

One factor, heavy tails and fragile pegs: evidence from a ten-year panel of cryptoassets and stablecoins

Nilton Gomes Furtado

`nilton.furtado@ua.pt`

University of Aveiro

Daniel Ferreira Polonia

University of Aveiro

Jullia Vasconcelos Furtado

University of Algarve

Research Article

Keywords: Cryptoassets, Stablecoins, Tail risk, Co-movement, Peg stability, Data quality

Posted Date: October 6th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11093115/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

Cryptoassets and stablecoins are now held by institutional investors and supervised by financial authorities, yet evidence on their risk rests largely on small samples that end before the failures of 2022. This paper examines a panel of 45 cryptoassets and 10 US dollar stablecoins, observed daily from January 2016 to April 2026 (155,221 observations) using an open-source dataset. Five hypotheses derived from the theory of speculative prices, asset pricing, and the monetary economics of pegs are assessed using rank-based tests, non-overlapping blocks, a break test at an unknown date, and a fixed cohort of assets. Normality is rejected for every series, and 42 of 45 cryptoassets lost more than 90% of their peak value. Volatility is not rewarded in the cross-section but is associated with deeper drawdowns. All 990 pairwise correlations are positive, and one component explains more than half of the variance. Mean pairwise correlation rose from about 0.2 to about 0.6 in mid-2017, a shift that is larger in a fixed cohort and therefore not an artefact of sample composition. Mean deviations from par differ by a factor of 48 across stablecoins, and maximum deviations reach 9 to 204 times the mean. An audit shows that single-source errors and single crash days can dominate estimates of kurtosis and correlation. A two-block model that links stress in cryptoasset markets to peg stability is proposed.

JEL Classification G11 · G12 · G15 · C58 · E42

1 Introduction

Cryptoassets have moved from an experiment to an asset class held by retail and institutional investors and supervised by financial authorities. At the end of the period studied here, the market capitalisation of Bitcoin alone was 1.5 trillion USD, and the two largest US dollar stablecoins, Tether and USD Coin, had a combined capitalisation of 267 billion USD. Stablecoins have become the settlement asset of the cryptoasset market and a link between that market and the banking system, and they have been brought under dedicated legal regimes in the European Union (European Union, 2023) and in the United States (United States Congress 2025). For those who hold, manage, or supervise these instruments, three questions are fundamental: how risky cryptoassets are, how well they diversify one another, and how stable stablecoins are.

A substantial body of literature has addressed these questions, and its main results are well known. Returns on major cryptocurrencies are heavy-tailed, and their volatility exhibits clustering (Zhang et al., 2018; Ghosh et al., 2023). Expected returns are captured by market-specific factors (Liu & Tsyvinski, 2021; Liu et al., 2022). Coins are strongly connected to one another, and this connection tightens during periods of stress (Ji et al., 2019; Corbet et al., 2020; Smales, 2021). Stablecoins are not uniformly stable, and their stability depends on arbitrage, collateral and design (Lyons & Viswanath-Natraj, 2023; Duan & Urquhart, 2023; Ma et al., 2023).

Four limitations of this literature motivate the present paper. First, most studies cover a small number of assets over periods ending before 2021 and therefore do not include the failures of Terra and FTX in 2022 or the banking stress of March 2023, which were the most severe tests these markets have faced.

Second, cryptoassets and stablecoins are usually studied separately, although they are linked through collateral, funding and the flight of investors between them. Third, the rise in co-movement among cryptoassets is documented but has not been dated, and it has not been separated from changes in the set of observed assets, since almost every panel grows over time. Fourth, data quality is a known weakness of the field. Fewer than half of the papers published since 2017 were found to use correct data (Alexander & Dakos, 2020), and vendors disagree on prices, capitalisation and volume (Schwenkler et al., 2026), yet audits of the data and the sensitivity of results to single observations are seldom reported.

The paper addresses these limitations using a panel of 45 cryptoassets and 10 US dollar stablecoins, observed daily from 1 January 2016 to 30 April 2026, yielding 155,221 asset-day observations drawn from an open-source dataset. It asks four questions. Are heavy tails a property of every cryptoasset, and is volatility rewarded in the cross-section? How strong is the common factor in cryptoasset returns, when did the market become integrated, and does the higher co-movement in recent years reflect a change in the market or in the sample? How does the stability of the peg differ across stablecoins, and does the average deviation from par describe its risk? Finally, how sensitive are the answers to choices about data and estimators that are rarely reported?

Five hypotheses are derived from three bodies of theory: the statistical theory of speculative prices, asset pricing together with the theory of market integration and contagion, and the monetary economics of pegs and runs. They are assessed using methods suited to heavy-tailed data: rank-based measures of association, inference on non-overlapping blocks, a test for a break at an unknown date, a fixed cohort of assets that keeps the sample composition constant, and a day-by-day inspection of extreme observations. The hypotheses were written down after an initial descriptive pass over the data, as stated in Sect. 2, and the tests are to be read as an examination of whether known regularities hold in a larger and longer panel.

All five hypotheses are supported. Normality is rejected for each of the 55 series, and 42 of the 45 cryptoassets lost more than 90% from a peak at some point. Volatility is not rewarded: its rank correlation with compound growth is -0.19 , and it is associated with deeper drawdowns and more extreme days. All 990 pairwise correlations are positive, and a single component accounts for more than half of the variance. The mean pairwise correlation rose from about 0.2 to about 0.6 in mid-2017 during the boom, but not in the subsequent collapse, and the shift is larger in the fixed cohort than in the full panel, so it cannot be attributed to changes in the market's composition. Mean deviations from par differ by a factor of 48 across stablecoins, and the largest deviation of a coin is between 9 and 204 times its mean. The audit also shows that a single scaling error in the source raises an asset's excess kurtosis from 20 to 236, and that a single crash day produces a correlation peak that disappears under a rank-based estimator.

The paper makes four contributions. It extends the stylised facts of cryptoasset returns to a ten-year panel spanning three market cycles and including stablecoins. It dates the integration of the cryptoasset

market without imposing the date, and separates it from the effect of composition. It shows that the quality of a peg has to be judged by the tail and the sign of its deviations and not by their mean. Finally, it documents a reproducible audit of a widely used open data source and, by example, shows which checks are needed before reporting results on tail risk and co-movement. On this basis, it proposes a two-block model, in which each hypothesis corresponds to a parameter and the transmission of stress from cryptoassets to pegs can be tested.

The remainder of the paper is organised as follows. Section 2 presents the theoretical background and the hypotheses. Section 3 describes the data and the methods. Section 4 reports the results and assesses the hypotheses. Section 5 discusses the findings and sets out the proposed model. Section 6 concludes.

2 Theoretical background and hypotheses

The analysis draws on three bodies of theory: the statistical theory of speculative prices, which describes the distribution of returns; asset pricing and the theory of market integration and contagion, which describe how risk is rewarded and why assets move together; and the monetary economics of pegs and runs, which describes when a promise of convertibility at par holds. A fourth strand, on the quality of cryptoasset data, conditions what can be inferred from any of them. Each of the following subsections reviews the theory and evidence on cryptoassets and concludes with the hypothesis assessed in the Subsequent Section. 4.8. The hypotheses were written down in their present form after a first descriptive pass over the data. The tests are therefore to be read as an examination of whether regularities established in smaller samples and shorter periods hold in a panel of 45 cryptoassets and 10 stablecoins over more than 10 years, rather than as tests of novel predictions formulated in advance (Kerr, 1998).

2.1 Speculative prices and the distribution of returns

That changes in speculative prices are not normally distributed is one of the oldest results in empirical finance. Mandelbrot (1963) and Fama (1965) showed that the tails of price changes are too heavy for the Gaussian model, and the accumulated evidence was summarised by Cont (2001) as a set of stylised facts shared across markets: heavy tails, the absence of linear autocorrelation in returns, volatility clustering and the slow decay of autocorrelation in absolute returns. These facts motivated the models of conditional heteroskedasticity (Engle, 1982) and their extension to heavy-tailed innovations (Bollerslev, 1987).

There are reasons to expect the same properties to hold in stronger form for cryptoassets. Most of them generate no cash flows, so that their value cannot be anchored by discounting. In the model of Cong et al. (2021), the value of a token derives from platform users' transactional demand and is amplified by network externalities, making it sensitive to expectations about adoption. Empirically, Bitcoin is held mainly as a speculative investment and not as a medium of exchange (Baur et al., 2018); its price has

contained a substantial bubble component (Cheah & Fry, 2015), and cryptocurrency returns are driven by factors specific to the market, such as network growth, momentum and investor attention, and not by the factors that price traditional assets (Liu & Tsyvinski, 2021). Prices that depend on beliefs about adoption and attention, in markets that trade continuously and have a large retail presence, are prone to large revisions.

The evidence agrees. Heavy tails, volatility clustering, and the rejection of normality have been documented for the main cryptocurrencies (Zhang et al., 2018) and for broader universes that include tokens (Ghosh et al., 2023). Conditional heteroskedasticity models with long-run and short-run components fit Bitcoin best (Katsiampa, 2017), extreme value analysis shows large differences in tail risk across the major coins (Gkillas & Katsiampa, 2018), tail risk is conditional on the state of the other cryptocurrencies (Borri, 2019), and models that allow for changes of regime in volatility produce more accurate Value-at-Risk and expected shortfall forecasts (Maciel, 2021). Most of these studies cover fewer than 10 assets and end before 2021. If heavy tails are a property of the asset class and not of individual coins, they should be found in every series of a larger and longer panel.

H1 Daily log returns of cryptoassets are heavy-tailed, and normality is rejected for every asset in the panel.

2.2 Risk, return and the pricing of volatility

In the intertemporal capital asset pricing model, the expected excess return on the market is proportional to its conditional variance (Merton, 1973). The empirical record is mixed even for equities: French et al. (1987) find a positive relation between the expected market premium and predictable volatility, whereas Glosten et al. (1993) find a negative one once asymmetries are allowed for. In the cross-section, the theory gives no role to total volatility, since only systematic risk is priced (Ross, 1976), and the evidence runs against a premium: stocks with high idiosyncratic volatility earn low subsequent returns (Ang et al., 2006). Behavioural theory offers an explanation. Under cumulative prospect theory, investors overweight small probabilities of large gains, so that positively skewed securities are overpriced and earn low average returns (Barberis & Huang, 2008), and stocks with extreme positive daily returns are indeed followed by low returns (Bali et al., 2011).

Cryptoassets, with large positive skewness and a retail investor base, are a natural setting for such preferences. The evidence is divided. Grobys and Junttila (2021) find that cryptocurrencies with the highest maximum daily returns underperform those with the lowest by more than 1.5% per week, in parallel with equity markets, whereas Ozdamar et al. (2021) report a positive relation between maximum daily returns and expected returns. The factors that capture the cross-section of expected cryptocurrency returns are market, size and momentum (Liu et al., 2022; see also Liu et al. 2020), and whether idiosyncratic volatility is priced has been examined separately (Zhang & Li, 2020). None of these results implies a premium for total volatility. A further consideration applies to any sample defined by survival. Brown et al. (1992) show that survivorship-based truncation induces a spurious positive relationship between volatility and returns, because volatile assets survive only if they perform well. A

sample of survivors therefore biases the analysis against the absence of a relation, which makes such an absence more informative when it is found.

H2 Across cryptoassets, higher total volatility is not associated with higher realised growth; it is associated with larger tail losses.

