Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

An Auditable Framework for Stablecoin Depeg Detection: Validation, Bias Quantification and Evidence from a Decade of Stablecoin Data

Nilton Gomes Furtado

nilton.furtado@ua.pt

University of Aveiro

Julia Vasconcelos Furtado

University of Aveiro

José Ricardo Filgueiras

Centro Universitário das Américas (FAM) - São Paulo

Research Article

Keywords: stablecoins, peg deviation, event detection, financial stability, reproducibility

Posted Date: July 21st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10083601/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

Stablecoins have become a central infrastructure of digital finance, yet the reliability of empirical conclusions about their stability depends on how peg deviations are measured. This article investigates the extent to which the frequency and severity of depegs depend on the detection rules adopted, proposing an auditable and reproducible framework for detection, validation and bias quantification. Drawing on a quantitative documentary approach with a longitudinal design, it analyses ten US dollar-pegged stablecoins with a daily price series in the Coin Metrics Community Data, between January 2016 and April 2026; the sample mixes active and discontinued coins and excludes UST and FDUSD for lack of price data. The event count is highly sensitive to the rules: it varies from 191 to 21 when the magnitude threshold rises from 1% to 5%, and from 191 to 29 when a minimum duration of seven days is required. A threshold-selection criterion that maximises Youden's J identifies a 2.5% daily-close threshold as optimal against documented episodes. In addition, 21.4% of the raw events are post-discontinuation illiquidity artefacts, and centralised fiat-collateralised stablecoins reveal greater peg fidelity than crypto-collateralised ones, while market size does not imply greater stability. A criterion-based validation reports a sensitivity of 1.00 at the any-event operating point and 0.375 at the severe operating point, with wide confidence intervals; a panel logit of depeg incidence is reported as an indicative, under-powered model. A second-source robustness test corroborates 13 of 21 severe episodes (61.9%) on an independent provider, leaving roughly two fifths uncorroborated. It is concluded that measured frequency and severity are not neutral properties of the market but functions of methodological and data-source choices.

JEL classification: G12; G23; G28; C55.

1. Introduction

Stablecoins, crypto-assets designed to maintain parity with a reference currency, typically the US dollar, have become a central infrastructure of contemporary digital finance, with a market capitalisation that peaked above USD 180 billion in 2022 and that remains in the region of USD 150 to 160 billion (Seegolam & Joseph, 2025). Their adoption energises payments, remittances and financial inclusion, especially in emerging economies, while at the same time raising concerns of monetary sovereignty, consumer protection and financial stability (Thomason & Umar, 2026). It is within this setting that the present study is located, its object being the detection and measurement of peg-deviation episodes, or depegs, and its scope falling on the ten US dollar-pegged stablecoins with a daily price series available in the Coin Metrics Community Data, over the period from January 2016 to April 2026. For operational purposes, a depeg is defined as a sustained departure of the market price from unit parity, namely an absolute deviation exceeding 1% for at least two consecutive days, or a single daily close exceeding 2%, with a severe episode defined as any daily deviation above 5%. The study does not address regulation in a normative sense, nor macro-financial transmission, concentrating on the empirical and methodological layer of stability measurement.

The literature has examined the stability of stablecoins from the standpoint of their economic design and of the trade-off between collateralisation and stability (De Blasis et al., 2023; Häfner et al., 2024), of fragility to runs and to flight-to-quality dynamics (Morgan, 2023; Oefele et al., 2024), and of the measurement of specific deviation episodes (Diop et al., 2024; Eichengreen & Viswanath-Natraj, 2022; Galati & Capalbo, 2023). The measurement procedure itself, however, remains under-scrutinised. How depegs are detected, with which magnitude thresholds, with which minimum duration and at which temporal resolution, is rarely made explicit, even though these choices condition the empirical conclusions about stability. From this gap arises the research problem that guides the work: to what extent do the measured frequency and severity of depegs depend on the detection rules adopted?

The relevance of the problem is at once social and academic. At the social and economic level, millions of users hold stablecoins as dollar substitutes, so that a peg deviation translates into a concrete financial loss, and regulators require reliable metrics to calibrate proportionate, risk-based regimes (Thomason & Umar, 2026). An inflated or understated measurement of the frequency and severity of deviations may induce disproportionate or complacent policy responses, all the more so given that these assets are ascribed risk analogous to that of unregulated money-market instruments (Seegolam & Joseph, 2025). At the academic level, the study shifts the focus from what is measured to how it is measured, and proposes a principled, performance-based criterion for choosing detection rules, a contribution transferable to other assets and markets where the reliability of measurement conditions the validity of inference (Eichengreen & Viswanath-Natraj, 2022).

Considering the stated problem, the general objective is to assess the extent to which the frequency and severity of depegs depend on the detection rules adopted, by proposing an auditable framework that validates the episodes and quantifies the biases of the magnitude threshold, minimum duration and temporal resolution. Five specific objectives operationalise it: to specify and document the detector within a reproducible pipeline; to quantify the sensitivity of frequency and severity to variation in the magnitude and minimum-duration thresholds; to measure the contamination of the event set by illiquidity artefacts and assess its effect on the peg-fidelity ranking; to examine the effect of temporal resolution on measured severity; and to validate the detector against documented episodes and model the determinants of depeg incidence. After this introduction, the article presents the theoretical framework, the methodology, the results and discussion, and the final considerations.

2. Theoretical Framework

2.1 Stablecoins as private money and the parity problem

The theoretical point of departure is the understanding of stablecoins as private liabilities that promise convertibility at par with a reference currency. On this reading, the maintenance of parity is not a given property but the endogenous outcome of a design of collateralisation, governance and redemption mechanisms (Häfner et al., 2024). The literature examining the monetary attribute of these instruments holds that their capacity to function as a medium of exchange and a store of value depends on the

credibility of the parity promise, which is subject to theoretical and empirical testing (Fan & Dai, 2023). This premise grounds the choice to treat the peg deviation, rather than the absolute price, as the theoretically relevant quantity.

2.2 The adopted model: collateral design and stability

The model chosen to guide interpretation links collateral design to peg stability. The degree of decentralisation and the nature of the collateral condition the stability that the issuance mechanism can sustain, such that crypto-collateralised arrangements tend to incur greater deviation risk than centralised fiat-collateralised ones, particularly under market turbulence (De Blasis et al., 2023). The economic design of stablecoins is, in this perspective, a problem of incentive engineering, in which the choice of collateral and of redemption rules determines the robustness of the peg in the face of shocks (Häfner et al., 2024). This model justifies the expectation, subsequently confirmed by the evidence, that centralised fiat-collateralised stablecoins exhibit greater peg fidelity than crypto-collateralised ones.

