

Tighter Robust Upper Bounds for Options via No-Regret Learning

Shan Xue,¹ Ye Du,^{2*} Liang Xu¹

¹ School of Business Administration, Southwestern University of Finance and Economics, Chengdu, China

² Southwestern University of Finance and Economics, Chengdu, China

xueshanads123@gmail.com, henry,duye@gmail.com, arecxuliang1112@gmail.com

Abstract

Classic option pricing models, such as the Black-Scholes formula, often depend on some rigid assumptions on the dynamics of the underlying asset prices. These assumptions are inevitably violated in practice and thus induce the model risk. To mitigate this, robust option pricing that only requires the no-arbitrage principle has attracted a great deal of attention among researchers. In this paper, we give new robust upper bounds for option prices based on a novel η -momentum trading strategy. Our bounds for European options are tighter for most common money, volatility, and expiration date setups than those presented in (DeMarzo, Kremer, and Mansour 2016). Our bounds for average strike Asian options are the first closed-form robust upper bounds for those options. Numerical simulations demonstrate that our bounds significantly outperform the benchmarks for both European and Asian options.

Introduction

Option pricing has long been one of the most intriguing problems in finance. An option is a financial contract that allows its holder to purchase or sell a given underlying asset, such as a stock, for a predetermined price on or before a predetermined date. The price and the date in the contract are known as the strike price and the expiration date, respectively. For instance, a European call (or put) option on a stock with strike price K and expiration date T gives its holder the right to buy (or sell) the stock for a price K at time T . Consequently, the European call (or put) option has a payoff of $\max(S_T - K, 0)$ (or $\max(K - S_T, 0)$) at time T , where S_T is the stock price at time T . Besides the standard European options, various customized exotic options are traded in the over-the-counter market. Options are widely used for hedging against potential price fluctuations of the underlying asset as well as speculation. According to the webpage of Chicago Board Options Exchange, the largest U.S. options exchange, more than 1.5 billion stock options have been traded in 2021.¹

The famous Nobel Prize-winning Black-Scholes formula (Black and Scholes 1973), which is based on the principle of no-arbitrage pricing and risk-neutral probability, was a milestone in the literature on option pricing. An implicit assumption in the Black-Scholes framework is that the underlying asset price follows a geometric Brownian motion, which is quite rigid. With this model setting, the stochastic calculus can be utilized to obtain the explicit pricing formulas for European options. To make the option pricing theory work in practice, the follow-up studies attempt to reduce or weaken the assumptions a pricing model depends on. In one line of existing research on option pricing, numerous models are introduced to deal with more general dynamics of the underlying assets, such as the Heston model (Heston 1993), the Bates model (Bates 1996), the jump-diffusion model (Kou 2002), among others. In the other line of the research, model-independent bounds are developed for various options, such as (Hobson 1998; Brown, Hobson, and Rogers 2001; Chen and Yeh 2002; Hobson, Laurence, and Wang 2005a, 2005b; Chung and Chang 2007; Chen et al. 2008; Laurence and Wang 2008; Laurence and Wang 2009; Hobson and Neuberger 2012; Kahalé 2016, 2017). DeMarzo, Kremer, and Mansour (2006) first introduce the techniques of no-regret learning to option pricing and provide upper bounds for European call options. Following their work, Gofer and Mansour (2011) apply the methodology to exotic option pricing. DeMarzo et al. (2016) exploit a new gradient trading strategy to derive upper bounds for European options. Their results are robust in that they only depend on the no-arbitrage principle. Based on the same trading strategy, Du, Xue, and Liu (2019) and Xue et al. (2022) develop robust upper bounds for American options and various exotic options, respectively. Since robust option pricing only requires the no-arbitrage principle, the robust bounds are usually not very tight. Therefore, in this study, we investigate a novel no-regret learning strategy, which we call the η -momentum trading strategy, to derive tighter robust upper bounds for options.

Pricing and Learning. In the classic finance literature, the seller (writer) of an option will use the underlying asset to hedge her risk (i.e., use the underlying asset to replicate the option). The cost of this replication process is actually the price of the option. This dynamic hedging process can be exactly reframed as a learning process. The intuition is that the

*Corresponding author.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹https://www.cboe.com/us/options/market_statistics (last accessed on November 27, 2022).

more loss incurred by the seller, the more hedging position she will take. Based on this simple observation, DeMarzo et al. (2006) invent a dramatically different approach to price a European call option from the classic finance literature. The basic idea comes from no-regret learning, a well-studied technique from machine learning. Consider an investor with an initial capital I , she has a dynamic strategy to allocate the capital between a stock and a bond during the subsequent periods. If her final capital at time T is no less than the payoff of the option on the stock plus the strike price K for any possible scenario of the stock price, the option price is bounded by $I - K$ from above based on the no-arbitrage principle. DeMarzo et al. (2006) provide an upper bound for the option, which only depends on two variables: the maximal sum of squared single-period return of the stock over N periods Q_T and the maximal absolute return in each period M . Although this result is groundbreaking, it has two drawbacks: On one hand, the two variables Q_T and M are not directly observable in the market. In order to make this work in practice, some estimation methods must be introduced for Q_T and M . On the other hand, in their model, M is assumed to be less than $1 - \frac{\sqrt{2}}{2}$, which is impractical. For instance, in the classic geometric Brownian motion model, the stock return in any time period follows a normal distribution. This implies that the maximal return in any period could potentially be infinitely large, which contradicts the assumption of $M < 1 - \frac{\sqrt{2}}{2}$. Gofer and Mansour (2011) study the pricing of exotic options by extending the methodology of (DeMarzo et al. 2006). Unfortunately, their results inherit the two drawbacks of the method in (DeMarzo et al. 2006). In the work of (DeMarzo et al. 2016), the assumption on the maximal possible absolute return of the stock in each period is dropped.

Contributions. Our contributions are threefold. Foremost, we exploit a novel trading strategy: the η -momentum strategy, in which the potential function is an exponential function with η as the exponent. We derive an upper bound on the regret of the η -momentum strategy. Based on that, new robust upper bounds for the prices of European call options, which is a function of moneyness, the maximum quadratic variation of the stock return paths $q^2(\Phi_T)$ and η , are achieved. Our optimal bound for the at-the-money option is in closed form and strictly tighter than the result presented in (DeMarzo et al. 2016). Although the optimal bounds in our framework do not have closed-form solutions for general cases, we prove that they are theoretically tighter for most common moneyness, volatility and expiration date setups than the benchmarks. This implies that by taking a specific value of η , our framework admits suboptimal closed-form upper bounds that are still superior.

Furthermore, we extend the methodology to derive robust and closed-form upper bounds for average strike Asian call options. This is the first result of its kind to the best of our knowledge. Although Xue et al. (2022) provide robust upper bounds for various exotic options, they do not cover average strike Asian options. We show that our bounds are strictly tighter than the corresponding results in (Gofer and Mansour 2011) when $q^2(\Phi_T) \leq Q_T$. This implies that in the

standard binomial-tree model, our bounds are guaranteed to be superior.

