Robust Dynamic Hedging

Henry Lam
Boston University
khlam@math.bu.edu

Zhenming Liu
Princeton University
zhenming@cs.princeton.edu

Abstract

We consider a robust approach to price European and American options by modeling the market dynamic as a repeated game between the nature (adversary) and the investor. This adversarial approach has been recently studied as a alternative to the traditional area of stochastic finance. In this paper, we build a systematic framework to study as well as to connect between two important questions in the adversarial setup: 1) the relation to classical models such as Black-Scholes and its extensions, and 2) the design of efficient pricing algorithms. The starting point of our analysis is an elementary characterization of the optimal strategy for each round of the option game, which gives explicit trading strategies in replicating options in the case of convex payoff, and allows to design and analyze a simple discretization-based approximation algorithm for non-convex payoffs. We further use the characterization to obtain limiting representation for options with general payoffs as diffusion processes with controlled volatilities, which reduce to the Black-Scholes price for convex payoff and a degenerate price for the concave counterpart. Our framework can also be adapted to American options, as well as a semi-adversary model that incorporates rare shocks in asset prices. The latter extension serves as the analog in robust pricing theory to so-called "jump processes" in stochastic studies on option pricing.

1 Introduction

In the financial market, an *option* is a contract that gives the holder the right, but not the obligation, to buy or sell an underlying asset or instrument [Hul09]. Consider the "vanilla" European call option as an example. This contract is controlled by three parameters: the underlying asset S, say a stock traded at the New York Stock Exchange, the strike price K, and the time to expiration T. A holder of a European call option has the right to purchase the stock at time T at the prefixed strike price K, regardless of the price of the stock at that time. Suppose that the stock price at time T, say $S(T)$ (we shall abuse notation to denote $S(t)$ as the price and S as the stock for convenience), exceeds K, the holder will exert his/her right, and *exercise* the option. In the opposite event that $S(T)$ submerges below K, a rational holder will initiate no action. Assuming the market is liquid, since the holder who decides to exercise the option can sell the stock immediately at market price, the payoff of this option at time T can be summarized as $\max\{S(T) - K, 0\}$.

In general, any single-instrument European-type option has a payoff function $g(S(T))$ on an underlying asset S, depending on the terms of the contract. An American-type option, on the other hand, gives the holder the right to exercise the option at *any* time before maturity.

The study of the fair prices of different options has not only been a core area in financial economics, but is also important for market practitioners, given the gigantic volume of options being traded at any trading day [FIA12]. The area bloomed after the groundbreaking discovery of Black and Scholes [BS73], which was later expanded by Merton [Mer73]. In order to find the fair price of an option, their main idea is to *hedge* against the movement of the underlying asset. Consider for example a European call option. If the purchaser of the option, at time 0, chooses to concurrently short a portfolio that consists of the underlying asset and cash, in such a way that the portfolio's value at time T is exactly $\max\{S(T) - K, 0\}$, then the portfolio must cost the same as the European option; otherwise the option holder can construct an *arbitrage* that generates positive profit at no downside risk, a scenario that is economically prohibited. In this way the fair price of the option is exactly the value of the *replicating* portfolio. We shall elaborate this concept further in Section 3.

Since then, there has been a huge body of literature on the pricing of options under various stochastic models of the underlying asset. In the original work in [BS73], it was shown that a unique price can be determined using the no-arbitrage principle if the underlying asset's price follows a geometric Brownian motion. Other stochastic models may or may not succumb to a unique price (the latter case is known as *incomplete* market). For example, a jump diffusion market, which we shall discuss later on in this paper, is incomplete, and a unique price can be obtained only under diversification [Mer76] or additional utility structure [REK00]. See for example [CT12, HF11] for comprehensive reviews of the literature.

In this paper we shall consider a non-stochastic framework to price and hedge options. We model it as a repeated game between the trader of the option and the "adversary" nature that controls the movement of stock price. Broadly speaking, our goal in this paper is to build a systematic framework that allows one to answer fundamental questions that have been around in the stochastic finance community. Two such questions are:

> *Q1. What is the feature of a non-stochastic model that allows one to relate to the traditional Black-Scholes model, the "Holy Grail" in option pricing?*
>
> *Q2. Can we design efficient pricing algorithms under the non-stochastic framework and what is the explicit hedging strategy implied by such algorithms?*

These interrelated questions are motivated from both theoretical and practical viewpoint. With an understanding of the similarity to the well-established Black-Scholes model, one can take advantage of existing hedging strategies to design good pricing algorithms. On the other hand, if the non-stochastic model is provably far from Black-Scholes, we want to be still able to design pricing algorithms, especially when the model is capable of capturing certain salient features of the financial market. We point out that, in practice, explicit hedging strategy is as important as the option price itself. It is only with such a strategy that the trader can replicate a portfolio, thereby realizing the value of the option.

While the posted questions are important and natural, there appears to exist limited amount of work. Moreover, the scopes of these works are each confined to quite specific settings and lack the ability to generalize; in fact, even for the simplest case of convex European options, the analysis involved can be complicated and non-constructive. For example, the work of Abernethy et al. [AFW12] provides a convergence result to Black-Scholes for convex options, but the proof is non-constructive, and hence is not easily generalizable to non-convex payoffs or American options. Regarding algorithmic results, DeMarzo et al. [DKM06] and more recently Gofer and Mansour [GM11] focus on the adaptive adversary setting and propose regret minimization algorithms; however, their technique is confined to convex payoff functions; the same holds for the line of work in [RES05, Ber05a, Ber05b, Kol11], which is also limited to analysis on convex options. Chen et al. [Che10] and Bandi and Bertsimas [BB12], on the other hand, consider general options (both European and American), but they model the nature as oblivious. Taking into account these previous limitations, our goal in this paper is to give the first comprehensive study on a non-stochastic model that can address both of the posted questions, for both European and American options, and with possibly non-convex payoff.

Our work and contribution. We consider a model in which the stock price movement at each step is restricted to a bounded deterministic set. Our choice of this model is two-fold. First, from the practical viewpoint, an estimate of the worst-case upper and lower bound of price movement is very attainable (relative to, say, the worst-case moments). Secondly, as we shall see, this simple model allows us to analyze options that are structurally difficult (such as non-convexity, and American-type), and to propose efficient pricing algorithms. It also allows us to draw insight from some previously considered settings through simplified analysis, which can be subsequently generalized to more complex situations. The following are a few key algorithmic results that we contribute:

1. We start with analyzing the equilibrium of the European option game *per round* based on *elementary* techniques. Such result enables us to connect with the standard Black-Scholes model at the limit and gives us explicit hedging strategy, when the payoff function is convex.
2. We further leverage this technique and design constructive algorithm for more complex financial derivatives. In particular, we analyze an approximation algorithm to robustly price any options whose payoff function is monotonic and Lipschitz continuous.
3. We show that, for non-convex option, convergence to Black-Scholes is not guaranteed. The limiting bound on the price, instead, is the optimal value of a continuous-time control problem with volatility constraint.
4. We generalize the above results to two other important settings. First, we extend our algorithmic and convergence results to American-type options. Second, we adapt our model and algorithmic result to incorporate rare jumps in the financial market. The latter is important because non-smooth price movement is ubiquitous in financial markets, *e.g.,* it can model "volatility smile" [Kou02].

In the supplementary material, we complement our algorithmic contribution with two hardness results: 1) exactly computing the option price's upper bound in a finite round hedging game is $\#P$-hard even for convex functions, when the uncertainty sets across different rounds are non-uniform; 2) our algorithm's approximation quality is near-optimal when we access the payoff function through an oracle. Figure 1 summarizes the contribution of this paper.

Our techniques. The subject of option pricing has traditionally been an interdisciplinary area, drawing

	Convergence	Algorithm	Explicit Hedging	Complexity
European Convex (Lipschitz)	Yes (simplified)	Approximation (new)	Yes (was open in [AFW12])	$\#P$-Hard (new)
European Non-Convex (Lipschitz)	Yes (new)	Approximation (new)	Yes (new)	Tight in the oracle model (new)
American	Yes (generalized)	Approximation (generalized)	Yes (generalized)	Same as above
Jump[1]	Yes (new)	Approximation (new)	Yes (new)	Same as above

Figure 1: Summary of our contribution.

interest from the finance, operations research, statistics, and more recently the computer science community. The development of algorithms and their analysis should, unsurprisingly, blend together ideas originating from inside and outside computer science, most notably from convex optimization and stochastic analysis. Consequently, we shall also follow this interdisciplinary nature in this paper, while taking an algorithmic-centric view. The advantage of this is two-way: algorithmic tools can help us gain better insight on the problem of option pricing; on the other hand, these "exotic" ideas can help us develop new methods for analyzing algorithms and formulating new and natural theoretical questions. Below we highlight a number of conceptually new techniques and problem formulations that result from such idea-synthesization:

1. **Construction of artificial probability measure in analyzing a deterministic algorithm.** As will be seen later on, we design a deterministic algorithm for pricing non-convex options, and the challenge lies in the analysis of the algorithm. Here, we attach an "artificial" probabilistic measure to the deterministic algorithm so that the movement of the algorithm can be characterized by the statistics under this artificial measure. At the end, analyzing the performance of this deterministic algorithm reduces to understanding the behavior of a standard martingale. This way of analyzing a deterministic algorithm appears to be new.
2. **An online algorithm where the adversary is allowed to withdraw.** In the case of American put option, the adversary is allowed to exercise the option early. From the viewpoint of online learning problem (which aims to maximize the payoff), this is equivalent to the feature that the adversary can "withdraw" from the game. Allowing the adversary to behave in this way appears to be new in the literature. Our formulation here gives the first natural problem under such setting and the first non-trivial analysis.
3. **Analysis of a limiting process inspired by the underlying combinatorial structure.** Understanding the limiting behavior of financial derivatives is not often a focus in computer science. Nevertheless, we present a major convergence result for non-convex options, for two reasons. First, such result is a coherent part of our robust dynamic hedging framework. Second, we want to convey the message that the combinatorial techniques and insight we develop can lead to "broader consequences" to other important areas.

Along the discussion above, we shall close our introduction by mentioning two important concepts in classical option pricing theory that are intimately related to this paper. The first is the celebrated binomial tree model [CRR79], as a discrete-time analog of the approach in Black-Scholes, whereby the market at each time step can only go to two levels, up or down. This model is computationally convenient and converges to the Black-Scholes price in continuous-time limit. Its extensions, such as trinomial and general multinomial tree models, are also well studied [MMS89, He90]. We will demonstrate how our robust framework reduces to these models under convexity conditions of the option payoffs. Next, the second important concept is the notion of *risk-neutral measure*, as a convenient probabilistic tool to compute option prices. The so-called fundamental theorem of asset pricing asserts that under no-arbitrage condition, the fair option price can be computed as the expectation under an artificial probability measure, the risk-neutral measure, when the market is complete [Duf05]. The concept has also been extended to incomplete market by imposing various risk/utility structures [Shr04]. As we shall see, this notion of risk-neutral measure will appear as the optimal solution of a dual problem in our formulation, and will also serve as an important proof vehicle in our *deterministic* setup.

Organization. We describe our model in Section 2. Section 3 analyzes the equilibrium of our minimax game and gives convergence results for convex payoff functions. Section 4 and Section 5 present algorithm and convergence result for pricing options with general payoff functions. Section 6 discusses how our model can incorporate price jumps. Section 7 generalizes our results to American options. In the appendices, we present hardness results as well as an alternative way to prove Black-Scholes convergence from a control formulation derived in Section 5.

2 Our model

Discrete game. We now describe our model, which has a similar spirit to [Ber05b, AFW12]. Throughout the paper we shall focus on a discrete-time setting, *i.e.,* the trader only has the chance to trade at discrete time points. Continuous-time implications of our models will be discussed in Section 3. Specifically, consider an option that expires T days from now. We denote time 0 as the time when a transaction of the option occurs, *i.e.,* a trader either buys or sells the option. We assume that, before the option expires, the trader has in total τ time points that allow trade execution. Let these time points be $t = \frac{T}{\tau}, \frac{2T}{\tau}, ..., T$. Notice that as soon as τ is decided, the value T is not a parameter in the game (but will reappear in our continuous-time limit later on). Throughout our analysis we assume no transaction costs and the market is liquid (*i.e.,* the trader can always buy or sell any volume of the asset at the market price at the τ time points).

We shall model the dynamic of the financial market from time 0 to T as a τ-round two-player game between the trader and nature. Consider an option on the underlying asset S, with initial price S_0 and the price at the t-th round denoted as $S_t (= S(tT/\tau))$. For each round t, where $1 \leq t \leq \tau$, the adversary has complete freedom to choose the return of S, given by $R_t \triangleq S_t/S_{t-1} - 1$, within a pre-specified *uncertainty set* $\mathcal{U}_t^\tau$ (we will often drop the τ in $\mathcal{U}_t^\tau$ and use the notation $\mathcal{U}_t$ when no confusion arises). On the other hand, at the beginning of each round, the trader can choose to long[2] Δ_t dollars' worth of the asset (a negative Δ_t will imply a short position). At this point we do not impose any capital capacity on the trader, *i.e.,* Δ_t can be as large or small as possible; however, we shall soon see from our analysis that the optimal Δ_t is bounded and can be explicitly found.

Let us first describe what should be the upper and lower bounds of the option price under no arbitrage assumption in our model. To better illustrate our ideas, all of our analysis will assume the risk-free interest rate is 0. But all our results can be adapted easily to non-zero interest rates.

Upper bound. Suppose the trader shorts the option at time 0. To hedge his/her position, at each round of the game, the trader decides to buy Δ_t dollars' worth of the underlying asset. The cumulated payoff to the trader at time T is then given by $\sum_{t=1}^{\tau}(R_t\Delta_t) - g\left(S_0 \prod_{t=1}^{\tau}(1+R_t)\right)$, plus the option price that he/she gets from selling the option at time 0. Since a rational trader will strive to maximize gain, against an adversary that strives to minimize so, the outcome of this game to the trader will be $\max_{\Delta_t, t\in[\tau]} \min_{R_t\in\mathcal{U}_t, t\in[\tau]} \sum_{t=1}^{\tau}(R_t\Delta_t) - g\left(S_0 \prod_{t=1}^{\tau}(1+R_t)\right)$, again plus the initial option price. Now, if the option price is strictly higher than $-\max_{\Delta_t, t\in[\tau]} \min_{R_t\in\mathcal{U}_t, t\in[\tau]} \sum_{t=1}^{\tau}(R_t\Delta_t) - g\left(S_0 \prod_{t=1}^{\tau}(1+R_t)\right)$, then shorting the option and carrying out the optimal hedging strategy gives the trader a positive gain at time 0 with no risk, *i.e.,* an arbitrage

[1] We will study two natural jump models, and the contributions depend on the specific models we study.