2.3 Currents in dialogue: fragility, runs and systemic risk

The adopted model dialogues with the literature on fragility and runs, which analyses stablecoins by analogy with money-market instruments and shadow banking and stresses that the parity promise is vulnerable to systemic dynamics (Morgan, 2023). During the March 2023 banking crisis, stablecoins exhibited flight-to-quality behaviour comparable to money-market mutual funds, which lends liquidity and reserve quality a decisive role in stability (Oefele et al., 2024). In parallel, the financial-stability and systemic-risk perspective frames the growth of these assets as a source of macroprudential vulnerabilities, advocating supervision and stress-testing (Seegolam & Joseph, 2025), within an international regulatory convergence around proportionate, risk-based approaches (Thomason & Umar, 2026).

2.4 Measurement as a theoretical layer

A distinctive premise of this study is that the measurement of the peg deviation is not a neutral act but a theory-laden operation. The empirical literature that measures devaluation risk and the impact of shocks on the peg demonstrates that the detection and quantification of episodes depend on choices of window, threshold and method (Eichengreen & Viswanath-Natraj, 2022; Fernández-Mejía, 2024). The use of local-projection and machine-learning approaches to study specific episodes evidences the sensitivity of results to analytical design (Diop et al., 2024), and the measurement of the effect of banking-crisis shocks on the peg illustrates how the operational definition of the event conditions its measure (Galati & Capalbo, 2023). From this body of theory derives the premise that structures the research problem: the fragility signal read in the data is inseparable from the rules with which it is detected.

3. Methodology

The research adopts a quantitative documentary approach with a longitudinal design, based on secondary data from a public source. In terms of depth it is predominantly descriptive and correlational,

with a methodological-evaluative core and an indicative inferential component; it is not causal. The choice of auditable secondary data and of a rule-based detector is consistent with the recent empirical tradition that studies peg-deviation episodes through local projections and machine learning (Diop et al., 2024) and that measures the impact of banking-crisis shocks with rigour (Galati & Capalbo, 2023).

The universe comprises the US dollar-pegged stablecoins with a daily price series available in the Coin Metrics Community Data (coinmetrics/data repository, master branch), from 1 January 2016 to 30 April 2026. The sample comprises the ten stablecoins with a usable price series: USDT, USDC, DAI, SAI, TUSD, PAX, GUSD, BUSD, HUSD and USDK. Sampling is non-probabilistic, by source-coverage criterion; UST and FDUSD were excluded for want of an accessible price series. It is acknowledged that the universe is not a point-in-time reconstitution of the Top 100, which introduces potential selection bias, partly mitigated by the retention of discontinued assets.

The central collection instrument is a reproducible pipeline that extracts daily closing prices at 00:00 UTC, building the price variable by hierarchical coalescence; calendar gaps are retained as missing values and flagged, without imputation. From the price, the daily peg deviation is computed as the difference from unity. A rule-based detector classifies an event as moderate when the absolute deviation exceeds 1% for two or more days, or when a single close exceeds 2%, and as severe when the deviation exceeds 5% on any day. Each event is flagged for occurrence in a residual post-discontinuation market. Because the study relies on a single data provider with reported volume of intermediate quality, a second-source robustness protocol is specified for the severe episodes: for each severe episode, the daily price series of an independent provider is retrieved and the day-by-day divergence between sources is computed, classifying as a market signal the days on which both sources record a deviation above the threshold and as a possible provider error the days on which only one does; this protocol was executed against an independent second source (CryptoCompare / CoinDesk Data API) and the results are reported in Section 4.7.

The analysis combines five procedures: descriptive statistics of events and monthly deviation metrics; sensitivity analysis recomputing frequency and severity under alternative magnitude and minimum-duration rules; a performance-based criterion that selects the magnitude threshold maximising Youden's J against documented episodes; criterion validation with classification metrics and Wilson confidence intervals; and an inferential panel logit of depeg incidence in the asset-month on log market capitalisation, log volume, a crypto-collateral dummy, a market-stress regime and asset age, with standard errors clustered by asset, complemented by an ordinary-least-squares robustness regression on the log mean absolute deviation. A battery of tests on daily logarithmic returns is also reported: the Augmented Dickey-Fuller test, whose formulation derives from Fuller (1976) and whose critical values are based on MacKinnon (1994, 2010), under constant (c) and constant-plus-trend (ct) specifications, the KPSS test under the same specifications, the Jarque-Bera test, the Ljung-Box test on returns and squared returns (at lags 10 and 20), and the ARCH-LM test (at lag 12), with control of the false discovery rate through the Benjamini-Hochberg procedure at a 5% level.

4. Results and Discussion

4.1 Contamination of the event set

The detector produces 243 raw events. Of these, 52 (21.4%) are artefacts of a residual illiquid market following discontinuation, entirely attributable to BUSD after its delisting in February 2023, with prices economically impossible for a stablecoin (Diop & Chevallier, 2026). This leaves 191 valid events (78.6%); Fig. 1 illustrates the composition. Any study using the raw output overstates the frequency of depegs by more than one fifth. Detections concentrate in GUSD (87 events) and SAI (20), low-liquidity stablecoins, and the largest-magnitude event (GUSD, an 82% deviation) coincides with the asset's first trading day, a launch artefact. Illiquidity therefore contaminates symmetrically, at launch and at discontinuation. All subsequent figures and tables exclude post-discontinuation and launch artefacts.

4.2 Chronology and distribution of valid events

Figure 2 arranges the valid events over time and Fig. 3 their annual frequency by severity. Events concentrate in 2016–2021 and become scarce thereafter for the main assets, partly reflecting the maturation of reserve-custody mechanisms and partly sample composition; severe events cluster in the early period.

4.3 Peg-fidelity hierarchy

Table 1

and Fig. 4 rank the stablecoins by mean absolute deviation, expressed in basis points (bps; one bp equals 0.01%), with bootstrap 95% confidence intervals. Two results stand out. First, centralised fiat-collateralised stablecoins (USDC, 2.4 bps) exhibit far greater fidelity than crypto-collateralised ones (DAI, 20.3 bps; SAI, 125.5 bps), corroborating the collateral-design view of stability (De Blasis et al., 2023; Häfner et al., 2024). Second, market size is not an indicator of fidelity: USDT, the largest asset, sits mid-table (29.2 bps), above smaller competitors, which dialogues with evidence on the asset-backing risk of the largest issuer (Jorcano Fernández et al., 2024) and on devaluation risk independent of scale (Eichengreen & Viswanath-Natraj, 2022). The confidence intervals show that the coarse ordering (best versus worst) is robust, whereas the middle of the ranking is not statistically distinguishable, so the hypothesis is supported only at the extremes.