Finally, we investigate the performance of our bounds by extensive numerical simulations. Consistent with the theoretical results, both of our bounds for European options and Asian options are significantly tighter than those benchmarks for most cases. In particular, our bounds for European options outperform the benchmarks better when they are out of the money than in the money. More interestingly, since we do not have closed-form solutions for the optimal bounds, numerical results suggest that our bounds have even better performance than the predictions of our theoretical results.

Related Literature. No-regret learning is an important model studied in the interdisciplinary field of computer science, game theory and statistics (Hannan 1957; Foster and Vohra 1999; Cesa-Bianchi and Lugosi 2006; Blum and Mansour 2007). Blackwell (1956) studies the approachability problem that is equivalent to no-regret learning in repeated games. It is well known that when players follow no-regret strategies in games, the empirical frequency of their plays converges to correlated equilibrium (Foster and Vohra 1997; Hart and Mas-Colell 2000). The idea of no-regret learning is also closely related to the concept of boosting in machine learning (Freund and Schapire 1996). Recently, the model of no-regret learning is further extended to the field of online convex optimization (Shalev-Shwartz 2011).

DeMarzo et al. (2006) are the first to study option pricing from the perspective of no-regret learning. Abernethy, Frongillo, and Wibisono (2012) consider the option pricing through an online learning game between Nature and an Investor, and show that the minimax option price converges to the price under the Black-Scholes model in the limit. Their results are further proved by (Abernethy et al. 2013) under much weaker assumptions.

The η -monmomentum Trading Strategy

Consider a discrete-time trading model with N periods from time 0 to time T , where a period is indexed by $n \in \{0, 1, \dots, N\}$ with equal length $\Delta t = \frac{T}{N}$. There are two basic assets: a stock and a bond, the prices of which at the end of the n th period are S_n and B_n , respectively. Denote the stock return in the n th period as $r_{n,s} = \frac{S_n}{S_{n-1}} - 1$ and its log return as $\pi_{n,s} = \ln(1 + r_{n,s})$. Let $R = \{r_{1,s}, r_{2,s}, \dots, r_{N,s}\}$ be a return path of the stock, and the set of all possible paths is denoted by ϕ_T . Denote by r_b the bond return in each period such that $B_n = B_{n-1}(1 + r_b)$. For simplicity, we assume that $B_0 = 1$ and $r_b = 0$, which yields $B_n = 1$ for all n .

There are two reference investment strategies: Strategy **I** and Strategy **II**, which are self-financing trading strategies on the stock and the bond. For each Strategy j with $j = \mathbf{I}$ or $\mathbf{II}$, $V_{0,j}$ is the initial capital invested in the strategy and its value is $V_{n,j}$ at the end of the n th period. Denote by $r_{n,j}$ the return of Strategy j in the n th period where $r_{n,j} = \frac{V_{n,j}}{V_{n-1,j}} - 1$, and its corresponding log return is $\pi_{n,j} = \ln(1 + r_{n,j})$.

An investor in this model manages a dynamic portfolio consisting of the stock and the bond. The portfolio has an initial value $G_0 = 1$ and its value at the end of the n th period is G_n . The investor aims at making the value of the portfolio

comparable to the ex-post best of the two reference strategies at time T . To accomplish the objective, She utilizes a no-regret learning strategy, which we call the η -momentum trading strategy, to allocate her capital between the two reference strategies. To illustrate the strategy, we first define a regret vector $L_n = (\ln \frac{x_1 V_{n,I}}{V_{0,I}} - \ln G_n, \ln \frac{x_2 V_{n,II}}{V_{0,II}} - \ln G_n)$, where $\ln x = (\ln x_1, \ln x_2)$ is a parameter representing an initial fictitious loss, and is used to optimize the performance of the dynamic portfolio. At the beginning of the n th period, the η -momentum trading strategy invests $w_{n,I}$ portion of G_{n-1} in Strategy **I** and $w_{n,II} = 1 - w_{n,I}$ portion in Strategy **II**, where $w_{n,I}$ is specified as

$$w_{n,I} = \frac{e^{\eta L_{n-1,1}}}{e^{\eta L_{n-1,1}} + e^{\eta L_{n-1,2}}} = \frac{(\frac{x_1 V_{n-1,I}}{V_{0,I}})^\eta}{(\frac{x_1 V_{n-1,I}}{V_{0,I}})^\eta + (\frac{x_2 V_{n-1,II}}{V_{0,II}})^\eta}.$$

The positive parameter η in the trading strategy is used to optimize the performance of the dynamic portfolio as well. The regret bound of the η -momentum trading strategy is summarized in the following Theorem.

Theorem 1. *The η -momentum trading strategy with initial loss $\ln x$ guarantees*

$$\max(\ln \frac{x_1 V_{N,I}}{V_{0,I}}, \ln \frac{x_2 V_{N,II}}{V_{0,II}}) - \ln G_N \leq \frac{\ln(x_1^\eta + x_2^\eta)}{\eta} + \frac{\eta q^2(R)}{8},$$

$$\text{where } q^2(R) = \sum_{n=1}^N (\pi_{n,I} - \pi_{n,II})^2.$$

Proof. We first define $W_n = (\frac{x_1 V_{n,I}}{V_{0,I}})^\eta + (\frac{x_2 V_{n,II}}{V_{0,II}})^\eta$ and $\bar{\pi}_n = w_{n,I} \pi_{n,I} + w_{n,II} \pi_{n,II}$. Direct calculation verifies

$$\begin{aligned} \ln \frac{W_n}{W_{n-1}} &= \ln \frac{(\frac{x_1 V_{n,I}}{V_{0,I}})^\eta + (\frac{x_2 V_{n,II}}{V_{0,II}})^\eta}{(\frac{x_1 V_{n-1,I}}{V_{0,I}})^\eta + (\frac{x_2 V_{n-1,II}}{V_{0,II}})^\eta} \\ &= \ln[w_{n,I} (\frac{V_{n,I}}{V_{n-1,I}})^\eta + w_{n,II} (\frac{V_{n,II}}{V_{n-1,II}})^\eta] \\ &= \ln(\sum_{i \in \{I, II\}} w_{n,i} e^{\eta \pi_{n,i}}) \\ &= \eta \bar{\pi}_n + \ln[\sum_{i \in \{I, II\}} w_{n,i} e^{\eta(\pi_{n,i} - \bar{\pi}_n)}] \\ &= \eta \bar{\pi}_n + \ln[w_{n,I} + w_{n,II} e^{\eta(\pi_{n,II} - \pi_{n,I})}] \\ &+ \eta(\pi_{n,I} - \bar{\pi}_n) \\ &= \eta \bar{\pi}_n + \ln[w_{n,I} + (1 - w_{n,I}) e^{\eta(\pi_{n,II} - \pi_{n,I})}] \\ &+ \eta(1 - w_{n,I})(\pi_{n,I} - \pi_{n,II}). \end{aligned}$$

