

## Stablecoins under stress in a national economy: transaction-level evidence from Austrian crypto-asset service providers

Questa è la versione preprint della seguente opera:

*Original*

Stablecoins under stress in a national economy: transaction-level evidence from Austrian crypto-asset service providers / Saggese, P., Sigmund, M., Raunig, B., Segalla, E., Haslhofer, B., Makridis, C.. - (2026). [10.48550/arXiv.2607.08524]

*Availability:*

This version is available at: 20.500.11771/43360

*Publisher:*

*Published*

DOI:10.48550/arXiv.2607.08524

*Terms of use:*

This publication is made accessible in accordance with the terms for deposit in the institutional repository, as defined by the IMT School for Advanced Studies Lucca's Open Access Policy. ([https://library.imtlucca.it/sites/default/files/regolamento-policy-open-access-imtlib\\_0.pdf](https://library.imtlucca.it/sites/default/files/regolamento-policy-open-access-imtlib_0.pdf)).

Si prega di consultare le pagine informative dell'editore relative alle politiche di autoarchiviazione.

(Article begins on next page)

# Stablecoins under Stress in a National Economy: Transaction-Level Evidence from Austrian Crypto-Asset Service Providers

Pietro Saggese<sup>\*1,2</sup>, Michael Sigmund<sup>3</sup>, Burkhard Raunig<sup>3</sup>, Esther Segalla<sup>3</sup>, Bernhard Haslhofer<sup>2</sup>, and Christos A. Makridis<sup>2,4</sup>

<sup>1</sup>IMT Scuola Alti Studi Lucca

<sup>2</sup>Complexity Science Hub

<sup>3</sup>Oesterreichische Nationalbank (OeNB)

<sup>4</sup>Arizona State University

July 10, 2026

## Abstract

Cryptoassets are increasingly entangled with the traditional financial system, and how this activity integrates into national economies and behaves under stress bears on financial stability and the design of public digital money. However, blockchain pseudonymity and the lack of geographic identifiers force existing work to rely on indirect proxies to infer and locate market participants. Here we use a regulatory registry that directly identifies the on-chain addresses of all crypto-asset service providers (CASPs) registered in Austria, reconstructing their on-chain transaction activity across Bitcoin, Ether, USDC, and USDT through May 2025, and separating retail-like from institutionally mediated flows. We find that Austrian CASPs intermediate roughly \$30 billion with external counterparties and are integrated globally rather than domestically. In value, this activity is dominated by a few institutional counterparties; in number, by retail-like ones. Around three major shocks, the Terra-Luna collapse, the FTX bankruptcy, and the Silicon Valley Bank failure, the two groups respond through different mechanisms, and stablecoins do not act as a uniform safe haven. The clearest case is SVB, where retail-like deposits and institutional withdrawals are consistent with USDC's two-tiered redemption mechanism. These patterns are invisible in aggregate data. Registry-based, transaction-level measurement thus offers a reproducible, cross-jurisdictional basis for monitoring how cryptoasset markets transmit risk.

**Keywords:** Cryptoasset, Stablecoin, Crypto Asset Service Providers, Blockchain, Bitcoin, Ethereum, DeFi, Silicon Valley Bank, FTX

**JEL codes:** C81, E42, E58, F31, G19, G23, O33

---

<sup>\*</sup>Corresponding author. E-mail: [pietro.saggese@imtlucca.it](mailto:pietro.saggese@imtlucca.it)

# 1 Introduction

How does the demand for digital assets respond to macro-financial shocks? The question is central to financial stability, monetary transmission, and the design of public digital money, yet it has resisted clean empirical answers. Cryptoassets are now a major asset class, with a total market capitalization of roughly \$2.2 trillion mid-2026,<sup>1</sup> and their growing entanglement with the traditional financial system, through stablecoin reserves, custodial intermediaries, and institutional participation, makes their behavior under stress a first-order concern for regulators and central banks.

However, measuring how these markets interact with national financial systems faces a binding data constraint: blockchain transactions are pseudonymous, public networks can generate fresh addresses by design, and on-chain activity carries no geographic identifier. Existing work therefore infers jurisdiction-level activity from indirect proxies such as web traffic, IP geolocation, off-chain exchange data, or survey indices (Auer et al., 2025; Parino et al., 2018; Von Luckner et al., 2023; Makridis, 2025; Abramova et al., 2023), all of which are noisy, partial, and ill-suited to fine-grained behavioral analysis. Basic questions thus remain open: how large crypto activity is within an economy, how individuals reallocate across assets and custody structures under stress, and whether stablecoins serve the safe-haven function their issuers claim.

In this paper, we address this gap by measuring, at the transaction level, how the cryptoasset activity intermediated by all licensed crypto-asset service providers (CASPs) in a national jurisdiction is structured and how it responds to stress. Our approach exploits a feature of Austrian financial reporting: under existing regulation, crypto-asset service providers must report the full set of blockchain addresses they control to the competent authorities (FMA, 2025). From the resulting pseudonymized address-to-entity registry, we reconstruct the complete on-chain activity of the 12 CASPs registered in Austria at the end of 2024, across Bitcoin (BTC), Ether (ETH), and the two largest stablecoins, Tether (USDT) and USD Coin (USDC), through May 2025. Because ownership is established directly from the registry rather than inferred through clustering heuristics, we observe these flows, 11,980,687 asset transfers totaling \$50.4 billion, with a reliability that proxy-based country-level studies cannot achieve.

We document four main findings. First, Austrian CASPs intermediate substantial economic activity, roughly \$30 billion in on-chain flows with external counterparties, but the domestic intermediation is economically negligible: flows between Austrian CASPs total just \$5.8 million, so the ecosystem is integrated globally rather than domestically. A further 41.1% of total volume reflects within-entity bookkeeping rather than counterparty activity, and would be misclassified as economic activity in aggregate on-chain data. Second, external activity is highly concentrated: a small set of counterparties transacting through CASP warm and cold wallets account for 57.8% of external volume while representing only 0.77% of external transfers, whereas about 98% of addresses on each chain exhibit low-frequency, retail-like behavior. Third, the response of retail-like and institutional actors to systemic shocks is substantially different. After the FTX collapse, retail-like withdrawals increased significantly and persistently, consistent with reduced custodial

---

<sup>1</sup><https://coinmarketcap.com/charts/>, accessed June 2026.

exposure following a counterparty failure, while institutional flows reflect liquidity management and cross-platform rebalancing rather than a one-way exit. The Silicon Valley Bank (SVB) failure produced a narrower response concentrated in USDC, the stablecoin with direct Circle reserve exposure, and did not propagate to BTC and ETH. The Terra-Luna crash generated broader but less persistent reallocation across assets. Fourth, and contrary to the stablecoin-as-safe-haven hypothesis, we find little evidence that stablecoins act as a systematic within-crypto safe haven during stress: rather than a uniform inflow, their responses were two-sided and shock-specific. This was clearest at SVB, where large retail-like deposits and institutional withdrawals are consistent with USDC’s asymmetric redemption mechanism, which restricts par redemption to institutions and forces retail investors to trade below par in secondary markets.

Our paper primarily contributes to a growing literature on the integration of digital assets into national economies. Empirical work has largely measured cryptoasset activity globally (Ron and Shamir, 2013; Makarov and Schoar, 2021) or approximated country-level flows through indirect proxies: IP-address relay data (Parino et al., 2018), exchange geolocation and web traffic (Feyen et al., 2022; Auer et al., 2025), off-chain fiat trading pairs in cryptocurrency exchanges (Von Luckner et al., 2023), and survey-based adoption measures (Makridis, 2025; Abramova et al., 2023). Most studies link pseudonymous addresses to real-world actors through clustering heuristics and attribution techniques (Meiklejohn et al., 2013). Saggese et al. (2024), for instance, evaluated Austrian CASP balance-sheet adequacy prior to the adoption of the new regulatory framework. Moreover, direct-observation studies remain rare. Alvarez et al. (2023) survey Bitcoin adoption in El Salvador after its 2021 legal-tender designation, and Hu et al. (2021), Cardozo et al. (2024), and Cerutti et al. (2024) use indirect attribution to study Bitcoin flows in China, Brazil, and globally, respectively. None of these studies identifies ownership directly, nor observes the full population of licensed crypto-asset service providers in a national jurisdiction at the transaction level. We do so for Austria, covering every CASP licensed at the end of 2024 and tracing their complete transaction history through May 2025 across Bitcoin, Ether, Tether, and USD Coin. To our knowledge, this is the first study of this kind.

Our paper also contributes to an active macro-finance debate over whether stablecoins function as safe-haven assets (Baur and McDermott, 2010).<sup>2</sup> Two strands of evidence reach opposing conclusions. The flight-to-quality view (Anadu et al., 2024; Oefele et al., 2024b) documents that aggregate stablecoin flows during stress episodes, like the March 2023 SVB failure, mirror the prime-to-government money-market fund (MMF) reallocation patterns familiar from traditional finance, sharing structural vulnerabilities to run risk (Oefele et al., 2024a; Gorton and Zhang, 2023; Awrey, 2020). The contrary view (Aldasoro et al., 2025) finds that stablecoins respond to crypto-specific and monetary-policy shocks differently from MMFs and concludes that they do not consistently serve as safe havens. A related literature reaches similarly mixed findings on

---

<sup>2</sup>In the broader literature on safe assets and convenience yields, flights-to-safety involve reallocation toward assets providing liquidity, collateral, and protection in bad states, especially government debt and money-like instruments (Baele et al., 2020; Krishnamurthy and Vissing-Jorgensen, 2012). These properties are not intrinsic to dollar denomination alone, but depend on institutional credibility, market depth, liquidity, and state-contingent hedging value (Caballero et al., 2008; He et al., 2019; Jiang et al., 2021). Stablecoins are an analogous but distinct case: dollar-denominated private money-like claims whose safe-asset properties depend on reserve composition, redemption rights, issuer credibility, and custodial structure.



whether stablecoins act as safe havens against other cryptoassets (Baur and Hoang, 2021; Xie et al., 2021; Wang et al., 2020; Baumöhl and Vyrost, 2020; Kołodziejczyk, 2023), equity markets (Conlon et al., 2020; Feng et al., 2024), and geopolitical risk (Dionysopoulos et al., 2025; Mo et al., 2025; Adediran et al., 2024; Belguith et al., 2024). A common constraint of these studies is reliance on aggregate price or flow data, which cannot distinguish heterogeneous responses.

We extend this strand of literature in three ways. First, we use address-level on-chain transactions rather than aggregate flows. Second, we study three distinct shock types: an algorithmic stablecoin de-pegging (Terra), a centralized exchange insolvency (FTX), and a banking-sector failure with direct stablecoin reserve exposure (SVB). Third, we find no evidence supporting the safe-haven hypothesis: Terra led to broad reallocation across groups and assets; FTX increased persistent retail-like withdrawals to self-custodial wallets; USDC retail-like deposits and institutional withdrawals increased after SVB’s collapse. In all shocks, the flight-to-quality reallocation predicted by Anadu et al. (2024) and Oefele et al. (2024b) is not visible. These findings bear directly on debates about whether stablecoins can compete with public digital monies under stress (Dionysopoulos and Makridis, 2026).

Finally, our paper contributes to the literature on retail versus institutional behavior in cryptoasset markets. Existing work on crypto investor characteristics relies primarily on brokerage account data (Hackethal et al., 2022) or survey-based behavioral measures (Almeida and Gonçalves, 2023; Fessler and Weber, 2024), neither of which can be matched to on-chain transactions. We import retail-identification techniques from fixed-income (Schultz, 2012; Bessembinder et al., 2020; O’Hara and Zhou, 2021; deHaan et al., 2023) and equity microstructure (Boehmer et al., 2021; Barber et al., 2024). Because Hund et al. (2025) show that the standard \$100,000 trade-size cutoff misclassifies a substantial share of trades, we replace the single threshold approach with two structural signals: warm/cold wallet contact and chain-wide transaction activity. The resulting classification feeds an event-study framework (Kothari and Warner, 2007) that builds on crypto-specific precedents (Abramova and Böhme, 2021; Marmora, 2022; Yousaf et al., 2023; Liu et al., 2023; Jalan and Matkovskyy, 2023).

Consistent with previous studies, retail and institutional actors respond to systemic shocks through different margins. Retail activity dominates the transaction-count response, while institutional activity dominates the volume response; FTX led to persistent retail-like reallocations to self-custodial wallets, but large institutional deposits consistent with cross-platform rebalancing; the Circle redemption mechanism documented by Ma et al. (2025) provides instead one behavioral explanation for the differences across groups after SVB: retailers cannot redeem USDC at par directly with the issuer and thus sell on secondary markets at a loss. This heterogeneity is invisible in aggregate flows. These results also speak to the literature on macroeconomic narrative transmission (Mertens et al., 2020; Ter Ellen et al., 2022; Binder, 2025; Feldkircher and Makridis, 2025): responses concentrate around shocks with simple, asset-relevant narratives and may be muted for technically complex events.

The findings carry implications for several stakeholders. For supervisors, the registry-based approach is a systematic and reproducible alternative to proxy- and heuristic-based measurement, pointing toward collecting address and counterparty data directly from supervised CASPs,

and could form the basis for a cross-jurisdictional observatory of cryptoasset flows. For policymakers, the limited and two-sided stablecoin response during stress bears on whether privately issued stablecoins can replicate the stabilizing role of traditional safe assets and how they will interact with public digital currencies. For CASPs, abnormal flows appear within days of each shock and with no jurisdictional lag, which argues for ex-ante liquidity and redemption safeguards rather than reactive ones. For retail users, episodes like SVB expose a structural disadvantage: when a peg is threatened, par redemption is reserved for institutional accounts, leaving retail holders to sell on secondary markets, often at a loss (Ma et al., 2025).

The remainder of the paper proceeds as follows. Section 2 reviews the related literature and background. Section 3 describes the dataset and how we select the events. Section 4 documents the structure of Austrian CASP activity and develops the retail and institutional classification, and Section 5 reports the event-study results for the Terra-Luna, FTX, and SVB shocks. Section 6 discusses the findings and their implications, and Section 7 concludes.

## 2 Background and Related Work

Measuring the on-chain financial activity intermediated by CASPs and customer responses to shocks requires connecting four threads: how CASPs custody funds on-chain and why their flows are observable yet hard to attribute (Section 2.1); the specific redemption mechanisms of fiat-backed stablecoins and the unresolved question of their role as a safe haven against other cryptoassets (Section 2.2); how investors respond selectively to interpretable, socially mediated shocks (Section 2.3); and how retail activity can be isolated when identity is not observable or under pseudonymity (Section 2.4).

### 2.1 CASPs and Their On-Chain Footprint

Crypto-Asset Service Providers (CASPs) are defined under the European Union’s Markets in Crypto-Assets (MiCA) Regulation as entities that provide one or more cryptoasset services such as custody, order execution, trading, portfolio management, and advisory services, on a professional basis (MiCA, 2023). They operate as centralized intermediaries connecting the cryptoasset and traditional financial systems, and include crypto exchanges, custodial services, and payment providers (Saggese et al., 2024).

A CASP’s wallet architecture refers to the set of on-chain addresses or wallets<sup>3</sup> and the key management scheme through which the provider custodies client and operational funds. CASPs operating a custodial model record customer balances in an internal, off-chain ledger, while funds are held on-chain in a smaller set of provider-controlled wallets. They typically control multiple types of wallets serving different functions: *hot wallets* are connected to the internet for client deposits and withdrawals, while *warm* and *cold wallets* follow higher security standards and are employed for offline storage and large-value transactions. Figure 1 shows a stylized representation of a CASP wallet architecture.

