Finance Research Letters 113 (2026) 110877

Contents lists available at ScienceDirect

Finance Research Letters

journal homepage: www.elsevier.com/locate/frl

When code fails: Operational trust shocks and liquidity scarring in decentralized finance

Michael Zouari [a,*], Ilan Alon [a,b], Zeev Shtudiner [a]

[a] Department of Economics and Business Administration, Ariel University, Ariel, 40700, Israel
[b] Department of Strategic Management, Corvinus University of Budapest, Budapest, 1093, Hungary

ARTICLE INFO

JEL classification:
G01
G14
G23
G32

Keywords:
Decentralized finance
Smart-contract exploits
Operational risk
Liquidity scarring
Bank runs
Market discipline

ABSTRACT

Smart-contract exploits are a defining operational risk in decentralized finance (DeFi), yet their protocol-level liquidity effects remain poorly measured. We assemble 197 exploit and bridge-attack episodes across 51 chains (November 2020–May 2026); 133 incidents enter a matched stacked event study of price-neutral cumulative abnormal flows. Relative to matched controls, exploited protocols experience abnormal price-neutral liquidity losses of 27.6 percentage points of pre-event total value locked (TVL) by day 7 and 42.7 points by day 90. At day 180, 37.5% of coverage-qualified affected incidents are severely impaired on the price-neutral measure, with an adjusted matched risk difference of 28.1 percentage points ($p < 0.0001$). Short-run impairment rises sharply with relative exploit severity; in the two lowest-severity quartiles, median outflows exceed median reported theft, consistent with reassessment beyond mechanical loss. Matched-peer spillovers are economically small and lie within a ±5-percentage-point equivalence band. The evidence is consistent with operational trust functioning as franchise capital: once compromised, liquidity impairment is immediate, protocol-specific, and often persistent, and severity relative to pre-event TVL signals subsequent impairment that aggregate TVL conceals.

1. Introduction

How does liquidity behave when an uninsured financial intermediary's operating code fails in public? In banks, deposit insurance, supervision, and resolution regimes govern how withdrawal runs unfold. In decentralized finance (DeFi) the intermediary is a smart contract: user capital sits visibly on-chain, exit is immediate, and no standardized public backstop separates a trust shock from the withdrawal it provokes. Smart-contract exploits therefore offer a revealing setting for studying market discipline when operational trust fails.

Prior work measures DeFi risk at the protocol level (Bertomeu et al., 2024), links trust to DeFi valuation (Grande & Borondo, 2025), and shows that hacking incidents move Bitcoin prices and volatility, with nonlinear cyberattack effects across market regimes (Caporale et al., 2020; Chen et al., 2023). Recent studies add governance-token price discovery and losses of decentralized autonomous organizations after hacks, stablecoin safety premia, and aggregate redemptions (Azar et al., 2025; Campello et al., 2026; Kitzler et al., 2025; Lin, 2026), and valuation premia for protocol audits (Bhambhwani & Huang, 2024; Bourveau et al., 2024). To our knowledge, this is the first large-sample, multi-chain estimate of the exploited protocol's own price-neutral liquidity impairment against matched controls, with its persistence, severity gradient, and peer spillovers.

* Corresponding author.
E-mail addresses: michael.zouari@msmail.ariel.ac.il (M. Zouari), ilan.alon@uni-corvinus.hu (I. Alon), zeevs@ariel.ac.il (Z. Shtudiner).

https://doi.org/10.1016/j.frl.2026.110877
Received 14 July 2026; Received in revised form 29 September 2026; Accepted 2 October 2026
Available online 2 October 2026
1544-6123/© 2026 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

Aggregate total value locked (TVL) can signal confidence and market prices can register cyber incidents, but neither reveals whether liquidity providers abandon the exploited protocol itself. Raw TVL mixes deposits and withdrawals with token-price movements: a protocol can look impaired because its token price fell or because investors withdrew, and only the latter is quantity-based market discipline.

This paper studies smart-contract exploits as operational trust shocks to DeFi protocols. We build protocol-level, price-neutral cumulative abnormal flows (CAF) and compare exploited protocols with matched, non-exploited controls observed over the same calendar dates in a matched stacked event study. The estimand is deliberately narrow: the abnormal liquidity response once an exploit is confirmed.