Stablecoin	Mean abs. dev. (bps)	95% CI (bootstrap)	Days > 50 bps (%)	Months
USDC	2.4	1.9–3.0	0.18	92
BUSD	2.8	2.3–3.4	0.00	51
USDK	10.3	7.2–13.6	1.25	35
HUSD	10.9	4.9–18.5	7.50	41
TUSD	14.8	10.8–20.1	3.95	94
PAX	19.0	15.3–23.1	11.14	90
DAI	20.3	11.7–30.7	10.60	78
USDT	29.2	20.2–39.8	12.64	124
GUSD	46.2	28.2–75.1	22.16	92
SAI	125.5	83.8–175.7	54.30	24

Table 1. *Peg-fidelity ranking with bootstrap 95% confidence intervals (post-discontinuation artefacts excluded). Source: authors' elaboration.*

4.4 Duration, magnitude and the sensitivity of measures to rules

Most valid events are short (median duration of 2 days) and of low amplitude (median maximum deviation of 2.2%), with a right tail concentrating the severe events (Fig. 5), consistent with the heavy-tailed peg behaviour reported elsewhere (Fernández-Mejía, 2024). The central result, answering the research problem, is in Fig. 6: the event count varies from 191 events at a 1% magnitude threshold to 21 at 5% and to 8 at 8%, and from 191 at a one-day minimum duration to 82 at three days and 29 at seven days. The measured frequency and severity are therefore convex functions of the measurement rules, not neutral properties of the market.

Beyond demonstrating sensitivity, the framework offers a principled criterion for choosing the threshold: selecting the value that maximises Youden's J against the documented ground truth. Table 2 shows that

the optimum is a 2.5% daily-close threshold, with sensitivity of 1.00 and specificity of 0.93. This criterion is operationally actionable and supersedes the two prevailing conventions: the 1% threshold generates excessive low-materiality false positives, whereas the 5% (severe) threshold misses most documented depegs. The recommendation is therefore to calibrate the detector by classification performance against recognised episodes rather than by convention.

Table 2
Threshold selection by Youden's J against documented episodes.

Source: authors' elaboration.

Threshold	Sensitivity	Specificity	Youden's J
1.0%	1.000	0.870	0.870
2.0%	1.000	0.907	0.907
2.5% (optimal)	1.000	0.927	0.927
3.0%	0.875	0.941	0.816
5.0%	0.375	0.973	0.348

Figure 7 shows that the monthly mean absolute deviation peaks during the COVID-19 stress (March 2020, DAI and GUSD) and the Silicon Valley Bank collapse (March 2023, USDC and the immediate dragging of DAI), corroborating the collateral-contagion channel (Diop et al., 2024) and the propagation of banking-crisis shocks to the peg (Galati & Capalbo, 2023).

4.5 Validation

Validation uses an independent ground-truth set of documented depegs assessable in the sample (USDT 2017 and 2018; DAI 2020 and 2023; USDC 2023; TUSD 2024), at the asset-month resolution (8 positives; 713 negatives). Table 3 reports the confusion matrix and metrics with Wilson 95% confidence intervals, and Fig. 8 the ROC curve. At the any-event operating point sensitivity is 1.00 (CI 0.68–1.00), but the wide interval reflects the small positive set, so the validation is under-powered and the study states this explicitly. At the severe operating point sensitivity falls to 0.375: the daily-close 5% rule misses five of the eight documented episodes, including USDC and DAI in March 2023, precisely because their intraday troughs are not reflected in the 00:00 UTC close. This provides systematic, although still indirect, evidence for the temporal-resolution hypothesis, which is therefore treated as partially confirmed pending intraday data. Precision against documented episodes is low (0.08–0.14) and constitutes a lower bound, since most flagged events are genuine deviations of lesser notoriety.

Table 3
Confusion matrix and classification metrics by operating point (asset-month; n = 721; documented positives = 8).

Source: authors' elaboration.

Operating point	TP	FP	FN	TN	Sensitivity (95% CI)	Specificity (95% CI)
Any event (≥ 1%)	8	93	0	620	1.000 (0.676–1.000)	0.870 (0.843–0.892)
Severe only (≥ 5%)	3	19	5	694	0.375 (0.137–0.694)	0.973 (0.959–0.983)

4.6 Determinants of depeg incidence

As an inferential layer, a panel logit models the probability of a depeg event in the asset-month on log market capitalisation, log volume, a crypto-collateral dummy, a market-stress regime and asset age, with standard errors clustered by asset (Table 4). The signs are economically plausible: crypto-collateralisation roughly doubles the odds of a depeg, while greater volume and greater age reduce them. No coefficient is statistically significant at conventional levels, which is not concealed but reflects a structural limitation: the panel has only nine clusters, so cluster-robust inference is itself unreliable, and events are relatively rare. An ordinary-least-squares robustness regression on the log mean absolute deviation confirms the pattern of signs without significance. The inferential layer is therefore reported as indicative, and the reliable estimation of determinants is left as an agenda conditional on a larger sample.

Table 4
Panel logit of depeg incidence (asset-month; n = 697; 9 clusters; pseudo-R^2 = 0.159; SE clustered by asset).

Source: authors' elaboration. Finally, the battery of tests on daily logarithmic returns confirms the stylised facts of crypto-asset returns (Table 5): the ADF test rejects a unit root in all series while the KPSS test rejects stationarity in only three, jointly indicating stationary returns; normality is universally rejected, and conditional heteroscedasticity and volatility clustering are pervasive.

Variable	Odds ratio	95% CI	p
log market cap (z)	0.79	0.26–2.42	0.674
log volume (z)	0.41	0.10–1.75	0.227
crypto collateral	2.07	0.58–7.41	0.264
market-stress regime	1.20	0.65–2.22	0.567
asset age (z)	0.61	0.27–1.37	0.230

Table 5
Summary of the test battery (rejections after Benjamini-Hochberg correction, q = 0.05; ADF and KPSS under c and ct specifications; Ljung-Box at lags 10 and 20; ARCH-LM at lag 12).

Source: authors' elaboration.

Test	n	Rejections (FDR-BH)
ADF (stationarity)	110	110
KPSS (stationarity)	110	3
Jarque-Bera (normality)	55	55
Ljung-Box on r (autocorrelation)	110	81
Ljung-Box on r^2 (clustering)	110	110
ARCH-LM (heteroscedasticity)	55	54

4.7 Second-source robustness

Because the primary analysis relies on a single provider, the 21 severe episodes were re-examined against an independent second source over a ± 5-day window, classifying an episode as confirmed when the second source also records a deviation above 5% on an overlapping day. Table 6 and Fig. 9 report the outcome: 13 of 21 severe episodes (61.9%; Wilson 95% CI 40.9–79.2%) are corroborated, while 8 (38.1%) are not, and the maximum deviations across the two sources correlate at 0.58. The non-corroborated cases include the DAI episode of April 2025 (9.1% on the primary source versus 0.01% on the second) and two large USDT episodes of January 2016 and December 2017, which are therefore plausibly provider-specific and should be read with caution. Conversely, the GUSD launch anomaly and the two SAI episodes are extreme on both sources, indicating shared data quirks rather than provider-specific error. This result quantifies, rather than merely acknowledges, the single-source limitation: roughly two fifths of severe episodes are not independently confirmed, which reinforces the central caution that severe-event measurement is sensitive to the data source as well as to the detection rules.

