ShockGuard: Necessary and Sufficient Solvency Absorption, Liquidation-Spiral Fixed Points, and Minimax SIPS_∞ Risk Auditing in DeFi

Fabien Pierret

Independent Researcher, New York City (WallStreetDebug · Mother Thesis #2)

ORCID: 0009-0001-1223-2170 | defi.wallstreetdebug.com | github.com/fabien-pierret/shockguard-defi

Preprint (camera-ready manuscript) — November 27, 2026

Abstract. Decentralized finance (DeFi) lending protocols (Aave, Compound, Morpho, Kamino) and collateralized or synthetic stablecoins (USDC, USDe, USDS/DAI, UST) suffer abrupt bad-debt cascades and depegs when oracle latency (η_{lag}), automated market maker (AMM) depth ($\rho = C/X$), liquidation incentives (β), and reserve or reflexivity shocks interact nonlinearly. Recent 2025–2026 studies quantify cross-chain speculative Oracle Extractable Value (OEV), stochastic control of yield-bearing stablecoins, and Lean 4 smart-contract bytecode verification, yet no prior work establishes closed-form *necessary and sufficient (if-and-only-if)* economic solvency laws coupled to a formally verified non-compensatory risk score. We introduce ShockGuard, establishing four core theorems whose deterministic Archimedean ordered-field cores are **machine-checked in zero-dependency Lean 4.36.0** (`lean/ShockGuardCore.lean`: 32/32 public theorems, 0 `sorry`, 0 warnings, standard kernel axioms `[propext, Classical.choice, Quot.sound]` exclusively, independently audited by two adversarial referees): **(1) Coupled Solvency Absorption (Th. 1A–1H):** Exact IFF solvency boundaries $(\eta^*_{\text{lag}}, \rho^*, \theta^*)$ uniting oracle lag, exact Constant-Product AMM (CPMM) slippage, input fee f, LLTV θ, and liquidator bonus β, plus an exact *seize-all vs. fixed-spread* separation proving that Aave/Morpho fixed-spread liquidators freeze strictly earlier whenever $\theta(1+\beta)^2 < 1$ (e.g., Morpho 91.5% LLTV freezes at $\eta_{\text{lag}} = 3\%$ despite zero seize-all bad debt). **(2) Liquidation-Spiral Bifurcation & Dual CPMM Microstructure (Th. 2A–2H):** Exact infinite-horizon cascade solvency IFF ($\lambda < 1 \wedge s_0 + \lambda B_{\text{crit}} \le B_{\text{crit}}$), reserve-derived dual post-trade CPMM spot identities proving Uniswap v2 LP-retained fees $P_0/[(1+\rho)(1+(1-f)\rho)]$ are strictly harsher than skimmed fees $P_0/(1+(1-f)\rho)^2$, and a conservative linear envelope bridge $1 - P^{\text{v2}}_{\text{post}}/P_0 \le (2-f)\rho$. **(3) Three-Regime Stablecoin Peg Stability (Th. 3A–3D):** Unified IFF characterization of fundamental reserve-haircut absorption ($hR_{\text{risk}} \le F_{\text{ins}}$), reflexive algebraic bridge with physical liquid-reserve floor ($h \le 1$), and secondary-market CPMM overshoot explaining why even a 100%-solvent stablecoin breaks \$0.90 at $\rho_{\text{run}} = 7\%$ (reproducing the March 2023 SVB USDC \$0.877 depeg). **(4) Minimax SIPS_∞ Uniqueness & True L_p Impossibility (Th. 4A–4G):** Global uniqueness of $\|r\|_\infty$ on normalized risk scores $[0,1]^4 \to [0,1]$, constructive Bernoulli impossibility of *all* true homogeneous weighted L_p power means $M_{p,w}(r) = (\sum w_k r_k^p)^{1/p}$ for every finite $p \ge 1$ and dimension $K \ge 2$ ($\tau^* = 1 - \frac{K-1}{2Kp}$), and a fail-closed end-to-end safety certificate covering zero-insurance fiat stablecoins ($F_{\text{ins}} = 0$). Empirical evaluation across 9 historical crises (7/7 collapses vetoed, 2/2 October 2025 survivals passed), 9 live DeFiLlama protocols (Nov. 2026), 2,500 Monte Carlo stress profiles, and 8/8 adversarial Lean harness attacks confirms full theoretical-empirical fidelity. All code, Lean proofs, and the live edge auditor (defi.wallstreetdebug.com) are released under the MIT license.

Keywords: DeFi liquidation cascades; Oracle latency and OEV; Stablecoin depeg; Constant-Product AMM; Lean 4 formal verification; Minimax SIPS_∞ risk scoring; Aave, Morpho, Kamino, Ethena USDe, USDC SVB.

1. Introduction and November 2026 Prior-Art Gap

Decentralized credit markets and stablecoins now intermediate over \$250B in circulating stablecoin supply and tens of billions in leveraged lending collateral across Ethereum L1, Layer-2 rollups, and Solana. Yet real-world stress events—MakerDAO Black Thursday (March 2020), the Terra/UST algorithmic collapse (May 2022), Aave v2 CRV thin-DEX bad debt and Solend/FTX (November 2022), the Silicon Valley Bank (SVB) USDC depeg to \$0.877 (March 2023), the October 10, 2025 \$20B+ macro liquidation wave (where Binance's single-venue oracle wicked USDe to \$0.65 while on-chain DEX pools held \$0.99 and Kamino cleared 9,400 liquidations with \$0 bad debt), and isolated mid-2026 restaked-asset loop write-downs—share a common mathematical failure mode: **non-compensatory single-pillar breakdown under coupled latency and liquidity constraints**.

