RugPull

Ethical Crypto through Randomized Unethical Behavior

Stella Nova Research Center

White Paper v2.0 - July 2026

RESEARCH OBJECT	Radically disclosed paired-token liquidity shocks
REFERENCE PRICE	USD 0.40 from the launch reserve ratio
REFERENCE YEAR	9 events; USD 2,346 modeled Stella Nova funding
PUBLIC SIGNAL	1.34x event-day probability signal lift
IMPLEMENTATION	Working PHP model and testnet-only Solidity protocol

Central proposition. Event timing can remain random while event risk becomes statistically predictable. The experiment therefore studies both the liquidity shock and the unequal ability of participants to interpret a public probability signal.

This document is a research specification, not an offer, solicitation, legal conclusion, or promise of profit. The included contracts reject mainnet deployment and use virtual market accounting.

Abstract

RugPull is a paired-token adversarial-market experiment developed by the Stella Nova Research Center. It asks whether a behavior normally associated with concealment and opportunistic extraction can become a legitimate object of study when the behavior is announced in advance, encoded in public rules, driven by independently verifiable randomness, and linked to a disclosed public-benefit funding stream.

Two tokens, RugPullA and RugPullB, begin with symmetric fixed supplies, symmetric liquidity, and an initial 50% protocol treasury allocation. Organic trading changes their relative prices and trading intensity. A public hazard function uses elapsed time, trading frequency, valuation, and A/B imbalance to calculate a daily event probability. A random draw then determines whether the probability becomes an event. When an event occurs, a bounded portion of the selected treasury position is sold into its virtual constant-product pool. A disclosed protocol percentage is recorded as Stella Nova funding, and the remaining virtual USD proceeds purchase the opposing token.

This release contains a working PHP simulator, a public probability display, participant-archetype diagnostics, no-database research comments, daily server snapshots, and testnet-only Solidity contracts. It does not contain a mainnet decentralized-exchange router or a real-value automated transaction signer. Real-value deployment remains contingent on legal, economic, security, custody, and institutional review.

1. Research proposition

The project is organized around a narrow question:

> Can radical transparency convert a concealed market exploit into informed participation without merely providing sophisticated actors with a better mechanism for extracting value from less sophisticated participants?

The project does not assume that disclosure makes harmful conduct ethical. Disclosure is instead an experimental variable. RugPull measures whether publishing the mechanism changes participant comprehension, willingness to accept risk, anticipatory trading, institutional legitimacy, and the distribution of gains and losses.

The central anthropological concern is moral economy. Participants may evaluate an identical financial loss differently depending on whether the beneficiary is an anonymous promoter, a commercial enterprise, a charitable organization, or an educational research institution. RugPull makes that institutional relationship visible rather than treating the beneficiary as an incidental feature.

2. Token structure and initial pricing

For each token i in {A, B}, the reference configuration uses:

- Total fixed supply S = 1,000,000 units.
- Initial virtual treasury inventory $Q_i(0) = 0.50S$.
- Initial token reserve $X_i(0) = 0.25S$.
- Initial virtual USD reserve $Y_i(0)$ = USD 100,000.
- Remaining token allocation = 0.25S outside the initial treasury and launch pool accounting.

The launch price is determined by the reserve ratio:

$$P_i(0) = Y_i(0) / X_i(0)$$

Under the reference configuration:

$$P_i(0) = 100{,}000 / 250{,}000 = \text{USD } 0.40$$

The fully diluted value is:

```
FDV_i(0) = P_i(0) x S = USD 400,000
```

If the 50% protocol treasury is excluded from public float, initial circulating market capitalization is:

```
Circulating market cap_i(0) = P_i(0) x 0.50S = USD 200,000
```

This pricing rule is auditable. A higher initial USD reserve raises the displayed price and increases the depth available to absorb a trade. A larger token reserve with unchanged USD lowers the price. Nominal price is therefore not an independent economic decision; it is a consequence of launch supply and liquidity.