2.3 Common factors and market integration

Factor models represent returns as exposures to a small number of common sources of risk plus idiosyncratic components (Ross, 1976), and the strength of the common component measures the degree of market integration. Several mechanisms point to a dominant common factor among cryptoassets. A market factor is the first of the three factors of Liu et al. (2022). Measures of connectedness based on variance decompositions (Diebold & Yilmaz, 2012) show that return and volatility spillovers among cryptocurrencies are large and that Bitcoin is a dominant transmitter (Koutmos, 2018), that a broad set of 52 cryptocurrencies is tightly interconnected (Yi et al., 2018), and that the degree of integration varies over time (Ji et al., 2019). Spillovers run from Bitcoin and Ether to the rest of the market, stablecoins included (Smales, 2021), and the equicorrelation literature treats a single average correlation as an adequate summary of the market (Bouri et al., 2021; Demiralay & Golitsis, 2021).

Behaviour and market structure reinforce the common component. Herding among cryptocurrency investors is significant, varies over time, and intensifies with uncertainty (Bouri et al., 2019; Kallinterakis & Wang, 2019). Rising retail attention to Bitcoin in particular makes the prices of different coins move more synchronously (Koch & Dimpfl, 2022). Many coins are quoted and funded against the same few instruments: purchases with Tether were followed by sizeable increases in the prices of Bitcoin and other cryptocurrencies during the boom of 2017 (Griffin & Shams, 2020), and price deviations across exchanges co-move, which points to common arbitrage capital (Makarov & Schoar, 2020). At the same time, cryptocurrencies were found to be relatively isolated from other financial assets (Corbet et al., 2018), suggesting a factor specific to the asset class.

H3: Returns of cryptoassets are driven by a dominant common factor: all pairwise correlations are positive, and the first principal component accounts for more than half of total variance.

2.4 Time-varying co-movement, contagion and liquidity

Theory predicts that co-movement is not constant and rises under stress. When the traders who provide liquidity are financed with margin, losses tighten their funding, force sales across all the assets they hold and produce liquidity spirals with commonality across securities (Brunnermeier & Pedersen, 2009). Contagion can also arise as a wealth effect, when convergence traders who suffer losses in one market liquidate positions in another (Kyle & Xiong, 2001). Both mechanisms are plausible in cryptoasset markets, where leverage is widely available, and the same participants are active across coins. For equities, correlations are higher in falling than in rising markets (Longin & Solnik, 2001; Ang & Chen, 2002).

Measuring this effect is delicate. Forbes and Rigobon (2002) show that correlation coefficients are upwardly biased when volatility rises, so higher measured correlation in a crisis need not indicate a change in transmission, and Longin and Solnik (2001) warn of a spurious relationship between correlation and volatility. For cryptoassets, conditional correlations are reported to be persistently high (Ghorbel & Jeribi, 2021), to peak during the March 2020 sell-off (Smales, 2021), and to exhibit contagion effects during the COVID-19 period (Corbet et al., 2020). What the literature has not established is when the market became integrated, and whether the higher co-movement in recent years reflects a change in the market or in the set of assets observed.

H4 Co-movement among cryptoassets is time-varying, moved to a persistently higher level after 2017 and is highest in windows that contain episodes of market stress.

2.5 Stablecoins as pegs: arbitrage, runs and design

A stablecoin is a promise of convertibility at par, and three classical theories describe when such a promise fails. In first-generation models of currency crises, a peg that is inconsistent with the issuer's policy collapses under a speculative attack once reserves fall to a critical level (Krugman, 1979). In second-generation models, the attack is self-fulfilling, so that a peg can fail even when fundamentals would have sustained it (Obstfeld, 1996). In the theory of bank runs, an issuer of liquid claims backed by less liquid assets is exposed to a run in which it is rational to withdraw because others withdraw (Diamond & Dybvig, 1983). Gorton and Zhang (2021) apply this reasoning to stablecoins, which they describe as privately issued money comparable to the notes of the free-banking era.

Recent work develops these ideas for stablecoins. Parity of reserve-backed coins is maintained by arbitrage and demand-side forces, not by issuance; premiums reflect stablecoins' role as a safe haven within the cryptoasset economy, whereas discounts reflect liquidity and collateral concerns (Lyons & Viswanath-Natraj, 2023). Arbitrage is deliberately concentrated in a few authorised agents, because more efficient arbitrage stabilises the secondary-market price but increases the risk of a run (Ma et al., 2023). Investors move from riskier to safer stablecoins during periods of stress, and redemptions accelerate once the price falls below 0.99 USD, analogous to money market funds that break the buck (Anadu et al., 2023). Deviations of the crypto-collateralised DAI reflect the risk of its collateral and fell sharply when conversion into fiat-backed collateral at par became possible (Kozhan & Viswanath-Natraj, 2021). Designs without safe collateral are more fragile in theory (d'Avernas et al., 2026) and in practice, as demonstrated by the collapse of TerraUSD in May 2022 (Briola et al., 2023; Liu et al., 2023). Empirically, stablecoins are not uniformly stable (Duan & Urquhart, 2023). Taken together, these theories predict a peg that holds tightly on ordinary days, when arbitrage is cheap, and fails discontinuously when redemption is in doubt. The distribution of deviations should then have a small mean and a heavy tail, and should differ across designs.

H5 Peg stability is heterogeneous across stablecoins and peg deviations are heavy-tailed, so that the mean deviation understates tail risk.

2.6 Data quality as a condition of inference

Empirical work on cryptoassets depends on data of uneven quality. Alexander and Dakos (2020) find that fewer than half of the papers published since 2017 use correct data; Schwenkler et al. (2026) document mislabelling and large discrepancies across vendors in prices, market capitalisation and volume, and wash trading has been estimated at more than 70% of the volume reported by unregulated exchanges (Cong et al., 2023). Heavy-tailed statistics are particularly exposed, because a single erroneous observation can dominate a fourth moment or a correlation coefficient. This literature motivates the audit of the data in Sect. 3.4, the preference for rank-based statistics and the day-by-day inspection of extreme observations, and it is the reason why the results of Sect. 4 are reported together with the checks that they survived.

3 Data and methods

3.1 Research design

The study is a descriptive and confirmatory analysis of a longitudinal panel of daily observations on cryptoassets and stablecoins. It proceeds in four steps: construction and audit of the panel (Sect. 3.2 to 3.4), preliminary diagnostics and description of returns, co-movement and peg stability (Sect. 3.5 to 3.7), assessment of the five hypotheses of Sect. 2 with tests that make few distributional assumptions (Sect. 3.8), and specification of a model for subsequent estimation (Sect. 5.3). No parametric model is estimated in this paper. Rank-based statistics and medians are preferred throughout, because the returns are heavy-tailed and a small number of observations were found to be source artefacts. All statistics are computed by scripts from a single versioned master file, so that every number in Sect. 4 can be regenerated (Sect. 3.9).

3.2 Data source and sample

Data are taken from the Coin Metrics Community Data repository (Coin Metrics, 2026), which publishes daily market and network metrics by asset under the Creative Commons Attribution-NonCommercial 4.0 licence. The data vintage is 11 June 2026, and the sample runs from 1 January 2016 to the cutoff of 30 April 2026. The panel contains 55 assets: 45 cryptoassets and 10 stablecoins pegged to the US dollar (USDT, USDC, DAI, SAI, TUSD, PAX, GUSD, BUSD, HUSD and USDK). The universe is defined by coverage in the source **[authors: state the inclusion rule of the pipeline specification, for example, the minimum history and the availability of a price series]**. The shortest history is 1,816 days for a cryptoasset (ICP) and 708 days for a stablecoin (SAI). A universe defined in this way does not contain assets that failed and were dropped by the source, a limitation addressed in Sect. 4.1 and 5.1.

The master file has 155,221 asset-day observations and is unbalanced by construction. Markets for these assets trade every calendar day, so the calendar has seven days a week, and each observation refers to the close at 00:00 UTC. Only observed days are stored, and no value is imputed. The variables used in this paper are defined in Table 1. The price is the source's USD price metric, with the reference rate as a fallback that was never needed. Market capitalisation is the current-supply measure of the source, with the estimated measure as a fallback, and applies to 11,467 asset-days (8.6% of cryptoasset

observations); it is missing for 922 asset-days of BNB, EOS, and TRX. Volume is the aggregate spot volume reported by exchanges. It is classified as lower-quality data due to potential wash trading (Cong et al., 2023) and is used only as auxiliary evidence in anomaly inspection.

Table 1
Variables of the daily master file used in the analysis

Variable	Definition	Unit	Source or formula
Price	Closing price at 00:00 UTC	USD	PriceUSD, fallback to reference rate
Market capitalisation	Price times current supply	USD	CapMrktCurUSD, fallback CapMrktEstUSD
Reported volume	Aggregate spot volume reported by exchanges	USD	volume_reported_spot_usd_1d
Circulating supply	Units of the asset in circulation	native units	SplyCur
Log return	Daily logarithmic return	fraction	Eq. 1
Simple return	Daily simple return	fraction	$P_t/P_{t-1} - 1$
Outlier flag	Absolute simple return above 50%; retained and flagged	indicator	computed
Peg deviation	Deviation of the price of a stablecoin from 1 USD	USD	$P_t - 1$

Metric names are those of the source. The master file also contains network metrics and technical indicators that are not used in this paper

3.3 Measures of return, risk and peg stability

The daily log return of asset i is

$$r_{it} = \ln(P_{it} / P_{it-1}) \quad (1)$$

and is undefined on the first observation of each series. With $\bar{r}_i$ the mean and s_i the sample standard deviation (with divisor $n - 1$) of the daily log returns of asset i over its full history, the compound annual growth rate and the annualised volatility are

$$\text{CAGR}_i = \exp(365 \bar{r}_i) - 1 \quad (2)$$

$$\sigma_i = s_i \sqrt{365} \quad (3)$$

The factor of 365 reflects continuous trading. The Sharpe ratio is computed with a zero risk-free rate as $365 \bar{r}_i / \sigma_i$, which equals $\ln(1 + \text{CAGR}_i) / \sigma_i$. Because it is based on log returns, it is lower than the ratio

based on arithmetic returns, and it is used only to rank assets. Skewness and excess kurtosis are the bias-adjusted sample estimators. The maximum drawdown is the largest loss from a running peak,

$$\text{MDD}_i = \min(\text{Pit} / \max_{\{s < t\}} \text{P}_{is} - 1) \quad (4)$$

Outlier days are days with an absolute simple return above 50%. They are retained because extreme returns are a genuine property of these markets and are excluded only where day-by-day inspection identifies a source artefact (Sect. 3.4). For the cross-sectional analysis, the 45 cryptoassets are sorted into quartiles of annualised volatility and terciles of market capitalisation at the cutoff. The second sort uses information from the end of the sample and is reported only to show that it cannot be read as a size effect (Sect. 4.3).

For stablecoin j , the absolute deviation from par, in basis points, is

$$d_{jt} = 10^4 |P_{jt} - 1| \quad (5)$$

Peg stability is summarised by the mean, the 95th percentile and the maximum of d_{jt} , by the share of days on which it exceeds 50 bps and 1%, and by the number of days on which it exceeds 5%, the threshold above which a deviation is classed as severe. The sign of the deviation is used to separate premiums from discounts on days above 50 bps. BUSD is measured up to 15 December 2023, the date on which Binance ceased supporting the coin (Binance, 2023), because quotes after that date come from a residual, illiquid market.