---

<sup>3</sup>In the following, we use the terms *wallet* and *address* synonymously.




Figure 1: **CASP on-chain wallet architecture.** CASPs use hot wallets for daily deposit and withdrawal operations, and warm and cold wallets for secure storage and large-value transfers. Typical custody patterns are: (i) users interact with one main (hot) wallet, either directly or through so-called deposit (hot) wallets; this wallet collects funds, manages daily operations, and occasionally interacts with internal secure wallets for rebalancing; (ii) a few large transactions occur directly with CASP’s warm or cold wallets.

CASPs are linked to institutions in the traditional financial system (e.g., banks) through fiat accounts. This allows customers to transfer funds denominated in a fiat currency (e.g., EUR, USD) from their bank accounts to a CASP and vice versa. Such transactions are not publicly visible and are recorded only in the CASP’s internal ledger.

Transfers to other entities in the cryptoasset ecosystem are executed on-chain and are therefore publicly visible asset transfers between addresses. Consequently, and this is the focus of this paper, it is in principle possible to measure on-chain cryptoasset flows *to* and *from* CASPs.

A major challenge lies in linking on-chain addresses to the CASP controlling the underlying private keys. This deanonymization task, commonly known as *attribution*, currently relies mainly on so-called clustering heuristics (Meiklejohn et al., 2013), which, by definition, produce false positives and false negatives (Fröwis et al., 2020). Blockchain transactions also carry no geographic information, so localizing on-chain financial activity directly is hardly possible (Auer et al., 2025). As a consequence, the role of CASPs as intermediaries integrating digital assets into domestic financial systems, and the implications for financial stability, remain understudied.

Since only the CASPs themselves know the addresses they control, this information is best retrieved directly from them. Where a responsible authority within a jurisdiction can compel its disclosure, at least partial geographic bindings become possible. This is the approach we are pursuing in this work.

## 2.2 Stablecoins: Redemption and Safe-Haven Dynamics

Stablecoins play a pivotal role in Decentralized Finance (DeFi) and crypto markets as an alternative asset class (Auer et al., 2024). By maintaining a stable value through reserved assets or algorithms, they offer a solution to the notorious volatility of cryptoassets like Bitcoin and Ethereum (Catalini et al., 2022). A large body of work examined their design choices, price stability, and impact on traditional and crypto markets, summarized in two recent reviews (Ante et al., 2023; Dionysopoulos and Urquhart, 2024).

Fiat-backed stablecoins like USDT and USDC, which are the focus of this paper, are issued by private firms (e.g., Tether, Circle) and maintain their dollar peg by promising a 1:1 (par-value) redemption and full collateralization with highly liquid fiat reserves. Thus, stablecoin owners hold the right to redeem one unit of stablecoin for one unit of fiat currency at any time. In practice, however, direct redemption is restricted to only a small set of vetted institutional counterparties, subject to KYC/AML screening, minimum thresholds, fees, and jurisdictional limits. Retail holders can interact only indirectly through secondary markets, where the price is set by supply and demand (Ashwanth et al., 2025; Ma et al., 2025). This produces a two-tiered structure in which the primary-market peg can hold for institutional actors even as the secondary-market peg can break for other participants, and the primary peg holds only as long as the issuer remains solvent and able to convert reserves into fiat quickly.

Stablecoins are thus exposed to known risks, such as the classic run risk on private money if confidence in their backing falters, and represent a channel for the transmission of shocks across the crypto and traditional finance systems (Briola et al., 2023; Griffin and Shams, 2020; Gorton and Zhang, 2023). Whether these risks are outweighed by their stability is at the heart of a debated policy question: *do stablecoins serve as a safe haven for investors during periods of market stress and extreme price movements?* Aggregate flows suggest a safe-haven role relative to Bitcoin and occasional flight-to-quality behavior similar to MMFs (Baur and Hoang, 2021; Anadu et al., 2024; Oefele et al., 2024b), but more recent evidence reports heterogeneous and inconsistent dynamics across events (Aldasoro et al., 2025), or contrasting evidence (see e.g. Wang et al., 2020; Kołodziejczyk, 2023; Feng et al., 2024). Existing research thus remains mixed and inconclusive.

## 2.3 Shocks, Social Media, and Investor Behavior

Individuals respond selectively when adjusting cryptoasset holdings: research on central bank communication offers a useful framework for explaining why. Most households do not process primary policy information directly but update their beliefs through media intermediaries that translate technical developments into salient narratives; when communication is complex, attention and behavioral responses weaken, and individuals substitute toward simplified interpretations supplied by journalists and platform-level curators (Mertens et al., 2020; Ter Ellen et al., 2022; Binder, 2025; Feldkircher and Makridis, 2025). Responses thus depend not only on the magnitude of a shock but also on whether it is interpretable.

Investors may therefore rebalance primarily in response to shocks that lend themselves to simple, asset-relevant narratives (Shiller, 2017; Jiang et al., 2022). Monetary policy surprises, inflation releases, and episodes of financial stress are more readily interpreted than shocks requiring technical expertise, such as changes in term premia or nuanced regulatory guidance, which rarely trigger widespread rebalancing. Media thus operates not merely as a conduit for shocks but as a filter determining which events become investor signals.

Although this evidence concerns traditional markets, the underlying mechanism is domain-general, and arguably it applies to a narrative-sensitive market such as crypto, where investors rely heavily on social media. Responsiveness to these narratives is in fact heterogeneous and stronger among younger and more educated individuals, consistent with their lower marginal cost of acquiring and processing information (Philippas et al., 2019; Li et al., 2025); rebalancing therefore concentrates among users embedded in fast-moving media ecosystems and attuned to narrative shifts. In crypto, these ecosystems are shaped both by traditional journalists and by Key Opinion Leaders (KOLs), community figures whose narratives carry material consequences for followers’ financial decisions and whose influence rests on perceived credibility rather than formal credentials (Kropiunig et al., 2026). Which shocks become investor signals thus depends in part on how such intermediaries interpret and transmit them. Selective responsiveness can thus be read as an equilibrium outcome of attention allocation, narrative production, and interpretability, rather than simple inattention or irrationality. Despite its plausibility, this channel has received little attention in the empirical literature on crypto, which has examined how specific events affect prices and demand (Abramova and Böhme, 2021; Marmora, 2022; Yousaf et al., 2023) without asking which shocks become salient, or for whom.

## 2.4 Identification of Retail Activity

Because blockchain activity is pseudonymous and a single entity can control multiple addresses, investor identity cannot be directly observed (Möser et al., 2013). Two strands of literature have developed standard methods for identifying retail trades under similar conditions in traditional finance. In municipal and corporate bond markets, the convention is a trade-size threshold: trades below \$100,000 par are treated as retail and trades above \$1 million as institutional (Schultz, 2012; O’Hara and Zhou, 2021; Bessembinder et al., 2020; deHaan et al., 2023). In U.S. equities, Boehmer et al. (2021) flag retail orders by the sub-penny price improvements that wholesalers apply when internalizing retail order flow, though Barber et al. (2024) document non-trivial misclassification in validation tests. Both reduce investor type to a single observable proxy, but this approach is fragile: using a structural model with latent investor type on municipal trades from 2005 to 2023, Hund et al. (2025) find that the \$100,000 cutoff separates informed from uninformed trades poorly, with roughly half of uninformed (retail-like) trades occurring above the threshold and sophisticated investors routinely placing retail-sized orders.

Comparable strategies have been applied to cryptoasset markets. Divakaruni and Zimmerman (2024) exploit trade-size thresholds, reflecting the assumption that retailers place smaller trades than professional investors; Baur et al. (2019) infer investor type from time-dependent

trading patterns; and John and Li (2025) proxy retail activity via aggregate user counts on trading platforms such as RobinHood. Other work relies on brokerage account data (Hackethal et al., 2022) or survey-based behavioral measures (Almeida and Gonçalves, 2023; Fessler and Weber, 2024), neither of which can be matched to on-chain transactions. Across domains, investor type is recovered through indirect proxies rather than observed directly.

Our setting shares the central constraint of this literature. While the CASP-controlled addresses in our data are attributed through disclosure rather than heuristics, the individual counterparties transacting with them remain pseudonymous and carry no strong identity information. Investor type therefore cannot be observed directly and must be inferred from observable transaction characteristics. We follow the established convention by relying on trade-size proxies and extend it via an additional structural signal - warm/cold wallet interaction - to distinguish retail from institutional activity, while accounting for the misclassification documented above.

## 3 Data

### 3.1 Address Reporting and CASP Identification

As motivated in Section 2.1, the set of addresses a CASP controls is reliably known only to the provider itself and is best obtained directly, wherever an authority can compel its disclosure. Austria’s reporting requirements establish such an obligation: crypto-asset service providers must report the blockchain addresses under their control to the competent financial authority (Austrian Financial Market Authority or FMA, 2025). Our analysis builds on the resulting registry of address-to-entity mappings, comprising pseudonymized Bitcoin and Ethereum addresses reported by the 12 Austrian CASPs registered with the FMA at the end of 2024.

The registry is pseudonymized to protect the privacy of both reporting entities and individual address holders: each service is represented solely by the set of addresses it controls, with no disclosure of user identities or other sensitive information. Because these mappings originate from a mandatory regulatory report rather than from voluntary disclosures or clustering heuristics, they avoid the false positives and false negatives inherent to heuristic attribution and can be treated as highly trustworthy data for the population of registered Austrian providers.<sup>4</sup> We therefore do not apply any address clustering to expand the registry. To the best of our knowledge, this is the first study to draw on a dataset of this kind.

Further details on the registry, including the share of addresses by CASP and asset type, are provided in Appendix A. The exact number of addresses is withheld for confidentiality reasons, as agreed with the Austrian financial authorities.

---

<sup>4</sup>We note however that reporting obligations bind on the CASP, not on the underlying blockchain, so the registry can in principle omit addresses a provider controls but fails to disclose. In the course of the analysis we identified and reported a small number of addresses likely controlled by Austrian CASPs that were absent from the original mapping (see Section 6.3 for further discussion).

### 3.2 Asset Flow Extraction

From this registry, we reconstruct the on-chain activity of the reported wallets. We parse the full Bitcoin and Ethereum ledgers and extract every transaction directed to or originating from a registered address, up to a cutoff date of 31 May 2025. We restrict attention to four major cryptoassets: Bitcoin (BTC) and Ether (ETH), the two largest by market capitalization, and Tether (USDT) and USD Coin (USDC), the two largest fiat-backed stablecoins by market capitalization. This yields 11,980,687 asset transfer records, corresponding to 3,214,643 unique transactions. To express values in a common currency, we convert each transfer using daily USD prices retrieved from the Coingecko API. The total volume routed through Austrian CASP wallets exceeds \$50.4 billion, distributed across assets as 65.53% BTC, 15.90% ETH, 15.47% USDT, and 3.10% USDC.

We note that our unit of observation is on-chain transfers to and from CASP-controlled addresses. We do not observe customers’ off-chain ledger balances, the trades executed on a CASP’s platform, or the ultimate disposition of withdrawn funds. Deposits and withdrawals are therefore best read as changes in on-chain custodial exposure rather than as direct evidence of trading decisions, redemption, or final asset custody. Where we describe a flow pattern as consistent with selling, redemption, or self-custody, we mean that the observed on-chain movement matches what such behavior would produce, not that we directly observe the underlying trade, redemption, or custody outcome.

### 3.3 Event Selection

We next identify events with a broad impact on the crypto ecosystem, restricting attention to crypto-related incidents, hacks, failures, and bankruptcies. As argued in Section 2.3, individuals respond most strongly to shocks that lend themselves to simple narratives and attract social media coverage; the practical consequences of such events are easy to grasp and communicate, even when the underlying technical mechanics are not, and they typically receive substantial media attention.

Starting from a published repository of 1,141 crypto crime events (Carpentier-Desjardins et al., 2025), supplemented with data from DeFiLlama and Coinmarketcap, we select 19 events that (i) occurred between 2019 and 2025, when all four assets were actively traded; (ii) caused estimated losses exceeding \$100 million; and (iii) affected CEX, DeFi, or centralized finance (CeFi) entities. Using Google Trends search volumes as a proxy for public attention, queried with event-specific keywords across two geographical areas, Austria and worldwide, three events stand out consistently: the May 2022 collapse of the algorithmic stablecoin Terra and its Luna token; the November 2022 solvency crisis of the crypto exchange FTX; and the March 2023 run on Silicon Valley Bank (SVB), a lender closely tied to the crypto sector. Table 1 provides a detailed timeline drawn from prior studies (Briola et al., 2023; Liu et al., 2023; Conlon et al., 2023; Akyildirim et al., 2023; Galati and Capalbo, 2024; Metrick, 2024). Further details on the selection procedure are reported in Appendix A.

---

**Terra (UST) and Luna (LUNA) collapse**

---

|            |                                                               |
|------------|---------------------------------------------------------------|
| 2022-05-05 | Selling pressure on LUNA aimed at destabilising markets       |
| 2022-05-07 | UST loses its peg to the dollar for the first time            |
| 2022-05-09 | UST loses its peg to the dollar for the second and final time |
| 2022-05-11 | Final attempt by Terra governance to restore the peg          |

---

**FTX bankruptcy**

---

|            |                                                                                       |
|------------|---------------------------------------------------------------------------------------|
| 2022-11-02 | Media report that Alameda, FTX’s sister firm, is heavily exposed to the FTX token FTT |
| 2022-11-06 | Binance announces sale of its FTT holdings                                            |
| 2022-11-08 | FTT price falls below the critical \$22 level                                         |
| 2022-11-09 | Binance withdraws from tentative deal to acquire FTX                                  |
| 2022-11-10 | Regulators freeze FTX Digital Markets assets                                          |
| 2022-11-11 | FTX and Alameda file for bankruptcy                                                   |

---

**Silicon Valley Bank failure**

---

|            |                                           |
|------------|-------------------------------------------|
| 2023-03-08 | SVB announces a \$1.8 bn loss             |
| 2023-03-10 | SVB fails; FDIC takes control             |
| 2023-03-11 | Circle discloses \$3.3 bn exposure to SVB |
| 2023-03-12 | FDIC guarantees uninsured depositors      |

---

Table 1: **Timeline of the three selected events:** the Terra–Luna collapse (May 2022), the FTX bankruptcy (November 2022), and the Silicon Valley Bank failure (March 2023).

## 4 Cryptoasset Flow Analysis

This section characterizes the on-chain activity intermediated by Austrian CASPs, from aggregate flows down to individual counterparties. We first decompose total volume in three components, separating market-facing activity from inter-CASP and internal flows (Section 4.1). We then examine transaction-level structure, identifying the recurring patterns through which CASPs route customer funds and the wallet tiers they use (Section 4.2). Finally, we classify the external counterparties that interact with Austrian CASPs into retail-like and institutional, using a proxy-based heuristic that combines wallet-tier interaction with overall on-chain activity (Section 4.3). These steps establish who transacts with Austrian CASPs, in what volumes, and through which channels, providing the basis for the event-study analysis that follows.