3. The 50% ownership correction

The original concept proposed that a 50% holder could sell the entire 50% position, use the proceeds to buy the opposing token, and continue holding 50% of both assets. Those conditions cannot all remain true in a fixed-supply market.

After selling the complete RugPullA treasury, the seller owns zero RugPullA unless it later repurchases RugPullA. Purchasing RugPullB does not restore the RugPullA position. Enforcing a continuous 50% ownership ratio would require minting and burning, rebasing, a secondary reserve, protocol-owned liquidity counted as ownership, or repeated market reacquisition.

RugPull v2 therefore treats 50% as the initial treasury allocation. Every event reports the actual post-event treasury inventory. This preserves fixed supply and prevents the interface from representing a mathematically impossible ownership state.

4. Automated-market-maker model

Each virtual market follows the constant-product relation:

```
x y = k
```

Here x is the pool's token reserve, y is its virtual USD reserve, and k is the invariant before fees. A token purchase adds virtual USD and removes tokens. A token sale adds tokens and removes virtual USD.

For USD input dY after the pool fee:

```
Token out = x - k / (y + dY)
```

For token input dX after the pool fee:

```
USD out = y - k / (x + dX)
```

The curve is nonlinear. A large treasury transaction receives progressively worse execution as it moves through the reserve curve. The source sale and target purchase therefore both incur slippage. Liquidity depth, rather than displayed market capitalization alone, governs the severity of the event.

5. Public hazard function

RugPull separates predictable risk from unpredictable realization. The public state determines the daily hazard score. Randomness determines whether the disclosed probability becomes an event.

For token i on day t:

```
H_i,t = 0.35T_i,t + 0.25F_i,t + 0.25V_i,t + 0.15I_i,t
```

Where:

- T is normalized elapsed time since the token was last selected as the source.
- F is normalized recent trading activity relative to current pool depth.
- V is normalized valuation relative to the launch price.
- I is the relative A/B price imbalance favoring the candidate source token.

After the cooldown period, the daily probability is:

```
p_i,t = min[p_max, b(0.35 + H_i,t)]
```

Where b is the disclosed base event rate and p_max is a public maximum probability. An independently generated value u_i,t is sampled from a uniform distribution:

```
u_i,t ~ U(0,1)
```

An event is eligible when:

```
u_i,t < p_i,t
```

If both tokens trigger, the model selects the token with the greater probability-to-draw margin. This rule is public and deterministic once the random values are known.

6. Why random events become statistically predictable

Randomness does not eliminate learning. The exact event day may remain unknown, but the conditional likelihood becomes increasingly estimable. Participants can observe the cooldown, trading intensity, relative price, valuation, and published probability.

As the event history grows, market participants can estimate:

- The probability distribution of waiting times.
- The relationship between valuation and event risk.
- Whether one token tends to become the source after relative overperformance.
- The frequency with which liquidity or volume caps constrain requested transactions.
- The average recovery period after a source-token shock.
- The degree to which event risk is incorporated into prices before execution.

RugPull reports the Brier score of the public probabilities and an event signal-lift statistic. Signal lift compares the average published probability on actual event days with the average probability across all eligible token-days. A value above 1.0 indicates that events tend to occur when the visible model already reports above-average risk.

7. Event-size selection and safety caps

For a standard event, the requested treasury fraction beta_t is selected from a public interval:

```
beta_t ~ U(beta_min, beta_max)
```

The requested source quantity is:

```
D_t = beta_t Q_source,t
```

The executable quantity is limited by treasury inventory, source-pool depth, and trailing organic volume:

$$D_{exec,t} = \min[D_t, C_{liquidity,t} / P_{source,t}, C_{volume,t} / P_{source,t}]$$

Where:

$$C_{liquidity,t} = \alpha Y_{source,t}$$

$$C_{volume,t} = \gamma ADV_{30,t}$$

Alpha is the published pool-depth cap, gamma is the trailing-volume multiplier, and ADV_30 is modeled 30-day average daily volume. These limits prevent the normal research profile from converting every event into near-total pool depletion.