[2] We adopt the terminology in finance: to "long" means to buy, and to "short" means to sell a product.

opportunity arises. In other words, the option price at time 0 cannot be higher than

$$\min_{\Delta_t, t\in[\tau]} \max_{R_t\in\mathcal{U}_t, t\in[\tau]} g\left(S_0\prod_{t=1}^{\tau}(1+R_t)\right) - \sum_{t=1}^{\tau}(R_t\Delta_t). \tag{1}$$

Lower bound. The no-arbitrage lower bound can be obtained by reversing the action of the trader from shorting to longing the option at time 0. Suppose the trader shorts Δ_t dollars' worth of the underlying asset at the t-th round, and strives to maximize $g\left(S_0\prod_{t=1}^{\tau}(1+R_t)\right) - \sum_{t=1}^{\tau}(R_t\Delta_t)$. It can be argued similarly that the option price cannot be lower than

$$\max_{\Delta_t, t\in[\tau]} \min_{R_t\in\mathcal{U}_t, t\in[\tau]} g\left(S_0\prod_{t=1}^{\tau}(1+R_t)\right) - \sum_{t=1}^{\tau}(R_t\Delta_t) \tag{2}$$

or otherwise arbitrage opportunity arises.

We note that this model is closely related to the one studied in [AFW12]. There, they impose a round-wise second moment constraint on the adversary. When the payoff function is convex, both constraints are equivalent algorithmically (see Section 3.4). When the payoff function is non-convex, the adversary may want to impose sparse jumps in its optimal strategy in [AFW12]'s model (and thus violating our uncertainty set constraints) but then second moment would not be a natural constraint for jumps (see discussions in [CT12]). Section 6 discusses how we incorporate price jumps in our model.

Interpretation of bounds. We remark that the definition of upper and lower bounds here is different from what one often sees in Computer Science. When there is a "gap" between upper and lower bounds, it does not mean the bounds are not "tight". Instead, we shall interpret our bounds as follows: let u be the upper bound from (1) and ℓ be the lower bound from (2). We have:

- When an option price does not fall into $[\ell, u]$, then there exists a trading strategy so that under *any* adversarial scenarios the overall payoff of the strategy is strictly positive (arbitrage exists). Economically speaking, this is a "wrong" price of the option.
- When an option price is in $[\ell, u]$, then for any trading strategy, there exists an adversary so that the payoff is non-positive (arbitrage does not exist). This price can be the "fair" price of the option.

Oracle model for payoff functions. We assume we have oracle access to the payoff function $g(\cdot)$ (which is possibly non-convex). Also, when saying $g(\cdot)$ is Lipschitz continuous, we mean that for any real values x, y, we have $|g(y) - g(x)| \leq L|y - x|$ for some constant L.

Continuous time limit. To take continuous-time limit, we fix T and take $\tau \to +\infty$. Let $\left\{\{\mathcal{U}_t^{\tau}\}_{t\leq\tau}\right\}_{\tau\geq 1}$ be a sequence of uncertainty set collections. We say an option's upper (lower) bound has continuous-time limit with respect to $\left\{\{\mathcal{U}_t^{\tau}\}_{t\leq\tau}\right\}_{\tau\geq 1}$ when the upper (lower) bound with uncertainty sets $\{\mathcal{U}_t^{\tau}\}_{t\leq\tau}$ converges as $\tau \to \infty$.

We remark that the user of the model should not interpret the discrete-time multi-round model as only a "discrete approximation" of the continuous time case. Instead, the user can choose an arbitrary τ and then estimate the uncertainty sets accordingly. So long as the uncertainty sets accurately bound the movement of the stock price, our upper and lower bounds will be legitimate. It is expected that the upper and lower bounds we derive here have the highest quality when τ is suitably large and $\mathcal{U}_t^{\tau}$ can be accurately estimated. When τ is too large, it might become more difficult to estimate $\mathcal{U}_t^{\tau}$, and the quality of the corresponding upper (lower) bound could be worse. The practical choice of such an optimal τ is out of the scope of this paper.

We also remark that we leave the choice of the uncertainty set $\mathcal{U}_t^{\tau}$ to depend on the round t. A reasonable approach to construct uncertainty sets is of course to set all $\mathcal{U}_t^{\tau}$ to be equal over time, in which case we shall merely denote $\mathcal{U}^\tau$ (or even $\mathcal{U}$). Putting time dependence on the uncertainty sets nevertheless allows modeling flexibility, and is indicative of state-dependent model, *i.e.,* the uncertainty set can depend on the current underlying asset price. The latter is the robust analog of continuous-time models that have state-dependent volatilities, such as the Heston model [Hes93].

3 Explicit characterization of equilibrium for the hedging game

In this section we will analyze the equilibrium between the trader and the nature in every single round. In particular, we will obtain an explicit optimal hedging strategy for the trader. It then follows from standard results that the continuous-time limit of the equilibrium is the Black-Scholes price.

We will start with the standard binomial tree model and connect its analysis to our adversary model. Doing so gives us insight on how we shall "battle" against the adversary in more general scenarios.

3.1 Binomial tree as a weak adversary model

Single-round case. Suppose there is only one round of the game, *i.e.,* $\tau = 1$. Here the trader needs only to decide Δ, the amount of the underlying asset S to hold for hedging. In the standard single-round binomial model (see Chapter 12 in [Hul09]), the stock price either goes up by a factor of $(1+u)$ or down by a factor of $(1-d)$ (see Figure 2). In the literature, it is typically assumed that the movement of S is stochastic, i.e. with certain probability S goes up and another probability it goes down. The idea of perfect hedging [CRR79] is to pick Δ such that the total payoff at time 1 is constant, regardless of the movement of S. In other words, set Δ that satisfies

$$g(S_0(1+u)) - \Delta u = g(S_0(1-d)) - \Delta(-d),$$

which gives $\Delta = (g(S_0(1+u)) - g(S_0(1-d)))/(u+d)$. Under this hedging strategy Δ, there is no risk for the trader to be compensated

$$g(S_0(1+u)) - \Delta u = g(S_0(1-d)) - \Delta(-d) = \frac{d}{u+d} g(S_0(1+u)) + \frac{u}{u+d} g(S_0(1-d)). \tag{3}$$

Suppose the option price is different from (3), then an arbitrage opportunity must exist. If the price is higher, the trader shorts the option and longs Δ dollars' worth of the underlying asset, whereas if the price is lower, the trader longs the option and shorts the same amount of the underlying asset. Both cases lead to risk-free gain to the trader.

Let us now go back to our model described in Section 2, with an uncertainty set $\mathcal{U} = \{u, -d\}$. This is the same as the standard binomial model except that stochasticity of the underlying asset price is now replaced by adversarial movement. The upper bound (1) becomes

$$\min_{\Delta} \max_{R \in \{u,-d\}} g(S_0(1+R)) - R\Delta. \tag{4}$$

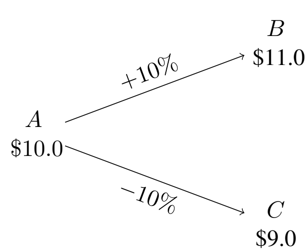

Figure 2: A one step binomial tree model. The uncertainty set in this example is $\{-10\%, 10\%\}$. The price of the underlying asset is $10. At the end of the game, the price can either move to $11 or $9.

It is easy to observe that (4) reaches optimum when we set Δ such that

$$g(S_0(1+u)) - \Delta u = g(S_0(1-d)) - \Delta(-d), \tag{5}$$

or $\Delta = (g(S_0(1+u)) - g(S_0(1-d)))/(u+d)$, leading to the same hedging strategy as the standard (stochastic) binomial model. The same argument works for the lower bound (2) and gives rise to the same hedging strategy. We thus have our first basic conclusion: If the uncertainty set in our hedging game is binomial, the upper bound of the option price *matches* the lower bound; moreover, this unique price is the same as the price concluded from the standard (stochastic) binomial model.

We also make another observation on the form of our optimal value *i.e.,* the equilibrium. Since the optimal hedging amount Δ for (1) and (2) are both equal to the standard binomial model, their corresponding optimal values are both given by (3), which can be written as

$$\min_{\Delta} \max_{R\in\{u,-d\}} g(S_0(1+R)) - R\Delta = \max_{\Delta} \min_{R\in\{u,-d\}} g(S_0(1+R)) - R\Delta = \mathrm{E}_{\Sigma}[g(S_0(1+R))] \tag{6}$$

where Σ assigns probability $d/(u+d)$ to upward movement u and probability $u/(u+d)$ to downward movement d. It is easily observed that Σ is "risk-neutral", *i.e.,* $\mathrm{E}_{R\leftarrow\Sigma}[R] = 0$. Hence in this particular case our upper and lower bounds of the option price are both characterized by the same risk-neutral probability measure on the option payoff. We will see that in more general scenarios, the option price bounds can still be characterized by risk-neutral measures, but the measures can be different from each others, and also they can be both different from the measure used in standard binomial pricing. Note that these risk-neutral measures act as analytical artifacts and do not have a real-world correspondence; they will play a key role in our analysis in the rest of this paper.

Multi-round case. Keeping in mind the result above for the single-round case, our price bounds for the multi-round setting can be obtained through straightforward backward induction (dynamic programming). Suppose the game has τ rounds and each round entails either an up or a down movement for the stock, i.e. $\mathcal{U}_t = \{u, -d\}$. The upper bound (1) can be written as

$$\begin{aligned}
&\min_{\Delta_t, t\in[\tau]} \max_{R_t\in\{u,-d\}, t\in[\tau]} g\left(S_0 \prod_{t=1}^{\tau}(1+R_t)\right) - \sum_{t=1}^{\tau}(R_t\Delta_t) \\
= \;& \min_{\Delta_1} \max_{R_1\in\{u,-d\}} \{\cdots \min_{\Delta_{\tau-1}} \max_{R_{\tau-1}\in\{u,-d\}} \{\cdots \min_{\Delta_\tau} \max_{R_\tau\in\{u,-d\}} \{g\left(S_{\tau-1}(1+R_\tau)\right) - R_\tau\Delta_\tau\} \\
& -R_{\tau-1}\Delta_{\tau-1}\} \cdots - R_1\Delta_1\}
\end{aligned} \tag{7}$$

The quantity (7) can be solved by iteratively computing $g_\tau(S) = g(S)$ and

$$g_{t-1}(S) = \min_{\Delta_t} \max_{R_t\in\{u,-d\}} g_t(S(1+R_t)) - R_t\Delta_t.$$

for $t = \tau, \tau-1, \ldots, 1$. By our result above, the solutions to each of these minimax problems are given by $\mathrm{E}_{\Sigma}[g_t(S(1+R))]$. Hence our upper bound again matches the lower bound, and they both match the price according to the standard binomial tree model.

Figure 3 illustrates an example with two rounds. A denotes the state at time 0, B and C at time 1 and so on. We first compute the price of the option at B, by analyzing the one-stage game assuming B is the initial time point. The same argument applies to state C. We then price the option with initial state A by

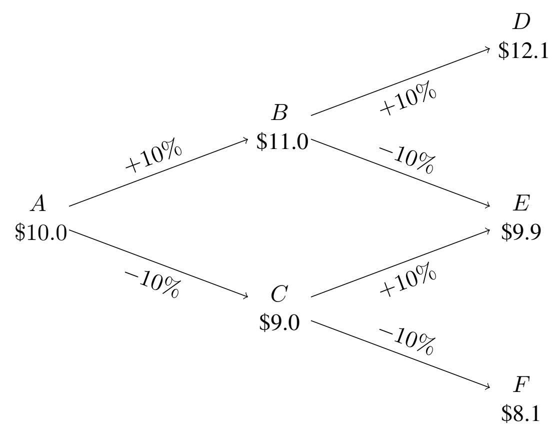

Figure 3: A two-step binomial tree model.

taking into account the maximal gain the trader can make in the future when the next states are in B or C. Observe that there are in total 3 states in this example, instead of $2^2 = 4$, at the end of the second round. In general, the number of states at the final level of a binomial tree grows linearly with the depth of the tree, which makes it a feasible device for option pricing.

Binomial tree at the limit. Since our adversary binomial tree model is in effect the same as the standard model, the continuous-time limit converges to the Black-Scholes price under appropriate scaling. The following result is a rephrase of the well-known result in the literature:

Proposition 3.1. *Consider the τ-round European option game. Let the uncertainty set for each round be $\mathcal{U}^\tau = \{u/\sqrt{\tau}, -d/\sqrt{\tau}\}$. Let $g(\cdot)$ be an arbitrary Lipschitz continuous payoff function. The upper and lower bounds of the option with respect to $\mathcal{U}^\tau$ are the same for any τ and they both converge to the Black-Scholes price as $\tau \to \infty$.*

3.2 A general characterization of equilibrium

We now state our characterization results for a single-round game under very general conditions on the uncertainty set and the payoff function. We will concentrate on the upper bound (1) in our analysis; the lower bound (2) can be obtained easily merely by replacing g by $-g$, and will be discussed at the end of this section.

Suppose $\tau = 1$, and let $\mathcal{U}$ be a (Borel) measurable set. Our goal is to find the optimal solution of $\min_\Delta \max_{R \in \mathcal{U}} g(S_0(1+R)) - R\Delta$. The following result provides an optimality characterization for any payoff function g and any uncertainty set $\mathcal{U}$:

Proposition 3.2. *Let $\tau = 1$ and S_0 be the initial price. Consider a bounded uncertainty set $\mathcal{U}$ and a continuous payoff function g. We have*

$$\min_{\Delta} \max_{R \in \mathcal{U}} g(S_0(1+R)) - R\Delta = \max_{\substack{P_f \in \mathcal{P}(\mathcal{U}), \\ \mathrm{E}_{R \leftarrow P_f}[R]=0}} \mathrm{E}_{P_f}[g(S_0(1+R))], \tag{8}$$

where $\mathcal{P}(\mathcal{U})$ denotes the set of all probability measures P_f that have support on $\mathcal{U}$. The maximization problem in the right hand side above is over all such probability measures that satisfy $\mathrm{E}_{P_f}[R] = 0$.

The optimal value of (8) is finite only when $\mathcal{U}$ contains at least a point larger than 0 and a point smaller than 0, *e.g.,* when $\mathcal{U}$ is an interval that covers 0; otherwise risk-neutral measure cannot be constructed.