Table 6
Second-source confirmation of severe episodes (n = 21; Wilson 95% CI for the confirmation rate: 40.9–79.2%).

Source: authors' elaboration.

Outcome	Episodes	Share	Notable cases
Confirmed by two sources	13	61.9%	USDT 2017 (banking stress); DAI 2020; SAI 2017–18
Not corroborated	8	38.1%	DAI 2025; USDT 2016-01 and 2017-12; four GUSD 2020

5. Final Considerations

The research answered the proposed problem and broadened it. The measured frequency and severity are convex functions of the detection rules, which confirms the central hypothesis, and the analysis revealed an additional problem rarely made explicit in the literature: illiquidity contamination, at discontinuation (21.4% of raw events) and at launch. The specific objectives were achieved: the detector was documented within an auditable pipeline; the sensitivity of the metrics to thresholds was quantified and a performance-based selection criterion proposed; the contamination was measured and shown not to alter the coarse ranking; the effect of temporal resolution was examined; and the detector was validated with classification metrics and complemented by an inferential model. The substantive hierarchy of peg fidelity is robust at the extremes, with centralised fiat-collateralised assets more faithful than crypto-collateralised ones, and market size unrelated to fidelity.

The findings carry practical implications for three audiences. For treasury and risk managers holding stablecoins as quasi-cash, the peg-fidelity hierarchy, robust at the extremes, offers a criterion for collateral selection and exposure limits: a preference for centralised fiat-collateralised stablecoins with a low historical deviation, and the avoidance of market size as a safety heuristic, since the largest asset is not the most faithful. For investors and platforms conducting issuer due diligence, the auditable framework and the threshold criterion provide a reproducible monitoring instrument that explicitly separates genuine deviation from illiquidity noise. For supervisors, the sensitivity of the metrics to detection rules implies that fixed, arbitrary regulatory thresholds may understate or exaggerate the incidence of depegs; thresholds should be calibrated by classification performance, and intraday data collected for supervisory purposes, lest fast episodes be underestimated.

The position defended is that the debate on stablecoin stability has privileged episodes and mechanisms at the expense of scrutinising measurement itself; such conclusions are reliable only when the measurement procedure is transparent and its errors are quantified. The method proved sufficient, with one circumscribed insufficiency: the full test of temporal resolution requires intraday data absent from the sample, so that effect was demonstrated systematically but indirectly. The reliance on a single provider was addressed by the second-source robustness test (Section 4.7), which corroborated 61.9% of severe episodes on an independent source and flagged the remaining 38.1% as plausibly provider-dependent, to be read with caution. Future work should report results under multiple rule configurations, treat illiquidity explicitly and symmetrically, incorporate intraday and multi-source data, broaden the sample to currently non-assessable stablecoins (UST, FDUSD), and integrate macro-financial factors to study transmission to the traditional financial system.

Declarations

Ethical Approval and Consent to Participate

Not applicable. This study uses publicly available secondary market data and does not involve human participants, their data, or animals.

Consent for Publication

Not applicable.

Funding

No funding was received for this research.

References

1. De Blasis, R., Galati, L., Webb, A., & Webb, R. I. (2023). Intelligent design: Stablecoins (in)stability and collateral during market turbulence. Financial Innovation, 9(1). https://doi.org/10.1186/s40854-023-00492-4
2. Diop, P. O., & Chevallier, J. (2026). Binance USD delisting and stablecoins repercussions: A local projections approach. Econometrics, 14(1), 6. https://doi.org/10.3390/econometrics14010006
3. Diop, P. O., Chevallier, J., & Sanhaji, B. (2024). Collapse of Silicon Valley Bank and USDC depegging: A machine learning experiment. FinTech, 3(4), 569–590. https://doi.org/10.3390/fintech3040030
4. Eichengreen, B., & Viswanath-Natraj, G. (2022). Stablecoins and central bank digital currencies: Policy and regulatory challenges. Asian Economic Papers, 21(1), 29–46. https://doi.org/10.1162/asep_a_00843
5. Fan, M., & Dai, J. (2023). Monetary attribute of stablecoins: A theoretical and empirical test. National Accounting Review, 5(3), 261–281. https://doi.org/10.3934/nar.2023016
6. Fernández-Mejía, J. (2024). Extremely stablecoins. Finance Research Letters, 62, 105268. https://doi.org/10.1016/j.frl.2024.105268
7. Galati, L., & Capalbo, F. (2023). Silicon Valley Bank bankruptcy and stablecoins stability. International Review of Financial Analysis, 91, 103001. https://doi.org/10.1016/j.irfa.2023.103001
8. Häfner, M., Henriques Pereira, M., Dietl, H., & Beccuti, J. (2024). The four types of stablecoins: A comparative analysis. Ledger, 9. https://doi.org/10.5195/ledger.2024.326
9. Jorcano Fernández, J., Echarte Fernández, M. Á., & Náñez Alonso, S. L. (2024). The asset-backing risk of stablecoin trading: The case of Tether. Economics and Business Review, 10(1). https://doi.org/10.18559/ebr.2024.1.1211
10. MacKinnon, J. G. (1994). Approximate asymptotic distribution functions for unit-root and cointegration tests. Journal of Business and Economic Statistics, 12, 167–176.
11. MacKinnon, J. G. (2010). Critical values for cointegration tests (Working Paper No. 1227). Queen's University, Department of Economics.
12. Morgan, J. (2023). Systemic stablecoin and the brave new world of digital money. Cambridge Journal of Economics, 47(1), 215–260. https://doi.org/10.1093/cje/beac060
13. Oefele, N., Baur, D. G., & Smales, L. A. (2024). Flight-to-quality—Money market mutual funds and stablecoins during the March 2023 banking crisis. Economics Letters, 234, 111464. https://doi.org/10.1016/j.econlet.2023.111464

14. Seegolam, H. K., & Joseph, L. (2025). Financial stability and systemic risks of stablecoins (pp. 141–188). https://doi.org/10.4018/979-8-3373-6655-5.ch006
15. Thomason, J., & Umar, M. (2026). Synthesizing global lessons for stablecoin regulation and financial stability (pp. 319–336). https://doi.org/10.4018/979-8-3373-9043-7.ch010

