Unified Liquidity for Continuous Prediction Markets

Current prediction markets represent continuous outcomes through collections of binary contracts. For a market forecasting Bitcoin's future price, traders may interact with contracts such as BTC > $100k, BTC > $125k, BTC > $150k, BTC > $175k, and so on. Each threshold becomes an independent market with its own liquidity pool, order book, and price discovery process.

This architecture introduces liquidity fragmentation. A trader forecasting Bitcoin between $140k and $180k cannot express that view directly. Instead, they must construct a synthetic position across multiple contracts. Liquidity becomes concentrated around a handful of popular strikes while neighboring strikes remain thinly traded. Information is distributed across many disconnected markets despite all contracts referencing the same underlying outcome.

The core insight behind Continuous Markets is that the outcome itself should be the market.

Rather than creating many contracts around a single variable, the protocol maintains one market, one liquidity pool, and one probability distribution over all possible outcomes. Traders express views as ranges over the outcome space. A forecast that Bitcoin will finish between $140k and $180k becomes a first-class market position rather than a synthetic combination of threshold contracts.

To support this architecture, the market requires a mechanism capable of pricing an effectively infinite number of possible ranges while maintaining a single shared liquidity pool. Continuous Logarithmic Market Scoring Rules (CLMSR) provide exactly this capability.

CLMSR extends Hanson's Logarithmic Market Scoring Rule (LMSR) from discrete outcomes to continuous outcome spaces. Instead of maintaining separate probabilities for individual contracts, the protocol maintains a single probability distribution over the entire outcome range. Every trade updates this distribution and therefore updates the market's collective forecast.

CLMSR preserves the desirable properties of LMSR, including continuous pricing, bounded loss, path independence, and automated market making while extending them to continuous outcomes. This allows a single liquidity pool to support an effectively infinite number of ranges without requiring separate order books or market makers for each possible position.

Prices emerge from changes in the market distribution itself. Traders do not purchase contracts against counterparties. Instead, they pay for the amount by which they reshape the market's forecast. Increasing confidence in one region of the outcome space necessarily reduces confidence elsewhere.

Because every position references the same underlying distribution, all traders interact with the same liquidity pool. Information discovered in one region of the outcome space immediately affects pricing across all other regions. The market therefore behaves as a single forecast rather than a collection of disconnected contracts.

This architecture naturally deepens liquidity. Traditional prediction markets divide capital across N independent contracts, each requiring its own liquidity. Continuous Markets concentrate liquidity into a single market state. Rather than requiring separate capital for every strike, one pool simultaneously backs all possible ranges. Market depth therefore scales with total participation rather than participation within individual contracts.

The result is a market structure where information, liquidity, and price discovery all converge into a single source of truth.

Trader returns are determined by the relationship between the purchased range and the realized outcome. If the final value falls inside a trader's selected range, the position receives a payout according to the market's settlement rules. Wider ranges contain more probability mass and are therefore more expensive to purchase. Narrower ranges contain less probability mass and are therefore cheaper, but generate larger returns if correct.

This creates a natural incentive structure. Traders are rewarded not only for being correct, but for being precise. A trader who correctly identifies a narrow, low-probability outcome contributes more information to the market than a trader who selects a broad interval that was already expected by consensus.

Continuous Markets therefore reward three properties simultaneously: correctness, precision, and disagreement with consensus.

The objective is not to replace binary prediction markets. Binary markets remain the optimal primitive for discrete events. Continuous Markets simply provide the appropriate market structure for scalar outcomes such as prices, revenue, inflation, weather, economic indicators, and dates.

TLDR;

Polymarket: scalar outcomes are being forced into binary structures

Therefore continuous markets are inevitable.

Pricing and Payoffs in a Continuous Range Market

A continuous market maintains a single market-implied distribution over all possible outcomes. Let the outcome be X, and let the outcome space be divided into small ticks or bins. Each bin i has a current market probability p_i, and all probabilities sum to one:

$$\sum_i p_i = 1$$

What the Distribution Represents

The distribution should be understood as the market's current forecast of reality. It is not a set of separate contract prices. It is a single object from which all range prices, tail probabilities, expected values, and confidence intervals can be derived. This is what makes the market forecast-centric rather than contract-centric.