Denote by $\pi_{n,g}$ the log return of the dynamic portfolio adjusted by the η -momentum trading strategy such that $\pi_{n,g} = \ln[1 + w_{n,I} r_{n,I} + w_{n,II} r_{n,II}]$. From the concavity of the $\ln$ function, we have

$$\bar{\pi}_n = w_{n,I} \pi_{n,I} + w_{n,II} \pi_{n,II} \leq \pi_{n,g}.$$

Define $f(s) := \ln[w_{n,I} + (1 - w_{n,I})e^{-\eta s}] + \eta(1 - w_{n,I})s$. By straightforward calculation we have

$$f'(s) = -\frac{\eta(1 - w_{n,I})e^{-\eta s}}{w_{n,I} + (1 - w_{n,I})e^{-\eta s}} + \eta(1 - w_{n,I}),$$

and

$$f''(s) = \frac{\eta^2 w_{n,I} (1 - w_{n,I}) e^{-\eta s}}{[w_{n,I} + (1 - w_{n,I})e^{-\eta s}]^2} \leq \frac{\eta^2}{4}.$$

Since $f(0) = f'(0) = 0$, by Taylor's theorem there exists some θ lying between 0 and s such that

$$f(s) = f(0) + f'(0)s + \frac{f''(\theta)}{2} s^2 = \frac{f''(\theta)}{2} s^2 \leq \frac{\eta^2 s^2}{8}.$$

It follows that

$$\ln \frac{W_n}{W_{n-1}} = \eta \bar{\pi}_n + f(\pi_{n,I} - \pi_{n,II}) \leq \eta \pi_{n,g} + \frac{\eta^2 (\pi_{n,I} - \pi_{n,II})^2}{8}.$$

For one thing,

$$\begin{aligned} \ln W_N &= \ln[(\frac{x_1 V_{N,I}}{V_{0,I}})^\eta + (\frac{x_2 V_{N,II}}{V_{0,II}})^\eta] \\ &\geq \eta \max(\ln \frac{x_1 V_{N,I}}{V_{0,I}}, \ln \frac{x_2 V_{N,II}}{V_{0,II}}). \end{aligned}$$

For another,

$$\begin{aligned} \ln W_N &= \ln W_0 + \sum_{n=1}^N \ln \frac{W_n}{W_{n-1}} \\ &\leq \ln(x_1^\eta + x_2^\eta) + \sum_{n=1}^N [\eta \pi_{n,g} + \frac{\eta^2 (\pi_{n,I} - \pi_{n,II})^2}{8}] \\ &= \ln(x_1^\eta + x_2^\eta) + \eta \ln G_N + \frac{\eta^2}{8} \sum_{n=1}^N (\pi_{n,I} - \pi_{n,II})^2. \end{aligned}$$

Define $q^2(R) = \sum_{n=1}^N (\pi_{n,I} - \pi_{n,II})^2$. Combining the lower bound with the upper bound of $\ln W_N$ and rearranging, we get the stated inequality. $\square$

The Robust Pricing of European Options

Consider a European call option on the stock with strike price K and expiration date T . It has a payoff of $\max(S_N - K, 0)$ at time T . In order to price the call option, we consider a static portfolio consisting of the call option and K shares of the bond. At time T , the portfolio earns $\max(S_N, K)$. If we can construct a dynamic portfolio such that its value at time T is always greater than $\max(S_N, K)$ for any $R \in \Phi_T$, then based on the no-arbitrage principle, the initial cost of the dynamic portfolio minus K could be viewed as the upper bound of the call option price. In fact, based on the regret bound in Theorem 1, we can construct such a dynamic portfolio by the η -momentum trading strategy as long as the two reference strategies are defined as:

- Strategy **I**: Buy one share of the stock at time 0 and hold it up to time T .
- Strategy **II**: Invest an initial capital of K in the bond at time 0 and hold it up to time T .

Robust Upper Bounds for European Call Options

Theorem 2. *The price of a European call option with stock price S_0 , strike price K and expiration date T satisfies*

$$C^E(S_0, K, T | \Phi_T) \leq \min_{\eta > 0} S_0 e^{\frac{\ln[1+(K/S_0)^\eta]}{\eta} + \frac{\eta q^2(\Phi_T)}{8}} - K, \quad (1)$$

where $q^2(\Phi_T)$ is the maximal quadratic variation of the excess log returns among all of the possible stock return paths

$$\text{in } \Phi: q(\Phi_T) = \sqrt{\sup_{R \in \Phi_T} q^2(R)} = \sqrt{\sup_{R \in \Phi_T} \sum_{n=1}^N \pi_{n,s}^2}.$$

Proof. To derive upper bounds for European call options, we define the two reference strategies in section as

- **Strategy I:** Buy one share of the stock at time 0 and hold it up to time T .
- **Strategy II:** Invest an initial capital of K in the bond at time 0 and hold it up to time T .

By Theorem 1, the dynamical portfolio adjusted by the η -momentum trading strategy at time T is worth

$$G_N \geq e^{-\frac{\ln(x_1^\eta + x_2^\eta)}{\eta} - \frac{\eta q^2(R)}{8}} \max\left(\frac{x_1 S_N}{S_0}, x_2\right),$$

where $q^2(R) = \sum_{n=1}^N \pi_{n,s}^2$. Define $q^2(\Phi_T) = \sup_{R \in \Phi_T} q^2(R)$. We get, for any $R \in \Phi_T$,

$$G_N \geq e^{-\frac{\ln(x_1^\eta + x_2^\eta)}{\eta} - \frac{\eta q^2(\Phi_T)}{8}} \max\left(\frac{x_1 S_N}{S_0}, x_2\right).$$

Now consider the following two portfolios at time 0:

Portfolio A: long one European call option on stock with strike price K and expiration date T and simultaneously invest an initial capital of K in the bond.

Portfolio B: invest I into the dynamically adjusted portfolio above.