Results are reported for two protocol universes: the expanded universe, the full protocol panel meeting the criteria of Section 3.1, and a core sample of established protocols fixed in advance, retained for robustness. In the expanded universe, exploited protocols lose about 28% of pre-event TVL within a week, and the effect persists in the core sample and under TVL floors up to $100 million.

The price-neutral measure links DeFi security risk to the economics of runs (Diamond & Dybvig, 1983; Goldstein & Pauzner, 2005) and to franchise value (Kamiya et al., 2021).

2. Literature and hypotheses

DeFi protocols combine financial intermediation with code-based execution. Smart contracts widen the contracting space (Cong & He, 2019), but they also make code failure a core source of institutional fragility.

We treat an exploit as an information event about a protocol's operational integrity: its code and security controls. Liquid claims turn fragile when confidence deteriorates (Diamond & Dybvig, 1983), and collateralized non-bank funding can face rapid withdrawal when its perceived safety is impaired (Gorton & Metrick, 2012). DeFi sharpens this fragility, because exit is immediate and no standardized resolution mechanism exists to rebuild trust after failure. An exploit is thus a durable, protocol-specific signal about implementation quality and serves as the treatment below.

With rapid exit and no backstop, the liquidity response should be immediate.

H1. Following an exploit, affected protocols experience more negative price-neutral liquidity changes than matched non-exploited protocols.

Because cyber losses reach franchise value beyond the direct theft (Kamiya et al., 2021), the response should persist.

H2. At 180 days, exploited protocols are less likely to recover and more likely to remain severely impaired than matched controls.

Belief updating predicts that the response scales with the signal (Goldstein & Pauzner, 2005): larger relative losses reveal deeper failure. Because mechanical theft also scales with relative loss, the behavioral reading holds only at low severity (Section 4.4).

H3. Short-run price-neutral liquidity impairment increases with exploit severity relative to pre-event TVL.

Peers may share code, chains, oracles, or liquidity, so contagion is possible (Gorton & Metrick, 2012); but the direct asset loss and the sharpest information update fall on the exploited protocol, so peer spillovers should be small.

H4. Average short-run price-neutral liquidity spillovers to matched peers are economically small.

3. Data and methodology

3.1. Data and event sample

The empirical setting is a protocol-day panel of DeFi intermediaries matched to a curated file of smart-contract exploits and bridge attacks. Protocol-level TVL and per-token balance histories, which underlie the price-neutral flows of Section 3.2, are obtained from DeFiLlama Pro, while dated exploit and bridge-attack records (event date, affected protocol, chain, reported dollar loss) are obtained from DeFiLlama Hacks. The sample spans November 2020 to May 2026 across 51 chains.

Incidents are classified from their attack descriptions into three types: code-vulnerability incidents, defects in smart-contract or protocol code (144 of 197 usable incidents); operational-security incidents, compromised private or administrator keys, phishing, DNS or front-end compromise, or access-control and governance failures (36); and bridge incidents, whose compromised infrastructure spans chains (17, including four key compromises). Section 4.2 uses this taxonomy; the attack-vector labels ship with the replication package.

An incident enters the sample when the protocol maps to the panel, the event date and dollar loss are observed, pre-event TVL is at least $1 million, and panel coverage is sufficient, yielding 197 usable incidents (133 matched) in the expanded universe and 111 (78 matched) in the core sample.

3.2. Price-neutral liquidity outcomes

The dependent variable is cumulative abnormal price-neutral flow scaled by pre-event TVL, following the logic of DeFi risk measurement (Bertomeu et al., 2024). Daily flow aggregates quantity changes across a protocol's TVL tokens:

$$Flow_{i,t} = \sum_{k \in K_{i,t}} \left(B_{i,k,t} - B_{i,k,t-1}\right) P_{k,t} \tag{1}$$

where $K_{i,t}$ is the set of protocol i's TVL tokens, $B_{i,k,t}$ the protocol-level balance of token k, and $P_{k,t}$ its dollar price on day t. The balances

are protocol-level holdings of all constituent tokens, not the governance token alone and not liquidity providers' wallet balances. Pure price revaluation generates no measured flow when token quantities are unchanged; quantity changes are valued at contemporaneous token prices. For incident e and event window [a,b],

$$\mathrm{CAF}_{i,e,[a,b]} = \frac{\sum_{\tau=a}^{b}\mathrm{Flow}_{i,e,\tau}}{\mathrm{TVL}_{i,e,-1}} \quad (2)$$

where t = 0 is the exploit date (UTC) and $\mathrm{TVL}_{i,e,-1}$ is pre-event TVL. The unit-level CAF is a cumulative flow ratio; its abnormal component is identified by the affected-versus-control contrast in Eq. (3).

