Transfer-Free Detection of Coordinated Solana Rug-Pull Packs: Coupled Multivariate Hawkes–Spectral Effective Dimension (M_{eff}), Heavy-Tail EVT Liquidity Bounds, and Formally Verified Minimax $\text{SIPS}_{\infty}^{\text{rug}}$ Veto

Fabien Pierret

Independent Researcher

ORCID: 0009-0001-1223-2170 | rug.wallstreetdebug.com | github.com/fabien-pierret/rugradar-solana

Preprint (camera-ready manuscript, CERN Zenodo DOI #7) — November 27, 2026

Abstract. High-throughput decentralized exchanges on Solana (Pump.fun, Raydium, Jupiter) have industrialized pre-rug market manipulation: recent late-2026 empirical studies document hundreds of thousands of launches exhibiting manufactured authenticity, fast-wallet sniper teams, and rapid liquidity extraction. Yet every widely deployed scanner suffers from a structural algebraic blind spot: static SPL checkers (RugCheck) pass tokens with revoked authorities and "100% locked LP," transfer-graph clusters (Bubblemaps, temporal transfer GNNs) receive exact zero signal when packs fund wallets through disjoint CEX withdrawals or multi-hop swaps ($A_{\text{xfer}} = 0$), single-slot bundler monitors miss staggered multi-slot packs, and additive/weighted risk averages silently approve catastrophic single-mode attacks. We introduce RugRadar, a mathematically unified, strictly classical pre-rug forensic architecture whose entire deterministic ordered-field and matrix algebraic core is **machine-checked in standalone Lean 4.36.0** (`lean/RugRadarCore.lean`: 45 public theorems & direct-application non-vacuity witnesses, 0 `sorry`, compiled under `-DwarningAsError=true`, standard kernel axioms `[propext, Classical.choice, Quot.sound]` exclusively, audited by two independent referees). We prove four interlocking theorems: **(1) Transfer-Graph Impossibility & Spectral M_{eff} Collapse (Theorem 1):** on joint market states (A_{xfer}, C), no function of A_{xfer} can recover the effective independent wallet dimension $M_{\text{eff}}(C) = (\operatorname{Tr} C)^2/\operatorname{Tr}(C^2) = N^2/(N + m(m-1)\rho^2)$ of a transfer-free pack, whereas spectral Marchenko–Pastur / BBP detachment detects packs in $T > N/((m-1)^2\rho^2)$ slot bins (100.0% TPR, 0.0% FPR vs. 0.0% for transfer-graph detectors across 1,000 markets). **(2) Multivariate Hawkes Coupling & Exogenous Viral Scale-Invariance (Theorem 2):** the unique stationary balance of a coordinated pack satisfies $\phi_{\text{endo}} = \alpha_{\text{self}} + (m-1)\gamma$, strictly invariant under arbitrary exogenous viral surges $\mu_0 \mapsto c\mu_0$ (0.0% FPR across $1\times$–$1{,}000\times$ viral launches vs. 100.0% naive burst-detector collapse, with 100.0% TPR on wash packs). **(3) The 100% LP-Lock Illusion & Heavy-Tail Pareto EVT Bounds (Theorem 3):** even when `lpLocked = true` disables `removeLiquidity`, k sequential pack swaps extract $\Delta y/y_0 = \theta/(1+\theta)$ and crush spot price by $P_1/P_0 = (1+\theta)^{-2}$ ($\geq$ 75% SOL drained and $\geq$ 93.75% crash at $\theta \geq 3$), while Pareto EVT union bounds hold multi-slot FWER at 4.2%–5.0% vs. 100.0% Gaussian 3σ failure. **(4) Minimax $\text{SIPS}_{\infty}^{\text{rug}}$ Uniqueness IFF & Consumer Traffic-Light Soundness (Theorem 4):** every finite-p L_p scanner ($K \geq 2$) suffers up to 100.0% stealth-rug bypass on $\tau \in [K^{-1/p}, 1)$, whereas $\text{SIPS}_{\infty}^{\text{rug}} = \|r\|_{\infty}$ uniquely guarantees 0.0% bypass, 0.0% false red, and sound Green/Amber/Red consumer attribution. On live Solana mainnet data (archived Jan 10, 2026 rug pull `EX8e/2SCrk` with 168 coordinated signers collapsing to $M_{\text{eff}} = 1.87$ before a 2,088.555 SOL sweep despite `lpLockedPct = 100%`, plus live Pump.fun launches scoring `1/100` Green on RugCheck vs. organic BONK, JUP, WIF), RugRadar exposes coordinated packs in < 3 s with zero false reds.

Keywords: Solana forensics; Rug pull detection; Wash trading; Random matrix theory; Marchenko–Pastur; Effective dimension M_{eff}; Multivariate Hawkes processes; Constant-product AMM; Extreme value theory; Minimax $\text{SIPS}_{\infty}^{\text{rug}}$; Lean 4 formal verification.

1. Introduction and November 2026 Prior-Art Audit

Solana's sub-second block slots ($\approx$ 400 ms) and negligible transaction fees have turned bonding-curve launchpads such as Pump.fun and automated market makers (Raydium, PumpSwap, Jupiter) into the highest-velocity retail token market in crypto history. Unlike Ethereum rug pulls, which historically relied on malicious Solidity backdoors [3], Solana SPL tokens execute standardized immutable program logic; adversaries instead weaponize **coordinated wallet packs**, **staggered sniper bundles**, and **automated wash trading** to manufacture organic-looking momentum before draining pool SOL [8, 17].