Figures

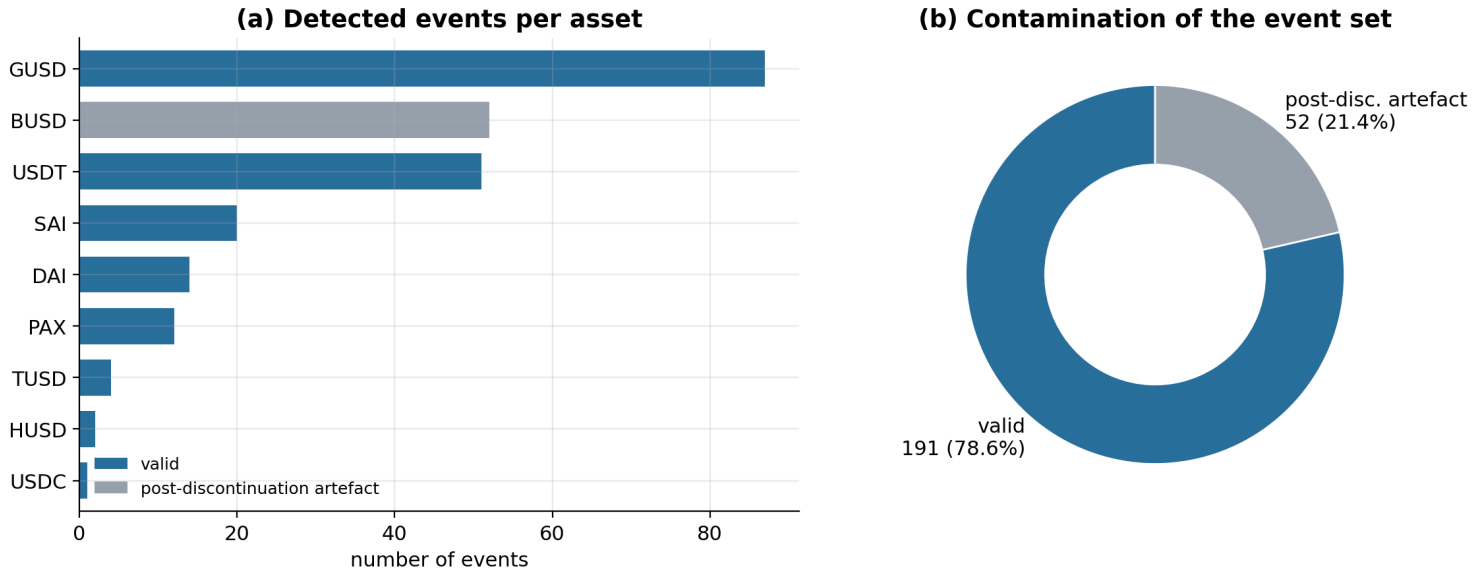


Figure 1

Detected events per asset (valid vs post-discontinuation artefacts) and overall contamination. Artefacts excluded from all subsequent analyses. Source: authors' elaboration.

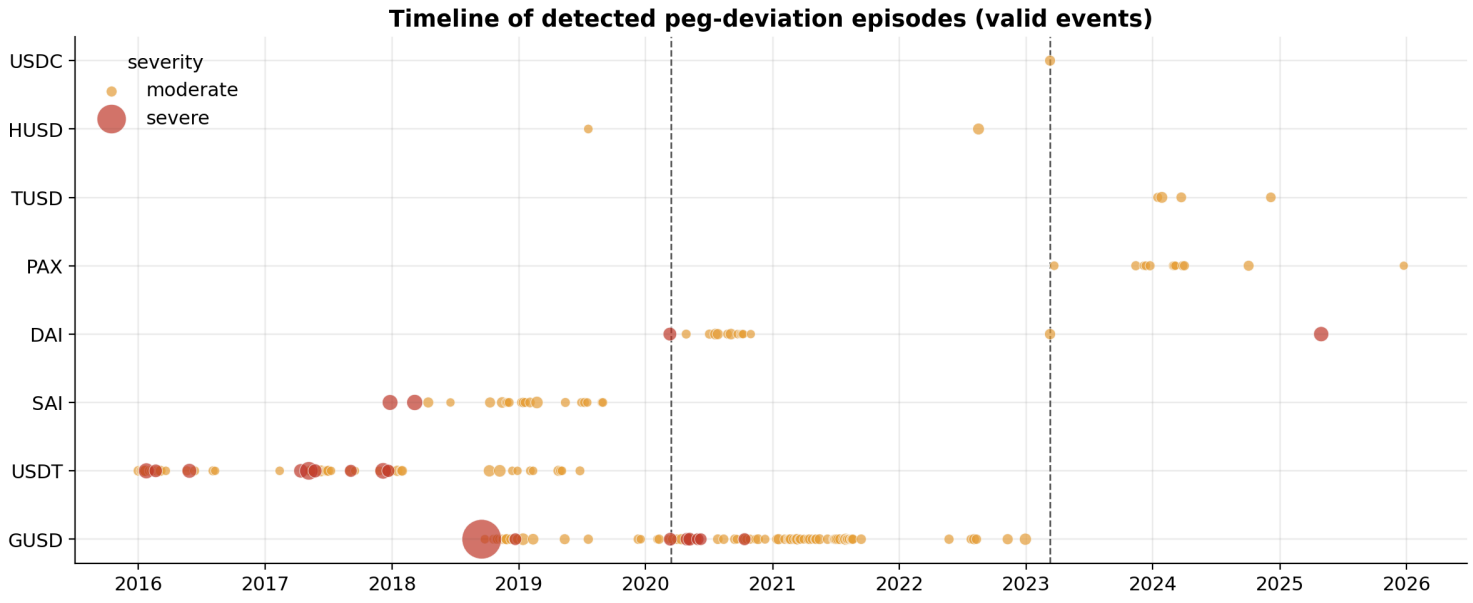


Figure 2

Timeline of detected peg-deviation episodes (valid events; marker area proportional to maximum deviation; artefacts excluded). Source: authors' elaboration.