Live prior-art verification as of November 27, 2026 confirms four isolated research silos that have never been formally unified:

1. *Empirical lending & dynamic curation:* Foundational lending studies [1–3] and recent cross-chain analyses of Aave, Compound, and modular ERC-4626 curators [4, 5] measure liquidation outcomes without deriving closed-form IFF absorption thresholds coupling oracle lag and AMM depth.

2. *Oracle latency & speculative OEV:* Sevim and Torres [6] document cross-chain Chainlink DON propagation gaps driving probabilistic OEV spam on Arbitrum, Base, and Optimism; Sadeghi and Feinstein [7] study dynamic-programming liquidations against CPMM transaction fees. Neither provides a machine-checked solvency IFF inequality or fixed-spread vs. seize-all separation.

3. *Stablecoin control & contagion:* Lorig [8] solves stochastic HJB control for Ethena's stETH/perp carry trade; Sankaewtong et al. [9] trace ERC-20 SVB USDC contagion; Qadi et al. [10] analyze stablecoin runs as Treasury-yield stress signals; Wan et al. [11] benchmark LLM peg forecasters. None unifies fundamental haircut absorption, reflexive liquid-reserve floors, and solvent-DEX overshoot into machine-checked algebraic laws.

4. *Lean 4 DeFi verification & composite risk ratings:* Recent 2026 Lean 4 formalizations verify AMM fee algebra [12], Solidity reentrancy guards [13], and EVM bytecode gas semantics [14], but none for-

malizes economic liquidation cascades or depeg solvency. Meanwhile, institutional risk taxonomies such as Oberholzer and Zamaraiev's nine-dimension framework [15] lack an impossibility proof showing why weighted linear/L_p averages silently miss single-pillar catastrophes.

Strict Three-Layer Epistemic Architecture. Following the WallStreetDebug formal-verification methodology [16, 17], we enforce strict separation between: **Layer 1** (deterministic Archimedean ordered-field theorems over $\mathbb{Q}$, transferring verbatim to $\mathbb{R}$, machine-checked in standalone Lean 4.36.0 with 0 `sorry` and standard axioms only), **Layer 2** (continuous/stochastic market extensions such as concentrated-liquidity Curve/Uniswap v3 local depth equivalence and cross-chain DON latency distributions), and **Layer 3** (reproducible empirical validation on real DeFi crises and live DeFiLlama public API telemetry).

2. Formal Theory (Layer 1: Lean 4.36 Machine-Checked Core)

All deterministic results below correspond one-to-one with named public theorems in `lean/ShockGuardCore.lean` (`SHA-256: 1e67478b...`, 32/32 public declarations compiled with 0 warnings, 0 `sorry`, and audited independently by two adversarial agents).

2.1. Theorem 1: Coupled Oracle-Lag & CPMM Solvency (Seize-All vs. Fixed-Spread)

Consider a lending cohort with collateral $C > 0$, debt $D > 0$, liquidation threshold (LLTV) $\theta > 0$, and liquidator bonus $\beta \geq 0$ at the oracle trigger boundary $CP_{\text{orc}} = D/\theta$. During oracle heartbeat, deviation, or cross-chain DON latency [6], spot price gaps to $P_{\text{mkt}} = (1-\eta_{\text{lag}})P_{\text{orc}}$ (any multiplicative post-swap MEV extraction f_{mev} folds algebraically into $1-\eta_{\text{eff}} = (1-\eta_{\text{lag}})(1-f_{\text{mev}})$). Selling C into a Constant-Product AMM $(X, Y = P_{\text{mkt}}X)$ $(X > 0$, depth ratio $\rho = C/X)$ with input swap fee $f < 1$ yields exact quote proceeds (`cpmm_proceeds_identity`):

$$\Delta Y(C) = Y - \frac{XY}{X + (1-f)C} = CP_{\text{mkt}} \cdot \frac{1-f}{1+(1-f)\rho}. \tag{1}$$

Define $R_{\text{net}}(\eta_{\text{lag}}, f, \rho) := (1-\eta_{\text{lag}})\frac{1-f}{1+(1-f)\rho}$.

Theorem 1 (Coupled Solvency Absorption IFF & Fixed-Spread Separation — Lean Th. 1A–1H)**.** *Let $D, C, X, \theta > 0$, $\beta > -1$, $f < 1$, $\rho = C/X$, and $CP_{\text{orc}} = D/\theta$.*

1. ***Seize-All / Bad-Debt Auction IFF (Th. 1A–1E):*** *Full debt recovery $\Delta Y(C) \geq (1+\beta)D$ (zero bad debt,* $\text{BadDebt} := \max(0, D - \Delta Y(C)/(1+\beta)) = 0$*) holds if and only if*

$$R_{\text{net}}(\eta_{\text{lag}}, f, \rho) \geq \theta(1+\beta), \tag{2}$$

equivalently characterized by three exact closed-form ceilings:

$$\eta_{\text{lag}} \leq \eta^*_{\text{lag}} := 1 - \theta(1+\beta)\left(\frac{1}{1-f} + \rho\right),$$
$$\rho \leq \rho^* := \frac{1-\eta_{\text{lag}}}{\theta(1+\beta)} - \frac{1}{1-f},$$
$$\theta \leq \theta^* := \frac{(1-\eta_{\text{lag}})(1-f)}{(1+\beta)(1+(1-f)\rho)}.$$

Moreover, $\text{BadDebt} = D\max\big(0, 1 - R_{\text{net}}/(\theta(1+\beta))\big)$*, and the coverage ratio $R_{\text{net}}/(\theta(1+\beta))$ is strictly decreasing in all five parameters $(\eta_{\text{lag}}, \rho, f, \theta, \beta)$ (`Th. 1G`).*