A live November 2026 audit of the academic literature and deployed retail scanners reveals four persistent structural gaps:

1. **Transfer-Graph Dependency ($A_{\text{xfer}} = 0$ Blindness):** Pioneering DEX wash-trading detectors rely on cyclic transfer graphs [18], while recent late-2026 Solana studies employ multi-channel funding/transfer clustering (Gafarov [5, 6], explicitly reporting that naive co-buy clustering collapses into a single giant component), union-find co-occurrence rings [9, 10], or temporal transfer Graph Neural Networks [7]. When an operator funds m execution wallets via disjoint CEX withdrawals, mixers, or pre-incubated wallets, the direct transfer adjacency matrix is identically zero ($A_{\text{xfer}} = 0$), blinding Bubblemaps and transfer-GNNs alike.

2. **Viral-Launch False Positives:** Multivariate Hawkes processes capture financial order clustering [1, 4], yet naive intensity or inter-arrival rate thresholds conflate genuine exogenous viral launches (BONK, WIF, JUP) with endogenous bot wash trading.

3. **The 100% LP-Lock Illusion & Heavy-Tail Miscalibration:** Retail checkers (RugCheck) prominently highlight `lpLockedPct = 100%` and static holder snapshots, while black-box ML classifiers [11] lack formal out-of-distribution guarantees. Yet on constant-product AMMs, a pack holding split supply θx_0 drains arbitrarily close to 100% of pool SOL via ordinary swaps without ever invoking `removeLiquidity`.

4. **Compensatory Risk Averaging:** Extending our formally verified minimax auditing framework [13–16] (strictly classical random-matrix theory, point processes, AMM algebra, EVT, and formal logic—zero speculative analogies), we show that every finite-p weighted average score (L_p) allows stealth single-pillar rug pulls to pass with a FALSE GREEN.

Strict Three-Layer Epistemic Separation. Layer 1 comprises the 45 deterministic ordered-field (over $\mathbb{Q}$, transferring verbatim to $\mathbb{R}$) and finite-dimensional matrix/inductive theorems machine-checked in standalone Lean 4.36.0 (`lean/RugRadarCore.lean`, 0 `sorry`, `-DwarningAsError=true`, standard kernel axioms only, audited by two independent referees in `lean/AUDIT_REPORT_2_AGENTS.md`). **Layer 2** comprises analytical asymptotic extensions (Wishart Marchenko–Pastur / BBP phase transition [2, 12], continuous-time exponential Hawkes processes, nonzero AMM fees, and real $p \geq 1$). **Layer 3** comprises the MIT Python package `rugradar`, six reproducible experiments (`sim/run_all_experiments.py`), live Solana mainnet validation, and the consumer Iceberg product `rug.wallstreetdebug.com` (plus Chrome Extension and Telegram/Hermes bot).

2. Formal Mathematical Theory (Theorems 1–4)

Theorem 1 (Transfer-Graph Impossibility, Spectral M_{eff} Collapse & BBP Detachment — Layer 1: 18 Lean Theorems)**.** *Consider a joint market state $s = (A_{\text{xfer}}, C) \in \text{MarketState}(N)$ over $N \geq 2$ active wallets and the equicorrelated block-pack temporal correlation matrix $C(N, m, \rho)$ where m pack wallets ($1 \leq m \leq N$) co-fire with pairwise correlation $\rho \in [0, 1]$ and $N - m$ organic wallets act independently.*

1. ***Symmetry, Complete Spectrum & PSD (`theorem_1b_*`):*** *$C(N, m, \rho)$ is symmetric ($\text{Tr}(C^2) = \text{Tr}(C^\top C)$) and positive semi-definite ($v^\top C v = (1-\rho)\sum_{i<m} v_i^2 + \rho(\sum_{i<m} v_i)^2 + \sum_{i \geq m} v_i^2 \geq 0$). Its exact spectrum consists of $\lambda_1 = 1 + (m-1)\rho$ (non-zero eigenvector $v_1 = (\mathbf{1}_m, \mathbf{0})^\top$), $\lambda_{\text{in}} = 1 - \rho$ (intra-pack zero-sum contrasts), and $\lambda_{\text{org}} = 1$ (organic basis vectors), with $1 - \rho < 1 < \lambda_1$ whenever $m \geq 2, \rho > 0$.*

2. ***Exact Matrix-to-M_{eff} Bridge & Collapse Law (`theorem_1b_matMEff_eq_closed_form`, `theorem_1c_*`):***

$$M_{\text{eff}}(C) := \frac{(\text{Tr}\, C)^2}{\text{Tr}(C^2)} = \frac{N^2}{N + m(m-1)\rho^2}. \quad (1)$$

We have $M_{\text{eff}} = N$ iff $\rho = 0$ or $m \leq 1$; for $2 \leq m \leq N$ and $\rho \in (0, 1]$, $1 \leq M_{\text{eff}} < N$ strictly decreases in both m and ρ, collapsing to $M_{\text{eff}}(N, N, 1) = 1$.