The simulator also includes an optional catastrophe condition. When enabled, the requested fraction can become 100% of the active treasury inventory. The same caps may still constrain execution unless the user deliberately relaxes them. This preserves the original full-pull hypothesis as a stress condition without treating it as the default live architecture.

8. Stella Nova funding mechanism

RugPull v2 separates protocol funding from treasury rotation. Gross virtual USD proceeds from the source sale are divided into a disclosed Stella Nova funding amount and the amount routed into the opposing token.

$$R_t = f E_t$$

Where:

- E_t is gross event value.
- f is the public Stella Nova funding rate.
- R_t is the recorded institutional funding amount.

The opposing-token purchase receives:

$$E_{route,t} = E_t - R_t$$

This design makes the institutional incentive visible. Any observer can calculate how much value is attributed to Stella Nova on every event. The model does not describe the full source-sale proceeds as institutional income.

In a real-value implementation, the legal and accounting treatment of R_t would require a formal determination. The simulation labels it protocol funding, not unrestricted revenue, donation income, investment return, or tax-exempt contribution.

9. Reference one-year simulation

The v2 reference run uses seed 240619 and the following conditions:

- 365 days.
- 1,000,000 units per token.
- 50% initial treasury allocations.
- 250,000 tokens and USD 100,000 per virtual pool.
- 60 average trades per token per day.
- USD 125 average trade size.
- 1.8% base daily event rate after a 14-day cooldown.
- Standard event requests between 8% and 22% of the active source treasury.

- 1.5% Stella Nova funding rate.
- 30% source-liquidity cap.
- 1.5 times trailing-average-volume cap.
- Catastrophe requests disabled.

This deterministic seed produces nine events. RugPullA ends at approximately USD 0.709 and RugPullB ends at approximately USD 0.433, compared with the USD 0.40 launch price. Maximum modeled drawdown is approximately 40.4% for A and 81.3% for B. Modeled organic volume is approximately USD 5.51 million. Disclosed Stella Nova protocol funding is approximately USD 2,346. The combined virtual treasury net asset value, including recorded funding, ends approximately 40.2% above its marked launch value.

These results are a path from one seed, not an expected return, forecast, or claim regarding a real market. Alternate seeds and parameter combinations can reverse the outcome, produce fewer events, or create substantially larger losses.

10. Participant-archetype diagnostics

The simulator follows three stylized USD 10,000 participant archetypes without market impact:

Passive 50/50 holder

The participant purchases equal USD values of RugPullA and RugPullB at launch and never trades. This participant is exposed to both long-term appreciation and both drawdown paths.

Hazard-aware rotator

The participant observes the public hazard scores, retains 20% cash, and allocates the remaining 80% to the lower-hazard token. Rebalancing incurs modeled friction. This archetype tests whether predictability creates an advantage even when the exact event remains random.

Post-event contrarian

The participant waits for an event, invests 60% of available cash into the shocked source token after the event, and exits after 14 days. This archetype tests the proposed lower-entry-point mechanism.

Under the reference seed, the passive archetype ends near USD 14,273, the hazard-aware diagnostic near USD 19,989, and the post-event diagnostic near USD 11,137. These values demonstrate path dependence and information asymmetry. They are not trading recommendations.

11. Expected market adaptation

The likely behavioral trajectory is staged.

Novelty phase

Early trading is dominated by narrative, publicity, launch liquidity, and Stella Nova's institutional credibility. Participants have insufficient observations to evaluate calibration.

Hazard-learning phase

Participants begin monitoring cooldown completion, relative valuation, and the daily probability display. Price movement may precede events as traders move away from the likely source asset or toward the likely target asset.

Strategic phase

Automated participants estimate the event process, liquidity providers alter exposure, and arbitrageurs react to reserve dislocations. Predictability can reduce surprise while increasing the advantage of participants able to model the system.

Partial-equilibrium phase

The market begins pricing event risk continuously. Post-event entry becomes an established behavior, and the probability function itself can influence the price and volume inputs that feed the next probability calculation. RugPull thus becomes reflexive: published risk changes behavior, and behavior changes published risk.