At time T , Portfolio A earns $\max(S_N, K)$, while portfolio B is worth greater than $I * e^{-\frac{\ln(x_1^\eta + x_2^\eta)}{\eta} - \frac{\eta q^2(\Phi_T)}{8}} \max\left(\frac{x_1 S_N}{S_0}, x_2\right)$. If (x_1, x_2, η, I) is carefully chosen such that the payoff of portfolio B at time T is greater than $\max(S_N, K)$, portfolio B should have a greater initial investment cost than portfolio A. As a result, $I - K$ should be viewed as an upper bound for the price of the European call option $C^E(S_0, K, T | \Phi_T)$. In order to obtain an upper bound as tight as possible, (x_1, x_2, η, I) is selected to satisfy:

$$\begin{aligned} \min_{x_1, x_2, \eta, I} \quad & I \\ \text{s.t.} \quad & x_1 e^{-\frac{\ln(x_1^\eta + x_2^\eta)}{\eta} - \frac{\eta q^2(\Phi_T)}{8}} I \geq S_0, \\ & x_2 e^{-\frac{\ln(x_1^\eta + x_2^\eta)}{\eta} - \frac{\eta q^2(\Phi_T)}{8}} I \geq K. \end{aligned}$$

By Karush-Kuhn-Tucker Conditions, the optimal solution (x_1^*, x_2^*, I^*) satisfies:

$$\begin{aligned} x_1^* &= \frac{S_0}{K} x_2^*, \\ \text{and} \quad I^* &= \min_{\eta > 0} S_0 e^{\frac{\ln[1+(K/S_0)^\eta]}{\eta} + \frac{\eta q^2(\Phi_T)}{8}}. \end{aligned}$$

Therefore, the price of the European call option $C^E(S_0, K, T | \Phi_T)$ satisfies

$$C^E(S_0, K, T | \Phi_T) \leq \min_{\eta > 0} S_0 e^{\frac{\ln[1+(K/S_0)^\eta]}{\eta} + \frac{\eta q^2(\Phi_T)}{8}} - K. \quad \square$$

Note that the inequality (1) holds for all $\eta > 0$. Specially, when $K = S_0$, the optimal η takes the value of $\eta^* = \sqrt{\frac{8 \ln 2}{q^2(\Phi_T)}}$. Substituting η by η^* in the inequality (1), we can obtain the optimal upper bound with a closed-form expression for an at-the-money call option.

Corollary 3. *The price of an at-the-money call option with stock price S_0 and expiration date T satisfies*

$$C^E(K = S_0, S_0, T | \Phi_T) \leq S_0 (e^{\sqrt{\frac{\ln 2}{2}} q(\Phi_T)} - 1),$$

where $q(\Phi_T)$ is defined in Theorem 2.

Comparison with (DeMarzo et al. 2016)

By exploiting the gradient trading strategy, DeMarzo et al. (2016) develop robust upper bounds for European call option prices that satisfy

$$C^E(S_0, K, T | \Phi_T) \leq S_0 e^{\frac{1}{2} \sqrt{2q^2(\Phi_T) + (\ln \frac{K}{S_0})^2} + \frac{1}{2} \ln \frac{K}{S_0}} - K,$$

where $q^2(\Phi_T)$ is defined in Theorem 2. Let U^E be our optimal upper bounds and U_{DKM}^E be the corresponding results in (DeMarzo et al. 2016) for European call options. Based on the above results, we can directly compare U^E with U_{DKM}^E for the at-the-money option.

Theorem 4. *If $K = S_0$, then $U^E < U_{DKM}^E$ for all positive $q^2(\Phi_T)$.*

Proof. When $K = S_0$, $U_{DKM}^E = S_0 (e^{\sqrt{\frac{1}{2}} q(\Phi_T)} - 1)$. Since $\frac{\ln 2}{2} < \frac{1}{2}$, we get $U^E < U_{DKM}^E$ for all positive $q^2(\Phi_T)$. $\square$

Specially, for the at-the-money option with small $q^2(\Phi_T)$,

$$\frac{U^E}{U_{DKM}^E} = \frac{e^{\sqrt{\frac{\ln 2}{2}} q(\Phi_T)} - 1}{e^{\sqrt{\frac{1}{2}} q(\Phi_T)} - 1} \approx \frac{\sqrt{\frac{\ln 2}{2}} q(\Phi_T)}{\sqrt{\frac{1}{2}} q(\Phi_T)} = \sqrt{\ln 2} \approx 0.83.$$

For general cases, our optimal bounds U^E do not have closed-form solutions. Instead of comparing U^E with U_{DKM}^E directly, we first construct suboptimal upper bounds by specifying η to take a fixed value of η^* in Theorem 2, and then compare U_{DKM}^E with the suboptimal bounds to evaluate the tightness of U^E . For a given option, define the mon-eyness, m , as the ratio of its strike price to the stock price. We have the following result.

Theorem 5. *Given any $q^2(\Phi_T)$, $U^E < U_{DKM}^E$ if $\frac{K}{S_0} \in [e^{-1.27 * q(\Phi_T)}, e^{1.27 * q(\Phi_T)}]$. In other words, given any mon-eyness interval $[1 - \epsilon, 1 + \epsilon]$ with $\epsilon \in (0, 1]$, if $q(\Phi_T) \geq q_{min} = -\frac{\ln(1-\epsilon)}{1.27}$, $U^E < U_{DKM}^E$ for all $\frac{K}{S_0} \in [1 - \epsilon, 1 + \epsilon]$.*

Proof. We first prove the following technical inequality: for $\forall y > 0$,

$$\left(\frac{y-1}{y+1}\right)^2 \geq \frac{\ln^2 y}{16 \ln 2 + \ln^2 y}, \quad (2)$$

Define $h(y) := \frac{2\sqrt{\ln 2}(y-1)}{\sqrt{y}} - \ln y$. Since $h(1) = 0$, and

$$\begin{aligned} h'(y) &= \sqrt{\frac{\ln 2}{y}} + \frac{\sqrt{\ln 2}}{y^{\frac{3}{2}}} - \frac{1}{y} \\ &= \sqrt{\frac{\ln 2}{y}} \left(1 + \frac{1}{y} - \frac{1}{\sqrt{y \ln 2}}\right) \\ &= \sqrt{\frac{\ln 2}{y}} \left[\left(\frac{1}{\sqrt{y}} - \frac{1}{\sqrt{4 \ln 2}}\right)^2 + 1 - \frac{1}{4 \ln 2}\right] > 0, \end{aligned}$$

we have $\frac{2\sqrt{\ln 2}(y-1)}{\sqrt{y}} \geq \ln y \geq 0$ for $y \geq 1$, and $\frac{2\sqrt{\ln 2}(y-1)}{\sqrt{y}} < \ln y < 0$ for $0 < y < 1$. It follows that for $\forall y > 0$, $\frac{4 \ln 2 * (y-1)^2}{\ln^2 y} \geq y$. Substituting y by $\frac{1}{4}[(y+1)^2 - (y-1)^2]$ in the right-hand side of the inequality above and rearranging, we obtain the inequality (2).