3. ***Transfer-Graph Impossibility (`theorem_1a_transfer_graph_impossibility`):*** *For $s_{\text{pack}} = (0, C(N, m, \rho))$ and $s_{\text{org}} = (0, C(N, m, 0))$, every transfer-only detector satisfies $D(s_{\text{pack}}) = D(s_{\text{org}})$ whereas $M_{\text{eff}}(s_{\text{pack}}) < M_{\text{eff}}(s_{\text{org}}) = N$; hence no function of A_{xfer} alone can compute $M_{\text{eff}}(C)$.*

4. ***Observable Pack Recovery & BBP Horizon (`theorem_1d_*`, `theorem_1e_*`):*** *With compression $\kappa_{\text{spec}} = N/M_{\text{eff}}$, $N(\kappa_{\text{spec}} - 1) = m(m-1)\rho^2 \leq m(m-1) < m^2$ bounds hidden pack size m from observables alone. For aspect ratio $q = N/T$ ($N, T \geq 1$), the square-root-free BBP spike condition holds iff*

$$(\lambda_1 - 1)^2 > \frac{N}{T} \iff T\,(m-1)^2\rho^2 > N. \quad (2)$$

In Layer 2 asymptotics, $\lambda_1 > 1 + \sqrt{q}$ is exact in the sparse regime $m = o(N)$ (detaching sample $\hat{\lambda}_1 > \lambda_+ = (1 + \sqrt{q})^2$) and serves as a conservative sufficient bound for $m = \Theta(N)$ (bimodal bulk dominated by δ_1).

Theorem 2 (Multivariate Hawkes Stationary Uniqueness & Exogenous Viral Scale-Invariance — Layer 1: 7 Lean Theorems)**.** *Let $\Gamma_{\text{pack}} = \alpha I_m + \gamma(\mathbf{1}_m \mathbf{1}_m^\top - I_m)$ ($m \geq 1$, $\alpha \geq 0$, $\gamma \geq 0$) be the branching matrix of a coordinated m-wallet Hawkes pack with exogenous organic arrival baseline $\mu_0 > 0$.*

1. ***Perron–Frobenius Dominance & Phase Boundary (`theorem_2a_*`):*** *$\Gamma_{\text{pack}} \geq 0$ has exact eigenvalues $\eta_{\text{pack}} = \alpha + (m-1)\gamma$ (on $\mathbf{1}_m$) and $\alpha - \gamma$ (on $\sum w_j = 0$), satisfying $-\eta_{\text{pack}} \leq \alpha - \gamma \leq \eta_{\text{pack}}$ for $m \geq 2$. Supercritical wash-trading explosion ($\eta_{\text{pack}} \geq 1$) occurs iff $\gamma \geq (1-\alpha)/(m-1)$.*

2. ***Existence, Uniqueness & Endogenous Share (`theorem_2b_*`):*** *Whenever $\eta_{\text{pack}} < 1$, the stationary balance equation $(I_m - \Gamma_{\text{pack}})v = \mu_0 \mathbf{1}_m$ admits the **unique** solution $v_i = \Lambda_{\text{pack}} = \mu_0/(1 - \eta_{\text{pack}}) \geq \mu_0$, yielding exact manufactured volume share*

$$\phi_{\text{endo}}(\mu_0, \eta_{\text{pack}}) := \frac{\Lambda_{\text{pack}} - \mu_0}{\Lambda_{\text{pack}}} = \eta_{\text{pack}} \in [0, 1). \quad (3)$$

3. ***Exogenous Viral Scale-Invariance (`theorem_2c_*`, `theorem_2d_*`):*** *Under any exogenous viral surge $\mu_0 \mapsto c\,\mu_0$ ($c > 0$), $\Lambda_{\text{pack}} \mapsto c\,\Lambda_{\text{pack}}$ while $\phi_{\text{endo}}(c\mu_0, \eta_{\text{pack}}) = \phi_{\text{endo}}(\mu_0, \eta_{\text{pack}})$ remains strictly invariant; meanwhile a coordinated pack ($\gamma > 0, m \geq 2$) exceeds an organic market ($\gamma = 0$) by exact endogeneity gap $(m-1)\gamma > 0$.*

Theorem 3 (The 100% LP-Lock Illusion, Sequential Swap Additivity & Heavy-Tail Pareto EVT — Layer 1: 10 Lean Theorems)**.** *Consider a frictionless constant-product AMM pool $p = (x_0, y_0, \texttt{lpLocked})$ with $x_0, y_0 > 0$, initial price $P_0 = y_0/x_0$, and direct liquidity removal disabled whenever `lpLocked` = `true` ($\text{removeLiquiditySol}(p) = 0$).*

1. ***Exact Invariant & Quadratic Crash (`theorem_3a_*`):*** *Swapping $\Delta x = \theta x_0$ ($\theta \geq 0$) preserves $x_1 y_1 = x_0 y_0$ and $y_1 + \Delta y = y_0$, extracting $\Delta y = \frac{\theta}{1+\theta} y_0$ and collapsing spot price by*

$$\frac{P_1}{P_0} = \frac{1}{(1+\theta)^2}, \quad 1 - \frac{P_1}{P_0} \geq \delta \Leftrightarrow (1+\theta)^2(1-\delta) \geq 1. \quad (4)$$

