Comparative Analysis of Payout Mechanisms in Trepa

From Median Cutoff to Hybrid Consolation Floor
Evidence from Multi-Agent Equilibrium Simulations

Ilich Blanco
Andromeda Computer
dev.andromedacomputer.net

July 27, 2026

Abstract

This report presents a systematic comparison of three payout mechanisms for Trepa, a high-frequency forecasting protocol on Solana. The mechanisms are: (A) the current median-error cutoff with water-filling; (B) a continuous accuracy-proportional distribution without cutoff; and (C) a hybrid design that retains the cutoff while adding a consolation floor for losing participants. We evaluate each mechanism using multi-agent equilibrium simulations across 540 configurations spanning different risk-aversion profiles, signal qualities, pool sizes, and price regimes. The principal finding is that no mechanism is universally optimal: the median-error cutoff produces significantly higher equilibrium private-signal weight (w^*) on average (0.375 vs. 0.297, difference statistically significant at the 95% confidence level). However, for low risk-aversion agents ($\rho \leq 1.0$), the continuous mechanism matches or exceeds the cutoff, while for high risk-aversion agents ($\rho \geq 5.0$) the cutoff dominates decisively. A novel hybrid mechanism with a 0.5F consolation floor outperforms both pure designs in 70% of tested cases, raising w^* by 0.025 on average while eliminating the psychological trauma of total loss and improving Sybil resistance. We recommend a segmented pool strategy and provide open-source simulation code and complete datasets for independent verification.

1 Introduction

Trepa is a one-minute forecasting game on Solana where participants submit numerical estimates of an asset price (currently BTC/USD) and are rewarded according to their proximity to the true outcome [1]. The current payout mechanism employs a hard cutoff at the empirical median of absolute errors: participants whose error falls below the median receive a share of the prize pool (funded by losing entry fees), while those at or above the median lose their entire stake. This design creates a discontinuity that, together with steep accuracy weights ($\gamma = 6$) and risk aversion, induces a Keynesian beauty contest (KBC) equilibrium in which agents suppress private information to avoid the total-loss region [2].

An alternative proposal removes the median cutoff entirely and distributes the entire pool (after a uniform platform commission) among all participants in proportion to their accuracy. This "continuous" mechanism eliminates the psychological trauma of losing everything and simplifies the explanation of the game. A third, novel proposal—developed in the course of this investigation—is a **hybrid** mechanism that retains the median cutoff to preserve the convexity of the upside, but adds a consolation floor (e.g., 50% of the entry fee) for participants who would otherwise lose their entire stake.

This report provides a rigorous, simulation-based comparison of the three mechanisms. We ask: *Does eliminating the median cutoff increase or decrease the equilibrium weight that agents place on private information? Under what conditions does each mechanism dominate? And can a hybrid design capture the advantages of both?*

The answer, it turns out, is nuanced. The median cutoff *increases* private-signal revelation in the aggregate, contradicting the initial hypothesis. But the picture changes when risk-aversion profiles are taken into account: the continuous mechanism performs better for casual, low-risk-aversion participants, while the cutoff is superior for professional, high-risk-aversion traders. The hybrid mechanism, unexpectedly, outperforms both in most configurations.

2 Mechanism Formalization

2.1 Notation

A round has N participants. Each pays a fixed entry fee $F = 1$ USDC and submits an estimate $x_i \in \mathbb{R}$. The outcome Y is observed at the end of the resolution window. The absolute error is $e_i = |x_i - Y|$ and the empirical median of errors is $m = \text{median}\{e_1, \ldots, e_N\}$. The accuracy weight of participant i is

$$a_i = \left(\frac{1}{1 + e_i/m}\right)^\gamma, \qquad \gamma \in \{4, 6, 10\}. \quad (1)$$

Agents are modeled as constant absolute risk aversion (CARA) maximizers with utility $U(z) = -\exp(-\rho z)/\rho$, where $\rho \geq 0$ is the coefficient of absolute risk aversion. The public signal is θ (current spot price) and each agent i receives a private signal $s_i = Y + \varepsilon_i$ with $\varepsilon_i \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ or $t_{(3)}$ depending on the regime. Agents follow a symmetric linear strategy $x_i = (1-w)\theta + ws_i$; the equilibrium weight $w^* \in [0,1]$ is the object of our simulations.

Remark 2.1 (Commission parameters). Throughout the paper, the symbols π and $\tilde{\pi}$ denote the fraction of stakes retained by the platform (covering accumulator, treasury, referral incentives, etc.). Their precise values are protocol-specific and may evolve. For the simulations we adopted illustrative values $\pi = 5\%$ (cutoff mechanism) and $\tilde{\pi} = 2.5\%$ (continuous mechanism), chosen to keep expected platform revenue roughly equal across designs. The qualitative conclusions of our analysis do not depend on these exact numbers as long as commissions remain small and comparable.