12. Research outputs

The system can produce evidence concerning:

- Whether participants understand a probability disclosure rather than merely noticing a warning label.
- How quickly event risk becomes incorporated into price.
- Whether public-benefit funding increases willingness to tolerate loss.
- Whether a transparent algorithm reduces or redistributes informational advantage.
- Which actor classes capture value from the source and target price movements.
- Whether sophisticated participants extract gains from participants who understand the branding but not the mathematics.
- Whether transparency changes perceptions of legitimacy without changing distributive harm.
- Whether the protocol's institutional funding becomes dependent on participant losses or on genuine transaction activity.

13. Technical architecture

The explanatory application is designed for PHP 8.2 or later on Hostinger shared hosting. It contains:

- A Stella Nova landing page.
- A deterministic simulator API.
- Price and public-probability visualizations.
- Event and funding ledgers.
- Participant-archetype diagnostics.
- No-database JSON comments with CSRF protection, file locking, input controls, a honeypot, and basic rate limiting.
- Daily cron snapshots.
- Password-protected backend status.

The Solidity package creates RugPullA Research and RugPullB Research on local Hardhat, Sepolia, or Base Sepolia. It uses virtual pool reserves, virtual treasury balances, a virtual Stella Nova funding ledger, and an asynchronous randomness-provider interface. The included mock randomness provider permits complete local testing. A reviewed implementation could replace it with a Chainlink VRF v2.5 adapter or another verifiable-randomness system.

Chainlink describes VRF as producing random values with an on-chain cryptographic proof, and its current documentation identifies VRF v2.5 as the supported generation. A production integration must follow the current network-specific coordinator, subscription, billing, confirmation, and callback requirements rather than relying on

remembered addresses or deprecated examples.

14. Legal and regulatory boundary

RugPull is not accurately described as operating in a market free from regulation. U.S. agencies assert different authorities over digital-asset conduct depending on the asset, activity, and facts.

The CFTC states that it has enforcement authority over fraud and manipulation in spot digital-commodity markets and has specifically warned about virtual-currency pump-and-dump activity. Advance disclosure is relevant to deception and participant knowledge, but it does not by itself answer whether planned price-impact conduct constitutes manipulation.

In February 2025, SEC Corporation Finance staff described a limited category of typical meme coins that may not involve an investment contract when they function principally as entertainment or cultural assets and do not create an expectation of profit derived from the managerial efforts of others. The staff statement is fact-specific, nonbinding, and expressly does not protect fraud or assets mislabeled to avoid securities law. RugPull's institutional funding purpose, managed treasury, published algorithm, and continuing Stella Nova role differ materially from a purely entertainment-oriented meme coin.

FinCEN guidance states that administrators and exchangers of convertible virtual currency may be money transmitters depending on their activities and applicable exemptions. Issuing tokens, exchanging value, transmitting value for others, accepting funds, or redeeming assets can change the analysis.

A real-value launch therefore requires fact-specific counsel concerning securities, commodities and manipulation, BSA/AML, state money-transmission law, consumer protection, tax, nonprofit restrictions, advertising, privacy, and institutional governance.

15. Live-value deployment gate

No real-value deployment should occur until the following are complete:

- Written U.S. crypto-market legal analysis.
- Commodity and manipulation analysis addressing the disclosed event mechanism.
- Securities analysis based on the actual marketing, purchaser expectations, treasury role, and funding purpose.
- FinCEN and state money-transmission analysis.
- Tax and nonprofit-purpose analysis for Stella Nova.
- Independent smart-contract and economic audits.
- Multisignature or institutional custody with role separation.
- Formal incident-response and emergency-pause procedures.
- Public terms, risk disclosures, fee disclosures, and event equations.
- Participant-comprehension testing that goes beyond a warning checkbox.
- Institutional authorization documenting the research purpose and permitted use of funds.
- A data-governance protocol for publishing anonymized research results.