The main windows are [0,7], [0,30], and [0,90]; the seven-day window captures the immediate response, and longer windows are read cautiously because DeFi accounting can generate non-economic jumps. Event windows [0,h] include both the incident day and day h. CAF is winsorized at the 1st and 99th percentiles. Each horizon uses incidents whose [0,h] window is complete at the June 28, 2026 snapshot. All 133 matched incidents have complete calendar follow-up through day 30; eight lack complete calendar follow-up through day 90, leaving 125. The corresponding core-sample counts are 78 and 71. Long-run impairment is classified through 180-day states on a price-neutral ratio, PNR_{180}, equal to one plus cumulative net flow over [0,180] scaled by pre-event TVL. A protocol is recovered at $\mathrm{PNR}_{180} \geq 0.80$, severely impaired at $\mathrm{PNR}_{180} \leq 0.20$, and scarred otherwise. The ratio is computed for incidents whose window has observed daily flows on at least half of its days (Online Appendix Table A.17); raw-TVL states serve as comparison.

3.3. Matched stacked event design

The empirical strategy combines event-study measurement (MacKinlay, 1997) with matching to improve treated-control comparability (Abadie & Imbens, 2006). Controls come from the affected protocol's category (26 groups; Online Appendix Table A.2), hold pre-event TVL of at least $1 million within a 0.5–2× caliper, have no recorded own exploit with a positive reported dollar loss within ±180 days, and show at least 30 post-event daily observations. Candidates are ranked by absolute log-TVL distance, with a penalty added when the primary chain differs, and the nearest 20 are retained. Incidents with fewer than five controls are dropped; the median incident retains 20 controls. Matching is with replacement, and all estimates adjust for log pre-event TVL.

Chain is a ranking preference rather than a restriction because joint category, size, and chain requirements leave little support: 55 of 133 matched incidents have no qualifying same-chain control, and requiring five same-chain controls would leave 18 incidents, all drawn from a single chain. Incident fixed effects absorb market-wide conditions but not shocks confined to the affected chain; Section 4.2 and Online Appendix Table A.16 test that margin directly.

The cross-sectional event-window specification is

$$Y_{i,e} = \beta\mathrm{Affected}_{i,e} + \alpha_e + \Gamma X_{i,e,-1} + \varepsilon_{i,e} \quad (3)$$

Here Y is CAF or a long-run state, affected equals one for the exploited protocol, α_e is an incident fixed effect, and X holds pre-event covariates where used; β estimates the abnormal response of the affected protocol relative to its matched controls. Because every unit in an incident stack shares the same calendar window, α_e absorbs market-wide crypto conditions over those dates.

The stacked design follows modern staggered-event estimators (Callaway & Sant'Anna, 2021; Sun & Abraham, 2021). Dynamics are traced by the nested windows [0,7], [0,30], and [0,90] (Table 1, Fig. 1) and by non-overlapping pre-event windows (Online Appendix Table A.7).

3.4. Inference and validity checks

Inference uses two-way clustering by incident and protocol (Cameron et al., 2011), with matched-set permutation p-values that reassign the affected label within incident sets as a finite-sample complement to clustered inference (Young, 2019). Validity checks cover pre-event leads (Table 1, notes), duplicate episodes, covariate adjustment, core-versus-expanded samples, TVL floors up to $100

Table 1
Matched estimates of abnormal price-neutral liquidity loss after exploits.

Sample	Matched incidents, complete horizon (7/30/90d)	CAF[0,7] (pp)	CAF[0,30] (pp)	CAF[0,90] (pp)
Expanded full universe	133/133/125	−27.6 (2.8)	−33.6 (3.3)	−42.7 (4.9)
Core established protocols	78/78/71	−19.6 (2.9)	−22.2 (3.7)	−29.2 (11.2)