Define $g(m) := \frac{\ln(1+m^{\eta^*})}{\eta^*} + \frac{\eta^* q^2(\Phi_T)}{8} - \frac{1}{2} \sqrt{2q^2(\Phi_T) + (\ln m)^2} - \frac{1}{2} \ln m$, where $m = \frac{K}{S_0}$ and $\eta^* = \sqrt{\frac{8 \ln 2}{q^2(\Phi_T)}}$. Since

$$\begin{aligned} g'(m) &= \frac{m^{\eta^*-1}}{1+m^{\eta^*}} - \frac{\ln m}{2m \sqrt{2q^2(\Phi_T) + (\ln m)^2}} - \frac{1}{2m} \\ &= \frac{1}{2m} \left[\frac{2m^{\eta^*}}{1+m^{\eta^*}} - \frac{\ln m}{\sqrt{2q^2(\Phi_T) + (\ln m)^2}} - 1 \right] \\ &= \frac{1}{2m} \left[\frac{m^{\eta^*} - 1}{m^{\eta^*} + 1} - \frac{\ln m^{\eta^*}}{\sqrt{2q^2(\Phi_T) \eta^{*2} + (\ln m^{\eta^*})^2}} \right] \\ &= \frac{1}{2m} \left[\frac{m^{\eta^*} - 1}{m^{\eta^*} + 1} - \frac{\ln m^{\eta^*}}{\sqrt{16 \ln 2 + (\ln m^{\eta^*})^2}} \right], \end{aligned}$$

based on the inequality (2), we have $g'(m) > 0$ for $m > 1$ and $g'(m) < 0$ for $0 < m < 1$. Since $g(1) = \sqrt{\frac{\ln 2}{2}} q(\Phi_T) - \sqrt{\frac{1}{2}} q(\Phi_T) < 0$ and $\lim_{m \rightarrow 0} g(m) = \lim_{m \rightarrow +\infty} g(m) = \frac{\eta^* q^2(\Phi_T)}{8} > 0$ for positive $q(\Phi_T)$, it follows that there must exist $0 < m_1 < 1$ and $m_2 > 1$ such that $g(m_1) = g(m_2) = 0$. To determine m_1 and m_2 , we guess they satisfy the following function form: $m_i = e^{\lambda_i q(\Phi_T)}$ with $i \in \{1, 2\}$. Plugging m_i in $g(m)$, we get

$$g(m_i) = \left[\frac{\ln(1 + e^{\lambda_i \sqrt{8 \ln 2}})}{\sqrt{2 \ln 2}} + \sqrt{\frac{\ln 2}{2}} - \sqrt{2 + \lambda_i^2} - \lambda_i \right] \frac{q(\Phi_T)}{2}.$$

Solving $g(m_i) = 0$ numerically, we get $\lambda_1 \approx -1.27$ and $\lambda_2 \approx 1.27$. Therefore, when $m \in [e^{-1.27 * q(\Phi_T)}, e^{1.27 * q(\Phi_T)}]$, we have $\frac{\ln(1+m^{\eta^*})}{\eta^*} + \frac{\eta^* q^2(\Phi_T)}{8} \leq \frac{1}{2} \sqrt{2q^2(\Phi_T) + (\ln m)^2} + \frac{1}{2} \ln m$. Noting that $\min_{\eta > 0} \frac{\ln(1+m^\eta)}{\eta} + \frac{\eta q^2(\Phi_T)}{8} \leq \frac{\ln(1+m^{\eta^*})}{\eta^*} + \frac{\eta^* q^2(\Phi_T)}{8}$, we thus have $U^E < U_{DKM}^E$ if $\frac{K}{S_0} \in [e^{-1.27 * q(\Phi_T)}, e^{1.27 * q(\Phi_T)}]$. It follows that given any moneyness interval $[1 - \epsilon, 1 + \epsilon]$ with

$\epsilon \in (0, 1]$, if $q(\Phi_T) \geq q_{min} = \max(\frac{\ln(1+\epsilon)}{1.27}, -\frac{\ln(1-\epsilon)}{1.27}) = -\frac{\ln(1-\epsilon)}{1.27}$, $U^E < U_{DKM}^E$ for all $\frac{K}{S_0} \in [1 - \epsilon, 1 + \epsilon]$. $\square$

In the Black-Scholes or standard binomial-tree framework with risk-free rate $r = 0$, $q^2(\Phi_T) = \sigma^2 T$.² This result provides us with an opportunity to examine the superiority of our bounds in common moneyness, volatility and expiration date setups. We set $\sigma \in \{0.15, 0.3, 0.45\}$ to cover low, medium, and high volatility, and $T \in \{1 \text{ month}, 3 \text{ months}, 6 \text{ months}, 1 \text{ year}\}$ to cover short-, medium- and long-term expiration date. Since options with moneyness $m \in [0.9, 1.1]$ are usually traded frequently, we are concerned if our bounds are superior within the prevalent moneyness interval for these common moneyness, volatility and expiration date setups. The moneyness interval within which our bounds are superior for various combination of σ and T is reported in Table 1. We find that, except for $(\sigma, T) = (0.15, 1 \text{ month})$ combination, our bounds outperform the benchmarks when $m \in [0.9, 1.1]$ for all volatility and expiration date setups. Therefore, although our optimal upper bounds do not have closed-form solutions, by taking a specific value of η , our framework admits sub-optimal closed-form upper bounds that are still superior for most common moneyness, volatility and expiration date setups.

σ	T			
	1 month	3 months	6 months	1 year
0.15	[0.95, 1.05]	[0.90, 1.10]	[0.88, 1.14]	[0.83, 1.20]
0.30	[0.90, 1.11]	[0.83, 1.21]	[0.76, 1.31]	[0.69, 1.46]
0.45	[0.85, 1.18]	[0.75, 1.33]	[0.67, 1.50]	[0.56, 1.77]

Table 1: Moneyness interval within which our bounds are superior.

Extension to Asian Options

By appropriately modifying the reference strategies to accommodate the payoff function of options, our methodology can be extended to obtain robust upper bounds for several prevalent exotic options. In this section, we mainly focus on pricing discretely monitored average strike Asian options with the arithmetic mean.

Robust Upper Bounds for Average Strike Asian Call Options

An average strike call option with expiration date T is an Asian option that pays its holder $\max(S_N - \bar{S}_N, 0)$ at time

T , where $\bar{S}_N = \frac{1}{N+1} \sum_{i=0}^N S_i$. To derive the upper bound, we define the two reference strategies as:

- **Strategy I**: Buy one share of the stock at time 0. From time 0 on, sell a fraction $\frac{1}{N+1-n}$ of the remaining stock and invest the proceeds in the bond at the end of the n th period for $n \in \{0, \dots, N-1\}$.

²Please refer to (DeMarzo et al. 2016) for the detail.

- **Strategy II:** Buy one share of the stock at time 0 and hold it up to time T .

Denote by $p_{n,s}$ the portion of $V_{n-1,I}$ invested in the stock at the beginning of the n th period. By definition,
$$p_{n,s} = \frac{\frac{1}{N+1}(N-n+1)S_{n-1}}{\frac{1}{N+1}\sum_{i=0}^{n-1}S_i + \frac{(N-n+1)S_{n-1}}{N+1}} = \frac{N-n+1}{\sum_{i=0}^{n-2}\frac{S_i}{S_{n-1}} + N-n+2}.$$
 We have the following result.