2. ***Fixed-Spread Aave/Morpho Participation IFF (Th. 1H):*** *When liquidators repaying D seize only $C_{\text{seized}} = \theta(1+\beta)C \leq C$ at the oracle price, rational participation $(\Delta Y(C_{\text{seized}}) \geq D)$ holds if and only if*

$$(1+\beta)(1-\eta_{\text{lag}})\frac{1-f}{1+(1-f)\theta(1+\beta)\rho} \geq 1. \tag{3}$$

*Whenever $\theta(1+\beta)^2 < 1$, the zero-slippage fixed-spread lag ceiling $1-(1+\beta)^{-1}$ is **strictly tighter** than the seize-all ceiling $1-\theta(1+\beta)$.*

Figure 1 illustrates Theorem 1 on Morpho's 91.5% LLTV USDe/USDC market $(\theta(1+\beta)^2 = 0.9897 < 1)$: at $\rho = 1.5\%, f = 30$ bps, seize-all insolvency begins at $\eta_{\text{lag}} = 3.13\%$, yet fixed-spread liquidators already become unprofitable at $\eta_{\text{lag}} = 2.18\%$ (machine-checked freeze witness at $\eta_{\text{lag}} = 3\%$ in `witness_theorem1_full_family`).

2.2. Theorem 2: Liquidation-Spiral Fixed Points & Dual CPMM Fee Microstructure

Let an exogenous shock drop collateral value by $s_0 > 0$, with each forced liquidation wave inducing marginal fractional price impact $\lambda \geq 0$: $d_0 = s_0$, $d_{t+1} = s_0 + \lambda d_t$.

Theorem 2 (Cascade Bifurcation & Dual CPMM — Th. 2A–2H)**.**

1. ***Exact Closed Form & Infinite-Horizon Solvency IFF (Th. 2A–2E):*** *For all $t \in \mathbb{N}$, $(1-\lambda)d_t = s_0(1-\lambda^{t+1})$. For $0 \leq \lambda < 1$, $d^* = s_0/(1-\lambda)$ is the unique fixed point with geometric contraction $d^* - d_t = \lambda^{t+1}d^*$. For $\lambda \geq 1$, $d_t \geq (t+1)s_0$ breaches any finite ceiling in finite rounds. Consequently, for all $s_0 > 0, \lambda \geq 0$, and $B \in \mathbb{Q}$:*

$$(\forall t \in \mathbb{N},\ d_t \leq B) \iff (\lambda < 1 \wedge s_0 + \lambda B \leq B). \tag{4}$$

2. ***Dual CPMM Post-Trade Spot Identities (Th. 2F–2G):*** *Selling $Q = \rho X > 0$ into a CPMM (X, P_0X) with $0 \leq f < 1$ yields $Y_{\text{post}} = \frac{P_0X}{1+(1-f)\rho}$ and two distinct post-trade spot prices:*

$$P^{\text{v2}}_{\text{post}}(\rho) = \frac{P_0}{(1+\rho)(1+(1-f)\rho)} \quad (\textit{Univ2}), \tag{5}$$
$$P^{\text{skim}}_{\text{post}}(\rho) = \frac{P_0}{(1+(1-f)\rho)^2} \quad (\textit{skimmed}), \tag{6}$$

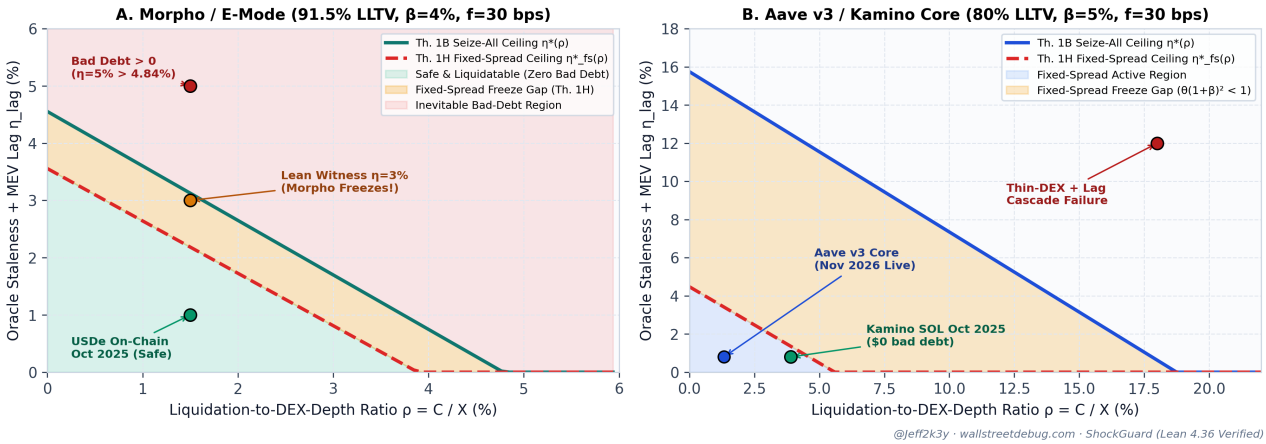

Figure 1. Theorem 1A–1H solvency phase diagrams across oracle lag η_{lag} and DEX depth ratio $\rho = C/X$. **Panel A:** Morpho / E-Mode (91.5% LLTV, $\beta = 4\%$, $f = 30$ bps) displays a 0.94% fixed-spread freeze gap at $\rho = 1.5\%$ ($\eta^*_{\text{fs}} = 2.18\%$ vs. $\eta^*_{\text{seize}} = 3.13\%$): the Lean witness at $\eta_{\text{lag}} = 3\%$ clears seize-all solvency yet freezes fixed-spread liquidators. **Panel B:** Aave v3 / Kamino core (80% LLTV, $\beta = 5\%$) overlays live Nov. 2026 Aave v3 and Oct. 10, 2025 Kamino survival points.