3.4 Data quality procedures

Three procedures were applied before any statistic was interpreted. First, the master file was audited for internal consistency: uniqueness of asset-date keys, continuity of the calendar inside the window of each asset, undefined returns on first observations only, reproduction of the stored log returns from prices to a tolerance of 10^{-9} , reproduction of derived fields, and absence of non-positive prices. Second, prices were compared with an independent second source for the five assets for which one was available (BTC, DAI, TUSD, USDC and USDT) **[authors: identify the provider of the second source]**, through the relative difference $|P_{\text{master}} - P_{\text{second}}| / P_{\text{second}}$ and the number of days on which it exceeds 1% and 5%. Third, extreme observations were reviewed day by day alongside reported volume and circulating supply. This was done for the two largest absolute log returns in the panel and for every stablecoin deviation above 5% after 2020. An observation is treated as a suspected source artefact when a shift in the price level is reversed exactly without an increase in volume, or when it is not corroborated by the second source. Suspected artefacts are kept in the master file, explicitly reported, and the affected statistics are provided with and without them.

3.5 Preliminary time-series diagnostics

The daily log returns of each asset were subjected to a battery of tests that motivates the choices made later: the augmented Dickey-Fuller test of a unit root, with a constant and the lag length chosen by the

Akaike criterion (Dickey & Fuller, 1979); the KPSS test of level stationarity with automatic bandwidth (Kwiatkowski et al., 1992); the Jarque-Bera test of normality (Jarque & Bera, 1987); the Ljung-Box test with 10 lags on returns and on squared returns (Ljung & Box, 1978); and the Lagrange multiplier test for autoregressive conditional heteroskedasticity with 10 lags (Engle, 1982). Because each test is repeated across 45 cryptoassets and 10 stablecoins, p-values are adjusted within each family of tests and each asset class using the Benjamini-Hochberg (1995) procedure, with a false discovery rate of 5%.

3.6 Static co-movement

Co-movement among the 45 cryptoassets is first measured using Pearson correlations of daily log returns, computed pairwise over all overlapping days, with a required minimum of 250 days, which is never binding (the shortest overlap is 1,779 days). A matrix estimated pairwise is not guaranteed to be positive semidefinite. The matrix is therefore also estimated jointly using Pearson and Spearman correlations, on the common sample of 1,779 days for which all 45 assets are observed (12 May 2021 to 25 March 2026). The factor structure is summarised by the eigenvalues of the correlation matrix. The share of total variance of the first principal component is λ_1 / N , where λ_1 is the largest eigenvalue and $N = 45$, and the loadings are the elements of the first eigenvector multiplied by $\sqrt{\lambda_1}$. The average correlation of an asset is its mean correlation with the other 44 assets, excluding the diagonal.

3.7 Dynamic co-movement

The evolution of co-movement is measured by the mean pairwise correlation in a trailing window of 90 days,

$$\bar{\rho}_t = 2 / [N_t (N_t - 1)] \sum_{i < j} \rho_{ij,t} \quad (6)$$

where $\rho_{ij,t}$ is the correlation between assets i and j over the 90 days that end on day t , and N_t is the number of cryptoassets with a complete window. Pairs are equally weighted, and the series is sampled weekly. A 90-day window balances the precision of each correlation with the ability to track changes. Two design choices address the threats to inference identified in Sect. 4.5. To separate a change in co-movement from a change in the composition of the panel, the series is computed for all assets and for the fixed cohort of the nine cryptoassets observed from 1 January 2016. To limit the influence of single extreme days, it is computed with Spearman rank correlations as well as with Pearson correlations.

Overlapping windows are serially dependent by construction. Inference is therefore based on 41 consecutive, non-overlapping 90-day blocks. Blocks that end on or before 31 December 2017 (8 blocks) are compared with blocks that begin on or after 1 July 2018 (30 blocks) using a one-sided Mann-Whitney test (Mann & Whitney, 1947); the three transition blocks are excluded. The shift date is not imposed. It is estimated by the test of Andrews (1993) for a single break in the mean at an unknown date,

$$\sup_k F = \max_k F(k), F(k) = (SSR_0 - SSR_k) / [SSR_k / (T - 2)] \quad (7)$$

where SSR_0 is the sum of squared deviations from the overall mean, SSR_k is the sum of squared deviations from the two sub-sample means when the break is placed at block k , $T = 41$, and k ranges over the central 70% of the blocks. The relation between co-movement and market conditions is measured, over the blocks from mid-2018, by the rank correlation of $\bar{\rho}$ with the annualised volatility and with the cumulative return of the equally weighted market portfolio of the assets in the block.

3.8 Assessment of the hypotheses

Table 2 maps each hypothesis to its variables, test and criterion. Associations in the cross-section of 45 assets and 10 stablecoins are measured using Spearman rank correlations, with confidence intervals derived from the Fisher transformation. The significance of an individual Pearson correlation r estimated on n days is judged by:

$$t = r \sqrt{(n - 2) / (1 - r^2)} \quad (8)$$

evaluated at the shortest overlap. Tests are two-sided at the 5% level unless stated otherwise. The five hypotheses are assessed separately, with no adjustment across them, whereas the Benjamini-Hochberg adjustment is applied whenever a test is repeated across assets. Each hypothesis is classed as supported, supported with qualification or not supported. The robustness of the results is examined by removing the two most influential assets (BTG and BNB) from the cross-sectional correlations, by the fixed cohort and the rank-based estimator in the dynamic analysis, and by the common sample in the static analysis.

Table 2
Operationalisation of the hypotheses

H1	Variables	Test	Criterion for support
	Daily log returns of each cryptoasset	Jarque-Bera test by asset, Benjamini-Hochberg adjustment; excess kurtosis	Normality rejected for every asset
H2	Annualised volatility, CAGR, maximum drawdown, outlier days, excess kurtosis	Spearman rank correlations across 45 assets, Fisher confidence interval; volatility quartiles	No positive relation with CAGR; significant relation with measures of tail loss
H3	Pairwise and common-sample correlation matrices	Sign and t statistic of 990 correlations (Eq. 8); share of variance and loadings of the first principal component	All correlations positive and significant; first component above 50%
H4	Mean pairwise correlation (Eq. 6), all assets and fixed cohort, Pearson and Spearman	Mann-Whitney test on non-overlapping blocks; break test at unknown date (Eq. 7); rank correlation with market volatility	Higher level after the transition in all four series; positive relation with volatility
H5	Absolute peg deviation (Eq. 5); daily log returns of stablecoins	Dispersion of mean deviation; ratio of maximum to mean deviation; Jarque-Bera test and excess kurtosis	Large dispersion across coins; heavy tails for every coin

CAGR is the compound annual growth rate. The criteria were fixed together with the hypotheses, after a first descriptive pass over the data (Sect. 2)

3.9 Software, reproducibility and use of AI tools

The analysis was carried out in Python 3.11 with pandas 3.0, NumPy 2.4, SciPy 1.17, statsmodels 0.15 and Matplotlib 3.10. The scripts that produce the figures, the co-movement tests, and the diagnostics are provided as Online Resources, together with the derived tables, and regenerate every statistic in Sect. 4 from the master file. A large language model (Claude, Anthropic) was used under the direction of the authors to write the analysis code, generate the figures from that code, recompute and cross-check the statistics, and draft parts of the text. The literature was searched with Elicit and scite, and every reference was checked against its bibliographic record. The authors reviewed and verified all content and take full responsibility for it.

4 Results

4.1 Data integrity and panel structure

The master file contains 155,221 asset-day observations across 55 assets: 45 cryptoassets (132,592 observations) and 10 US dollar stablecoins (22,629 observations), between 1 January 2016 and 30 April 2026. The panel passes every internal consistency check applied to it. There are no duplicate asset-date

keys, no missing calendar days inside the observation window of any asset, and no non-positive prices. Stored log returns reproduce $\ln(P_t/P_{t-1})$ with a maximum absolute residual of 5.6×10^{-17} , the moving-average and Bollinger fields are reproduced exactly from the price series, and returns are undefined on exactly the first 55 observations, so that no return is computed across an asset boundary. The full checklist is reported in Online Resource 1.

Three features of the file are informative rather than erroneous. First, the price provenance field takes a single value in all rows. The fallback chain was therefore never exercised, and the source-divergence flag, which is false throughout, should be read as inert rather than as evidence of agreement between sources. Second, the gap flag is uniformly false because only observed days are materialised; gap accounting must rely on the first and last dates of each series. Third, 139 asset-days spread over 37 assets show an absolute simple return above 50%. These observations were retained rather than winsorised, which is the conservative default for a research panel but has consequences for the moment estimates discussed in Sect. 4.2.

Because published work on cryptoassets has been shown to depend heavily on the choice and quality of the price source (Alexander & Dakos, 2020; Schwenkler et al., 2026), prices were compared with an independent second source for the five assets for which one was available (Table 3). The comparison covers 15,515 overlapping asset-days, or 10.0% of the panel. Median relative differences range from 1.4×10^{-4} (USDC) to 1.3×10^{-3} (BTC), and no day differs by more than 5% after 2017. Two disagreements require action. For USDT, the 200-day values in 2016 and 2017 differ by more than 5%, with a maximum relative difference of 9.2%. The second source records a constant value of 0.0999 for most of 2016, while the master file records values close to 1.00, indicating that the comparison file is defective. For DAI, the master file records 1.0908 and 1.0724 on 30 April and 1 May 2025, both within 10 basis points (bps) of par, whereas the second source remains at 1.0000. Circulating supply is unchanged over the two days (3.25 billion), so the parallel move in market capitalisation is a mechanical product of price and supply, carrying no independent information. Reported volume is unremarkable on 30 April. Here the master file is the suspect series. Until the two observations are verified against a third source, they are treated as unconfirmed and are excluded from all statements about recent peg breaches.

Table 3
Agreement between the master price series and an independent second source

Asset	Overlapping days	Median rel. diff.	95th pct. rel. diff.	Days > 1%	Days > 5%
BTC	3,773	1.31×10^{-3}	1.05×10^{-2}	206	0
DAI	2,354	2.85×10^{-4}	3.50×10^{-3}	12	2 ^a
TUSD	2,856	5.63×10^{-4}	7.83×10^{-3}	106	0
USDC	2,759	1.39×10^{-4}	7.24×10^{-3}	90	0
USDT	3,773	4.35×10^{-4}	5.48×10^{-2}	463	200 ^b

Relative difference is $|P_{\text{master}} - P_{\text{second}}| / P_{\text{second}}$.^a 30 April and 1 May 2025; the master series is the suspect one. ^b All in 2016 (159 days) and 2017 (41 days); the second source is the suspect one

Figure 1 shows the coverage of each series. The panel is unbalanced by construction: 10 series are observed from the first day of the sample; 3 more enter in 2016, 14 in 2017, 9 in 2018, 7 in 2019, 11 in 2020, and 1 in 2021. No series has an interior gap. Four series end before the cutoff: SAI on 30 November 2019, following its migration to multi-collateral DAI; HUSD on 17 November 2022; USDK on 17 April 2023; and MKR on 25 March 2026. The last of these ends 36 days before the cutoff and is more plausibly a truncation of source coverage than a delisting. As a consequence, the 55 series are never observed simultaneously. The maximum is 54, between 11 May 2021 and 17 November 2022, while all 45 cryptoassets are jointly observed from 11 May 2021 to 25 March 2026.