Theorem 6. *The price of an average strike call option with initial stock price S_0 and expiration date T satisfies:*

$$C_{AS}^E(S_0, T | \Phi_T) \leq S_0(e^{\sqrt{\frac{\ln 2}{2}}q(\Phi_T)} - 1),$$

where $q(\Phi_T)$ is defined in Theorem 2.

Proof. By Theorem 1, the dynamical portfolio adjusted by the η -momentum trading strategy at time T is worth $G_N \geq e^{-\frac{\ln(x_1^\eta + x_2^\eta)}{\eta} - \frac{\eta q^2(R)}{8}} \max(\frac{x_1 \bar{S}_N}{S_0}, \frac{x_2 S_N}{S_0})$, where $q^2(R) = \sum_{n=1}^N (\pi_{n,I} - \pi_{n,s})^2$. Since $\frac{\partial(\pi_{n,I} - \pi_{n,s})^2}{\partial p_{n,s}} = \frac{2 \ln[1 - \frac{(1-p_{n,s})r_{n,s}}{1+r_{n,s}}] \cdot \frac{r_{n,s}}{1+r_{n,s}}}{1 - \frac{(1-p_{n,s})r_{n,s}}{1+r_{n,s}}} \leq 0$, namely, $(\pi_{n,I} - \pi_{n,s})^2$ decreases in $p_{n,s}$, we have $q^2(R) = \sum_{n=1}^N (\pi_{n,I} - \pi_{n,s})^2 \leq \sum_{n=1}^N \pi_{n,s}^2$. It follows that for any $R \in \Phi_T$,

$$G_N \geq e^{-\frac{\ln(x_1^\eta + x_2^\eta)}{\eta} - \frac{\eta q^2(\Phi_T)}{8}} \max(\frac{x_1 \bar{S}_N}{S_0}, \frac{x_2 S_N}{S_0}).$$

Consider the following two portfolios at time 0:

Portfolio A: Buy one average strike call option with expiration date T and simultaneously buy one share of the stock following Strategy I.

Portfolio B: Invest an initial capital of I into the dynamic portfolio, which will be adjusted by the η -momentum strategy later on.

At time T , Portfolio A earns $\max(S_N, \bar{S}_N)$. On the other hand, Portfolio B is worth $I * G_N$, which satisfies:

$$I * G_N \geq I * e^{-\frac{\ln(x_1^\eta + x_2^\eta)}{\eta} - \frac{\eta q^2(\Phi_T)}{8}} \max(\frac{x_1 \bar{S}_N}{S_0}, \frac{x_2 S_N}{S_0}).$$

If (x_1, x_2, η, I) is carefully chosen such that the payoff of Portfolio B at time T is greater than $\max(S_N, \bar{S}_N)$, the initial value of Portfolio B should be greater than that of Portfolio A. Consequently, $I - S_0$ should be an upper bound on $C_{AS}^E(S_0, T)$. To obtain an optimal upper bound, (x_1, x_2, η, I) should satisfy:

$$\begin{aligned} \min_{x_1, x_2, \eta, I} \quad & I \\ \text{s.t.} \quad & x_1 e^{-\frac{\ln(x_1^\eta + x_2^\eta)}{\eta} - \frac{\eta q^2(\Phi_T)}{8}} I \geq S_0, \\ & x_2 e^{-\frac{\ln(x_1^\eta + x_2^\eta)}{\eta} - \frac{\eta q^2(\Phi_T)}{8}} I \geq S_0. \end{aligned}$$

By the Karush-Kuhn-Tucker Conditions, we obtain the optimal solution $(x_1^*, x_2^*, \eta^*, I^*)$ as:

$$\begin{aligned} x_1^* &= x_2^*, \\ \eta^* &= \sqrt{\frac{8 \ln 2}{q^2(\Phi_T)}}, \\ \text{and } I^* &= S_0 e^{\sqrt{\frac{\ln 2}{2}}q(\Phi_T)}. \end{aligned}$$

The stated result is thus obtained by plugging in I^* . $\square$

Given a constant $\underline{R}$, if $\frac{S_n}{S_{n-1}} \geq \underline{R}$, the above upper bound can be further tightened by reducing the value of $q^2(\Phi_T)$. It is easy to verify that $p_{n,s} \geq \frac{N-n+1}{\frac{1-\underline{R}^{1-n}}{\underline{R}-1} + N-n+2}$. Denote $\frac{N-n+1}{\frac{1-\underline{R}^{1-n}}{\underline{R}-1} + N-n+2}$ as $\underline{p}_{n,s}$, we have the following Corollary according to Theorem 6.

Corollary 7. *If $\frac{S_n}{S_{n-1}} \geq \underline{R}$, then the price of an average strike call option with initial stock price S_0 and expiration date T satisfies:*

$$C_{AS}^E(S_0, T | \Phi_T) \leq S_0(e^{\sqrt{\frac{\ln 2}{2}}q_c(\Phi_T)} - 1),$$

where $q_c^2(\Phi_T) = \sup_{R \in \Phi_T} \sum_{n=1}^N \ln^2(1 + \underline{p}_{n,s} r_{n,s})$.

Comparison with (Gofer and Mansour 2011)

By extending the methodology of (DeMarzo et al. 2006), Gofer and Mansour (2011) derive upper bounds for average strike Asian call options that satisfy

$$C_{AS}^E(S_0, T | \Phi_T) \leq \min_{1 \leq \mu \leq \frac{1-2M}{2M(1-M)}} S_0(e^{(\mu-1)Q_T + \frac{\ln 2}{\mu}} - 1),$$

where $Q_T = \sup_{R \in \Phi_T} \sum_{n=1}^N r_{n,s}^2$ and M is the maximal absolute return of the stock in each period. Note that their bounds are not robust, since they require $M \leq 1 - \frac{\sqrt{2}}{2}$ except for the no-arbitrage principle. Let U^{AS} and U_c^{AS} be our bounds in Theorem 6 and Theorem 7, and U_{GM}^{AS} be the corresponding results in (Gofer and Mansour 2011) for average strike Asian call options. Since $q^2(\Phi_T)$ and Q_T have distinct definitions, it is hard to determine which bounds are lower theoretically. However, for the special case of $q^2(\Phi_T) \leq Q_T \leq 2 \ln 2$, our upper bounds for average strike Asian options are definitely superior.