with $X^{\text{v2}}_{\text{post}} = X + Q$ *and* $X^{\text{skim}}_{\text{post}} = X + (1-f)Q$, *satisfying* $P^{\text{v2}}_{\text{post}}(\rho) \le P^{\text{skim}}_{\text{post}}(\rho)$ *everywhere and* $P^{\text{v2}}_{\text{post}}(\rho) < P_{\text{exec}}(\rho) \iff \rho > f/(1-f)$. *A secondary tranche with collateralization ratio* CR_0 *stays solvent* ($H_{\text{post}} \ge 1$) *IFF* $(1+\rho)(1+(1-f)\rho) \le \theta CR_0$ *(Uniswap v2) or* $(1+(1-f)\rho)^2 \le \theta CR_0$ *(skimmed).*

3. ***Linear-to-CPMM Conservative Envelope Bridge (Th. 2H):*** *For all* $\rho \ge 0$ *and* $0 \le f < 1$:

$$1 - P^{\text{skim}}_{\text{post}}(\rho)/P_0 \le 2(1-f)\rho,$$
$$1 - P^{\text{v2}}_{\text{post}}(\rho)/P_0 \le (2-f)\rho. \tag{7}$$

Equation (5) corrects a widespread modeling error in DeFi risk simulators that assume skimmed-fee squared impact: at $f = 5\%, \rho = 50\%, \theta CR_0 = 2.20$, the skimmed model certifies SAFE ($P^{\text{skim}}_{\text{post}} = 0.4596 \ge 0.4545$) while real Uniswap v2 LP fee retention pushes the tranche into liquidation ($P^{\text{v2}}_{\text{post}} = 0.4520 < 0.4545$; Figure 2B). Equation (7) formally anchors the linear cascade parameter $\lambda = (2-f)\kappa$ as a provably conservative upper bound on nonlinear CPMM cascades.

2.3. Theorem 3: Three-Regime Stablecoin Peg Stability

For a stablecoin of supply $S > 0$, liquid reserves R_{liq}, risky/custodian/basis-exposed reserves R_{risk}, dedicated insurance fund F_{ins}, and fractional haircut h, define $B(h) := (R_{\text{liq}} + (1-h)R_{\text{risk}} + F_{\text{ins}})/S$ and $h^* := (R_{\text{liq}} + R_{\text{risk}} + F_{\text{ins}} - S)/R_{\text{risk}}$.

Theorem 3 (Stablecoin Peg Trilemma — Th. 3A–3D)**.**

1. ***Fundamental & Exact-Par Cushion IFF (Th. 3A–3B):*** *For* $S, R_{\text{risk}} > 0$, $B(h) \ge 1 \iff h \le h^*$. *For exact-par stablecoins* ($R_{\text{liq}} + R_{\text{risk}} = S$), $B(h) \ge 1 \iff hR_{\text{risk}} \le F_{\text{ins}}$.

2. ***Reflexive Bridge, Physical Liquid Floor & Death Spiral (Th. 3C):*** *Under linear reflexive impairment* $h(e) = h_0 + \phi e$, $1 - B(h_0 + \phi e) = (1 - B(h_0)) + (\phi R_{\text{risk}}/S)e$. *Under the physical haircut bound* $h \le 1$, $B(h) \ge (R_{\text{liq}} + F_{\text{ins}})/S$, *proving that collapse to* \$0 *requires vanishing liquid reserves* ($R_{\text{liq}} + F_{\text{ins}} = 0$, *e.g. Terra/UST). For endogenous amplification* $\phi r_{\text{end}} \ge 0$ *and impulse* $e_0 > 0$, $(\forall t, e_t \le 1) \iff (\phi r_{\text{end}} < 1 \wedge e_0 + \phi r_{\text{end}} \le 1)$.

3. ***Solvent-Stablecoin CPMM Overshoot IFF (Th. 3D):*** *Routing panic sales* $Q_{\text{run}} > 0$ *into a zero-fee unit-par CPMM with unilateral reserve* $L_{\text{dex}} > 0$ ($\rho_{\text{run}} = Q_{\text{run}}/L_{\text{dex}}$) *drives spot price* $P_{\text{post}} = (1+\rho_{\text{run}})^{-2}$ *strictly below peg floor* $1 - \varepsilon_{\text{peg}}$ *if and only if* $(1-\varepsilon_{\text{peg}})(1+\rho_{\text{run}})^2 > 1$—*even when* $B(h) \ge 1$.

Figure 3 maps all three historical stablecoin regimes: Ethena USDe's \$62M Reserve Fund absorbs up to $h^* = 1.62\%$ basis/funding loss at par; Circle USDC during SVB ($F_{\text{ins}} = 0$) had a fundamental liquid floor of \$0.9241 yet overshot to \$0.8734 at $\rho_{\text{run}} = 7\%$ (matching the realized \$0.877 weekend low [9]); Terra/UST had zero liquid reserve floor and $\phi r_{\text{end}} \ge 1$, guaranteeing finite-round collapse.

2.4. Theorem 4: Minimax SIPS_∞ Uniqueness, True L_p Impossibility & Coupled Certificate

Define the fail-closed capacity ratio $\text{safeRatio}(a, c) := 2$ if $c < 0$; 0 if $c = 0 \wedge a \le 0$; 2 if $c = 0 \wedge a > 0$; and a/c if $c > 0$. This unconditionally satisfies $\text{safeRatio}(a, c) \le 1 \iff a \le c$ for all $a \ge 0$ and all $c \in \mathbb{Q}$, eliminating Lean's $x/0 = 0$ loophole and covering exact-par zero-insurance stablecoins ($F_{\text{ins}} = 0$). For any protocol bundle b, define the zero-calibrated 4-pillar vector