Proof. To illustrate the key idea in our analysis, let us start with analyzing a "discrete" version of the problem, *i.e.,* let $\mathcal{U} = \{r_1, r_2, ..., r_n\}$ be a discrete set on $\mathbb{R}$. We can write $\min_\Delta \max_{R\in\mathcal{U}} g(S_0(1+R))-R\Delta$ as the following linear program (LP):

$$\begin{array}{ll} \min & p \\ \text{s.t.} & g(S_0(1+r_i)) - r_i\Delta \le p \;\text{ for } i \in [n] \end{array} \equiv \begin{array}{ll} \min & p \\ \text{s.t.} & p + r_i\Delta \ge g(S_0(1+r_i)) \;\text{ for } i \in [n], \end{array} \tag{9}$$

where the decision variables are p and Δ. The formulations in (9) follow simply by the definition of minimax problem, with the optimal p representing the option price's upper bound. Now we invoke the standard primal-dual theorem for LP on (9) and obtain the following equivalent LP in dual form:

$$\begin{array}{ll} \text{maximize} & \sum_{i\in[n]} w_i g(S_0(1+r_i)) \\ \text{subject to} & \sum_{i\in[n]} w_i = 1 \\ & \sum_{i\in[n]} w_i r_i = 0 \\ & w_i \ge 0 \;\text{ for } i \in [n] \end{array} \tag{10}$$

Here w_i's are the decision variables. Observe that $\{w_i\}_{i\in[n]}$ can be interpreted as a probability distribution on the uncertainty set $\mathcal{U}$ since $\sum w_i = 1$. Call this distribution P_f. This is a risk-neutral probability distribution since the expected return $\mathrm{E}_{P_f}[R] = \sum w_i r_i = 0$. Moreover, note that the objective function under this probability interpretation can be rewritten as $\sum_{i\in[n]} w_i g(S(1+r_i)) = \mathrm{E}_{R\leftarrow P_f}[g(S(1+R))]$. We thus have proved (8).

It is straightforward to generalize this argument to general uncertainty set, with slightly more function space technicalities. For general uncertainty set $\mathcal{U}$, we can generalize (9) as:

$$\begin{array}{ll} \text{minimize} & p \\ \text{subject to} & p + r\Delta \ge g(S_0(1+r)) \;\text{ for } r \in \mathcal{U} \end{array} \tag{11}$$

Next, note that the dual cone of $\mathcal{C}^+(\mathcal{U})$, the set of non-negative continuous functions on $\mathcal{U}$, is $P^+(\mathcal{U})$, the set of positive measures on $\mathcal{U}$. Since g is assumed to be continuous, the Lagrangian of (11) is

$$L(p, \Delta, w) = p + \int_{\mathcal{U}} (g(S_0(1+r)) - r\Delta - p)dw(r)$$

where $w(\cdot) \in \mathcal{P}^+(\mathcal{U})$ (see [Lue97]). The dual function is defined as

$$\begin{aligned} \ell(w) &= \inf_{p,\Delta} \left\{ p + \int_{\mathcal{U}} (g(S_0(1+r)) - r\Delta - p)dw(r) \right\} \\ &= \inf_{p,\Delta} \left\{ \int_{\mathcal{U}} g(S_0(1+r))dw(r) + (1 - w(\mathcal{U}))p + \int_{\mathcal{U}} r dw(r)\Delta \right\} \end{aligned} \tag{12}$$

Suppose $w(\cdot)$ does not satisfy either $w(\mathcal{U}) = 1$ or $\int_{\mathcal{U}} r dw(r) = 0$, then one can always find p or Δ that gives arbitrarily large objective value in (12). Hence the dual problem $\max_{w\in\mathcal{P}^+(\mathcal{U})} \ell(w)$ can be written as

$$\begin{array}{ll} \text{maximize} & \int_{\mathcal{U}} g(S_0(1+r))dw(r) \\ \text{subject to} & \int_{\mathcal{U}} r dw(r) = 0 \\ & w(\mathcal{U}) = 1 \\ & w \in \mathcal{P}^+(\mathcal{U}) \end{array} \tag{13}$$

Finally, it is easy to see that the constraint set in (11) has non-empty interior (by picking large enough p for example). Hence strong duality holds and the dual optimal value in (13) equals the primal counterpart (see *e.g.,* Chapter 8 in [Lue97]). By identifying w as a probability measure on $\mathcal{U}$, we conclude that (13) is the same as (8). This completes our proof. □

3.3 Convex payoff function

When the payoff function $g(\cdot)$ is convex and the uncertainty set is an interval, *i.e.,* $\mathcal{U} = [-\underline{\zeta}, \overline{\zeta}]$ with $\underline{\zeta}, \overline{\zeta} > 0$, we are able to find an analytic form for (8). Specifically, we will show that it is sufficient to consider risk-neutral probability distributions that have point masses concentrated only on the extremes, namely $-\underline{\zeta}$ and $\overline{\zeta}$, *i.e.,*

Corollary 3.3. *When the payoff function $g(\cdot)$ is convex and $\mathcal{U} = [-\underline{\zeta}, \overline{\zeta}]$, we have*

$$\min_{\Delta} \max_{R \in \mathcal{U}} g(S_0(1+R)) - R\Delta = \max_{\substack{P_f \in \mathcal{P}(\{-\underline{\zeta}, \overline{\zeta}\}), \\ \mathrm{E}_{R \leftarrow P_f}[R] = 0}} \mathrm{E}_{P_f}[g(S_0(1+R))] \tag{14}$$

where $\mathcal{P}(\{-\underline{\zeta}, \overline{\zeta}\})$ is the set of probability distributions that have support only on $-\underline{\zeta}$ and $\overline{\zeta}$. Furthermore, since $P_f \in \mathcal{P}(\{-\underline{\zeta}, \overline{\zeta}\})$ and $\mathrm{E}_{P_f}[R] = 0$ uniquely defines P_f, the max *operator is redundant. Hence (14) can be rewritten as*

$$\mathrm{E}_{P_f}[g(S_0(1+R)], \quad \textit{where } P_f \in \mathcal{P}(\{-\underline{\zeta}, \overline{\zeta}\}) \textit{ and } \mathrm{E}_{P_f}[R] = 0. \tag{15}$$

An immediate implication of Corollary 3.3 is that the upper bound of the option price collides with the binomial model for a single-round game.

Proof. We can prove (14) by analyzing either the primal program (11) in the proof of Proposition 3.2 or the characterization (8) directly. Let us consider the former as this is more elementary. We argue that, in the case of convex g and $\mathcal{U} = [-\underline{\zeta}, \overline{\zeta}]$, the program (11) is equivalent to

$$\begin{array}{ll} \text{minimize} & p \\ \text{subject to} & p - \underline{\zeta}\Delta \geq g(S(1-\underline{\zeta})) \\ & p + \overline{\zeta}\Delta \geq g(S(1+\overline{\zeta})) \end{array} \tag{16}$$

In other words, all other constraints $p + r\Delta \geq g(S(1+r))$ for $r \in (-\underline{\zeta}, \overline{\zeta})$ are redundant. To prove this, consider any $-\underline{\zeta} < r < \overline{\zeta}$. One can write $r = \underline{q}(-\underline{\zeta}) + \overline{q}\overline{\zeta}$ where $\underline{q} + \overline{q} = 1$, $\underline{q}, \overline{q} > 0$. Suppose the inequalities $p - \underline{\zeta}\Delta \geq g(S(1-\underline{\zeta}))$ and $p + \overline{\zeta}\Delta \geq g(S(1+\overline{\zeta}))$ hold. Then

$$\begin{aligned} p + r\Delta &= \underline{q}(p - \underline{\zeta}\Delta) + \overline{q}(p + \overline{\zeta}\Delta) \\ &\geq \underline{q}g(S(1-\underline{\zeta})) + \overline{q}g(S(1+\overline{\zeta})) \\ &\geq g(S(1+r)) \;\; \text{(by the convexity of } g) \end{aligned}$$

Therefore, all other constraints are redundant. Now by the same argument as the proof of Proposition 3.2 (for discrete uncertainty set), we immediately get (14). The other statement in the corollary follows trivially. □

3.4 Analysis for the multi-round model

For our multi-round game, the trader has the discretion to choose τ rounds of hedging amount $\{\Delta_t\}_{t \in [\tau]}$ against the nature who controls the τ rounds of returns $\{R_t\}_{t \in [\tau]}$. We assume a uniform uncertainty set

$\mathcal{U} = [-\underline{\zeta}, \overline{\zeta}]$ across time. The upper bound of the option price is then

$$\begin{aligned} & \min_{\Delta_t, t\in[\tau]} \max_{R_t\in\mathcal{U}, t\in[\tau]} g\left(S_0 \prod_{t=1}^{\tau}(1+R_t)\right) - \sum_{t=1}^{\tau}(R_t\Delta_t) \\ = & \min_{\Delta_1} \max_{R_1\in\mathcal{U}} \{\cdots \min_{\Delta_{\tau-1}} \max_{R_{\tau-1}\in\mathcal{U}} \{\cdots \min_{\Delta_\tau} \max_{R_\tau\in\mathcal{U}} \{g\left(S_{\tau-1}(1+R_\tau)\right) - R_\tau\Delta_\tau\} \\ & \qquad -R_{\tau-1}\Delta_{\tau-1}\}\cdots - R_1\Delta_1\} \end{aligned} \tag{17}$$

The following lemma is a consequence of the result for the single-round game in Proposition 3.2.

Lemma 3.4. *Consider the τ-round hedging game with the same uncertainty set $\mathcal{U} = [-\underline{\zeta}, \overline{\zeta}]$ across time and convex payoff function $g(\cdot)$. The upper bound of the option price is $\mathrm{E}_{P_f}[g(S_0 \prod_{t=1}^{\tau}(1+R_t))]$, where P_f is the unique risk-neutral probability distribution on $\{-\underline{\zeta}, \overline{\zeta}\}$ for all $\{R_t\}_{t\in[\tau]}$, i.e., $P_f \in (\{-\underline{\zeta}, \overline{\zeta}\})$ and $\mathrm{E}_{P_f}[R_t] = 0$ for all $t \in [\tau]$. In other words, it coincides with the price computed from a binomial tree model with $u = \overline{\zeta}$ and $d = \underline{\zeta}$.*

The proof is a direct application of dynamic programming, coupled with the preservation of convexity across iterations of the value functions. First, observe the following:

Lemma 3.5. *Let $h(\cdot)$ be an arbitrary convex function. Then $\mathrm{E}_P[h(S(1+R))]$ is convex in S, where P is an arbitrary distribution for R.*

Proof. The statement is immediate by using linearity of expectations and the assumption that $h(\cdot)$ is convex. □

Proof of Lemma 3.4. We can write (17) as a dynamic program, given by $g_\tau(x) = g(x)$ and

$$g_{t-1}(x) = \min_{\Delta_t} \max_{R_t\in\mathcal{U}} g_t(x(1+R_t)) - R_t\Delta_t$$

for $t = \tau, \tau-1, \ldots, 1$. The price upper bound is then given by $g_0(S_0)$. We prove by induction that $g_t(\cdot)$ are all convex and $g_{t-1}(x) = \mathrm{E}_{P_f}[g_t(x(1+R_t))]$. The statement is obvious for $g_\tau(\cdot)$. Now, supposing $g_t(\cdot)$ is convex, we have from Corollary 3.3 that $g_{t-1}(x) = \mathrm{E}_{P_f}[g_t(x(1+R_t))]$, and from Lemma 3.5 that $g_{t-1}(\cdot)$ is convex. Hence the induction holds. □

Explicit hedging strategy. When the uncertainty sets are uniform intervals and the payoff function is convex, the optimal hedging strategy is straightforward (and is identical to the binomial model): $\Delta_t = \frac{g(S_{t-1}(1+\overline{\zeta})) - g(S_{t-1}(1-\underline{\zeta}))}{\underline{\zeta}+\overline{\zeta}}$ dollar on S for each round t.

Non-uniform uncertainty sets. When the uncertainty sets are non-uniform, say $\mathcal{U}_t \triangleq [-\underline{\zeta}_t, \overline{\zeta}_t]$, Corollary 3.3, Lemma 3.5 and the form of the hedging strategy all still hold with the natural modification. This means at each round we need only consider the two points $\{-\underline{\zeta}_t, \overline{\zeta}_t\}$. However, from a computational point of view, the number of states we need to keep track of in the backward induction could grow exponentially in τ (See Figure 4 for an illustration). Thus, a naive application of dynamic programming algorithm will not be efficient. In Appendix A we show that exact computation of the option's upper bound with non-uniform uncertainty sets is $\#P$-hard, and in Section 4 we shall design an approximation algorithm to solve the problem.

Algorithmic equivalence to [AFW12]. We remark that this backward induction approach is also applicable to the model in [AFW12], in which second moment constraints are imposed on R_t. In particular,

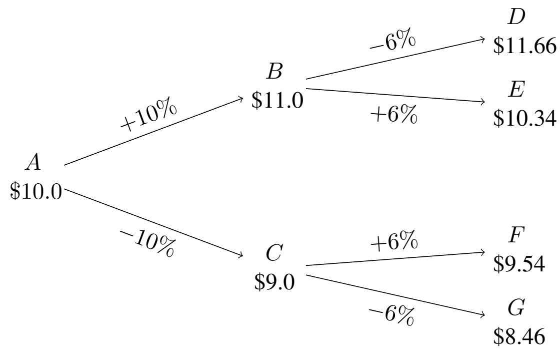

Figure 4: An example of an option hedging game with non-uniform uncertainty sets. $\mathcal{U}_1 = [-10\%, 10\%]$ and $\mathcal{U}_2 = [-6\%, 6\%]$. The number of states for the dynamic program can grow exponentially in τ.

our stepwise dual characterization can be derived in a similar manner under their moment constraint, by applying (with a small modification) the primal-dual theorem to find the optimal solution for one-round game, followed by using Lemma 3.5 to show the preservation of convexity over rounds (so that the primal-dual theorem can be applied recursively). Thus, the results here and techniques developed in Section 4 also give efficient pricing algorithm and an explicit hedging strategy for the model and problems considered in [AFW12]. But we also remark that the convergence result developed below does not directly imply convergence in [AFW12]'s model.

