Non-Custodial Prediction Markets on Bitcoin

James Pierog

July 2024

1 Introduction

Prediction markets have emerged as powerful tools for aggregating dispersed information and forecasting future events. By allowing participants to trade contracts whose payoffs are contingent on the occurrence of specific outcomes, these markets harness the wisdom of crowds to produce accurate predictions across various domains, including politics, economics, and technology. However, traditional prediction markets often rely on centralized, custodial systems, which can introduce risks of censorship, manipulation, and counterparty default.

In this paper, we propose a novel approach to creating a non-custodial prediction market leveraging an automated market maker based on the logarithmic scoring rule (LMSR) utilizing Discreet Log Contracts (DLCs). Our system combines the information aggregation capabilities of prediction markets with the censorship resistance and security guarantees of Bitcoin, enabling truthful and robust event outcome forecasting without counterparty risk.

2 Background Information

2.1 Prediction Markets

Prediction markets are speculative exchanges where participants can buy and sell contracts whose payoffs depend on the occurrence of future events. The market prices of these contracts can be interpreted as collective forecasts of the probability of the events occurring [1]. For example, a contract that pays 1 USD if a specific candidate wins an election and 0 USD otherwise will theoretically trade at a price equal to the market's estimate of that candidate's probability of winning. Prediction markets have shown remarkable accuracy in various settings, often outperforming traditional forecasting methods [3]. Their success is attributed to their ability to aggregate diverse information from many participants and provide continuous, real-time updates as new information becomes available.

2.2 Automated Market Makers

To address liquidity challenges in prediction markets, especially for events with less trading activity, automated market makers (AMMs) have been developed. AMMs are algorithmic agents that provide continuous liquidity by always being willing to buy or sell contracts at some price. One of the most widely used AMM mechanisms is the Logarithmic Market Scoring Rule (LMSR), introduced by Hanson

[4]. The LMSR uses a cost function to determine the prices of contracts and ensures bounded loss for the market maker.

2.3 Discreet Log Contracts (DLCs)

Bitcoin, introduced by Satoshi Nakamoto in 2008 [5], is a decentralized digital currency and payment system. Bitcoin provides a secure, transparent, and censorship-resistant ledger for financial transactions. Discreet Log Contracts (DLCs), proposed by Dryja [2], enable non-custodial conditional payments on Bitcoin. This is achieved through the use of cryptographic commitments and signatures provided by an oracle [3].

In a DLC, an oracle commits to possible outcomes by publishing a signature. When the event occurs, the oracle reveals the signature which enables either participant to sign the corresponding contract execution transaction (CET) and collect the reward corresponding to the outcome. This allows the winning party to claim the funds without the oracle needing to interact directly with the contract or the blockchain.

3 A Framework for Shares, Price, and Cost with DLCs

Traders place bets by committing the share cost K BTC as collateral to a DLC. The market maker commits the other half of the contract, the reward R, betting against the trader at the market price. After the event outcome occurs, the oracle releases the signature attesting to the outcome, which unlocks the CETs corresponding to the correct agent receiving the total value locked (TVL $= K + R$). Either the trader or the market maker will receive the TVL.

3.1 Shares and Prices

The market maker always keeps track of shares and price, and uses a cost function to offer instant prices to traders[1]. All shares have Arrow–Debreu prices where $p \in (0, 100)$ sats before the event outcome occurs. After the event outcome occurs, shares convert to either 0 or 100 sats. Before the outcome occurence, the share price fluctuates based on the the number of shares issued in the market. The price per share is determined with a modified version of the LMSR [6].

We define the share price for yes on some outcome given a market state $q = (q_1 \dots q_i) = (y\ n)$ as

$$\begin{aligned} p_y(q) &= 100 \cdot \alpha \cdot \log(e^{y/b(q)} + ne^{n/b(q)}) + \frac{ye^{y/b(q)} + ne^{y/b(q)} - (ye^{y/b(q)} + ne^{n/b(q)})}{ye^{y/b(q)} + ne^{n/b(q)} + ne^{y/b(q)} + ye^{n/b(q)}} \\ &= 100 \cdot \alpha \cdot \log(e^{y/b(q)} + e^{n/b(q)}) + \frac{ne^{y/b(q)} - ne^{n/b(q)}}{ye^{y/b(q)} + ne^{n/b(q)} + ne^{y/b(q)} + ye^{n/b(q)}} \end{aligned}$$