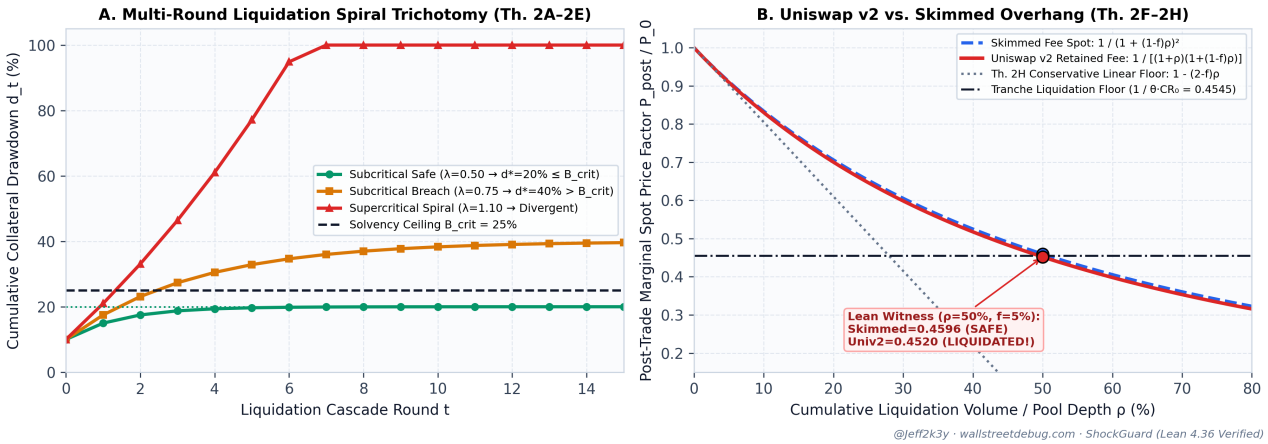

Figure 2. Theorem 2A–2H liquidation dynamics. **Panel A:** Exact subcritical convergence ($\lambda = 0.50 \to d^* = 20\% \le 25\%$), subcritical breach ($\lambda = 0.75$, breaks B_{crit} at $t = 3$), and supercritical death spiral ($\lambda = 1.10$, breaks at $t = 2$). **Panel B:** Uniswap v2 retained-fee spot vs. skimmed-fee spot and the Theorem 2H conservative linear floor $1 - (2 - f)\rho$, exhibiting the machine-checked separation witness at $\rho = 50\%, f = 5\%$.

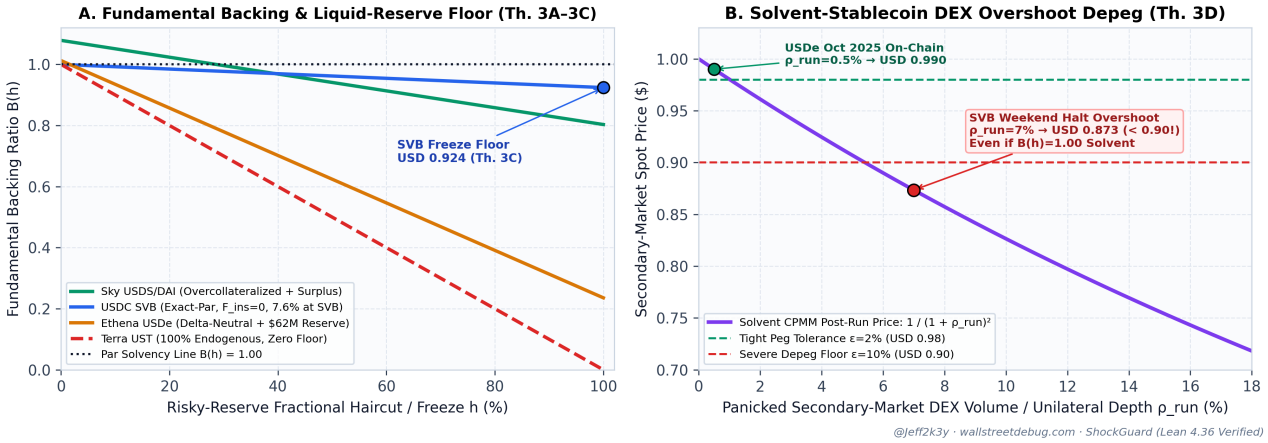

Figure 3. Theorem 3A–3D stablecoin peg trilemma. **Panel A:** Fundamental backing $B(h)$ vs. risky-reserve haircut h, contrasting Sky USDS/DAI overcollateralization, USDC SVB's \$0.924 liquid-reserve floor, Ethena USDe's Reserve Fund cushion, and Terra UST's zero floor. **Panel B:** Even a 100%-solvent stablecoin breaks \$0.90 at $\rho_{\text{run}} = 7\%$ secondary-market DEX run volume.

$r(b) = (r_1, r_2, r_3, r_4) \in \mathbb{Q}^4_{\ge 0}$:

$$
\begin{aligned}
r_1 &= \max\big(\text{safeRatio}(\text{conc}, \text{conc}_{\text{cap}}),\\
&\qquad (s_0 + \lambda B_{\text{crit}})/B_{\text{crit}}\big),\\
r_2 &= \text{safeRatio}(\rho, \rho^*), \quad r_3 = \text{safeRatio}(\eta_{\text{lag}}, \eta^*_{\text{lag}}),\\
r_4 &= \max\Big(\text{safeRatio}(h, h^*),\ e_{0,\text{ref}} + \phi_{\text{ref}} r_{\text{end}},\\
&\qquad \tfrac{((1+\rho_{\text{run}})^2 - 1)(1-\varepsilon_{\text{peg}})}{\varepsilon_{\text{peg}}}\Big),
\end{aligned}
$$

and the ShockGuard minimax risk and Shock Survival Score:

$$
\begin{aligned}
R_\infty(r) &:= \|r\|_\infty = \max_{k \in \{1,2,3,4\}} r_k,\\
\text{Score}(r) &:= 100\big(1 - \min(1, \|r\|_\infty)\big).
\end{aligned}
$$