Two implications follow for the remainder of the section. Cross-sectional statistics are computed over asset-specific windows, so that they are not measured over common periods, and the number of cryptoassets available rises from 9 at the start of the sample to 12 at the end of 2016 and 25 at the end of 2017. In addition, the universe is defined by the source's coverage at the data vintage rather than by a point-in-time reconstitution. Assets that failed and were dropped by the source are absent, biasing the average performance upward. Studies that include defunct coins do so precisely to avoid this problem (Liu et al., 2022). The performance figures reported below should therefore be read as upper bounds for the investable universe.

4.2 Return distributions and drawdowns

Table 4 summarises the cross-sectional distribution of the per-asset statistics for the 45 cryptoassets; the asset-level values are given in Online Resource 2. Annualised volatility of daily log returns has a median of 111.7% and ranges from 67.5% (BTC) to 188.1% (BTG). Excess kurtosis exceeds 3 for every asset, with a median of 12.1 and a range from 3.4 (AAVE) to 235.8 (BTG). Heavy tails are therefore a property of the entire cross-section, not of a few assets, which is consistent with the stylised facts documented in smaller samples and earlier periods (Zhang et al., 2018; Ghosh et al., 2023). Skewness is positive for 29 of the 45 assets (median 0.26), so the heavy tail is not predominantly a left tail. Any model fitted to these returns requires a heavy-tailed or conditionally heteroskedastic specification, because Gaussian standard errors would be severely understated.

The diagnostics of Sect. 3.5 confirm this picture. After adjustment for multiple testing, the unit root is rejected, and level stationarity is not rejected for all 55 series, so that log returns can be treated as stationary. Normality is rejected for all 55 series, with Jarque-Bera statistics ranging from 955 (AAVE) to 7.2×10^6 (BTG) for the cryptoassets. The Ljung-Box test on squared returns rejects the null of no dependence for all 55 series, and the ARCH test for 54 of them (the exception is HT, $p = 0.06$), thereby establishing volatility clustering across the panel. Returns themselves are serially correlated for 32 of the 45 cryptoassets and for all 10 stablecoins, the latter being the expected consequence of mean reversion towards the peg. The full results are given in Online Resource 5.

The extreme values deserve a qualification. The largest positive and negative daily log returns in the panel, + 2.43 and – 2.38, both belong to BTG. Its price moves from 0.55 to 6.29 USD on 16 January 2026, stays near 6.3 USD for nine days, and returns to 0.58 USD on 25 January 2026. Over the same days, reported volume falls from about 2.8 million USD to about 0.6 million USD, then to less than 2,000 USD. An elevenfold shift in level, accompanied by falling volume, that is reversed exactly is the signature of a scaling or stale-feed error in the source, not of trading. The two days account for almost all of the reported kurtosis: without them, the excess kurtosis of BTG falls from 235.8 to 20.0 and its annualised volatility from 188.1% to 147.7%. The nine days should be flagged and excluded before the BTG moments are used; therefore, the summary statistics in Table 4 are based on medians.

Drawdowns are close to total. The median maximum drawdown is – 97.6%, and 42 of the 45 assets lost more than 90% from a running peak at some point in their history. The mildest drawdowns are for TRX (– 79.1%), BNB (– 80.0%), and BTC (– 83.8%). Any buy-and-hold interpretation of the returns in Sect. 4.3 has to be set against these figures.

The summary file also reports return-based statistics for the 10 stablecoins. These are computed mechanically and describe a peg rather than an investment; for instance, BUSD shows an annualised volatility of 74.4% that is entirely attributable to residual trading after Binance ended its support for it. They are not interpreted here. The excess kurtosis of the stablecoin series (between 15 and 445) is nevertheless meaningful, since a peg that holds on most days and breaks violently on a few is by construction a high-kurtosis process. Peg quality is analysed with dedicated measures in Sect. 4.6.

Table 4

Cross-sectional distribution of per-asset statistics for the 45 cryptoassets, full history of each asset

Statistic	Minimum	1st quartile	Median	3rd quartile	Maximum
CAGR (%)	–63.7	–26.8	–4.8	29.8	177.8
Annualised volatility (%)	67.5	105.3	111.7	124.0	188.1
Sharpe ratio (zero risk-free rate)	–0.93	–0.29	–0.04	0.22	0.97
Skewness	–3.25	–0.16	0.26	1.09	3.61
Excess kurtosis	3.4	8.2	12.1	25.3	235.8
Worst daily log return	–2.38	–0.58	–0.51	–0.47	–0.33
Best daily log return	0.22	0.43	0.56	0.70	2.43
Maximum drawdown (%)	–99.9	–98.8	–97.6	–94.5	–79.1
Outlier days	0	1	2	4	12
Days observed	1,816	2,446	3,110	3,471	3,773

CAGR is the compound annual growth rate. Volatility is the annualised standard deviation of daily log returns. The Sharpe ratio is computed using a zero risk-free rate as the annualised mean log return divided by annualised volatility, yielding a lower value than the ratio based on arithmetic returns. Outlier days are days with an absolute simple return above 50%

4.3 Cross-sectional relation between risk and return

Figure 2 plots annualised volatility against the compound annual growth rate (CAGR) for the 45 cryptoassets. There is no positive relation between the two. The Spearman rank correlation is -0.19 ($p = 0.22$, $n = 45$) and the Pearson correlation is -0.26 ($p = 0.08$). Sorting the assets into volatility quartiles (Table 5, panel a) produces no monotonic ordering of median CAGR, which is 16.3% in the lowest quartile and -19.9% in the highest. What increases with volatility is tail risk: median excess kurtosis rises from 12.5 to 25.6, and the median maximum drawdown deepens from -94.3% to -99.1% between the lowest and highest quartiles.

Dispersion is more informative than the average. A total of 25 of the 45 assets have a negative full-sample CAGR, and only five reach a Sharpe ratio above 0.5 at a zero risk-free rate: BNB (0.97), ETH (0.78), BTC (0.74), XMR (0.62) and DOGE (0.50). BNB is the return outlier, with a CAGR of 177.8%, about 65 percentage points above ETH's (112.8%). BTC is the least volatile asset in the panel.

The flat relation should not be read as evidence on the pricing of risk. It is a statement about total volatility measured ex post over unequal windows, whereas the asset-pricing literature on this market works with factor exposures and portfolio sorts. That literature finds that market, size and momentum factors capture the cross-section of expected cryptocurrency returns (Liu et al., 2022; see also Liu et al. 2020) and has examined separately whether idiosyncratic volatility is priced (Zhang & Li, 2020). The present result is compatible with those findings. It shows only that, among assets covered by the source until 2026, higher realised volatility was associated with heavier tails and deeper drawdowns and not with higher realised growth.

Size does order the cross-section (Table 5, panel b). Median CAGR is 47.4% in the largest tercile by market capitalisation, -6.9% in the middle tercile and -25.6% in the smallest, and the rank correlation between log market capitalisation and CAGR is 0.69 ($p < 0.001$). This association is mechanical and carries a look-ahead bias, because market capitalisation is measured at the end of the sample: assets are large because they compounded. It points in the opposite direction to the size premium reported by Liu et al. (2022), who find that small coins earn higher returns based on point-in-time sorts, and it cannot be interpreted as a size factor. A test of that claim requires point-in-time market capitalisation, which the present summary files do not contain.

Table 5
Median statistics of the 45 cryptoassets by volatility quartile and by size tercile

Group	<i>n</i>	Volatility (%)	CAGR (%)	Excess kurtosis	Max. drawdown (%)
<i>Panel a: volatility quartiles</i>					
Q1 (lowest volatility)	12	101.1	16.3	12.5	−94.3
Q2	11	109.5	−11.8	8.5	−97.9
Q3	11	118.8	3.0	16.9	−96.2
Q4 (highest volatility)	11	133.4	−19.9	25.6	−99.1
<i>Panel b: size terciles</i>					
Large	15	105.0	47.4	14.6	−94.0
Medium	15	111.5	−6.9	9.0	−97.1
Small	15	124.0	−25.6	11.1	−99.2

All entries are group medians. Size terciles are formed on market capitalisation on 30 April 2026 and are therefore subject to look-ahead bias

4.4 Structure of return co-movement

Figure 3 reports the Pearson correlations of daily log returns for the 20 largest cryptoassets; the full 45 × 45 matrix is given in Online Resource 3. Within the block of 20 assets, the median pairwise correlation is 0.56; the highest value is 0.74 (XLM and ADA), and the lowest is 0.38 (DOGE and XMR). Across all 990 pairs of the 45 assets, the median is 0.53; the interquartile range is 0.46 to 0.60; the maximum is 0.77 (LTC and EOS); and the minimum is 0.17 (HT and BTG). Every pair has at least 1,779 overlapping days. No pair is uncorrelated or negatively correlated, so no combination of these assets offers diversification within the asset class over the full sample.

The matrix is close to a single-factor structure. The first principal component of the correlation matrix accounts for 54.3% of the total variance, compared with 3.5% and 2.3% for the second and third, respectively. The average correlation of an asset with the other 44 ranges from 0.61 (ETH and EOS) and 0.60 (DOT, NEO and ADA) down to 0.45 (DCR), 0.43 (DOGE), 0.41 (FTT), 0.35 (HT) and 0.28 (BTG). The assets with low loadings are idiosyncratic for identifiable reasons, not by design: BTG is a largely inactive fork; HT and FTT are exchange tokens whose prices followed the fortunes of their venues; and DOGE was driven by retail episodes. The closest co-movers of BTC are EOS (0.70), LTC (0.69), ADA, DOT and NEO (0.66 each).

Two caveats apply. The correlations are unconditional and estimated over unequal windows spanning different regimes, and Pearson correlation on heavy-tailed returns is sensitive to a small number of crash

days, a point that becomes apparent in Sect. 4.5. A rank-based version of the matrix is the natural robustness check. In addition, because each correlation is estimated on its own overlapping window, the 45×45 matrix is not positive semidefinite (its smallest eigenvalue is -0.35) and cannot be used as it stands in portfolio optimisation or as the target of a dynamic correlation model. The matrix was therefore re-estimated jointly on the 1,779 days on which all 45 assets are observed (12 May 2021 to 25 March 2026; Online Resource 4). This matrix is positive definite, with a smallest eigenvalue of 0.08. All 990 correlations remain positive, between 0.11 and 0.84 with a median of 0.60, and the first principal component accounts for 59.8% of total variance (63.7% for rank correlations). The single-factor reading is thus stronger, not weaker, on the common sample. The static matrix is best regarded as the baseline against which a dynamic specification should be judged and not as a result in its own right.

4.5 Dynamics of co-movement

Figure 4 shows the mean of all pairwise correlations of daily log returns within a trailing 90-day window, sampled weekly. Two series are plotted: one for all cryptoassets with a complete window, whose number rises from 9 in 2016 to 45 in 2021, and one for the fixed cohort of the nine cryptoassets observed from 1 January 2016 (BTC, DASH, DGB, DOGE, ETH, LTC, XLM, XMR and XRP). Panel a uses Pearson and panel b Spearman correlations.