### 4.1 Decomposing Flows by Counterparty

We decompose transaction volumes into three components according to the counterparties involved. *CASP-to-market* activity comprises all transactions between an Austrian CASP and an external (non-Austrian-CASP) actor. *CASP-to-CASP* activity isolates transactions in which the sender address belongs to one Austrian CASP and the recipient to another Austrian CASP. *Within-CASP* activity captures transactions in which sender and recipient are controlled by the same Austrian CASP.

**CASP-to-Market** Separating CASP-to-Market flows from internal bookkeeping isolates the economic activity that Austrian CASPs conduct with the wider market, which amounts to roughly \$30 billion in total.



Austrian CASPs operate primarily in Bitcoin and Ether. Figure 2 shows the number of asset transfers over time, broken down by asset and month. Bitcoin has played a major role since 2016, while Ether and stablecoins first appeared in CASP activity in 2017 and 2019, respectively. Activity peaked around 2021 and the Bitcoin share fell relative to other assets thereafter, whereas stablecoin activity grew over the full period. Figure 3 presents cumulative inflows and outflows by asset, respectively in the left and right panels, and shows that total inflows amount to \$15.1 billion, while the total outflows to \$14.6 billion.

CASP counterparties are highly heterogeneous: both transaction sizes and activity levels follow heavy-tailed distributions. Transaction values (aggregated by unique transaction and separated by asset) are strongly right-skewed, with median values of only a few hundred dollars (\$206 for ETH to \$544 for BTC) and 95th percentiles from roughly \$5,000 (ETH) to \$21,000 (BTC). Aggregate volumes are driven by a few very large transactions: the largest Bitcoin transaction (\$287.5M) exceeds the largest transaction in any other asset (\$13.5M–\$20.8M) by an order of magnitude. Activity is similarly concentrated: most addresses transact fewer than ten times, with the corresponding bin holding 95–98% of addresses across all assets (see Table 2); only a small fraction are highly active, a pattern that holds across assets and both blockchains.

**CASP-to-CASP** CASP-to-CASP activity captures volume flows between distinct Austrian CASPs, identified as transactions whose sender and recipient addresses are associated with two different Austrian CASPs. We find that such domestic intermediation is limited in both frequency and volume: just 3,040 transfers, totaling \$5.774 million. Most CASP-to-CASP flows are in Bitcoin, though its share has declined over time; transaction counts peaked in 2021, while transferred value was highest in 2022 and 2023.

These magnitudes are economically negligible relative to the \$30 billion of CASP-to-Market activity, indicating that financial integration among Austrian CASPs is marginal and that the network is instead integrated primarily with the global market. This is consistent with the relatively small size of the Austrian market. Appendix B reports CASP-to-CASP activity broken down by asset and year, together with quarterly CASP-to-CASP interaction networks.

**Within-CASP** Within-CASP activity captures asset flows between a CASP’s own hot, warm, and cold wallets. We identify 798,196 such interactions, which account for only 6.7% of all asset transfers but ~\$20.7 billion, or 41.1% of total volume. Internal bookkeeping thus involves comparatively few transfers, each moving large amounts on average. Makarov and Schoar (2021) report that within-entity transfers make up around 90% of global Bitcoin on-chain volume. The two figures are not directly comparable,<sup>5</sup> but both indicate that within-entity activity accounts for a substantial share of total on-chain volume.

---

<sup>5</sup>Makarov and Schoar (2021) consider a broader set of global entities, including the largest centralized exchanges, focus exclusively on Bitcoin, and aggregate within-entity volumes together with peeling chains and change-of-address transactions.




Figure 2: **Number of on-chain asset transfers by month and asset.** Bitcoin plays a predominant role; stablecoins first appear in CASP activity in 2019. Activity on the Ethereum chain grew steadily from 2017 on.




Figure 3: **Cumulative deposits and withdrawals over time by asset.** Panels (a) and (b) respectively show the cumulative inflows directed to known CASP addresses and outflows for BTC, ETH, USDT, and USDC. Values are denominated in US dollars at historical prices.

| Asset | Number of transactions |              |                |                |                |             |
|-------|------------------------|--------------|----------------|----------------|----------------|-------------|
|       | $[1, 10)$              | $[10, 10^2)$ | $[10^2, 10^3)$ | $[10^3, 10^4)$ | $[10^4, 10^5)$ | $\geq 10^5$ |
| BTC   | 1,259,786              | 57,560       | 2,833          | 398            | 64             | 7           |
| ETH   | 390,036                | 13,752       | 598            | 40             | 8              | 0           |
| USDT  | 134,123                | 4,135        | 140            | 24             | 0              | 0           |
| USDC  | 33,709                 | 951          | 38             | 1              | 0              | 0           |

Table 2: **Distribution of non-CASP addresses by transaction activity.** For each asset, addresses are grouped into bins by their total number of transactions, from  $[1, 10)$  up to  $\geq 10^5$ . Low-activity addresses dominate.

## 4.2 Transaction-level Analysis

Having examined aggregate flows, we now turn to the transaction level to identify how CASPs internally manage customer activity. Prior work shows that CASPs may use deposit addresses to credit customer funds to the correct account before forwarding them to a main exchange wallet under the same CASP’s control (Victor, 2020). Building on this, we develop tracing algorithms, adapted to heterogeneity across CASPs and blockchains, that proceed in two steps: we first construct network-based indicators (degree, betweenness, PageRank, eigenvector centrality) to detect hub wallets within a CASP, then match flows by transaction identifier, asset, and transferred value to trace chained movements across same-CASP addresses that terminate at a hub wallet. We restrict the analysis to CASPs with sufficient within-CASP activity, excluding those controlling few on-chain addresses. This sample covers more than 94% of the Austrian BTC market, 91% ETH, 98% USDC, and 99% USDT.

We identify two patterns in transaction-level flows. *Pattern 1*, followed by 89% of analyzed transactions, routes funds either to a deposit hot wallet and immediately on to the main hot wallet, or directly to the main wallet; main wallets occasionally interact with the same CASP’s warm or cold wallets for rebalancing. This resembles the “fan-in” transactions that merge funds from many wallets into one (Baquer et al., 2016). *Pattern 2* is a distinct mode in which a restricted set of market-controlled addresses interacts directly with the warm or cold wallets of Austrian CASPs. It represents only 0.75% of asset records yet accounts for over 55.8% of volume directed from external addresses to Austrian CASPs.<sup>6</sup>

Table 3 reports statistics for the wallet-tier interaction on the full sample of CASP-to-Market interactions: for every asset, the large majority of interactions target hot wallets (99.23% overall), yet those reaching warm or cold wallets account for more than 57.8% of total volume. Taken together, these findings suggest that interactions are concentrated among a small set of actors and are associated with high-value or operational flows, whereas routine activity is mediated through hub wallets. Appendix B describes the tracing algorithms, shows representative within-CASP networks and per-asset breakdowns, and documents minor additional patterns among Bitcoin-operating CASPs.

## 4.3 Separating Retail-like from Institutional Activity

The preceding analysis showed that activity across external addresses is highly skewed: a few addresses account for most volume and transactions, while the remainder interact with Austrian CASPs only a handful of times, typically for a few hundred to a few thousand dollars. We ask whether the body of this distribution corresponds to retail-like on-chain activity and the head to institutional counterparties such as other CASPs, market makers, custodians, and OTC desks.

Standard methods identify retail trades through a single observable proxy (Section 2.4); such a classifier misassigns addresses that resemble retail on that one margin. We therefore combine

---

<sup>6</sup>*Pattern 1* and *Pattern 2* do not cover the full sample; the remaining  $\sim 10\%$  could not be classified into a specific pattern. We also note that while addresses in the second group can interact with CASPs through the first identified pattern, we do not observe this for addresses in the first group.

| Asset | Wallet type | Transfers  | Volume (\$M) | % of volume |
|-------|-------------|------------|--------------|-------------|
| BTC   | hot         | 9,149,248  | 8,796        | 52.7        |
|       | warm/cold   | 39,844     | 7,887        | 47.3        |
| ETH   | hot         | 1,464,273  | 2,612        | 53.3        |
|       | warm/cold   | 31,720     | 2,284        | 46.7        |
| USDC  | hot         | 87,857     | 263          | 23.8        |
|       | warm/cold   | 1,866      | 841          | 76.2        |
| USDT  | hot         | 395,336    | 872          | 12.4        |
|       | warm/cold   | 12,347     | 6,157        | 87.6        |
| Total | hot         | 11,096,714 | 12,543       | 42.2        |
|       | warm/cold   | 85,777     | 17,170       | 57.8        |

Table 3: **CASP-to-Market interactions by wallet tier and asset.** Most asset transfers involve hot wallets, regardless of chain and asset; the few interactions with warm or cold wallets account for a large share of total volume.

two signals. The first is wallet-tier interaction: warm/cold flows originate almost exclusively from large counterparties (0.77% of external transfers, 57.8% of external volume), so an address interacting with a warm or cold wallet is unlikely to be retail. The second is overall on-chain activity, measured by transaction count rather than dollar volume to remove denomination effects (Makarov and Schoar, 2021; Auer et al., 2025); retail addresses fall in the body of the chain-wide count distribution, whereas exchanges, market makers, and other CASPs occupy the upper tail.

#### 4.3.1 Definitions and Assumptions

Our classification approach is proxy-based and heuristic. We do not observe whether a non-Austrian-CASP address belongs to a retail or an institutional actor; instead, we classify addresses by two observable proxies for institutional behavior. The signals are proxies for an unobserved actor type, and the thresholds that turn them into binary labels are conventional rather than theoretically derived. The definitions below should therefore be read as operational rules over these proxies, not as identification of actor types.

Let  $A_c$  denote the set of external (non-CASP) addresses<sup>7</sup> that interact with at least one Austrian CASP wallet on chain  $c \in \{\text{BTC}, \text{ETH}\}$ . For each element  $i \in A_c$ , we define two binary signals. The first,  $\text{cold}_i \in \{0, 1\}$ , equals 1 if  $i$  ever transacts with a warm or cold Austrian CASP wallet, and 0 otherwise. The second,  $\text{active}_i \in \{0, 1\}$ , flags whether  $i$  ranks among the most active addresses on chain  $c$ . Let  $n_i$  be the total number of transactions of  $i$  on chain  $c$ , and define its log-transformed count  $\ell_i = \log(1 + n_i)$ . Because transaction counts are heavy-tailed,

<sup>7</sup>The CASP addresses in our registry require no clustering; the external counterparties, however, are not in the registry and thus require clustering. Consistent with existing approaches (Makarov and Schoar, 2021; Capponi et al., 2026; Cong et al., 2025), we thus apply standard multi-input clustering heuristic to Bitcoin counterparties, while we treat each Ethereum counterparty address as a persistent entity under the account-based model and do not cluster. So each  $i \in \{\text{BTC}\}$  is a clustered set of addresses rather than a single address, and the resulting clusters are used to compute network statistics such as transaction count and degree.

spanning several orders of magnitude, we measure activity on the log scale, where the bulk of the distribution is approximately symmetric and a dispersion-based threshold is meaningful. With

$$\mu_c = \frac{1}{|A_c|} \sum_{j \in A_c} \ell_j, \quad \sigma_c = \sqrt{\frac{1}{|A_c| - 1} \sum_{j \in A_c} (\ell_j - \mu_c)^2},$$

the mean and standard deviation of log activity across  $A_c$ , we compute the z-score and set

$$\text{active}_i = \begin{cases} 1 & \text{if } z_i = \frac{\ell_i - \mu_c}{\sigma_c} > 3, \\ 0 & \text{otherwise,} \end{cases}$$

so that  $i$  is active when its log transaction count exceeds the chain-wide mean by more than three standard deviations. The threshold of three is the conventional outlier cutoff; we assess sensitivity to it in the robustness checks in Appendix C. We thus define:

**Definition 1** (Retail-like (RL)). An external address  $i$  is *retail-like* if  $\text{cold}_i = 0$  and  $\text{active}_i = 0$ . Such an address never interacts with a warm or cold wallet and is not among the most active addresses on its chain, the on-chain footprint expected of individual retail activity.

**Definition 2** (High-confidence institutional (HI)). An external address  $i$  is *high-confidence institutional* if  $\text{cold}_i = 1$  and  $\text{active}_i = 1$ . Both signals are present: the address transacts at the warm/cold tier and ranks among the most active on its chain, a combination characteristic of exchanges, market makers, or other large intermediaries.

**Definition 3** (Low-confidence institutional (LI)). An external address  $i$  is *low-confidence institutional* if  $\text{cold}_i + \text{active}_i = 1$ . We partition this group into  $\text{LI}_c$  ( $\text{cold}_i = 1$ ,  $\text{active}_i = 0$ ) and  $\text{LI}_a$  ( $\text{cold}_i = 0$ ,  $\text{active}_i = 1$ ). Exactly one signal is present: the address shows a single institutional marker, either warm/cold contact ( $\text{LI}_c$ ) or high activity ( $\text{LI}_a$ ), but not both.

Together, the HI,  $\text{LI}_c$ , and  $\text{LI}_a$  groups constitute institutionally mediated activity. The two signals are jointly informative because each recovers a way an institutional address can resemble retail on the other margin. An institutional counterparty with modest aggregate activity, such as an OTC desk or a custodian that transacts on-chain only when settling large flows, is not flagged by the activity signal but is recovered through warm/cold contact ( $\text{LI}_c$ ). A highly active intermediary that never reaches the warm/cold tier, such as a hot-wallet-only aggregator or arbitrageur, is not flagged by the wallet-tier signal but is recovered through the activity threshold ( $\text{LI}_a$ ). The two subgroups are therefore not residual cases but addresses that a single-signal classifier would have assigned to retail. The classification rests on two assumptions:

**Assumption 1** (Wallet-tier segmentation). The warm/cold wallet tier of Austrian CASPs serves high-value, low-frequency flows that originate predominantly from institutional rather than retail counterparties.

**Assumption 2** (Activity dispersion). Retail on-chain activity does not reach the upper tail of the chain-wide activity distribution; addresses in the upper tail are therefore not retail.

| Group           | Chain | Addresses | Avg. transactions |           | Avg. degree | Avg. USD  | Hot wallet<br>(fraction) |
|-----------------|-------|-----------|-------------------|-----------|-------------|-----------|--------------------------|
|                 |       |           | overall           | in sample | overall     | in sample |                          |
| RL              | ETH   | 534,062   | 29                | 2.6       | 11          | 1,568     | 1.00                     |
|                 | BTC   | 1,286,555 | 10                | 2.9       | 15          | 1,516     | 1.00                     |
| HI              | ETH   | 472       | 6,114,749         | 644       | 292,149     | 53,539    | 0.16                     |
|                 | BTC   | 82        | 5,044,178         | 25,425    | 2,636,324   | 80,604    | 0.53                     |
| LI <sub>c</sub> | ETH   | 552       | 88                | 4.8       | 25          | 72,170    | 0.03                     |
|                 | BTC   | 2,815     | 3.7               | 1.8       | 46          | 378,333   | 0.02                     |
| LI <sub>a</sub> | ETH   | 6,957     | 349,525           | 44        | 32,886      | 4,458     | 1.00                     |
|                 | BTC   | 31,196    | 17,479            | 108       | 10,749      | 975       | 1.00                     |

Table 4: **External addresses by group and chain.** Addresses are grouped into retail-like (RL), high-confidence institutional (HI), and low-confidence institutional (LI<sub>c</sub>, LI<sub>a</sub>), separately for each chain (Bitcoin, Ethereum).