Convergence. Since the price from the binomial tree model converges to Black-Scholes [CRR79], an immediate consequence is the convergence of our upper bound, with uncertainty sets $\mathcal{U}_t^\tau = [-\underline{\zeta}/\sqrt{\tau}, \overline{\zeta}/\sqrt{\tau}]$, also to the Black-Scholes price. Here we state a convergence result that is more general: as long as the (possibly non-uniform) collection of uncertainty sets follow a "bounded quadratic variation" condition, we obtain convergence to the Black-Scholes price for European-type options:

Corollary 3.6. *Consider the τ-round hedging game with Lipschitz continuous convex payoff function $g(\cdot)$. Let $\{\{\mathcal{U}_t^\tau\}_{t\le\tau}\}_{\tau\ge 1}$ be the sequence of uncertainty sets and let $\mathcal{U}_t^\tau = [-\underline{\zeta}_t^\tau, \overline{\zeta}_t^\tau]$. If $\lim_{\tau\to\infty}\sum_{t=1}^{\tau}\underline{\zeta}_t^\tau\overline{\zeta}_t^\tau = \nu$ for a positive number ν, and $\sup_{t\in[\tau]}\max\{\underline{\zeta}_t^\tau, \overline{\zeta}_t^\tau\} \to 0$, the upper bound of the European option price converges to $\mathrm{E}[g(S_0\exp\{\sqrt{\nu}N(0,1) - \nu/2\})]$, where $N(0,1)$ is standard Gaussian variable,* i.e., *it converges to the option price for a geometric Brownian motion with zero drift.*

We remark that if $\nu = \sigma^2 T$ for a positive constant σ^2, then the condition $\lim_{\tau\to\infty}\sum_{t=1}^{\tau}\underline{\zeta}_t^\tau\overline{\zeta}_t^\tau = \nu$ imitates the quadratic variation of a Brownian motion, and the result recovers the Black-Scholes price. The uniform convergence condition $\sup_{t\in[\tau]}\max\{\underline{\zeta}_t^\tau, \overline{\zeta}_t^\tau\} \to 0$ is necessary; there is no guarantee of Gaussian convergence in the limit if one uncertainty set keeps constant size as $\tau \to \infty$.

Proof. From Lemma 3.4, the upper bound of the option price, for any τ, is $\mathrm{E}_{P_f}[g(S_0\prod_{t=1}^{\tau}(1+R_t^\tau))]$ where R_t^τ is the t-th round return in a τ-round game, and P_f (which also depends on τ) is the unique probability measure that satisfies $\mathrm{E}_{P_f}[R_t^\tau] = 0$ and has support $\{-\underline{\zeta}_t^\tau, \overline{\zeta}_t^\tau\}$ on R_t^τ for any $t \in [\tau]$. Simple calculation reveals that P_f puts weights $\overline{\zeta}_t^\tau/(\underline{\zeta}_t^\tau + \overline{\zeta}_t^\tau)$ on $-\underline{\zeta}_t^\tau$ and $\underline{\zeta}_t^\tau/(\underline{\zeta}_t^\tau + \overline{\zeta}_t^\tau)$ on $\overline{\zeta}_t^\tau$. This implies that $\mathrm{E}_{P_f}[(R_t^\tau)^2] = \underline{\zeta}\overline{\zeta}$.

We shall prove that

$$\log(S_0\prod_{t=1}^{\tau}(1+R_t^\tau)) = \log S_0 + \sum_{t=1}^{\tau}R_t^\tau - \sum_{t=1}^{\tau}\frac{(R_t^\tau)^2}{2} + \sum_{t=1}^{\tau}\frac{\xi(R_t^\tau)}{3} \tag{18}$$

where $\xi(R_t)$ satisfies $|\xi(R_t^\tau)| \le C|R_t^\tau|^3$ for a constant C, converges in distribution to $\log S_0 + N(0,1) - \nu/2$.

Now consider each term in (18), and we start with $\sum_{t=1}^\tau R_t^\tau$. Since $\sum_{t=1}^\tau \mathrm{E}_{P_f}[(R_t^\tau)^2] = \sum_{t=1}^\tau \underline{\zeta}_t^\tau \overline{\zeta}_t^\tau \to \nu$, and $\sum_{t=1}^\tau \mathrm{E}_{P_f}[(R_t^\tau)^2; |R_t^\tau| > \epsilon]$ is eventually zero as $\tau \to \infty$, for any $\epsilon > 0$, by Lindeberg-Feller Theorem (p. 114, (4.5) in [Dur10]), we have $\sum_{t=1}^\tau R_t^\tau \to \sqrt{\nu} N(0,1)$ in distribution.

Next consider the term $\sum_{t=1}^\tau \frac{(R_t^\tau)^2}{2}$. By our condition $\sup_{t\in[\tau]} \max\{\underline{\zeta}_t^\tau, \overline{\zeta}_t^\tau\} \to 0$, since $\sum_{t=1}^\tau \Pr(|R_t^\tau| > \epsilon)$ is eventually zero as $\tau \to \infty$, for any $\epsilon > 0$, and also $\sum_{t=1}^\tau \mathrm{E}_{P_f}(R_t^\tau)^4 \le \sum_{t=1}^\tau \mathrm{E}_{P_f}(R_t^\tau)^2 \cdot \sup_{t\in[\tau]}(R_t^\tau)^2 \to 0$, the Weak Law for Triangular Arrays hold (p. 40, (5.5) in [Dur10]), and $\sum_{t=1}^\tau (R_t^\tau)^2 \to \sum_{t=1}^\tau \mathrm{E}_{P_f}[(R_t^\tau)^2] = \nu$ in probability.

For the last term, we have $|\sum_{t=1}^\tau \xi(R_t^\tau)| \le C\sum_{t=1}^\tau (R_t^\tau)^2 \sup_{t\in[\tau]} |R_t^\tau| \to 0$ in probability. Combining all these terms, by Slutsky's Theorem (see *e.g.,* p. 19 in [Ser80]), we conclude that $\log(S_0 \prod_{t=1}^\tau (1 + R_t^\tau)) \to \sqrt{\nu} N(0,1) - \nu/2$ in distribution.

Lastly, we will conclude our result by checking a uniform integrability condition (see *e.g.,* p. 14 in [Ser80]). First, since g is continuous, the Continuous Mapping Theorem [Bil09] stipulates that $g(S_0 \prod_{t=1}^\tau (1+ R_t^\tau))$ converges to $g(\exp\{\sqrt{\nu} N(0,1) - \nu/2\})$ in distribution. We shall show that $\sup_\tau E_{P_f}[g(S_0 \prod_{t=1}^\tau (1 + R_t^\tau))^2] < \infty$. This will imply that $g(S_0 \prod_{t=1}^\tau (1 + R_t^\tau))$ is uniformly integrable, which will then conclude the convergence in L_1 of $g(S_0 \prod_{t=1}^\tau (1 + R_t^\tau))$ into $g(S_0 \exp\{\sqrt{\nu} N(0,1) - \nu/2\})$ and conclude our result. To this end, note that

$$
\begin{aligned}
& \mathrm{E}_{P_f}\left[g(S_0 \prod_{t=1}^\tau (1 + R_t^\tau))^2 \right] \\
\le\ & C_1 \mathrm{E}_{P_f} \left| S_0 \prod_{t=1}^\tau (1 + R_t^\tau) \right|^2 + C_2 \\
& \text{(for some constants } C_1, C_2 > 0, \text{ since } g \text{ is assumed to be Lipschitz continuous)} \\
=\ & C_1 S_0 \prod_{t=1}^\tau E_{P_f}(1 + R_t^\tau)^2 + C_2 \ \ \text{(by independence of } R_t^\tau) \\
=\ & C_1 S_0 \prod_{t=1}^\tau (\underline{\zeta}_t^\tau \overline{\zeta}_t^\tau + 1) + C_2 \\
\le\ & C_1 S_0 \exp\{\sum_{t=1}^\tau \underline{\zeta}_t^\tau \overline{\zeta}_t^\tau\} + C_2 \\
<\ & C_3
\end{aligned}
$$

for some constant $C_3 > 0$, by our assumption that $\sum_{t=1}^\tau \underline{\zeta}_t^\tau \overline{\zeta}_t^\tau \to \nu$.

□

Concave payoffs. We have a simple characterization of the hedging game's equilibrium when the payoff function is concave, under general conditions on $\mathcal{U}_t$:

Corollary 3.7. *Consider a τ-round game. When the payoff function $g(\cdot)$ is concave, with uncertainty sets $\{\mathcal{U}_t\}_{t\in[\tau]}$ each of which contains the point 0, the option price's upper bound is*

$$
\min_{\Delta_t, t\in[\tau]} \max_{R_t\in\mathcal{U}_t, t\in[\tau]} g\left(S_0 \prod_{t=1}^\tau (1 + R_t) \right) - \sum_{t=1}^\tau (R_t \Delta_t) = g(S_0) \tag{19}
$$

Proof. Consider a single-round game *i.e.,* $\tau = 1$. Recall Proposition 3.2, which states that the upper bound is $\max_{P_f \in \mathcal{P}(\mathcal{U}): \mathrm{E}_{P_f}[R]=0} E_{P_f}[g(S_0(1 + R))]$. By Jensen's inequality, $\mathrm{E}_{P_f}[g(S_0(1 + R))] \le g(S_0(1 + \mathrm{E}_{P_f}[R])) = g(S_0)$ for any $P_f \in \mathcal{P}(\mathcal{U})$ such that $\mathrm{E}_{P_f}[R] = 0$. The result is then immediate for $\tau = 1$.

The conclusion from multi-round game follows exactly the same as the argument for Lemma 3.4 (now concavity is preserved in every step in the backward induction). □

Lower bounds. By replacing g with $-g$ in all analysis above, we immediately get results for lower bounds. The following is analogous to Proposition 3.2:

Proposition 3.8. *Let $\tau = 1$ and S_0 be the initial price. Consider a bounded uncertainty set $\mathcal{U}$ and a continuous payoff function g. The lower bound of the option price is*

$$\max_{\Delta} \min_{R \in \mathcal{U}} g(S_0(1+R)) - R\Delta = \min_{\substack{P_f \in \mathcal{P}(\mathcal{U}), \\ \mathrm{E}_{R \leftarrow P_f}[R]=0}} \mathrm{E}_{P_f}[g(S_0(1+R))], \tag{20}$$

where $\mathcal{P}(\mathcal{U})$ denotes the set of all probability measures P_f that have support on $\mathcal{U}$. The maximization problem in the right hand side above is over all such probability measures that satisfy $\mathrm{E}_{P_f}[R] = 0$.

The following summarizes the characterizations for convex and concave payoffs:

Corollary 3.9. *Consider the τ-round European option hedging game, with uncertainty sets $\{\mathcal{U}_t\}_{t \in [\tau]}$. The following results hold:*

1. *Suppose the payoff function g is concave. If the uncertainty sets $\mathcal{U}_t = [-\underline{\zeta}_t, \overline{\zeta}_t]$ for all $t \in [\tau]$, then the lower bound is $\mathrm{E}_{P_f}[g(S_0 \prod_{t=1}^{\tau}(1+R_t))]$, where P_f is the unique risk-neutral measure supported on $\{-\underline{\zeta}_t, \overline{\zeta}_t\}$ for each R_t, i.e. $E_{P_f}[R_t] = 0$.*
2. *Suppose the payoff function g is convex, and the uncertainty sets $\mathcal{U}_t$ all contain the point 0. Then the lower bound is $g(S_0)$.*

4 Algorithms for non-convex payoffs

This section presents an approximation algorithm for computing the price upper bounds for general payoff functions under the oracle model; the lower bound's algorithm and its analysis is similar and so is omitted here. Throughout this section, we will assume the size of the uncertainty set $\mathcal{U} = [-\underline{\zeta}, \overline{\zeta}]$ is uniform across time steps and $\underline{\zeta}, \overline{\zeta} > 0$ are polynomials in τ; at the end of this section we will discuss the non-uniform uncertainty set case. We also assume that the payoff function g is Lipschitz continuous and monotonically non-decreasing, and without loss of generality that $g(0) = 0$.

Our main result is an approximation algorithm that has δS_0 additive error with running time being polynomial in $\frac{1}{\delta}$ and τ (we also say such algorithm is *δ-additive approximation*).

Our algorithm. We use a fairly natural algorithm to approximate the upper bound: we discretize the uncertainty set $\mathcal{U}$, *i.e.,* instead of allowing the adversary to choose an arbitrary value from $\mathcal{U}$, we only allow it to choose from the discrete set $\hat{\mathcal{U}} \triangleq \{-\underline{\zeta}, -\underline{\zeta}+\epsilon, -\underline{\zeta}+2\epsilon, ..., \overline{\zeta}\}$, where ϵ is a parameter of our algorithm. In other words, we consider a *multinomial tree approximation* of the problem.

We shall briefly address two important properties of our algorithm regarding the discretization. First, since at each step the adversary may choose multiple ways to move the price, it could be worrying that the number of possible states under consideration is exponential in τ. However, so long as ϵ remains uniform over the rounds, the number of states we need to keep track of at the t-th round is $t(\underline{\zeta} + \overline{\zeta})/\epsilon$ (and thus linear in t). This is also an elegant property of multinomial tree (as well as binomial tree) that makes these models useful in practice. However, note that at the final round, the price S_τ could be any value in an interval of *exponential length* (*i.e.,* in the range $[(1-\underline{\zeta})^\tau S, (1+\overline{\zeta})^\tau S]$), while our multinomial tree algorithm only "samples" polynomial number of points from the function $g(x)$. This implies that on average the distance between any two sampled points is exponential. This is a surprising feature of our algorithm: while a large portion of internal states in the multinomial tree have additive errors being $\gg \delta S_0$, the error *shrinks* over time to give an accurate final output. This feature is a consequence of the probabilistic interpretation in our dual formulation, which we will further elaborate when we discuss the challenges of our analysis shortly.