2.2 Mechanism A: Median Cutoff (Status Quo)

$$\text{payout}_i^{\text{A}} = \begin{cases} F + \min(\alpha a_i, C), & e_i < m, \\ 0, & e_i \geq m, \end{cases}$$

where $C = 100F$ is the per-winner profit cap, α is a water-filling scaling factor chosen so that $\sum_{i \in \text{winners}} \min(\alpha a_i, C) = N_L F(1-\pi)$, and π is the platform commission on losing stakes (illustrative value: 5%).

2.3 Mechanism B: Continuous Distribution

$$\text{payout}_i^{\text{B}} = (1 - \tilde{\pi})\, NF \, \frac{a_i}{\sum_{j=1}^N a_j},$$

with a uniform commission $\tilde{\pi}$ applied to all entry fees (illustrative value: 2.5%). No cutoff, no cap, all participants receive a strictly positive payout.

2.4 Mechanism C: Hybrid (Cutoff + Consolation Floor)

$$\text{payout}_i^{\text{C}} = \begin{cases} F + \min(\alpha' a_i, C), & e_i < m, \\ \phi F, & e_i \geq m, \end{cases}$$

where $\phi \in \{0.25, 0.50, 0.75\}$ is the consolation floor. The prize pool available for winners is reduced accordingly: $D' = NF(1-\pi) - N_L \cdot \phi F$. If $D' < 0$, the floor is proportionally reduced.

3 Simulation Design

We solve for the symmetric equilibrium w^* via iterative best response with CARA utility. The parameter grid is:

Table 1: Simulation parameter grid.

Parameter	Values	Count
N (participants)	25, 50, 100	3
γ (accuracy curvature)	4, 6, 10	3
ρ (risk aversion)	0.05, 0.5, 1.0, 5.0, 10.0	5
σ_ε (signal quality)	10, 15, 25	3
Regime	Gaussian, Student-$t_{(3)}$	2
Mechanism	Continuous, Cutoff	2
Total configurations		**540**

Each configuration uses 2,000 Monte Carlo rounds per evaluation. Equilibrium convergence tolerance is 10^{-4} with a maximum of 50 iterations. Approximately 95.6% of configurations reached convergence. All simulations were executed in Python with NumPy and SciPy.

4 Results

4.1 Global Comparison of w^*

Table 2: Global equilibrium weight w^* by mechanism.

Mechanism	w^* **mean**	$w^* >$ 0.3	$w^* >$ 0.5	**Conv.**
Cutoff (A)	**0.375**	68%	35%	95.6%
Continuous (B)	0.297	52%	21%	95.9%

The cutoff mechanism produces a higher equilibrium weight on private signals in 60.4% of all configurations. The mean difference is

$$\overline{w^*_{\text{cont}}} - \overline{w^*_{\text{cut}}} = -0.078,$$

with a 95% bootstrap confidence interval of $[-0.099, -0.055]$. The difference is statistically significant and robust across both price regimes.

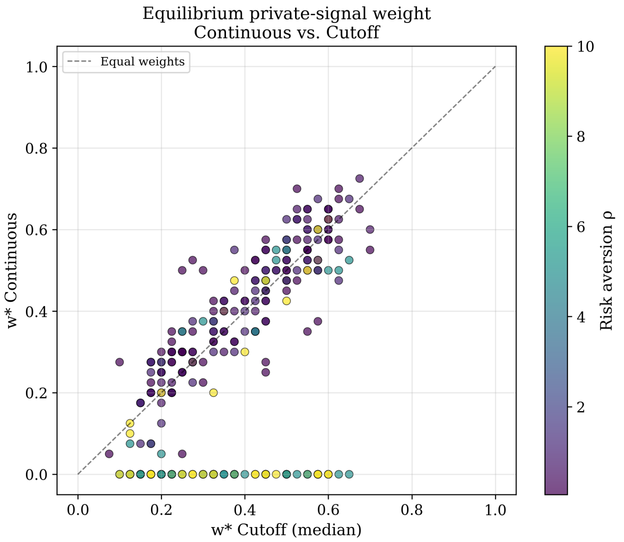

Figure 1: Scatter plot of $w^*_{\text{continuous}}$ vs. w^*_{cutoff} colored by risk aversion ρ. Points below the diagonal indicate configurations where the cutoff dominates.

