Designing Expressive and Liquid Financial Options Markets via Linear Programming and Automated Market Making

Xintong Wang
Rutgers University
US
xintong.wang@rutgers.edu

David M. Pennock
Rutgers University
US
dpennock@dimacs.rutgers.edu

David M. Rothschild
Microsoft Research
US
davidmr@microsoft.com

Nikhil R. Devanur
Amazon
US
Iam@nikhildevanur.com

Abstract

We propose a two-stage market making strategy to improve the expressivity and liquidity of financial options markets, supporting investors to trade options of *any* strike price at the most competitive prices. In the first stage, we formulate a linear program (LP) to identify arbitrage opportunities and in the absence of arbitrage, find the tightest price bounds for a family of *option series* associated with the same security. Given a target option, the LP works by constructing a portfolio of outstanding options across different types and strikes to maximize immediate profit, while ensuring a nonnegative payoff in the future. By using the technique, a market maker can quote at any requested strike and provide liquidity to its maximum extent at no loss. Evaluation on real-market options data shows that the proposed LP can reduce bid-ask spreads significantly. In the second stage, through building a connection between options markets and prediction markets, we utilize a variant of the *constant-utility market maker* to further improve liquidity at a constant bounded loss. The loss can be considered as the market designer's subsidy to facilitate trading and elicit information. The market maker works by always quoting at risk-neutral prices to keep its utility constant, and thus can simultaneously recover the market's collective estimation on the probability distribution of the underlying security. We demonstrate on simulated options data that our proposed market maker can recover the option-implied probability distribution fairly well in comparison to alternative static optimization methods. We showcase several examples of probability distributions recovered from real-market options data, revealing the market's aggregated belief behind observed option prices.

ACM Reference Format:
Xintong Wang, David M. Pennock, David M. Rothschild, and Nikhil R. Devanur. 2024. Designing Expressive and Liquid Financial Options Markets via Linear Programming and Automated Market Making. In *5th ACM International Conference on AI in Finance (ICAIF '24), November 14–17, 2024, Brooklyn, NY, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3677052.3698687

ICAIF '24, November 14–17, 2024, Brooklyn, NY, USA

ACM ISBN 979-8-4007-1081-0/24/11
https://doi.org/10.1145/3677052.3698687

1 Introduction

Financial markets, broadly speaking, are places where people gather to trade assets and their derivatives. Depending on the ultimate purpose, markets are designed differently with outcomes interpreted accordingly. Traditional markets aim to match supply and demand, and price is considered the revealed value of the underlying asset. Prediction markets are designed to predict future events through eliciting and aggregating information. Speculators bet on an event by buying a contract for a one-dollar reward if the event occurs. Therefore, the market price reflects a collective belief in the event, often interpreted as a consensus probabilistic forecast [33].

With flavors of both, *financial options markets* provide investors the opportunities to trade on contracts that specify agreements upon potential future transactions. There are two basic types of options: *a call* and *a put* that gives the respective right to buy and sell an underlying asset at a given *strike price* and *expiration date*. An option is a *derivative* instrument, as its value is a function of the underlying asset. For example, an investor can buy a call option which specifies the right to buy 100 shares of the S&P 500 index for \$5000 each at the end of 2024. On the expiration date, the investor will exercise this right only if the S&P 500 is worth more than \$5000; meanwhile, the option seller who has *written* the option is obligated to sell at \$5000.

Many investors trade options to achieve a certain return pattern or hedge risks, but other investors, called speculators, act solely on their belief about the movement of the underlying asset price, buying an option when price falls below their estimate of its expected value. Thus, option prices reveal investors' collective risk-neutral belief distribution of the underlying asset's future price, which is often recovered for both academic and practical purposes [18]. This option-implied probability distribution has been further studied to infer systematic disaster concerns [19] and future inflation [28].

Options derived from the same underlying asset, of the same type, with the same strike price and expiration date are referred to as an *option series*. Current options market treats different option series separately, having each publicly traded in a distinct and independent continuous double auction (CDA). Consider options on a security offering both calls and puts, ten expiration dates and twenty strike prices. All possible option series render a total of

400 markets, among which investments get dispersed, even though participants are interested in the same security. This can often lead to the *thin market* problem, where few trades happen and bid-ask spreads become wide.

A *market maker* (MM) who maintains quotes on both buy and sell sides is often introduced to break the *no-trade reticence* and improve price discovery. However, falling trading volumes in options market are of raising concerns in recent years, as this drained liquidity can even spur market makers to retrench, trapping the market in a dangerous loop [12, 13]. Even for some of the most actively-traded option families, empirical evidence has shown that liquidity can vary much across option types and strikes. Cao and Wei (2010) studied eight years of options trading data and find consistently lower liquidity in puts and deep in-the-money options. As a result, catering to the interest of majority, exchanges (e.g., CBOE) offer options only at a limited number of strikes.

In this paper, we propose a two-stage market making strategy to initiate an expressive and liquid financial options market, where investors have the privilege to trade options at any strike price for competitive prices. In the first stage, we formulate an efficient linear program (LP) that can identify any arbitrage opportunity across a family of options associated with a single security, and derive the tightest possible pricing bound for options at arbitrary strike price. Given a specified option, the LP works by finding the minimum price of a portfolio with existing options that has a higher payoff than that of the target option, and the maximum price of a portfolio that has a lower payoff. By using such a technique, a market maker can consolidate options written on the same underlying security, quote at any requested strike price, and provide liquidity to its maximum extent *at no loss*, regardless of final outcomes.

In the second stage, by establishing a connection between options markets and prediction markets, we employ a *constant-utility market maker* to further improve liquidity. A constant-utility market maker has several intriguing properties: (1) It has a worst-case bounded loss, which can be considered as a preset subsidy to encourage trading and improve price discovery; (2) By quoting at risk-neutral prices to keep utility constant, it can simultaneously recover the market's risk-neutral probability distribution of the underlying security.