Second, our discretization scheme and backward induction is related to the so-called stochastic mesh method in the area of financial engineering [BG04]. First, to compute the price upper bound for the multinomial tree, one can use a dynamic program based on the recursive formula discussed in Section 3. Specifically, let $\hat{g}_t(x)$ be the approximate price upper bound of the option at the t-th round. Also, let $b = (\overline{\zeta} + \underline{\zeta})/\epsilon$ be the total number of choices an adversary has for each move, and the choices of the return are $r_i = -\underline{\zeta} + i\epsilon$ for $i \in [b]$. We compute $\hat{g}_t(x)$ by finding the optimal solution of the LP:

$$\begin{array}{ll} \text{minimize} & p \\ \text{subject to} & \hat{g}_{t+1}(x(1+r_i)) - r_i\Delta \leq p \ \text{ for } i \in [b] \end{array} \tag{21}$$

This idea resembles the stochastic mesh method in pricing options, most notably American-type, when the underlying asset's price movement is assumed to be stochastic. Assuming a well-defined risk-neutral measure, the stochastic mesh method calculates each $g_t(x)$ through backward induction via Monte Carlo simulation on the return R_t. However, since the state space of R_t is typically assumed to be unbounded in stochastic models, the stochastic mesh algorithm does not use ϵ-spacings to construct a tree but rather generates the nodes randomly from some convenient distributions. The number of nodes at each level (*i.e.,* time point) is kept constant to avoid exponential computational burden, and each node at each level is "connected" to all the nodes at the neighboring two levels. The backward induction step then utilizes all the nodes in the successive level, and relies on importance sampling to calculate unbiased probability weights at each step. Hence, in a sense, our algorithm replaces this importance sampling with an LP in obtaining the weights. As such, our analysis also differs and, in fact, it avoids some hard-to-verify moment conditions on the transition probability thanks to our assumption of bounded uncertainty sets.

Our key contribution in this section is to analyze the difference between $\hat{g}_t(x)$ and $g_t(x)$. We face the following two challenges.

Local error due to discretization. The first challenge is to understand how much error we make *per step*. For instance, let us assume that we are able to accurately compute $g_{t+1}(x)$ for the $(t+1)$-st round. If we use (21) to find an approximation function $\hat{g}_t(x)$, we need to quantify the difference between $g_t(x)$ and $\hat{g}_t(x)$.

The propagation of errors. In our backward induction, the value of $g_0(S_0)$ depends on the values of $g_1(x_1)$ for $x_1 \in [(1-\underline{\zeta})S_0, (1+\overline{\zeta})S_0]$ and the value of $g_1(x_1)$ for any x_1 depends on the values of $g_2(x_2)$ for $x_2 \in [(1-\underline{\zeta})x_1, (1+\overline{\zeta})x_1]$. Thus, if we look two steps forward, $g_0(S_0)$ depends on $g_2(x)$ for $x \in [(1-\underline{\zeta})^2 S_0, (1+\overline{\zeta})^2 S_0]$. In general, for any t, $g_0(S_0)$ depends on $g_t(x)$ for $x \in [(1-\underline{\zeta})^t S_0, (1+\overline{\zeta})^t S_0]$, which is an exponentially growing interval. This could cause the error to propagate at an exponential rate: if our solution $g_0(S_0)$ (indirectly) depends on $g_t((1+\overline{\zeta})^t S_0)$, then even an ϵ-additive approximation of $g_t((1+\overline{\zeta})^t S_0)$, *i.e.,* $\hat{g}_t((1+\overline{\zeta})^t S_0) \approx g_t((1+\overline{\zeta})^t S_0) \pm \epsilon(1+\overline{\zeta})^t S_0$ can potentially lead to the same error of $\epsilon(1+\overline{\zeta})^t S_0$ for computing $g_0(S_0)$ *i.e.,* an $\epsilon(1+\overline{\zeta})^t$-additive error, which is prohibitively huge as τ becomes large. Thus, the challenge here is to accurately analyze and control the error when the algorithm uses information from a lot of highly noisy internal states. Moreover, as discussed before, the fact that a polynomial number of points are sampled in the multinomial tree algorithm from an exponential length of uncertainty set as time progresses can also result in huge error. The combination of the magnitude of errors from the sources and their propagation needs to be addressed.

Our techniques. Our analysis consists of four major components to address the above two issues. First, we need to show that Lipschitz continuity of $g_t(x)$ is preserved in each backward induction step over time;

this property will be used later on. Second, we need to show that when $g_{t+1}(x)$ is Lipschitz and when we discretize the uncertainty set, the approximate solution we get at each step is sufficiently accurate. This can be done by analyzing the geometric structure of both the primal and the dual LP (for the non-discretized version). This will give us a convenient form of the solutions in terms of their binding constraints, which helps quantify the difference between approximate solutions and the exact solutions locally.

Next, we need to control the error propagation rate. As discussed above, if we only focus on the "dependency graph" of the dynamic program, there will be no way for us to control the error propagation rate. The crux to our analysis is to view the solutions in the dual spaces and express our recursive solutions probabilistically. These probability measures, attached "artificially" to our deterministic algorithm, would bestow martingale property on $\{R_t\}_{t\le\tau}$ that helps to address $g_0(S_0)$'s pathological dependency on $g_\tau((1+\overline{\zeta})^\tau S_0)$: while $g_0(S_0)$ could potentially depend on $g_\tau((1+\overline{\zeta})^\tau S_0)$, it is unlikely that the random walk $\{R_t\}_{t\le\tau}$ will hit the price $(1+\overline{\zeta})^\tau S_0$, and hence $g_0(S_0)$'s dependence on $g_\tau((1+\overline{\zeta})^\tau S_0)$ will also be weak probabilistically. Our *deterministic* error can be expressed in terms of the expectation of the price process under one of these probability measures. We will see that our choice of the most suitable probability measure for analysis is, interestingly, neither the maximal risk-neutral measure for the original problem nor the discretized problem, but in some sense a problem in between.

Finally, we need to "glue" our local analysis and the control of error propagation together. We use a standard telescoping technique to do so.

Now we explain all our building blocks in detail.

4.1 Preservation of Lipschitz continuity for the value functions

We start with the first building block. We have the following lemma.

Lemma 4.1. *Suppose the payoff function $g(\cdot)$ is Lipschitz continuous and monotonically non-decreasing, i.e. $g(x)-g(y)\le L(x-y)$ for $x\ge y$. Then the value function $g_t(\cdot)$ at the t-th round is also Lipschitz continuous with the same Lipschitz constant for all t.*

Proof. We prove by (backward) induction. Obviously $g_\tau(x)=g(x)$ satisfies the Lipschitz condition. Suppose $g_{t+1}(x)$ satisfies $g_{t+1}(x)-g_{t+1}(y)\le L(x-y)$ for $x\ge y$. We prove $g_t(x)-g_t(y)\le L(x-y)$ for $x\ge y$ by contradiction.

Assuming this is not true, then

$$\max_{P\in\mathcal{P}(\mathcal{U}):\mathrm{E}[R]=0}\mathrm{E}[g_{t+1}((x+\delta)(1+R))] > \max_{P\in\mathcal{P}(\mathcal{U}):\mathrm{E}[R]=0}\mathrm{E}[g_{t+1}(x(1+R))]+L\delta$$

for some $x\in\mathcal{U}$ and $\delta>0$. Now let $\tilde{P}_f$ be an optimal solution for $\mathrm{E}[g_{t+1}((x+\delta)(1+R))]$ (the existence of an optimal solution will be seen immediately in the next lemma). Then

$$\begin{aligned}\mathrm{E}_{\tilde{P}_f}\left[g_{t+1}((x+\delta)(1+R))\right] &= \max_{P_f\in\mathcal{P}(\mathcal{U}):\mathrm{E}[R]=0}\mathrm{E}[g_{t+1}((x+\delta)(1+R))]\\ &> \max_{P_f\in\mathcal{P}(\mathcal{U}):\mathrm{E}[R]=0}\mathrm{E}[g_{t+1}(x(1+R))]+L\delta\\ &\ge \mathrm{E}_{\tilde{P}_f}[g_{t+1}(x(1+R))]+L\delta.\end{aligned}$$

But since $g_{t+1}(x)$ is assumed to be Lipschitz continuous, we have

$$\mathrm{E}_{\tilde{P}_f}[g_{t+1}((x+\delta)(1+R))]-\mathrm{E}_{\tilde{P}_f}[g_{t+1}(x(1+R))]\le \mathrm{E}_{\tilde{P}_f}[L\delta(1+R)]=L\delta$$

by the risk-neutral property of $\tilde{P}_f$. This leads to a contradiction. □

4.2 Analysis of local errors through binding constraints

Next we analyze the local error of each induction step. We shall first show that the primal formulation

$$\begin{array}{ll} \text{minimize} & p \\ \text{subject to} & p + r\Delta \geq g_{t+1}(x(1+r)) \text{ for all } r \in \mathcal{U} \end{array} \tag{22}$$

has the following property in terms of "binding constraints":

Lemma 4.2. *For any t and x, there exists an optimal solution for* (22)*, say (p^*, Δ^*), such that either:*

1. *$p^* = g_{t+1}(x)$ or*
2. *There are exactly two binding constraints, corresponding to $r^{(1)}$ and $r^{(2)}$, such that $r^{(1)} > 0$ and $r^{(2)} < 0$. These two constraints uniquely define (p^*, Δ^*).*

Proof. First, there must be at least one binding constraint for (22), because if not, one can always decrease p to achieve lower objective value while preserving all constraints. On the other hand, there must be at most two binding constraints for (22); otherwise, there will be three linearly independent equations $p^* + r_i\Delta^* = g_{t+1}(x(1+r_i))$ for some $r_i, i = 1, 2, 3$ that (p^*, Δ^*) satisfies, which is impossible.

Next, let us move to the case where there are two binding constraints. We need to show that $r^{(1)} < 0 < r^{(2)}$. Suppose $r^{(1)}, r^{(2)} > 0$. We can see that the dual program of (22) is infeasible, which means the primal LP is either infeasible or unbounded. Infeasibility is easily ruled out since one can put $\Delta = 0$ and a large enough p to construct a feasible solution. To show that unboundedness is also impossible, we will show that all p that are smaller than a negative threshold are infeasible, implying the minimization problem (22) is finite. To this end, consider any $\tilde{r}_1 > 0$ and $-\tilde{r}_2 < 0$ that lie in $\mathcal{U}$. Suppose $p + \tilde{r}_1\Delta \geq g_{t+1}(x(1+\tilde{r}_1))$. This implies $\Delta \geq (g_{t+1}(x(1+\tilde{r}_1)) - p)/\tilde{r}_1$. Now, if we choose p to be very negative, then

$$p - \tilde{r}_2\Delta \leq p - \tilde{r}_2 \frac{g_{t+1}(x(1+\tilde{r}_1)) - p}{\tilde{r}_1} = \left(1 + \frac{\tilde{r}_1}{\tilde{r}_2}\right)p - \frac{\tilde{r}_2}{\tilde{r}_1} g_{t+1}(x(1+\tilde{r}_1)) < g_{t+1}(x(1-\tilde{r}_2))$$

and hence (p, Δ) does not satisfy $p - \tilde{r}_2\Delta \geq g_{t+1}(x(1-\tilde{r}_2))$. Hence the set of feasible p must be bounded from below. Similarly, we can show that it is impossible to have $r^{(1)}, r^{(2)} < 0$. It is trivial to see that the two binding constraints uniquely define (p^*, Δ^*).

Finally, suppose there is exactly one binding constraint. When r in the corresponding binding constraint is non-zero, the dual LP is again infeasible, which will result in a contradiction again as above. □

Now we are ready to present the major building block in our local analysis. To encapsulate the effect of local errors due to discretization, we need to introduce an intermediate quantity $g_t^m(x)$, for each $t \in [\tau]$, defined as the optimal solution of the following linear program:

$$\begin{array}{ll} \text{minimize} & p \\ \text{subject to} & g_{t+1}(x(1+r_i)) - r_i\Delta \leq p \;\text{ for } i \in [b] \end{array} \tag{23}$$

The difference between $g_t^m(x)$ and $\hat{g}_t(x)$ is that the calculation of $g_t^m(x)$ assumes accurate access to the function $g_{t+1}(\cdot)$. Thus, we may view $g_t^m(\cdot)$ as a "hybrid variable" that sits in between $g_t(\cdot)$ and $\hat{g}_t(\cdot)$. The following gives a bound for $g_t(\cdot) - g_t^m(\cdot)$:

Lemma 4.3. *With $g_t^m(x)$ defined in* (23)*, we have*

$$0 \leq g_t(x) - g_t^m(x) \leq \sqrt{2Lg_t(x)x\epsilon} + Lx\epsilon \tag{24}$$

where ϵ is the discretization parameter in the algorithm.

The main device we use in the proof is a convenient dual characterization of the optimal solution obtained through the binding constraints in Lemma 4.2. From this characterization, Lipschitz continuity is then used to bound the magnitude of the local errors.

Proof. Recall that we use the following LP to compute $g_t(x)$:

$$\begin{array}{ll} \text{minimize} & p \\ \text{subject to} & g_{t+1}(x(1+r)) - r\Delta \leq p \ \ \text{for } r \in \mathcal{U} \end{array} \tag{25}$$

Let (p^*, Δ^*) be the optimal solution. By Lemma 4.2, we have either 1) $p^* = g_{t+1}(x)$ or 2) there exists a pair $r^{(1)}$ and $r^{(2)}$ such that $r^{(1)} < 0 < r^{(2)}$ and

$$p^* + r^{(i)}\Delta^* = g_{t+1}(x(1 + r^{(i)}) \text{ for } i = 1, 2.$$

Let us focus on the second case, as the first case follows similarly. In this case, we can express $p^* = g_t(x)$ in terms of $r^{(1)}$ and $r^{(2)}$ as follows:

$$g_t(x) = \frac{r^{(2)}}{r^{(2)} - r^{(1)}} g_{t+1}(x(1 + r^{(1)})) - \frac{r^{(1)}}{r^{(2)} - r^{(1)}} g_{t+1}(x(1 + r^{(2)})) \tag{26}$$

Since the discretization length is ϵ, there exist $\hat{r}^{(1)} < 0 < \hat{r}^{(2)} \in \hat{\mathcal{U}}$ such that $0 \leq r^{(1)} - \hat{r}^{(1)} < \epsilon$ and $0 \leq \hat{r}^{(2)} - r^{(2)} < \epsilon$, *i.e.,* we can define $\hat{r}^{(1)}$ as the closest r_i that is at least as large as $r^{(1)}$ and $\hat{r}^{(2)}$ as the closest r_i that is at least as small as $r^{(2)}$. Note that a lower bound of $g_t^m(x)$ is given by the optimal solution of the following linear program:

$$\begin{array}{ll} \text{minimize} & p \\ \text{subject to} & g_{t+1}(x(1 + \hat{r}^{(1)})) - x\hat{r}^{(1)}\Delta \leq p \\ & g_{t+1}(x(1 + \hat{r}^{(2)})) - x\hat{r}^{(2)}\Delta \leq p \end{array} \tag{27}$$