16. Falsification and ethical failure conditions

RugPull should not be treated as validated merely because participants voluntarily acquire the tokens. The ethical proposition is weakened or falsified if:

- Participants systematically misunderstand the event probability.
- Marketing emphasizes potential gains more prominently than modeled losses.
- Stella Nova retains discretionary control over event timing or random selection.
- Material parameters can be changed without delay and public notice.
- Institutional revenue depends primarily on transferring pool reserves from uninformed participants.
- Automated actors capture most of the event-related value.
- The public-benefit narrative functions as moral licensing rather than informed consent.
- Research publication is withheld when results are unfavorable.

The experiment is most credible when it is structured to discover that its own ethical premise may be wrong.

17. Development sequence

Phase 1 - Explanatory release: landing page, white paper, simulator, no-database comments, and public algorithm.

Phase 2 - Testnet protocol: dual tokens, virtual reserves, asynchronous randomness, daily observations, event evidence hashes, and funding accounting.

Phase 3 - Research validation: Monte Carlo analysis, parameter sensitivity, participant-comprehension instruments, bot and liquidity-provider models, and independent economic review.

Phase 4 - Legal and institutional review: written opinions, custody design, governance, audit, tax analysis, and formal Stella Nova authorization.

Phase 5 - Limited controlled study, only if approved: restricted participants, capped exposure, research consent, immutable parameters, transparent data publication, and no unsupported return claims.

References

Chainlink. 2026. "Chainlink VRF." Official documentation. https://docs.chain.link/vrf

Chainlink. 2026. "Getting Started with Chainlink VRF V2.5." Official documentation. https://docs.chain.link/vrf/v2-5/getting-started

Commodity Futures Trading Commission. "Beware Virtual Currency Pump-and-Dump Schemes." Customer advisory. https://www.cftc.gov/LearnAndProtect/AdvisoriesAndArticles/beware_virtual_currency_pump_dump.html

Commodity Futures Trading Commission. 2026. "Digital Asset Frauds." https://www.cftc.gov/LearnAndProtect/digitalassetfrauds

Financial Crimes Enforcement Network. 2013. "Application of FinCEN's Regulations to Persons Administering, Exchanging, or Using Virtual Currencies." FIN-2013-G001. https://www.fincen.gov/resources/statutes-regulations/guidance/application-fincens-regulations-persons-administering

Financial Crimes Enforcement Network. 2019. "Application of FinCEN's Regulations to Certain Business Models Involving Convertible Virtual Currencies." FIN-2019-G001. https://www.fincen.gov/resources/statutes-regulations/guidance/application-fincens-regulations-certain-business-models

Securities and Exchange Commission, Division of Corporation Finance. 2025. "Staff Statement on Meme Coins." https://www.sec.gov/newsroom/speeches-statements/staff-statement-meme-coins

Securities and Exchange Commission, Commissioner Caroline A. Crenshaw. 2025. "Response to Staff Statement on Meme Coins: What Does it Meme?" https://www.sec.gov/newsroom/speeches-statements/crenshaw-response-staff-statement-meme-coins-022725

Reference model visualization

The reference chart combines token prices, published event probabilities, actual event days, and cumulative disclosed Stella Nova funding for seed 240619.