[1]Shares are an abstraction over Bitcoin locked in DLCs over time. No digital representation of shares are issued (like tokenized shares or ecash). All shares are purchased by committing the amount of collateral corresponding to the cost of the number of shares to the DLC. The resultant DLC which is constructed is effectively a payment to the trader conditional on some event outcome occurring, representing the reward for the quantity of shares purchased.

and we define the share price for no on some outcome as

$$p_n(q) = 100 \cdot \alpha \cdot \log(e^{y/b(q)} + ne^{n/b(q)}) + \frac{ye^{n/b(q)} + ne^{n/b(q)} - (ye^{y/b(q)} + ne^{n/b(q)})}{ye^{y/b(q)} + ne^{n/b(q)} + ne^{y/b(q)} + ye^{n/b(q)}}$$
$$= 100 \cdot \alpha \cdot \log(e^{n/b(q)} + e^{n/b(q)}) + \frac{ye^{y/b(q)} - ye^{n/b(q)}}{ye^{y/b(q)} + ne^{n/b(q)} + ne^{y/b(q)} + ye^{n/b(q)}}$$

where

$$b(q) = \alpha \sum_i q_i = \alpha \cdot (y + n).$$

The value α reflects the market maker commission, q_i reflect the quantity of shares in outcome i. DLCs are most naturally structured as binary options, so the only outcomes are yes or no. Thus, the quantity of shares in outcome yes is y and the quantity of shares in outcome no is n. Prices are public and reflect the aggregated market belief in the probability of the event outcome occurring. Thus, share prices enable the public to forecast for future event outcomes.

3.2 Purchasing (or Selling) Shares with the Cost Function

The cost function for a trade given a state $q = (q_1 \dots q_i) = (y\ n)$ is defined as follows

$$C(q) = -100 \cdot b(q) \log\left(\sum_i \exp\left(\frac{q_i}{b(q)}\right)\right) = -100 \cdot b(q) \log(e^{y/b(q)} + e^{n/b(q)})$$

In practice, if a trader purchases Q quantity of shares on one side, then that trader must pay some cost K to the market maker. If the trader purchases shares of yes on the outcome, the cost is

$$K_y = C(q + Q) - C(q) = -b(q_n) \log(e^{(y+Q)/b(q_n)} + e^{n/b(q_n)}) - b(q_0) \log(e^{y/b(q_0)} + e^{n/b(q_0)})$$

where

$$b(q_0) = \alpha \cdot (y + n) \quad \text{and} \quad b(q_n) = \alpha \cdot (y + Q + n).$$

By symmetry, we can see that if a trader wants to purchase Q shares on outcome no, the cost is

$$K_n = C(q + Q) - C(q) = -b(q_n) \log(e^{y/b(q_n)} + e^{(n+Q)/b(q_n)}) - b(q_0) \log(e^{y/b(q_0)} + e^{n/b(q_0)}).$$

If the output of the cost function is positive, that means the traders are paying the market maker (for making bets/purchasing shares). If the output is negative, then it means the market maker is paying the traders (for selling shares/exiting bets).

In general the cost K yields Q shares for an average of $p = K/Q$ sats per share. Since the trader purchased Q shares, the total value locked in the DLC must be $100Q$ (as every share returns either 0 or 100 sats upon event occurrence). Thus, the reward R that the market maker must commit to the DLC is

$$R = 100Q - K.$$

Thus, the DLC has a liquidity ratio

$$\frac{Y}{N} = \frac{K}{R} = \frac{K}{100Q - K}.$$

Rewriting $K = Q \cdot p$ where p is the price per share, we see that

$$\frac{Y}{N} = \frac{Qp}{100Q - Qp} = \frac{p}{100 - p}.$$

Thus, when a trader purchases shares in an event outcome with a DLC, they essentially commit the cost K of those shares as collateral to a DLC. The market maker commits the reward R as collateral, based on the number of shares purchased by the trader at the time of commitment. The ratio of $Y : N$ in the DLC matches the price per share at time of purchase $p : 100 - p$. This ratio reflects the underlying probability/odds which the market prices the risk of the event outcome.

3.3 Example: Share Price and Costs with DLCs

To illustrate the concepts of share prices and costs in our DLC-based prediction market, let's consider a concrete example. We'll examine a market for a binary outcome event, with initial equal quantities of Yes and No shares, and observe how prices change with trades.