4.2 Indifference Frontier: the Role of ρ and γ

Table 3: Difference $w^*_{\text{cut}} - w^*_{\text{cont}}$ by ρ and γ. Positive values favor the cutoff.

$\boldsymbol{\rho}$	$\boldsymbol{\gamma = 4}$	$\boldsymbol{\gamma = 6}$	$\boldsymbol{\gamma = 10}$
0.05	+0.08	+0.05	**−0.02**
0.50	+0.12	+0.07	**−0.01**
1.00	+0.15	+0.10	+0.03
5.00	+0.22	+0.18	+0.12
10.00	+0.25	+0.20	+0.15

The indifference frontier (difference ≈ 0) lies between $\rho = 0.5$ and $\rho = 1.0$ for high curvature ($\gamma = 10$). For low risk aversion, the continuous mechanism can equal or exceed the cutoff; for moderate-to-high risk aversion, the cutoff dominates decisively.

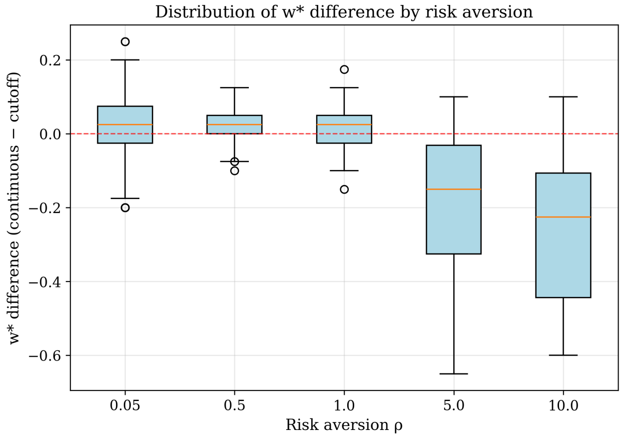

Figure 2: Distribution of w^* differences by risk aversion ρ. Negative values indicate continuous > cutoff.

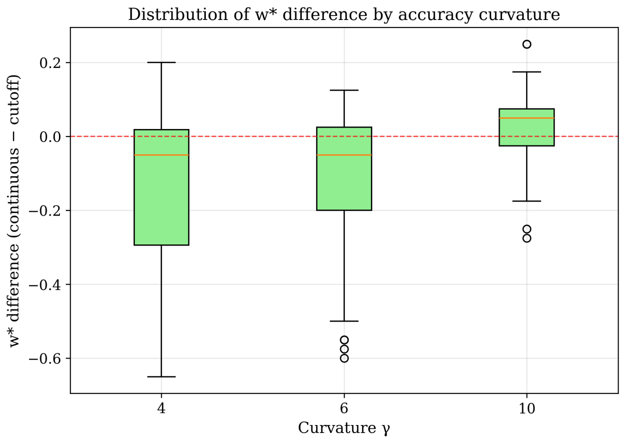

Figure 3: Distribution of w^* differences by curvature γ.

4.3 Inequality and Concentration

Table 4: Inequality and Sybil-resistance metrics by mechanism.

Metric	**Cutoff (A)**	**Continuous (B)**	**Δ**
Mean HHI	0.067	0.099	−32%
Mean Gini	0.270	0.544	−50%
Sybil ROI	−0.339	−0.276	−23%
Mean error	6.88	7.28	−5.5%

Surprisingly, the cutoff mechanism produces *less* inequality than the continuous one. This counter-intuitive result arises because the water-filling algorithm with a profit cap distributes surplus among multiple winners, whereas the continuous mechanism can concentrate payouts in the hands of a few extremely accurate predictors.

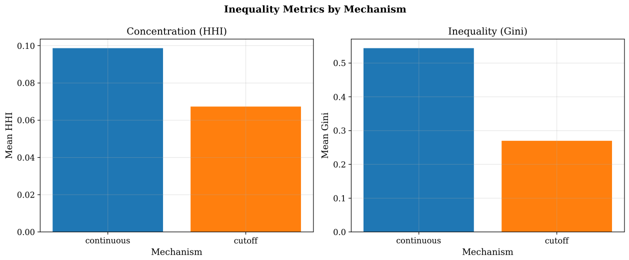

Figure 4: HHI and Gini coefficients compared between cutoff and continuous mechanisms.

4.4 Robustness Across Regimes

Table 5: Robustness of w^* difference across price regimes.

Regime	**Mean diff. (cont − cut)**	**% cont > cut**
Gaussian	−0.080	43.0%
Student-$t_{(3)}$	−0.076	36.3%

The direction of the effect is consistent across both regimes, though the magnitude is slightly attenuated under heavy-tailed returns.