Theorem 4 (Minimax SIPS$_\infty$ & True L_p Impossibility — Th. 4A–4G)**.**

1. ***Global Uniqueness (Th. 4A–4B):*** *Among normalized risk scores $S : [0,1]^4 \to [0,1] \cap \mathbb{Q}$, $S(r) \le \tau \iff \forall k, r_k \le \tau$ holds across all thresholds $\tau \in [0,1]$ if and only if $S(r) = \|r\|_\infty$ identically.*

2. ***Universal Impossibility of True Weighted L_p Power Means (Th. 4C–4E):*** *Let $M_{p,w}(r) := (\sum_{j=1}^K w_j r_j^p)^{1/p}$ ($M_{p,w}(r) \le \tau \iff \sum w_j r_j^p \le \tau^p$ for $\tau \ge 0$). For every dimension $K \ge 2$, every finite integer $p \ge 1$, and every affine-sum-1 weight vector w, there exists a single-pillar catastrophe e_i ($\|e_i\|_\infty = 1$) satisfying*

$$
M_{p,w}(e_i)^p \le \frac{1}{K} < (\tau^*)^p, \quad \tau^* := 1 - \frac{K-1}{2Kp} \in (0,1). \tag{8}
$$

For symmetric $K = 4$ weights ($w_k = 1/4$), every single-pillar crisis slips past any threshold $\tau \in (4^{-1/p}, 1)$, whereas SIPS$_\infty$ vetoes 100% of them. In

9-*dimension frameworks [15], linear* L_1 *caps a total single-pillar failure at* $1/9 \approx 11.1\%$ *risk (*88.9/100 *deceptive safety!).*

3. ***End-to-End Coupled Safety Certificate (Th. 4G):*** *Whenever a well-posed protocol bundle satisfies* $\|r(b)\|_\infty \le 1$, ***all six invariants hold simultaneously****: (i)* $r(b) \in [0,1]^4$*, (ii)* $\text{conc} \le \text{conc}_{\text{cap}}$*, (iii) zero liquidation bad debt (Th. 1), (iv) infinite-horizon cascade* $\forall t, d_t \le B_{\text{crit}}$ *(Th. 2), (v) par backing* $B(h) \ge 1$ *and non-collapsing reflexivity* $\forall t, e_t \le 1$ *(Th. 3), and (vi) secondary CPMM peg floor* $P_{\text{post}} \ge 1 - \varepsilon_{\text{peg}}$ *(Th. 3).*

3. Layer 2 Extensions and Empirical Validation (Layer 3)

Layer 2 Analytical Extensions. While Layer 1 verifies exact CPMM algebra over $\mathbb{Q} \hookrightarrow \mathbb{R}$, real DeFi liquidity also spans Uniswap v3 concentrated ticks and Curve StableSwap pools ($An^n \sum x_i + D = ADn^n + D^{n+1}/(n^n \prod x_i)$). Locally within an active tick or near-par StableSwap band, effective unilateral virtual reserve $X_{\text{eff}} = L/\sqrt{P}$ (Uniswap v3) or $X_{\text{eff}} \approx (An^n + 1)X_0$ (Curve) maps directly into Theorems 1–3; outside the active tick or once Curve imbalance exhausts amplification, virtual depth collapses sharply to CPMM tail behavior—making our Theorem 1–3 CPMM bounds conservative tail guardrails. Similarly, stochastic oracle lag $\eta_{\text{lag}}(\Delta t)$ under heavy-tail jump-diffusion drawdowns calibrates the extreme quantile plugged into Pillar 3.

3.1. Experiment 4: Monte Carlo True L_p Blindness ($N = 2{,}500$)

We evaluate 2,000 single-pillar catastrophic protocols (one pillar $\in [0.95, 1.0]$, three pillars $\in [0.02, 0.25]$) and 500 genuinely safe protocols (all four pillars $\in [0.05, 0.55]$) at institutional threshold $\tau = 0.65$ (Figure 4). True homogeneous power means L_1, L_2, and L_3 miss 100.0% of catastrophic single-pillar failures because $4^{-1/3} = 0.630 < 0.65$, matching Theorem 4 verbatim; ShockGuard SIPS_∞ achieves 0.0% missed crises and 0.0% false vetoes.

3.2. Experiment 5: Nine Historical Crises & Nine Live DeFiLlama Protocols

We benchmark ShockGuard across 9 real historical DeFi stress episodes (7 collapses + 2 battle-tested October 10, 2025 survivals) and 9 live protocols enriched via the public DeFiLlama API on November 27, 2026 (9/9 live fetches succeeded; Table 1 and Figure 5). ShockGuard vetoes 7/7 (100%) real historical collapses with Score = 0.0 while passing 2/2 (100%) real October 2025 survivals (on-chain Ethena USDe at 44.4/100 and Kamino Lend at 28.0/100). On the March 2023 SVB USDC depeg, naive linear averaging yields a misleading 45.6/100 score despite total Pillar 4 failure, whereas SIPS_∞ fires an immediate RED HARD VETO. Live Nov. 2026 telemetry places blue-chip overcollateralized/Treasury-backed systems in GREEN (Circle USDC 64.0, Compound v3 56.4, Tether USDT 56.0, Aave v3 Core 54.0, Sky USDS/DAI 52.0) and high-LLTV or yield-loop structures in AMBER (Kamino 41.2, Morpho 91.5% LLTV 37.8, Ethena USDe 32.1, Pendle PT loops 15.6).