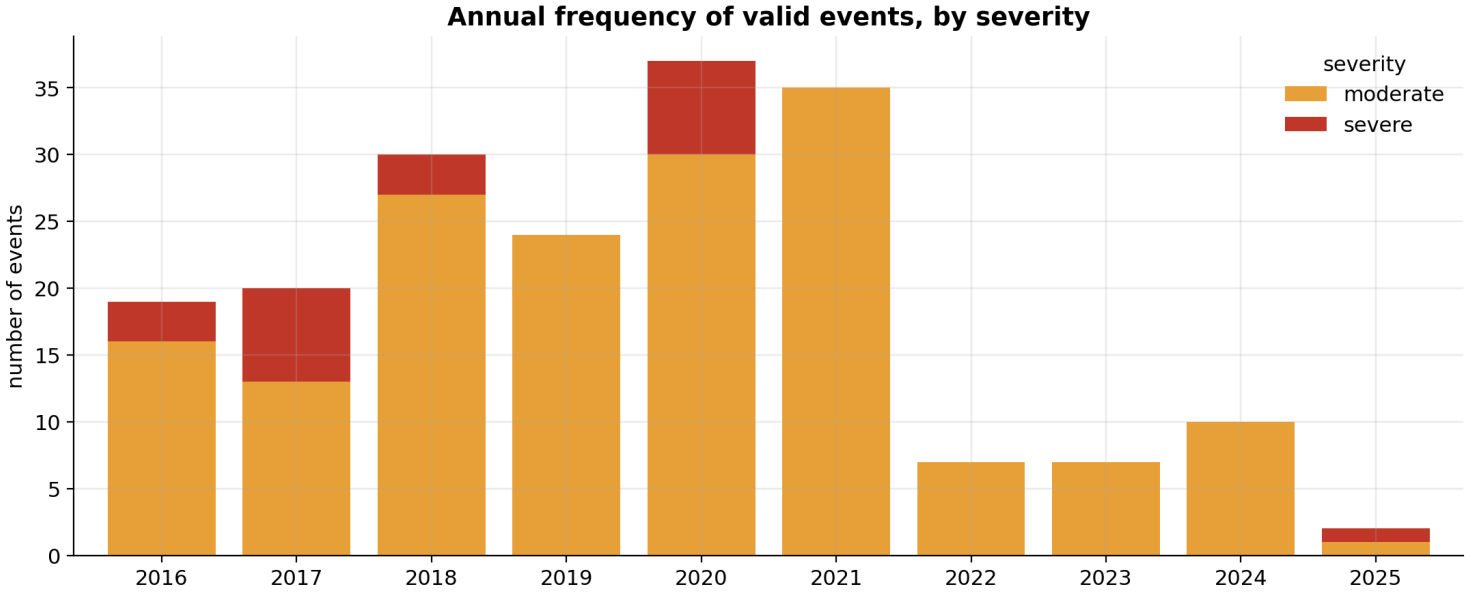


Figure 3

Annual frequency of valid events, by severity (artefacts excluded). Source: authors' elaboration.

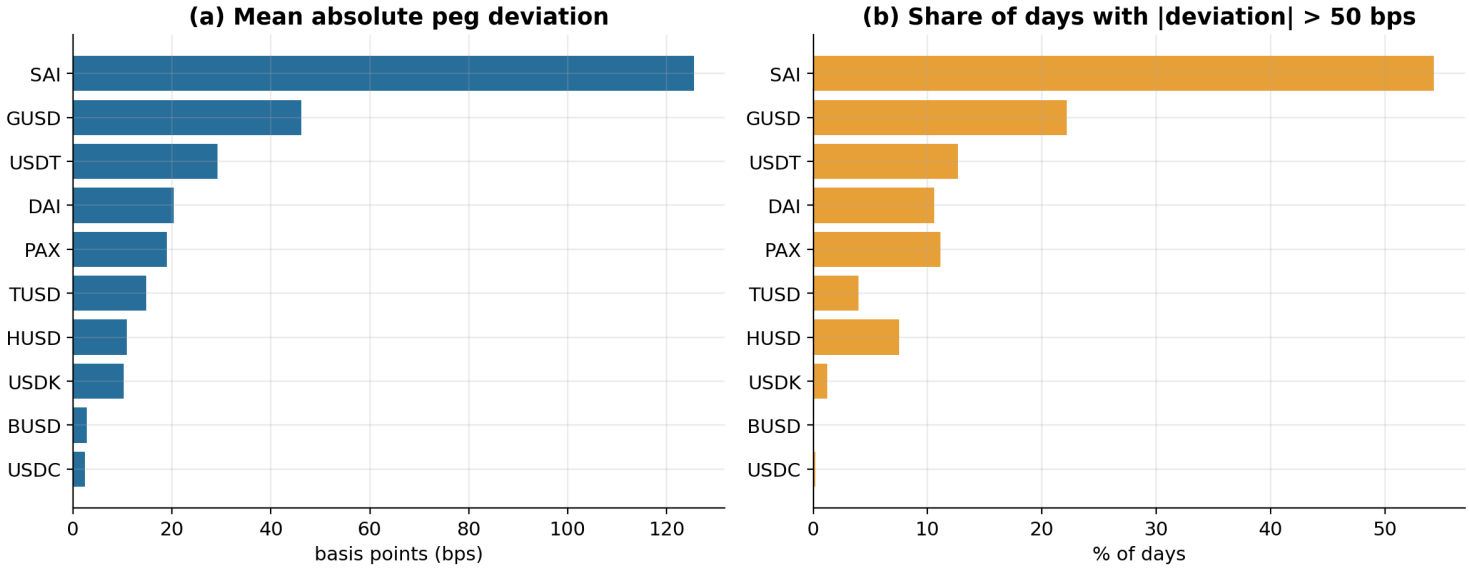


Figure 4

Peg-fidelity ranking: (a) mean absolute deviation (bps); (b) share of days with deviation above 50 bps. Source: authors' elaboration.

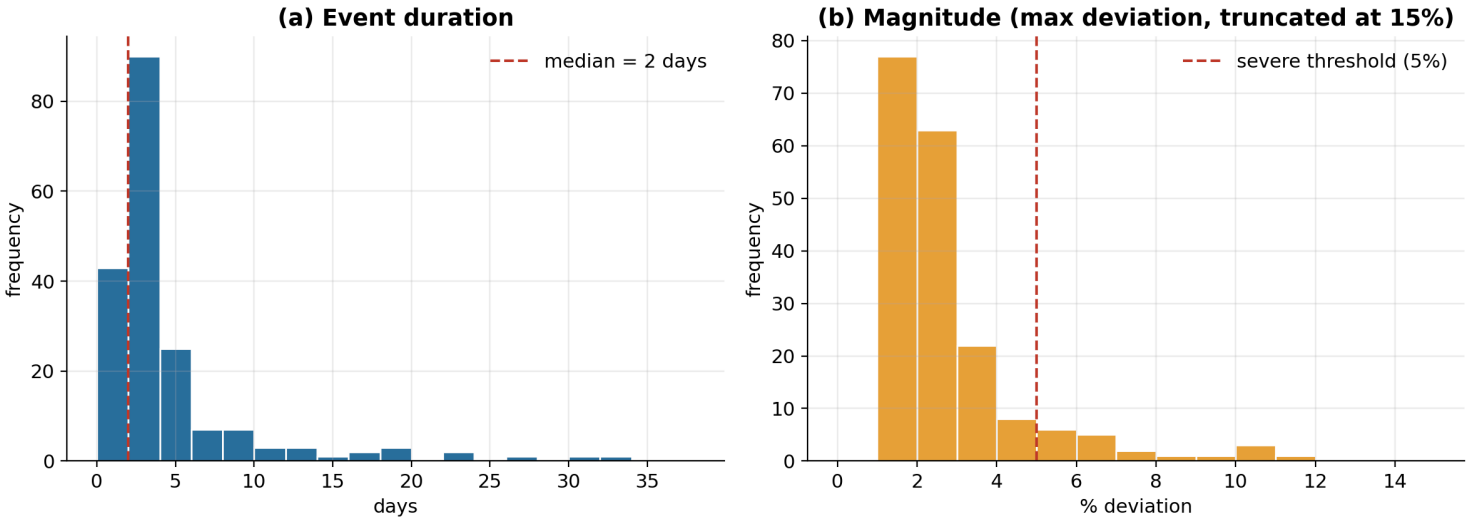


Figure 5

Distributions of (a) duration and (b) magnitude of valid events. Source: authors' elaboration.