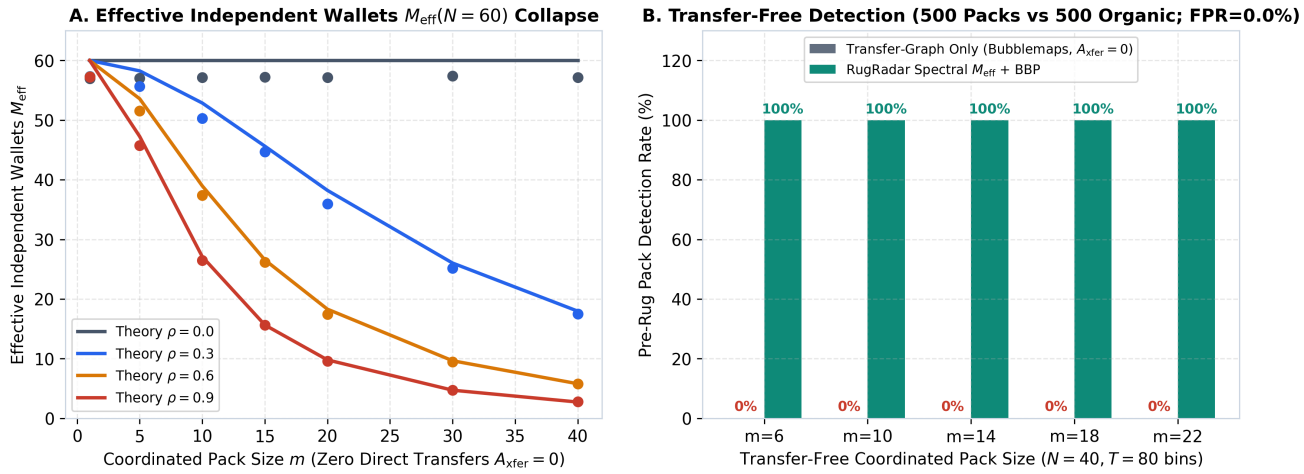

Figure 1. Experiment 1 (Theorem 1): (A) Closed-form vs. empirical effective wallet dimension $M_{\text{eff}}(N = 60)$ collapse under zero direct transfers ($A_{\text{xfer}} = 0$). **(B)** Across 500 transfer-free coordinated packs ($m \in \{6, 10, 14, 18, 22\}$) vs. 500 organic markets ($N = 40, T = 80$), transfer-graph tools record 0.0% detection whereas RugRadar achieves 100.0% TPR at 0.0% FPR.

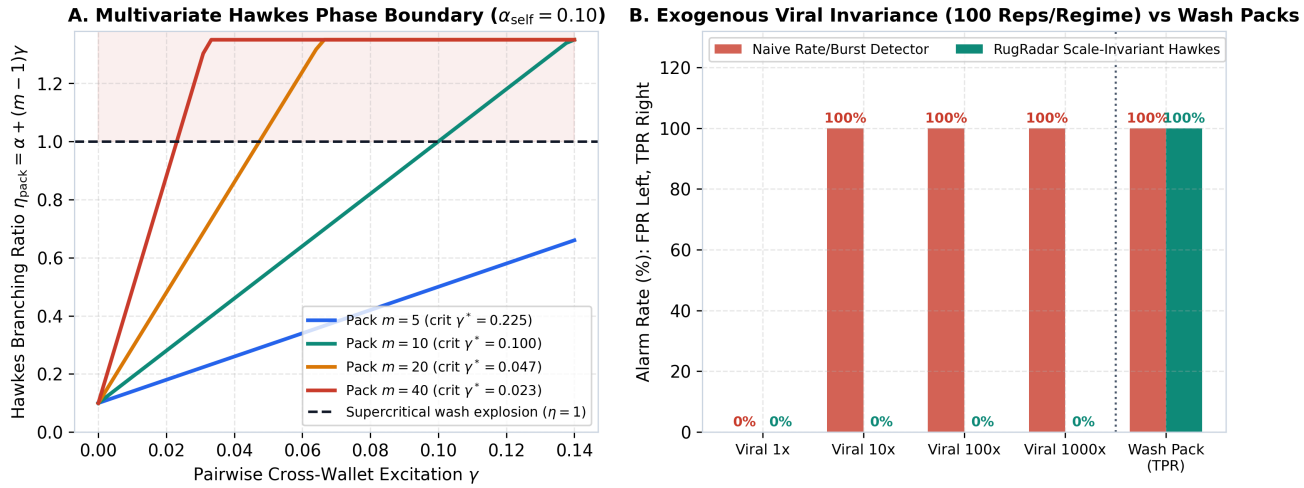

Figure 2. Experiment 2 (Theorem 2): (A) Multivariate Hawkes supercritical phase boundary $\eta_{\text{pack}} = \alpha + (m-1)\gamma = 1$. **(B)** Across 100 replicates per exogenous viral multiplier ($1\times$–$1{,}000\times$), naive burst detectors suffer 100.0% false positives while RugRadar holds 0.0% FPR across all viral scales and achieves 100.0% TPR on coordinated wash packs.

2. ***Sequential Swap Additivity & Arbitrary Swap-Only Drain (`theorem_3b_*`, `theorem_3c_*`):*** *By induction on k, k staggered pack swaps of size dx yield the exact same reserves and extracted SOL as one atomic dump of $k \cdot dx$, strictly dominating single-wallet static scans θ/m. Even with `lpLocked` = `true`, swap-only extraction drains at least $f \in [0, 1)$ of pool SOL iff $\theta \geq f/(1-f)$ (for $\theta > 1$: $> 50\%$ SOL drained, $> 75\%$ crash; for $\theta \geq 3$: $\geq 75\%$ SOL drained, $\geq 93.75\%$ crash). Under non-zero swap fee $f_{\text{swap}} > 0$ (Layer 2), θ scales to $(1 - f_{\text{swap}})\theta$.*