Theorem 8. *If $q^2(\Phi_T) \leq Q_T \leq 2 \ln 2$, then we have $U_c^{AS} \leq U^{AS} \leq U_{GM}^{AS}$.*

Proof. If $\ln 2 \leq q^2(\Phi_T) \leq Q_T \leq 2 \ln 2$, $U_{GM}^{AS} = S_0$ and $U^{AS} \leq S_0$. We thus have $U^{AS} \leq U_{GM}^{AS}$. If $q^2(\Phi_T) \leq Q_T < \ln 2$, it is easy to verify that $\min_{1 \leq \mu \leq \frac{1-2M}{2M(1-M)}} (\mu - 1)Q_T + \frac{\ln 2}{\mu} \geq \min_{\mu > 0} (\mu - 1)Q_T + \frac{\ln 2}{\mu} = 2\sqrt{\ln 2} * \sqrt{Q_T} - Q_T \geq (2\sqrt{\ln 2} - 1)\sqrt{Q_T} \geq (2\sqrt{\ln 2} - 1)q(\Phi_T)$. Since $2\sqrt{\ln 2} - 1 > \sqrt{\frac{\ln 2}{2}}$, we have $\min_{1 \leq \mu \leq \frac{1-2M}{2M(1-M)}} (\mu - 1)Q_T + \frac{\ln 2}{\mu} \geq \sqrt{\frac{\ln 2}{2}}q(\Phi_T)$. It follows that $U^{AS} \leq U_{GM}^{AS}$. Recalling that $U_c^{AS} \leq U^{AS}$, we conclude the proof. $\square$

Note that, when $Q_T > 2 \ln 2$, U_{GM}^{AS} is equal to S_0 , which is a trivial upper bound, while U^{AS} and U_c^{AS} may exceed S_0 . However, by taking the minimum of U^{AS} (or U_c^{AS}) and S_0 , we can guarantee our new upper bounds are always lower than U_{GM}^{AS} as long as $q^2(\Phi_T) \leq Q_T$.

Consider a N -time step binomial tree model from time 0 to time T with step size $\Delta t = \frac{T}{N}$, the gross return of the stock in each step is either $u = e^{\sigma\sqrt{\Delta t}}$ or $d = \frac{1}{u}$. Since $u - 1 \geq 1 - d > 0$, we have $Q_T = N(u - 1)^2$ and $q^2(\phi_T) = N(\ln u)^2$ by definition. It follows that $q^2(\Phi_T) < Q_T$. Therefore, by Theorem 8, our bounds for Asian options are superior in the standard binomial-tree model.

Numerical Simulations

In this section, the performance of our upper bounds for European options and average strike Asian options is examined by comparing them with their corresponding benchmarks.

For a given option, let U and U_b be our upper bound and the benchmark upper bound for the option, respectively. To evaluate the performance of our bounds relative to the benchmarks, we define the tightness of our bounds, T_u , as $T_u = 1 - \frac{U}{U_b}$. We enumerate all reasonable combinations of m and $q^2(\phi)$. The range for m is $[0.6, 1.4]$ with stepsize 0.1 while the range for $q^2(\phi)$ is $[0, 0.5]$ with stepsize 10^{-4} .




Figure 1: Tightness of our bounds for European call options.

For European options, the benchmarks we use are the bounds in (DeMarzo et al. 2016). Our optimal upper bounds are generated by solving the optimization problems with η numerically, while the benchmarks are calculated by plugging in the parameters directly. The magnitude of T_u with varying $q^2(\Phi_T)$ and m for European call options is illustrated in Figure 1. In the heatmap, the red area represents our bounds underperform the benchmarks, while the other color area represents our bounds outperform the benchmarks. As the color changes from yellow to green, and then to blue, T_u increases gradually, which indicates the performance of our bounds relative to the benchmarks is getting better.

For a given $q^2(\Phi_T)$, denote by $[m^l, m^u]$ the moneyness interval within which $U \leq U_b$ through numerical simulations. Denote by $\hat{q}_{min}^2$ the minimum $q^2(\Phi_T)$ to guarantee $U \leq U_b$ within a given moneyness interval $[1 - \epsilon, 1 + \epsilon]$ through numerical simulations. There are at least four observations worth noting. First, our bounds outperform the benchmarks for most cases because of the quite small proportion of the red area. In line with Theorem 5, the interval $[m^l, m^u]$ becomes wider as $q^2(\Phi_T)$ increases. Second,

ϵ	Moneyness interval	$\hat{q}_{min}^2$	q_{min}^2
0.05	[0.95, 1.05]	$5.0 * 10^{-4}$	$1.63 * 10^{-3}$
0.1	[0.9, 1.1]	$2.2 * 10^{-3}$	$6.88 * 10^{-3}$
0.15	[0.85, 1.15]	$5.1 * 10^{-3}$	$1.64 * 10^{-2}$
0.2	[0.8, 1.2]	$9.5 * 10^{-3}$	$3.09 * 10^{-2}$

Table 2: q_{min}^2 versus $\hat{q}_{min}^2$ for common moneyness interval.




Figure 2: Upper bounds versus Black-Scholes prices.

as reported in Table 2, $\hat{q}_{min}^2$ is one third smaller than q_{min}^2 for common moneyness interval. This indicates that our upper bounds are lower than the benchmarks within a much wider moneyness spread than that implied by Theorem 5 for a given $q^2(\Phi_T)$. Third, T_u tends to take the largest value as m approaches to one, i.e., when options are near at-the-money. With m departure from one, T_u exhibits a decreasing tendency. Specially, for the at-the-money call options, T_u increases and ranges from 17% to 22% with increasing $q^2(\Phi_T)$. Last but not the least, It's interesting to see that T_u tends to take larger value when $m > 1$ than $m < 1$ for a given $q^2(\Phi_T)$, which means our upper bounds perform better relative to the benchmarks for out-of-the-money options than for in-the-money options, especially for deep-out-of-the-money options.

Besides the general comparisons of the European call option pricing above, we also make comparisons in the Black-Scholes framework. We set $S_0 = 100$, $\sigma = 0.15$, $T = 1$ year, $m \in [0.8, 1.2]$, and $r = q = 0$. It follows that $q^2(\phi_T) = \sigma^2 T = 0.0225$. The results are illustrated in Figure 2. Both our bound and the bound in (DeMarzo et al. 2016) are a convex function of the strike price, which are akin to the shape of the Black-Scholes option price. Furthermore, our bounds are significantly tighter than those in (DeMarzo et al. 2016) for all $m \in [0.8, 1.2]$. Specially, for the at-the-money call option, our bound is 54% larger than the Black-Scholes price, whereas the value is 87% for the




Figure 3: Upper bounds for Asian call options.

result in (DeMarzo et al. 2016).