We evaluate the proposed LP on real-market options data of 30 stocks that compose the Dow Jones Industrial Average (DJI), and find it can substantially tighten the market bids and asks. This indicates that the market efficiency of less-traded option series can be improved by its well-discovered siblings. We further demonstrate on two simulated options data, one with exact options prices and one with noisily generated bids and asks, that a constant-log-utility market maker can recover option-implied distributions fairly well, comparing with other optimization methods. Examples of probability distributions implied from real-world options data suggest the proposed LP, when using to preprocess market bids and asks, can help recovering a smoother distribution.

2 Related Work

Rational Option Pricing. Arbitrage conditions have been studied extensively in financial economics, with the importance firstly addressed in corporation valuation [22]. Merton (1973) extend it to rational option pricing by proving the necessity of convexity in option prices. Other relevant works examine no arbitrage conditions to derive pricing bounds under different scenarios, including evaluating in discrete-time model [25], with the presence of zero-coupon bond [11] and risk-aversion [27].

Among those, the most relevant work to our approach in stage one is by Herzel (2005), where a LP approach is used to examine the existence of arbitrage by checking the convexity between every strike pair. Wang et al. (2021) propose the use of linear and mixed-integer linear programs to design matching functions for financial exchanges to facilitate trading standard options as well as *combinatorial financial options*. Our work here focuses on the role of a market maker, utilizing the LP approach to provide the most competitive price quotes, which further help in recovering a more accurate option-implied probability distribution of the underlying security.

Implied Probability from Options Prices. Two major approaches have been proposed to recover the option-implied probability distributions. Parametric methods assume an option pricing model, such as the Black-Scholes model [3] or some distribution family, with the observed prices inverted accordingly to obtain parameters of the distribution [2, 20]. However, stock market crashes has proved that the assumed pricing model together with parameters estimated from historical data usually fit market options prices poorly. Therefore, the nonparametric methods are promoted [18, 23]. It works by choosing probabilities to optimize an objective function, subject to the constraint that the chosen probabilities yield prices consistent with the observed market option prices. In order to get a feasible solution, it is necessary to preprocess the observed option prices and filter out any arbitrage violation. Our proposed constant-utility market making approach falls into the nonparametric methods. Instead of optimizing for expected utility, it keeps the utility constant and ensures a bounded loss while recovering the implied distribution.

Design of Prediction Markets. Interested in predicting future events, prediction markets are designed to recover probabilities via information aggregation. Though in a truly efficient market, prices will be the best predictor of the event and no other information can improve this aggregated forecast [33], many prediction markets are thinly traded with little revealed information. Much prior work dedicates to solving the problem by designing *automated market makers* to improve liquidity and aggregate information [1, 6, 7, 14, 15]. The market maker constantly quotes according to some designed scoring rules, and adjust the prices in response to the trading quantity. A market designer can subsidize the market maker to elicit and aggregate information, but would prefer a bounded-loss regardless of the final outcome. Some common techniques include the *logarithmic market scoring rule* market maker [15] and the constant-utility market maker [6], the framework we extend to options market in this work. Recent works examining decentralized finance has introduced an axiomatic framework that connects general *constant function market makers* (CFMMs), which form the core implementation of *Uniswap v2*, to cost-function-based prediction markets [26]. Strategic liquidity provision is studied in [5, 8, 9], where liquidity providers hold stochastic beliefs about the evolution of market prices, and how contract prices may change along with market prices via arbitrage and non-arbitrage trades.

3 Preliminaries

We denote *a call option* as $Call(S, K, T)$ and *a put option* as $Put(S, K, T)$, which respectively grants the option buyer the right to *buy* and *sell* an underlying security S at a specified *strike price* K on the expiration date T.[1] That is, the option buyer decides whether to exercise an option contract.

Suppose that a buyer spends \$18 and purchases a call option, $Call(S\&P\ 500, 5800, 20241231)$. If the *S&P* 500 is \$6000 at expiration (i.e., the last day of 2024), the buyer will pay the agreed strike \$5800, receive the index as cash, and get a *payoff* of \$200 and a *net profit* of \$182 (assuming no time value). If the *S&P* 500 is \$5600, the buyer will walk away without exercising the option. Therefore, the payoff of a purchased option is

$$\Psi := \max\{\chi(S-K), 0\} \tag{1}$$

where S is the value of underlying asset at expiration and $\chi \in \{-1, 1\}$ equals 1 for calls and -1 for puts. As the payoff for a buyer is always non-negative, the seller receives a *premium* now (e.g., \$18) to compensate for future obligations.

Options written on the same underlying, type, strike, and expiration are referred to as an *option series.* A standard option exchange (e.g., CBOE) treats different option series separately, having each publicly traded in a distinct and independent continuous double auction (CDA), though their values depend commonly on the underlying security. Investors can achieve a certain payoff pattern through constructing a *portfolio* of options — a combination of long and short positions of options with different strike prices.

Definition 1. Portfolio A is *(weakly) dominant* over portfolio B, if the payoff of A is larger than (or equal to) that of B for all possible states of the underlying asset in the future. Portfolio B is said to be *(weakly) dominated* by A.

A rational investor will always prefer A when the two portfolios have the same cost. However, if portfolio A is strictly cheaper than B, an *arbitrage* opportunity arises.

No Arbitrage Principle [31]. Portfolios which are guaranteed to have nonnegative payoffs must have a nonnegative cost.

4 Stage 1: A Linear Program to Remove Arbitrage and Improve Bids and Asks

In this section, we propose a linear program (LP) to remove any arbitrage opportunity and tighten market bids and asks, leveraging portfolio dominance and the interconnectedness of option markets written on the same underlying security. The model is simple without making assumptions on the option's pricing model and the stochastic behavior of underlying security. We further demonstrate that our proposed LP is computationally efficient, in which key market operations (i.e., arbitrage removal and price quotes) can be computed in polynomial time.