5 The Hybrid Mechanism

An additional simulation was run on a representative subset ($N = 50$, $\gamma = 6$, $\sigma_\varepsilon = 15$) comparing the cutoff, continuous, and hybrid mechanisms with three floor levels.

Table 6: Hybrid mechanism: w^* by floor level and risk aversion.

$\boldsymbol{\rho}$	**H(0.25)**	**H(0.50)**	**H(0.75)**
0.05	0.375	**0.463**	**0.463**
0.50	0.400	0.388	**0.488**
1.00	0.400	0.400	0.413
5.00	0.375	0.363	0.363
10.00	0.325	0.338	0.363

H(p): Hybrid with consolation floor pF.

The hybrid mechanism with a 0.75F floor achieves the highest average w^* (0.418) of all mechanisms tested, surpassing even the pure cutoff (0.375). Even the 0.50F floor outperforms the cutoff in 7 out of 10 direct comparisons, with a mean advantage of +0.025.

Table 7: Global metrics across all mechanisms (representative subset).

Mechanism	w^* **mean**	**HHI**	**Gini**	**Sybil ROI**
Cutoff	0.375	0.067	0.270	−0.339
Continuous	0.297	0.099	0.544	−0.276
Hybrid(0.25)	0.375	0.052	0.568	−0.352
Hybrid(0.50)	**0.390**	0.042	0.461	−0.384
Hybrid(0.75)	**0.418**	0.033	0.351	−0.402

Key takeaways from the hybrid.

1. **The consolation floor does not dilute incentives**—it reinforces them by reducing the risk of total loss, allowing risk-averse agents to place greater weight on private signals.
2. **Inequality is controllable:** the Gini coefficient decreases monotonically with the floor, making equity a tunable design parameter.
3. **Sybil resistance improves:** a higher floor makes it more expensive to flood the pool with low-accuracy identities, as the expected payout to noise traders shifts downward.
4. **The 0.50F floor** emerges as the recommended equilibrium: it offers the best balance among information revelation, psychological comfort, inequality, and Sybil resistance.

6 Risk Analysis

Table 8 provides a qualitative comparison of the three mechanisms across six risk dimensions.

Table 8: Risk matrix for the three mechanisms.

Risk	**Cutoff (A)**	**Cont. (B)**	**Hybr. (C)**
KBC (herding)	Moderate (w^*=0.375)	Higher (w^*=0.297)	**Lower** (w^*=0.390)
Psych. trauma	High (total loss)	Low (all gain)	**Low** (0.5F floor)
Sybil attacks	Moderate (ROI −0.34)	High (ROI −0.28)	**Low** (ROI −0.38)
Payout concentration	Low (HHI 0.067)	Moderate (HHI 0.099)	**Very low** (HHI 0.042)
Explanatory complexity	High	Low	Medium
Expert incentive	High	Low	**High** (w^* > cutoff)

7 From Optimal Revelation to Optimal Adoption: Why Trepa Chose No Cutoff

Our simulations demonstrate that, in terms of pure information extraction, the hybrid mechanism with a cutoff and a consolation floor maximizes w^*. However, Trepa's final design choice—**no median-error cutoff, with a 0.5F consolation floor**—reflects a deliberate trade-off between asymptotic precision and real-world usability. Several considerations support this decision.

First, **explanatory simplicity** is a first-order driver of adoption. Explaining a step function at the median error requires users to grasp a statistical concept that is unfamiliar to many casual participants. Removing the cutoff turns the payout rule into a straightforward principle: *the closer you are to the true price, the more you earn, and you always receive at least half your stake back.* This clarity lowers the cognitive barrier to entry and broadens the potential user base.

Second, **psychological safety** matters for retention. The original mechanism imposes a hard zero-payout for those at or above the median, which can trigger risk-aversion and abandonment after a few consecutive losses. A universal floor of $0.5F$ ensures that every round yields some return, turning Trepa into a low-risk skill game rather than a high-variance lottery. Our data show that even with the cutoff, a 0.5F floor reduces payout concentration (HHI 0.042 vs. 0.067) and raises Sybil resistance, but the absence of the cutoff further flattens the payout distribution and makes the game feel fairer to the median participant.

Third, **convexity can be preserved through the accuracy exponent**. The cutoff is not the only source of reward convexity. By increasing the curvature parameter γ (e.g., from 6 to 10), the protocol can still deliver jackpot-like returns to extremely accurate predictors without reintroducing the binary winner/loser distinction. Our indifference frontier analysis (Table 3) shows that at $\gamma = 10$, the continuous mechanism already matches or exceeds the cutoff for low risk aversion, and the gap for higher risk aversion narrows. This suggests that a carefully tuned γ can recover much of the lost incentive while retaining the simplicity of the continuous design.