For average strike Asian option, the results in (Gofer and Mansour 2011) are the benchmarks. To accommodate their requirement that the maximal absolute return of the stock in each period is no more than $1 - \frac{\sqrt{2}}{2}$, we use the binomial tree as the model of stock returns. Consider a N -time step binomial tree model from time 0 to time T with step size $\Delta t = \frac{T}{N}$, where $r = q = 0$. During each step, the gross return of the stock is either $u = e^{\sigma\sqrt{\Delta t}}$ or $d = \frac{1}{u}$. Since $u - 1 \geq 1 - d > 0$, by definition, it follows that $M = u - 1$, $\underline{R} = d$, $Q_T = N(u - 1)^2$, $q^2(\phi_T) = \sigma^2 T$, and $q_c^2(\phi_T) = \sum_{n=1}^N \ln[1 + \frac{u - u_n}{1 - u} + \frac{N - n + 1}{N - n + 2}(u - 1)]^2$. Both of our bounds and the benchmarks are calculated by plugging in these parameters directly. The comparison of our bounds with the benchmarks is illustrated in Figure 3. In this figure, the blue line represents the benchmarks scaled by S_0 , whereas the dashed and solid red line represent our bounds with or without the constraint of $d \leq \frac{S_n}{S_{n-1}}$ scaled by S_0 , respectively. Consistent with Theorem 8, our upper bounds for Asian call options are always lower than the benchmarks across the whole spectrum of $q^2(\phi_T)$. As $q^2(\phi_T)$ increases, T_u decreases monotonically with range from 46.9% to 64.6%. With the constraint of $d \leq \frac{S_n}{S_{n-1}}$, our bounds are even tighter. As $q^2(\phi_T)$ increases, T_u increases monotonically with range from 80% to 86.3%.

Conclusions

By exploiting the η -momentum trading strategy, we develop new robust upper bounds for European options and average strike Asian options. Our upper bounds are proved to be tighter than the benchmarks under mild conditions. Numerical simulations demonstrate that both of our bounds significantly outperform the benchmarks for most cases. Except for average strike Asian options, our methodology can be applied in other prevalent exotic option pricing, including average price Asian options, shout options, forward start options as well as exchange options.

It is worth noting that, although our bounds significantly outperform the benchmarks for most common moneyess,

volatility and expiration date setups. It would be interesting to investigate the optimal upper bound in the robust setting. In addition, we do not provide lower bounds in this paper, which is left as an attractive open question.

Acknowledgments

The authors are grateful for many valuable comments from the anonymous reviewers. This study was partially supported by the National Natural Science Foundation of China (Grant Numbers 11501464, 11761141007, 71971177).

References

- Abernethy, J.; Bartlett, P.; Frongillo, R.; and Wibisono, A. 2013. How to hedge an option against an adversary: Black-Scholes pricing is minimax optimal. *Advances in Neural Information Processing Systems*, 26: 2346–2354.
- Abernethy, J.; Frongillo, R. M.; and Wibisono, A. 2012. Minimax option pricing meets Black-Scholes in the limit. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing*, 1029–1040.
- Bates, D. S. 1996. Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options. *The Review of Financial Studies*, 9(1): 69–107.
- Black, F.; and Scholes, M. 1973. The pricing of options and corporate liabilities. *The Journal of Political Economy*, 81(3): 637–654.
- Blackwell, D. 1956. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1): 1–8.
- Blum, A.; and Mansour, Y. 2007. From external to internal regret. *Journal of Machine Learning Research*, 8(6): 1307–1324.
- Brown, H.; Hobson, D.; and Rogers, L. C. 2001. Robust hedging of barrier options. *Mathematical Finance*, 11(3): 285–314.
- Cesa-Bianchi, N.; and Lugosi, G. 2006. *Prediction, learning, and games*. Cambridge University Press.
- Chen, R.-R.; and Yeh, S.-K. 2002. Analytical upper bounds for American option prices. *Journal of Financial and Quantitative Analysis*, 37(1): 117–135.
- Chen, X.; Deelstra, G.; Dhaene, J.; and Vanmaele, M. 2008. Static super-replicating strategies for a class of exotic options. *Insurance: Mathematics and Economics*, 42(3): 1067–1085.
- Chung, S.-L.; and Chang, H.-C. 2007. Generalized analytical upper bounds for American option prices. *Journal of Financial and Quantitative Analysis*, 42(1): 209–227.
- DeMarzo, P.; Kremer, I.; and Mansour, Y. 2006. Online trading algorithms and robust option pricing. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, 477–486.
- DeMarzo, P. M.; Kremer, I.; and Mansour, Y. 2016. Robust option pricing: Hannan and Blackwell meet Black and Scholes. *Journal of Economic Theory*, 163: 410–434.
- Du, Y.; Xue, S.; and Liu, Y. 2019. Robust upper bounds for American put options. *Journal of Futures Markets*, 39(1): 3–14.

- Foster, D. P.; and Vohra, R. V. 1997. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1-2): 40.
- Foster, D. P.; and Vohra, R. V. 1999. Regret in the on-line decision problem. *Games and Economic Behavior*, 29(1-2): 7–35.
- Freund, Y.; and Schapire, R. E. 1996. Game theory, on-line prediction and boosting. In *Proceedings of the 9th Annual Conference on Computational Learning Theory*, 325–332.
- Gofer, E.; and Mansour, Y. 2011. Pricing exotic derivatives using regret minimization. In *International Symposium on Algorithmic Game Theory*, 266–277.
- Hannan, J. 1957. *Approximation to Bayes risk in repeated play*, *Contributions to the Theory of Games III*. Princeton University Press.
- Hart, S.; and Mas-Colell, A. 2000. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5): 1127–1150.
- Heston, S. L. 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6(2): 327–343.
- Hobson, D. 1998. Robust hedging of the lookback option. *Finance and Stochastics*, 2(4): 329–347.
- Hobson, D.; Laurence, P.; and Wang, T.-H. 2005a. Static-arbitrage optimal subreplicating strategies for basket options. *Insurance: Mathematics and Economics*, 37(3): 553–572.
- Hobson, D.; Laurence, P.; and Wang, T.-H. 2005b. Static-arbitrage upper bounds for the prices of basket options. *Quantitative Finance*, 5(4): 329–342.
- Hobson, D.; and Neuberger, A. 2012. Robust bounds for forward start options. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, 22(1): 31–56.
- Kahalé, N. 2016. Model-independent lower bound on variance swaps. *Mathematical Finance*, 26(4): 939–961.
- Kahalé, N. 2017. Superreplication of financial derivatives via convex programming. *Management Science*, 63(7): 2323–2339.
- Kou, S. G. 2002. A jump-diffusion model for option pricing. *Management Science*, 48(8): 1086–1101.
- Laurence, P.; and Wang, T.-H. 2008. Distribution-free upper bounds for spread options and market-implied antimonotonicity gap. *The European Journal of Finance*, 14(8): 717–734.
- Laurence, P.; and Wang, T.-H. 2009. Sharp distribution free lower bounds for spread options and the corresponding optimal subreplicating portfolios. *Insurance: Mathematics and Economics*, 44(1): 35–47.
- Shalev-Shwartz, S. 2011. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2): 107–194.
- Xue, S.; Du, Y.; Zhou, M.; and Xu, L. 2022. Robust exotic option pricing via no-regret learning. *Working Paper*.