Notes: TVL denotes total value locked; CAF denotes cumulative abnormal price-neutral flow; pp denotes percentage points. CAF[0,h] cumulates daily price-neutral flows from the exploit day (day 0) through day h and divides by pre-event TVL, the last TVL observed in the seven days before the exploit. Entries are affected-minus-matched-control coefficients in percentage points of pre-event TVL, not mean losses among affected protocols. Estimates include incident fixed effects and log pre-event TVL adjustment. Controls are matched on protocol category and pre-event TVL, with a same-chain ranking preference. The expanded sample includes all eligible protocols; the core sample restricts the design to a predefined established-protocol universe. CAF outcomes are winsorized at the 1st and 99th percentiles. Standard errors, clustered by incident and protocol, appear in parentheses. $p < 0.001$ for all coefficients except the core day-90 estimate ($p \approx 0.01$). Matched-set permutation (add-one) $p = 0.00005$ for all three expanded-sample windows (Online Appendix Table A.3). Each horizon uses incidents with complete follow-up at the June 28, 2026 snapshot; eight expanded-sample (seven core) incidents are right-censored before day 90. The joint null that all three pre-event leads equal zero is not rejected ($p = 0.24$), although the nearest lead is positive and individually significant (7.2 pp, $p = 0.039$); see Online Appendix Table A.7.

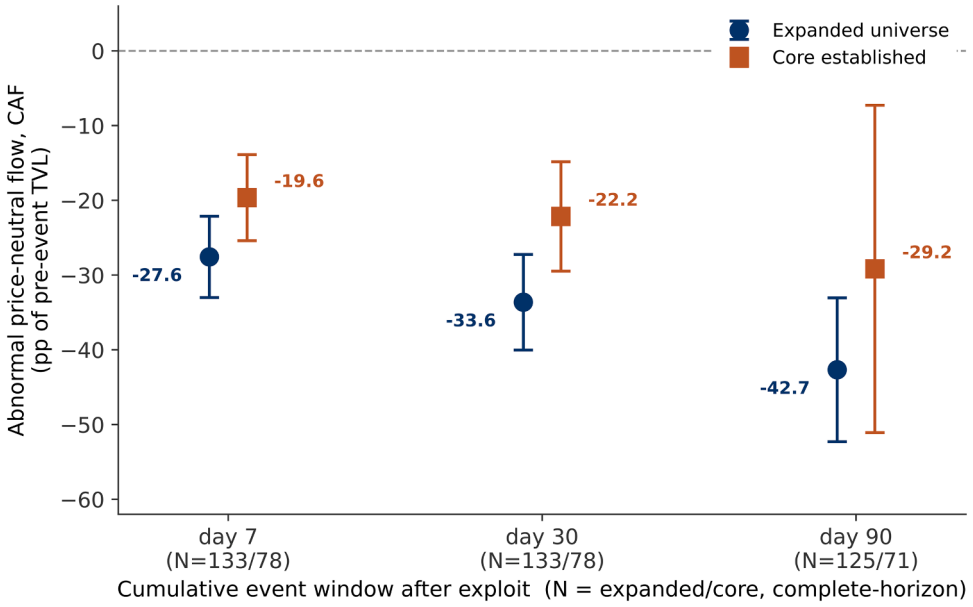


Fig. 1. Matched event-window estimates of abnormal price-neutral liquidity (CAF) after smart-contract exploits.
Notes: TVL denotes total value locked; CAF denotes cumulative abnormal price-neutral flow; pp denotes percentage points. CAF[0,h] cumulates daily price-neutral flows from the exploit day (day 0) through day h, divided by pre-event TVL. Points are affected-minus-matched-control event-window coefficients with incident fixed effects and log pre-event TVL adjustment. The expanded sample includes all eligible protocols, whereas the core sample is the predefined established-protocol subset. Whiskers show 95% confidence intervals based on standard errors clustered by incident and protocol. Incident counts at days 7/30/90 are 133/133/125 for the expanded sample and 78/78/71 for the core sample. CAF outcomes are winsorized at the 1st and 99th percentiles.

million, matched-versus-unmatched selection (Table A.14), and the chain diagnostics of Table A.16.

4. Results

4.1. Immediate abnormal liquidity loss

Table 1 and Fig. 1 report the matched estimates. In the expanded matched sample, exploited protocols lose 27.6 percentage points (pp) of pre-event TVL in abnormal price-neutral liquidity by day 7, 33.6 points by day 30, and 42.7 points by day 90 (125 incidents with complete follow-up). The core sample shows the same pattern at smaller magnitudes: a seven-day loss of 19.6 percentage points with comparable persistence. The seven-day estimate is statistically robust, with an add-one permutation p-value below 0.0001 and significance under two-way clustering.