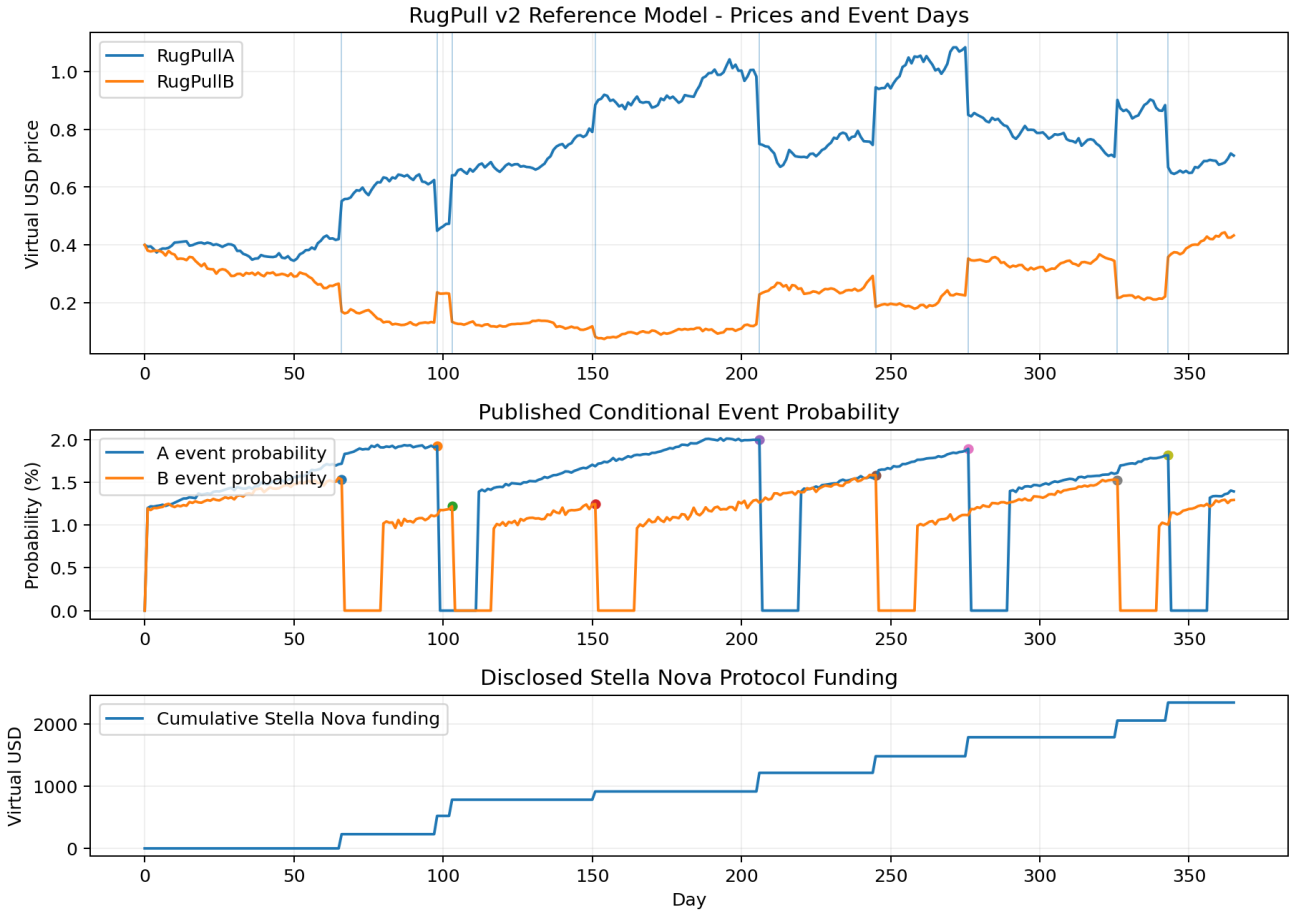

One seeded path is not a forecast. The complete JSON and CSV outputs accompany this document so each value can be inspected and reproduced.

Reference metrics

Events	9	**Organic volume**	USD 5,511,854
Ending A	USD 0.709	**Ending B**	USD 0.433
A drawdown	-40.4%	**B drawdown**	-81.3%
SN funding	USD 2,346	**Treasury NAV return**	40.2%
Signal lift	1.34x	**Brier score**	0.0121

Implementation boundary

Included	Landing page, public equations, 365-day model, probability chart, event and funding ledger, participant diagnostics, JSON comments, daily snapshots, admin status, testnet token creation, asynchronous randomness, and virtual reserve events.
Excluded	Mainnet deployment, decentralized-exchange router, real stablecoin custody, public-market transaction signer, automated real-value liquidation, redemption, and return promises.
Required before value	Independent legal opinions, institutional authorization, economic and contract audits, multisignature custody, public terms, comprehension testing, tax analysis, BSA/AML analysis, and a controlled-study protocol.

The strongest version of RugPull is one designed to discover that its own ethical proposition may fail. Unfavorable findings, participant misunderstanding, and institutional conflicts must be publishable outcomes rather than reasons to suppress the study.