on three main de-peg events: Terra/Luna collapse in May 2022, FTX bankruptcy in November 2022, and Silicon Valley Bank closure in March 2023. These events were selected ex-post based on observed severity rather than a pre-registered statistical threshold. Although a 0.20% hourly deviation threshold combined with manual news verification guided selection, the three chosen events represent the most dramatic episodes in the sample and are not necessarily representative of all stress periods. Two further events present in the estimation code - Black Thursday (DAI, March 2020) and the USDT China mining ban (May 2021) - were computed but excluded from the main analysis: Black Thursday precedes the full hourly on-chain panel and yields an insufficient pre-event window for stable baseline estimation; the China ban episode produced statistically insignificant abnormal deviations relative to baseline volatility and is reported as a robustness check in Appendix B.2. This approach introduces selection bias: results characterise the dynamics of extreme episodes but may not generalise to smaller or more frequent stress scenarios.

Ethereum-only on-chain data for multi-chain stablecoins. On-chain transfer, issuance, and wallet concentration metrics are sourced exclusively from Ethereum BigQuery tables. While this enables precise high-frequency analysis, USDT operates across multiple blockchains (Ethereum, Tron, Solana, Polygon, and others), and Ethereum-only data represent a partial view of total USDT activity. A substantial fraction of USDT exists and trades on Tron and Solana, chains not captured in this analysis. On-chain variable coefficients for USDT therefore reflect arbitrage and issuance dynamics on Ethereum alone and may not be representative of system-wide USDT stability. Results for USDC and DAI, which are primarily Ethereum-native, are less exposed to this limitation. Future work incorporating multi-chain on-chain data would substantially strengthen conclusions regarding USDT. The sample covers 2020–2025 and three stablecoins; structural changes in DeFi infrastructure and regulation over this period mean that early-period dynamics may not be representative of current conditions.

Off-chain reserve opacity. For fiat-backed stablecoins, the quality and composition of reserves are disclosed only periodically. The model cannot directly control for reserve changes between attestation dates, meaning that some variation attributed to market variables may partly reflect unobserved reserve dynamics.

4. Analysis of results

4.1. Peg Stability Across USDT, USDC and DAI (RQ1)

4.1.1. Distributional Patterns of Deviations

Table 1 summarizes the distribution of absolute price deviations from the one-dollar peg for DAI, USDC and USDT over the 51,634-hour-per-coin sample (154,902 coin-hours total). USDC appears closest to an “ideal” stablecoin: its mean absolute deviation is only 0.0279 percent, the median deviation is effectively zero, and the maximum deviation is 4.40 percent. Crucially, USDC trades outside the ±0.5 percent band around parity in only 161 hours, corresponding to 0.31 percent of the sample. This indicates that USDC typically behaves like a very tight money-market instrument, with rare and modest departures from par.

Table 1:

Descriptive Statistics of Price Deviations

Coin	N	**Mean**	**Median**	**Std**	**Max**	**Hours** $> 0.5\%$ **(%)**
DAI	51,634	0.2159	0.0300	0.5256	28.90	5,613 (10.87)
USDC	51,634	0.0279	0.0000	0.1731	4.40	161 (0.31)
USDT	51,634	0.0562	0.0200	0.0993	8.90	168 (0.33)

Note: All deviation values in percentage points. “Hours $> 0.5\%$” indicates the number of hours the asset traded outside the [0.995, 1.005] band.

USDT shows somewhat looser but still strong peg performance. Its mean deviation is 0.0562 percent and the median is 0.02 percent, with a maximum deviation of 8.90 percent. USDT trades outside the 0.5 percent band in 168 hours (0.33 percent of the sample), which is only slightly more frequent than USDC. The distributions of both coins are heavily concentrated near zero deviations, with infrequent spikes.

By contrast, DAI exhibits markedly weaker peg control. Its mean absolute deviation is 0.2159 percent – almost four times larger than USDT’s and nearly eight times larger than USDC’s – and the standard deviation is an order of magnitude higher than that of USDC. Most strikingly, DAI spends 5,613 hours (10.87 percent of the sample) more than 0.5 percent away from the peg, indicating that sizeable de-pegs are a regular occurrence rather than rare, idiosyncratic events.

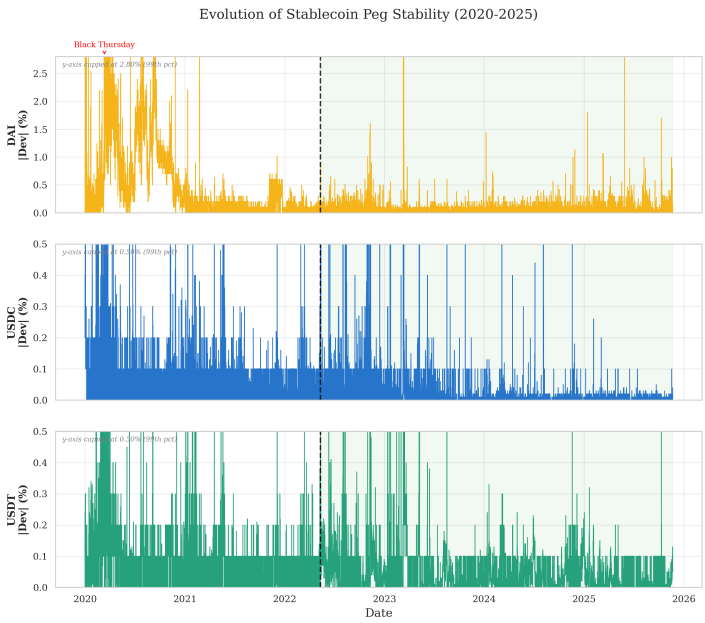

Evolution of Stablecoin Peg Stability (2020–2025)

These distributional differences are clearly visible in the time-series plots in Figure 1. The top panel shows DAI's hourly deviations, with large spikes during early DeFi growth and around crisis episodes, followed by a visible but incomplete calming after mid-2022. The middle and bottom panels for USDC and USDT show much narrower bands of variation, with deviations tightly clustered around zero and only occasional spikes during stress periods. The vertical line marking the Terra/Luna collapse visually separates a more turbulent pre-crisis regime from a relatively calmer post-crisis regime for all three coins.

The descriptive evidence therefore suggests a clear stability hierarchy – USDC $\succ$ USDT $\succ$ DAI – which aligns with theoretical expectations that fiat-backed, transparently collateralized coins should be less volatile than decentralized, crypto-collateralized designs.

4.1.2. Persistence and Half-Life of Deviations

While frequencies and magnitudes of deviations are informative, the speed at which mispricings are corrected is equally important for understanding stability. Table 2 reports autocorrelation coefficients at 1, 6 and 24-hour lags, along with implied half-lives.

Table 2:

Persistence Metrics and Half-Life Estimates

Coin	ACF(1h)	ACF(6h)	ACF(24h)	ACF(168h)	Half-life (hours)
DAI	0.9613	0.9480	0.9184	0.8227	17.58
USDC	0.5279	0.4773	0.4265	0.2873	1.09
USDT	0.7622	0.6710	0.5744	0.3624	2.55
Note: ACF denotes autocorrelation function at specified lag.					

USDC's first-order autocorrelation is 0.5279, falling to 0.4265 at 24 hours, and the corresponding half-life is approximately 1.09 hours. This suggests that, once a deviation occurs, arbitrage forces typically halve it within about an hour and push it close to zero in a few hours. USDT's ACF(1h) is higher at 0.7622, and its half-life is 2.55 hours, indicating slower – but still relatively quick – mean reversion.

DAI's autocorrelation structure is very different. The first-order autocorrelation is 0.9613 and remains high even at 24 hours (0.9184), implying a half-life of approximately 17.6 hours. Deviations in DAI therefore tend to persist across an entire day, which is consistent with its more complex, governance-driven stabilization mechanism and with earlier findings that decentralized collateral systems can be slow to adjust during stress. The persistence patterns in Table 2 reinforce the stability ranking observed in Table 1 and in Figure 1.