This LP can be solved analytically, which gives us

$$g_t^m(x) \geq \frac{\hat{r}^{(2)}}{\hat{r}^{(2)} - \hat{r}^{(1)}} g_{t+1}(x(1 + \hat{r}^{(1)})) - \frac{\hat{r}^{(1)}}{\hat{r}^{(2)} - \hat{r}^{(1)}} g_{t+1}(x(1 + \hat{r}^{(2)})). \tag{28}$$

Now let $\eta > \epsilon$ be a parameter to be decided later. We consider two cases.
Case 1. $r^{(2)} - r^{(1)} \geq \eta$. First we can see that (detailed calculation appears in Appendix C)

$$\frac{\hat{r}^{(2)}}{\hat{r}^{(2)} - \hat{r}^{(1)}} \geq \frac{r^{(2)}}{r^{(2)} - r^{(1)}} \left(1 - \frac{\epsilon}{\eta}\right). \tag{29}$$

Similarly, we have

$$-\frac{\hat{r}^{(1)}}{\hat{r}^{(2)} - \hat{r}^{(1)}} \geq -\frac{r^{(1)}}{r^{(2)} - r^{(1)}} \left(1 - \frac{\epsilon}{\eta}\right). \tag{30}$$

Hence we may continue (28) and get:

$$\begin{aligned}
g_t^m(x) &\geq \frac{\hat{r}^{(2)}}{\hat{r}^{(2)}-\hat{r}^{(1)}} g_{t+1}(x(1+\hat{r}^{(1)})) - \frac{\hat{r}^{(1)}}{\hat{r}^{(2)}-\hat{r}^{(1)}} g_{t+1}(x(1+\hat{r}^{(2)})) \\
&\geq \frac{r^{(2)}}{r^{(2)}-r^{(1)}}\left(1-\frac{\epsilon}{\eta}\right)\left(g_{t+1}(x(1+r^{(1)})) - xL(r^{(1)}-\hat{r}^{(1)})\right) \\
&\qquad -\frac{r^{(1)}}{r^{(2)}-r^{(1)}}\left(1-\frac{\epsilon}{\eta}\right) g_{t+1}(x(1+r^{(2)})) \\
&\qquad \text{(by Lipschitz continuity of } g_{t+1}\text{, and Eq. (29) and (30))} \\
&= \left(1-\frac{\epsilon}{\eta}\right)\left\{\frac{r^{(2)}}{r^{(2)}-r^{(1)}} g_{t+1}(x(1+r^{(1)})) - \frac{r^{(1)}}{r^{(2)}-r^{(1)}} g_{t+1}(x(1+r^{(2)}))\right\} - xL\epsilon\left(1-\frac{\epsilon}{\eta}\right) \\
&= g_t(x)(1-\frac{\epsilon}{\eta}) - xL\epsilon \ \text{ (by Eq. (26)).}
\end{aligned}$$

Case 2. When $r^{(2)} - r^{(1)} < \eta$. This implies $r^{(2)} < \eta$ and $r^{(1)} > -\eta$. We shall first show that $|g_{t+1}(x) - g_t(x)|$ is small. Specifically,

$$\begin{aligned}
& |g_{t+1}(x) - g_t(x)| \\
=\ & \left| g_{t+1}(x) - \left(\frac{r^{(2)}}{r^{(2)}-r^{(1)}} g_{t+1}(x(1+r^{(1)})) - \frac{r^{(1)}}{r^{(2)}-r^{(1)}} g_{t+1}(x(1+r^{(2)}))\right)\right| \\
\leq\ & \frac{r^{(2)}}{r^{(2)}-r^{(1)}}\left|g_{t+1}(x) - g_{t+1}(x(1+r^{(1)}))\right| - \frac{r^{(1)}}{r^{(2)}-r^{(1)}}\left|g_{t+1}(x) - g_{t+1}(x(1+r^{(2)}))\right| \\
\leq\ & 2Lx\frac{|r^{(1)}\cdot r^{(2)}|}{r^{(2)}-r^{(1)}} \quad \text{(by Lipschitz continuity of } g_{t+1}(\cdot)\text{)} \\
\leq\ & 2Lx|r^{(1)}| \quad \text{(since } \tfrac{r^{(2)}}{r^{(2)}-r^{(1)}} \leq 1\text{)} \\
\leq\ & 2Lx\eta
\end{aligned}$$

We next use the above inequality to compute a lower bound for $g_t^m(x)$:

$$\begin{aligned}
& g_t^m(x) \\
\geq\ & \frac{\hat{r}^{(2)}}{\hat{r}^{(2)}-\hat{r}^{(1)}} g_{t+1}(x(1+\hat{r}^{(1)})) - \frac{\hat{r}^{(1)}}{\hat{r}^{(2)}-\hat{r}^{(1)}} g_{t+1}(x(1+\hat{r}^{(2)})) \quad \text{(by Eq.(28))} \\
\geq\ & \frac{\hat{r}^{(2)}-\hat{r}^{(1)}}{\hat{r}^{(2)}-\hat{r}^{(1)}} g_{t+1}(x) - \frac{2Lx|\hat{r}^{(1)}\hat{r}^{(2)}|}{\hat{r}^{(2)}-\hat{r}^{(1)}} \quad \text{(Lipschitz condition)} \\
\geq\ & g_{t+1}(x) - 2Lx|\hat{r}^{(1)}| \quad \text{(again use the fact that } \tfrac{\hat{r}^{(2)}}{\hat{r}^{(2)}-\hat{r}^{(1)}} \leq 1\text{)} \\
\geq\ & g_t(x) - (2\eta+2\epsilon)Lx. \quad \text{(since } |\hat{r}^{(1)}| \leq |\hat{r}^{(1)} - r^{(1)}| + |r^{(1)}| \leq \epsilon + \eta\text{)}
\end{aligned}$$

Summarizing both cases, we have

$$\begin{aligned}
g_t(x) - g_t^m(x) &\leq \min_{\eta>\epsilon}\max\left\{g_t(x)\frac{\epsilon}{\eta} + xL\epsilon, (2\eta+2\epsilon)Lx\right\} \\
&\leq \min_{\eta>\epsilon}\max\left\{g_t(x)\frac{\epsilon}{\eta}, 2Lx\eta\right\} + 2\epsilon Lx \\
&\leq \sqrt{g_t(x)2\epsilon Lx} + 2L\epsilon x
\end{aligned}$$

by setting $\eta = \sqrt{g_t(x)\epsilon/(2Lx)}$ when $\epsilon \leq \sqrt{g_t(x)/(2Lx)}$, and $\eta = \epsilon$ when $\epsilon > \sqrt{g_t(x)/(2Lx)}$. □

4.3 Error propagation control and telescoping

We now move to the final stage of our analysis. The following is our main result.

Theorem 4.4. *The approximation* $g_0(S_0) - \hat{g}_0(S_0)$ *satisfies*

$$0 \leq g_0(S_0) - \hat{g}_0(S_0) \leq C_1\sqrt{\epsilon}L\tau S_0. \tag{31}$$

for some constant C_1.

The proof of this theorem relies heavily on the "artificial" risk-neutral probability measures that define the optimal dual solutions of the several different value functions, $g_t(x)$, $g_t^m(x)$ and $\hat{g}_t(x)$, for each step of backward induction. Although these probability measures have no real-world correspondence, as discussed in Section 3, they confer "artificial" martingale properties on the underlying asset's price movement. Moreover, we are in fact granted with some freedom in choosing the measure to work under, and we will see that the one associated with $g_t^m(x)$ is the most effective in truncating the propagation of error. We will now lay out the arguments precisely.

Proof. We first apply a standard telescoping trick (see, *e.g.,* [BG04]) to "decouple" the local error from the global error, *i.e.,* let $d_t(x) = g_t(x) - \hat{g}_t(x)$ and we have

$$d_t(x) = (g_t(x) - g_t^m(x)) + (g_t^m(x) - \hat{g}_t(x)).$$

The first term represents the local error at each induction step, whereas the second term comes from error propagation from the future. The first term $g_t(x) - g_t^m(x)$ is handled by Lemma 4.3. For the second term $g_t^m(x) - \hat{g}_t(x)$, we recall the dual characterization in Proposition 3.2 to write

$$g_t^m(x) - \hat{g}_t(x) = \max_{\substack{P\in\mathcal{P}(\hat{\mathcal{U}})\\ \mathrm{E}_P[R_{t+1}]=0}} \mathrm{E}\left[g_{t+1}(x(1+R_{t+1}))\right] - \max_{\substack{P\in\mathcal{P}(\hat{\mathcal{U}})\\ \mathrm{E}_P[R_{t+1}]=0}} \mathrm{E}\left[\hat{g}_{t+1}(x(1+R_{t+1}))\right]$$

Notice that R_{t+1} in both expectations above share the same support $\hat{\mathcal{U}}$, the discretized uncertainty set. Let us write $\mathrm{E}_{x,t}^m$ as the expectation under $P_{x,t}^m := \arg\max\{\mathrm{E}\left[g_{t+1}(x(1+R_{t+1}))\right] \mid P \in \mathcal{P}(\hat{\mathcal{U}}), \mathrm{E}_P[R_{t+1}] = 0\}$. The above equation becomes

$$\begin{aligned}
&\mathrm{E}_{x,t}^m\left[g_{t+1}(x(1+R_{t+1}))\right] - \max_{\substack{P\in\mathcal{P}(\hat{\mathcal{U}})\\ \mathrm{E}_P[R_{t+1}]=0}} \left[\hat{g}_{t+1}(x(1+R_{t+1}))\right]\\
\leq\ &\mathrm{E}_{x,t}^m\left[g_{t+1}(x(1+R_{t+1})) - \hat{g}_{t+1}(x(1+R_{t+1}))\right]\\
&\quad \text{(since } P_{x,t}^m \text{ is feasible for the maximization in the second term above)}\\
=\ &\mathrm{E}_{x,t}^m[d_{t+1}(x(1+R_{t+1}))]
\end{aligned}$$

Hence, $d_t(x)$ can be bounded recursively under the probability measure $P_{x,t}^m$:

$$d_t(x) \leq \sqrt{2Lg_t(x)x\epsilon} + Lx\epsilon + \mathrm{E}_{x,t+1}^m[d_{t+1}(x(1+R_{t+1}))]. \tag{32}$$

Now, let us define a probability measure $P_{S_0}^m$ on the process $\{S_t\}_{t\in[\tau]}$, where $S_t = S_0\prod_{i=1}^t(1+R_i)$ is the price of the underlying asset at the t-th round. This measure $P_{S_0}^m$ is defined by the stepwise transition probability $P^m_{x,t}$ on each R_t, $t \in [\tau]$. From (32), we can expand $d_0(S_0)$ recursively:

$$\begin{aligned}
d_0(S_0) &\leq \sqrt{2Lg_t(S_0)S_0\epsilon} + LS_0\epsilon + \mathrm{E}^m_{S_0,1}[d_1(S_1)] \\
&\leq \sqrt{2Lg_t(S_0)S_0\epsilon} + LS_0\epsilon + \mathrm{E}^m_{S_0,1}[\sqrt{2Lg_1(S_1)S_1\epsilon} + LS_1\epsilon] + \mathrm{E}^m_{S_0,1}\mathrm{E}^m_{S_1,2}[d_2(S_2)] \\
&\qquad \text{(by expanding } d_1(S_1) \text{ by (32) again)} \\
&\vdots \\
&\leq \mathrm{E}^m_{S_0}\left[\sum_{t=1}^{\tau}\left\{\sqrt{2Lg_t(S_t)S_t\epsilon} + LS_t\epsilon\right\}\right]. \qquad (33)
\end{aligned}$$

Next, observe that $\{S_t\}_{t\leq\tau}$ is a martingale under $P^m_{S_0}$ and the filtration $\mathcal{F}_t = \sigma(R_1, \ldots, R_t)$, since $\mathrm{E}^m_{S_0}[R_{t+1}|\mathcal{F}_t] = \mathrm{E}^m_{t,S_t}[R_{t+1}] = 0$. We leverage this fact to bound both terms in (33). First, notice that

$$\mathrm{E}^m_{S_0}[\sum_{t=1}^{\tau} L\epsilon S_t] = L\epsilon\sum_{t=1}^{\tau}\mathrm{E}^m_{S_0}[S_t] = L\tau\epsilon S_0$$

For the other term in Eq. (33), note that by Lemma 4.1 we have $g_t(x) \leq Lx$ for any t and x. Hence we can wrap up the analysis for (33):

$$\begin{aligned}
\mathrm{E}^m_{S_0}\left[\sum_{t=1}^{\tau}\left\{\sqrt{2\epsilon Lg_t(S_t)S_t} + L\epsilon S_t\right\}\right] &\leq \sum_{t=1}^{\tau}\mathrm{E}^m_{S_0}[\sqrt{2\epsilon L^2 S_t^2}] + L\tau\epsilon S_0 \\
&= \sqrt{2\epsilon}L\sum_{t=1}^{\tau 1}\mathrm{E}^m_{S_0}[S_t] + L\tau\epsilon S \\
&= \sqrt{2\epsilon}L\tau S_0 + L\tau\epsilon S_0 \leq C_1\sqrt{\epsilon}L\tau S_0,
\end{aligned}$$

for some constant C_1. This completes the proof of Theorem 4.4. □

Thus we have the following corollary:

Corollary 4.5. *Let δ be an arbitrary constant. Consider using the multinomial tree approximation algorithm to find the price upper bound. When $\epsilon = c\delta^2/(L^2\tau^2)$ for some constant c, the algorithm gives a $\hat{g}_0(S_0)$ such that $g_0(S_0) - \delta S_0 \leq \hat{g}_0(S_0) \leq g_0(S_0)$.*

Tightness of the performance. We shall show in Appendix A.2 that as long as ϵ is a polynomial in τ, an additive error term that is linear in L and S_0 will be unavoidable under the oracle model. This means that the running time of our algorithm necessarily depends on L and the additive dependency on S is essentially tight.

Non-uniform uncertainty set. When the uncertainty sets are non-uniform, we can still use the multinomial tree algorithm to find the approximate solution so long as the largest uncertainty set is still polynomial in τ. We remark, though, that the parameter ϵ has to remain unchanged even if the uncertainty sets change over the time, *i.e.,* our ϵ is still set to $\epsilon = c\delta^2/(L^2\tau^2)$.