The series exhibits a shift in level and not a trend. For all assets, the median of the Pearson series is 0.20 up to the end of 2017 and 0.59 from mid-2018 onwards; its minimum since mid-2018 is 0.40, which is twice the earlier median, and 86% of the weekly observations since mid-2018 exceed 0.5. The shift is not an artefact of the panel's changing composition. In the fixed cohort, the median rises from 0.21 to 0.67, a larger increase than in the full panel, which shows that assets that entered later lowered the average rather than raising it. Annual medians reach their maxima in 2018 (0.67 for all assets and 0.73 for the fixed cohort) and in 2022 (0.67 for all assets and 0.75 for the fixed cohort), the year of the Terra and FTX failures. The most recent readings are lower for the full panel, with a median of 0.57 in 2025 and a final value of 0.55 on 29 April 2026, while the fixed cohort ends at 0.73.

Because overlapping windows are serially dependent by construction, the shift was tested on 41 non-overlapping 90-day blocks (Table 7). The block correlations up to the end of 2017 are lower than those from mid-2018 onwards under both estimators and in both panels (one-sided Mann-Whitney tests, $p < 0.001$ in all four cases). A test for a break in the mean at an unknown date (Andrews, 1993), with 15% trimming, places the break at the block that begins on 25 June 2017 in all four cases, with sup- F statistics between 85 and 136, against an asymptotic 5% critical value below 9. The block correlations remain autocorrelated, so the statistics are indicative, but the margin leaves little doubt about the existence or the timing of the shift. The break precedes the peak of the boom in initial coin offerings and the collapse of early 2018. The level and persistence of the series are consistent with studies on equicorrelation and dynamic conditional correlation in this market (Bouri et al., 2021; Demiralay & Golitsis, 2021; Ghorbel & Jeribi, 2021).

The comparison of the two panels of Fig. 4 shows how much the Pearson estimator owes to single days. In panel a, the series jumps from 0.54 to 0.78 when 12 March 2020 enters the window, then plateaus, peaking at 0.84 on 3 June 2020, and falls back when that day leaves the window about 90 days later. Similar plateaus follow 19 May 2021 and 8 November 2022. In panel b, the same episode raises the rank-based series from 0.56 to a local maximum of 0.74, with the overall peak of that series at 0.76 in July 2022. The Pearson peak of 0.84 is therefore a property of the estimator and should not be reported as a finding. An estimator with declining weights, or Engle's (2002) dynamic conditional correlation model, avoids the problem by design.

Co-movement is higher when the market is more volatile. Across the 30 blocks from mid-2018 onwards, the rank correlation between the mean pairwise Spearman correlation and the annualised volatility of the equally weighted market return is 0.48 ($p = 0.007$) for all assets and 0.43 ($p = 0.018$) for the fixed cohort. The relationship with the market return is negative but not significant ($-0.21, p = 0.26$; $-0.34, p = 0.07$). This is consistent with the evidence of contagion during the COVID-19 period (Corbet et al., 2020), with one reservation. Forbes and Rigobon (2002) show that correlation coefficients are upwardly biased when volatility rises, so an increase in measured correlation during turmoil does not, by itself, establish a change in the transmission mechanism. Rank correlations are less exposed to this bias than Pearson correlations, but are not immune to it. The statement supported by the evidence is the narrower one: that measured co-movement was highest during the periods when diversification would have been most valuable.

Table 7
Mean pairwise correlation in non-overlapping 90-day blocks before and after the transition

Panel	Estimator	Median to 2017	Median from mid-2018	Mann-Whitney p	Sup- F
All assets	Pearson	0.23	0.60	3.4×10^{-5}	90.8
All assets	Spearman	0.25	0.61	1.0×10^{-4}	135.9
Fixed cohort	Pearson	0.23	0.65	9.5×10^{-6}	84.9
Fixed cohort	Spearman	0.24	0.64	2.1×10^{-5}	126.3

There are 8 blocks up to the end of 2017 and 30 blocks from 1 July 2018. The Mann-Whitney test is one-sided. Sup- F is the largest F statistic for a single break in the mean over the central 70% of the 41 blocks (Andrews, 1993); in all four cases, it is attained at the block that begins on 25 June 2017

4.6 Peg stability of stablecoins

Table 6 and Fig. 5 report the absolute deviation of the 10 stablecoins from 1 US dollar. The mean absolute deviation differs by a factor of 48 across coins (17 if SAI is excluded), from 2.4 bps (USDC) and

2.7 bps (BUSD, measured until the end of support by Binance) to 29.2 bps (USDT), 40.2 bps (GUSD) and 114.4 bps (SAI). The ordering follows neither the age nor the size of the coins. The two tightest pegs belong to fully reserved tokens with direct redemption and deep arbitrage, consistent with evidence that arbitrage and demand-side forces, rather than issuance, keep reserve-backed stablecoins at par (Lyons & Viswanath-Natraj, 2023). The loosest peg is that of SAI, the single-collateral, crypto-backed predecessor of DAI. This is consistent with the finding that peg deviations of DAI reflect the risk of its collateral and fell sharply once conversion into fiat-backed collateral at par became possible (Kozhan & Viswanath-Natraj, 2021), and with the theoretical result that designs without safe collateral are the more fragile ones (d'Avernas et al., 2026). With 10 coins and the design type not coded as a variable, the association between mechanism and peg quality is descriptive rather than tested.

The mean and the tail of the distribution convey different information. USDC has the lowest mean deviation in the panel, yet its largest deviation is 293.5 bps, recorded on 11 March 2023 during the Silicon Valley Bank failure. GUSD is the mirror image, with a mean of 40.2 bps and a maximum of 8,223 bps, driven by illiquid quotes in the weeks following its September 2018 launch. A measure of peg quality based on the mean alone would place these two coins close together in risk, which they are not. The share of days spent away from par adds the dimension of exposure. GUSD is more than 50 bps from par on 21.9% of days and more than 1% from par on 9.9% of days, whereas USDC exceeds 50 bps on 0.2% of days. DAI and PAX exceed 50 bps on about 11% of days but 1% on only 6.3% and 2.2% of days, so their deviations are frequent but small. These results are consistent with evidence that stablecoins are not uniformly stable (Duan & Urquhart, 2023). The sign of the deviations differs by design. Among days with a deviation above 50 bps, DAI trades at a premium on 92% of them, whereas PAX and HUSD trade at a premium on only 14% and 6%, and USDT and USDC lie in between (48% and 40%). This matches the distinction between premiums arising from safe-haven demand and discounts arising from liquidity and collateral concerns (Lyons & Viswanath-Natraj, 2023). The daily log returns of the stablecoins are nearly orthogonal to those of BTC, with correlations of -0.14 (USDK) and 0.08 (USDT).

Breaches of the 5% threshold are historical. All 23 breach days of USDT fall in 2016 (3) and 2017 (20); the 16 breach days of SAI fall in 2017 (4) and 2018 (12); and the 11 breach days of GUSD fall in 2018 (5) and 2020 (6). DAI has three: one in March 2020 and two unverified observations in 2025 discussed in Sect. 4.1. USDC, TUSD, PAX, HUSD and USDK never breach 5%. If the two DAI observations prove to be a source artefact, no stablecoin in the panel breaches 5% after 2020. This statement is conditional on the universe: the panel does not include TerraUSD or any other failed algorithmic coin, so it cannot be extended to the stablecoin market as a whole.

One curation decision changes the ranking. In the raw series, BUSD is the worst performer in the panel, with 36 days exceeding 5% and a maximum deviation of 168%. Every one of those days falls after 15 December 2023, the end of Binance's support, and reflects a residual illiquid market. Truncated at that date, BUSD never deviates by more than 39 bps and ranks second. Because the daily master file contains no discontinuation flag, this truncation must be applied explicitly in every downstream analysis. Finally,

deviations are measured using daily closing prices at 00:00 UTC, which attenuates episodes resolved within the day.

Table 6
Peg stability of the 10 US dollar stablecoins, ordered by mean absolute deviation

Asset	Window	Days	Mean (bps)	95th pct. (bps)	Maximum (bps)	> 50 bps (%)	> 1% (%)	Days > 5%
USDC	2018-09-28 to 2026-04-30	2,772	2.4	7.5	293.5	0.2	0.0	0
BUSD ^a	2019-09-20 to 2023-12-15	1,548	2.7	8.5	38.6	0.0	0.0	0
HUSD	2019-07-18 to 2022-11-17	1,219	9.7	74.4	409.4	6.3	0.4	0
USDK	2020-06-13 to 2023-04-17	1,039	9.9	40.6	105.3	1.3	0.1	0
TUSD	2018-07-06 to 2026-04-30	2,856	14.7	44.9	402.6	3.9	2.1	0
PAX	2018-11-30 to 2026-04-30	2,709	19.0	72.8	292.4	11.3	2.2	0
DAI ^b	2019-11-20 to 2026-04-30	2,354	20.4	116.2	907.5	10.7	6.3	3
USDT	2016-01-01 to 2026-04-30	3,773	29.2	131.8	1,506.8	12.6	7.1	23
GUSD	2018-09-16 to 2026-04-30	2,784	40.2	156.4	8,222.9	21.9	9.9	11
SAI	2017-12-23 to 2019-11-30	708	114.4	301.5	1,061.6	54.0	32.3	16

Deviations are the absolute differences between the daily closing price and 1 USD, expressed in basis points. ^a Measured until 15 December 2023, the end of support by Binance. ^b Includes the unverified observations of 30 April and 1 May 2025, which also determine the maximum

4.7 Summary and limits of inference

Five statements are supported by the evidence. The panel is internally consistent, and its derived fields are exactly reproducible. Returns on cryptoassets are heavy-tailed throughout the cross-section, and drawdowns above 90% are the rule. Diversification within the asset class is unavailable over the full sample: all 990 pairwise correlations are positive and a single component accounts for more than half of the variance. Co-movement shifted to a higher level in mid-2017, in both the fixed cohort and the full panel. Peg quality differs across stablecoins by a factor of 48 in mean absolute deviation, and breaches of 5% are concentrated before 2021.

Four statements are not supported. The flat cross-sectional relation between volatility and growth is not evidence of a risk premium because it is unconditional, measured over unequal windows, and conditional on survival in the source. The association between size and growth is not evidence of a size factor, because size is measured after the fact. The peak of 0.84 in the Pearson co-movement series is not a finding because it is driven by a single day, and the higher correlation during volatile periods is not, by itself, evidence of contagion. No statement can be made about the DAI episode of 2025 until it has been verified.

Several elements were not assessed. They include the estimation of conditional variance and correlation models, the monthly panel, conditioning on macro-financial variables and a point-in-time reconstitution of the universe. Volume is reported by exchanges and is known to be inflated by wash trading, which has been estimated to account for more than 70% of reported volume on unregulated exchanges (Cong et al., 2023); therefore, it is not used as a measure of liquidity in this section.

4.8 Assessment of the hypotheses

Table 8 sets the five hypotheses of Sect. 2 against the evidence. All five are supported, with H2 showing moderate robustness in its second part. H1 is supported without exception: the Jarque-Bera statistic computed on the daily log returns of each series is at least 955 (AAVE), compared with a 1% critical value of 9.21, so normality is rejected for all 45 cryptoassets. H2 is supported in its first part, since the 95% confidence interval for the rank correlation between volatility and growth (-0.46 to 0.11) excludes any positive relationship with economic size. Its second part is supported with moderate robustness: volatility is related to deeper maximum drawdowns ($\rho = -0.37, p = 0.013$) and to more days with absolute returns above 50% ($\rho = 0.60, p < 0.001$), but the relation with drawdowns weakens to $\rho = -0.30$ ($p = 0.054$) when BTG and BNB are removed, and the relation with excess kurtosis is marginal ($\rho = 0.29, p = 0.057$).