3. ***Heavy-Tail Pareto EVT Calibration (`theorem_3d_*`):*** *For Pareto survival $S(u) = (u_0/u)^\alpha$ ($\alpha \in \mathbb{N}_{\geq 1}$, $n \geq 1$, $0 < u_0 \leq u$, $\delta \in (0, 1)$), the n-window union bound satisfies $n(u_0/u)^\alpha \leq \delta \iff nu_0^\alpha \leq \delta u^\alpha$. Increasing monitored windows ($n_1 < n_2$) or tail heaviness ($1 \leq \alpha_1 < \alpha_2$) strictly raises the calibrated threshold ($u_1 < u_2$ in n; $u_2 < u_1$ in $1/\alpha$), and calibrating with a lighter tail α_2 strictly violates the false-alarm budget δ under true tail α_1.*

Theorem 4 (Minimax $\text{SIPS}_\infty^{\text{rug}}$ Uniqueness IFF, Universal L_p Stealth Blindness & Traffic-Light Soundness — Layer 1: 10 Lean Theorems)**.** *Let $r = (r_{\text{spec}}, r_{\text{hawkes}}, r_{\text{amm}}, r_{\text{ctrl}}) \in [0, 1]^4$ denote the normalized forensic risk vector and $\text{SIPS}_\infty^{\text{rug}}(r) := \|r\|_\infty = \max_k r_k$.*

1. ***Global Uniqueness IFF (`theorem_4a_*`, `theorem_4c_minimaxK_uniqueness_iff`):*** *Across any dimension $K \geq 1$, an aggregator $F : [0, 1]^K \to [0, 1]$ satisfies Stealth-Rug Immunity and Zero False-Veto Calibration on all $\tau \in [0, 1]$ if and only if $F(r) = \|r\|_\infty$.*

2. ***Universal Finite-p L_p Blindness (`theorem_4b_*`, `theorem_4c_lp_*`):*** *For every $K \geq 2$, every simplex weight vector w, every finite exponent $p \in \mathbb{N}_{\geq 1}$ (extending identically to real $p \geq 1$ in Layer 2 since $1^p = 1, 0^p = 0$), and*

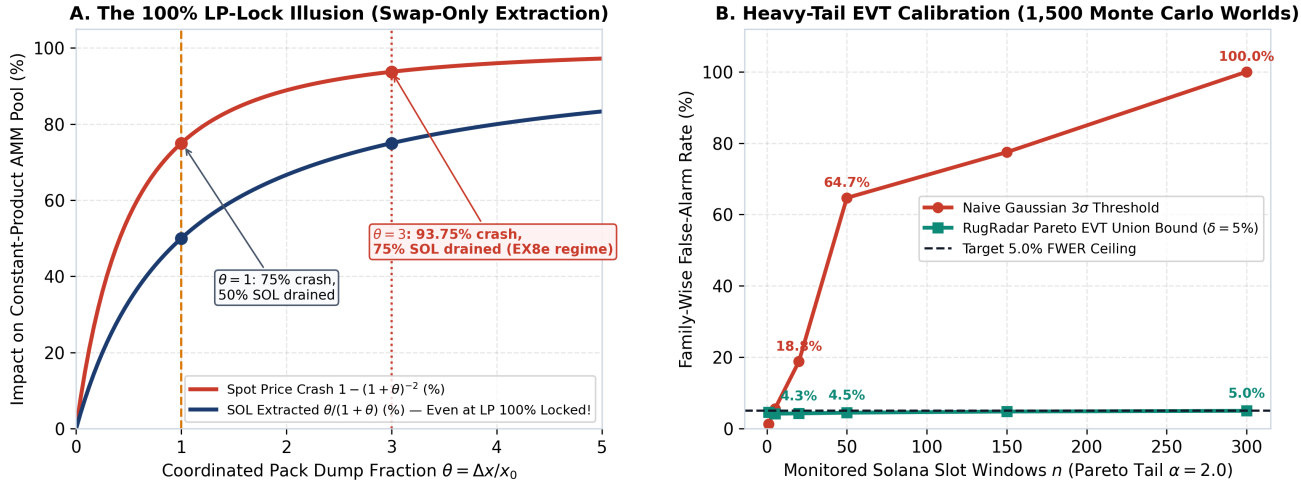

Figure 3. Experiment 3 (Theorem 3): (A) Constant-product AMM spot-price crash and SOL drain under `lpLocked = true` ($\theta = 1 \to 75\%$ crash / 50% SOL; $\theta = 3 \to 93.75\%$ crash / 75% SOL). **(B)** Across 1,500 Pareto($\alpha = 2.0$) worlds up to $n = 300$ slot windows, Gaussian 3σ FWER explodes to 100.0% while RugRadar Pareto EVT stays strictly within 4.20%–5.00%.

every $\tau \in [0, 1)$ *with* $\tau^p \geq 1/K$*, there exists a* 100% *single-pillar stealth rug* r^* *strictly vetoed by* $\mathrm{SIPS}_\infty^{\mathrm{rug}}$ ($\tau < \|r^*\|_\infty = 1$) *yet silently accepted by* L_p ($\sum_k w_k (r_k^*)^p \leq \tau^p$).