A trader does not buy an isolated binary contract. Instead, they choose a range R, such as:

$$R = [70k, 75k]$$

The market price of that range is determined by the amount of probability mass currently assigned to it:

$$P(R) = \sum_{i \in R} p_i$$

This is the continuous-market equivalent of a binary contract price. In a binary prediction market, a YES share priced at 0.30 implies a 30% probability. In a continuous range market, a range priced at 0.30 means the market currently assigns 30% probability to the outcome landing inside that range.

When a trader buys a range, they increase exposure to that region of the outcome space. The market updates the shared distribution by increasing probability mass inside the purchased range and decreasing probability mass elsewhere. Under a CLMSR-style mechanism, the trader pays the change in the market's convex cost function:

$$Cost = C(q') - C(q)$$

where q is the current market state and q' is the market state after the trader's range position has been added. In the discrete LMSR form, the cost function is:

$$C(q) = b \cdot \ln\left(\sum_i e^{q_i/b}\right)$$

Where b controls market depth. A larger b means deeper liquidity and a lower price impact. A smaller b means shallower liquidity and higher price impact.

The implied probability of each bin is:

$$p_i = \frac{e^{q_i/b}}{\sum_j e^{q_j/b}}$$

A range is therefore priced by summing the probabilities of all bins inside the range:

$$P(R) = \sum_{i \in R} \frac{e^{q_i/b}}{\sum_j e^{q_j/b}}$$

Derived Forecasts

Because the market maintains a full probability distribution rather than a collection of isolated contract prices, a variety of forecasts can be derived directly from the market state.

The expected outcome is given by:

$$E[X] = \sum_i x_i p_i$$

The probability that the outcome exceeds a threshold a is:

$$P(X > a) = \sum_{i:x_i > a} p_i$$

The probability that the outcome lands within a range $[a, b]$ is:

$$P(a \leq X \leq b) = \sum_{i:a \leq x_i \leq b} p_i$$

Similarly, confidence intervals, tail probabilities, implied volatility, and other forecasting metrics can be derived from the same distribution.

This is an important distinction between Continuous Markets and traditional prediction markets. Binary markets produce prices for individual contracts. Continuous Markets produce a complete forecast surface from which many different market views can be extracted without requiring additional liquidity or separate markets.

At settlement, the market observes the realized outcome x^*. A range position pays if an only if the realized outcome falls inside the selected range:

$$\text{Payout}(R, x^*) = \begin{cases} 1 & \text{if } x^* \in R \\ 0 & \text{otherwise} \end{cases}$$

This means that traders are rewarded for both correctness and precision. A wide range is more likely to contain the final outcome, so it carries more probability mass and costs more to buy. A narrow range is less likely to contain the final outcome, so it is cheaper to buy, but pays a higher return if correct.

For example, suppose five traders enter positions:

$A = [50k, 60k]$

$B = [60k, 80k]$

$C = [70k, 75k]$

$D = [100k, 105k]$

$E = [50k, 80k]$

If the final outcome is:

$x^* = 72k$

then traders B, C, and E are correct, while A and D receive no payout.

However, the correct traders do not necessarily earn the same return. C selected the narrowest correct range, so the position was more precise and likely cheaper to buy. If correct, it should earn the highest return on capital. B selected a medium-width range and earns a moderate return. E selected the widest correct range, making the position safer but more expensive, so it earns the lowest return among the winners.

In general:

$$\text{Return} \propto \frac{\text{Payout}}{\text{Purchase Cost}}$$

and purchase cost is determined by the market-implied probability mass of the selected range:

$$\text{Purchase Cost} \approx P(R)$$

Thus, the market rewards three things:

$$\text{Correctness} + \text{Precision} + \text{Disagreement with Consensus}$$

A trader who selects a narrow, low-probability range and is correct earns a large return. A trader who selects a wide, high-probability range and is correct earns a smaller return. A trader whose range does not contain the final outcome earns nothing.

