Palpitada: A Protocol for Decentralized Forecasting

Vinicius Santiago Tenente
vinicius.santiago@gmail.com

Max Nieuwoudt Hayes
maxhayeshome@gmail.com

Abstract

Palpitada is a noncustodial decentralized forecasting protocol built for the Ethereum Virtual Machine (EVM). Inspired by automated market maker (AMM) innovations, it allows users to take positions on the outcomes of real-world events via a novel implementation of event contracts. It introduces a new paradigm for quantifying belief and collective intelligence by allowing users to take positions on real-world outcomes through tokenized event markets. Palpitada builds an original, decentralized architecture with fully on-chain transparency, mathematically sound pricing via LMSR, and long-term community governance through the PPT token. This document details the underlying math, algorithms, and smart contract architecture that define Palpitada's technical foundation.

1. Introduction

Forecasting is one of humanity's oldest and most important tools—used to anticipate elections, economic events, weather, and geopolitical developments. Palpitada proposes a disruptive innovation in this space: a fully decentralized, tokenized framework for capturing and quantifying subjective beliefs on-chain.

Prediction markets have historically provided insight into public opinion, often outperforming polls and analysts in forecasting real-world outcomes. However, centralized platforms suffer from high fees, manipulation risks, and opaque operations. Palpitada aims to solve these issues by introducing an open, decentralized protocol where event outcomes are tokenized and recorded.

At its core, Palpitada is an engine of prediction powered by incentive-aligned cryptoeconomics. It enables users to express confidence in outcomes by purchasing outcome shares and leverages blockchain's trustless architecture to transparently track, resolve, and reward correct predictions.

At launch, all markets will be created and curated by the Palpitada core team to ensure clarity, fairness, and compliance. These curated markets will cover a wide range of relevant topics and follow strict standards for resolution sources and timelines. Over time, market creation will be progressively decentralized, allowing PPT token holders to propose, curate, and approve new markets via governance mechanisms.

Palpitada's business model is sustained by a 9% fee on each prediction placed. This fee funds protocol development, operations, and community incentives.

1.1 Why Crowd Forecasting Works: Theory and Evidence

The core hypothesis behind Palpitada is that **aggregated predictions from a large, diverse group of individuals can outperform predictions made by a few experts**. This principle — often referred to as the **"wisdom of crowds"** — is supported by both statistical theory and decades of empirical research.

Aggregation & Accuracy

Under assumptions of **diversity**, **independence**, and **decentralized input**, a group's collective prediction tends to cancel out individual errors and converge toward the true value — a phenomenon rooted in the **Law of Large Numbers (LLN)**.

> "Under the right circumstances, groups are remarkably intelligent, and are often smarter than the smartest people in them."
> — Surowiecki [1]

This is especially true in binary event forecasting, where Palpitada's design enables thousands of users to participate, generating a **collective probability** that is dynamically updated by the market's AMM pricing function.

Historical Evidence: Galton's Ox Experiment

The statistical roots of this phenomenon trace back to **Francis Galton's 1907 experiment** [2], where 787 fair-goers guessed the weight of an ox. The **average guess was 1,207 pounds**, while the actual weight was **1,198 pounds** — just **0.75% error**.

> Galton, F. (1907). *Vox Populi.* Nature, 75, 450–451.
> DOI: 10.1038/075450a0

Modern Research: Expert vs. Crowd

In his two-decade study of over 80,000 geopolitical predictions, **Philip Tetlock** found that **crowds of ordinary people consistently outperformed experts**.

> "The average expert was only slightly more accurate than a dart-throwing chimpanzee."
> — Tetlock [3]

Prediction markets — such as those analyzed by Wolfers and Zitzewitz — show that aggregating many small predictions results in highly accurate probabilities [4].

Additionally, Scott Page highlights that cognitive diversity plays a central role in the collective problem-solving ability of large groups [5].

2. Event Contracts: The Core Primitive

Each prediction market in Palpitada revolves around an Event Contract. These contracts represent binary outcomes (e.g., YES/NO) and allow users to take positions by purchasing outcome tokens specific to that event.