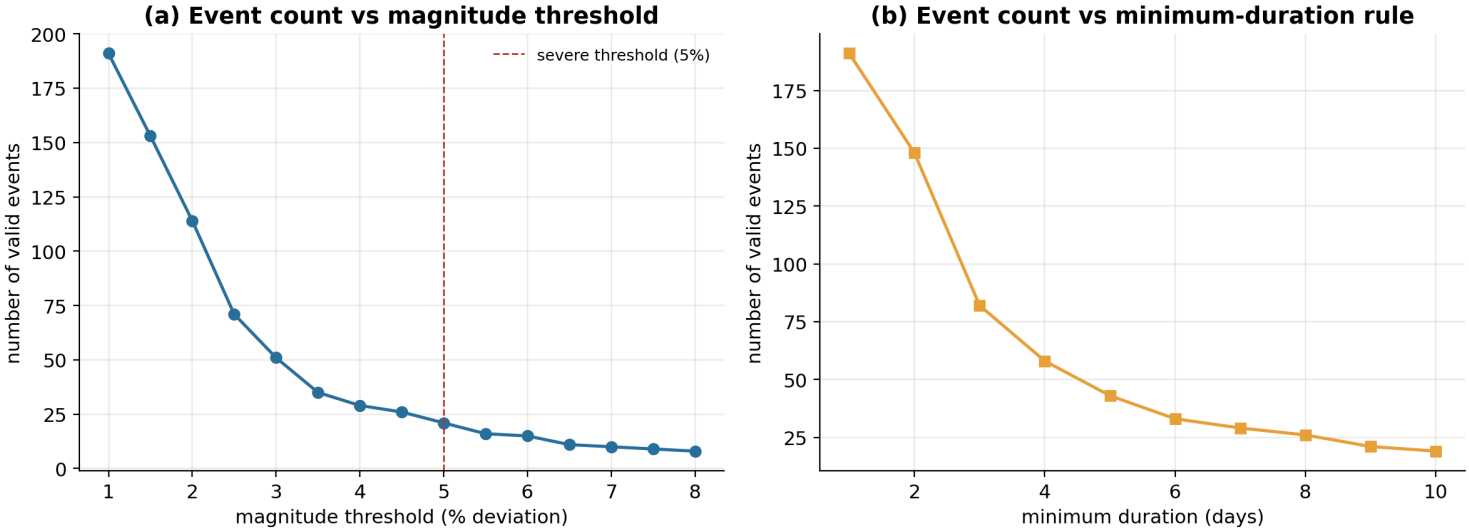


Figure 6

Sensitivity of the valid-event count to (a) the magnitude threshold and (b) the minimum-duration rule. Source: authors' elaboration

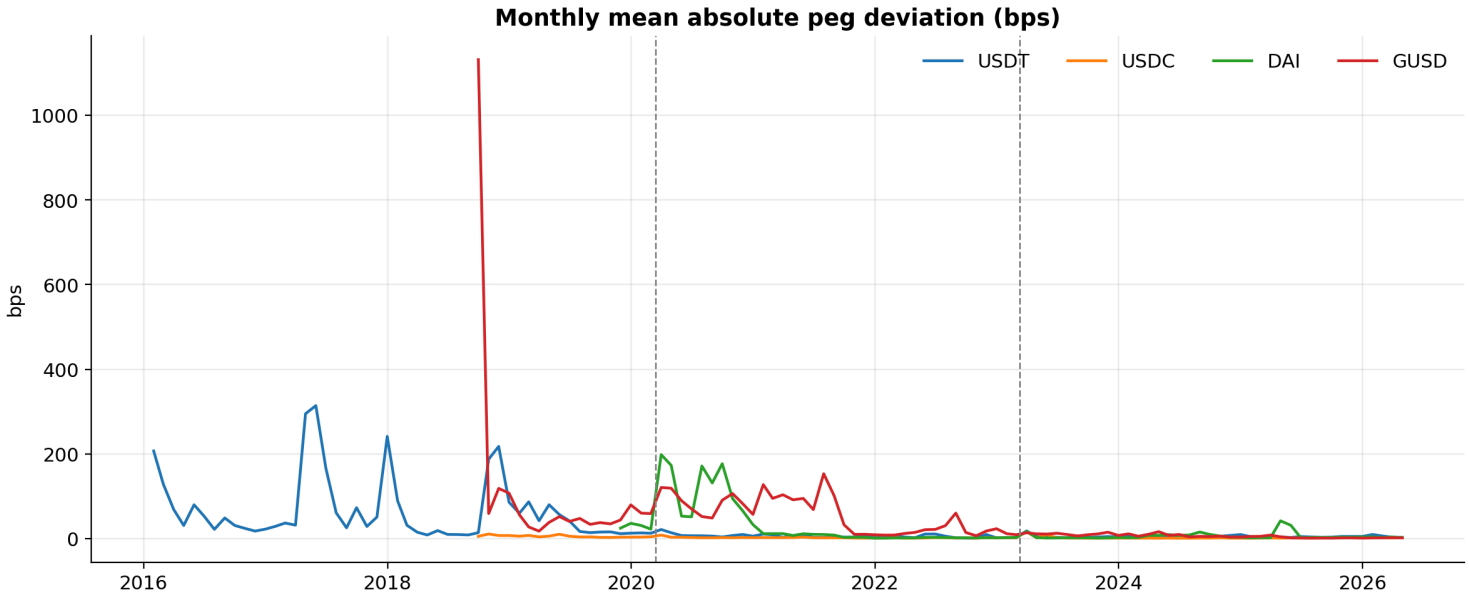


Figure 7

Monthly mean absolute peg deviation (bps), with markers of the stress episodes. Source: authors' elaboration.