4.2. Robustness, scale, and chain conditions

Raising the pre-event TVL floor from $1 million to $100 million shrinks the sample but does not eliminate the effect: 64 usable incidents remain (33 matched), and the seven-day loss is still 14.6 percentage points. The estimate strengthens in code-vulnerability incidents excluding bridges (−33.3 pp, 98 matched incidents) and is far smaller for operational-security incidents (−7.4 pp, 25 matched incidents including three bridge key compromises; Online Appendix Table A.10). Magnitudes attenuate with scale, but severity is measured relative to TVL, so the floor gradient partly reflects large protocols' lower relative severity. The seven-day estimate uses the pre-specified 1/99 winsorization and stays negative and significant under every regression tail treatment in Table A.12 (none through 5/95 winsorization and 1/99 trimming).

Chain-specific liquidity conditions do not drive the estimate; three diagnostics in Online Appendix Table A.16 support this conclusion. First, under the paper's estimator, cumulative chain-day conditions differ between affected protocols and their matched controls by −0.18 percentage points (95% confidence interval −0.60 to +0.24). Second, netting each unit's own chain conditions out of the outcome, on identical samples, moves the seven-day estimate by 0.15 points and the 30-day estimate by 0.02 points. Third, among the 78 incidents with same-chain support, restricting the original matched sets to same-chain controls changes the seven-day estimate by 1.7 percentage points.

4.3. Liquidity scarring and severe impairment

The short-run loss remains evident at day 180. Under the price-neutral classification, 37.5% of coverage-qualified affected

incidents are severely impaired, 41.2% are scarred, and 21.3% recover (n = 136; Table 2, Panel A). Among matched control observations, 11.8% are severely impaired and 52.0% recover. The adjusted matched severe-impairment difference is 28.1 percentage points (SE 5.0, 92 matched incident sets), and the recovery difference is −35.7 points (SE 4.1, p < 0.0001; Table A.15). Raw-TVL classification yields a higher severe-impairment share (51.2%; Table 2, notes), but it also incorporates token-price revaluation. Among 65 raw-severe incidents that also meet the flow-coverage rule, 38 remain severe under the price-neutral classification. The matched severe-impairment gap survives excluding negative price-neutral ratios (26.2 points; Table A.13).

4.4. Severity and limited spillovers

Short-run price-neutral losses and raw-TVL impairment rise sharply with relative exploit severity. Under the coverage-qualified price-neutral classification, the share of severely impaired incidents also rises monotonically across severity quartiles (Table A.9). Across quartiles the median seven-day loss widens from 3.1% to 48.0% and the raw-TVL impairment rate climbs from 23.7% to 73.2% (Table 2, Panel B; Fig. 2). A matched Affected × severity interaction confirms the gradient (−36.5 pp, $p < 0.0001$, capped severity; negative under log and raw; Table A.15).

Comparing median net outflow with median reported loss (Table 2, Panel B) helps gauge outflows beyond the mechanical loss. In the two lowest-severity quartiles the seven-day outflow exceeds the reported loss (Q1: 3.1% versus 0.6%; Q2: 9.5% versus 7.8%). In the core affected-only regression, the CAF[2,7] intercept is −5.8 pp (p = 0.001), controlling for CAF[0,1] and relative severity; both patterns are consistent with reassessment beyond the theft (Online Appendix E, Table A.5). In the two highest quartiles the reported loss meets or exceeds the outflow (Q3: 26.7% versus 21.1%; Q4: 89.1% versus 48.0%), so reassessment cannot be separated from theft there.

Peer effects are far smaller. The average matched-peer response is about −2.4%, against −23.4% for affected protocols (core-sample group means), and an equivalence test based on two one-sided tests (TOST) bounds it within the pre-specified ±5-percentage-point margin ($p < 0.0001$; Table A.6). Together the results support H1 through H4. Protocol scale is treated as a robustness dimension, not as a separately identified effect.

Table 2
Post-exploit scarring, severity gradient, and scale sensitivity.