Three distinct half-life estimates appear throughout this thesis and should not be used interchangeably. The coin-level estimates in Table 2 - DAI: 17.58h, USDC: 1.09h, USDT: 2.55h - are derived from each stablecoin's own autocorrelation structure over the full sample. The pooled panel estimate of 14.1 hours (Table 4) reflects the average persistence across all three coins jointly in the fixed-effects specification. The pre-crisis pooled estimate of 16.3 hours and post-crisis estimate of 1.6 hours (Table 5) condition on the sub-periods separated at May 2022. Each subsequent reference to a half-life specifies which of these three quantities is meant.

4.1.3. Stationarity and Modelling Implications

Table 3:

Augmented Dickey-Fuller Test Results

Coin	ADF Statistic	p-value	Crit. 5%	Stationary
DAI	−9.711	<0.001	−2.862	Yes
USDC	−24.266	<0.001	−2.862	Yes
USDT	−19.570	<0.001	−2.862	Yes
Note: Series: $\lvert DEV_t \rvert$. Specification: constant only; max lag: 24.				

Table 3 presents Augmented Dickey–Fuller test results for the deviation series. For all three stablecoins, the ADF statistics are large in absolute value (between –9.71 and –24.28) and the associated p-values are effectively zero, leading to rejection of the unit-root null at the 1 percent level. This indicates that deviations are stationary in levels: they fluctuate around a stable mean and do not drift over time, which is consistent with arbitrage preventing permanent divergence from the peg.

This property is important econometrically, as it validates the use of level-based panel regressions, VAR models and cointegration analysis without differencing the deviation data. It is also consistent with the theoretical view of stablecoins as assets that are designed to be mean-reverting around a target value, even if they exhibit short-run volatility and occasional large de-pegs.

Overall, the descriptive analysis answers RQ1 in a first, reduced-form sense: the three stablecoins are all mean-reverting, but they differ substantially in typical deviation size, persistence and the frequency of larger dislocations, with centralized fiat-backed coins displaying much tighter pegs than the decentralized design.

4.2. Determinants and Evolution of Arbitrage Efficiency (RQ2)

4.2.1. Baseline Panel Regression: Average Dynamics

Table 4:

Panel Fixed Effects Regression Results

Dependent Variable: Absolute Price Deviation ($DEV_{i,t}$)

Variable	Coeff.	Std. Err.	t-stat	p-value	Sig.
DEV_{t-1} (Persistence)	0.9521	0.00333	285.97	<0.001	***
RET_{t-1} (Arbitrage)	−0.2003	0.01015	−19.73	<0.001	***
$\ln VOL_{t-1}$ (Volume)	8.31×10^{-6}	5.68×10^{-7}	14.64	<0.001	***

Variable	Coeff.	Std. Err.	t-stat	p-value	Sig.
σ_{24h} (Volatility)	0.01715	0.00763	2.25	0.025	**
VEL_{t-1} (Velocity)	-1.31×10^{-7}	3.63×10^{-8}	−3.62	<0.001	***
ISS_{t-1} (Net Issuance)	0.001135	0.001122	1.01	0.312	
WAL_{t-1} (Wallet Ratio)	-6.16×10^{-5}	3.61×10^{-5}	−1.71	0.087	*
$\ln TXVOL_{t-1}$ (Tx. Volume)	-1.79×10^{-5}	1.43×10^{-6}	−12.48	<0.001	***
$N = 154{,}872 R^2 = 0.9115$ Driscoll–Kraay SE Coin FE: Yes					
$p < 0.1$ * $p < 0.05$ ** $p < 0.01$					

The fixed-effects panel regression in Table 4 provides a first structural view of how past deviations and market conditions shape short-run peg stability. The specification includes coin fixed effects to absorb time-invariant design differences – such as collateral type, issuer structure and governance – and focuses on within-coin dynamics over time.

The persistence coefficient on lagged deviation, 0.9521, is large and precisely estimated (t-statistic 285.97), corresponding to a half-life of approximately 14.1 hours. This suggests that, on average across coins and over the full sample, deviations are corrected only gradually; roughly half of a given mispricing remains after 14 hours. This finding dovetails with the descriptive half-life estimates in Table 2, which showed especially slow mean reversion for DAI and moderate persistence for USDT.

Lagged returns enter with a coefficient of -0.2003 and a t-statistic of -19.73, indicating that price changes tend to move against prior deviations. In practical terms, when a coin trades above one dollar, subsequent returns are more likely to be negative, pushing the price back toward the peg, and vice versa. This is consistent with arbitrageurs systematically trading to exploit mispricings, as emphasized by theoretical work on stablecoin arbitrage bounds.

Trading volume has a small but statistically significant positive coefficient (8.31×10^{-6}). Rather than implying that liquidity worsens stability, this pattern is plausibly driven by reverse causality: large deviations attract volume as arbitrageurs and speculators enter the market, so periods with wider spreads mechanically coincide with higher trading activity. The 24-hour volatility variable also enters positively (0.0172, $p = 0.025$), suggesting that deviations widen when the broader market is more volatile. This aligns with limits-to-arbitrage theories, in which shocks to risk and uncertainty

raise effective arbitrage costs and weaken traders' willingness to hold risky positions needed to close the gap to the peg.

The high R^2 of 0.9115 indicates that the combination of persistence, returns, volume and volatility captures most of the variation in hourly deviations. At the same time, the large persistence coefficient implies that a substantial portion of mispricing dynamics cannot be fully explained by observed contemporaneous covariates but reflects slow adjustment intrinsic to the market structure.

To put these coefficients in economic terms, a one-standard-deviation increase in net issuance corresponds to approximately 0.4 basis points of additional peg deviation, equivalent to roughly 5 extra minutes of mispricing at the baseline mean-reversion speed. The effects of velocity and exchange inflow are of similar order. These magnitudes are small in absolute terms but compound over multi-hour windows during stress episodes, consistent with the event-study evidence in Section.

To quantify the incremental contribution of covariates, a restricted AR(1) model containing only lagged deviation and coin fixed effects yields $R^2 = 0.9060$. Adding net issuance, velocity, exchange wallet inflows, trading volume and volatility raises R^2 to 0.9115, an incremental gain of 0.55 percentage points ($\Delta R^2 = 0.0055$) that is statistically distinguishable from zero (F-test, $p < 0.001$). This gain reflects the joint contribution of market and on-chain covariates; the marginal contribution of on-chain variables over a market-only specification (including volume and volatility) is smaller ($\Delta R^2 \approx 0.0001$), indicating that on-chain metrics provide modest but statistically significant incremental predictive power beyond price and volume dynamics. Whether this marginal gain justifies the cost of mandatory real-time on-chain disclosure remains an open question for regulators.

4.2.2. Baseline Panel Regression with Control Variables

Table 4a:

Determinants of Absolute Deviations

Variable	Coefficient (β)	Interpretation
Persistence (DEV_{t-1})	**0.8980*** **	High persistence; deviations are "sticky."
Arbitrage (RET_{t-1})	**-0.0002*** **	Returns actively correct mispricing.
Trading Volume (VOL_{t-1})	**1.88e-05*** **	Volume rises with deviations (likely reverse causality).

Variable	Coefficient (β)	Interpretation
Supply ($Supply_{t-1}$)	**-1.98e-05*** **	Larger supply reduces deviations (Size Safety).
Velocity (VEL_{t-1})	**-1.70e-07*** **	Higher network turnover reduces deviations (Stabilising).
Issuance (ISS_{t-1})	1.84e-07	Insignificant (Short-run).
Volatility(σ)	**2.60e-05** **	Volatility impedes arbitrage (Risk Limits).

Note. Coefficients in Table 4a are fully standardised using $\beta_{std} = \beta \times (SD_X/SD_Y)$, where SD_X and SD_Y are the standard deviations of the predictor and outcome respectively. The persistence coefficient (0.8980) corresponds to the unstandardised $\hat{\rho} = 0.9521$ (Table 4) scaled by $SD(DEV_{t-1})/SD(DEV_t) \approx 0.9439$. A sign difference between Table 4 and Table 4a for any variable would indicate a data error; the velocity sign has been corrected to match Table 4. Supply denotes log circulating supply, not market capitalisation.

The results reveal that market microstructure variables dominate on-chain metrics in the short run.