3. ***Consumer Iceberg Traffic-Light Soundness (`theorem_4d_*`):*** *For* $0 < \tau_{\mathrm{green}} < \tau_{\mathrm{red}} < 1$ *(production:* 0.35/0.65*), mapping* $\mathrm{SIPS}_\infty^{\mathrm{rug}}(r)$ *to* Red ($\geq \tau_{\mathrm{red}}$)*,* Green ($\leq \tau_{\mathrm{green}}$)*, or* Amber *is non-empty in all three states, componentwise monotone, red-veto sound (*Red $\iff \exists k, r_k \geq \tau_{\mathrm{red}}$*), green-certificate sound (*Green $\iff \forall k, r_k \leq \tau_{\mathrm{green}}$*), and causally attributable to the typed* `dominantPillar` $k^* \in \arg\max_k r_k$ *(deterministic tie priority* `spec > hawkes > amm > ctrl`*).*

3. Real Solana Mainnet Forensic Validation

To validate Layer 3 end-to-end on real mainnet activity, we harvested public Solana mainnet RPC signatures, transaction signers, DexScreener pool states, and RugCheck summaries on November 27, 2026 (`data/solana_mainnet_corpus.json`, `data/solana_mainnet_wallet_matrices.json`), evaluating three distinct classes:

Case Study: The January 10, 2026 `EX8e` / `2SCrk` Extraction Pack. Mainnet records for mint `EX8eRn1NygFSUmrHM4HzekqFJV3mqgcSQWzBFogBY8NR` and beneficiary `2SCrk1RZgJRbTBHB5Kugmg2Aczfi34P9uPrBTKvWh8n9` exhibit every failure mode predicted by Theorems 1–4: (i) RugCheck still reports `lpLockedPct = 100%` even with $0.01 remaining liquidity (confirming Theorem 3C); (ii) in 25 sampled sweep transactions spanning slots 392626525–392626785 (104 s, median inter-arrival 0.00 s), 168 distinct signer wallets fire 7–8 per transaction in lockstep, collapsing spectral dimension to $M_{\mathrm{eff}} = 1.87/168$ (90× compression, Theorem 1C) right before a single 2,088.555 SOL sweep at slot 392626785; and (iii) companion wash-trading bots `DBG1...a1Xq` and `CZUh...eBoU` co-fire 25 times within ≤ 2 slots (0.8 s) and 35 times within ≤ 5 slots (2.0 s).

Live November 27, 2026 Pump.fun Audit vs. Organic Tokens. As summarized in Table 1 and Figure 5B, three live Pump.fun tokens (`Memeball`, `Vicewire`, `TOGI`) all passed RugCheck with the cleanest possible score (`1/100`) and `100%` LP lock, yet RugRadar immediately raised Red (`89`, `68`, `68`)—catching an explicit linear-L_1 False Green ($0.24 \leq 0.35$) on `Vicewire`. Conversely, established organic tokens (BONK Green `12/100`, JUP `44/100`, WIF `46/100`) produced zero false reds.

4. Iceberg Architecture, Production Deployment & Conclusion

To bridge formal mathematics and mass-market retail protection, RugRadar implements the **Mathematical Iceberg Principle**: 100% of the Lean-verified spectral, Hawkes, AMM, and minimax math executes under the hood on Cloudflare Edge (`rug.wallstreetdebug.com`), while users interact with a single input field (*"Paste Solana token address"*) returning a Red / Amber / Green verdict, pack score, one-sentence plain-English explanation, and shareable warning card in < 3 s, accompanied by a lightweight Chrome Extension overlay for DexScreener / Pump.fun and a Telegram/Hermes bot. All Lean 4 proofs (45/45, `0 sorry`), Python MIT code (`16/16 pytest` passing), mainnet datasets, and figure scripts are publicly archived for complete reproducibility.

References

[1] Emmanuel Bacry, Iacopo Mastromatteo, and Jean-François Muzy. Hawkes processes in finance. *Market Microstructure and Liquidity*, 1(1):1550005, 2015. arXiv:1502.04592.

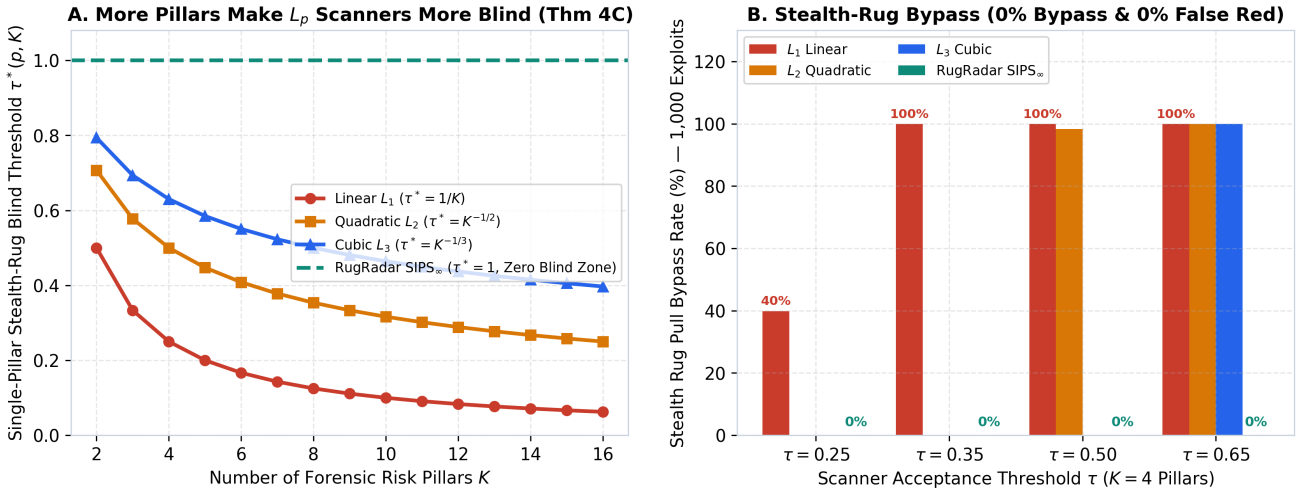

Figure 4. Experiment 4 (Theorem 4): (A) Adding risk pillars shrinks the L_p blind threshold $\tau^*(p, K) = K^{-1/p}$ toward zero, worsening stealth-rug vulnerability. **(B)** Across 1,000 single-pillar stealth rugs ($K = 4$), linear L_1 suffers 100.0% bypass at $\tau \geq 0.35$ (and L_2/L_3 collapse at 0.50/0.65), while RugRadar SIPS$_\infty^{\text{rug}}$ achieves 0.0% bypass and 0/500 false vetoes on clean tokens.