4.1 Problem Formulation

To use portfolio dominance, we consider all option series that relate to the same underlying security and expiration (e.g., call and put options on `DJI` across different strike prices). We represent an option series offered in the market as (χ_i, K_i, b_i, a_i), with $\chi_i \in \{-1, 1\}$ denoting its type (put or call), K_i its strike value, and (b_i, a_i) its corresponding best bid and ask in the market. Throughout this paper, we consider *European options* which can be exercised only at expiration. The settlement value is calculated as the opening value of the index on the expiration date or the last business day (usually a Friday) before the expiration date.

Remove Arbitrage. We aim to find arbitrage opportunities that may exist across option series. Specifically, the market maker (also an arbitrageur in this case) decides the fraction $\gamma_i \in [0, 1]$ to sell to each buy order at its bid b_i and the fraction $\delta_i \in [0, 1]$ to buy from each sell order at its ask a_i, with the objective of maximizing net profit at the time of order transaction subject to no payoff loss in the future at expiration for all possible states of the underlying security S:

$$\begin{aligned}
&\max_{\boldsymbol{\gamma},\boldsymbol{\delta}} \sum_i \gamma_i b_i - \sum_i \delta_i a_i && (2)\\
&\text{s.t.} \sum_i \gamma_i \max(\chi_i(S-K_i), 0) \le \sum_i \delta_i \max(\chi_i(S-K_i), 0)\\
&\qquad\qquad\qquad\qquad\qquad \forall S \in [0, \infty)
\end{aligned}$$

We denote options sold as Portfolio Γ and options bought as Portfolio Δ, and the constraint enforces that Portfolio Γ is *weakly dominated* by Portfolio Δ. In other words, it guarantees that regardless of the value of S at expiration, the liability of the market maker from sold options will not exceed the payoff gained from bought ones.

We note that as both sides of the constraint is a piecewise linear function of S, it suffices to solve LP (2) by satisfying constraints defined by S at each breakpoint. In our case, those breakpoints are the defined strike values in the market, plus two endpoints: $\boldsymbol{K} \cup \{0, \infty\}$. Therefore, the proposed LP has $|\boldsymbol{K}|+2$ payoff constraints and requires time polynomial in the size of the problem instance to solve.

4.2 Improving Price Quotes for Options

By utilizing LP (2), the market maker can improve overall market liquidity by facilitating trades among orders that may *not* match under current independent-market design. After all arbitrage opportunities are removed,[2] we can further use variants of the LP to improve the market bids and asks of existing options, and even construct price quotes for options at strike values that are not offered by the exchange. Specifically, we aim to derive the tightest possible price bounds using portfolio dominance, given observed market bids and asks on option series written on the same underlying security. Given an option at query, (χ, S, K, T), we would like to derive

- The lowest ask by finding *the minimum cost to buy a portfolio that weakly dominates* (χ, S, K, T)

$$\begin{aligned}
&\min_{\boldsymbol{\gamma},\boldsymbol{\delta}} \sum_i \delta_i a_i - \sum_i \gamma_i b_i && (3)\\
&\text{s.t.} \sum_i \delta_i \max(\chi_i(S-K_i), 0) - \sum_i \gamma_i \max(\chi_i(S-K_i), 0)\\
&\qquad \ge \max(\chi(S-K), 0) \qquad\qquad \forall S \in [0, \infty)
\end{aligned}$$

[1] We omit T from the tuples for simplicity when consider options of the same expiration.

[2] If there exists arbitrage violation, prices may be negative or unbounded.

- The highest bid by finding *the maximum profit to sell a portfolio that is weakly dominated by* (χ, S, K, T)

$$\max_{\gamma,\delta} \sum_i \gamma_i b_i - \sum_i \delta_i a_i \tag{4}$$
$$\text{s.t.} \sum_i \gamma_i \max(\chi_i(S-K_i),0) - \sum_i \delta_i \max(\chi_i(S-K_i),0) \leq \max(\chi(S-K),0) \quad \forall S \in [0,\infty)$$

Solution to the above LPs yields the best price quotes, without introducing arbitrage violation and any loss to the market maker regardless of the realization of S. Specifically, from an automated market maker's perspective, it offers an ask at the lowest feasible price; if the target option (χ, S, K) is sold, the market maker can always cover the short position by purchasing the weakly dominant portfolio in LP (3) with the ask it received. Similarly, the market maker quotes a bid at the highest possible price; if the option (χ, S, K) is bought, it can get rid of the long position by selling the portfolio with a lower payoff in LP (4) and getting the paid bid back. By using such a technique, the automated market maker can keep its position at zero, provide liquidity to the maximum possible extent, and ensure no loss regardless of the underlying security price at expiration. Below we provide two motivating examples based on real-market options data to illustrate the usefulness of the LPs and the improved price quotes.

Example 1 (Quote the best ask)**.** We use LP (3) to find the best asks for options of Dow Jones Industrial Average (DJI) that are priced on August 28, 2023 and expire on March 15, 2024. The best ask for Call(DJI, 350) can be improved from the market quote of $20 to $17.9, given by the following portfolio that weakly dominates Call(DJI, 350):

- Buy 0.5 Call(DJI, 335) at the market ask $26.6,
- Buy 0.5 Call(DJI, 365) at the market ask $9.2.

Example 2 (Quote the best bid)**.** We use LP (4) to find the best bids for options of Dow Jones Industrial Average (DJI) that are priced on August 28, 2023 and expire on March 15, 2024. The best bid for Call(DJI, 120) can be improved from the market quote of $214 to $223.10, given by the following portfolio that is weakly dominated by Call(DJI, 120):

- Sell 1.0 Call(DJI, 155) at the market bid $189.9
- Sell 0.125 Call(DJI, 160) at the market bid $185.1
- Sell 0.125 Put(DJI, 440) at the market bid $82.85
- Buy 0.125 Call(DJI, 440) at the market ask $0.25,
- Buy 1.0 Put(DJI, 160) at the market ask $0.26.