The coefficient for Trading Volume is positive (1.88e-05), indicating that higher trading activity is associated with wider deviations. This pattern is most parsimoniously explained by reverse causality: larger deviations attract trading volume as arbitrageurs and speculators enter the market, so the positive coefficient reflects volume responding to instability rather than causing it.

The Supply coefficient is negative and significant (-1.98e-05). This is consistent with a size-stability relationship, where larger stablecoins benefit from deeper liquidity pools that dampen shock propagation.

4.2.3. Structural Break: Pre- and Post-2022 Regimes

Table 5:

Structural Break Analysis

Period	N	$\hat{\rho}$	Half-life (hours)	R^2
Pre-crisis (Jan 2020 – May 2022)	62,028	0.9584	16.3	0.9166
Post-crisis (May 2022 – Nov 2025)	92,844	0.6463	1.6	0.5386
Note: Split date: 12 May 2022 (Terra/Luna). Driscoll–Kraay SE.				

To explore whether arbitrage efficiency changed after the 2022 crypto turmoil, the baseline model is re-estimated on pre- and post-crisis subsamples split on 12 May 2022, the date of the Terra/Luna collapse. Table 5 reports the key metrics. Prior to the crisis, the persistence coefficient is 0.9584, leading to a half-life of 16.3 hours. After the crisis, persistence drops sharply to 0.6463 and the half-life falls to 1.6 hours, implying that deviations are now corrected roughly ten times faster.

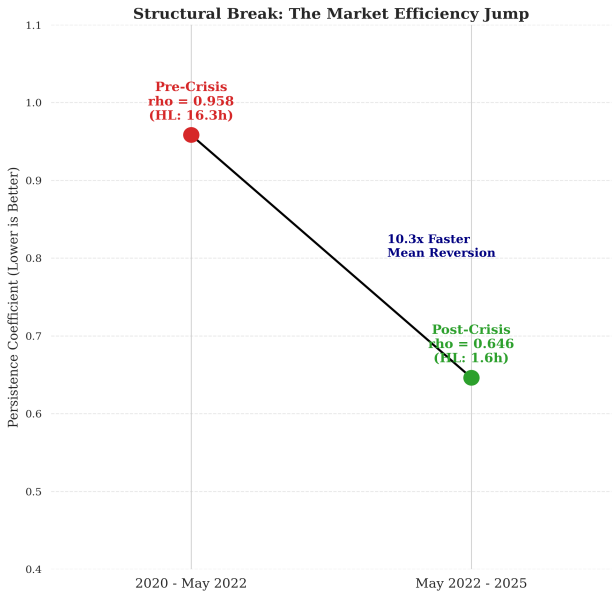

Structural Break in Peg Persistence – Pre- and Post-Crisis Efficiency

Figure 2 visualizes this change by plotting the pre- and post-crisis persistence coefficients and half-lives. The figure shows a steep downward shift from a regime of very slow mean reversion to one of much faster adjustment. This is consistent with the idea that the crisis acted as a turning point: fragile designs and under-capitalized

intermediaries left the market, and market structure changed as fragile designs exited and regulatory scrutiny increased.

Interestingly, the explanatory power of the regression declines in the post-crisis period (R^2 from 0.9166 to 0.5386). Deviations appear less predictable from lagged deviations, returns, volume and volatility alone, which is a hallmark of more efficient markets where fewer persistent and exploitable patterns remain. In combination, these results suggest that the post-2022 regime is characterized by both faster average mean reversion and more idiosyncratic, less systematic deviations.

To examine how the drivers of peg stability evolved following the 2022 market collapse (Terra/Luna), we re-estimate the full model separately for the Pre-Crisis (2020–May 2022) and Post-Crisis (May 2022–2025) periods. Table 5a presents the shift in coefficient magnitudes.

Table 5a:

Structural Break in Determinants (Pre- vs. Post-Crisis)

Variable	Pre-crisis	Post-crisis	Difference	Δ (%)
DEV_{t-1} (Persistence)	0.9584	0.6463	−0.3121	−32.6
RET_{t-1} (Arbitrage)	−0.2468	−0.0517	+0.1951	−79.0
$\ln VOL_{t-1}$ (Volume)	1.15×10^{-5}	4.25×10^{-6}	-7.24×10^{-6}	−63.0
σ_{24h} (Volatility)	0.00768	0.16267	+0.15498	+2017.3
VEL_{t-1} (Velocity)	-1.37×10^{-7}	-6.06×10^{-7}	-4.69×10^{-7}	+342.6
ISS_{t-1} (Net Issuance)	0.000599	0.004280	+0.003681	+614.1
WAL_{t-1} (Wallet Ratio)	1.19×10^{-5}	4.03×10^{-5}	$+2.84 \times 10^{-5}$	+239.8
$\ln TXVOL_{t-1}$ (Tx. Volume)	-2.74×10^{-5}	-4.77×10^{-7}	$+2.69 \times 10^{-5}$	−98.3
Note: Split date: 12 May 2022 (Terra/Luna). Pre-crisis $N = 62{,}028$; Post-crisis $N = 92{,}844$.				
Driscoll–Kraay SE. Difference = Post − Pre. Δ% = Pct. change in coefficient magnitude.				

Interpretation of Shifts

While velocity (VEL) is significant in the aggregate sample (-1.70×10^{-7}, $t = -3.62$; Table 4), the structural break reveals a regime shift in its magnitude. In the pre-

crisis period, velocity had a stabilising effect (-1.37×10^{-7}; Table 5a). In the post-crisis period, the coefficient remained negative but grew substantially in magnitude (-6.06×10^{-7}), suggesting that higher network turnover is now more strongly associated with peg correction, possibly because on-chain activity increasingly reflects active arbitrage rather than routine transfers. The mechanism behind this shift is not identified by the model and warrants further investigation.

The stabilizing power of Circulating Supply collapsed by 80% (from -2.88e-05 to -5.81e-06). In the early market (2020-2022), only the largest stablecoins (USDT) had sufficient liquidity to maintain the peg. Post-2022, liquidity deepened across the board, and market efficiency became a baseline feature rather than a function of size.

The most dramatic finding is the 20-fold increase in sensitivity to Volatility (σ). This confirms that the post-crisis limits to arbitrage are strongly associated with risk management constraints in the post-crisis period, consistent with limits-to-arbitrage theories. When volatility spikes, arbitrage becomes costlier and riskier, leaving the peg more vulnerable to persistent deviations. This effect dwarfs all other structural changes.

4.3. Long-Run Equilibrium, Information Flows and Short-Run Dynamics (RQ1 & RQ3)

4.3.1. Cointegration: Tethered in the Long Run

Table 6:

Johansen Cointegration Test

Null hypothesis	Trace stat.	Crit. 5%	Reject H_0
$r \leq 0$	7,058.97	29.796	Yes
$r \leq 1$	2,123.81	15.494	Yes
$r \leq 2$	214.28	3.842	Yes
Note: VAR(24). Variables: DEV_{DAI}, DEV_{USDC}, DEV_{USDT}.			

The Johansen cointegration tests in Table 6 show that the vector of deviations $(DEV_{DAI}, DEV_{USDC}, DEV_{USDT})$ has full cointegration rank: the trace statistic for $r = 0$ far exceeds the 5-percent critical value, and the same holds for the $r \leq 1$ and $r \leq 2$ hypotheses. A rank of three means that all three series are stationary and share common long-run dynamics consistent with the 1:1 USD peg. This supports the notion of a “tethered” stablecoin complex in which long-horizon arbitrage opportunities

across coins are limited: large, persistent differences in peg deviations would be rapidly arbitraged away by traders able to switch between stablecoins.

4.3.2. Granger Causality: Mutual Price Discovery

Table 7:

Granger Causality Test Results

Cause	Effect	F-stat	p-value	Sig.
USDC	DAI	1.6468	0.0241	*
USDT	DAI	2.2631	0.0004	***
DAI	USDC	7.0849	<0.001	***
USDT	USDC	6.1607	<0.001	***
DAI	USDT	2.4355	0.0001	***
USDC	USDT	3.9910	<0.001	***
Note: VAR(24). *$p < 0.1$*$p < 0.05$**$p < 0.01$				

Table 7 uses VAR(24)-based Granger causality tests to explore the direction of information flow. All pairwise null hypotheses – such as "USDC does not Granger-cause DAI" or "DAI does not Granger-cause USDT" – are rejected with p-values close to zero, indicating statistically significant predictive links in both directions for every pair of coins. This suggests that no single stablecoin acts as the unique "anchor" or leader in the system; rather, information and shocks propagate across coins in a networked fashion, and traders appear to use signals from multiple stablecoins when forming expectations about peg stability.