H3 is supported. All 990 correlations are positive; the smallest (0.17) has a t-statistic of 7.4 at the shortest overlap of 1,779 days; the first principal component accounts for 54.3% of the total variance, and all 45 loadings on it are positive, ranging from 0.39 (BTG) to 0.85 (ETH). H4 is supported. The shift in level is significant in non-overlapping blocks under both estimators; it is larger in the fixed cohort than in the full panel; a break test places it in mid-2017; and co-movement from mid-2018 onwards rises with market volatility (Sect. 4.5, Table 7). The Pearson peak of 0.84 is not supported, which reflects a single day. H5 is supported. Mean absolute deviations differ by a factor of 48; the largest deviation of a coin is between 9 and 204 times its mean deviation (median 35); and the excess kurtosis of the daily returns of

the 10 stablecoins lies between 15 and 445, with normality rejected for all of them. The rank correlation between mean and maximum deviation is 0.77 ($p = 0.009$). Coins with a loose peg tend to have larger extremes on average, but the ordering is not the same, as the cases of USDC and GUSD show.

Table 8
Assessment of the hypotheses

H1	Test	Evidence	Assessment
	Jarque-Bera test on daily log returns, by asset, with Benjamini-Hochberg adjustment	Minimum 955 (1% critical value 9.21); rejected for 45 of 45; excess kurtosis above 3 for 45 of 45 (median 12.1)	Supported
H2	Spearman rank correlation of annualised volatility with CAGR, maximum drawdown, outlier days and excess kurtosis ($n = 45$)	CAGR – 0.19 ($p = 0.22$; 95% CI – 0.46 to 0.11); drawdown – 0.37 ($p = 0.013$); outlier days 0.60 ($p < 0.001$); kurtosis 0.29 ($p = 0.057$)	Supported; tail-loss part moderately robust
H3	Sign and significance of the 990 pairwise correlations; principal components of the correlation matrix	990 of 990 positive; minimum 0.17 ($t = 7.4$); first component 54.3% of variance; 45 of 45 loadings positive (0.39 to 0.85)	Supported
H4	Mean pairwise correlation in 41 non-overlapping 90-day blocks: Mann-Whitney test, break test at unknown date, rank correlation with market volatility	Median 0.23 to 0.25 up to 2017 and 0.60 to 0.65 from mid-2018 ($p < 0.001$, all assets and fixed cohort, Pearson and Spearman); break at June 2017; correlation with volatility 0.48 ($p = 0.007$)	Supported ^a
H5	Dispersion of mean absolute deviation; ratio of maximum to mean deviation; excess kurtosis of daily returns ($n = 10$)	Means from 2.4 to 114.4 bps (factor 48); ratio from 9 to 204 (median 35); excess kurtosis from 15 to 445; normality rejected for 10 of 10	Supported ^b

Confidence interval from the Fisher transformation. ^a Block correlations are serially dependent, so the break statistic is indicative. ^b The association between design and peg quality is descriptive and is not tested; the DAI ratio includes the two unverified observations of 2025

5 Discussion

5.1 Interpretation of the findings

The five results describe one market from different angles. Returns are heavy-tailed in every series (H1), extending the stylised facts reported for eight coins by Zhang et al. (2018) and for a broader universe by Ghosh et al. (2023) to a ten-year panel spanning three market cycles. The practical consequence is that second moments are incomplete measures of risk in this market. The same volatility can be associated with an excess kurtosis of 3 or 80, and an asset's risk is better described by its tails.

This helps to interpret H2. If volatility in this cross-section mainly reflects exposure to extreme days, as its correlation of 0.60 with the number of outlier days suggests, there is no reason for it to carry a premium, and none is found. The result does not contradict the factor literature. Liu et al. (2022) show that market, size and momentum capture expected returns, and total volatility is not among the priced characteristics. Two qualifications remain. The test has little power with 45 assets, and the sample is conditioned on coverage by the source in 2026, thereby removing failed assets that were likely both volatile and unprofitable. Survivorship therefore runs counter to the finding (Brown et al. 1992), and a universe that included defunct coins would be expected to show a more negative relationship.

H3 and H4 are two readings of the same structure. The first component accounts for 54.3% of the variance (59.8% on the common sample), with positive loadings for every asset, which means that a single market factor dominates, in line with the first factor of Liu et al. (2022) and with the spillovers from Bitcoin and Ether documented by Smales (2021). The rolling correlation indicates that the strength of this factor is not constant. The market moved from a segmented state, with a mean pairwise correlation near 0.2, to an integrated state near 0.6, and did not return to the former. The fixed cohort shows that the shift is not an artefact of the panel's composition, and the break test dates it to mid-2017, before the collapse in early 2018. The data cannot identify its cause. A candidate explanation, which is not tested here, is the growth from 2018 onwards of common funding in stablecoins, of derivatives markets and of leveraged participants active across many coins, which would create a common liquidity factor of the kind described by Brunnermeier and Pedersen (2009). The date of the break coincides with the 2017 boom, the period during which Griffin and Shams (2020) document the influence of Tether flows on the prices of Bitcoin and other coins. The assets with the lowest loadings (BTG, HT, FTT, DOGE) are those whose prices depended more on venue or retail episodes than on that common factor.

H5 shows that the term "stablecoin" covers instruments whose deviations from par differ by a factor of 48 in the mean, and whose extremes range from 9 to 204 times their mean. The economics of the peg explains this pattern. Arbitrage and demand-side forces keep reserve-backed coins close to par on ordinary days (Lyons & Viswanath-Natraj, 2023), so the mean is small. When redemption is in doubt, as with USDC in March 2023, or when markets are thin, as with GUSD after its launch, the deviation is limited only by arbitrageurs' willingness to commit capital, so the tail is long. Crypto-collateralised designs introduce the risk of their collateral (Kozhan & Viswanath-Natraj, 2021), which is a common factor in H3. This is the channel that links the two halves of the panel.

5.2 Implications

For investors, diversification across cryptoassets has offered little protection since 2018. A portfolio of many coins is, to a first approximation, a leveraged or deleveraged position in one factor, and the near-total drawdowns of Sect. 4.2 were experienced jointly. For risk management, Gaussian Value-at-Risk is unsuitable for every asset in the panel, and measures of peg quality based on mean deviation understate the risk that matters, which lies in the tail. For supervisors, the results support monitoring stablecoins using tail measures and their sensitivity to stress in the cryptoasset market, rather than average deviations.

5.3 A proposed model

The findings suggest a two-block model linked by the state of the cryptoasset market (Fig. 6). It is proposed as a specification to be estimated on the daily master file and has not been estimated here.

In the first block, the return of cryptoasset i on day t loads on a common factor,

$$r_{it} = \mu_i + \beta_i f_t + \varepsilon_{it}, \varepsilon_{it} = \sigma_{it} z_{it} \quad (9)$$

where f_t is the return of the value-weighted market portfolio or the first principal component, and it follows a skewed Student t distribution with v_i degrees of freedom. The conditional variance follows the asymmetric specification of Glosten et al. (1993) with the t innovations of Bollerslev (1987),

$$\sigma^2_{it} = \omega_i + (\alpha_i + \gamma_i I[\varepsilon_{it-1} < 0]) \varepsilon^2_{it-1} + \delta_i \sigma^2_{it-1} \quad (10)$$

Heavy tails (H1) enter through v_i and the single factor (H3) through β_i . The correlation of the standardised residuals is restricted to a single time-varying number, as in the dynamic equicorrelation model of Engle and Kelly (2012),

$$R_t = (1 - \rho_t) I + \rho_t J \quad (11)$$

where I is the identity matrix, and J is a matrix of ones. Three results argue for equicorrelation over an unrestricted dynamic conditional correlation model. The static matrix is close to a single-factor structure with a narrow interquartile range. A scalar correlation remains defined as assets enter the panel, whereas the pairwise matrix of Sect. 4.4 is not even positive semidefinite. Finally, equicorrelation has already been applied to this market (Bouri et al., 2021; Demiralay & Golitsis, 2021). To represent H4, the parameters that govern ρ_t depend on an unobserved state S_t that takes two values, segmented and integrated, and follows a first-order Markov chain (Hamilton, 1989). The filtered probability of the integrated state dates the transition without imposing a break date and can be compared with the mid-2017 break found in Sect. 4.5.

The second block models the absolute deviation from par of stablecoin j , $d_{jt} = |P_{jt} - 1|$, in two parts. The probability that the deviation exceeds a threshold u is,

$$\Pr(d_{jt} > u) = \Lambda(\theta_0 + \theta_1 \sigma_{f,t} + \theta_2 \rho_t + \theta_3 D_j + \theta_4 L_{jt}) \quad (12)$$

where Λ is the logistic function, $\sigma_{f,t}$ the conditional volatility of the common factor, ρ_t the equicorrelation of the first block, D_j indicators of design (fiat-reserved or crypto-collateralised) and L_{jt} indicators of the life cycle (launch window and wind-down). Conditional on exceedance, the size of the deviation follows a generalised Pareto distribution, in the spirit of the conditional extreme value approach of McNeil and Frey (2000),

$$d_{jt} - u \mid d_{jt} > u \sim \text{GPD}(\xi_j, \psi_{jt}). \quad (13)$$

The tail index ξ_j measures the heaviness of the tail of each coin (H5); θ_3 tests the association between design and peg quality that Sect. 4.6 could only describe, and θ_1 and θ_2 test whether stress in the cryptoasset market is transmitted to the pegs. The direction of this link follows evidence that spillovers run from Bitcoin and Ether to Tether, not the other way (Smales, 2021), and that cryptocurrency returns Granger-cause USDC depegging events, while depegs do not predict returns (Boido & Jones, 2026). Premiums and discounts should be modelled separately because premiums reflect safe-haven demand, whereas discounts reflect liquidity and collateral concerns (Lyons & Viswanath-Natraj, 2023).

H2 is tested outside the likelihood using cross-sectional regressions of returns on the estimated loadings, idiosyncratic volatility, and point-in-time size (Fama & MacBeth, 1973), thereby removing look-ahead bias in Table 5. The model can be estimated in steps: univariate variance equations, equicorrelation with regime switching, and then the stablecoin block, using the filtered quantities from the first block as regressors, with standard errors that account for the generated regressors. It should be judged out-of-sample through Value-at-Risk and expected shortfall backtests against a constant-correlation Gaussian benchmark and an unrestricted dynamic conditional correlation model (Engle, 2002).

5.4 Limitations and further research

The limitations of Sect. 4.7 carry over to the discussion. The most consequential ones are the survivorship of the universe, the absence of failed algorithmic stablecoins, source artefacts of the kind found for BTG and suspected for DAI, and the small number of stablecoins, which limits the power of any test of design effects. The model of Sect. 5.3 requires a positive semidefinite treatment of the unbalanced panel and the point-in-time market capitalisation series contained in the master file. Extending the stablecoin block to failed coins is the natural next step, because failures identify the tail.

6 Conclusions

This paper examined a panel of 45 cryptoassets and 10 US dollar stablecoins, observed daily from January 2016 to April 2026, a period that covers three market cycles, the 2022 failures, and the March 2023 banking stress. Every statistic was recomputed from a single audited master file, and five hypotheses derived from the literature were assessed with tests that make few distributional assumptions. All five are supported, and the findings can be stated in four points.