This creates a natural incentive structure. Traders are not simply rewarded for being correct; they are rewarded for being correct in a way that was informative to the market. The more precise and less consensus a correct forecast was at the time of purchase, the greater the return.

This is the core distinction between binary prediction markets and continuous range markets. In binary markets, each threshold has its own price and liquidity. In continuous range markets, every range is priced from the same shared distribution. All traders interact with one market, one liquidity pool, and one source of truth.

Liquidity Providers

A Continuous Market requires capital to quote prices, absorb trades, and pay winning positions at settlement. This capital is supplied by liquidity providers. Unlike traditional prediction markets, where liquidity providers must choose which individual contract or strike to support, liquidity providers in a Continuous Market deposit into a single pool that backs the entire outcome space.

Let the total liquidity supplied to the market be L. This liquidity determines the market's depth parameter b. In an LMSR-style market, b controls how sensitive prices are to trading activity. A larger b means deeper liquidity, lower price impact, and more stable prices. A smaller b means shallower liquidity, higher price impact, and faster price movement.

The relationship can be expressed simply as:

$$b = \alpha L$$

where α is a protocol-defined scaling parameter.

This means that as more liquidity enters the pool, the market becomes deeper. Traders can express larger views with less slippage, and probability mass moves more gradually in response to trades. Liquidity is not assigned to one strike or one contract. It increases depth across the entire forecast.

In a binary prediction market, capital is fragmented across many independent contracts. If a market has N thresholds and each threshold requires liquidity L_i, then total liquidity is distributed as:

$$L_{\text{total}} = \sum_{i=1}^{N} L_i$$

Each L_i only supports one contract. Liquidity deposited into BTC > \$150k does not directly support BTC > \$175k, even though both contracts reference the same underlying outcome.

In a Continuous Market, the same total liquidity supports every possible range simultaneously:

$$L_{\text{total}} \to b \to p(x)$$

Liquidity is transformed into market depth, and market depth supports the entire probability distribution. A trader buying $[70k, 75k]$, another buying $[60k, 80k]$, and another buying $[100k, 105k]$ all interact with the same liquidity base.

This is the source of the capital efficiency improvement. The market does not need separate liquidity for every threshold. One pool supports all expressions of the same outcome.

LP risk comes from settlement. When traders buy exposure to ranges, the market collects payments upfront. At resolution, the market pays positions whose ranges contain the realized outcome x^*. The LP pool earns from fees and from correctly priced order flow. It loses when settlement obligations exceed the premiums and fees collected against the realized outcome. It loses when the realized outcome falls in a region where the market has sold too much underpriced exposure.

Let trader j purchase a range R_j with position size s_j. The payoff of that position at settlement is:

$$\text{Payoff}_j(x^*) = s_j \cdot \mathbf{1}\{x^* \in R_j\}$$

where $1\{x^* \in R_j\}$ equals 1 if the realized outcome falls inside the trader's range and 0 otherwise.

The total payout owed by the market is:

$$\text{Payout}(x^*) = \sum_j s_j \cdot \mathbf{1}\{x^* \in R_j\}$$

The LP pool's settlement profit and loss can be written as:

$$\text{LP PnL}(x^*) = \text{Premiums Collected} + \text{Fees Collected} - \text{Payout}(x^*)$$

This makes LP risk explicit. LPs are not exposed to one binary outcome. They are exposed to the full shape of the market's sold payoff function across the outcome space.

Define the market's aggregate sold exposure at outcome x as:

$$S(x) = \sum_j s_j \cdot \mathbf{1}\{x \in R_j\}$$

Then the LP pool's worst-case payout is:

$$\max_x S(x)$$

This represents the maximum amount the market may owe if the final outcome lands at the most crowded point in the outcome space.