Remark. We note that solutions returned by the linear programs may involve fractional shares of options and stocks. Usually, we can naturally scale it to integer by incorporating the number of orders available for the best bids and asks. For fractional shares of stocks, ideally, an exchange that allows cash settlement can handle this. However, if the exchange requires physical settlement and does not allow fractional shares, we would need to normalize or round.

4.3 Empirical Evaluations

We evaluate the proposed LP(s) on real-market options data on DJI index and the 30 stocks that compose the DJI, as these stocks have actively traded options that cover a wide range of moneyness and

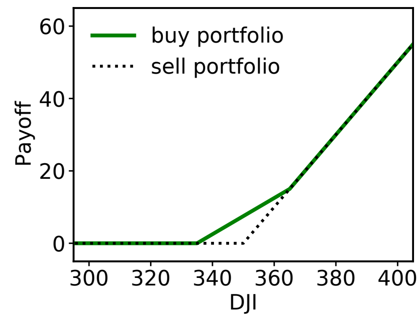

Figure 1: Payoffs of DJI options bought and sold in Example 1 that finds the lowest ask price for Call(DJI, 350).We zoom to $S_{\mathtt{DJI}} \in [300, 400]$ to visualize portfolio dominance.

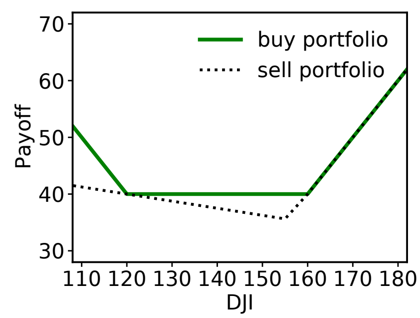

Figure 2: Payoffs of DJI options bought and sold in Example 2 that finds the highest bid price for Call(DJI, 120). We zoom to $S_{\mathtt{DJI}} \in [110, 180]$ to visualize portfolio dominance.

maturity levels. For each stock, we obtain contemporaneous closing bids and asks of all available option series on the OptionMetric dataset provided by the Wharton Research Data Services (WRDS). Market prices used for this experiment are dated August 28, 2023, at the close of market. There are a total of 28,638 distinct options markets for the 30 stocks[3] in DJI on August 28, 2023, covering around 14 expiration dates for each stock.

We use the proposed LP to find arbitrage opportunities and tighten price quotes by considering the outstanding buy orders and sell orders from independently-traded options markets that associated with the same security and expiration date. Out of a total of 356 such consolidated markets, the market maker spots 296 arbitrage opportunities, yielding an average profit of $2.28 (per arbitrage). Detailed statistics for options of each stock are available in the supplementary materials.[4]

For the other arbitrage-free markets, we find that the bid-ask spreads can be reduced from an average of 78 cents for each option series to 40 cents, by using the market maker to tighten the price quotes. Specifically, for option series written on DJI, one of the most actively traded financial options, the LP achieves a bid-ask spread improvement of 33%. These results show that the market maker, by considering orders on options markets across types and

[3]Our data includes American options that allow early exercise before expiration. In practice, American options are almost always more profitable to sell than to exercise early. In experiments, we ignore early exercise and treat them as European options.
[4]Supplementary materials can be found at https://shorturl.at/bqmjR

strike prices, can potentially help to achieve a higher economic efficiency, matching orders that the current independent design cannot and providing more competitive bid and ask prices.

5 Stage 2: A Constant-Utility MM for Options

A market maker using our Stage 1 LP(s) based on portfolio dominance provides liquidity without incurring any risk. To improve liquidity further, some risk of loss is necessary, and the loss (ideally bounded) can be viewed as the market designer's subsidy to facilitate trading and price discovery.

Therefore, in Stage 2, drawing inspiration from prediction market designs, we propose a bounded-loss automated market maker for the options market. The most popular prediction market maker adopts the *logarithmic market scoring rule*, but its loss can be unbounded, growing with the number of outcomes, so technically infinite in the case of stock prices [14, 15]. We thus consider a class of *constant utility market makers* [6], which achieves constant expected utility and bounded loss under common classes of utility functions, regardless of the number of outcomes. As the market maker trades, its *risk-neutral prices* reflect an option-implied probability distribution of the underlying stock price, consistent with the bid-ask bounds that we derive in Stage 1.

5.1 Background

We start by giving an overview of the *utility-based market makers*. Consider predicting a discrete random variable with N mutually exclusive and exhaustive outcomes (e.g., the opening price of `MSFT` on the last day of 2024, ranging from 0 to 600, discretized at intervals of 0.01).[5] The market maker offers to trade N securities, and for each share of security i sold, it pays \$1 to the security holder when outcome i happens. The market maker has a utility function for wealth across all outcome states, denoted $u(\mathbf{w})$, and offers to buy or sell an infinitesimal share of security i at price equal to its risk-neutral probability of state i [17], calculated as the normalized product of its subjective probability π_i and marginal utility $u'(w_i)$:

$$p_i = \frac{\pi_i u'(w_i)}{\sum_j \pi_j u'(w_j)}. \tag{5}$$

In other words, p_i is the instantaneous price for security i, at which the market maker is indifferent between buying and selling.

Throughout trading, the market maker maintains a quantity vector $\mathbf{q}$ to keep track of total shares held by traders across states. The quantity at state i, q_i, also indicates the amount the market maker must pay to traders if state i happens. To facilitate implementation, a *cost function* $C(\mathbf{q})$ is often introduced to record the total amount of money traders have spent as a function of $\mathbf{q}$. Therefore, the market maker's wealth at state i (i.e., w_i) is $C(\mathbf{q}) - q_i$, and the cost function is defined implicitly by:

$$\sum_i \pi_i u(w_i) = \sum_i \pi_i u(C(\mathbf{q}) - q_i) = U, \tag{6}$$

where U denotes the constant expected utility level, so called a *constant utility market maker*. When we have a continuous, differentiable, and monotonically increasing utility function $u(\cdot)$, which is typically the case, there exists a unique $C(\mathbf{q})$ by *implicit function theorem*; it can be found using binary search or other root-finding methods (if no explicit form). The cost of a trade that moves the quantity from $\mathbf{q}$ to $\mathbf{q'}$ is

$$C(\mathbf{q'}) - C(\mathbf{q}). \tag{7}$$

It is proven that a bounded loss is guaranteed when the market maker uses any utility function in the non-linear, *hyperbolic absolute risk aversion* (HARA) class [6, 24], which contains most popular parametric families of utilities, including logarithmic utility and negative exponential utility that we will later use for our market maker in options market.