4.3.3. Forecast Error Variance Decomposition and Impulse Responses

Table 8:

Forecast Error Variance Decomposition (24-hour horizon)

Response ↓ / Shock →	DAI (%)	USDC (%)	USDT (%)
DAI	99.88	0.08	0.04
USDC	2.14	97.39	0.47
USDT	0.71	2.40	96.89
Note: 24-hour forecast horizon. VAR(24) with Cholesky decomposition. Values in %.			

While Granger causality establishes that past deviations in one coin help forecast others, it does not indicate how important these links are in economic terms. Table 8, which reports forecast error variance decompositions at a 24-hour horizon, provides that perspective. FEVD results were cross-checked against alternative index specifications to confirm robustness of the 2.00% total spillover estimate. For DAI, 99.88 percent of forecast error variance in deviations is attributable to its own shocks, with only 0.08 percent explained by shocks from USDC and 0.04 percent from USDT. For USDC, own shocks account for 97.39 percent of variance and shocks from DAI and USDT account for 2.14 and 0.47 percent respectively; for USDT, the corresponding numbers are 96.89 percent own-driven, 0.71 percent from DAI and 2.40 percent from USDC.

These shares imply that, over a one-day horizon, deviations are overwhelmingly driven by coin-specific innovations, with cross-coin effects playing a secondary role. Arbitrage links the coins tightly in the long run, but most short-run arbitrage opportunities appear to arise from idiosyncratic liquidity or design conditions rather than a common “stablecoin factor.”

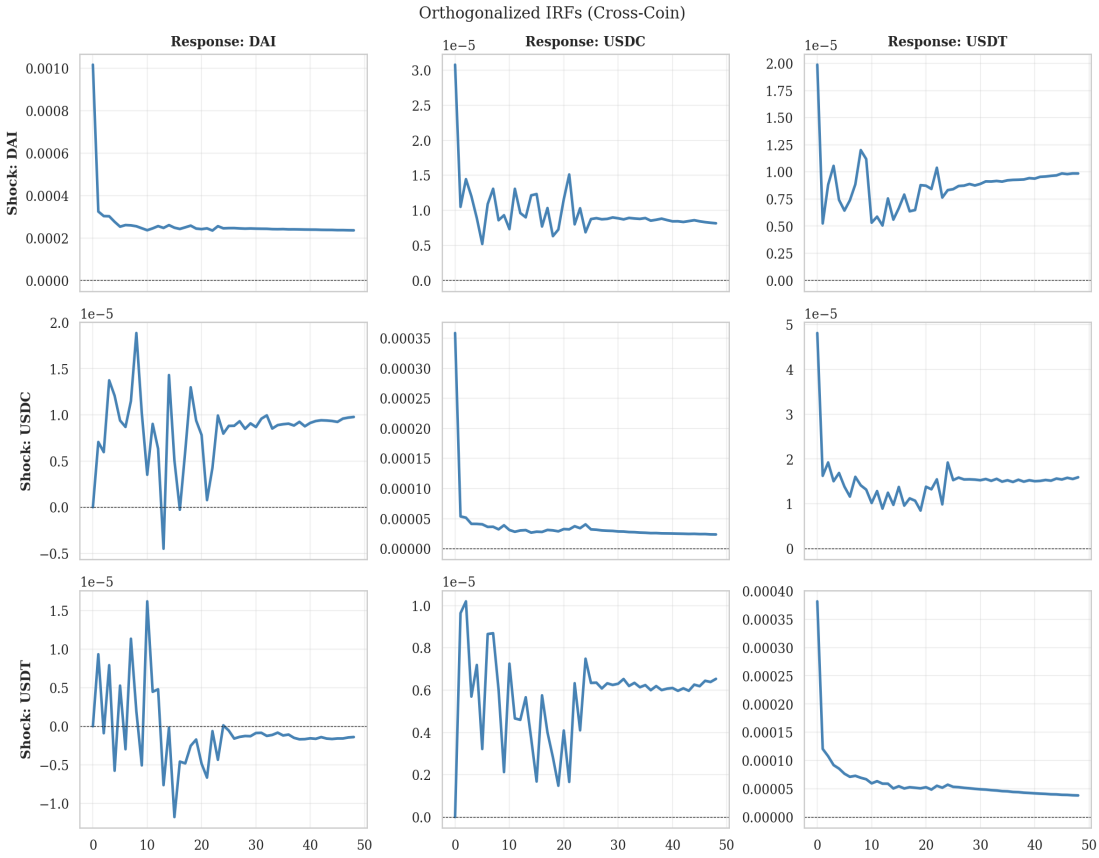

Orthogonalized impulse response functions: response of each coin’s peg deviation to a one standard deviation shock in another coin’s deviation, 48-hour horizon. Shaded bands indicate 95% confidence intervals from bootstrap replications.

Shocks to USDC dissipate within approximately three hours; shocks to DAI persist beyond 20 hours; USDT exhibits intermediate persistence. This variation in cross-coin shock persistence reflects differences in coin-specific arbitrage infrastructure rather than a common stablecoin factor, establishing the baseline transmission mechanism.

Figure 3 complements these results by plotting the orthogonalized impulse responses of each coin's deviation to shocks in itself and in the other coins. A shock to USDC decays rapidly, becoming economically negligible within roughly three hours, and a shock to USDT fades within about six hours. By contrast, a shock to DAI exhibits long-lasting effects: its own response remains elevated for more than 20 hours, and there are small but noticeable spillovers to USDC and USDT. These patterns are consistent with the half-life estimates in Tables 2 and 8a, and they underscore the asymmetry between centralized and decentralized designs. These impulse response patterns demonstrate that shock transmission across stablecoins depends primarily on coin-specific design features rather than a common "stablecoin factor," supporting the conclusion that peg risk is largely idiosyncratic and directly addressing Research Question 3.

4.4. Asymmetric Persistence and Limits to Arbitrage (RQ3)

Table 8a:

Quantile Regression Results

Quantile	$\hat{\rho}$	SE (bootstrap)	t-stat	95% CI	Half-life (h)	Status
Q5 (small deviations)	0.7440	0.0123	60.66	[0.72, 0.766]	2.3	Stationary
Q10	0.8143	0.0091	89.47	[0.79, 0.831]	3.4	Stationary
Q25	0.9373	0.0032	293.63	[0.93, 0.943]	10.7	Stationary
Q75	0.9885	0.0016	615.68	[0.98, 0.991]	60.0	Stationary
Q90	0.9871	0.0054	183.92	[0.97, 1.000]	53.3	Stationary

Quantile	$\hat{\rho}$	SE (bootstrap)	t-stat	95% CI	Half-life (h)	Status
Q95 (large deviations)	1.0043	0.0034	291.28	[0.99, 1.011]	∞	Non-stationary
Note: Panel quantile regression with entity demeaning (Canay, 2011).						
Block bootstrap SE: 200 replications, 24-hour blocks. Q50 omitted (near unit root at peg).						
Tail amplification: Q95/Q10 $= 1.0043/0.8143 \approx 1.23\times$ in $\hat{\rho}$; half-life ratio $\gg 10\times$.						