Panel A. 180-day price-neutral states

180-day state	Affected incidents	Share (%)
Recovered	29	21.3
Scarred	56	41.2
Severely impaired	51	37.5

Panel B. Outcomes by exploit severity quartile

Severity quartile	Median severity (%)	Median affected-incident CAF[0,7] (% of pre-event TVL)	Median raw TVL_{180}/TVL_{-1} (%)	Raw-TVL severe impairment at day 180 (%)
Q1	0.6	−3.1	49.9	23.7
Q2	7.8	−9.5	32.1	45.0
Q3	26.7	−21.1	15.8	60.0
Q4	89.1	−48.0	1.9	73.2

Panel C. Selected TVL-floor sensitivity

TVL floor	Usable incidents	Matched incidents	Matched CAF[0,7] coefficient (pp)	Raw-TVL severe impairment at day 180 (%)
$1M	197	133	−27.6	51.2
$10M	143	90	−23.9	49.6
$50M	89	52	−15.0	40.8
$100M	64	33	−14.6	35.2

Notes: TVL denotes total value locked; CAF denotes cumulative abnormal price-neutral flow; PNR_{180} denotes the price-neutral ratio, one plus cumulative net price-neutral flow over [0,180] divided by pre-event TVL; pp denotes percentage points; SE denotes standard error. Panel A reports price-neutral liquidity states for the 136 affected incidents whose 180-day window has observed daily flows on at least half of its days (Online Appendix Table A.17): recovered, $PNR_{180} \geq 0.80$; severely impaired, $PNR_{180} \leq 0.20$; scarred, a ratio strictly between these thresholds. Among the 1,586 coverage-qualified matched control observations, 11.8% are severely impaired; the adjusted matched risk difference is +28.1 pp (SE 5.0, p < 0.0001). Descriptive shares and adjusted matched risk differences are distinct statistics. Affected-state shares use all of these incidents, whereas adjusted risk differences use 92 matched incident sets with incident fixed effects and log pre-event TVL adjustment. The adjusted estimates therefore need not equal the differences between the displayed descriptive shares. In Panel B, CAF[0,7] is price-neutral, whereas the day-180 TVL ratio and severe-impairment share use raw TVL. In Panel C, CAF[0,7] is the matched price-neutral coefficient, whereas the severe-impairment share uses raw TVL. Fig. 2 also reports raw-TVL states. Raw-TVL states apply the same cutoffs to TVL_{180}/TVL_{-1}, where TVL_{180} is TVL on day 180 and TVL_{-1} is pre-event TVL, the last TVL observed in the seven days before the exploit. Among the 164 incidents with analytic raw-TVL outcomes at day 180, 12.8% are recovered, 36.0% are scarred, and 51.2% are severely impaired; of the 65 raw-severe incidents that also meet the flow-coverage rule, 38 remain severe under PNR_{180}. In Panel B, severity is reported dollar loss divided by pre-event TVL, with Q1 the lowest-severity quartile (cut points set on all 197 usable incidents). Median severity and median CAF[0,7] (the affected incident's own seven-day flow, not control-differenced) use all 197 usable incidents (50, 49, 49, and 49 by quartile); the two day-180 columns use the 164 incidents with a day-180 raw-TVL observation (38, 40, 45, and 41). In Panel C, usable incidents meet the pre-event TVL floor, matched incidents retain at least five controls, and the day-180 severe-impairment shares use 164, 117, 71, and 54 incidents at the four floors. See Online Appendix Tables A.4, A.9, and A.17.