[5] Though we cap the price at 600, we can treat the last outcome as "the opening price of `MSFT` is no less than 600".

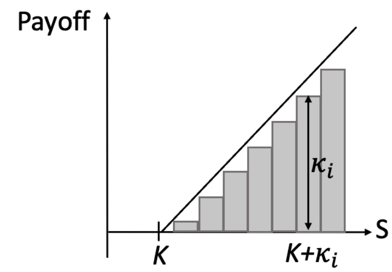

Figure 3: Approximating the payoff of a call option with a series of Arrow-Debreu securities.

5.2 Market Making to Recover Option-Implied Probability

In Stage 2, we aim to extend the constant-utility market maker to options market and use options prices to recover traders' collective estimation of state probabilities. We first approximate the payoff of an option with a series of discretized Arrow-Debreu securities. Figure 3 provides an illustration: A long position of a call option with a strike price K, i.e., $Call(S, K)$, is equivalent to buying $\kappa_i = S_i - K$ shares of security i for every state where $S_i > K$, betting that outcome i will happen or equivalently, the value of underlying stock at expiration is exactly S_i. Similarly, buying a put option at strike K is equivalent to buying $\kappa_i = K - S_i$ shares of Arrow-Debreu security i for every $S_i < K$.

Following the cost of a trade in Eq.(7), the cost of buying a call option at strike K is $C(\mathbf{q} + \boldsymbol{\kappa}) - C(\mathbf{q})$, which can be calculated by solving Eq.(6). With the setup, a market maker initialized with a subjective probability $\boldsymbol{\pi}$ (usually uniform distribution), a constant utility U, and a quantity vector $\mathbf{q} = \mathbf{0}$ can offer to trade options of *any type and strike* and achieve a worst-case loss of $C(\mathbf{0})$, regardless of the security's outcome at expiration. When trading, the market maker also facilitates price discovery by recovering prices in Eq.(5), which reflect the market's aggregated information about the underlying security (e.g., the probability that `MSFT` will have a specific value at option expiration).

Leveraging the payoff approximation, we utilize a constant-utility market maker to recover the option-implied probability distribution of the underlying stock price. We enforce that price quotes provided by the market maker align with the tightened bids and asks derived from market data in Stage 1. Algorithm 1 describes the market making procedure. Specifically, for every strike, we query the constant-utility market maker for its corresponding call and

Algorithm 1 Recover option-implied probabilities by trading with a constant-utility market maker.

Input: A market of N outcome events, indexed by i;
A vector of strike prices offered in the market: $\mathbf{K}$;
Best bids and asks on call and put options across all strikes: $(\mathbf{b^c}, \mathbf{a^c})$ and $(\mathbf{b^p}, \mathbf{a^p})$;
The utility function and constant utility: $u(\cdot)$ and U.
Output: Implied probability (price) for each state (security) i.

1: Initialize $\boldsymbol{\pi} \leftarrow \mathbf{1/N}, \mathbf{q} \leftarrow \mathbf{0}, \mathbf{q}' \leftarrow \mathbf{1}$
   cost $\leftarrow$ solution to $\sum_i \pi_i u(C(\mathbf{q})) = U$
2: **while** $\mathbf{q} \neq \mathbf{q}'$ **do** ▷ market is not at quiescence
3: $\mathbf{q}' \leftarrow \mathbf{q}$
4: **for** k in $\mathbf{K}$ **do**
5: $\boldsymbol{\kappa} \leftarrow \max(\mathbf{S} - \mathbf{k}, \mathbf{0})$ ▷ check call options
6: $c \leftarrow$ solution to $\sum_i \pi_i \log(C(\mathbf{q} + \boldsymbol{\kappa}) - q_i - \kappa_i) = U$
7: $\text{call}_{\text{buy}} \leftarrow c -$ cost ▷ MM quoted call buy
8: **if** $b^c_k \geq \text{call}_{\text{buy}}$ **then**
9: $\mathbf{q} \leftarrow \mathbf{q} + \boldsymbol{\kappa}$, cost $\leftarrow c$
10: **else**
11: $c \leftarrow$ solution $\sum_i \pi_i \log(C(\mathbf{q} - \boldsymbol{\kappa}) - q_i + \kappa_i) = U$
12: $\text{call}_{\text{sell}} \leftarrow$ cost $- c$ ▷ MM quoted call sell
13: **if** $a^c_k \leq \text{call}_{\text{sell}}$ **then**
14: $\mathbf{q} \leftarrow \mathbf{q} - \boldsymbol{\kappa}$, cost $\leftarrow c$
15: $\boldsymbol{\kappa} \leftarrow \max(\mathbf{k} - \mathbf{S}, \mathbf{0})$ ▷ check put options
16: $c \leftarrow$ solution to $\sum_i \pi_i \log(C(\mathbf{q} + \boldsymbol{\kappa}) - q_i - \kappa_i) = U$
17: $\text{put}_{\text{buy}} \leftarrow c -$ cost ▷ MM quoted put buy
18: **if** $b^p_k \geq \text{put}_{\text{buy}}$ **then**
19: $\mathbf{q} \leftarrow \mathbf{q} + \boldsymbol{\kappa}$, cost $\leftarrow c$
20: **else**
21: $c \leftarrow$ solution $\sum_i \pi_i \log(C(\mathbf{q} - \boldsymbol{\kappa}) - q_i + \kappa_i) = U$
22: $\text{put}_{\text{sell}} \leftarrow$ cost $- c$ ▷ MM quoted put sell
23: **if** $a^p_k \leq \text{put}_{\text{sell}}$ **then**
24: $\mathbf{q} \leftarrow \mathbf{q} - \boldsymbol{\kappa}$, cost $\leftarrow c$
25: **return** $p_i = \frac{\pi_i u'(C(\mathbf{q}) - q_i)}{\sum_j \pi_j u'(C(\mathbf{q}) - q_j)} \quad \forall i$