5 Convergence to controlled diffusion process

In this section, we show that our minimax upper bound of the option price, with possibly non-convex payoff, converges to the price based on a controlled diffusion process when the uncertainty sets are appropriately scaled. To facilitate discussion, for this section we let the time steps be $0, \delta, 2\delta, \ldots, T$ (for simplicity let T be a multiple of δ), and thus $\tau = \frac{T}{\delta}$. Let $\mathcal{U}^\delta = [-\underline{\zeta}^\delta, \overline{\zeta}^\delta]$, where $-\underline{\zeta}^\delta \triangleq -\underline{\zeta}\sqrt{\delta}$ and $\overline{\zeta}^\delta \triangleq \overline{\zeta}\sqrt{\delta}$. Let us recall that we write $S(t)$ as the continuous-time price of the underlying asset at time t, with $S(0) = S_0$ (this process is decided by nature). We have the following theorem:

Theorem 5.1. *Let g be a Lipschitz continuous payoff function. As $\delta \to 0$, the upper bound of the option price defined by*

$$\min_{\Delta_1, \Delta_2, \ldots, \Delta_\tau} \max_{R_1^\delta, \ldots, R_\tau^\delta \in \mathcal{U}^\delta} g\left(S_0 \prod_{i=1}^{\tau}(1 + R_i^\delta)\right) - \sum_{i=1}^{\tau} \Delta_i R_i^\delta \tag{34}$$

converges to

$$G(0, S_0) := \max_{\xi} \mathrm{E}[g(S(T))]. \tag{35}$$

Here $S(t)$ follows the dynamic

$$dS(t) = \sigma(u)S(t)dw(t) \tag{36}$$

where $w(t)$ is a standard Brownian motion, $\sigma(u) = u$ is the controlled volatility, and $\xi = \{u(t), 0 \le t \le T : u(t) \in U\}$ is an adapted control sequence. The domain of control $U = [0, \underline{\zeta}\overline{\zeta}]$.

Intuitively, this theorem asserts that the continuous-time limit of the option price's upper bound for non-convex payoff is Gaussian in nature, similar to the case of convex payoff and Black-Scholes model. The crucial difference with those cases, however, is that nature now has the additional power to choose (to reduce) the volatility at any point of time, and this is only helpful to nature in the case of non-convex payoff. Since we know from Section 4 that in the discrete-time setting with non-convex payoff, the nature does not necessarily choose the extremal point in the uncertainty set, it is not surprising that there is reduction in volatility at certain points of time in the continuous-time counterpart. In Appendix B we will also demonstrate how Theorem 5.1 can be easily reduced to Black-Scholes model in the case of convex payoff.

The way to show Theorem 5.1 is to view the diffusion process (36) as the starting object, and argue that the discrete-time game is an *approximating Markov chain* of (36) [KD01]. The following is the key to prove Theorem 5.1:

Theorem 5.2. *Consider $x(t) = x_0 + \int_0^t \sigma(w(s), u(s))dw(s)$ where $w(t)$ is a standard Brownian motion, and $u(t)$ is an adapted control sequence on a compact set U. Let $V(t, x) = \max \mathrm{E}_{(t,x)}[h(x(T))]$, where h is continuous and $\mathrm{E}_{(t,x)}$ denotes the expectation conditional on $x(t) = x$.*

Consider a Markov chain approximation as follows. Divide the time into steps of size δ. Define a Markov chain $\{X^\delta(t)\}_{t=0,\delta,2\delta,\ldots,T}$ with transition $p^\delta(x, y|\alpha)$. Let $V^\delta(T, x) = h(x)$. For each step backward, solve

$$V^\delta(t, x) = \max_{\alpha \in U} \mathrm{E}^{\delta,\alpha}_{(t, X^\delta(t)=x)}[V^\delta(t + \delta, X^\delta(t + \delta))]$$

where $\mathrm{E}^{\delta,\alpha}_{(t,X^\delta(t)=x)}$ is the expectation taken using the transition $p^\delta(x, y|\alpha)$ for $X^\delta(t+\delta)$. Then we interpolate $X^\delta(\cdot)$ and $V^\delta(\cdot, x)$ such that they are piecewise constant on $t \in [0, T]$, i.e., *$X^\delta(s) = X^\delta(i\delta)$ and $V^\delta(s, x) = V^\delta(i\delta, x)$ for $i\delta \le s < (i + 1)\delta$.*

Suppose that

1. *$\sigma(x, \alpha)$ is Lipschitz continuous in x, uniformly in U.*

2. *The transition is locally consistent,* i.e., $\mathrm{E}^{\delta,\alpha}_{(t,X^\delta(t)=x)}[X^\delta(t+\delta)-X^\delta(t)]=o(\delta)$, $\mathrm{E}^{\delta,\alpha}_{(t,X^\delta(t)=x)}(X^\delta(t+\delta)-X^\delta(t))^2=\sigma(x,\alpha)^2\delta+o(\delta)$ *for any* $\alpha\in U$, *and* $\sup_{0\le t\le T}|X^\delta(t+\delta)-X^\delta(t)|\overset{a.s.}{\to}0$ *as* $\delta\to 0$.
3. *h is uniformly integrable,* i.e., $\lim_{\eta\to\infty}\sup_\delta \mathrm{E}^{\delta,\alpha}_{(t,x)}[|h(X^\delta(T))|;|h(X^\delta(T))|>\eta]<\infty$ *for any* t,x.

Then $V^\delta(t,x)\to V(t,x)$ *pointwise on* $t\in[0,T],x\in\mathbb{R}^+$.

Outline of proof. The proof is adapted from that of [KD01], p. 356, Theorem 1.4. As such we only provide an outline here. The proof calls on machinery in weak convergence analysis. Since there is no regularity condition on the control $u(t)$ in (36), the first step is to enlarge the space of the process to measure-valued space. Namely, for any given control sequence $\{u(t)\}_{t\in[0,T]}$, we can write

$$x(t)=S_0+\int_0^t\int_U\sigma(\alpha)x(v)M(d\alpha dv)\tag{37}$$

where $M(d\alpha dv)$ is a real measure-valued continuous random process defined with, roughly speaking, the following properties: For any fixed measurable subset A in U, $M(A,t)$ is a martingale that has quadratic variation $m(A,t)$ (so-called martingale measure; see [KD01], p. 352, (1.9)). The quantity $m(\cdot,\cdot)$ is the so-called relaxed control, and is a measure that puts a delta mass if the control $u(t)\in A$ at time t ([KD01], p. 263, (5.1)). Under Lipschiz continuity of $\sigma(\cdot,\cdot)$ and compact control set U, there is a unique weak sense solution to (37) ([KD01], p. 353, discussion under A1.1).

Now, one can similarly define $x^\delta(t)=S_0+\int_0^t\int_U\sigma(\alpha)x^\delta(v)M^\delta(d\alpha dv)$ under the relaxed control m^δ that is discretized and interpolated through the time steps $0,\delta,2\delta,\dots,T$, and $M^\delta(A,t)$ has quadratic variation $m^\delta(A,t)$. By p. 356, Theorem 1.3 in [KD01], under the conditions of Lipshitz continuity of $\sigma(\cdot,\cdot)$ and local consistency (*i.e.,* conditions 1) and 2) in our theorem), there exists a subsequence (M^δ,m^δ) that converges weakly to (M,m). By uniform integrability (*i.e.,* condition 3)), we then have $\liminf_\delta V^\delta(t,x)\ge V(t,x)$ for any t,x [AK94]. The final step is to argue that the inequality is in fact matched by some control. This consists of choosing a so-called ϵ-optimal solution of the relaxed control by using a finite-dimensional Wiener process and a finite-valued and piecewise constant control, and arguing that this solution provides an ϵ-approximation to the optimal value ([KD01], p. 355, Theorem 13.1.2).

□

Proof of Theorem 5.1. We will show that our hedging game is an approximating Markov chain to the control problem given in (35), and has the three listed properties in Theorem 5.2. First, the $\sigma(x,\alpha)$ in Theorem 5.2 equals $x\alpha$ in (36), where α is the control lying in the compact set $U=[0,\underline{\zeta}\overline{\zeta}]$.

Second, we check the local consistency property. Define $\{G^\delta(t,x)\}_{t=0,\delta,2\delta,\dots,T}$ as the value function in each step of the backward induction in the discrete-time hedging game defined in (34). Note that by Proposition 3.2, $G^\delta(i\delta,S^\delta(i\delta))=\max_{P\in\mathcal{P}([-\underline{\zeta}^\delta,\overline{\zeta}^\delta])}\mathrm{E}[G^\delta((i+1)\delta,S^\delta(i\delta)(1+R^\delta_i))]$, where we define $S^\delta(i\delta)=S_0\prod_{j=1}^i(1+R^\delta_j)$ for $i\in[\tau]$, and piecewise constant interpolation of $S^\delta(t)$ for other t. In fact, we can rewrite this as

$$\begin{array}{ll}\max_{p_1,p_2,r_1,r_2} & p_1G(t+\delta,S^\delta(t)(1+r_1))+p_2G(t+\delta,S^\delta(t)(1+r_2))\\ \text{subject to} & p_1r_1+p_2r_2=0\\ & p_1+p_2=0\\ & p_1,p_2\ge 0,-\underline{\zeta}^\delta\le r_1\le r_2\le\overline{\zeta}^\delta\end{array}\tag{38}$$

where we encode the maximal probability distribution in (38) by $\{p_1,p_2,r_1,r_2\}$. Let $\alpha^2\delta=p_1r_1^2+p_2r_2^2$. Then the range of α is $[0,\underline{\zeta}^\delta\overline{\zeta}^\delta/\delta]=[0,\underline{\zeta}\overline{\zeta}]$. We then have $\mathrm{E}^{\delta,\alpha}_{(t,S^\delta(t)=x)}[S^\delta(t+\delta)-S^\delta(t)]=0$ and

$\mathrm{E}^{\delta,\alpha}_{(t,S^\delta(t)=x)}(S^\delta(t+\delta)-S^\delta(t))^2 = \mathrm{E}^{\delta,\alpha}_{(t,S^\delta(t)=x)}[(S^\delta(t))^2(R^\delta)^2] = (S^\delta(t))^2\alpha^2\delta$ (here R^δ denotes i.i.d copy of R^δ_i in (34)). Moreover, $|S^\delta(t+\delta)-S^\delta(t)| \le \max\{\underline{\zeta}^\delta, \overline{\zeta}^\delta\} = \max\{\underline{\zeta}, \overline{\zeta}\}\cdot\delta$ and so $\sup_{0\le t\le T}|S^\delta(t+\delta)-S^\delta(t)| \overset{a.s.}{\to} 0$. (One can see the policy defined in terms of α above is not unique; nevertheless it does not affect local consistency.)

Finally, we claim that g is uniformly integrable. We use an argument similar to the proof of Corollary 3.6, and conclude that $\sup_\delta \mathrm{E}^{\delta,\alpha}_{(t,x)}[g(S^\delta(T))^2] \le C_1 S_0 \exp\{T\underline{\zeta}\overline{\zeta}\} + C_2 < \infty$ for some constants $C_1, C_2 > 0$, by Lipschitz continuity of g and that $\alpha \in [0, \underline{\zeta}\overline{\zeta}]$. This will imply that g is uniformly integrable. □

6 Integration with jumps

A natural extension to the Black-Scholes framework is to allow "shocks" in the movement of the stock price, typically modeled by the stochastic community as Poisson arrivals of jumps on top of a continuous geometric Brownian motion on the asset price. In this section we will adopt our adversary framework to incorporate price jumps.

There can be various ways to model when and how much the adversary can control the price to jump. Below we analyze two natural examples that can be extended from our framework in the previous sections. In the first example, we assume the adversary has no control over the occurrence and magnitude of jumps. We will show that the price upper bound is exactly the same as the price in the standard jump diffusion model, when the payoff function is convex. In the second example, the adversary can control when the jump happens, subject to a constraint on the total number of jumps. The magnitude of the jumps can also be assumed to be controllable by the adversary, and we will see that only small modifications to the algorithm presented in Section 4 are needed.

Example 1. The random jump model. This model assumes a convex payoff function and that the nature performs a jump with a prefixed small probability q at each step. If a jump occurs, the nature moves according to a return Y that is random (whose distribution is given). Otherwise, with $1-q$ probability, the nature will have freedom to choose its path inside the uncertainty set $\mathcal{U}$. We also assume q is sufficiently small (this will be specified precisely in the sequel).

Consider first a single-round game, with initial price of the underlying asset S_0. The upper bound formulation is

$$\min_\Delta \max_R (1-q)[(g(S_0(1+R))) - \Delta R] + q\mathrm{E}_Y[g(S_0(1+Y)) - \Delta Y], \tag{39}$$

where R is the return that nature can choose, if a jump does not occur. The expectation E_Y is with respect to the jump magnitude variable Y.

The formulation (39) can be written as

$$\begin{array}{ll} \text{minimize} & p \\ \text{subject to} & (1-q)[(g(x(1+r))) - \Delta r] + q\mathrm{E}_Y[g(x(1+Y)) - \Delta Y] \le p \ \text{ for all } r \in \mathcal{U} \end{array} \tag{40}$$

or

$$\begin{array}{ll} \text{minimize} & p \\ \text{subject to} & p + \Delta((1-q)r + q\mathrm{E}_Y[Y]) \ge (1-q)g(x(1+r)) + q\mathrm{E}_Y[g(x(1+Y))] \ \text{ for all } r \in \mathcal{U} \end{array} \tag{41}$$

The dual of (41) is

$$\begin{array}{ll} \max & (1-q)\mathrm{E}[g(x(1+r))] + q\mathrm{E}_Y[g(x(1+Y))] \\ \text{subject to} & (1-q)\mathrm{E}[r] + q\mathrm{E}_Y[Y] = 0 \\ & P \in \mathcal{P}(\mathcal{U}) \end{array} \tag{42}$$

where $\mathcal{P}(\mathcal{U})$ denotes the set of all probability measures supported on $\mathcal{U}$.