Assumption 1 is supported by the evidence above that few transfers reach warm or cold wallets, yet those transfers carry a disproportionate share of volume (0.77% of external transfers, 57.8% of external volume), a profile incompatible with retail use. Assumption 2 follows Makarov and Schoar (2021), who identify exchanges and large Bitcoin entities through centrality- and activity-based heuristics.

Three limitations apply. First, the rule misclassifies any retail address that interacts with a warm or cold wallet under atypical circumstances; such cases are bounded by the small set of bookkeeping transactions in which CASPs route funds through these wallets (Table 5, Appendix A), whose volume is economically negligible. Second, a single actor’s activity may be fragmented across many addresses, which mechanically lowers per-address activity and biases classification *toward* retail rather than away from it; this bias operates in the same direction across all subsamples. Third, as a robustness check we re-classify under varying thresholds and using network degree centrality in place of transaction count, obtaining comparable group assignments (see Appendix C).

#### 4.3.2 Classification Results

Table 4 reports the four groups by chain. First, retail-like addresses dominate the count but not the value: they make up 98.6% (Ethereum) and 97.4% (Bitcoin) of external addresses, each transacting only a few times with Austrian CASPs, for roughly \$1,500 per transaction on average. The cross-chain symmetry is notable: despite the UTXO/account-based distinction between the two blockchains, the retail-like profile is the same.

Second, the HI tier is small and disproportionately large in value: 554 addresses across both chains, with overall transaction counts in the millions, network degree from the hundreds of thousands to the millions, and average per-transaction values between \$53,539 (ETH) and \$80,604 (BTC). A representative BTC-HI address transacts roughly 25,000 times with Austrian CASPs but about 5 million times on Bitcoin overall, a ratio consistent with exchanges, market makers, or large bridges for which Austrian CASPs are one node in a far larger network.




Figure 4: **Transaction count and volumes by group.** Retail-like activity is shown in orange, institutional is in blue. Panel (a): Transaction count. Panel (b): Volume flows.

Third, the two LI subgroups capture distinct profiles.  $LI_c$  addresses reach the warm/cold tier but are not highly active; they are large, infrequent counterparties consistent with OTC desks, custodians, or other CASPs operating at modest on-chain frequency, with average per-transaction values one to two orders of magnitude above retail-like addresses.  $LI_a$  addresses are highly active but never reach the warm/cold tier; they are consistent with the lightly regulated exchange intermediaries that Auer et al. (2022) term the “shadow crypto financial system”, plausibly including aggregators, arbitrageurs, or hot-wallet-only services, with small per-transaction values but large aggregate activity. That the two subgroups are populated by economically distinct entities is consistent with the value of the bivariate scheme: a single-signal rule would collapse them into one residual category.

Figure 4 aggregates the non-retail groups ( $HI + LI$ ) and plots daily counts and volumes on a log scale for retail-like and institutional addresses. In Panel (a), institutional activity leads in the early period (2016–2018), exceeding retail-like counts by up to an order of magnitude and peaking near  $10^4$  daily transactions at the end of 2017. From roughly 2019 the two series track each other closely, settling near  $10^3$  through 2022. After 2023 they separate again, but in the opposite order: retail-like counts hold near  $1.5 \times 10^3$  while institutional counts decline to below  $10^3$ , so retail-like activity overtakes institutional on the count margin. This comovement over 2019–2022 is consistent with institutionally mediated flow partly intermediating retail demand, though it does not rule out autonomous institutional activity such as cross-venue settlement or rebalancing.

Panel (b) shows a different pattern for volume. The two series move together until 2021 and then diverge sharply from late 2023: institutional daily volume rises by roughly two orders of magnitude, reaching  $\sim 10^7$  USD by 2025, while retail-like volume remains near  $10^5$ – $10^6$  USD. Two observations follow. First, the marginal dollar of CASP-mediated activity in recent years is institutional rather than retail. Second, this divergence follows in time the shocks we study (Terra/Luna, May 2022; FTX, November 2022; SVB, March 2023). Institutional addresses dominate on volume while retail-like addresses dominate on count in the recent period, and the two margins can move independently.

## 4.4 Summary

Three findings stand out. First, market-facing activity dominates: of the volume intermediated by Austrian CASPs, roughly \$30 billion is conducted with external actors, while inter-CASP flows are negligible (\$5.8 million) and the remainder is internal bookkeeping that moves few transfers but a large share of value. Second, both the size and the activity of CASP counterparties are heavy-tailed, and a structural divide separates them: warm and cold wallets account for only 0.77% of external transfers yet 57.8% of external volume, indicating that a small set of large counterparties handles most of the value. Third, classifying external addresses on these two margins (wallet-tier and activity) lets us examine flows separately and yields a population that is overwhelmingly retail-like by count but institutional by value.

Retail activity is informationally and behaviorally distinct from the institutional (Boehmer et al., 2021; Barber et al., 2024); aggregate data, pooling the two, obscures any group-specific pattern. The stablecoin-redemption channel (Section 2.2) is the cleanest case: Ma et al. (2025) document that direct par redemption of USDC is restricted to verified institutional accounts, so under stress retail users can only sell USDC at secondary-market prices while institutions arbitrage par redemption. Separating the two groups therefore lets us examine flow patterns consistent with such a mechanism. This and the patterns documented above motivate the next question: how retail and institutional flows respond differentially to market-wide shocks.

## 5 Event Study: Responses to Major Crypto Shocks

This section presents the paper’s main behavioral evidence. We estimate whether major crypto-market shocks produce abnormal changes in CASP flows, and whether retail-like and institutionally mediated activity respond differently across assets and flow directions. It also tests the within-crypto safe-haven hypothesis that market stress drives a systematic reallocation into stablecoins, as aggregate-flow studies have suggested (Anadu et al., 2024; Oefele et al., 2024a). We study three episodes identified in Section 3.3, in chronological order: the Terra-Luna crash (Section 5.2), the FTX collapse (Section 5.3), and the Silicon Valley Bank failure (Section 5.4).

### 5.1 Event Study Design

We estimate abnormal CASP flows using daily time series of on-chain activity in BTC, ETH, USDC, USDT. For each asset and group, we construct three outcome measures: transferred volumes in native asset units, transaction counts, and USD-denominated values. We focus on the first since they measure the on-chain quantity response directly. Additional results on other outcomes are reported in Appendix C.

Following Kothari and Warner (2007), abnormal activity is the difference between the observed outcome  $x_t$  and the normal outcome expected absent the event:

$$x_t^* = x_t - E[x_t \mid X_t],$$

where  $X_t$  denotes conditioning information at time  $t$ .



We model normal activity using a constant-mean specification. We use a 45-day baseline pre-event estimation window  $L_1$  and a 10-day post-event window  $L_2$  to estimate coefficients. As a robustness check, we consider alternative pre-event windows ranging from 30 to 60 days. We test the null hypothesis that abnormal activity equals zero with OLS regressions over the combined sample  $L_1 + L_2$ . Since the figures report daily post-event coefficients, we estimate specifications with separate dummy variables for each post-event day:

$$x_t = \alpha + \beta r_t + \sum_{k=0}^{10} \gamma_k D_t^k + \varepsilon_t, \quad (1)$$

where  $x_t$  denotes the respective outcome variable for a given cryptoasset, flow direction, and group;  $\alpha$  is a constant;  $r_t$  is an asset-specific return control (daily BTC returns for BTC, USDC, and USDT, and daily ETH returns for Ether, expressed in percentage points), accounting for normal variation associated with broader crypto-market movements; and  $D_t^k$  is an indicator equal to one on day  $k$  after the event and zero otherwise. The post-event coefficients  $\gamma_k$  estimate the abnormal activity for each day  $k$  after the event.<sup>8</sup> The horizontal line in the event-study figures corresponds to  $\gamma_k = 0$ . Inference is based on a bootstrap procedure with 1,000 resamples of the regression residuals (Kothari and Warner, 2007).

We estimate Eq. (1) separately for each asset, decomposing deposit and withdrawal flows as well as retail-like and institutionally mediated activity. Rather than reporting every possible combination of asset, group, and flow measure, we focus on the margins that best summarize the economic response to each event. The full set of results is reported in Appendix C.

## 5.2 Terra Luna Crash

The Terra-Luna crash in May 2022 was a crypto-native confidence shock: the algorithmic stablecoin broke its dollar peg after markets started exploiting its design vulnerabilities on May 5, entering a destabilizing spiral that culminated in a failed attempt to restore the peg on May 11. The event undermined the credibility of algorithmic stablecoins as a class and triggered a broader confidence shock within crypto markets (see Table 1).

Figure 5 shows the estimated responses, which are pronounced but heterogeneous across assets, groups, and flow directions. The clearest response is in BTC among institutionally mediated flows: both deposits and withdrawals rise well above their bootstrap intervals after May 11, after the final attempt to restore Terra’s peg. Because both directions move together, this points to active portfolio rebalancing rather than a one-sided movement into or out of BTC. Retail-like BTC flows show a weaker version of the same pattern, with deposits rising above their interval around May 9, when Terra lost its peg for the last time, and withdrawals staying largely within their (wide) bands. The BTC response is therefore driven mainly by institutional

---

<sup>8</sup>These coefficients should be interpreted, for all three shocks, as reactions to an unfolding crisis rather than to a single isolated announcement. This is by design: the daily-coefficient specification (Eq. 1) traces the dynamic response day by day and does not presume a one-time jump. The design requires the event to be unanticipated as of the pre-event estimation window, so that normal activity is estimated on uncontaminated data, rather than instantaneous. Our events satisfy this condition, so the pre-event baseline predates any destabilization, while the post-event dummies absorb the crisis as it unfolds.




Figure 5: **Terra Luna crash: selected asset-flow responses.** *Notes:* Selected post-event coefficients from the Terra Luna crash event-study regressions, in asset units. Each panel shows coefficients for days 0 to 10 after the event, using a 45-day pre-event estimation window. The dashed horizontal line marks zero. Shaded areas indicate bootstrap intervals based on 1,000 replications. Rows correspond to BTC and USDC, each split into institutionally mediated and retail-like flows; the left column reports deposits and the right column withdrawals.

actors, consistent with high-volume entities repositioning during the shock, and its lagged timing matches Terra’s progressive depeg.

Stablecoin flows are more mixed. Institutional USDC deposits are volatile, with positive spikes around May 9, when UST lost its peg for the second and final time, and again in the following days, coinciding with the last governance attempt to restore the peg, but show no persistent inflow. Institutional withdrawals remain limited for most of the window but rise strongly at the end, suggesting delayed liquidity adjustment. Retail-like USDC deposits spike on the event day and show smaller positive movements around May 11, while withdrawals lack a clear trend until a clear surge after May 13. Overall, the USDC results suggest short-run liquidity management around key moments in the unfolding Terra crisis rather than a stable and persistent safe-haven movement.


In sum, the Terra Luna results indicate broad reallocation and increased volatility after a crypto-native confidence shock. The most systematic response appears in BTC flows among institutional actors, where deposits and withdrawals move together, pointing to intensified trading or rebalancing. Stablecoin flows are more volatile and less directional, especially for USDC. In this sense, the Terra Luna crash did not generate a uniform flight into stablecoins, but rather triggered heterogeneous adjustments across groups and asset-flow margins.

### 5.3 FTX Collapse

The FTX collapse in November 2022 was a custodial counterparty shock, caused by the failure of one of the largest centralized exchanges active at the time, raising broad concerns about the safety of assets held on custodial platforms. The crisis unfolded on November 2, as doubts about FTX solvency and ties to its affiliated trading firm Alameda triggered a rapid loss of confidence and large withdrawal requests. The post-event window captures the escalation of the crisis, unfolding with a sharp fall in FTT price on November 8, Binance’s withdrawal from the tentative acquisition of FTX on November 9, and the bankruptcy filing of FTX and Alameda on November 11.

Figure 6 reports a selected set of results. The left column shows withdrawals by retail-like entities, while the right column shows deposits by institutional actors. The figure reveals a clear heterogeneous response to the FTX collapse. Retail-like withdrawals increase after the event, especially for BTC and ETH. The response becomes most visible around Nov 8-12, coinciding with the intensification of the FTX crisis. This timing is consistent with a reduction in custodial exposure as the severity of the crisis became fully visible.

Institutional activity displays a different pattern. Their response is concentrated in large deposit spikes occurring around days Nov 8-10. This suggests that institutional actors or large ecosystem participants engaged in substantial reallocation after the shock. One plausible mechanism is that some of these flows reflect CASP liquidity or reserve-management operations: customer trades and withdrawal requests may be settled internally first, while CASPs adjust their on-chain inventories only periodically. During stress episodes such as the FTX collapse, such rebalancing may become more pronounced.



**Figure 6: FTX collapse: selected asset-flow responses.** *Notes:* Selected post-event coefficients from the FTX collapse event-study regressions, in asset units. Each panel shows coefficients for days 0 to 10 after the event, using a 45-day pre-event estimation window. The dashed horizontal line marks zero. Shaded areas indicate bootstrap intervals based on 1,000 replications. The left column reports retail-like withdrawals. The right column reports institutional deposits. Rows correspond to BTC, ETH, USDC, and USDT, respectively.

Overall, the selected event-study results suggest that the FTX collapse increased cryptoasset mobility at Austrian CASPs, but through different adjustment margins. Retail-like activity reacted mainly through withdrawals consistent with reallocation to self-custody wallets, particularly once the crisis escalated after November 8. Institutional actors generated large deposit flows, consistent with high-volume reallocation or platform-level liquidity management. The evidence therefore points to both a broad withdrawal response among retail-like actors and large reallocations by institutional actors following the collapse of FTX.

## 5.4 Silicon Valley Bank failure

The Silicon Valley Bank (SVB) failure was a banking-sector shock that reached crypto markets through the stablecoin channel. It began on 8 March 2023, when SVB announced a \$1.8 billion loss, escalated when the bank failed on 10 March, and hit the crypto ecosystem in particular on 11 March, when Circle disclosed \$3.3 billion of exposure to SVB; tensions eased after uninsured depositors were guaranteed on 12 March. Because Circle issues USDC, the episode is especially relevant for stablecoins, with a weaker and less systematic response expected for BTC and ETH.

Figure 7 shows the estimated responses, which are consistent with a shock transmitted mainly through stablecoins rather than through broad crypto reallocation. BTC retail-like deposits stay within their bootstrap intervals throughout, with no persistent post-event movement. BTC institutional withdrawals are noisier, with several points above the interval, but show no sustained direction. ETH is similar: retail-like deposits and institutional withdrawals fluctuate and occasionally exceed their intervals on single days, but no response persists and the intervals remain wide. As anticipated, the non-stablecoin response is weak.

The clearest reactions appear in USDC, and align with the timing of the sub-events. Retail-like USDC deposits show large spikes well above their interval on March 9 and 11, around Circle’s exposure disclosure, before falling back within the band. At the same time, institutional USDC withdrawals exhibit a large spike on March 11, when Circle’s exposure became public, against a tight interval, with a further above-interval spike later in the window. Because retail-like USDC deposits are consistent with preparing to sell within the CASP, while institutional withdrawals are consistent with pulling USDC off-venue to redeem at par, this two-sided pattern matches the redemption asymmetry documented by Ma et al. (2025): direct par redemption of USDC is restricted to verified institutional accounts, so under stress retail users sell USDC at secondary-market prices while institutions can arbitrage par redemption. The USDC panels therefore point to strong two-sided activity rather than a one-directional safe-haven inflow.