put prices by updating $\mathbf{q}$ and solving for C as in Eq. (6) (lines 4-24). When the market maker's quoted buy (sell) price is lower (higher) than the market bid (ask), the trading agent would buy from (sell to) the constant-utility market maker (lines 8, 13, 18, 23). As the quoted price results in a trade, the market maker will update the cost and quantity accordingly. To recover the option-implied probability distribution, trades will continue until quiescence where the market maker's quoted prices fall in the range of market bids and asks across all strikes, and the risk-neutral prices as calculated in Eq. (5) are the recovered probability for each state (line 25).

5.3 Empirical Evaluation

We evaluate the option-implied probability distributions recovered by a constant-utility market maker that adopts the logarithmic utility function. We first conduct experiments on simulated options data, where a ground-truth distribution is set and known.[6] We are interested in the market maker's performance on options data with

[6]For real-market options data, the ground-truth belief distribution of the population is unknown.

Data type	Objective function	KL-Divergence
Exact Prices	Constant-log-utility MM	0.0025
	Quadratic	0.0435
	Goodness of fit	0.0129
	Max entropy	0.0024
Bids and Asks	Constant-log-utility MM	0.0104
	Quadratic	0.0666
	Goodness of fit	0.0180
	Max entropy	0.0065

Table 1: Average KL-divergence between probability distributions recovered by different approaches and the ground truth distributions.

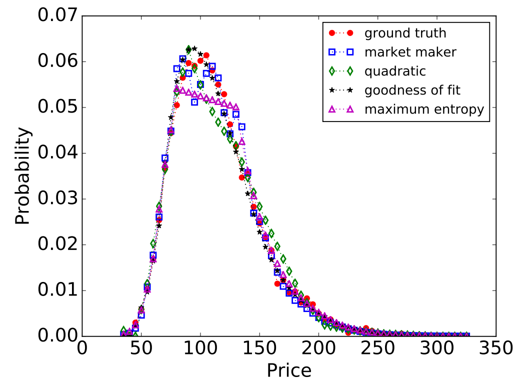

Figure 4: Option-implied probability distributions recovered by different approaches from simulated bids and asks with $\mu = 0.15$ **and** $\sigma = 0.3$**.**

(1) exact prices and with (2) noisily generated bids and asks around the exact prices. On both datasets, we compare our market making approach with common optimization methods using different objective functions. Finally, we demonstrate price distributions recovered by the constant-log-utility market maker from real-world options data across time, and qualitatively compare the distributions implied from observed market data and from the LP-tightened bids and asks in Stage 1.

5.3.1 Simulated Options Data. We generate exact option prices by first sampling stock outcomes, and then calculating the expected payoff across different strikes. Specifically, for each run, we simulate a stock with current price S_0 at \$100, an expected return of $\mu \sim N(0.15, 0.02)$ and a volatility of $\sigma \sim N(0.30, 0.05)$ per annum. We sample 10,000 one-year-later stock outcomes from the corresponding stochastic process and bin them into strike prices every \$5 apart to get the ground truth distribution. Both call and put option prices are calculated at each strike as its expected payoff given the sampled stock outcomes.

We further generate bids and asks around the exact price. Since the width of a bid-ask spread is highly correlated with the absolute value of the price, we fit the proportional bid-ask ratio—the ratio of absolute difference to the sum—from real options data to further generate bids and asks around exact option prices. We group options into *in-the-money*, *around-the-money* and *out-the-money*

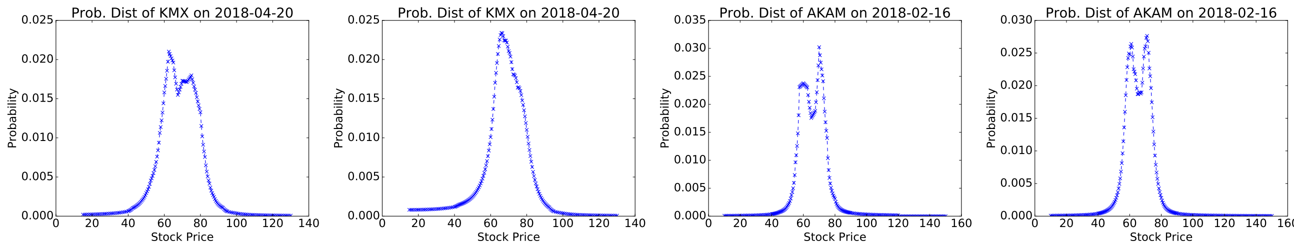

(a) Market bids and asks. **(b) LP-tightened bids and asks.** **(c) Market bids and asks.** **(d) LP-tightened bids and asks.**

Figure 5: Comparisons of option-implied probability distributions recovered from market option prices on December 26, 2017 and from LP tightened bids and asks.