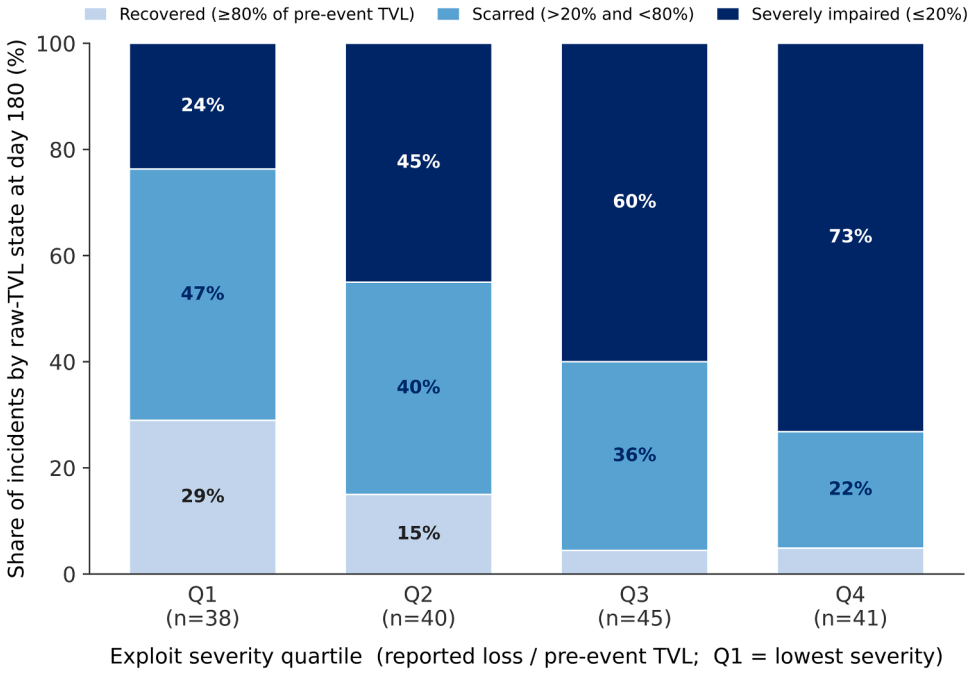


Fig. 2. Raw-TVL 180-day outcome composition by exploit-severity quartile (expanded sample; the principal price-neutral states appear in Table 2, Panel A).

Notes: This Fig. reports raw-TVL states, not price-neutral states. TVL denotes total value locked; TVL_{180} is TVL on day 180 and TVL_{-1} is pre-event TVL, the last TVL observed in the seven days before the exploit. Recovered means $TVL_{180}/TVL_{-1} \geq 0.80$; severely impaired means $TVL_{180}/TVL_{-1} \leq 0.20$; scarred means a ratio strictly between these thresholds. Severity is reported dollar loss divided by pre-event TVL, with Q1 denoting the lowest-severity quartile; quartile cut points are set on all 197 usable incidents. Labels show rounded within-quartile percentages; shares below 6% are not labeled. n denotes incident observations with analytic raw-TVL outcomes at day 180 (164 in total).

5. Discussion and conclusion

Smart-contract exploits generate large, immediate, and persistent protocol-level liquidity losses. The measure excludes pure price revaluation of unchanged token balances, while valuing quantity changes at contemporaneous prices. The estimates therefore capture balance-sheet quantities rather than valuation effects, and the matched design supports an incident-specific interpretation.

The findings connect DeFi operational risk to run-like liquidity discipline (Diamond & Dybvig, 1983; Gorton & Metrick, 2012). An exploit is a public fundamental shock, not a purely self-fulfilling run; yet immediate exit, mobile liquidity, and absent backstops let operational failure translate rapidly into withdrawals. This behavior matches the franchise-capital reading of operational trust: once code, security controls, or governance credibility are compromised, liquidity impairment can exceed the direct loss and persist well beyond the immediate event.

Impairment rises sharply with relative exploit severity. In the two lowest-severity quartiles, median net outflows exceed median reported losses, consistent with reassessment beyond mechanical theft; transaction-level tracing is needed to isolate the behavioral component. Matched-peer effects, by comparison, stay within the pre-specified equivalence band, indicating a predominantly protocol-specific short-run response (Eisenbach et al., 2022).

The analysis identifies post-exploit outcomes, not exploit selection, and long-run trajectories also incorporate reimbursements, rescue capital, migrations, and governance responses. Code-based intermediation does not eliminate trust; it relocates trust to software, security, and remediation capacity. Relative exploit severity and price-neutral liquidity flows are therefore useful post-incident indicators of persistent impairment that aggregate TVL does not capture.

CRediT authorship contribution statement

Michael Zouari: Writing – original draft, Visualization, Validation, Methodology, Investigation, Data curation, Conceptualization. **Ilan Alon:** Writing – review & editing. **Zeev Shtudiner:** Writing – review & editing.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.frl.2026.110877.

Data availability

Replication code and processed data are on OSF: https://osf.io/4mbdf/?view_only=afea56e0be2a4ce7ad7a64e220d4c3c5 (view-only for review)

References

Abadie, A., Imbens, G.W., 2006. Large sample properties of matching estimators for average treatment effects. Econometrica 74 (1), 235–267. https://doi.org/10.1111/j.1468-0262.2006.00655.x.

Azar, P.D., Olivas, S., Sinha, N.D., 2025. The price of processing: Information frictions and market efficiency in DeFi (Staff Report No. 1153). Federal Reserve Bank of New York. https://doi.org/10.59576/sr.1153.