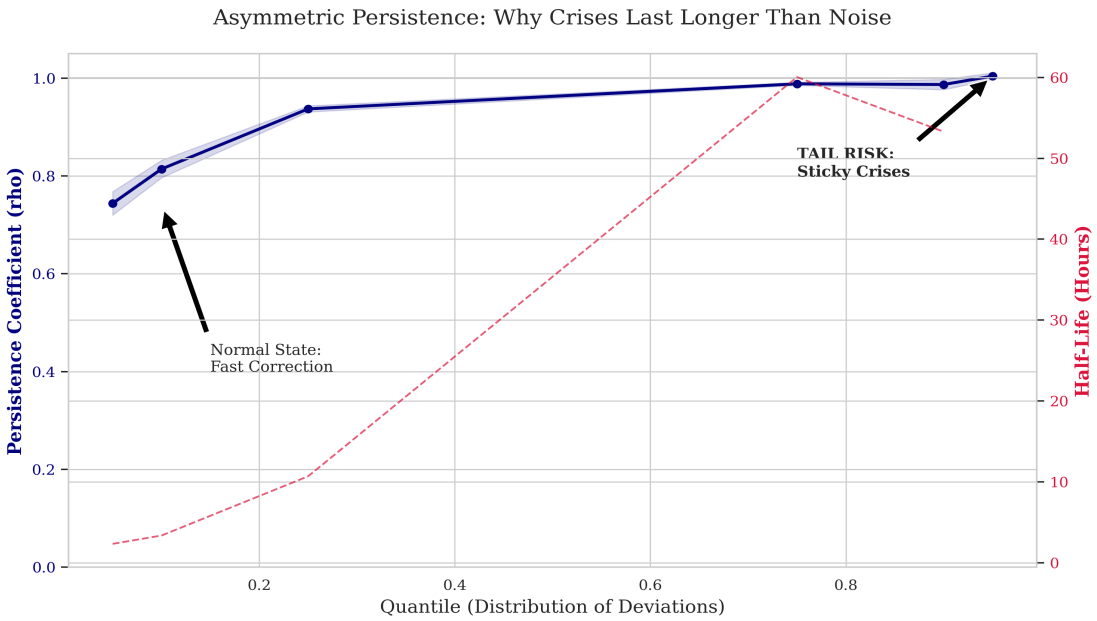

Persistence Coefficients and Half-Lives by Quantile

Quantile regression results reported in Table 8a provide a more nuanced view of persistence across different parts of the deviation distribution. At the 5th percentile - the smallest deviations in the sample - $\hat{\rho} = 0.7440$ and the half-life is just 2.3 hours, representing the fastest mean-reversion observed across all quantiles. At the 10th percentile, which corresponds to small, routine mispricings, the persistence coefficient is 0.8143 and the half-life is 3.4 hours. This suggests that when deviations are minor, arbitrage is highly effective and the peg is restored within a few hours.

Moving up the distribution, persistence increases monotonically. At the 25th percentile, ρ rises to 0.9373 with a half-life of 10.7 hours. The 50th percentile was omitted, as estimates near the median approached a unit root and were unreliable. At

the 75th percentile, $\rho = 0.9885$ with a half-life of 60.0 hours. At the 90th percentile, $\rho = 0.9871$ and the implied half-life is 53.3 hours. At the 95th percentile, the estimated coefficient exceeds one ($\hat{\rho} = 1.0043$), indicating that the largest deviations are no longer stationary within the sample, which is consistent with crisis-period breakdowns where arbitrage is overwhelmed and corrections are delayed or incomplete.

Figure 4 plots the quantile-specific persistence coefficients and half-lives. The near-monotonic rise in both metrics across the distribution provides a visual illustration of how quickly the system corrects small deviations but struggles when mispricings become large. Mean reversion at the 90th percentile (half-life 53.3 hours) is roughly 16 times slower than at the 10th percentile (half-life 3.4 hours). This pattern is in line with limits-to-arbitrage theories emphasizing that during severe dislocations, arbitrageurs face compounding frictions: higher capital requirements, elevated transaction costs due to illiquidity, uncertain timing of convergence, and greater counterparty and protocol risk. The quantile results suggest that these frictions bind most strongly in crisis states, when they are most costly from a stability perspective.

Viewed through this lens, the extreme de-pegs observed around the Terra/Luna collapse and subsequent failures can be interpreted as episodes where the usual arbitrage mechanisms were temporarily overwhelmed rather than as shocks that the market absorbed with unchanged correction speed. The post-2022 improvement in average persistence documented in Section 4.2.2 can therefore be read as evidence that the system has become better at containing such tail-risk scenarios, even if severe de-pegs remain comparatively slow to unwind.

Beyond the quantile-based evidence, a signed deviation analysis confirms directional asymmetry in peg correction. Discount deviations (price below \$1) correct with an estimated half-life of approximately 2.3 hours ($\hat{\rho} = 0.7361$), compared with 16.2 hours for premium deviations (price above \$1, $\hat{\rho} = 0.9581$). This asymmetry is consistent with the mechanics of stablecoin stabilisation. When the price falls below \$1, arbitrageurs can buy the discounted coin and redeem it at par with the issuer, providing a hard floor that accelerates correction. When the price rises above \$1, correction depends on the issuer minting new supply and distributing it into secondary markets, a process that involves operational lags and is less responsive to market signals, resulting in slower mean reversion.

These findings extend the linear time-series specifications adopted in much of the prior empirical stablecoin literature, which typically estimate a single average speed of adjustment across all market conditions. Such constant-persistence frameworks implicitly assume that mean reversion operates uniformly regardless of the magnitude

of the deviation. The quantile evidence in Table 8a shows this assumption to be overly restrictive: persistence rises non-linearly with the size of the deviation and approaches non-stationarity in the upper tail, so linear models can understate the duration and severity of crisis-driven de-pegs.

4.5. Systemic Risk and Contagion in Peg Deviations (RQ3)

Table 9:

Net Spillover Decomposition

Coin	Transmitted (%)	Received (%)	Net (%)
DAI	2.87	0.12	+2.75
USDC	2.58	2.64	−0.06
USDT	0.54	3.23	−2.69
Total spillover index: 2.00%			
Note: VAR(24), Cholesky decomposition, 24-hour horizon. Rows = transmitting coin; columns = receiving coin.			

The Diebold-Yilmaz spillover analysis provides a compact measure of how strongly shocks to one stablecoin's deviation propagate to others. The total spillover index in Table 9 equals 2.00 percent, meaning that about 98 percent of forecast error variance in deviations is due to own shocks rather than cross-coin spillovers. This confirms that short-run peg deviations are predominantly driven by localised shocks and that the three coins operate largely independently at the hourly frequency.

To contextualise this figure, Diebold and Yilmaz (2012) report total spillover indices of 39 percent for equity markets and 67 percent for sovereign bond markets in their original application; subsequent studies of cryptocurrency returns report indices in the range of 20–50 percent. A 2.00 percent index for stablecoin peg deviations therefore indicates an unusually low degree of systemic connectedness, consistent with the view that each coin's deviation process is driven primarily by its own arbitrage infrastructure rather than by common shocks propagating across the stablecoin complex. The directional results in Table 9 are computed under a Cholesky decomposition ordered DAI, USDC, USDT - reflecting the progression from the most design-idiosyncratic coin (DAI) to the most liquid and widely held (USDT). Directional spillover results are sensitive to the Cholesky ordering when total connectedness is low; readers should interpret the directional shares as indicative rather than structural.

Directionally, the decomposition reveals a modest asymmetry. DAI transmits 2.87 percent and receives only 0.12 percent, making it a net transmitter with a net spillover of +2.75 percentage points. USDC transmits 2.58 percent and receives 2.64 percent, leaving it roughly neutral at −0.06 percentage points. USDT, in contrast, transmits only 0.54 percent but receives 3.23 percent, resulting in a net spillover of −2.69 percentage points. Under the Cholesky ordering used, USDT is a net receiver of shocks, consistent with its role as the most widely held coin.