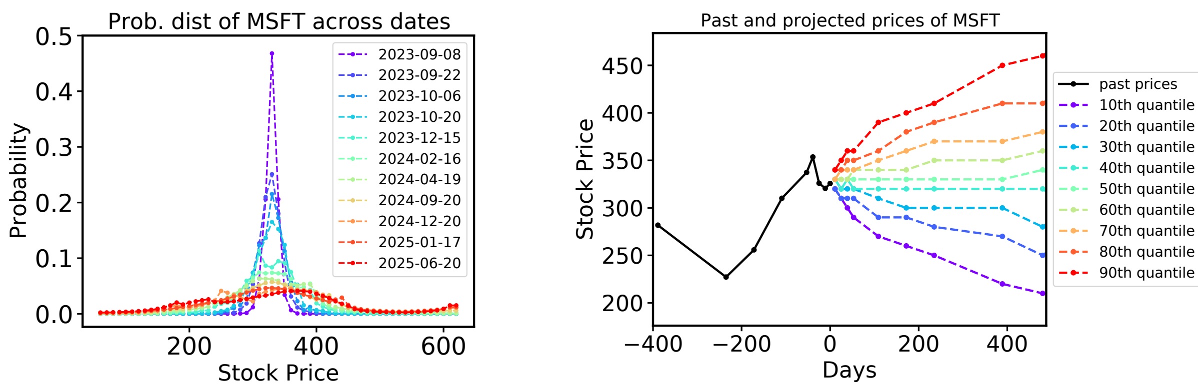

(a) Option-implied distributions across expiration dates. **(b) Past prices (negative Days) and projected future price quantiles (positive Days) from recovered distributions.**

Figure 6: Different representations of recovered probability distributions across expirations dates. Observed option prices are dated August 28, 2023 at the close of market.

options, and assign them a ratio of 0.03, 0.07 and 0.25 respectively.[7] Gaussian noise is further added to the exact bids and asks. We run the proposed LP over generated bids and asks across all strikes to ensure no arbitrage violation, so that a risk-neutral distribution can be recovered.

We compare the market making approach to optimization methods that minimize the following objective functions, where P' indicates a prior required by certain methods:

- Quadratic: $\sum_j (P_j - P'_j)^2$
- Goodness of fit: $\sum_j (P_j - P'_j)^2 / P'_j$
- Maximum Entropy: $\sum_j P_j \log(P_j)$

For the case with bids and asks, each method needs to satisfy the constraints that expected option payoffs calculated from the recovered distribution are consistent with the generated bids and asks.

Table 1 reports the average KL-divergence between the recovered and the ground truth distribution over 500 simulations of each method across the two datasets. As expected, the constant-log-utility market maker can recover a better distribution from exact option prices, as in the case of noisy bids and asks the quoted price can stop anywhere within the range. On both datasets, it statistically significantly outperforms optimization methods with quadratic and goodness-of-fit objective functions. Figure 4 further displays an example of option-implied distributions of different approaches from noisily generated bids and asks with $\mu = 0.15$ and $\sigma = 0.3$. Despite different priors, all of the recovered distributions exhibit fairly consistent shapes of posterior.

[7]This is a rough estimation of proportional bid-ask ratio for simulation purpose. A more rigorous model can fit the spread as a function of the corresponding strike price and option price.

5.3.2 Real-Market Options Data. Finally, we demonstrate implied probability distributions recovered from real-market options data. Figure 5 provides some representative examples to qualitatively compare distributions recovered from raw market option prices and from the LP-tightened bids and asks in Stage 1. We notice the distribution recovered from tightened bids and asks are usually smoother. This may be explained by the proposed LP reducing much of the noise in observed market bids and asks, leaving more precise bounds on prices. Interestingly, we also find many option-implied distributions are bimodal or even multi-modal, a phenomenon also reported in other studies [10, 29, 30]. Plausible explanations for this multi-modality include heterogeneity in investor's beliefs, and jumps and correlation between volatility and returns.

Figure 6 shows the implied distributions recovered by the constant-log-utility market maker (with $U = 10$) from the LP-tightened bid and ask prices of MSFT options observed on August 28, 2023 from WRDS. The implied distribution becomes flatter with larger variance, representing a higher level of uncertainty, as we look further ahead to future expiration dates (Fig. 6(a)). Fig. 6(b) provides a visual representation of Microsoft's stock prices over a span of 800 days, split into historical and projected periods. The quantiles summarize what is implied by hundreds of options prices at every strike across expiration date in one option family.

6 Conclusion

We introduced a two-stage market making strategy to improve expressivity and liquidity in financial options markets, offering investors the opportunity to trade options at any desired strike for the most competitive prices. In the first stage, we formulate an efficient linear program that can identify arbitrage opportunities across the family of options associated with a single security. In the absence of arbitrage, the LP finds the tightest possible bid-ask spread for existing options, and constructs pricing bounds for non-existing options at arbitrary strike prices. The LP works by exploiting the pricing information embodied in all options and finds a pricing range where no arbitrage opportunity arises. By using such a technique, a market maker can accept trades of options on arbitrary strike prices, consolidate distinct and independent options markets interested in the same security and provide liquidity to its maximum possible extent at no loss, regardless of final outcomes.

In the second stage, by building the connection between options markets and prediction markets, we adopt a constant-utility market maker to further provide liquidity at a bounded loss. By continuously updating its risk-neutral prices according to market activities, the constant-utility market maker can simultaneously recover the market's risk-neutral probability distribution for the underlying security that is consistent with the tight pricing bounds derived in the first stage. We further demonstrate on two simulated options data, one with exact options prices and one with noisily generated bids and asks, that a constant log utility market maker can recover option-implied distributions fairly precisely. Examples of distributions implied from real world options data suggest the proposed LP, when using to preprocess market bids and asks, can help recovering a smoother distribution.

References

[1] Jacob Abernethy, Yiling Chen, and Jennifer Wortman Vaughan. 2011. An Optimization-Based Framework for Automated Market-Making. In *12th ACM Conference on Electronic Commerce (EC)*.

[2] Bhupinder Bahra. 1997. Implied Risk-Neutral Probability Density Functions from Option Prices: Theory and Application. *Bank of England Working Paper No 66. Available at SSRN:77429* (1997). http://dx.doi.org/10.2139/ssrn.77429

[3] Fischer Black and Myron Scholes. 1973. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* 81, 3 (1973), 637–654. http://www.jstor.org/stable/1831029

[4] Melanie Cao and Jason Wei. 2010. Option market liquidity: Commonality and other characteristics. *Journal of Financial Markets* 13, 1 (2010), 20 – 48.