3.3. Experiment 6: Lean 4.36 Kernel & Adversarial Harness Verification

Full kernel compilation and axiom auditing via `lean/verify_lean.sh` completes in a mean of 4.44 s on Apple Silicon (M4 Max), checking all 32/32 public theorems with 0 warnings and strictly standard axioms. Against 8 adversarial proof-smuggling payloads—including `sorry` injection, `Char` quote smuggling (`def c1 : Char := '"' ... sorry`), `native_decide`, custom `axiom`, multiline `theorem\n name`, Unicode identifiers, `elab` output forging, and unused-variable warnings—our hardened audit harness blocks 8/8 (100%) attacks (Figure 6).

4. Conclusion and Open-Source Availability

ShockGuard closes the formal-economic verification gap in DeFi risk management by coupling four machine-checked Lean 4.36 theorems—seize-all vs. fixed-spread liquidation solvency, liquidation-spiral fixed points with dual Uniswap v2 vs. skimmed CPMM microstructure, three-regime stablecoin peg stability, and non-compensatory minimax SIPS_∞ uniqueness—with an instant 5-second web stress tester (defi.wallstreetdebug.com) and open-source MIT Python toolkit (`shockguard`). Protocol risk committees, DAO curators, and institutional desks can verify every solvency ceiling directly against Lean 4 kernel receipts while screening live protocols via free public on-chain APIs.

References

[1] Lewis Gudgeon, Sam M. Werner, Daniel Perez, and William J. Knottenbelt. Defi protocols for loanable funds: Interest rates, liquidity and market efficiency, 2020. URL https://arxiv.org/abs/2006.13922.

[2] Kaihua Qin, Liyi Zhou, Pablo Gamito, Philipp Jovanovic, and Arthur Gervais. An empirical study of defi liquidations: Incentives, risks, and instabilities, 2021. URL https://arxiv.org/abs/2106.06389.

[3] Liyi Zhou, Xihan Xiong, Jens Ernstberger, Stefanos Chaliasos, Zhipeng Wang, Ye Wang, Kaihua Qin, Roger Wattenhofer, Dawn Song, and Arthur Gervais. Sok: Decentralized finance (defi) attacks, 2022. URL https://arxiv.org/abs/2208.13035.

[4] Tarun Chitra. A curationary tale: Logarithmic regret in defi lending via dynamic pricing, 2025. URL https://arxiv.org/abs/2503.18237.

[5] Erum Iftikhar, Wei Wei, and John Cartlidge. Automated risk management mechanisms in defi lending protocols: A crosschain comparative analysis of aave and compound, 2025. URL https://arxiv.org/abs/2506.12855.

[6] Hasret Ozan Sevim and Christof Ferreira Torres. Signals and spoils: Speculative oracle extractable value in the era of cross-chain interoperability, 2026. URL https://arxiv.org/abs/2606.03434.



Figure 4. Theorem 4A–4E validation. **Panel A:** At $\tau = 0.65$ ($N = 2{,}000$ single-pillar crises, 500 safe protocols), L_1, L_2, and L_3 miss 100.0% of crises while SIPS$_\infty$ achieves 0.0% false-safe and 0.0% false-veto. **Panel B:** Exact blind-spot ceiling $K^{-1/p}$ for $K = 4$ and $K = 9$ [15] alongside our constructive Lean 4 Bernoulli witness $\tau^*(4, p)$.

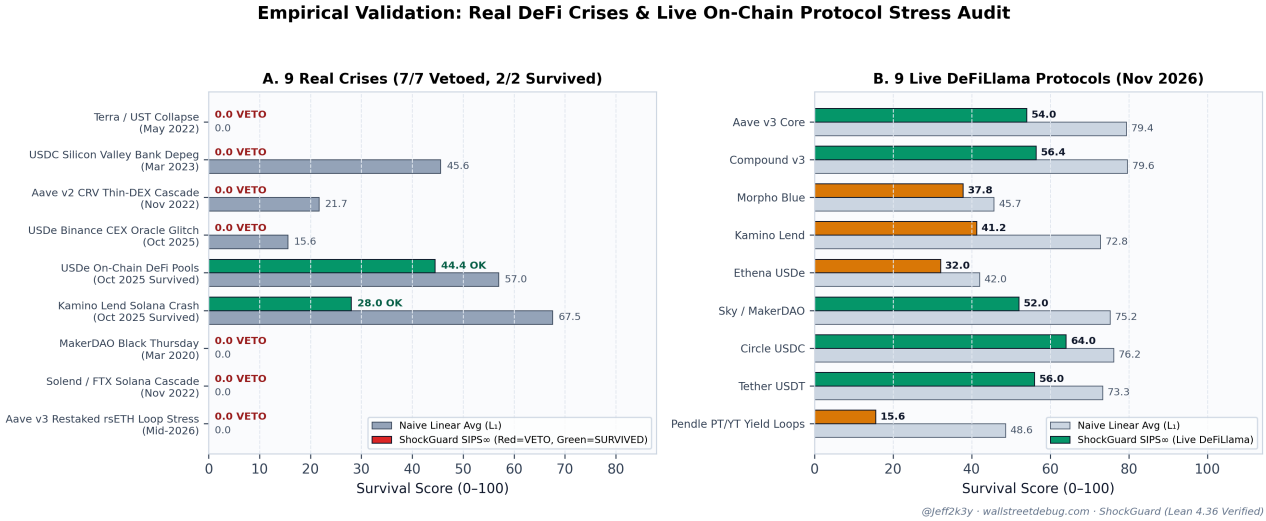

Figure 5. Empirical validation on real on-chain data. **Panel A:** ShockGuard SIPS$_\infty$ vetoes 7/7 historical collapses (0.0/100) and passes 2/2 Oct. 2025 survivals, exposing up to +45.6-point linear L_1 optimism gaps. **Panel B:** Live Nov. 2026 audit of 9 DeFi protocols enriched from the public DeFiLlama API.

Table 1. Empirical separation across 9 real historical DeFi crises (7 collapses, 2 Oct. 2025 survivals): ShockGuard SIPS$_\infty$ vs. naive linear average (L_1).