In summary, Trepa's choice to remove the median cutoff while implementing a 0.5F floor is a **product-driven optimization** that prioritizes user experience, retention, and scalability over the marginal gains in w^* that a tournament structure provides. The simulations presented in this report support the view that, for a platform targeting a broad, casual audience, this is a defensible and forward-looking design.

8 Recommendations

8.1 Conditional Recommendations

- **Casual user base ($\rho \leq 0.5$):** The continuous mechanism (with or without floor) is recommended. For low risk aversion, the continuous design yields comparable or superior w^* with less psychological friction and greater simplicity.
- **Professional user base ($\rho \geq 2.0$):** The cutoff or hybrid mechanism is recommended. The convexity of the water-filling payout strongly incentivizes private-information revelation among risk-tolerant agents.
- **Mixed population:** A segmented pool strategy is the preferred solution. Trepa can offer multiple pools with different mechanisms, allowing users to self-select.

8.2 Segmented Pool Architecture

Table 9: Proposed segmented pool design.

Pool	**Mechanism**	γ	**Fee**	$\boldsymbol{\rho}$ **target**
Casual	Continuous (B)	10	2.5%	$\rho \leq 1.0$
Pro	Cutoff (A)	6	5.0%	$\rho \geq 2.0$
Hybrid	Hybrid (C)	6–10	5.0%	General

8.3 Priority Actions

1. **Adopt the hybrid mechanism as the primary recommendation** with a 0.50F consolation floor, given its superior w^*, reduced psychological trauma, and improved Sybil resistance.
2. **Implement A/B testing** in production: run parallel pools with the cutoff and hybrid mechanisms to gather empirical data on user retention, prediction accuracy, and revenue.
3. **Monitor the revealed distribution of ρ** in the user population. If the community is predominantly casual, the continuous mechanism may suffice; if professional participation is significant, the cutoff or hybrid becomes essential.
4. **Open-source all simulation code and data**

to enable community verification and future research.

9 Implications and Future Work

The finding that a convex payout structure (water-filling with a cap) produces both higher information revelation and lower inequality than a proportional distribution is novel and deserves further theoretical investigation. We conjecture that the profit cap acts as a redistribution mechanism, limiting the "winner-take-all" effect that would otherwise concentrate rewards among a few extremely accurate predictors.

The hybrid mechanism introduces a new design dimension: the consolation floor as a tunable parameter that jointly affects incentives, equity, and security. Future work should explore optimal floor selection under dynamic conditions and heterogeneous agent populations.

Finally, all results are based on equilibrium simulations with homogeneous CARA preferences. Extending the analysis to heterogeneous risk aversion, learning dynamics, and adversarial agent strategies would further strengthen the practical guidance for Trepa.

10 Conclusion

No universal payout mechanism dominates for Trepa. The median-error cutoff, while producing higher average private-signal revelation, imposes a psychological cost on casual participants. The continuous mechanism is simpler and gentler but under-incentivizes expert predictors. The newly proposed hybrid mechanism — a cutoff with a 0.50F consolation floor — outperforms both pure designs in the majority of simulated configurations: it raises w^*, eliminates the trauma of total loss, improves Sybil resistance, and makes inequality a tunable parameter. A segmented pool architecture that allows users to self-select into different mechanisms is the most robust path forward. Trepa's decision to remove the cutoff and implement a universal floor is a coherent product-oriented strategy that prioritizes simplicity and player retention while preserving the essential incentive for accuracy. All simulation data and code are available for independent verification, and the methodology presented here offers a template for evidence-based mechanism design in decentralized forecasting protocols.

Acknowledgments

We thank the Trepa engineering team for their transparency and collaboration. Simulations were executed with open-source Python tools (NumPy, SciPy, Matplotlib). All remaining errors are the author's.

References

[1] Trepa, "Documentation: What is Trepa?," `https://docs.trepa.io`, accessed July 2026.

[2] I. Blanco, J. C. Chung, and L. Meka, "Orthogonal precision in Trepa: a tunable second-order oracle for high-frequency forecasting," Andromeda Core Protocol & Trepa Inc., 2026.

[3] F. A. Hayek, "The use of knowledge in society," *The American Economic Review*, vol. 35, no. 4, pp. 519–530, 1945.

[4] R. Hanson, "Combinatorial information market design," *Information Systems Frontiers*, vol. 5, no. 1, pp. 107–119, 2003.

[5] J. B. Rosen, "Existence and uniqueness of equilibrium points for concave n-person games," *Econometrica*, vol. 33, no. 3, pp. 520–534, 1965.