USDT behaves differently. Retail-like USDT deposits rise above their interval over the first post-event days, consistent with the pattern observed with USDC. Institutional USDT withdrawals, instead, are volatile, with a single above-interval spike, but the pattern is less clear than for USDC, where both margins move sharply against tight intervals.

Overall, the SVB failure operated through a narrower stablecoin channel. As expected, BTC and ETH responses are weak, noisy, and do not show persistent directional movements. The strongest reactions appear in USDC, and to a lesser extent USDT, around the key SVB and




Figure 7: **SVB failure: selected asset-flow responses.** *Notes:* Selected post-event coefficients from the SVB failure event-study regressions, in asset units. Each panel shows coefficients for days 0 to 10 after the event, using a 45-day pre-event estimation window. The dashed horizontal line marks zero. Shaded areas indicate bootstrap intervals based on 1,000 replications. The left column reports deposits by retail-like entities, while the right column reports withdrawals by institutional actors. Rows correspond to BTC, ETH, USDC, and USDT, respectively.

Circle-related disclosures. USDC retail-like deposits and institutional withdrawals both increase, consistent with Circle’s restriction of USDC direct redemption to verified institutional accounts.

## 5.5 Summary

The three event studies show that each shock triggered distinct and economically interpretable flow responses. The Terra Luna crash, a crypto-native confidence shock, triggered broad reallocation and elevated volatility, most clearly in the comovement of institutional BTC deposits and withdrawals, rather than uniform flights into stablecoins. These patterns are consistent with active portfolio rebalancing. After the FTX collapse, a custodial counterparty shock, retail-like actors reduced custodial exposure through sustained withdrawals once the crisis escalated, while institutional actors generated large episodic deposits consistent with high-volume reallocation or CASP inventory management. The SVB failure operated through the stablecoin channel, with BTC and ETH largely unresponsive and the strongest reaction in USDC, where retail-like deposits and institutional withdrawals are consistent with the restriction of par redemption to institutional accounts (Ma et al., 2025). Taken together, the nature of the shock determines the margin of response, and the opposing retail and institutional flows seen most sharply at SVB would largely cancel in aggregate data, underscoring the value of separating customer groups and flow directions. Results from alternative pre-event windows and outcome variables are comparable (see Appendix C).

## 6 Discussion

### 6.1 Key Findings

This paper studies how cryptoasset activity intermediated by Austria’s registered CASPs integrates into the national financial system, using a regulatory dataset covering all crypto-asset service providers registered in Austria at the end of 2024. Under Austrian anti-money-laundering law, CASPs must report the full set of blockchain addresses they control (FMA, 2025). From the resulting pseudonymized address-to-entity registry, we reconstruct the on-chain activity of the 12 registered CASPs across Bitcoin (BTC), Ether (ETH), and the two largest stablecoins (USDT and USDC) through May 2025: 11.98 million asset transfers and \$50.4 billion in cumulative volume. Because ownership is established directly from the registry rather than inferred through clustering heuristics, which by construction produce false positives and negatives (Meiklejohn et al., 2013; Fröwis et al., 2020), we observe these flows with a completeness and reliability that country-level studies relying on indirect proxies have not achieved (Auer et al., 2025).

We document four main findings. First, Austrian CASPs intermediate substantial economic activity, roughly \$30 billion in on-chain flows with external counterparties, but the domestic network is only weakly interconnected: flows between Austrian CASPs total just \$5.8 million. The Austrian ecosystem is integrated globally rather than domestically, consistent with the cross-border orientation of crypto flows documented at the aggregate level (Auer et al., 2025). A further \$20.7 billion, or 41.1% of total volume, reflects within-entity bookkeeping rather than

counterparty activity. This echoes, at a jurisdiction-specific scale, the much larger share of network-wide Bitcoin volume that Makarov and Schoar (2021) attribute to within-entity and protocol-related transfers, and it confirms that internal wallet management may be misclassified as economic activity in aggregate on-chain data.

Second, external activity is highly concentrated. A small set of transactions conducted directly with the warm and cold wallets of Austrian CASPs account for 57.8% of external volume while representing only 0.77% of external transfers, whereas the overwhelming majority of addresses, about 98% on each chain, exhibit low-frequency, retail-like behavior, transacting two to three times for a few hundred to a few thousand dollars. Exploiting this structural divide, we classify external counterparties using two signals, warm/cold-wallet contact and chain-wide transaction activity, rather than the single trade-size threshold standard in the microstructure literature (Schultz, 2012; Boehmer et al., 2021). This responds to the documented fragility of one-dimensional proxies: trade-size cutoffs misclassify a substantial share of trades, as informed and uninformed counterparties increasingly transact at similar sizes (Barber et al., 2024).

Third, in event studies around three broadly covered shocks, namely a run on an algorithmic stablecoin (Terra-Luna), the failure of a centralized exchange (FTX), and the failure of a bank intertwined with the crypto sector (SVB), retail-like and institutional actors respond through different margins. After the FTX collapse, retail-like withdrawals increased persistently and across assets, consistent with reduced custodial exposure following a counterparty failure, while institutional flows are consistent with liquidity management and cross-platform rebalancing. The SVB episode produced a narrower response concentrated in USDC, consistent with Circle’s two-tiered redemption mechanism, and split sharply by group, with retail-like depositing and institutional withdrawing; instead, BTC and ETH were largely unaffected. The Terra-Luna crash generated broader but less persistent reallocation across assets, leading to heightened volatility. This group-level heterogeneity is the behavioral counterpart to the distinction between retail and institutional order flow established in equity markets (Boehmer et al., 2021; Barber et al., 2024), and would be invisible in the aggregate flows on which prior crypto event studies rely (Abramova and Böhme, 2021; Yousaf et al., 2023).

Fourth, we find little evidence that stablecoins function as a systematic within-crypto safe haven during stress. Rather than a uniform inflow, stablecoin responses were two-sided and shock-specific. After FTX, retail-like outflows involved all assets, including stablecoins, and after SVB, retail-like deposits and institutional withdrawals dominated flows. The divergence after SVB is consistent with USDC’s redemption asymmetry (Ma et al., 2025): facing a peg break, retail users cannot redeem USDC at par and sell on secondary markets, while par redemption is reserved for institutional accounts. These findings complicate the safe-haven and flight-to-quality interpretations drawn from aggregate flows (Baur and Hoang, 2021; Anadu et al., 2024; Oefele et al., 2024a) and align instead with the more recent evidence of heterogeneous, inconsistent safe-haven dynamics across events (Aldasoro et al., 2025).



## 6.2 Impact

Our findings speak to several groups of stakeholders.

**Financial supervisors and regulators.** The registry-based approach offers a systematic and repeatable way to measure cryptoasset activity bound to a jurisdiction. Because it derives flows directly from addresses reported by supervised CASPs, rather than from third-party attribution or clustering heuristics, it yields precise and reproducible measurements once the reporting pipeline is in place. This suggests a concrete supervisory practice: collecting address-registry data directly from supervised entities, and updating it as their wallet infrastructure changes. Supervisors could further require CASPs to report available attribution on their counterparties, since reporting entities can often distinguish retail customers from institutional ones through their own KYC and account records. Such information would sharpen the retail-versus-institutional distinction that we can only approximate from on-chain behavior, and would let supervisors observe heterogeneous responses to stress directly rather than through proxies. The same design could, in principle, be extended across jurisdictions to support a broader observatory of cryptoasset flows, European or even global in scope, allowing supervisors to compare on-chain activity on a consistent basis rather than through heterogeneous national proxies.

**Policymakers.** The results bear on debates over financial stability, monetary transmission, and public digital money (e.g., central bank digital currencies). The limited and two-sided stablecoin response during stress, rather than a clean flight to quality, raises the question of how far privately issued stablecoins can replicate the stabilizing role of traditional safe assets, and how they will interact with future public digital currencies. At the same time, the pronounced withdrawal response following the custodial failure of FTX is consistent with how quickly confidence in custodial intermediaries can shift, a channel relevant to any framework that treats CASPs as points of systemic exposure.

**Crypto-asset service providers.** For CASPs, the speed of the responses is the salient feature: abnormal flows appear within days of each shock, with no jurisdictional lag visible in the data. Stress propagates through the on-chain system rapidly, leaving little time for operational adjustment once an event unfolds. This argues for ex-ante safeguards rather than reactive ones, in particular liquidity and redemption arrangements robust to sudden two-sided stablecoin activity of the kind observed around the SVB and Circle disclosures.

**Retail users.** The evidence indicates that retail-like and institutional flows operate on different margins, with retail dominating by transaction count and institutions by volume, and that the two can move in opposite directions during stress. For retail users specifically, the SVB episode illustrates a structural disadvantage: when a stablecoin’s peg is threatened, par redemption is available only to a restricted set of institutional accounts, so retail holders are left to sell on secondary markets, often at a loss (Ma et al., 2025). This asymmetry is a consumer-protection consideration that aggregate stability metrics do not capture.

### 6.3 Limitations and Future Work

Several limitations qualify our results, each pointing to a direction for future work. First, treating the dataset as a trustworthy source depends on the accuracy and completeness of CASP reporting. In the course of the analysis we identified a few addresses likely controlled by Austrian CASPs that were absent from the original mapping, which indicates that reporting is not fully exhaustive. Future work could pair the registry with independent attribution or clustering as a cross-check, quantifying reporting gaps rather than assuming completeness. Relatedly, because the registry reflects only CASPs registered at the end of 2024, our analysis excludes providers that exited the market earlier. While this survivorship concern is real, it is mitigated by the fact that entry and exit concentrate among small CASPs, while the large providers driving most volumes are stable over the sample period; furthermore, the Austrian CASP market itself was negligible before 2017-2018. One similar concern is that address control may change over time. This issue is mitigated by the fact that the registry is retroactive, that is, it includes both addresses controlled at the collection date and at earlier dates.

Second, the retail-versus-institutional classification is proxy-based: it cannot fully account for a single actor’s activity being fragmented across many addresses, and the activity threshold ( $z > 3$ ) is conventional rather than theoretically derived. As we argue above, the most direct remedy lies with supervisors, who could attach verified counterparty attribution from CASP KYC records; absent that, future work could validate and refine the classification against partial ground-truth labels and assess sensitivity to alternative thresholds and address-linkage methods.

Third, the analysis covers a single jurisdiction and may not generalize to larger or structurally different markets. Because the approach depends only on registry data that other supervisors could in principle collect, the natural extension is cross-jurisdictional replication, which would also enable the comparative observatory discussed above.

Fourth, we cover four major cryptoassets, omitting the long tail of smaller assets, and we observe USDT and USDC only on Ethereum although both are deployed on several chains. Extending coverage to additional assets and to the other chains on which stablecoins circulate would capture cross-chain substitution that our single-chain view of stablecoins may miss, an aspect directly relevant to the safe-haven question.

Fifth, we observe only Austrian CASPs. We therefore do not see Austrian holders who use non-Austrian providers or who transact solely through self-custodial wallets, so our flows are a partial view of national cryptoasset activity. Linking the registry to cross-border or self-custody flows, where data-sharing arrangements permit, would help place the Austrian segment, or that of another jurisdiction, within the wider ecosystem.

Beyond addressing these limitations, the decomposition we introduce, by counterparty, flow direction, and group, offers a tool for identifying which contagion channels are most active during stress. Applying it systematically across events and jurisdictions could help supervisors assess where cryptoasset markets transmit risk into the wider financial system, which is central to evaluating financial-stability risk.

## 7 Conclusions

Using a regulatory registry that identifies the on-chain addresses of all crypto-asset service providers (CASPs) registered in Austria, we provide a transaction-level, group-resolved view of how CASP-intermediated cryptoasset activity integrates into a national financial system and how it reacts under stress. The registry lets us measure these flows directly, without the clustering heuristics or indirect proxies on which prior country-level work has relied. We also propose a simple approach to separate retail-like from institutionally mediated activity using on-chain behavior rather than a single trade-size threshold. Two pictures emerge. Structurally, Austrian CASPs intermediate tens of billions of dollars but are integrated globally rather than with one another, and their external activity is dominated in value by a small set of institutional counterparties and in number by a large base of retail-like ones. Dynamically, the two groups respond to systemic shocks (Terra-Luna, FTX, and SVB) through different margins, and stablecoins do not act as a uniform within-crypto safe haven: the clearest stress response, around the SVB failure, was two-sided, with retail-like users depositing USDC into CASPs while institutions withdrew them. These patterns are visible only once we separate flows by counterparty and direction, so aggregate on-chain data would obscure them. Overall, registry-based, transaction-level measurement is a promising foundation for monitoring how the cryptoasset activity intermediated by a national jurisdiction’s registered CASPs integrates into the financial system.

## References

- Abramova, S. and Böhme, R. (2021). Out of the dark: The effect of law enforcement actions on cryptocurrency market prices. In *2021 APWG Symposium on Electronic Crime Research (eCrime)*, pages 1–11. IEEE.
- Abramova, S., Böhme, R., Elsinger, H., Stix, H., and Summer, M. (2023). What can central bank digital currency designers learn from asking potential users? In *Nineteenth Symposium on Usable Privacy and Security (SOUPS 2023)*, pages 151–170.
- Adediran, A., Babajide, B., and Mafimisebi, O. (2024). Are stablecoins safe haven asset during conflict and political crisis? evidence from natural experiment. *Available at SSRN: <https://ssrn.com/abstract=4869383>*.
- Akyildirim, E., Conlon, T., Corbet, S., and Goodell, J. W. (2023). Understanding the ftx exchange collapse: A dynamic connectedness approach. *Finance Research Letters*, 53:103643.
- Aldasoro, I., Cornelli, G., Minesso, M. F., Gambacorta, L., and Habib, M. M. (2025). Stablecoins, money market funds and monetary policy. *Economics Letters*, 247:112203.
- Almeida, J. and Gonçalves, T. C. (2023). A systematic literature review of investor behavior in the cryptocurrency markets. *Journal of Behavioral and Experimental Finance*, 37:100785.
- Alvarez, F., Argente, D., and Van Patten, D. (2023). Are cryptocurrencies currencies? bitcoin as legal tender in el salvador. *Science*, 382(6677):eadd2844.