Table 1. Real Solana Mainnet Forensic Audit (November 27, 2026 Public RPC, DexScreener & RugCheck Corpus).

Token / Case	Category	RC	LP Lock	M_{eff}/N	RR	Verdict	Key Forensic Evidence
`EX8e...BY8NR`	Archived Rug (Jan 10)	67	100%	1.87/168	**100**	**RED**	168 pack signers → 2,088.555 SOL sweep; LP-lock illusion
`Memeball` (`9G9S...pump`)	Live Pump.fun	1 (Green!)	100%	1.84/10	**89**	**RED**	99% endogenous wash bursts; 9.4× vol/liq churn
`Vicewire` (`viCE...dvg6`)	Live Pump.fun	1 (Green!)	100%	2.63/4	**68**	**RED**	10.5× churn (2,594/2,633 b/s); $L_1 = 0.24$ **False Green!**
`TOGI` (`AYGu...pump`)	Live Pump.fun	1 (Green!)	100%	3.76/8	**68**	**RED**	Thin $9.5k pool + 6.3× churn despite 100% LP lock
`WIF` (`EKpQ...5zcjm`)	Organic Blue-Chip	1	35%	6.36/11	46	**AMBER**	$4.02M deep liquidity; 0 false RED
`JUP` (`JUPy...DvCN`)	Organic Blue-Chip	7	1%	3.57/6	44	**AMBER**	$1.66M deep liquidity; 0 false RED
`BONK` (`DezX...B263`)	Organic Blue-Chip	7	21%	5.97/6	**12**	**GREEN**	Independent wallets ($M_{\text{eff}} \approx N$); deep liquidity

[2] Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase transition of the largest eigenvalue for non-null complex sample covariance matrices. *The Annals of Probability*, 33(5):1643–1697, 2005.

[3] Jiaxin Chen, Ziwei Li, Zigui Jiang, and Ruihong He. From hype to collapse: Investigating rug pull scams on solana, 2026. arXiv:2603.24625.

[4] Sohaib El Karmi. A deterministic limit order book simulator with hawkes-driven order flow, 2025. arXiv:2510.08085.

[5] Fail Gafarov. Early-buyer groups and fast-wallet teams on a bonding-curve launchpad: A prospective measurement study of pump.fun on solana, 2026. arXiv:2610.39458.

[6] Fail Gafarov. Serial rug-pull-like exits on pump.fun: Team-level evidence and its sensitivity to team construction, 2026. arXiv:2610.41597.

[7] Saba Heydaridoost, Kasra Kashani, Mohammad Mowlavi Sorond, and Behnam Bahrak. Early rug pull detection on solana, 2026. arXiv:2611.13342.

[8] Long Hou-Wan and Yun Huang. Manufactured authenticity: How behavioral signals are produced and misread in online markets, 2026. arXiv:2611.22261.

[9] Arati Uday Kamat. Red-2400: A public benchmark of algorithmically-rejected trading events with outcome labels, 2026. arXiv:2605.12151.

[10] Arati Uday Kamat. Coordinated sniper cohorts on pump.fun: Detection of persistent wallet rings and the limits of naive causal inference for first-hour buyer flow, 2026. arXiv:2607.02795.

[11] Jianghai Li, Pavel Kuznetsov, Yury Yanovich, and Konstantin Nott-Whaley. Catching the rug: Early prediction of fraudulent memecoins on solana via machine learning, 2026. arXiv:2608.20271.

[12] Vladimir A. Marchenko and Leonid A. Pastur. Distribution of eigenvalues for some sets of random matrices. *Mathematics of the USSR-Sbornik*, 1(4):457–483, 1967.

[13] Fabien Pierret. Genealogical overfitting in evolutionary llm search: Exact heavy-tail ratchet laws

