Litepaper: Externalized Price Discovery in Decentralized Prediction Markets

Abstract

This litepaper outlines a design for decentralized prediction markets that achieve live, truth-seeking price discovery without requiring on-chain order books or automated market makers (AMMs).

By separating payout accounting (the *primary pari-mutuel pool*) from price formation (the *secondary NFT market*), we enable market-derived probabilities to emerge naturally from open trading behaviour. The system maintains full decentralization, minimizes capital requirements, and retains the informational value of LMSR-style markets without introducing bounded-loss mechanisms or liquidity parameters.

1. Overview

Traditional prediction markets rely on one of two models:

- **Order-book-based systems**, where bids and asks are matched manually or algorithmically.

- **Cost-function AMMs**, such as the Logarithmic Market Scoring Rule (LMSR), which maintain continuous prices through a mathematical function and a pre-funded liquidity parameter b.

Both models achieve continuous prices, but each introduces operational challenges:

- Order books require continuous liquidity and market-maker participation.
- AMMs require an upfront subsidy and expose creators to bounded loss.

Our approach decouples these roles entirely.
We retain a simple **pari-mutuel vault** on-chain and shift price formation to an **externalized secondary market** for position receipts.

2. System Design

2.1 Core mechanism

Each market is defined by two vaults (**YES**, **NO**), representing binary outcomes.
Users deposit a stake on one side, receiving a **compressed NFT (cNFT)** as a transferable receipt.

Each cNFT includes:

- market_id
- outcome_side
- shares (quantity of $1-payout units)
- Stake_amount (USDC)
- creation_time

At resolution, the vault redistributes losing stakes to winning participants.
The holder of each cNFT at settlement can redeem *$1 × shares* if the outcome matches their side; otherwise, the NFT expires worthless.

2.2 Secondary market trading

Because each cNFT represents a claim paying $1 conditional on an outcome, its **market price** reflects the probability of that outcome occurring.

If an NFT with shares=10 trades for 6.40USDC, the implied probability is:

$$p_{\text{implied}} = 6.40/10 \ = \ 0.64$$

Aggregating these implied probabilities across all trades yields a **live market estimate** identical in interpretation to an LMSR-derived price.

Unlike an AMM, this mechanism does not enforce liquidity or pricing internally — those emerge from open market demand on secondary marketplaces (e.g. Tensor, Magic Eden).

2.3 Aggregation and Indexing

To construct a continuous probability feed, an indexer or oracle performs:

1. Collection of recent cNFT trades (per market and side).
2. Computation of per-share prices.
3. Normalization across outcomes to ensure $\sum ipi = 1$.
4. Publication of aggregated values as an open, permissionless data feed (e.g., via Switchboard or Chainlink).

This creates a verifiable on-chain “probability oracle” derived entirely from decentralized trading activity.

3. Market Dynamics and Predictive Accuracy

3.1 Emergent price curve

Secondary markets function as **distributed order books.** Traders reprice cNFTs in response to information changes, resulting in a smooth, continuous probability curve over time — the same behaviour that LMSR AMMs generate via their internal cost function.

3.2 Arbitrage feedback

If secondary NFT prices diverge from vault-implied odds (based on current YES/NO vault ratios), arbitrageurs can exploit the difference by:

- Buying undervalued NFTs and redeeming at resolution, or
- Entering the vault directly to mint new receipts and sell them at a premium.

This feedback loop aligns the vault’s notional probabilities with secondary-market consensus.

3.3 Robustness

- **Thin markets:** Use time-weighted averages (TWAP) and liquidity-weighted medians to smooth early volatility.
- **Price manipulation:** Open, multi-venue data aggregation dilutes the influence of any single trade.
- **Settlement integrity:** Metadata is program-signed; counterfeit cNFTs cannot claim payouts.

4. Advantages of Externalized Discovery

Category	Externalized cNFT Model	LMSR / AMM Model
Capital requirements	None beyond user stakes	Requires pre-funded bounded-loss parameter b
Complexity	Simple pari-mutuel vault logic	Continuous math, gradient updates
Liquidity	Shared across all NFT marketplaces	Confined to each AMM pool
Decentralization	Emergent from open trading	Defined by single program instance
Composability	NFTs integrate with standard Solana marketplaces	AMM requires dedicated front-end
Regulatory clarity	Vault = pool NFTs = transferable claims	AMM may constitute managed market-making

5. Transition Path

1. **Phase 1 (Complete):**
   Pari-mutuel vaults with direct stake deposits; odds visible only post-resolution.
2. **Phase 2 — cNFT issuance (Complete):**
   Each stake mints a transferable cNFT claim.
3. **Phase 3 — External aggregation (Current):**
   Indexers compute implied probabilities from secondary-market trades.
4. **Phase 4 — Optional hybridization (Future):**
   Introduce LS-LMSR mechanics to smooth thin markets using fee-funded liquidity parameters.

At every stage, the protocol remains non-custodial, permissionless, and transparent.

6. Broader Applications

- **Forecasting infrastructure:** Live probabilities can serve as inputs to DeFi hedging tools or insurance contracts.
- **Reputation systems:** Aggregated signals represent collective belief in public events.
- **Data provenance:** Prediction signals can inform dynamic pricing for verifiable data or outcomes.

7. Conclusion

By externalizing price discovery, we remove the need for on-chain market-making while preserving the core informational value of prediction markets.

Each NFT receipt functions as a standardized, verifiable claim on an outcome, and the collective pricing of these receipts across secondary markets yields real-time, decentralized probability data.

This approach maintains the simplicity of pari-mutuel settlement, achieves prediction-market accuracy, and scales without capital inefficiency — providing a practical, composable foundation for decentralized forecasting on Solana.

Summary

> A decentralized prediction market can function without order books or AMMs by letting cNFT receipts trade freely on secondary markets. The aggregated trade prices of these NFTs encode the same probabilities LMSR markets compute internally, providing continuous, decentralized insight with minimal on-chain complexity.