- Anadu, K., Azar, P., Huang, C., Cipriani, M., Eisenbach, T. M., La Spada, G., Landoni, M., Macchiavelli, M., Malfroy-Camine, A., and Wang, J. C. (2024). Runs and flights to safety: Are stablecoins the new money market funds? Staff Report 1073, Federal Reserve Bank of New York.
- Ante, L., Fiedler, I., Willruth, J. M., and Steinmetz, F. (2023). A systematic literature review of empirical research on stablecoins. *FinTech*, 2(1):34–47.
- Ashwanth, S., Aronoff, D., and Narula, N. (2025). 1:1 redemptions for some, not all. *Digital Currency Initiative (DCI). MIT Media Lab*.
- Auer, R., Farag, M., Lewrick, U., Orazem, L., and Zoss, M. (2022). Banking in the shadow of bitcoin? the institutional adoption of cryptocurrencies. BIS Working Paper 1013, Bank for International Settlements.
- Auer, R., Haslhofer, B., Kitzler, S., Saggese, P., and Victor, F. (2024). The technology of decentralized finance (defi). *Digital Finance*, 6(1):55–95.
- Auer, R., Lewrick, U., and Paulick, J. (2025). Defying gravity? an empirical analysis of cross-border bitcoin, ether and stablecoin flows. Technical Report 1265, Bank for International Settlements.
- Awrey, D. (2020). Bad money. *Cornell L. Rev.*, 106:1.
- Baele, L., Bekaert, G., Inghelbrecht, K., and Wei, M. (2020). Flights to safety. *The Review of Financial Studies*, 33(2):689–746.
- Baqer, K., Huang, D. Y., McCoy, D., and Weaver, N. (2016). Stressing out: Bitcoin “stress testing”. In *International conference on financial cryptography and data security*, pages 3–18. Springer.
- Barber, B. M., Huang, X., Jorion, P., Odean, T., and Schwarz, C. (2024). A (sub)penny for your thoughts: Tracking retail investor activity in TAQ. *Journal of Finance*, 79(4):2403–2427.
- Baumöhl, E. and Vyrost, T. (2020). Stablecoins as a crypto safe haven? not all of them!
- Baur, D. G., Cahill, D., Godfrey, K., and Liu, Z. F. (2019). Bitcoin time-of-day, day-of-week and month-of-year effects in returns and trading volume. *Finance Research Letters*, 31:78–92.
- Baur, D. G. and Hoang, L. T. (2021). A crypto safe haven against bitcoin. *Finance Research Letters*, 38:101431.
- Baur, D. G. and McDermott, T. K. (2010). Is gold a safe haven? international evidence. *Journal of Banking & Finance*, 34(8):1886–1898.
- Belguith, R., Manzli, Y. S., Bejaoui, A., and Jeribi, A. (2024). Can gold-backed cryptocurrencies have dynamic hedging and safe-haven abilities against defi and nft assets? *Digital Business*, 4(2):100077.

- Bessembinder, H., Spatt, C., and Venkataraman, K. (2020). A survey of the microstructure of fixed-income markets. *Journal of Financial and Quantitative Analysis*, 55(1):1–45.
- Binder, C. (2025). Central banks as educators. Working paper Nr. 34002, National Bureau of Economic Research.
- Boehmer, E., Jones, C. M., Zhang, X., and Zhang, X. (2021). Tracking retail investor activity. *Journal of Finance*, 76(5):2249–2305.
- Briola, A., Vidal-Tomás, D., Wang, Y., and Aste, T. (2023). Anatomy of a stablecoin’s failure: The terra-luna case. *Finance Research Letters*, 51:103358.
- Caballero, R. J., Farhi, E., and Gourinchas, P.-O. (2008). An equilibrium model of “global imbalances” and low interest rates. *American Economic Review*, 98(1):358–393.
- Capponi, A., Jia, R., and Yu, S. (2026). Price discovery on decentralized exchanges. *The Review of Financial Studies*, page hhag002.
- Cardozo, P., Fernández, A., Jiang, J., and Rojas, F. D. (2024). On cross-border crypto flows: Measurement drivers and policy implications. Technical report, International Monetary Fund.
- Carpentier-Desjardins, C., Paquet-Clouston, M., Kitzler, S., and Haslhofer, B. (2025). Mapping the defi crime landscape: an evidence-based picture. *Journal of Cybersecurity*, 11(1):tyae029. The dataset containing a list of crime-related events is available at: <https://zenodo.org/records/14706760>. Accessed on 2025-06-20.
- Catalini, C., de Gortari, A., and Shah, N. (2022). Some simple economics of stablecoins. *Annual Review of Financial Economics*, 14(1):117–135.
- Cerutti, E., Chen, M. J., and Hengge, M. (2024). *A Primer on Bitcoin Cross-Border Flows: Measurement and Drivers*. International Monetary Fund.
- Coingecko (2026). Coingecko api documentation. Available at: <https://www.coingecko.com/en/api>.
- Coinmarketcap (2026). Coinmarketcap timeline of cryptoasset-specific events. Available at: <https://coinmarketcap.com/currencies/bitcoin/>.
- Cong, L. W., Rabetti, D., Wang, C. C., and Yan, Y. (2025). Centralized governance in decentralized organizations. *Available at SSRN 5168660*.
- Conlon, T., Corbet, S., and Hu, Y. (2023). The collapse of the ftx exchange: The end of cryptocurrency’s age of innocence. *The British Accounting Review*, page 101277.
- Conlon, T., Corbet, S., and McGee, R. J. (2020). Are cryptocurrencies a safe haven for equity markets? an international perspective from the covid-19 pandemic. *Research in International Business and Finance*, 54:101248.

- DeFiLlama (2025). Defi hacks apis. Available at: <https://api.llama.fi/hacks>. Accessed on 2025-06-20.
- deHaan, E., Li, J., and Watts, E. M. (2023). Retail bond investors and credit ratings. *Journal of Accounting and Economics*, 76(1):101587.
- Dionysopoulos, L. and Makridis, C. A. (2026). Central bank digital currency rhetoric and stablecoin market responses. *Journal of the Royal Statistical Society: Series A*, 189(2):630–645.
- Dionysopoulos, L., Marra, M., and Urquhart, A. (2025). Stablecoins as digital havens: Cryptocurrencies in conflict and sanctioned economies. *Available at SSRN 5259143*.
- Dionysopoulos, L. and Urquhart, A. (2024). 10 years of stablecoins: Their impact, what we know, and future research directions. *Economics Letters*, 244:111939.
- Divakaruni, A. and Zimmerman, P. (2024). Uncovering retail trading in bitcoin: The impact of covid-19 stimulus checks. *Management Science*, 70(4):2066–2085.
- Feldkircher, M. and Makridis, C. A. (2025). Words matter: Central bank communication and household expectations in a global panel. *SSRN working paper*.
- Feng, J., Yuan, Y., and Jiang, M. (2024). Are stablecoins better safe havens or hedges against global stock markets than other assets? comparative analysis during the covid-19 pandemic. *International Review of Economics & Finance*, 92:275–301.
- Fessler, P. and Weber, B. (2024). Oenb bulletin, june 2024: Crypto assets in austria: an assessment of their prevalence and the motives of their holders. *OeNB Bulletin*, (Q2/2024-2).
- Feyen, E., Kawashima, Y., and Mittal, R. (2022). Crypto-assets activity around the world. *Evolution and Macro-Financial Drivers, World Bank, Washington, DC*, (9962).
- FMA (2025). Financial markets anti-money laundering act. Technical report.
- Fröwis, M., Gottschalk, T., Haslhofer, B., Rücker, C., and Pesch, P. (2020). Safeguarding the evidential value of forensic cryptocurrency investigations. *Forensic Science International: Digital Investigation*.
- Galati, L. and Capalbo, F. (2024). Silicon valley bank bankruptcy and stablecoins stability. *International Review of Financial Analysis*, 91:103001.
- Gorton, G. B. and Zhang, J. Y. (2023). Taming wildcat stablecoins. *U. Chi. L. Rev.*, 90:909.
- Griffin, J. M. and Shams, A. (2020). Is bitcoin really untethered? *The Journal of Finance*, 75(4):1913–1964.
- Hackethal, A., Hanspal, T., Lammer, D. M., and Rink, K. (2022). The characteristics and portfolio behavior of bitcoin investors: Evidence from indirect cryptocurrency investments. *Review of Finance*, 26(4):855–898.

- He, Z., Krishnamurthy, A., and Milbradt, K. (2019). A model of safe asset determination. *American Economic Review*, 109(4):1230–1262.
- Hu, M. R., Lee, A. D., and Putniņš, T. J. (2021). Evading capital controls via cryptocurrencies: Evidence from the blockchain. *Available at SSRN 3956933*.
- Hund, J., Lundblad, C. T., Makridis, C. A., and Nguyen, G. (2025). The rise of investor sophistication and the decline of underwriting profits in the municipal bond market. *SSRN Working Paper*, (4867270). Available at SSRN: <https://ssrn.com/abstract=4867270>.
- Jalan, A. and Matkovskyy, R. (2023). Systemic risks in the cryptocurrency market: Evidence from the ftx collapse. *Finance Research Letters*, 53:103670.
- Jiang, L., Liu, J., Peng, L., and Wang, B. (2022). Investor attention and asset pricing anomalies. *Review of Finance*, 26(3):563–593.
- Jiang, Z., Krishnamurthy, A., and Lustig, H. (2021). Foreign safe asset demand and the dollar exchange rate. *The Journal of Finance*, 76(3):1049–1089.
- John, K. and Li, J. (2025). Bitcoin price volatility: Effects of retail traders, illegal users, and sentiment. *Journal of Corporate Finance*, 94:102837.
- Kołodziejczyk, H. (2023). Stablecoins as diversifiers, hedges and safe havens: A quantile coherency approach. *The North American Journal of Economics and Finance*, 66:101912.
- Kothari, S. P. and Warner, J. B. (2007). Econometrics of event studies. In Eckbo, B. E., editor, *Handbook of Corporate Finance: Empirical Corporate Finance*, volume 1, pages 3–36. Elsevier.
- Koutrouli, E., Manousopoulos, P., Theal, J., and Tresso, L. (2025). Crypto asset markets vs. financial markets: Event identification, latest insights and analyses. *AppliedMath*, 5(2):36.
- Krishnamurthy, A. and Vissing-Jorgensen, A. (2012). The aggregate demand for treasury debt. *Journal of Political Economy*, 120(2):233–267.
- Kropiunig, A., Kremer, S., and Haslhofer, B. (2026). Credibility matters: Motivations, characteristics, and influence mechanisms of crypto key opinion leaders. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, CHI ’26, pages 1–17, New York, NY, USA. Association for Computing Machinery.
- Li, T., Shin, D., and Wang, B. (2025). Cryptocurrency pump-and-dump schemes. *Journal of Financial and Quantitative Analysis*, 60(8):3622–3659.
- Liu, J., Makarov, I., and Schoar, A. (2023). Anatomy of a run: The terra luna crash. Technical report, National Bureau of Economic Research.
- Ma, Y., Zeng, Y., and Zhang, A. L. (2025). Stablecoin runs and the centralization of arbitrage. Nber working paper, National Bureau of Economic Research.

- Makarov, I. and Schoar, A. (2021). Blockchain analysis of the Bitcoin market. *NBER Working Paper*, (29396).
- Makridis, C. A. (2025). The rise of central bank digital currencies: Exploring adoption determinants and early macroeconomic and well-being impacts. *Available at SSRN 4932810*.
- Marmora, P. (2022). Does monetary policy fuel bitcoin demand? event-study evidence from emerging markets. *Journal of International Financial Markets, Institutions and Money*, 77:101489.
- Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., and Savage, S. (2013). A fistful of bitcoins: characterizing payments among men with no names. In *Proceedings of the 2013 conference on Internet measurement conference*, pages 127–140.
- Mertens, K., Lewis, D. J., and Makridis, C. (2020). Do monetary policy announcements shift household expectations? CEPR Discussion Paper DP14360, Centre for Economic Policy Research.
- Metrick, A. (2024). The failure of silicon valley bank and the panic of 2023. *Journal of Economic Perspectives*, 38(1):133–152.
- MiCA (2023). Regulation (eu) 2023/1114 of the european parliament and of the council of 31 may 2023 on markets in crypto-assets. Technical report.
- Mo, B., Chen, J., Shi, Q., and Zeng, Z. (2025). Cryptocurrencies as safe havens for geopolitical risk? a quantile analysis approach. *The North American Journal of Economics and Finance*, page 102439.
- Möser, M., Böhme, R., and Breuker, D. (2013). An inquiry into money laundering tools in the bitcoin ecosystem. In *2013 APWG eCrime researchers summit*, pages 1–14. Ieee.
- Oefele, N., Baur, D. G., and Smales, L. A. (2024a). Are stablecoins the money market mutual funds of the future? *Journal of Empirical Finance*, 79:101557.
- Oefele, N., Baur, D. G., and Smales, L. A. (2024b). Flight-to-quality—money market mutual funds and stablecoins during the march 2023 banking crisis. *Economics Letters*, 234:111464.
- O’Hara, M. and Zhou, X. A. (2021). The electronic evolution of corporate bond dealers. *Journal of Financial Economics*, 140(2):368–390.
- Parino, F., Beiró, M. G., and Gauvin, L. (2018). Analysis of the bitcoin blockchain: socio-economic factors behind the adoption. *EPJ Data Science*, 7(1):1–23.
- Philippas, D., Rjiba, H., Guesmi, K., and Goutte, S. (2019). Media attention and bitcoin prices. *Finance Research Letters*, 30:37–43.



- Ron, D. and Shamir, A. (2013). Quantitative analysis of the full bitcoin transaction graph. In *Financial Cryptography and Data Security: 17th International Conference, FC 2013, Okinawa, Japan, April 1-5, 2013, Revised Selected Papers 17*, pages 6–24. Springer.
- Saggese, P., Segalla, E., Sigmund, M., Raunig, B., Zangerl, F., and Haslhofer, B. (2024). Assessing the solvency of virtual asset service providers: are current standards sufficient? *Applied Economics*, pages 1–16.
- Schultz, P. (2012). The market for new issues of municipal bonds: The roles of transparency and limited access to retail investors. *Journal of Financial Economics*, 106(3):492–512.
- Shiller, R. J. (2017). Narrative economics. *American Economic Review*, 107(4):967–1004.
- Ter Ellen, S., Larsen, V. H., and Thorsrud, L. A. (2022). Narrative monetary policy surprises and the media. *Journal of Money, Credit and Banking*, 54(5):1525–1549.
- Victor, F. (2020). Address clustering heuristics for ethereum. In *International conference on financial cryptography and data security*, pages 617–633. Springer.
- Von Luckner, C. G., Reinhart, C. M., and Rogoff, K. (2023). Decrypting new age international capital flows. *Journal of Monetary Economics*, 138:104–122.
- Wang, G.-J., Ma, X.-y., and Wu, H.-y. (2020). Are stablecoins truly diversifiers, hedges, or safe havens against traditional cryptocurrencies as their name suggests? *Research in International Business and Finance*, 54:101225.
- Xie, Y., Kang, S. B., and Zhao, J. (2021). Are stablecoins safe havens for traditional cryptocurrencies? an empirical study during the covid-19 pandemic. *Applied Finance Letters*, 10:2–9.
- Yousaf, I., Riaz, Y., and Goodell, J. W. (2023). What do responses of financial markets to the collapse of ftx say about investor interest in cryptocurrencies? event-study evidence. *Finance Research Letters*, 53:103661.

## A Data - Supplemental material

This section provides additional information on the address-to-entity mapping and the analysis on Google Trends data conducted to select important crypto-related events.

### A.1 Address Registry

Table 5 provides a breakdown of the addresses by CASP and type (hot, warm, and cold). The vast majority are hot wallets and are used to facilitate deposit/withdrawal interactions with customers. Addresses are not equally distributed among CASPs nor by wallet type, with a limited number of CASPs controlling the largest number of addresses.

| CASP  | Hot    | Warm   | Cold   | Total   |
|-------|--------|--------|--------|---------|
| 1     | 0      | 0      | 0.001  | 0.001   |
| 2     | 81.846 | 0.033  | 0.252  | 82.132  |
| 3     | 0.148  | 0      | <0.001 | 0.149   |
| 4     | 0.515  | 0      | 0      | 0.515   |
| 5     | 7.751  | 0      | <0.001 | 7.752   |
| 6     | 8.826  | 0      | <0.001 | 8.827   |
| 7     | <0.001 | 0      | 0      | <0.001  |
| 8     | 0.130  | 0      | 0.006  | 0.135   |
| 9     | 0      | <0.001 | <0.001 | 0.002   |
| 10    | 0.483  | 0      | 0.002  | 0.485   |
| 11    | <0.001 | 0      | 0      | <0.001  |
| 12    | 0      | 0.002  | 0      | 0.002   |
| Total | 99.702 | 0.036  | 0.263  | 100.000 |

Table 5: **Relative amount of wallets controlled by Austrian CASPs.** Wallets are divided in hot, warm, and cold; values are in percentage as agreed with the Austrian financial authorities.