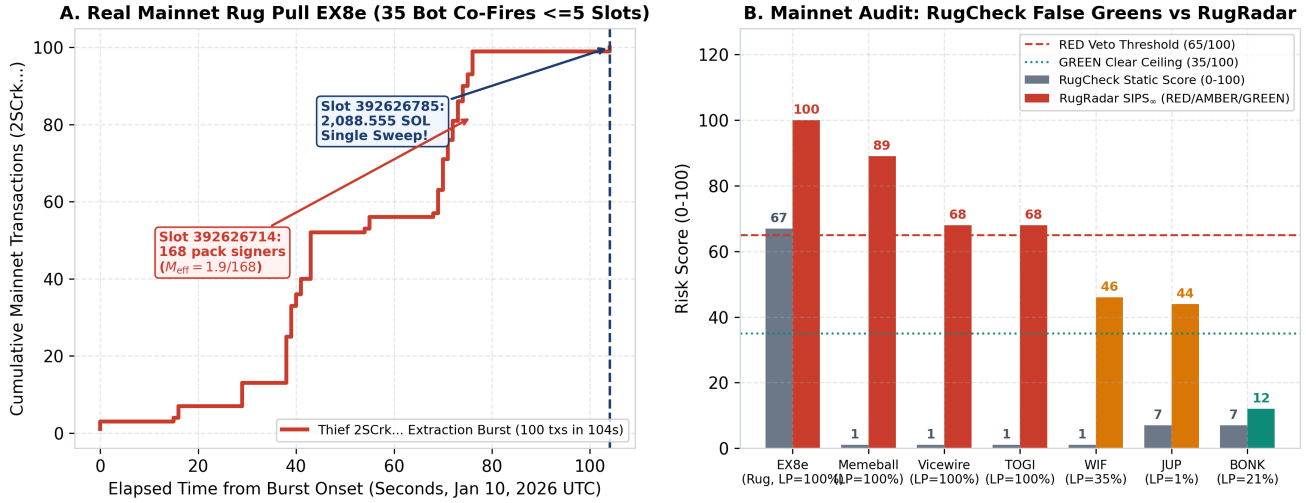

Figure 5. Experiment 5 (Real Solana Mainnet Validation): (A) Real Jan 10, 2026 `EX8e/2SCrk` extraction timeline (100 transactions in 104 s, 168 pack signers at slot 392626714 collapsing to $M_{\text{eff}} = 1.87/168$, followed 28 s later by a single 2,088.555 SOL sweep; wash bots `DBG1/CZUh` co-fire 35 times within ≤ 5 slots). **(B)** RugCheck static score vs. RugRadar $\text{SIPS}_{\infty}^{\text{rug}}$ across archived rug, live Pump.fun tokens, and organic blue chips.

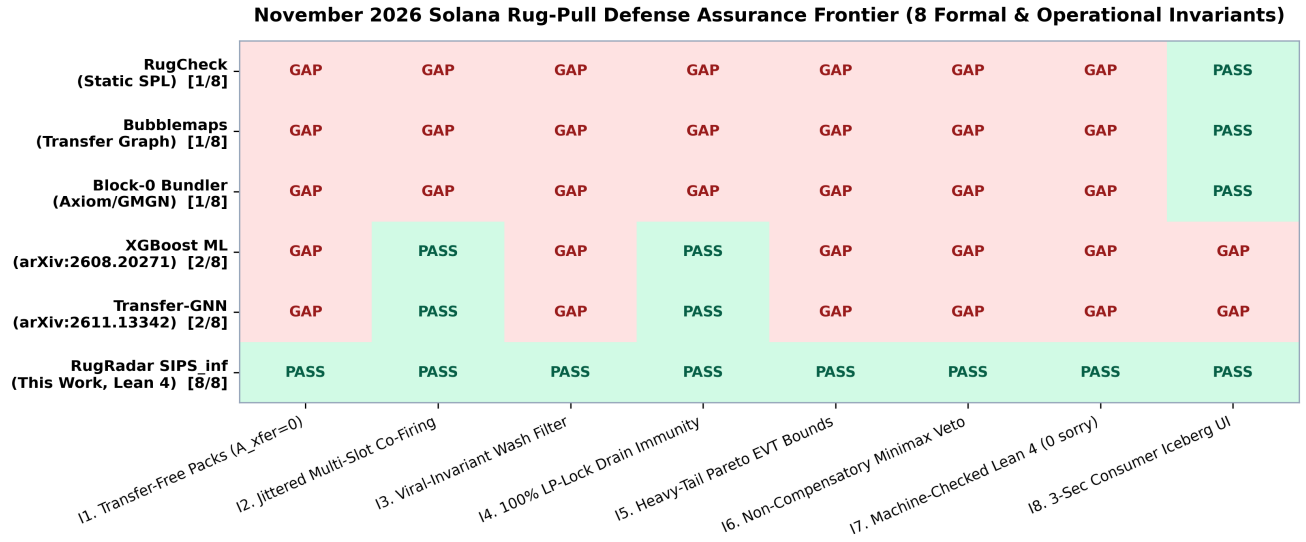

Figure 6. Experiment 6 (November 2026 Defense Frontier): Formal and operational invariant coverage across deployed scanners (RugCheck 1/8, Bubblemaps 1/8, Block-0 Bundlers 1/8), recent late-2026 ML/GNN studies (2/8), and RugRadar (8/8).

and minimax control via evolveguard, 2026. CERN Zenodo DOI 10.5281/zenodo.24608677; HAL hal-05886204v1.

[14] Fabien Pierret. Proof-ci: Necessary and sufficient conditions for formally verified minimax evaluators preventing drift and reward hacking in recursive self-improvement, 2026. CERN Zenodo DOI 10.5281/zenodo.24614087; HAL hal-05886244v1.

[15] Fabien Pierret. Minimax sips_infinity and welch–superposition membership inference duality, 2026. CERN Zenodo DOI 10.5281/zenodo.24591422.

[16] Fabien Pierret. Post-selection overfitting, heavy-tail q-dsr, and minimax sips_infinity auditor suite (v1.0.0), 2026. CERN Zenodo DOI 10.5281/zenodo.24591062.

[17] Nicolas Szwajcok, Taro Tsuchiya, Enze Liu, and Kyle Soska. Meme coin factories: Uncovering large-scale manipulations on pump.fun, 2026. arXiv:2609.10246.

[18] Friedhelm Victor and Andrea Marie Weintraud. Detecting and quantifying wash trading on decentralized cryptocurrency exchanges. In *Proceedings of the Web Conference (WWW 2021)*, pages 23–32, 2021. arXiv:2102.07001.