Consider a prediction market with the following initial parameters:

- Market maker commission $\alpha = 0.05$
- Initial quantities: $q = (10000, 10000)$ for (Yes, No) shares
- Trade size: $Q = 100$ shares

Initially, the market prices are:

$$p_{\text{Yes}} = 53.47 \text{ sats}$$
$$p_{\text{No}} = 53.47 \text{ sats}$$

These prices reflect an almost even split in probability between the two outcomes, with a slight deviation from 50 sats due to the market maker's commission.

Now, let's examine what happens when a trader buys 100 Yes shares:

- Cost of the trade: $K_{\text{Yes}} = 5470.90$ sats
- Average price per share: 54.71 sats
- New prices after the trade:

$$p_{\text{Yes}} = 55.95 \text{ sats}$$
$$p_{\text{No}} = 50.97 \text{ sats}$$

We observe that the cost of the trade is higher than the initial price would suggest ($5470.90 > 100 \times 53.47$). This difference represents the price impact of the trade. The average price per share (54.71 sats) falls between the initial price and the new Yes price, reflecting this impact. After the trade, the Yes price increases and the No price decreases, as expected. This shift represents the market's updated probabilities given the new information implied by the trade.

Similarly, if a trader buys 100 No shares instead:

- Cost of the trade: $K_{\text{No}} = 5470.90$ sats
- Average price per share: 54.71 sats
- New prices after the trade:

$$p_{\text{Yes}} = 50.97 \text{ sats}$$
$$p_{\text{No}} = 55.95 \text{ sats}$$

The symmetry in these results demonstrates the balanced nature of the market maker mechanism.

Finally, let's examine the liquidity ratios:

$$\text{Liquidity Ratio}_{\text{Yes}} = 1.2079$$
$$\text{Liquidity Ratio}_{\text{No}} = 1.2079$$

These ratios, calculated as $K/(100Q - K)$, represent the proportion of the total contract value that the trader is paying upfront. They match the theoretical ratio of $p/(100 - p)$, where p is the average price per share, confirming the correct implementation of our pricing model.

In the context of DLCs, these calculations determine the structure of the contract. For a Yes share purchase:

- The trader commits 5470.90 sats as collateral.
- The market maker commits $100Q - K_{\text{Yes}} = 4529.10$ sats as collateral.
- If the Yes outcome occurs, the trader receives the entire 10,000 sats (100 sats per share).
- If the No outcome occurs, the market maker receives the entire 10,000 sats.

This example demonstrates how our modified LMSR mechanism prices shares, calculates costs, and structures DLCs in a way that reflects market probabilities while maintaining liquidity and allowing for price discovery.

4 Share Price Speculation Before Event Occurrence

DLCs require a swap mechanism so traders can sell their shares for Bitcoin. Without exit liquidity, traders are locked into their position until the event occurs. If traders are unable to exit their position, then they cannot engage in trading to facilitate price discovery (and therefore probability forecasting). Since the value proposition of a prediction market is the elicitation and aggregation of opinions to determine the odds of future outcomes, traders must be able to sell their positions for a profit to to be able to speculate on share price.

References

[1] Yiling Chen and David M. Pennock. "Designing Markets for Prediction". In: *AI Magazine* 31.4 (Dec. 2010), pp. 42–52. DOI: 10.1609/aimag.v31i4.2313. URL: https://ojs.aaai.org/aimagazine/index.php/aimagazine/article/view/2313.

[2] Thaddeus Dryja. "Discreet Log Contracts". In: 2017. URL: https://api.semanticscholar.org/CorpusID:5522647.

[3] Lloyd Fournier. "One-Time Verifiably Encrypted Signatures A.K.A. Adaptor Signatures". In: 2019. URL: https://api.semanticscholar.org/CorpusID:212743751.

[4] Robin Hanson. "Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation". In: *Journal of Prediction Markets* 1 (May 2003). DOI: 10.5750/jpm.v1i1.417.

[5] Satoshi Nakamoto. "Bitcoin: A peer-to-peer electronic cash system". In: *Decentralized Business Review* (2008), p. 21260.

[6] Abraham Othman et al. "A practical liquidity-sensitive automated market maker". In: vol. 1. Dec. 2010, pp. 377–386. DOI: 10.1145/1807342.1807402.