Key features of Event Contracts:

- **Binary Outcomes:** Markets resolve to one of two mutually exclusive outcomes.
- **Initialization:** Markets are initialized with neutral prices (e.g., $P(YES) = 0.5$, $P(NO) = 0.5$).
- **Dynamic Pricing:** Market prices for YES and NO shares fluctuate between 0 and 1 based on the volume of purchases for each outcome, as determined by the LMSR AMM.
- **Purchase-Only Interaction:** A distinctive feature of Palpitada is that users cannot freely sell or transfer their positions (outcome tokens) in a secondary market. They can only purchase a position (YES or NO) directly from the Event Contract's AMM.
- **Resolution and Payout:** Once the event's outcome is determined and the market is resolved, users holding outcome tokens corresponding to the correct outcome can withdraw their pro-rata share of the contract's collateral (e.g., USDC). Incorrect outcome tokens expire worthless.

Event contracts are deployed as immutable smart contracts on an EVM-compatible blockchain, ensuring transparency, finality, and censorship resistance. Users interact with them through standardized functions such as `predict()` (to purchase outcome tokens) and `withdraw()` (to claim winnings), which are validated by on-chain conditions and off-chain data inputs for resolution.

3. Market Architecture

3.1 LMSR-Based AMM

Palpitada's automated pricing for Event Contracts relies on an Automated Market Maker (AMM) governed by the Logarithmic Market Scoring Rule (LMSR). LMSR, introduced by Robin Hanson, is a well-regarded cost function specifically designed for use in prediction markets.

LMSR Cost Function

The foundation of the LMSR AMM is its cost function, C. This function determines the total cost the market maker has incurred to sell the current portfolio of outstanding shares. For a binary market with q_1 shares of "YES" and q_2 shares of "NO" outstanding, the cost function is:

$$C(q_1, q_2) = b \cdot \ln(e^{q_1/b} + e^{q_2/b})$$

Where:

- q_1: The total quantity of "YES" shares purchased.
- q_2: The total quantity of "NO" shares purchased.
- b: A liquidity parameter. A higher b value makes the market "deeper," meaning prices move less for a given trade size, while a lower b value makes prices more sensitive.

When a user purchases shares, they pay the difference in the cost function before and after their trade: $Cost_{trade} = C(q_{new}) - C(q_{old})$.

LMSR Price Function

The instantaneous price of a share for a particular outcome is derived from the partial derivative of the cost function with respect to the quantity of shares for that outcome. For outcome 1 ("YES" shares), the price P_1 is:

$$P_1(q_1, q_2) = \frac{\partial C}{\partial q_1} = \frac{e^{q_1/b}}{e^{q_1/b} + e^{q_2/b}}$$

Similarly, for outcome 2 ("NO" shares), the price P_2 is:

$$P_2(q_1, q_2) = \frac{\partial C}{\partial q_2} = \frac{e^{q_2/b}}{e^{q_1/b} + e^{q_2/b}}$$

Where:

- q_1 and q_2 are the current total quantities of YES and NO shares purchased, respectively.
- b is the liquidity parameter.

This price function ensures that $P_1 + P_2 = 1$ (if no fees are considered at this level), and the market price dynamically reflects the relative demand for each outcome, interpretable as the collective probability estimate. (See Appendix for the detailed derivation of the price function.)

LMSR Price Curve Chart

(This chart illustrates how the market probability (price of a YES share) shifts based on the number of YES shares purchased.)

3.2 Categories and Market Types

Markets on Palpitada can span various real-world categories, allowing for diverse forecasting opportunities. Examples include:

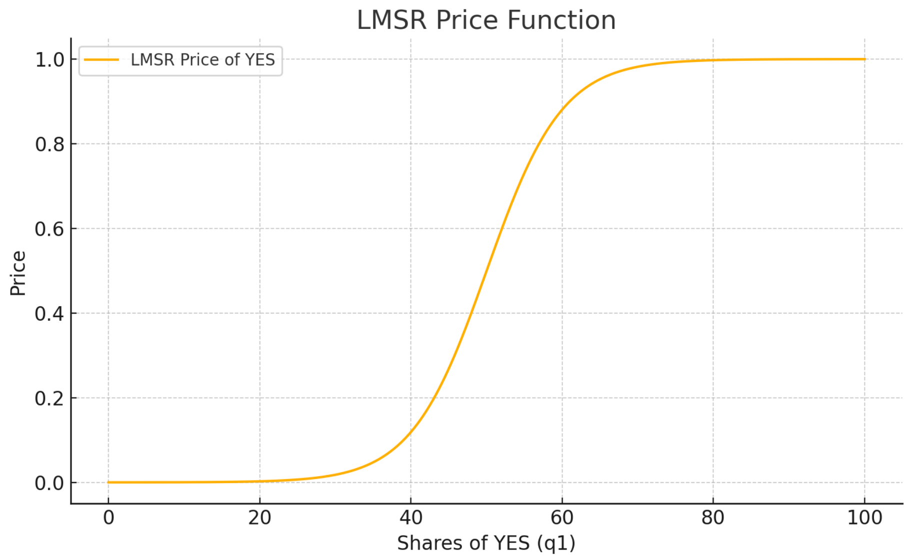

Figure 1: LMSR Price Curve

- **Sports:** “Will Team A win the upcoming match against Team B?”
- **Economics:** “Will the national inflation rate exceed 5% in Q3 of this year?”
- **Current Events:** “Will it rain more than 5mm in São Paulo tomorrow according to the National Weather Service?”
- **Technology:** “Will Company X release its new flagship product by December 31st?”

4. Tokenomics

4.1 PPT Token

Palpitada introduces its native governance token, PPT, which plays a central role in the protocol’s decentralization and long-term development.

- **Total Supply:** 1,000,000 PPT
- **Utility:** Used for protocol governance, potentially staking for rewards, and distributing certain incentives.
- **Governance:** PPT holders can propose and vote on protocol upgrades and parameter changes, such as fee structures, new market categories, oracle configurations, or treasury fund allocations. Governance contracts may utilize mechanisms like ERC20Votes to tally votes from token holders, potentially employing quadratic voting to promote broader consensus.

4.2 Token Allocation

Category	Allocation	Vesting / Release Schedule
Founders	20%	4-year vesting, 1-year cliff
Team	20%	4-year vesting, 1-year cliff
Private Sale	20%	1-year lockup, then 2-year vesting
Public Sale	20%	Quarterly release over 3 years
Treasury/Warchest	8%	Strategic reserve for ecosystem
Exchange Liquidity	6%	As needed for CEX/DEX liquidity
Airdrops/Growth	6%	Marketing and user incentives

4.3 Staking & Fee Distribution

A protocol fee is levied on market participation (i.e., when users purchase outcome tokens). For Palpitada, this fee is set at **3%** of the transaction value. These collected fees are accumulated into a smart contract vault. Periodically, these funds are distributed to PPT token holders who actively stake their tokens in the protocol, aligning long-term holders with the platform's success.

5. Market Resolution and Oracles

The integrity of a prediction market hinges on accurate and trustworthy market resolution. At launch, Palpitada will primarily rely on a curated set of highly reputable off-chain data sources and established news outlets (e.g., CNN, BBC, Reuters, official government statistics agencies, major sports league results) to determine the outcomes of events. The Palpitada core team will oversee this initial resolution process to ensure consistency and fairness.

Over time, Palpitada aims to transition towards a more decentralized oracle system. Potential components of this future system include:

- **Chainlink Data Feeds:** Utilizing Chainlink's established oracle networks for verifiable real-world data.
- **Voting-Based Resolution Protocols:** Implementing systems where PPT token holders (or a specialized group of staked resolvers) can vote on market outcomes, potentially using Schelling point mechanisms.
- **Fallback Mechanisms:** Employing a multisig committee of reputable individuals or organizations as a fallback resolver for contentious or ambiguous markets, with all decisions and reasoning made publicly verifiable.

This hybrid and phased approach is designed to ensure reliability and user trust from day one, while progressively building towards a more robust, decentralized, and censorship-resistant oracle system aligned with the protocol's core principles.

6. Regulatory Considerations

Palpitada is designed to operate as a decentralized information market protocol. Key aspects of its design relevant to the regulatory landscape include:

- **User Interaction with Smart Contracts:** Users interact directly with the deterministic logic of the Event Contract AMMs rather than a centralized operator.
- **Absence of Fixed Odds:** Market prices (odds) are determined dynamically by the LMSR algorithm based on user participation, not set by a central "house."
- **Algorithmic Counterparty:** The AMM, governed by transparent and immutable smart contract code, acts as the counterparty to user predictions. There is no discretionary "house" taking the other side of bets.
- **Crypto-Asset Settlement:** Transactions and payouts are conducted using cryptographic assets on the blockchain.
- **Publicly Auditable Criteria for Resolution:** Market terms and resolution sources are intended to be clearly defined and publicly verifiable.