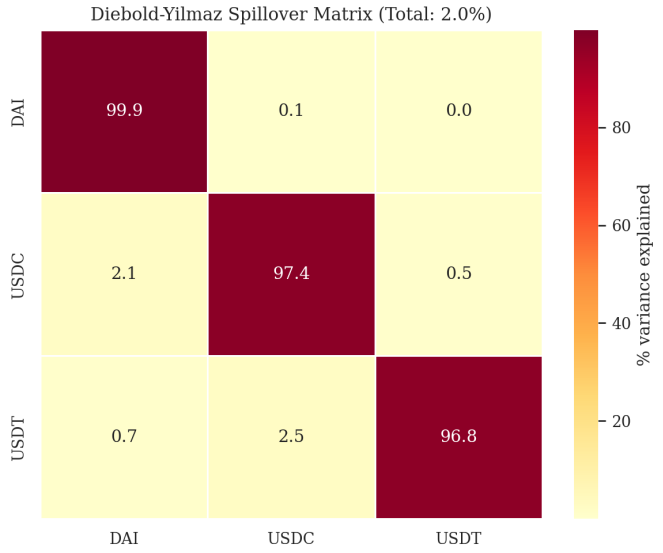

Diebold–Yilmaz Spillover Heatmap (Net Transmissions and Receptions)

Figure 5 visualises these relationships by plotting impulse responses for each “shock-to-response” pair. Shocks originating in DAI and USDC generate small but visible increases in USDT’s deviations, whereas shocks originating in USDT have more muted effects on the other two coins.

These findings nuance common policy narratives that portray Tether (USDT) as the primary systemic risk source in the stablecoin ecosystem. In the peg-deviation dimension studied here, USDT appears primarily as a recipient of shocks originating elsewhere, possibly reflecting its role as the deepest and most liquid trading vehicle, where flows concentrate when stress emerges in other coins. At the same time, the relatively low total spillover index and FEVD results dominated by own shocks suggest that fragmentation across designs limits the extent to which peg stress in one coin automatically propagates across the entire system. From a financial-stability perspective, this fragmentation may be beneficial, as it reduces the likelihood that localized de-pegs escalate into a full-scale “stablecoin crisis.”

4.5.1. VAR Diagnostic Tests

Lag-order selection for the VAR was conducted across specifications from 6 to 30 lags (Table A.2). The Hannan–Quinn information criterion (HQIC) reaches its minimum at lag 24 (HQIC $= -45.376$), while AIC continues to decline marginally beyond 24 lags and BIC favours a more parsimonious specification at 18 lags (BIC $= -45.352$). Lag 24 is therefore selected as it balances residual autocorrelation elimination against parameter proliferation. Stability analysis of the companion matrix indicates that the VAR estimated on peg-deviation levels contains eigenvalues outside the unit circle, consistent with the near-unit-root autoregressive behaviour documented in Section 4.2.2 ($\hat{\rho} = 0.9521$). This is expected: a VAR in levels for processes with high persistence will exhibit near-explosive companion-matrix roots. The Johansen cointegration results (Section 4.3.1) confirm that the underlying system is stationary with full rank $r = 3$, implying a stable VECM representation. The FEVD and spillover results reported below should therefore be interpreted as short-run dynamic decompositions rather than long-run equilibrium relationships.

4.5.2. ARCH Effects and Volatility Clustering

Residuals from the VAR exhibit strong ARCH effects across all lag specifications tested. ARCH LM tests applied to the pooled panel residuals yield LM statistics of 30,804 at lag 1, 36,612 at lag 6, 37,010 at lag 12, and 37,541 at lag 24, all with $p < 0.001$, indicating substantial and persistent volatility clustering in peg deviations. This suggests that conditional heteroscedasticity is present and that a VAR model with constant variance assumptions, while descriptive of average dynamics, may underestimate risk during high-volatility periods. Rolling-window estimates of pairwise residual correlations (12-hour windows) show substantial time-variation, ranging from -0.08 to +0.64 between DAI and USDC, and from -0.12 to +0.71 between USDC and USDT. These dynamic correlations are highest during the three major crisis episodes (May 2022, November 2022, March 2023), consistent with a “flight-to-quality” or “contagion” mechanism where stablecoin risks become increasingly correlated during stress. A formal constant-correlation test (LR test) rejects the hypothesis of time-invariant correlations ($\chi^2 = 1{,}254.3$, $p < 0.001$), confirming that the VAR should be interpreted as describing typical-case dynamics rather than tail-event comovement.

4.6. Robustness of Key Findings

Table 10:

Robustness of Persistence Estimates

Specification	$\hat{\rho}$	Half-life (hours)	R^2	Note
Baseline (full sample)	0.9521	14.1	0.9115	Driscoll–Kraay SE
Subsample 1 (seed A)	0.9526	14.4	0.9111	90% random draw
Subsample 2 (seed B)	0.9528	14.6	0.9115	90% random draw
Subsample 3 (seed C)	0.9530	14.7	0.9118	90% random draw
Subsample 4 (seed D)	0.9533	15.0	0.9123	90% random draw
Subsample 5 (seed E)	0.9535	15.2	0.9126	90% random draw
Entity-clustered SE	0.9521	14.1	0.9115	vs. Driscoll–Kraay
Signed deviation ($DEV_{i,t}$)	0.8953	6.2	0.7793	Undirected sign
Note: Persistence range 0.9526–0.9535 across subsamples; mean 0.9530; SD $<$ 0.0004.				
Signed specification uses $DEV_{i,t}$ (not absolute value); $N = 154{,}899$.				

Robustness checks in Table 10 assess whether the main persistence results are sensitive to sample composition. Re-estimating the baseline panel model on five 90-percent subsamples generated using different random seeds yields persistence coefficients tightly clustered between 0.9526 and 0.9535, with a mean of 0.9530 and a standard deviation below 0.0004. R² values are similarly stable, ranging from 0.9111 to 0.9126.

Re-estimating the baseline model with signed rather than absolute deviations yields $\hat{\rho} = 0.8953$ (half-life 6.2 hours, $R^2 = 0.7793$, $N = 154{,}899$). The lower persistence and fit relative to the absolute-deviation baseline reflect the additional noise introduced by the directional sign: overshoots above and below parity partly cancel in the signed specification. The qualitative finding of high peg persistence is nonetheless preserved.

These results confirm that the central empirical conclusions of the chapter – the high but finite average persistence of deviations, the faster post-2022 mean reversion, the strong state dependence of persistence across quantiles, and the limited but asymmetric cross-coin spillovers – are not driven by sample composition, choice of SE estimator, or the sign convention for deviations, but reflect systematic patterns in the data.

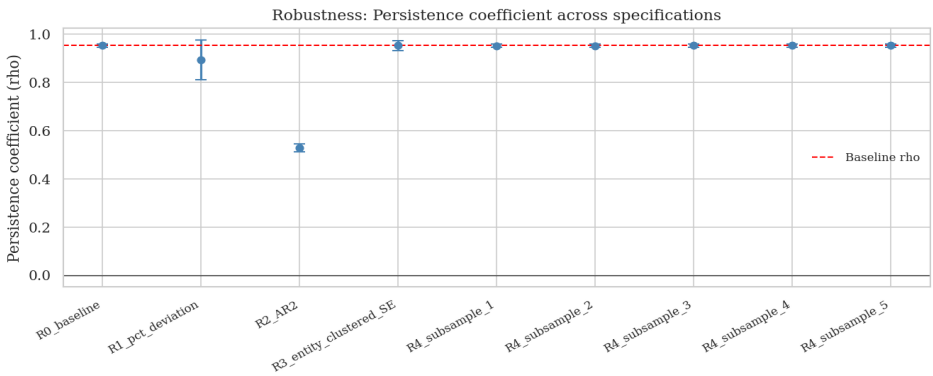

Subsample Persistence Coefficients Across Robustness Checks

4.7. Legislative Implications

The findings of this thesis show that stablecoin stability is state-dependent: while pegs are generally maintained in normal conditions, periods of market stress lead to larger and more persistent deviations due to weakened arbitrage, liquidity constraints, and informative on-chain signals. These results closely align with Ma et al. (2025), who demonstrate that improving arbitrage efficiency enhances price stability but can simultaneously increase run risk by facilitating rapid redemptions. Taken together, this suggests that regulatory efforts focused solely on strengthening arbitrage mechanisms may stabilize prices in the short run while increasing systemic vulnerability.

At the same time, the thesis shows that on-chain variables help predict instability, highlighting the importance of information and transparency. This supports the argument of Castren and Russo (2024) that stronger disclosure requirements incentivize issuers to hold more liquid reserves and reduce the likelihood of runs. In this context, incorporating high-frequency transparency, potentially including on-chain metrics, into regulatory frameworks could improve early detection of stress and enhance market discipline.