Historical Episode	SIPS$_\infty$	Lin L_1	Light	Outcome
Terra / UST (May 2022)	**0.0**	0.0	RED	Collapsed
USDC SVB (Mar 2023)	**0.0**	45.6	RED	Depeg \$0.88
Aave v2 CRV (Nov 2022)	**0.0**	21.7	RED	Bad Debt
USDe Binance (Oct 2025)	**0.0**	15.6	RED	CEX \$0.65
Maker Black Thu (2020)	**0.0**	0.0	RED	\$8.3M BD
Solend / FTX (Nov 2022)	**0.0**	0.0	RED	Bad Debt
Aave rsETH Loop (2026)	**0.0**	0.0	RED	Write-down
USDe On-Chain (Oct 2025)	**44.4**	57.0	AMBER	**Survived**
Kamino SOL (Oct 2025)	**28.0**	67.5	AMBER	**Survived**

[7] Agathe Sadeghi and Zachary Feinstein. Liquidation dynamics in defi and the role of transaction fees, 2026. URL https://arxiv.org/abs/2602.12104.

[8] Matthew Lorig. Optimal control of the ethena yield-bearing stablecoin, 2026. URL https://arxiv.org/abs/2605.11263.

[9] Krongtum Sankaewtong, Stefan Kitzler, Bernhard Haslhofer, and Yuichi Ikeda. Tracing stablecoin contagion during the usdc depeg after the silicon valley bank collapse, 2026. URL https://arxiv.org/abs/2606.07442.

[10] Ahmad Qadi, Hui Gong, and Francesca Medda. Stablecoin runs as safe-asset stress signals: Episode-level evidence for central bank surveillance, 2026. URL https://arxiv.org/abs/2611.26180.

[11] Sean Wan, Dongping Liu, and Luyao Zhang. Stableeval arena: A cost-aware agentic benchmark for stablecoin price stability prediction, 2026. URL https://arxiv.org/abs/2609.18949.

[12] Marco Dessalvi, Massimo Bartoletti, and Alberto Lluch-Lafuente. A formal approach to amm fee mechanisms with lean 4, 2026. URL https://arxiv.org/abs/2602.00101.

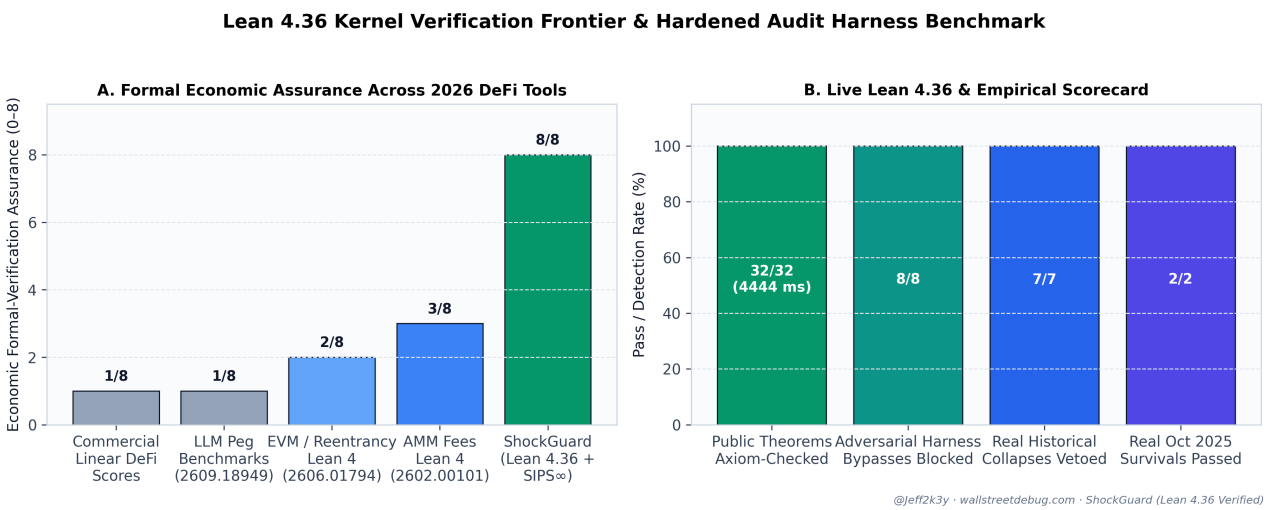

Figure 6. **Panel A:** Economic formal-verification assurance (0–8 criteria) comparing ShockGuard against commercial linear scores, LLM peg benchmarks [11], and smart-contract/AMM Lean 4 formalizations [12, 13]. **Panel B:** Live Lean 4.36 verification (32/32 theorems, 0 warnings), adversarial bypass blocking (8/8), and historical crisis separation (7/7 vetoed, 2/2 passed).

[13] Ray Iskander. Tridirectional discriminating-power formal verification of smart contract reentrancy defense against production-deployed solidity source, 2026. URL https://arxiv.org/abs/2606.01794.

[14] Seulkee Baek. Verifying deployed bytecode against executable evm semantics, with exact gas and no added axioms, 2026. URL https://arxiv.org/abs/2610.19764.

[15] Eva Oberholzer and Valeriy Zamaraiev. Toward a risk assessment framework for institutional defi: A nine-dimension approach, 2026. URL https://arxiv.org/abs/2605.05145.

[16] Fabien Pierret. Evolveguard: Genealogical selection bias and non-compensatory minimax control in evolutionary quantitative research, 2026. URL https://zenodo.org/records/24608677.

[17] Fabien Pierret. Proof-ci: Necessary and sufficient conditions for formally verified minimax evaluators preventing drift and reward hacking in recursive self-improvement, 2026. URL https://zenodo.org/records/24614087.