The protocol is designed with the aim of complying with frameworks such as those outlined by the Brazilian CVM for crypto-assets and similar evolving international guidelines, positioning Palpitada outside the scope of traditional gambling or financial derivatives regulations where applicable. However, users are responsible for understanding and complying with the laws of their respective jurisdictions.

7. Conclusion

Palpitada represents a novel approach to market-based forecasting, combining the robust infrastructure of decentralized systems, the engagement of tokenized governance, and the mathematical rigor of proven pricing models like LMSR. By initiating with a curated market environment and a hybrid resolution model, the protocol aims for practical usability and trustworthiness from launch, while steadily progressing towards greater decentralization in market creation and outcome verification.

With its foundational focus on fairness, transparency, and community-driven evolution, Palpitada is poised to redefine how collective intelligence is harnessed and how future events are anticipated. Join us in building the future of programmable belief systems on-chain.

References

[1] J. Surowiecki, *The Wisdom of Crowds: Why the Many Are Smarter Than the Few*, Anchor Books, 2004.
[2] F. Galton, "Vox Populi," *Nature*, vol. 75, no. 1949, pp. 450–451, 1907. https://doi.org/10.1038/075450a0
[3] P. E. Tetlock, *Expert Political Judgment: How Good Is It? How Can We Know?*, Princeton University Press, 2005.
[4] J. Wolfers and E. Zitzewitz, "Prediction Markets," *Journal of Economic Perspectives*, vol. 18, no. 2, pp. 107–126, 2004. https://doi.org/10.1257/0895330041371321
[5] S. E. Page, *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*, Princeton University Press, 2007.

Appendix: Derivation of LMSR Price Function

The price P_1 of an outcome share (e.g., "YES") in an LMSR market is the partial derivative of the LMSR cost function $C(q_1, q_2)$ with respect to the quantity of shares for that outcome, q_1.

Let the cost function be $C(q_1, q_2) = b\ln(e^{q_1/b} + e^{q_2/b})$. To find $\frac{\partial C}{\partial q_1}$, we'll use the chain rule.

Let $u = e^{q_1/b} + e^{q_2/b}$. Then $C = b\ln(u)$. The derivative of $\ln(u)$ with respect to u is $\frac{1}{u}$. The derivative of u with respect to q_1 is:

$$\frac{\partial u}{\partial q_1} = \frac{\partial}{\partial q_1}(e^{q_1/b} + e^{q_2/b})$$

Since $e^{q_2/b}$ is treated as a constant when differentiating with respect to q_1:

$$\frac{\partial u}{\partial q_1} = \frac{\partial}{\partial q_1}(e^{q_1/b}) + 0$$

So,

$$\frac{\partial u}{\partial q_1} = e^{q_1/b} \cdot \frac{\partial}{\partial q_1}\left(\frac{q_1}{b}\right) = e^{q_1/b} \cdot \frac{1}{b}$$

Now, applying the chain rule for $\frac{\partial C}{\partial q_1}$:

$$\frac{\partial C}{\partial q_1} = b \cdot \frac{1}{u} \cdot \frac{\partial u}{\partial q_1}$$

Substitute $u = e^{q_1/b} + e^{q_2/b}$ and $\frac{\partial u}{\partial q_1} = \frac{1}{b}e^{q_1/b}$:

$$\frac{\partial C}{\partial q_1} = b \cdot \frac{1}{e^{q_1/b} + e^{q_2/b}} \cdot \left(\frac{1}{b}e^{q_1/b}\right)$$

The b in the numerator and the $\frac{1}{b}$ term cancel out:

$$\frac{\partial C}{\partial q_1} = \frac{e^{q_1/b}}{e^{q_1/b} + e^{q_2/b}}$$

This is the instantaneous price P_1 for outcome 1 (e.g., "YES" shares).

Note: You can use the same logic to find the partial derivative with respect to q_2 for outcome 2 (e.g., "NO" shares):

$$P_2 = \frac{\partial C}{\partial q_2} = \frac{e^{q_2/b}}{e^{q_1/b} + e^{q_2/b}}$$

This represents the instantaneous price P_2 for outcome 2.