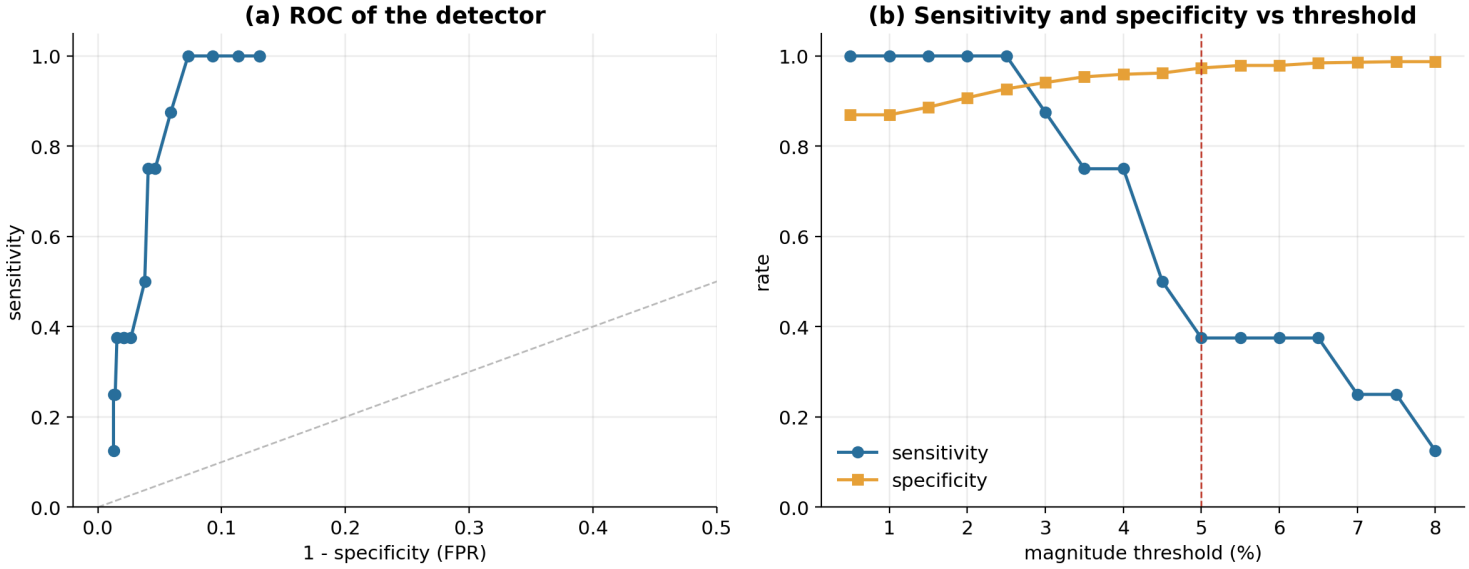


Figure 8

(a) ROC curve against the documented ground truth; (b) sensitivity and specificity by magnitude threshold. Source: authors' elaboration.

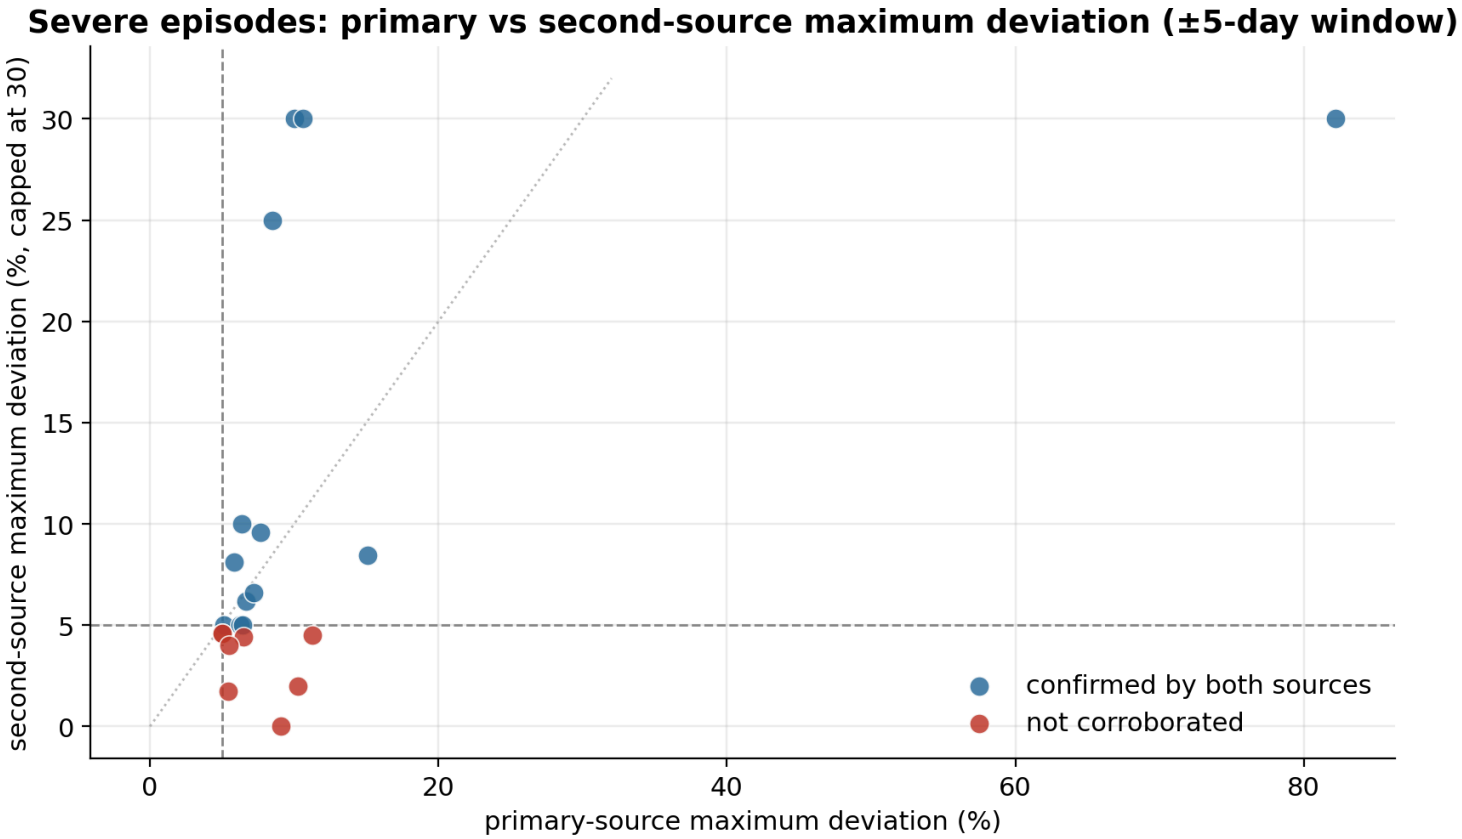


Figure 9

Severe episodes: primary versus second-source maximum deviation (±5-day window); points below the horizontal 5% line are not corroborated by the second source. Source: authors' elaboration.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- cryptomonthlypanel.xlsx