Bertomeu, J., Martin, X., Sall, I., 2024. Measuring DeFi risk. Finance Research Letters 63, 105321. https://doi.org/10.1016/j.frl.2024.105321.

Bhambhwani, S.M., Huang, A.H., 2024. Auditing decentralized finance. The British Accounting Review 56 (2), 101270. https://doi.org/10.1016/j.bar.2023.101270.

Bourveau, T., Brendel, J., Schoenfeld, J., 2024. Decentralized finance (DeFi) assurance: Early evidence. Review of Accounting Studies 29 (3), 2209–2253. https://doi.org/10.1007/s11142-024-09834-8.

Callaway, B., Sant'Anna, P.H.C., 2021. Difference-in-differences with multiple time periods. Journal of Econometrics 225 (2), 200–230. https://doi.org/10.1016/j.jeconom.2020.12.001.

Cameron, A.C., Gelbach, J.B., Miller, D.L., 2011. Robust inference with multiway clustering. Journal of Business & Economic Statistics 29 (2), 238–249. https://doi.org/10.1198/jbes.2010.07136.

Campello, M., Gallo, A., Mota, L., Terracciano, T., 2026. Demand for safety in the crypto ecosystem (NBER Working Paper No. 35557). National Bureau of Economic Research. https://doi.org/10.3386/w35557.

Caporale, G.M., Kang, W.-Y., Spagnolo, F., Spagnolo, N., 2020. Non-linearities, cyber attacks and cryptocurrencies. Finance Research Letters 32, 101297. https://doi.org/10.1016/j.frl.2019.09.012.

Chen, Y.-L., Chang, Y.T., Yang, J.J., 2023. Cryptocurrency hacking incidents and the price dynamics of Bitcoin spot and futures. Finance Research Letters 55, 103955. https://doi.org/10.1016/j.frl.2023.103955.

Cong, L.W., He, Z., 2019. Blockchain disruption and smart contracts. The Review of Financial Studies 32 (5), 1754–1797. https://doi.org/10.1093/rfs/hhz007.

Diamond, D.W., Dybvig, P.H., 1983. Bank runs, deposit insurance, and liquidity. Journal of Political Economy 91 (3), 401–419. https://doi.org/10.1086/261155.

Eisenbach, T.M., Kovner, A., Lee, M.J., 2022. Cyber risk and the U.S. financial system: A pre-mortem analysis. Journal of Financial Economics 145 (3), 802–826. https://doi.org/10.1016/j.jfineco.2021.10.007.

Goldstein, I., Pauzner, A., 2005. Demand–deposit contracts and the probability of bank runs. The Journal of Finance 60 (3), 1293–1327. https://doi.org/10.1111/j.1540-6261.2005.00762.x.

Gorton, G., Metrick, A., 2012. Securitized banking and the run on repo. Journal of Financial Economics 104 (3), 425–451. https://doi.org/10.1016/j.jfineco.2011.03.016.

Grande, M., Borondo, J., 2025. Trust as a driver in the DeFi market: Leveraging TVL/MCAP bands as confidence indicators to anticipate price movements. Finance Research Letters 75, 106705. https://doi.org/10.1016/j.frl.2024.106705.

Kamiya, S., Kang, J.-K., Kim, J., Milidonis, A., Stulz, R.M., 2021. Risk management, firm reputation, and the impact of successful cyberattacks on target firms. Journal of Financial Economics 139 (3), 719–749. https://doi.org/10.1016/j.jfineco.2019.05.019.

Kitzler, S., Paquet-Clouston, M., Haslhofer, B., 2025. The economic impact of DeFi crime events on decentralized autonomous organizations (DAOs). The Journal of Finance and Data Science 11, 100171. https://doi.org/10.1016/j.jfds.2025.100171.

Lin, T. (2026). A blessing in disguise? DeFi exploits and short-horizon responses in U.S. commercial paper spreads (Working paper). arXiv. https://arxiv.org/abs/2601.08263.

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

Sun, L., Abraham, S., 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. Journal of Econometrics 225 (2), 175–199. https://doi.org/10.1016/j.jeconom.2020.09.006.

Young, A., 2019. Channeling fisher: randomization tests and the statistical insignificance of seemingly significant experimental results. Q. J. Econ. 134 (2), 557–598. https://doi.org/10.1093/qje/qjy029.