Table 6 reports the evolution in time of the number of asset transfers recorded, broken down by cryptoasset and year. Bitcoin, as expected, plays a major role, while Ether and stablecoins entered the market in 2017 and 2019, respectively. The CASP activity peaked in terms of number of transactions in 2021 and decreased thereafter, with the Bitcoin share decreasing compared to other assets after 2022.

| Year Asset | 2016   | 2017      | 2018      | 2019      | 2020      | 2021      | 2022      | 2023      | 2024      | 2025    | Total      |
|------------|--------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|---------|------------|
| BTC        | 52,040 | 1,087,398 | 1,310,338 | 1,038,242 | 1,378,311 | 1,864,144 | 1,366,961 | 870,018   | 531,972   | 189,021 | 9,688,445  |
| ETH        | 0      | 2,080     | 39,405    | 137,319   | 165,967   | 377,384   | 246,493   | 272,073   | 335,149   | 126,034 | 1,701,904  |
| USDC       | 0      | 0         | 0         | 1,851     | 5,516     | 19,094    | 19,211    | 14,071    | 16,943    | 31,855  | 108,541    |
| USDT       | 0      | 0         | 0         | 3,441     | 33,743    | 100,183   | 86,160    | 112,035   | 125,635   | 20,600  | 481,797    |
| Tot.       | 52,040 | 1,089,478 | 1,349,743 | 1,180,853 | 1,583,537 | 2,360,805 | 1,718,825 | 1,268,197 | 1,009,699 | 367,510 | 11,980,687 |

Table 6: **Number of asset transfers recorded on-chain by year and asset.**

### A.2 Event Selection

Table 7 reports the keywords used to measure via Google Trends the media coverage for each of the 19 major events detected as discussed in Section 3. We note that, since the names of DeFi


| Event                        | Keywords searched in Google Trends                                                |
|------------------------------|-----------------------------------------------------------------------------------|
| CeFi events                  |                                                                                   |
| Bitfinex                     | Bitfinex                                                                          |
| Kucoin                       | Kucoin                                                                            |
| Bitmart                      | Bitmart                                                                           |
| FTX                          | FTX, Sam Bankman-Fried, FTX bankruptcy                                            |
| Poloniex                     | Poloniex                                                                          |
| Silicon Valley Bank Collapse | <i>SVB</i> , Silicon Valley Bank                                                  |
| DMM Bitcoin                  | DMM Bitcoin                                                                       |
| WazirX                       | WazirX, WazirX India                                                              |
| Bybit                        | Bybit                                                                             |
| DeFi events                  |                                                                                   |
| Poly Network                 | Poly Network                                                                      |
| Compound V2                  | <i>Compound V2</i> , Compound protocol, Compound DeFi                             |
| Badger DAO                   | <i>Badger DAO</i> , Badger, BadgerDAO                                             |
| Portal                       | <i>Portal protocol</i> , Portal, Portal DeFi                                      |
| Ronin Bridge                 | <i>Ronin Bridge</i> , Ronin, Ronin DeFi                                           |
| Terra depeg                  | <i>Terra Luna</i> , UST Terra, Terra stablecoin, Terra crash, UST                 |
| Binance Bridge               | <i>Binance chain</i> , BNBChain, BSC, BSC Token Hub, Binance Bridge, Binance hack |
| Euler V1                     | <i>Euler protocol</i> , Euler, Euler DeFi                                         |
| Multichain                   | <i>Multichain</i> , Multichain DeFi, Multichain protocol                          |
| Cetus CLMM                   | Cetus CLMM                                                                        |

Table 7: **List of keywords analyzed with Google Trends.** When two or more were used, the one selected is in italic.


protocols often correspond with terms widely used in spoken English (an obvious example is the lending protocol Compound), we searched for multiple keywords when ambiguities could arise. Instead, the choice for CeFi actors was more straightforward and led to lesser ambiguity.

Figure 8 reports the results of the Google Trends analysis. For each event, we show the evolution in time of the search volume and a vertical line in correspondence of each event date (colors of the time series and the vertical line correspond). It is evident that Terra-Luna, FTX, and SVB received much broader media attention and public interest at the time of the shock with respect to the other events.


Next, we describe the series of sub-events that led to each of the three major shocks selected: the Terra-Luna crash, the FTX bankruptcy, and the SVB collapse. The first one refers to the de-pegging from the US dollar of Terra-Luna, the largest algorithmic stablecoin at the time (May 2022). TerraUSD (UST) maintained its peg through an arbitrage mechanism with its associated token LUNA, rather than through collateral backing, making the system inherently fragile under stress. In early May 2022, significant selling pressure on both Bitcoin and LUNA contributed to destabilizing market conditions. On May 7, UST lost its peg to the US dollar for the first time, reportedly following large withdrawals from liquidity pools in the Terra ecosystem, triggering a wave of redemptions. Despite temporary recovery attempts, UST de-pegged again on May 9, leading to a self-reinforcing ‘death spiral’ in which the minting of LUNA to defend the peg led to hyperinflation and a collapse in its price. By May 11, efforts by the Terra founders to restore the peg had failed, marking the effective breakdown of the ecosystem.




(a) CeFi, World



(b) DeFi, World



(c) CeFi, Austria



(d) DeFi, Austria

Figure 8: **Google Trends results for CeFi (left) and DeFi events (right).** Top panels report results on web searches conducted worldwide, bottom panels refer to searches in Austria.

The second large shock we investigate is the failure of FTX, one of the most prominent cryptocurrency exchanges at the time of its failure in November 2022. Its founder also owned Alameda Research, a trading firm linked with FTX, to which FTX had lent billions of its customers' assets and who held significant amount of FTT, an exchange token issued by FTX. On November 2 2022, balance-sheet details of Alameda Research were published, showing heavy investments in FTT. On the 6<sup>th</sup> Binance, a competitor exchange, announced it would sell its remaining FTT holdings; a few days later, on November 8, Binance communicated its intent to acquire FTX, only to back out of the rescue deal one day later. On November 10, the Bahamas Securities Commission froze FTX Digital Markets, assets, and finally on November 11, FTX, FTX US, and Alameda filed for bankruptcy.

The third large shock we investigate is the failure of the Silicon Valley Bank. SVB was the sixteenth-largest bank in the United States at the time of its bankruptcy in March 2023 (Metrick, 2024). Focusing on the tech and startup sector, SVB had several interlinkages with crypto-related entities, among which Circle, the company that manages USDC. We identify four sub-events taking place during the failure of Silicon Valley Bank. On March 8 2023, SVB announced a USD 1.8 billion loss on its securities; within the next day, depositors had withdrawn around 25% of total deposits. On March 10, the Federal Deposit Insurance Corporation (FDIC) took SVB into receivership, i.e. the FDIC took control of SVB, promising immediate access only to

insured deposits (\$250,000 per depositor). On the day after, Circle revealed that it had \$3.3 billion in deposits at SVB, causing USDC to depeg for several hours, reaching a value as low as \$0.87. Lastly, on the 12<sup>th</sup>, FDIC promised to protect uninsured depositors, which led to the end of the run on SVB and USDC.

Finally, in addition to the events described in Section 3, we construct a dataset of other events that occurred within the time window of interest and that arguably may have impacted the cryptoasset markets. We extract data on sector-specific events and broader large-scale shocks to the global economy from Koutrouli et al. (2025) and Coinmarketcap. We also gather data from official sources (US government, Federal Reserve, ECB, EU, WHO). For each event, we record the date, a brief description, additional information on the type of event, the reference, and an indicator whether it is considered a major or minor event for the crypto sector. The dataset (Table 8) covers economic shocks and monetary policy decisions, such as interest rate changes and policy interventions; geopolitical events like conflicts or elections; announcements and milestones, e.g., Bitcoin halvings or price thresholds; and regulatory events, such as the implementation of MiCA regulations. The three shocks we analyze do not overlap with any of these events.

| Date                                   | Identified Event                        | Major | Type       | Reference      |
|----------------------------------------|-----------------------------------------|-------|------------|----------------|
| Crypto-related announcements, events   |                                         |       |            |                |
| 2020-05-11                             | Bitcoin halving                         | ✗     | Technical  | Blockpit       |
| 2020-06-30                             | USDT reaches 10 bln\$ mkt. cap          | ✗     | Media      | Coinmarketcap  |
| 2020-10-21                             | Paypal enables BTC payments             | ✗     | Media      | ECB et al.     |
| 2021-02-07                             | 50,000USD/BTC breached                  | ✗     | Media      | Coinmarketcap  |
| 2021-02-15                             | Tesla, Mastercard support BTC           | ✗     | Media      | ECB et al.     |
| 2021-10-15                             | SEC approves BTC futures ETFs           | ✗     | Media      | Coinmarketcap  |
| 2021-11-11                             | BTC falls 65,000USD/BTC                 | ✗     | Media      | Coinmarketcap  |
| 2023-06-23                             | BlackRock files for Spot BTC ETF        | ✗     | Media      | ECB et al.     |
| 2024-01-10                             | 11 BTC ETF approved by SEC              | ✗     | Media      | ECB et al.     |
| 2024-04-20                             | Bitcoin halving                         | ✗     | Technical  | Blockpit       |
| 2024-07-03                             | USDT reaches 100 bln\$ mkt.cap          | ✗     | Media      | Coinmarketcap  |
| 2024-12-13                             | 100,000USD/BTC breached                 | ✗     | Media      | Coinmarketcap  |
| 2025-01-23                             | US EO supporting cryptoassets           | ✗     | Regulation | US govt.       |
| Econ. shocks, stimuli, monetary policy |                                         |       |            |                |
| 2020-03-15                             | Major FOMC Rate Cut (COVID-19)          | ✓     | Monetary   | Fed.Res.       |
| 2020-03-27                             | CARES Act (US) approved                 | ✓     | Policy     | US govt.       |
| 2020-07-21                             | EU agrees 750 bln recovery package      | ✓     | Policy     | EU             |
| 2021-03-11                             | American Rescue Plan Act approved       | ✗     | Policy     | US govt.       |
| 2021-07-13                             | First EU recovery disbursements         | ✗     | Policy     | EU             |
| 2022-03-16 -<br>2023-07-26             | FOMC rate increases                     | ✓     | Monetary   | Fed.Res.       |
| 2022-07-27 -<br>2023-09-20             | ECB interest rate increases             | ✗     | Monetary   | ECB            |
| 2024-09-18 -<br>2024-12-18             | FOMC rate cuts                          | ✓     | Monetary   | Fed.Res.       |
| 2024-06-11 -<br>2025-05-31             | ECB interest rate cuts                  | ✗     | Monetary   | ECB            |
| 2025-03-02                             | US announces ‘Strategic crypto reserve’ | ✗     | Monetary   | US govt.       |
| Geopolitical events                    |                                         |       |            |                |
| 2020-03-02                             | COVID-19: EC fully activates IPCR       | ✗     | Covid-19   | EU             |
| 2020-03-11                             | COVID-19 as pandemic (WHO)              | ✓     | Covid-19   | WHO            |
| 2020-11-09                             | COVID-19 vaccines are 90% effective     | ✓     | Covid-19   | Media          |
| 2021-09-07                             | BTC legal tender in El Salvador         | ✓     | Regulation | Alvarez et al. |
| 2021-09-24                             | China crypto ban                        | ✓     | Regulation | Coinmarketcap  |
| 2022-02-24                             | Russian invasion of Ukraine             | ✓     | War        | Media          |
| 2022-04-27                             | BTC legal tender in CAR                 | ✓     | Regulation | Coinmarketcap  |
| 2023-10-07                             | Gaza terrorist attack to Israel         | ✓     | War        | Media          |
| 2023-10-27                             | Israel invades Gaza                     | ✗     | War        | Media          |
| 2024-11-05                             | Trump elected US President              | ✓     | Regulation | Media          |
| 2025-01-29                             | El Salvador amends BTC law              | ✗     | Regulation | Media          |
| Regulatory events                      |                                         |       |            |                |
| 2020-09-01                             | Zug announces tax payments in BTC       | ✗     | Regulation | Media          |
| 2021-07-01                             | Germany approves Fund Locations Act     | ✗     | Regulation | ECB et al.     |
| 2023-04-20                             | MiCa EU Parliament approval             | ✓     | Regulation | EU             |
| 2024-06-30                             | MiCa applies for stablecoins            | ✓     | Regulation | EU             |
| 2024-12-30                             | MiCa applies for all cryptoassets       | ✓     | Regulation | EU             |

Table 8: **Timeline of key crypto-related events in the cryptocurrency sector, 2020-2025.** Economic shocks, stimuli, monetary policy; geopolitical events; regulatory events.

## B CASP Volume Flows

This section provides additional information on the cryptoasset flow analysis in Section 4.

### B.1 CASP-to-market activity

For each asset (BTC, ETH, USDT, and USDC), we construct hourly and daily time series of inflows, outflows, net flows, and number of transactions. Figure 9 shows net flows by cryptoasset. Three patterns emerge. First, Bitcoin dominates trading volume across cryptoassets: its daily net flows are larger than those of any other asset. Second, stablecoins played a marginal role until 2023, after which their daily variation increased substantially. Third, the series are stationary and exhibit no discernible trend, which serves as a quality check on the underlying data.

Table 9 reports descriptive statistics for values aggregated by unique transaction identifier and separated by asset. The distributions are highly right-skewed: most transactions range from hundreds to thousands of dollars, while aggregate volumes are driven by a limited number of very large transactions. The pattern is particularly pronounced for Bitcoin, with the maximum transaction value exceeding by one order of magnitude that of the other assets.




Figure 9: **Daily net flows of BTC, ETH, USDC, USDT in Austrian CASPs.** Each plot reports the net flows for one cryptoasset. The y-axis is in millions of US dollars.