[5] Álvaro Cartea, Fayçal Drissi, and Marcello Monga. 2024. Decentralized Finance and Automated Market Making: Predictable Loss and Optimal Liquidity Provision. *SIAM Journal on Financial Mathematics* 15, 3 (2024), 931–959.

[6] Yiling Chen and David M. Pennock. 2007. A utility framework for bounded-loss market makers. In *Proceedings of the 23rd Annual Conference on Uncertainty in Artificial Intelligence (UAI)*. 49–56.

[7] Miroslav Dudík, Xintong Wang, David M. Pennock, and David M. Rothschild. 2021. Log-time Prediction Markets for Interval Securities. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 465–473.

[8] Zhou Fan, Francisco Marmolejo-Cossio, Daniel Moroz, Michael Neuder, Rithvik Rao, and David C. Parkes. [n. d.]. Strategic Liquidity Provision in Uniswap V3. In *5th Conference on Advances in Financial Technologies (AFT 2023)*. 25:1–25:22.

[9] Zhou Fan, Francisco J. Marmolejo-Cossío, Ben Altschuler, He Sun, Xintong Wang, and David Parkes. 2022. Differential Liquidity Provision in Uniswap v3 and Implications for Contract Design. In *Proceedings of the Third ACM International Conference on AI in Finance* (New York, NY, USA) *(ICAIF '22)*. Association for Computing Machinery, New York, NY, USA, 9–17.

[10] Matthias R. Fengler and Lin-Yee Hin. 2015. Semi-nonparametric estimation of the call-option price surface under strike and time-to-expiry no-arbitrage constraints. *Journal of Econometrics* 184, 2 (2015), 242 – 261.

[11] Mark B. Garman. 1976. An algebra for evaluating hedge portfolios. *Journal of Financial Economics* 3, 4 (1976), 403–427.

[12] Banerji Gunjan. 2017. Traders Are Fleeing the Options Market. *Wall Street Journal* (2017).

[13] Banerji Gunjan. 2018. Caution Hits Options Markets as LIquidity Fears Flare. *Barron's* (2018).

[14] Robin Hanson. 2003. Combinatorial information market design. *Information Systems Frontiers* 5, 1 (2003), 107–119.

[15] Robin Hanson. 2007. Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation. *Journal of Prediction Markets* 1, 1 (2007), 1–15.

[16] Stefano Herzel. 2005. Arbitrage opportunities on derivatives: A linear programming approach. *Dynamics of Continuous, Discrete and Impulsive Systems Series B: Application and Algorithms* (2005).

[17] Jens Carsten Jackwerth. 2000. Recovering Risk Aversion from Option Prices and Realized Returns. *The Review of Financial Studies* 13, 2 (2000), 433–451. https://doi.org/10.1093/rfs/13.2.433

[18] Jens Carsten Jackwerth and Mark Rubinstein. 1996. Recovering Probability Distributions from Option Prices. *The Journal of Finance* 51, 5 (1996), 1611–1631.

[19] Fang Liu. 2018. Option-Implied Systematic Disaster Concern. *Available at SSRN 2915906* (2018). https://ssrn.com/abstract=2915906

[20] William R. Melick and Charles P. Thomas. 1997. Recovering an Asset's Implied PDF from Option Prices: An Application to Crude Oil during the Gulf Crisis. *The Journal of Financial and Quantitative Analysis* 32, 1 (1997), 91–115.

[21] Robert C. Merton. 1973. Theory of Rational Option Pricing. *The Bell Journal of Economics and Management Science* 4, 1 (1973), 141–183.

[22] Franco Modigliani and Merton H. Miller. 1958. The Cost of Capital, Corporation Finance and the Theory of Investment. *The American Economic Review* 48, 3 (1958), 261–297.

[23] Ana Margarida Monteiro, Reha H. Tütüncü, and Luís N. Vicente. 2008. Recovering risk-neutral probability density functions from options prices using cubic Splines and ensuring nonnegativity. *European Journal of Operational Research* 187, 2 (2008), 525–542.

[24] Abraham Othman and Tuomas Sandholm. 2012. Automated Market Makers That Enable New Settings: Extending Constant-Utility Cost Functions. In *Auctions, Market Mechanisms, and Their Applications*, Peter Coles, Sanmay Das, Sébastien Lahaie, and Boleslaw Szymanski (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 19–30.

[25] Stylianos Perrakis and Peter J. Ryan. 1984. Option Pricing Bounds in Discrete Time. *The Journal of Finance* 39, 2 (1984), 519–525.

[26] Maneesha Papireddygari Rafael Frongillo and Bo Waggoner. 2024. An axiomatic characterization of CFMMs and equivalence to prediction markets.. In *Proceedings of the 15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*.

[27] Peter H. Ritchken. 1985. On Option Pricing Bounds. *The Journal of Finance* 40, 4 (1985), 1219–1233.

[28] Tom Smith. 2012. Option-Implied Probability Distributions for Future Inflation. *Bank of England Quarterly Bulletin 2012 Q3. Available at SSRN: 2150092* (2012). https://ssrn.com/abstract=2150092

[29] Marco Taboga. 2016. Option-implied probability distributions: How reliable? How jagged? *International Review of Economics & Finance* 45 (2016), 453 – 469.

[30] Stephen J. Taylor, Chi Feng Tzeng, and Martin Widdicks. 2012. Bankruptcy Probabilities Inferred from Option Prices. *25th Australasian Finance and Banking Conference. Available at SSRN 2135340* (2012).

[31] Hal R. Varian. 1987. The Arbitrage Principle in Financial Economics. *The Journal of Economic Perspectives* 1, 2 (1987), 55–72.

[32] Xintong Wang, David M. Pennock, Nikhil R. Devanur, David M. Rothschild, Biaoshuai Tao, and Michael P. Wellman. 2021. Designing a Combinatorial Financial Options Market. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*. 864–883.

[33] Justin Wolfers and Eric Zitzewitz. 2004. Prediction Markets. *Journal of Economic Perspectives* 18, 2 (2004), 107–126.