A healthy Continuous Market should therefore manage LP risk by monitoring the relationship between collected premiums, outstanding exposure, and worst-case settlement liability:

$$\text{Solvency Buffer} = \text{Pool Capital} + \text{Premiums Collected} - \max_x S(x)$$

The market is solvent if:

$$Solvency\ Buffer \ \geq \ 0$$

This framing makes liquidity provision similar to underwriting a portfolio of range-based claims. LPs earn fees and pricing edge in exchange for absorbing the risk that the realized outcome lands in a heavily purchased region.

The important difference is that this risk is managed at the level of the full distribution rather than at the level of isolated contracts. In a binary market, each pool only sees the risk of its own contract. In a Continuous Market, the protocol can observe the entire sold exposure curve $S(x)$, identify crowded regions, and adjust prices accordingly.

As traders buy more exposure to a region, $S(x)$ rises in that region. The market responds by increasing the cost of additional exposure there. This makes crowded outcomes progressively more expensive and protects LPs from selling unlimited exposure at stale prices.

The aggregate exposure curve $S(x)$ can be interpreted as the market's liability profile. Regions with little trader activity produce low exposure, while heavily purchased regions produce peaks in liability. The role of CLMSR is to continuously increase the cost of purchasing additional exposure in crowded regions, flattening the liability curve and protecting LP solvency.

This is the economic role of CLMSR. It converts liquidity into a convex pricing surface. Traders can still buy any range, but the cost of doing so increases as the market becomes more exposed to that region. LPs are compensated through trading fees and the spread between the prices traders pay and the final settlement obligations of the pool.

The LP is therefore not simply providing passive liquidity to isolated contracts. The LP is underwriting a continuously updated forecast. Their risk is the shape of the aggregate payout curve. Their return comes from fees, order flow, and the market's ability to price probability mass efficiently.

Design Considerations

LP incentive problem

LP compensation should scale with the risk being added to the pool. In a Continuous Market, LP risk is not uniform across the outcome space. Risk concentrates where many trader ranges overlap. We define the aggregate sold exposure curve as:

$$S(x) = \sum_j s_j \cdot \mathbf{1}\{x \in R_j\}$$

A trade that increases exposure in a low-risk region should pay a lower fee. A trade that increases exposure in an already crowded region should pay a higher fee. This creates a dynamic fee model where fees are a function of marginal exposure concentration.

Let a new trade add exposure $\Delta S(x)$. The concentration-adjusted fee can be expressed as:

$$\text{Fee}(R) = \phi \cdot \int_R S(x) \Delta S(x)\, dx$$

where ϕ is a protocol-defined risk fee parameter.

Under this model, LPs are compensated more when traders add risk to regions where the pool is already exposed. This makes liquidity provision more rational because fee income scales with the settlement risk being underwritten.

The LP's expected return can therefore be written as:

$$\text{LP Return} = \text{Base Trading Fees} + \text{Concentration Fees} - \text{Settlement Losses}$$

This does not eliminate LP risk, but it aligns LP compensation with the shape of the market's liability curve.

Adversarial analysis

A Continuous Market must be robust against attempts to manipulate the distribution before settlement. The simplest manipulation attempt is for a trader to buy a large range position in order to artificially shift probability mass toward that region.

Under CLMSR, this attack is not free. Any attempt to move the distribution requires paying the cost difference:

$$\text{Cost} = C(q') - C(q)$$

The farther the attacker attempts to move probability mass, and the deeper the market parameter b, the more expensive the manipulation becomes. In this sense, CLMSR converts manipulation into an economic cost rather than a free signaling attack.

However, manipulation cannot be ignored. A market is most vulnerable when liquidity is shallow, settlement is near, or the oracle outcome itself can be influenced. The protocol should therefore use basic safeguards:

1. higher liquidity requirements for markets approaching settlement,
2. dynamic fees for concentrated last-minute trades,
3. settlement windows rather than single-block oracle reads,
4. reputable oracle sources or committees for non-crypto outcomes,
5. market-specific caps on maximum exposure near resolution.

The goal is not to make manipulation impossible. No market can do that. The goal is to make manipulation economically irrational relative to the payoff gained.