First, heavy tails are a property of the asset class. Normality is rejected for each of the 55 series, excess kurtosis exceeds 3 for every cryptoasset, and 42 of 45 assets lost more than 90% from a peak at some point. Volatility is not rewarded in the cross-section: the rank correlation between annualised volatility and compound growth is -0.19 , with a confidence interval that excludes any sizeable positive relation,

whereas volatility is associated with deeper drawdowns and with more extreme days. Twenty-five of the 45 assets lost value over their full histories, in a universe already biased towards survivors.

Second, the market has been dominated by one factor since 2017. All 990 pairwise correlations are positive, and the first principal component accounts for 54.3% of the variance, or 59.8% on the common sample. The mean pairwise correlation moved from about 0.2 to about 0.6, and a break test places the shift in mid-2017, during the boom, not in the collapse of early 2018. The shift is larger in a fixed cohort of nine assets than in the full panel, which rules out the changing composition of the market as its explanation. Since then, co-movement has risen with market volatility, so that diversification across cryptoassets has been least available when it was most needed.

Third, the term "stablecoin" encompasses instruments of very different quality. Mean deviations from par differ by a factor of 48 across coins, and the largest deviation of a coin is between 9 and 204 times its mean. A peg that is tight on average can therefore carry a long tail, as USDC showed in March 2023, and the sign of the deviations differs by design: on days with a deviation above 50 bps, DAI trades at a premium on 92% of them, whereas PAX and HUSD trade at a discount on 86% and 94%. Breaches of 5% are concentrated before 2021, a statement that holds only for a panel that contains no failed algorithmic coin.

Fourth, conclusions in this field are sensitive to data and estimator choices that are rarely reported. One scaling error in the source, lasting 9 days, increases an asset's excess kurtosis from 20 to 236. One crash day produces a Pearson correlation peak of 0.84 that lasts exactly as long as the estimation window and disappears under a rank-based estimator. A correlation matrix estimated pairwise on an unbalanced panel is not positive semidefinite and cannot be used in portfolio optimisation as it stands. An apparent confirmation of a suspect price by market capitalisation is mechanical, because capitalisation is price times supply. None of these problems is visible in summary tables, and each would have changed a headline result.

The relevance of the paper follows from these points. For research, it offers a documented and reproducible audit of a widely used open data source, extends the stylised facts of cryptoasset returns to a ten-year panel that includes stablecoins, dates the integration of the market with a test that does not impose the break, and shows by example which robustness checks are needed before results on co-movement and tail risk are reported. For investors and risk managers, it shows that a portfolio of many cryptoassets is close to a single-factor position, that Gaussian risk measures are unsuitable for every asset in the panel, and that realised volatility has indicated the risk of ruin more than the prospect of return. For supervisors, it supports the assessment of stablecoins using tail measures, the sign of their deviations, and their sensitivity to stress in the cryptoasset market, rather than average deviations. The two-block model of Sect. 5.3 turns these findings into a specification in which each hypothesis corresponds to a parameter and in which the transmission of stress from cryptoassets to pegs can be tested.

The conclusions are bounded by the design. The universe is defined by coverage in the source and excludes failed assets; the number of stablecoins is too small to test the effect of design; the evidence is descriptive rather than causal; and the proposed model has not been estimated. Further work should estimate the model on the master file, reconstruct the universe at each point in time, including defunct assets, extend the stablecoin block to include failed coins, and condition co-movement on interest rates, equity markets, and the US dollar.

Statements and Declarations

Funding: The authors did not receive support from any organisation for the submitted work.

Competing interests: The authors have no relevant financial or non-financial interests to disclose.

Author contributions: Conceptualisation: Nilton Gomes Furtado; Methodology: Nilton Gomes Furtado, Daniel Ferreira Polónia; Software: Nilton Gomes Furtado; Data curation: Nilton Gomes Furtado; Investigation: Nilton Gomes Furtado, Daniel Ferreira Polónia; Writing, original draft preparation: Nilton Gomes Furtado; Writing, review and editing: Nilton Gomes Furtado, Julia Vasconcelos Furtado; Supervision: Julia Vasconcelos Furtado. All authors read and approved the final manuscript.

Data availability: The raw data are publicly available from the Coin Metrics Community Data repository (<https://github.com/coinmetrics/data>). The derived tables and the code that reproduce the figures, tables and tests of this paper are provided as Online Resources. The daily master file is available from the corresponding author upon reasonable request.

Ethics approval: Not applicable. The study uses publicly available market data and involves no human participants or animals.

Consent to participate and consent to publish: Not applicable.

References

1. Alexander C, Dakos M (2020). A critical investigation of cryptocurrency data and analysis. *Quant Finance* 20(2):173–188. <https://doi.org/10.1080/14697688.2019.1641347>
2. Anadu K, Azar P, Cipriani M et al. (2023). Runs and flights to safety: are stablecoins the new money market funds? Federal Reserve Bank of New York Staff Reports 1073. <https://doi.org/10.59576/sr.1073>
3. Andrews DWK (1993). Tests for parameter instability and structural change with unknown change point. *Econometrica* 61(4):821–856. <https://doi.org/10.2307/2951764>
4. Ang A, Chen J (2002). Asymmetric correlations of equity portfolios. *J Financ Econ* 63(3):443–494. [https://doi.org/10.1016/S0304-405X\(02\)00068-5](https://doi.org/10.1016/S0304-405X(02)00068-5)

5. Ang A, Hodrick RJ, Xing Y, Zhang X (2006). The cross-section of volatility and expected returns. *J Finance* 61(1):259–299. <https://doi.org/10.1111/j.1540-6261.2006.00836.x>
6. Bali TG, Cakici N, Whitelaw RF (2011). Maxing out: stocks as lotteries and the cross-section of expected returns. *J Financ Econ* 99(2):427–446. <https://doi.org/10.1016/j.jfineco.2010.08.014>
7. Barberis N, Huang M (2008). Stocks as lotteries: the implications of probability weighting for security prices. *Am Econ Rev* 98(5):2066–2100. <https://doi.org/10.1257/aer.98.5.2066>
8. Baur DG, Hong K, Lee AD (2018). Bitcoin: medium of exchange or speculative assets? *J Int Financ Mark Inst Money* 54:177–189. <https://doi.org/10.1016/j.intfin.2017.12.004>
9. Benjamini Y, Hochberg Y (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Stat Soc Series B Stat Methodol* 57(1):289–300. <https://doi.org/10.1111/j.2517-6161.1995.tb02031.x>
10. Binance (2023) Notice regarding the removal of BUSD and conversion of BUSD to FDUSD, 29 November 2023. <https://www.binance.com/en/support/announcement/notice-regarding-the-removal-of-busd-and-conversion-of-busd-to-fdusd-1c98ce7bb464422dbbaeda7066ae445b>. Accessed 19 September 2026
11. Boido C, Jones L (2026). The relationship between stablecoin and cryptocurrency returns during periods of market stress. *FinTech* 5(3):68. <https://doi.org/10.3390/fintech5030068>
12. Bollerslev, T. (1987). A conditionally heteroskedastic time series model for speculative prices and rates of return. *Rev Econ Stat* 69(3):542–547. <https://doi.org/10.2307/1925546>
13. Borri N (2019). Conditional tail-risk in cryptocurrency markets. *J Empir Finance* 50:1–19. <https://doi.org/10.1016/j.jempfin.2018.11.002>
14. Bouri E, Gupta R, Roubaud D (2019). Herding behaviour in cryptocurrencies. *Finance Res Lett* 29:216–221. <https://doi.org/10.1016/j.frl.2018.07.008>
15. Bouri E, Vo XV, Saeed T (2021). Return equicorrelation in the cryptocurrency market: analysis and determinants. *Finance Res Lett* 38:101497. <https://doi.org/10.1016/j.frl.2020.101497>
16. Briola A, Vidal-Tomás D, Wang Y, Aste T (2023). Anatomy of a stablecoin's failure: the Terra-Luna case. *Finance Res Lett* 51:103358. <https://doi.org/10.1016/j.frl.2022.103358>
17. Brown SJ, Goetzmann W, Ibbotson RG, Ross SA (1992). Survivorship bias in performance studies. *Rev Financ Stud* 5(4):553–580. <https://doi.org/10.1093/rfs/5.4.553>
18. Brunnermeier MK, Pedersen LH (2009). Market liquidity and funding liquidity. *Rev Financ Stud* 22(6):2201–2238. <https://doi.org/10.1093/rfs/hhn098>
19. Cheah ET, Fry J (2015). Speculative bubbles in Bitcoin markets? An empirical investigation into the fundamental value of Bitcoin. *Econ Lett* 130:32–36. <https://doi.org/10.1016/j.econlet.2015.02.029>
20. Coin Metrics (2026). Coin Metrics community data. <https://github.com/coinmetrics/data>. Accessed 11 June 2026
21. Cong LW, Li Y, Wang N (2021). Tokenomics: dynamic adoption and valuation. *Rev Financ Stud* 34(3):1105–1155. <https://doi.org/10.1093/rfs/hhaa089>

22. Cong LW, Li X, Tang K, Yang Y (2023). Crypto wash trading. *Manag Sci* 69(11):6427–6454. <https://doi.org/10.1287/mnsc.2021.02709>
23. Cont R (2001). Empirical properties of asset returns: stylised facts and statistical issues. *Quant Finance* 1(2):223–236. <https://doi.org/10.1080/713665670>
24. Corbet S, Meegan A, Larkin C, Lucey B, Yarovaya L (2018). Exploring the dynamic relationships between cryptocurrencies and other financial assets. *Econ Lett* 165:28–34. <https://doi.org/10.1016/j.econlet.2018.01.004>
25. Corbet S, Larkin C, Lucey B (2020). The contagion effects of the COVID-19 pandemic: evidence from gold and cryptocurrencies. *Finance Res Lett* 35:101554. <https://doi.org/10.1016/j.frl.2020.101554>
26. d'Avernas A, Maurin V, Vandeweyer Q (2026) Can stablecoins be stable? *Manag Sci*. <https://doi.org/10.1287/mnsc.2024.06992>
27. Demiralay S, Golitsis P (2021). On the dynamic equicorrelations in the cryptocurrency market. *Q Rev Econ Finance* 80:524–533. <https://doi.org/10.1016/j.qref.2021.04.002>
28. Diamond DW, Dybvig PH (1983). Bank runs, deposit insurance, and liquidity. *J Polit Econ* 91(3):401–419. <https://doi.org/10.1086/261155>
29. Dickey DA, Fuller WA (1979). Distribution of the estimators for autoregressive time series with a unit root. *J Am Stat Assoc* 74(366):427–431. <https://doi.org/10.1080/01621459.1979.10482531>
30. Diebold FX, Yilmaz K (2012). Better to give than to receive: predictive directional measurement of volatility spillovers. *Int J Forecast* 28(1):57–66. <https://doi.org/10.1016/j.ijforecast.2011.02.006>
31. Duan K, Urquhart A (2023). The instability of stablecoins. *Finance Res Lett* 52:103573. <https://doi.org/10.1016/j.frl.2022.103573>
32. Engle RF (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* 50(4):987–1007. <https://doi.org/10.2307/1912773>
33. Engle RF (2002). Dynamic conditional correlation: a simple class of multivariate generalised autoregressive conditional heteroskedasticity models. *J Bus Econ Stat* 20(3):339–350. <https://doi.org/10.1198/073500102288618487>
34. Engle RF, Kelly B (2012). Dynamic equicorrelation. *J Bus Econ Stat* 30(2):212–228. <https://doi.org/10.1080/07350015.2011.652048>
35. European Union (2023). Regulation (EU) 2023/1114 of the European Parliament and of the Council of 31 May 2023 on markets in crypto-assets. *Off J Eur Union L* 150:40–205. <https://eur-lex.europa.eu/eli/reg/2023/1114/oj>. Accessed 19 September 2026
36. Fama EF (1965). The behaviour of stock-market prices. *J Bus* 38(1):34–105. <https://doi.org/10.1086/294743>
37. Fama EF, MacBeth JD (1973). Risk, return, and equilibrium: empirical tests. *J Polit Econ* 81(3):607–636. <https://doi.org/10.1086/260061>
38. Forbes KJ, Rigobon R (2002). No contagion, only interdependence: measuring stock market comovements. *J Finance* 57(5):2223–2261. <https://doi.org/10.1111/0022-1082.00494>