### B.2 CASP-to-CASP

Table 10 reports the Austrian CASP-to-CASP interactions broken down by year and the relative Bitcoin dominance (where 1 indicates that all flows are in Bitcoin and 0 that no flows are in Bitcoin). Tables 11 and 12 respectively report the number of interactions and the volumes moved,

| Curr. | Count     | Mean      | St. Dev.   | 10%   | 25%    | 50%    | 75%      | 90%       | Max            |
|-------|-----------|-----------|------------|-------|--------|--------|----------|-----------|----------------|
| BTC   | 1,351,986 | 12,339.75 | 333,377.74 | 33.58 | 111.34 | 544.03 | 2,616.32 | 20,749.47 | 287,525,662.21 |
| ETH   | 1,383,287 | 3,539.92  | 84,402.38  | 22.00 | 65.03  | 206.44 | 712.44   | 5,292.05  | 20,843,506.00  |
| USDC  | 85,255    | 12,951.17 | 175,071.75 | 21.58 | 99.99  | 443.76 | 1,530.84 | 12,305.15 | 14,899,712.11  |
| USDT  | 394,115   | 17,834.19 | 163,148.29 | 42.13 | 111.58 | 449.89 | 1,384.37 | 13,582.27 | 13,509,941.63  |

Table 9: **Transaction values aggregated by transaction identifier.** Descriptive statistics. Values are in US dollars.

broken down by asset and year. To produce a visual representation, we construct a network where the nodes are the CASPs, and the edges are the cryptoasset transfers, aggregated at the quarterly level. Figure 10 shows illustrative examples for four different quarters: 2021Q1, 2022Q2, 2023Q1, and 2023Q3.

| Year | Number of interactions | BTC dominance | USD Value | BTC dominance |
|------|------------------------|---------------|-----------|---------------|
| 2019 | 70                     | (1.00)        | 22,420    | (1.00)        |
| 2020 | 717                    | (0.98)        | 373,578   | (0.97)        |
| 2021 | 851                    | (0.96)        | 439,860   | (0.94)        |
| 2022 | 712                    | (0.97)        | 1,719,969 | (0.99)        |
| 2023 | 334                    | (0.84)        | 1,558,845 | (0.80)        |
| 2024 | 277                    | (0.68)        | 1,181,264 | (0.76)        |
| 2025 | 79                     | (0.53)        | 477,751   | (0.81)        |

Table 10: **Austrian CASP-to-CASP interactions broken down by year.** Numbers in parenthesis indicate the fraction that can be attributed to Bitcoin flows.

| Year  | BTC  | ETH | USDC | USDT | Total |
|-------|------|-----|------|------|-------|
| 2019  | 70   | 0   | 0    | 0    | 70    |
| 2020  | 708  | 8   | 0    | 1    | 717   |
| 2021  | 824  | 27  | 0    | 0    | 851   |
| 2022  | 694  | 12  | 4    | 2    | 712   |
| 2023  | 281  | 28  | 10   | 15   | 334   |
| 2024  | 190  | 67  | 5    | 15   | 277   |
| 2025  | 42   | 26  | 11   | 0    | 79    |
| Total | 2809 | 168 | 30   | 33   | 3040  |

Table 11: **Number of CASP-to-CASP interactions broken down by cryptoasset and year.**



| Year  | BTC       | ETH     | USDC   | USDT   | Total     |
|-------|-----------|---------|--------|--------|-----------|
| 2019  | 22,420    | 0       | 0      | 0      | 22,420    |
| 2020  | 362,795   | 9,551   | 0      | 1,232  | 373,578   |
| 2021  | 416,329   | 23,531  | 0      | 0      | 439,860   |
| 2022  | 1,703,211 | 8,972   | 7,443  | 342    | 1,719,969 |
| 2023  | 1,253,409 | 268,179 | 13,918 | 23,339 | 1,558,845 |
| 2024  | 908,742   | 209,927 | 17,665 | 44,930 | 1,181,264 |
| 2025  | 391,626   | 62,387  | 23,739 | 0      | 477,751   |
| Total | 5,058,532 | 582,546 | 62,765 | 69,842 | 5,773,686 |

Table 12: **CASP-to-CASP** volumes broken down by cryptoasset and year.




Figure 10: **CASP-to-CASP interaction network.** Illustrative examples of interactions across Austrian CASPs aggregated at quarterly level. Top left: 2021, Q1; top right: 2022, Q2; bottom left: 2023, Q1; bottom right: 2023, Q3. Nodes represent CASPs, and their size is proportional to their relevance in the network. Edges represent value flows; the magnitude is in thousands of USD, and the arrow dimension is proportional to the magnitude of the flow. Arrow heads indicate flow direction and can be bidirectional (in this case, the magnitude is the sum of the two directions).

### B.3 within-CASP

In this section, we include additional details on the within-CASP analysis. We begin by describing the algorithms used to identify the patterns described in Section 4, focusing in particular on the identification of chained movements across addresses controlled by the same CASP. These algorithms follow a two-step process. In a first step, we construct network-based indicators to identify the presence of hub wallets within a CASP. Algorithm 1 describes the procedure. We first construct for each CASP and asset a directed graph  $G(E, V, W, C)$  where  $V$  is the set of wallet addresses, and  $E$  is the set of edges; a link exists if two addresses are recipient and sender in the same asset transfer.  $W$  and  $C$  represent the strength of the link and respectively measure the overall volume moved among address pairs and the number of interactions. We then compute a number of network statistics on  $G$  and identify, after manual inspection, the list of candidate hub wallets.

In a second phase, we detect chained movements across addresses controlled by the same CASP, where funds are ultimately routed to a hub wallet. Algorithm 2 describes the procedure for deposit addresses. For each CASP and asset, we first separate asset transfers directed to hub wallets (set  $T_1$ ) from other asset transfers (set  $T_2$ ); then, we match rows in the two sets such that the recipient of the rows in  $T_1$  are the same address sending funds to the hub wallet in the set  $T_2$ , under the condition that the two legs have the same transaction identifier and volume. We then flag with a binary indicator the rows flagged through this approach. We also develop a variant of the main algorithm to account for small time differences and minimal price changes when the chained flows are not executed within the same transaction. Finally, we note that Bitcoin and Ethereum chains work differently (the former is UTXO-based, while the latter is account-based). Since one Bitcoin typically consists of multiple input and output UTXOs, we first preprocess the asset transfer list by aggregating deposits with the same transaction identifier and originating from the same address, as well as withdrawals with the same transaction identifier and directed to the same address.

---

**Algorithm 1:** Hub Wallet Identification

---

**Data:** (1) Asset transfer list  $T = \{(s_i, d_i, v_i)\}$

where  $s_i$  = sender address,  $d_i$  = destination address,  $v_i$  = volume flow;

**Result:** Set of hub wallets  $H$

```
1  $P \leftarrow \text{Aggregate}(s_i, d_i);$  // Aggregate transactions by ordered pairs
2 foreach  $P(s, d)$  do
3    $W \leftarrow \text{Sum}(v_i);$ 
4    $C \leftarrow \text{Count}(v_i);$ 
5  $G(E, V, W, C);$  // Construct directed graph G
6 foreach  $v \in V$  do
7    $N_{ind} \leftarrow$  in-degree centrality;
8    $N_{outd} \leftarrow$  out-degree centrality;
9    $N_{pr} \leftarrow$  PageRank centrality;
10   $N_{ev} \leftarrow$  eigenvector centrality;
11   $N_{bw} \leftarrow$  betweenness centrality;
12   $N \leftarrow [N_{ind}, N_{outd}, N_{pr}, N_{ev}, N_{bw}]$ 
13  $H \leftarrow \text{filter}(V, \text{by}=N);$  // Identify hub wallets based on high values in  $N$ 
14 return  $H$ 
```

---

---

**Algorithm 2:** Matching Transaction Identification: ‘pattern 1’

---

**Data:** (1) Asset transfer list  $T = \{(tx_i, s_i, d_i, v_i, t_i)\}$ , where  $tx_i$  = transaction identifier,

$s_i$  = sender address,  $d_i$  = destination address,  $v_i$  = volume flow,  $t_i$  = time;

(2)  $H$  set of hub wallets

**Result:** Label set  $L = \{l_i\}, \quad l \in \{0, 1\}$


```
1  $L \leftarrow \{0\};$  // Init. label for all rows
2 foreach  $h \in H$  do
3    $T_1 \leftarrow \text{filter}(T, \text{by}=\text{recipient}, d_i = h \in H);$  // Select rows: recipient is hub
   wallet
4    $T_2 \leftarrow \text{filter}(T, \text{by}=\text{recipient}, d_i \neq h \in H);$  // Select other rows
5    $M \leftarrow \text{merge}(T_1, T_2, \text{on} = [tx_i, v_i]);$  // Match pairs
6    $M' \leftarrow \text{filter}(M, \text{by}=\text{recipient-sender}, d_{i1} = s_{i2});$  // Filter chained rows
7    $Tx \leftarrow \{tx_i \in M'\};$  // Obtain list of  $tx_i$  identified
8 for  $t_i \in T$  do
9   if  $tx_i \in Tx$  then
10     $l_i \leftarrow 1;$ 
11   else
12     $l_i \leftarrow 0;$ 
13 return  $L;$ 
```

---


Our analysis did not highlight changes of address usage from hot to warm or cold over time. We observed only one case of hub wallet that changed its use from hot to cold wallet used for large transactions. However, the change in its transaction activity is abrupt, and therefore it was possible to separate its activity before and after the change. This pattern is documented in the codebase and related documentation.

We also note that the algorithms described above enabled us to identify a pattern in Bitcoin transactions whereby funds are withdrawn immediately after being deposited by the same actors, for a total of \$300 million. We interpret this behavior as the result of the use of change addresses.


Finally, Figure 11 reports illustrative examples of the within-CASP networks of different CASPs and across different assets. In most of the reported examples, it is possible to identify one hot wallet acting as a hub with many other intermediary hot wallets, and a network of cold or warm wallets connected to the intermediary wallets through the hub wallet (violet and pink indicate hot wallets; orange and green respectively indicate warm and cold wallets). We also observed that in minor cases, Bitcoin assets were managed internally following peeling-chain-like patterns (see Panel 11b).




(a) Bitcoin




(b) Bitcoin




(c) Ethereum



(d) Ethereum



(e) USDC



(f) USDT

Figure 11: **Network visualization of within-CASP flows.** Violet and pink indicate hot wallets; orange and green respectively indicate warm and cold wallets.

## C Robustness checks

| Group           | Chain | Addresses | Avg. transactions |           | Avg. degree | Avg. USD  | Hot wallet<br>(fraction) |
|-----------------|-------|-----------|-------------------|-----------|-------------|-----------|--------------------------|
|                 |       |           | overall           | in sample | overall     | in sample |                          |
| RL              | ETH   | 528,400   | 29                | 2.6       | 9           | 1,566     | 1.00                     |
|                 | BTC   | 1,274,001 | 10                | 2.7       | 10          | 1,532     | 1.00                     |
| HI              | ETH   | 477       | 6,049,674         | 632       | 289,096     | 47,339    | 0.16                     |
|                 | BTC   | 149       | 2,775,969         | 13,991    | 1,451,378   | 38,903    | 0.34                     |
| LI <sub>c</sub> | ETH   | 547       | 942               | 10.1      | 16          | 77,747    | 0.02                     |
|                 | BTC   | 2,748     | 5.0               | 1.9       | 19          | 387,853   | 0.02                     |
| LI <sub>a</sub> | ETH   | 12,619    | 192,699           | 25        | 18,208      | 3,216     | 1.00                     |
|                 | BTC   | 43,750    | 12,467            | 83        | 7,814       | 680       | 1.00                     |

Table 13: **Classification robustness to threshold and activity measure.** Addresses are grouped into retail-like (RL), high-confidence institutional (HI), and low-confidence institutional (LI<sub>c</sub>, LI<sub>a</sub>), separately for each chain (Bitcoin, Ethereum). We use an alternative specification with a looser activity threshold (z-score > 2.5) and network degree centrality instead of transaction count. The classification is consistent with the results in Table 4.

This appendix reports robustness checks for the classification in Section 4 and the event study in Section 5. First, we report the classification results under an alternative specification, using a looser activity threshold (z-score > 2.5 instead of > 3) and network degree centrality in place of transaction count as the activity signal. Table 13 shows that the resulting groups are comparable to the baseline classification in Table 4: retail-like (RL) addresses still dominate by count, and the low-confidence institutional subgroups (LI<sub>c</sub>, LI<sub>a</sub>) retain distinct profiles. The main difference is that the looser threshold and degree-based measure flag more addresses as active; this is consistent with a mechanical effect rather than a change in the underlying classification logic, and the qualitative conclusions in Section 4 are unaffected.

Next, we report the full set of asset-flow responses for each event, completing the figures shown in the main text (Figures 5 to 7), so that every combination of asset, group, and flow direction is covered. For the Terra-Luna crash, Figure 12 adds ETH and USDT, again split by group and flow direction; the ETH response echoes BTC, with institutional deposits and withdrawals both rising after May 10, while USDT, like USDC, is volatile and lacks a persistent direction. For the FTX collapse, Figure 13 reports the complementary margins (retail-like deposits and institutional withdrawals) across all four assets; these are weaker and less systematic than the ones emphasized in the main text, confirming that the adjustment concentrated in retail-like withdrawals and institutional deposits. For the SVB failure, Figure 14 reports retail-like withdrawals and institutional deposits; BTC and ETH remain largely unresponsive on these margins as well, apart from an increase in ETH retail-like withdrawals on March 12, and the clean two-sided redemption pattern remains specific to the retail-deposit/institutional-withdrawal margins in USDC. In all three cases, the additional margins are consistent with the interpretation given in Section 5.

We then vary the length of the pre-event estimation window. Equation 1 estimates normal

activity over a fixed 45-day pre-event window. Figure 15 re-estimates the response using windows of 30, 35, 40, 50, 55, and 60 days, taking retail-like BTC withdrawals around the FTX collapse as the representative case. The estimated abnormal response is essentially unchanged across all six windows: withdrawals stay within their bootstrap interval through November 7 and rise sharply above it from November 8, peaking a few days later. The result is therefore not an artifact of the baseline window length.

Finally, we re-estimate the responses using transaction counts as the outcome variable in place of native asset units (Figure 16). Across all checks, the patterns documented in Section 5 are stable. The count results reproduce the asset-unit findings: at FTX, retail-like BTC withdrawal counts rise above their interval from November 8, while institutional deposit counts spike around November 9–10; at SVB, retail-like USDC withdrawal counts spike around Circle’s disclosure on March 11. Because the two measures tell the same story, the responses are not driven by a small number of unusually large transfers.




Figure 12: **Terra Luna crash: alternative asset-flow responses.** *Notes:* post-event coefficients from the Terra Luna event-study regressions, in asset units. Each panel shows coefficients for days 0 to 10 after the event, using a 45-day pre-event estimation window. The dashed horizontal line marks zero. Shaded areas indicate bootstrap intervals based on 1,000 replications.






Figure 13: **FTX collapse: alternative asset-flow responses.** *Notes:* post-event coefficients from the FTX collapse event-study regressions, in asset units. Each panel shows coefficients for days 0 to 10 after the event, using a 45-day pre-event estimation window. The dashed horizontal line marks zero. Shaded areas indicate bootstrap intervals based on 1,000 replications.




Figure 14: **SVB failure: alternative asset-flow responses.** *Notes:* post-event coefficients from the SVB failure event-study regressions, in asset units. Each panel shows coefficients for days 0 to 10 after the event, using a 45-day pre-event estimation window. The dashed horizontal line marks zero. Shaded areas indicate bootstrap intervals based on 1,000 replications.




Figure 15: **FTX collapse: BTC retail-like withdrawals, varying pre-event window.** *Notes:* post-event coefficients from the SVB failure event-study regressions, in asset units. Each panel shows coefficients for days 0 to 10 after the event, using a 45-day pre-event estimation window. The dashed horizontal line marks zero. Shaded areas indicate bootstrap intervals based on 1,000 replications. The response is stable across all six windows, indicating that the baseline result is not an artifact of the choice of pre-event window length.




Figure 16: **FTX and SVB: results with alternative outcomes (transaction count).** *Notes:* Each panel shows coefficients for days 0 to 10 after the event, using a 45-day pre-event estimation window. The dashed horizontal line marks zero. Shaded areas indicate bootstrap intervals based on 1,000 replications. These results reproduce the patterns obtained with asset-unit outcomes.