Cold start problem

A continuous distribution with no traders begins as a prior. The simplest initialization is a uniform distribution across the valid outcome range:

$p_i = \frac{1}{N}$

For N bins.

However, a uniform prior may be too naive for markets with obvious external reference prices, such as BTC, inflation, or interest rates. For these markets, the protocol may initialize the distribution using an external reference prior. For example, a BTC price market may initialize around the current spot price with a wide variance, while a future inflation market may initialize around current analyst consensus.

A practical default is:

$$p_i^0 = \lambda p_i^{\text{external}} + (1 - \lambda)\frac{1}{N}$$

where p_i^{external} is an external prior and λ controls how strongly the market trusts that prior.

Initial liquidity should also be conservative. The market can begin with a smaller b, allowing early information to move prices more easily, and increase b as LP capital and trading volume grow.

This gives the market a reasonable starting point without pretending that the prior is the truth. Once trading begins, the distribution becomes market-driven.

Empirical simulation

A complete evaluation of Continuous Markets should include simulation and backtesting. The simplest simulation would compare three structures over the same scalar outcome:

1. independent binary threshold markets,
2. a continuous range market with fixed fees,
3. a continuous range market with concentration-adjusted fees.

The simulation should measure liquidity depth, trader slippage, LP profit and loss, worst-case settlement liability, and forecast accuracy over time.

For historical testing, markets can be simulated on outcomes such as BTC price, inflation, oil prices, or weather data. The goal is to compare whether a shared distribution produces better capital efficiency and more coherent forecasts than a ladder of binary contracts.

This paper defines the mechanism. Empirical validation remains future work.

Token Utility

The Continuous Market mechanism should not depend on a token to function. Pricing, settlement, and liquidity provision should work independently of token incentives.

If a protocol token exists, its role should be tied to governance and risk management rather than artificial demand. Token holders may govern parameters such as:

α

which maps LP capital to market depth,

ϕ

which controls concentration-adjusted fees,

oracle selection,

market listing standards,

and risk limits near settlement.

A token can also be used to backstop protocol-level risk through an insurance or safety module. In this model, token stakers absorb a portion of tail losses in exchange for a share of protocol fees.

The key principle is that token utility should emerge from governing and securing the forecasting layer, not from forcing traders or LPs to use the token as a medium of exchange.

Open questions to still answer

Continuous Markets provide a framework for expressing and pricing uncertainty over continuous outcome spaces. However, several important questions remain open.

The first concerns liquidity providers. While CLMSR provides bounded-loss market making and unified liquidity, the optimal relationship between liquidity, market depth, and LP compensation remains an area for further research. How should liquidity providers be rewarded for underwriting an entire probability distribution rather than a single contract? What forms of risk management are necessary as markets scale?

The second concerns distribution design. This paper assumes a discretized representation of a continuous outcome space, but alternative approaches may exist. Future implementations may explore adaptive binning, continuous parameterizations, or other representations that improve capital efficiency while preserving market integrity.

The third concerns market resolution and payout structures. While range-based positions provide an intuitive mechanism for expressing forecasts, there may be alternative settlement functions that better reward information, precision, or confidence. Determining the optimal balance between simplicity and expressiveness remains an open design question.

The fourth concerns composability. A continuously updated probability distribution contains significantly more information than a single contract price. Expected values, confidence intervals, tail probabilities, and derivative instruments can all be derived from the same underlying forecast. Understanding how these forecasting objects interact with broader financial systems remains largely unexplored.

More broadly, Continuous Markets raise a fundamental question about the future of prediction markets themselves. If a market can maintain a complete probability distribution over an outcome, should contracts remain the primary unit of exchange, or should forecasts become the primitive that markets are built around?

This paper proposes one possible architecture. Whether Continuous Markets become the dominant structure for scalar forecasting remains an empirical question. What is clear, however, is that continuous outcomes continue to be forced into architectures originally designed for binary events. Exploring alternatives may prove necessary if prediction markets are to become a general-purpose mechanism for forecasting the real world.