39. French KR, Schwert GW, Stambaugh RF (1987). Expected stock returns and volatility. *J Financ Econ* 19(1):3–29. [https://doi.org/10.1016/0304-405X\(87\)90026-2](https://doi.org/10.1016/0304-405X(87)90026-2)
40. Ghorbel A, Jeribi A (2021). Investigating the relationship between volatilities of cryptocurrencies and other financial assets. *Decis Econ Finance* 44(2):817–843. <https://doi.org/10.1007/s10203-020-00312-9>
41. Ghosh B, Bouri E, Wee JB, Zulfiqar N (2023) Return and volatility properties: stylised facts from the universe of cryptocurrencies and NFTs. *Res Int Bus Finance* 65:101945. <https://doi.org/10.1016/j.ribaf.2023.101945>
42. Gkillas K, Katsiampa P (2018). An application of extreme value theory to cryptocurrencies. *Econ Lett* 164:109–111. <https://doi.org/10.1016/j.econlet.2018.01.020>
43. Glosten LR, Jagannathan R, Runkle DE (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *J Finance* 48(5):1779–1801. <https://doi.org/10.1111/j.1540-6261.1993.tb05128.x>
44. Gorton GB, Zhang J (2021). Taming wildcat stablecoins. SSRN working paper. <https://doi.org/10.2139/ssrn.3888752>
45. Griffin JM, Shams A (2020). Is Bitcoin really untethered? *J Finance* 75(4):1913–1964. <https://doi.org/10.1111/jofi.12903>
46. Grobys K, Junttila J (2021). Speculation and lottery-like demand in cryptocurrency markets. *J Int Financ Mark Inst Money* 71:101289. <https://doi.org/10.1016/j.intfin.2021.101289>
47. Hamilton JD (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* 57(2):357–384. <https://doi.org/10.2307/1912559>
48. Jarque CM, Bera AK (1987). A test for normality of observations and regression residuals. *Int Stat Rev* 55(2):163–172. <https://doi.org/10.2307/1403192>
49. Ji Q, Bouri E, Lau CKM, Roubaud D (2019). Dynamic connectedness and integration in cryptocurrency markets. *Int Rev Financ Anal* 63:257–272. <https://doi.org/10.1016/j.irfa.2018.12.002>
50. Kallinterakis V, Wang Y (2019). Do investors herd in cryptocurrencies – and why? *Res Int Bus Finance* 50:240–245. <https://doi.org/10.1016/j.ribaf.2019.05.005>
51. Katsiampa P (2017). Volatility estimation for Bitcoin: a comparison of GARCH models. *Econ Lett* 158:3–6. <https://doi.org/10.1016/j.econlet.2017.06.023>
52. Kerr NL (1998). HARKing, hypothesising after the results are known. *Pers Soc Psychol Rev* 2(3):196–217. https://doi.org/10.1207/s15327957pspr0203_4
53. Koch S, Dimpfl T (2022). Attention and retail investor herding in cryptocurrency markets. SSRN working paper. <https://doi.org/10.2139/ssrn.4251242>
54. Koutmos D (2018). Return and volatility spillovers among cryptocurrencies. *Econ Lett* 173:122–127. <https://doi.org/10.1016/j.econlet.2018.10.004>
55. Kozhan R, Viswanath-Natraj G (2021). Decentralised stablecoins and collateral risk. SSRN working paper. <https://doi.org/10.2139/ssrn.3866975>

56. Krugman P (1979). A model of balance-of-payments crises. *J Money Credit Bank* 11(3):311–325.
<https://doi.org/10.2307/1991793>
57. Kwiatkowski D, Phillips PCB, Schmidt P, Shin Y (1992). Testing the null hypothesis of stationarity against the alternative of a unit root. *J Econom* 54(1–3):159–178. [https://doi.org/10.1016/0304-4076\(92\)90104-Y](https://doi.org/10.1016/0304-4076(92)90104-Y)
58. Kyle AS, Xiong W (2001). Contagion as a wealth effect. *J Finance* 56(4):1401–1440.
<https://doi.org/10.1111/0022-1082.00373>
59. Liu Y, Tsyvinski A (2021). Risks and returns of cryptocurrency. *Rev Financ Stud* 34(6):2689–2727.
<https://doi.org/10.1093/rfs/hhaa113>
60. Liu W, Liang X, Cui G (2020). Common risk factors in the returns on cryptocurrencies. *Econ Model* 86:299–305. <https://doi.org/10.1016/j.econmod.2019.09.035>
61. Liu Y, Tsyvinski A, Wu X (2022). Common risk factors in cryptocurrency. *J Finance* 77(2):1133–1177.
<https://doi.org/10.1111/jofi.13119>
62. Liu J, Makarov I, Schoar A (2023). Anatomy of a run: the Terra Luna crash. NBER working paper 31160. <https://doi.org/10.3386/w31160>
63. Ljung GM, Box GEP (1978). On a measure of lack of fit in time series models. *Biometrika* 65(2):297–303. <https://doi.org/10.1093/biomet/65.2.297>
64. Longin F, Solnik B (2001). Extreme correlation of international equity markets. *J Finance* 56(2):649–676. <https://doi.org/10.1111/0022-1082.00340>
65. Lyons RK, Viswanath-Natraj G (2023). What keeps stablecoins stable? *J Int Money Finance* 131:102777. <https://doi.org/10.1016/j.jimonfin.2022.102777>
66. Ma Y, Zeng Y, Zhang AL (2023). Stablecoin runs and the centralisation of arbitrage. SSRN working paper. <https://doi.org/10.2139/ssrn.4398546>
67. Maciel L (2021). Cryptocurrencies value-at-risk and expected shortfall: do regime-switching volatility models improve forecasting? *Int J Finance Econ* 26(3):4840–4855.
<https://doi.org/10.1002/ijfe.2043>
68. Makarov I, Schoar A (2020). Trading and arbitrage in cryptocurrency markets. *J Financ Econ* 135(2):293–319. <https://doi.org/10.1016/j.jfineco.2019.07.001>
69. Mandelbrot B (1963). The variation of certain speculative prices. *J Bus* 36(4):394–419.
<https://doi.org/10.1086/294632>
70. Mann HB, Whitney DR (1947). On a test of whether one of two random variables is stochastically larger than the other. *Ann Math Stat* 18(1):50–60. <https://doi.org/10.1214/aoms/1177730491>
71. McNeil AJ, Frey R (2000). Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *J Empir Finance* 7(3–4):271–300.
[https://doi.org/10.1016/S0927-5398\(00\)00012-8](https://doi.org/10.1016/S0927-5398(00)00012-8)
72. Merton RC (1973). An intertemporal capital asset pricing model. *Econometrica* 41(5):867–887.
<https://doi.org/10.2307/1913811>

73. Obstfeld M (1996). Models of currency crises with self-fulfilling features. *Eur Econ Rev* 40(3–5):1037–1047. [https://doi.org/10.1016/0014-2921\(95\)00111-5](https://doi.org/10.1016/0014-2921(95)00111-5)
74. Ozdamar M, Akdeniz L, Sensoy A (2021). Lottery-like preferences and the MAX effect in the cryptocurrency market. *Financ Innov* 7(1). <https://doi.org/10.1186/s40854-021-00291-9>
75. Ross SA (1976). The arbitrage theory of capital asset pricing. *J Econ Theory* 13(3):341–360. [https://doi.org/10.1016/0022-0531\(76\)90046-6](https://doi.org/10.1016/0022-0531(76)90046-6)
76. Schwenkler G, Shah A, Yang D (2026). Aggregate confusion in crypto market data. *Manag Sci*. <https://doi.org/10.1287/mnsc.2025.00611>
77. Smales LA (2021). Volatility spillovers among cryptocurrencies. *J Risk Financ Manag* 14(10):493. <https://doi.org/10.3390/jrfm14100493>
78. United States Congress (2025) Guiding and Establishing National Innovation for U.S. Stablecoins Act (GENIUS Act), Public Law 119-27. <https://www.congress.gov/119/plaws/publ27/PLAW-119publ27.pdf>. Accessed 19 September 2026
79. Yi S, Xu Z, Wang GJ (2018). Volatility connectedness in the cryptocurrency market: is Bitcoin a dominant cryptocurrency? *Int Rev Financ Anal* 60:98–114. <https://doi.org/10.1016/j.irfa.2018.08.012>
80. Zhang W, Li Y (2020). Is idiosyncratic volatility priced in cryptocurrency markets? *Res Int Bus Finance* 54:101252. <https://doi.org/10.1016/j.ribaf.2020.101252>
81. Zhang W, Wang P, Li X, Shen D (2018). Some stylised facts of the cryptocurrency market. *Appl Econ* 50(55):5950–5965. <https://doi.org/10.1080/00036846.2018.1488076>

Figures




Figure 1

Coverage of the 55 series between 1 January 2016 and 30 April 2026. Each bar runs from the first to the last priced day of an asset; assets are ordered by entry date. No series contains a missing calendar day inside its own window




Figure 2

Annualised volatility of daily log returns and compound annual growth rate for the 45 cryptoassets, each measured over the full history of the asset. The marker area is proportional to the logarithm of market capitalisation as of 30 April 2026. Spearman rank correlation: -0.19 ($p = 0.22$)




Figure 3

Pearson correlations of daily log returns for the 20 largest cryptoassets by market capitalisation on 30 April 2026. Correlations are computed pairwise over all overlapping days, with a minimum overlap of 1,779 days. Leading zeros are omitted inside the cells




Figure 4

Mean pairwise correlation of daily log returns of cryptoassets in a trailing 90-day window, sampled weekly. **a** Pearson correlation. **b** Spearman rank correlation. Solid line, all cryptoassets with a complete window; dashed line, fixed cohort of the nine cryptoassets observed from 1 January 2016. Vertical dotted lines mark 12 March 2020, 19 May 2021 and 8 November 2022




Figure 5

Mean absolute deviation (filled markers) and 95th percentile of the absolute deviation (open markers) from the 1 USD peg, in basis points. Annotations give the number of days on which the deviation exceeds 5%. The asterisk indicates that BUSD is measured until 15 December 2023, the end of support by Binance




Figure 6

Proposed two-block model. Block A describes cryptoasset returns using a common factor, asymmetric conditional variances with heavy-tailed innovations, and regime-dependent equicorrelation. Block B describes stablecoin peg deviations in terms of the probability and size of exceedances. Grey arrows show the transmission from the cryptoasset market to the pegs. Labels H1 to H5 indicate where each hypothesis is tested