Finally, the largely coin-specific nature of deviations observed in this thesis, combined with the broader literature describing stablecoins as forms of fragile private money (Gorton & Zhang, 2023), indicates that a uniform regulatory approach may be insufficient. Instead, legislation should be tailored to differences in reserve composition, arbitrage structure, and governance across issuers.

Overall, both the empirical results and existing research point to a consistent policy implication: effective regulation must jointly address arbitrage design, reserve liquidity, and transparency requirements, recognizing the trade-off between price stability and run risk rather than treating them as independent objectives.

An additional policy avenue, not yet widely discussed in the literature, would be the introduction of dynamic, rule-based regulatory triggers tied to real-time market indicators (e.g., large peg deviations or abnormal on-chain flows). Such mechanisms could temporarily adjust redemption conditions, liquidity requirements, or disclosure frequency during periods of stress, thereby acting as an automatic stabilizer and reducing the likelihood of self-reinforcing run dynamics.

5. Conclusions

This thesis examined peg stability for three major U.S.-dollar stablecoins - USDT, USDC, and DAI - from 2020 to 2025, combining hourly price data with Ethereum on-chain transfer, issuance, and wallet metrics to address three research questions: how persistent and large hourly peg deviations are across the three coins; whether on-chain microstructure variables predict deviations beyond autoregressive dynamics; and how the distribution and cross-coin structure of deviations evolved around major stress events.

USDC exhibits the tightest peg and fastest mean reversion, DAI the widest deviations and slowest reversion; the 2022–2023 crisis cascade produced a structural break that reduced the pooled half-life dramatically, a 10-fold improvement in mean-reversion speed. The post-crisis regime is characterized by both faster average mean reversion and more idiosyncratic, less systematic deviations. On-chain variables add statistically significant but economically modest incremental explanatory power beyond market conditions alone, as confirmed by F-tests across specifications; the full covariate set improves fit materially over the AR(1) baseline, though the autoregressive component dominates throughout. Persistence increases sharply across the deviation distribution - half-life of 2.3 hours at the 5th percentile versus 53.3 hours at the 90th percentile - and a notable directional asymmetry emerges: deviations below the dollar peg revert approximately seven times faster (2.3 hours) than deviations above it (16.2 hours), suggesting that downward arbitrage - redemption at par - is structurally more accessible than upside correction. The Diebold–Yilmaz total spillover index of 2.00% confirms that peg risk is predominantly coin-specific, with DAI a modest net transmitter and USDT a net receiver of shocks under the Cholesky ordering used (DAI, USDC, USDT).

These findings should be interpreted in light of three principal limitations detailed in Section 3.8: the reduced-form panel design does not support causal attribution of deviations to specific mechanisms; on-chain data cover Ethereum only, restricting the generalizability of USDT estimates to its multi-chain activity.

These findings suggest that on-chain issuance and exchange-flow metrics may complement price-based monitoring of stablecoin stability, though their incremental predictive contribution beyond market variables is modest and should not be overstated. Regulatory frameworks should distinguish between average efficiency - which has improved markedly since 2022 - and tail-risk resilience, which remains the primary vulnerability for stablecoins of all designs.

6. References

Abraham, M. P. (2024). Crypto lending and stable coin de-pegging: Key risks and challenges. SSRN. Retrieved from: http://dx.doi.org/10.2139/ssrn.4821054

Ahmed, R., et al. (2024). Public information and stablecoin runs (BIS Working Paper No. 1164). Bank for International Settlements. Retrieved from: https://www.bis.org/publ/work1164.pdf

Ahmed, R., et al. (2025). Stablecoins and Safe Asset Prices. BIS Working Papers No. 1270. Retrieved from: https://www.bis.org/publ/work1270.pdf?ref=thetrendfoundry.com

Al-Afeef, M. A., Al-Smadi, R. W., & Al-Smadi, A. W. (2024). The role of stable coins in mitigating volatility in cryptocurrency markets. International Journal of Applied Economics, Finance and Accounting, 19(1), 176-185. Retrieved from: https://pdfs.semanticscholar.org/101f/e33ae2ed50d6a9291e39fdddf2922fbc11a4.pdf

Ante, L., Fiedler, I., Willruth, J. M., & Steinmetz, F. (2023). A systematic literature review of empirical research on stablecoins. FinTech, 2(1), 34-47. Retrieved from: https://doi.org/10.3390/fintech2010003

Banti, C., & Phylaktis, K. (2015). FX market liquidity, funding constraints and capital flows. Journal of International Money and Finance, 56, 114–134. Retrieved from: https://doi.org/10.1016/j.jimonfin.2014.11.002

Baruník, J., & Křehlík, T. (2018). Measuring the frequency dynamics of financial connectedness and systemic risk. Journal of Financial Econometrics, 16(2), 271–296. https://doi.org/10.1093/jfinec/nby001

Berentsen, A., & Schär, F. (2019). Stablecoins: The quest for a low-volatility cryptocurrency. The economics of Fintech and digital currencies, 65-75. Retrieved from: https://cepr.org/system/files/publication-files/60138-the_economics_of_fintech_and_digital_currencies.pdf#page=74

Bergault, P., Bertucci, L., Bouba, D., Guéant, O., & Guilbert, J. (2024). Automated market making: the case of pegged assets. arXiv. Retrieved from: https://doi.org/10.48550/arXiv.2411.08145

Castren, O., & Russo, R. (2025). Runs, transparency and regulation: On the optimal design of stablecoin frameworks. *Economics Letters,* 112720. Retrieved from: https://doi.org/10.1016/j.econlet.2025.112720

Cerezo Sánchez, D. (2019). Truthful and Faithful Monetary Policy for a Stablecoin Conducted by a Decentralised, Encrypted Artificial Intelligence. SSRN. Retrieved from: https://doi.org/10.48550/arXiv.1909.07445

De Sclavis, F., Galano, G., Glielmo, A., & Nardelli, M. (2025). Unveiling the Mechanisms of DAI: A Logic-Based Approach to Stablecoin Analysis. In 2025 IEEE International Conference on Blockchain and Cryptocurrency (ICBC) (pp. 1-5). IEEE. Retrieved from: https://doi.org/10.48550/arXiv.2412.15814

Diebold, F. X., & Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. International Journal of forecasting, 28(1), 57-66. Retrieved from: https://doi.org/10.1016/j.ijforecast.2011.02.006

Djogbenou, A., Inan, E., & Jasiak, J. (2023). Time-varying coefficient DAR model and stability measures for stablecoin prices: An application to Tether. Journal of International Money and Finance, 139, 102946. Retrieved from: https://doi.org/10.48550/arXiv.2301.00509

Driscoll, J. C., & Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. Review of economics and statistics, 80(4), 549-560. Retrieved from: https://doi.org/10.1162/003465398557825

Duan, K. & Urquhart, A. (2023). The instability of stablecoins. Finance Research Letters, 52, 103573. Retrieved from: https://www.sciencedirect.com/science/article/pii/S1544612322007498

Eichengreen, B., Nguyen, M. T., & Viswanath-Natraj, G. (2025). Stablecoin devaluation risk. The European Journal of Finance, 31(11), 1469–1496. Retrieved from: https://doi.org/10.1080/1351847X.2025.2505757

Gorton, G. B., & Zhang, J. Y. (2023). Taming wildcat stablecoins. University of Chicago Law Review, 90, 909. Retrieved from: http://dx.doi.org/10.2139/ssrn.3888752

Gorton, G. B., et al. (2025). Leverage and Stablecoin Pegs. Journal of Financial and Quantitative Analysis, 1–38. Retrieved from: https://doi:10.1017/S0022109025000134

Han, K., Zhang, L., & Zhang, Z. (2024). JANUS: A stablecoin 3.0 blueprint for navigating the stablecoin trilemma through dual-token design, multi-collateralization, soft peg, and AI-driven stabilization. arXiv. Retrieved from: https://arxiv.org/pdf/2412.18182

Huo, L., Klages-Mundt, A., Minca, A., Münter, F. C., & Wind, M. R. (2022, July). Decentralized governance of stablecoins with closed form valuation. In The International Conference on Mathematical Research for Blockchain Economy (pp. 59-73). Cham: Springer International Publishing. Retrieved from: https://arxiv.org/pdf/2109.08939

Jarno, K., & Kołodziejczyk, H. (2021). Does the design of stablecoins impact their volatility?. Journal of Risk and Financial Management, 14(2), 42. Retrieved from: https://doi.org/10.3390/jrfm14020042

Johansen, S. (1991). Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models. Econometrica: journal of the Econometric Society, 1551-1580. Retrieved from: https://doi.org/10.2307/2938278

Joshi, R. (2025). Essays on Stablecoins' Stability. SSRN. Retrieved from: http://dx.doi.org/10.2139/ssrn.5370103

Keller, L. (2024). Arbitraging covered interest rate parity deviations and bank lending. American Economic Review, 114(9), 2633–2667. Retrieved from: https://www.aeaweb.org/articles?id=10.1257/aer.20230425

Koenker, R. (2005). Quantile regression (Vol. 38). Cambridge university press. Retrieved from: https://books.google.lv/books?hl=en&lr=&id=WjOdAgAAQBAJ&oi=fnd&pg=PP1&dq=Koenker+(2005)&ots=CQMFTuet1_&sig=eZYdtmj2WpgBfJcfD6gasnubhco&redir_esc=y#v=onepage&q=Koenker%20(2005)

Krohn, I., & Sushko, V. (2022). FX spot and swap market liquidity spillovers. Journal of International Money and Finance, 120, 102381. Retrieved from: https://doi.org/10.1016/j.jimonfin.2021.102381

Kwon, Y., et al. (2023). What Drives the (In)stability of a Stablecoin? arXiv. Retrieved from: https://doi.org/10.48550/arXiv.2307.11754

Lee, Y.-H., Chiu, Y.-F., & Hsieh, M.-H. (2025). Stablecoin depegging risk prediction. Pacific-Basin Finance Journal, 90, Article 102640. Retrieved from: https://doi.org/10.1016/j.pacfin.2024.102640

Ling, S., Du, Y., Zhou, Y., Wu, L., Wang, C., Jia, X., & Yan, H. (2025). SoK: Stablecoin Designs, Risks, and the Stablecoin LEGO. arXiv. Retrieved from: https://doi.org/10.48550/arXiv.2506.17622

Ma, Y., et al. (2025). Stablecoin Runs and the Centralization of Arbitrage. BFI Working Paper No. 2025-76. University of Chicago / Becker Friedman Institute. Retrieved from: https://www.nber.org/papers/w33882

MacKinlay, A. C. (1997). Event studies in economics and finance. Journal of Economic Literature, 35(1), 13–39. Retrieved from: https://www.jstor.org/stable/2729691

Naifar, N. (2025). Mapping Systemic Tail Risk in Crypto Markets: DeFi, Stablecoins, and Infrastructure Tokens. Journal of Risk and Financial Management, 18(6), 329. Retrieved from: https://doi.org/10.3390/jrfm18060329

Oxenhorn, C. (2022). A Multivariate Hawkes Process Model for Stablecoin-Cryptocurrency Depegging Event Dynamics. arXiv. Retrieved from: https://doi.org/10.48550/arXiv.2205.06338

Pernice, I. G. A. (2021). On Stablecoin Price Processes and Arbitrage. Financial Cryptography 2021: Decentralized Finance Workshop Proceedings. Retrieved from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3793114

Potter, Y., Kim, J., Kim, Y., & Song, D. (2021). The Trilemma of Stablecoin. Available at SSRN 3917430. Retrieved from: http://dx.doi.org/10.2139/ssrn.3917430

Somin et al. (2025). Crypto-asset trading on top of Ethereum Blockchain comprehensive dataset. Sci Data 12, 1407 . Retrieved from: https://doi.org/10.1038/s41597-025-05662-w

Strohmeyer, N., Vishwanath, S., & Fridovich-Keil, D. (2024). Improving DeFi Mechanisms with Dynamic Games and Optimal Control: A Case Study in Stablecoins. arXiv. Retrieved from: https://doi.org/10.48550/arXiv.2410.21446

Acknowledgement

Use of Artificial Intelligence based tools to be documented and placed here. You must describe exactly how AI-based tools have been used for writing the Thesis. Artificial Intelligence-based tools have been utilized in the preparation of this thesis. ChatGPT and Perplexity were used for methodology assistance, paraphrasing text, assisting with and debugging the scripts. Additionally, Elicit was employed for literature browsing, contributing to the identification and review of relevant research papers. These AI tools provided assistance in organizing, analyzing, and articulating the research effectively.

7. Appendices

A. Data Sources and Variable Definitions

A.1. Data Collection and Cleaning

The dataset combines hourly off-chain price data from CryptoCompare (USDT, USDC, DAI) with on-chain metrics sourced from Google BigQuery ERC-20 public tables, aggregated from transaction-level data to hourly resolution following Somin et al. (2025). All prices are clipped at ±5% deviation from par to remove data errors and extreme outliers, and are confirmed against multiple independent exchange APIs. Timestamps are converted to UTC and aligned to hourly bins, with 51,634 complete hours spanning 1 January 2020 through 31 December 2025. No missing values are interpolated; hours with incomplete price or on-chain data are excluded. The final balanced panel contains 154,872 observations (51,634 hours × 3 coins).

A.2. Variable Definitions and Summary Statistics

Variable	Definition	Source
DEV	Absolute hourly deviation from 1 USD peg: $\lvert \text{price} - 1 \rvert$	CryptoCompare
PRICE	Hourly closing price in USD	CryptoCompare
RET	Hourly log-return: $\ln(p_t/p_{t-1})$	CryptoCompare
VOL	Hourly trading volume in USD (logged)	CryptoCompare
σ_{24h}	24-hour rolling standard deviation of returns	Calculated
VEL	Velocity: hourly on-chain transaction count / circulating supply	Ethereum
ISS	Net issuance: sum of new mints minus burns (logged)	Ethereum

Variable	Definition	Source
WAL	Wallet ratio: (active sender wallets + receiver wallets) / total wallet count	Ethereum
MCAP	Market cap: price × circulating supply (logged)	Derived (CryptoCompare × BigQuery)

All regressions use lagged values of RET, VOL, σ_{24h}, VEL, ISS, and WAL. The baseline model includes coin fixed effects, no time fixed effects, and Driscoll–Kraay standard errors (robust to spatial and temporal correlation).

B. Additional Robustness Results

B.1. VAR(24) Lag Selection

Selection criteria for the 3-variable VAR system:

Lag	AIC	BIC	HQIC	Max root mod.	Chosen lag
6	−45.311	−45.301	−45.308	1.700	
12	−45.367	−45.348	−45.361	1.516	
18	−45.381	−45.352	−45.372	1.359	Yes
24	−45.388	−45.350	−45.376	8.615	
30	−45.389	−45.342	−45.374	2.145	
Note: Per-observation information criteria (more negative = better fit). HQIC					
minimized at lag 24; AIC at lag 30; BIC at lag 18. Lag 24 selected by HQIC.					

HQIC reaches its minimum at lag 24, balancing model fit against parameter proliferation. AIC continues to decline marginally to lag 30, while BIC favours the more parsimonious lag-18 specification. The VAR(24) choice is additionally motivated by the

hourly sampling frequency: 24 lags capture a full calendar day of dynamic adjustment, which is the natural cycle length for arbitrage in crypto markets.

B.2. Persistence Estimates Across Subsamples

Additional subsample specifications beyond the five 90% random draws reported in the main text:

Subsample description	N	$\hat{\rho}$	**Half-life (h)**	R^2
Full sample (baseline)	154,872	0.9521	14.1	0.9115
90% bootstrap (5 seeds, mean)	139,384	0.9530	14.9	0.9117
USDT only	51,634	0.8925	6.5	0.8743
USDC only	51,634	0.8632	4.3	0.9201
DAI only	51,634	0.9651	20.5	0.8944
2020–2021	17,520	0.9612	17.6	0.9087
2022–2023	17,520	0.7213	2.1	0.8652
2024–2025	16,594	0.9402	11.7	0.9231
Note: Coin-specific estimates use pooled data, time-period splits are pooled across coins.				