We can use the same convexity argument as in Section 3.2 to argue that the optimal measure must be concentrated at the two extreme points of $\mathcal{U}$, namely $-\underline{\zeta}$ and $\overline{\zeta}$, when q is small enough *i.e.,* $(1-q)(-\overline{\zeta}) + q\mathrm{E}_Y[Y] > 0$ and $(1-q)\underline{\zeta} + q\mathrm{E}_Y[Y] < 0$. Let the weights of these two extreme points be w_1 and w_2. The constraints $(1-q)(w_1(-\underline{\zeta}) + w_2\overline{\zeta}) + q\mathrm{E}_Y[Y] = 0$ and $w_1 + w_2 = 1$ completely determine

$$w_1 = \frac{\overline{\zeta} + \frac{q\mathrm{E}_Y[Y]}{1-q}}{\overline{\zeta} + \underline{\zeta}},\ w_2 = \frac{\underline{\zeta} + \frac{-q\mathrm{E}_Y[Y]}{1-q}}{\overline{\zeta} + \underline{\zeta}} \tag{43}$$

Therefore, the upper bound is

$$(1-q)\left[\frac{\overline{\zeta} + \frac{q\mathrm{E}_Y[Y]}{1-q}}{\overline{\zeta} + \underline{\zeta}} g(x(1+\underline{\zeta})) + \frac{\frac{\underline{\zeta} - q\mathrm{E}_Y[Y]}{1-q}}{\overline{\zeta} - \underline{\zeta}} g(x(1+\overline{\zeta}))\right] + q\mathrm{E}_Y[g(x(1+Y))]. \tag{44}$$

For a τ-round game, the formulation follows analogously as in Section 3.4, with the optimal value being the value function of a dynamic program, with $g_\tau(x) = g(x)$ and $g_t(x)$ equal to (44) but with g replaced by g_{t+1}. This characterization is exactly the same as a discrete jump diffusion process, studied in [Ami93], which also demonstrated the continuous-time limit under appropriate scaling of the jump probability and the binomial return rates. This leads immediately to the following:

Proposition 6.1. *Consider the random jump model with jump probability $q^\tau = q/\tau$ for a constant q, jump magnitude random variable Y (not scaled with τ), and uncertainty set $U^\tau = [-\underline{\zeta}/\sqrt{\tau}, \overline{\zeta}/\sqrt{\tau}]$ at each step $t \in [\tau]$. Assume a convex payoff $g(\cdot)$. The upper bound of the option price converges to* $\mathrm{E}[g(S(T))]$, *where $S(t)$ follows a jump diffusion process given by*

$$S(t) = S_0 \exp\{\nu w(t) - \nu/2 + \sum_{j=1}^{J(t)} \log Y(j)\}.$$

Here $\nu = \underline{\zeta}\overline{\zeta}$, $J(t)$ is a Poisson process with rate p, and $Y(j)$ are i.i.d. copies of Y.

Example 2. The adversarial jump model. In this model, we allow the adversary to make no more than ℓ jumps throughout the τ rounds of game. When the adversary decides to make a jump, it can choose a return from the uncertainty set $\mathcal{W}$; Otherwise, it can only choose from the ordinary uncertainty set $\mathcal{U}$. We assume the sizes of $\mathcal{U}$ and $\mathcal{W}$ are both polynomial in τ. Typically $\mathcal{W} \supset \mathcal{U}$, but it is not required in the analysis.[3]

We now explain how our analysis in Section 3 and our algorithm in Section 4 can be extended to this scenario. Define $g_t(x, \ell)$ as the price upper bound at the t-th round when the underlying asset's price is x and the nature still has ℓ number of jump quota. At the t-th round, if $\ell \geq 1$, the adversary may choose to use a jump, in which case the relevant price upper bound at the next round will be $g_{t+1}(x(1+R), \ell-1)$,

[3] For simplicity, we assume the uncertainty sets when jumps are present and absent are both uniform. This condition can easily be relaxed.

where $R \in \mathcal{W}$; suppose the adversary chooses not to jump, then the relevant price upper bound becomes $g_{t+1}(x(1+R), \ell)$, where $R \in \mathcal{U}$. The LP formulation for each round is thus

$$\begin{array}{ll} \text{minimize} & p \\ \text{subject to} & g_{t+1}(x(1+r), \ell) - r\Delta \leq p \;\text{ for } r \in \mathcal{U} \\ & g_{t+1}(x(1+r), \ell-1) - r\Delta \leq p \;\text{ for } r \in \mathcal{W}, \text{if } \ell \geq 1. \end{array} \tag{45}$$

By an argument similar to Proposition 3.2, we can introduce a risk-neutral measure that characterizes the optimal dual solution for (45). The additional feature here is that the risk-neutral measure comprises of a mixture between the ordinary uncertainty set and the enlarged uncertainty set, depending on whether a jump is initiated. The dual formulation can be written as

$$\max_{\substack{q\in[0,1], P_f\in\mathcal{P}(\mathcal{U}), P_J\in\mathcal{P}(\mathcal{W}): \\ q\mathrm{E}_f[R_{t+1}]+(1-q)\mathrm{E}_J[R_{t+1}]=0}} q\mathrm{E}_{P_f}[g_{t+1}(x(1+R_{t+1}), \ell)] + (1-q)\mathrm{E}_{P_J}[g_{t+1}(x(1+R_{t+1}), \ell-1)]$$

Here q is the mixture probability of the occurrence of jump, and P_f and P_J are the conditional distributions supported on $\mathcal{U}$ and $\mathcal{W}$ respectively, depending on whether a jump occurs. When $\ell = 0$, then the dual formulation reduces to the case in Section 4, and we merely have

$$\max_{P_f\in\mathcal{P}(\mathcal{U})\mathrm{E}_f[R_{t+1}]} \mathrm{E}_{P_f}[g_{t+1}(x(1+R_{t+1}), 0)]$$

The dynamic program has the terminal value $g_\tau(x, \ell) = g(x)$ for all ℓ and $g_t(x, 0) = g_t(x)$ for all t, where $g_t(x)$ is defined in Section 4.

From this characterization, we can use a multinomial approximation scheme similar to that in Section 4 to compute the price upper bound for general payoff functions. In this algorithm we discretize the uncertainty sets $\mathcal{U}$ and $\mathcal{W}$ into $\hat{\mathcal{U}}$ and $\hat{\mathcal{W}}$ so that the step length in each discrete set is ϵ. The new feature is that in each backward induction step, we compute $\hat{g}_t(x, m)$ for $m \in \{0, 1, \ldots, \ell\}$, where x can take on polynomial number of distinct values (we can discretize $\mathcal{U}$ and $\mathcal{W}$ in a coherent way to achieve this). Here, we pay a factor of ℓ in the running time and the space comlexity compared to the algorithm in Section 4 because of the exhaustive enumeration regarding ℓ. We have the following analogous performance bound:

Corollary 6.2. *Let δ be an arbitrarily small constant. By setting $\epsilon = cL^2\tau^2/\delta^2$ for some constant c, our algorithm gives $\hat{g}_0(S_0)$ such that $g_0(S_0) - \delta S_0 \leq \hat{g}_0(S_0) \leq g_0(S_0)$.*

Proof. The proof can be adapted easily from that of Theorem 4.4, hence we shall highlight the main steps here. First, we extend, in a straightforward manner, the argument in Lemma 4.1 to prove that $g_t(x, \ell)$ is Lipschitz continuous for any $t \in [\tau]$ and ℓ. Next, the same binding constraint argument as in Lemma 4.3 will reveal that $0 \leq g_t(x, \ell) - g_t^m(x, \ell) \leq \sqrt{2L \max\{g_t(x, \ell), g_t(x, \ell-1)\} S\epsilon} + Lx\epsilon$ for any t and ℓ, where for convenience we use the convention that $g_t(x, -1) = g_t(S, 0)$. Then, since $g_t(x, \ell) \leq Lx$ for any ℓ by an argument similar to Lemma 4.1, we can use the "artificial" probabilistic machinery to arrive at the conclusion. ☐

7 American options

This section generalizes our results to American-type option. We will give a dual characterization using risk-neutral measure, and from there we will obtain convergence and algorithmic results similar to European-type options. As before, let us first consider the single-round game.

7.1 Single-round American option game

The upper bound of an American option can be expressed as:

$$\min_{\Delta} \max_{\substack{\theta\in\{0,1\}\\ R\in\mathcal{U}}} ((g(S_0(1+R)) - R\Delta)(1-\theta) + g(S_0)\theta)\,, \tag{46}$$

where θ is the decision made by the adversary in exercising the option prematurely: $\theta = 1$ if early exercise is prompted, otherwise $\theta = 0$. We remark that it is the adversary, not the trader, to have the right to exercise early in upper bound evaluation. It is because the upper bound comes from a short-option argument (see Section 2) that endows nature as the holder of the option and hence the early exercise right. For lower bound evaluation, the bound is

$$\max_{\Delta,\theta\in\{0,1\}} \min_{R\in\mathcal{U}} ((g(S_0(1+R)) - R\Delta)(1-\theta) + g(S_0)\theta)\,. \tag{47}$$

Let us focus on the upper bound (46), which can be rewritten as:

$$\min_{\Delta} \max_{\theta\in\{0,1\}} \max_{R\in\mathcal{U}} ((g(S_0(1+R)) - \Delta R)(1-\theta) + g(S_0)\theta)\,. \tag{48}$$

We shall argue that the $\min$ and the first $\max$ operators can be interchanged, namely

$$\begin{aligned}&\min_{\Delta} \max_{\theta\in\{0,1\}} \max_{R\in\mathcal{U}} ((g(S_0(1+R)) - \Delta R)(1-\theta) + g(S_0)\theta)\\ =\;& \max_{\theta\in\{0,1\}} \min_{\Delta} \max_{R\in\mathcal{U}} ((g(S_0(1+R)) - \Delta R)(1-\theta) + g(S_0)\theta)\,.\end{aligned}$$

To simplify notation, let us write $\Phi(\Delta,\theta) \triangleq \max_{R\in\mathcal{U}}((g(S_0(1+R)) - \Delta R)(1-\theta) + g(S_0)\theta)$. Thus, we need to show $\min_\Delta \max_\theta \Phi(\Delta,\theta) = \max_\theta \min_\Delta \Phi(\Delta,\theta)$. The direction $\min_\Delta \max_\theta \Phi(\Delta,\theta) \geq \max_\theta \min_\Delta \Phi(\Delta,\theta)$ is straightforward. To show that $\min_\Delta \max_\theta \Phi(\Delta,\theta) \leq \max_\theta \min_\Delta \Phi(\Delta,\theta)$, the main observation is that Δ is only influential if $\theta = 0$. More precisely, consider $\max_\theta \min_\Delta \Phi(\Delta,\theta)$. When $\theta = 1$, $\min_\Delta \Phi(\Delta,\theta) = g(S_0)$, independent of the choice of Δ; when $\theta = 0$, we have $\min_\Delta \Phi(\Delta,\theta) = \max_{\substack{P\in\mathcal{P}(\mathcal{U})\\ \mathrm{E}[R]=0}} \mathrm{E}[g(S_0(1+R))]$ by Proposition 3.2, with $\Delta^* = \arg\max_\Delta g(S_0(1+R)) - \Delta R$. Hence

$$\max_{\theta} \min_{\Delta} \Phi = \max \left\{ \max_{\substack{P\in\mathcal{P}(\mathcal{U})\\ \mathrm{E}[R]=0}} \mathrm{E}[g(S_0(1+R))], g(S_0) \right\} \tag{49}$$

Consider using Δ^* in $\max_\theta \Phi(\Delta,\theta)$. We then have

$$\min_{\Delta} \max_{\theta} \Phi(\Delta,\theta) \leq \max_{\theta} \Phi(\Delta^*,\theta) = \max \left\{ \max_{\substack{P\in\mathcal{P}(\mathcal{U})\\ \mathrm{E}[R]=0}} \mathrm{E}[g(S_0(1+R))], g(S_0) \right\}. \tag{50}$$

Thus, we have $\min_\Delta \max_\theta \Phi(\Delta,\theta) = \max_\theta \min_\Delta \Phi(\Delta,\theta)$, which renders the single-round price upper bound (46) as

$$\min_{\Delta} \max_{\substack{\theta\in\{0,1\}\\ R\in\mathcal{U}}} ((g(S_0(1+R)) - R\Delta)(1-\theta) + g(S_0)\theta) = \max \left\{ \max_{\substack{P\in\mathcal{P}(\mathcal{U})\\ \mathrm{E}[R]=0}} \mathrm{E}[g(S_0(1+R))], g(S_0) \right\} \tag{51}$$

Analogously, the lower bound is

$$\max_{\Delta,\theta\in\{0,1\}} \min_{R\in\mathcal{U}} ((g(S_0(1+R)) - R\Delta)(1-\theta) + g(S_0)\theta) = \max\left\{ \min_{\substack{P\in\mathcal{P}(\mathcal{U})\\ \mathrm{E}[R]=0}} \mathrm{E}[g(S_0(1+R))], g(S_0) \right\}$$

7.2 Multi-round game

With Eq. (51), we can get characterization and convergence results for American options for convex payoff for multi-round game:

Lemma 7.1. *Consider a τ-round American option hedging game with convex payoff function $g(\cdot)$ and uniform uncertainty set $\mathcal{U} = [-\underline{\zeta}, \overline{\zeta}]$. The option price's upper bound is the same as the τ-round binomial tree model. Moreover, if U has the same scaling as Corollary 3.6 as $\tau \to \infty$, the upper bound converges to the Black-Scholes price.*

Proof. The proof follows similarly as that for European-type options. We will argue that $g_t(x)$ is convex for any t and

$$g_t(x) = \max\left\{ \max_{\substack{P\in\mathcal{P}(\{-\underline{\zeta},\overline{\zeta}\})\\ \mathrm{E}[R]=0}} \mathrm{E}[g_{t+1}(x(1+R))], g(x) \right\}. \tag{52}$$

by induction. First, as in Corollary 3.3, when $g_{t+1}(\cdot)$ is convex, the optimal risk-neutral measure for the dual problem at the t-th round, assuming the option is not exercised, has probability masses only at $\{-\underline{\zeta}, \overline{\zeta}\}$. Second, when $g_{t+1}(\cdot)$ is convex, $\max_{\substack{P\in\mathcal{P}(\{-\underline{\zeta},\overline{\zeta}\})\\ \mathrm{E}[R]=0}} \mathrm{E}[g_{t+1}(S(1+R))]$ is also convex by Lemma 3.5. Since the maximum of convex functions is still convex, we conclude the induction. □

The algorithm for computing American options over binomial trees can be found in, for example, [Hul09]. It is known that under the binomial tree model, it is always optimal to exercise at the terminal time for American call option, but not the case for put option. These results immediately carry over to our model.

For non-convex payoff in upper bound calculation, or any payoff in lower bound calculation, the reduction to binomial tree does not necessarily hold. One can resort to our multinomial tree algorithm presented in Section 4 to approximate, for instance the upper bound, using the following recursive formula:

$$\hat{g}_t(x) = \max\left\{ \max_{\substack{P\in\mathcal{P}(\hat{\mathcal{U}})\\ \mathrm{E}[R]=0}} \mathrm{E}[\hat{g}_{t+1}(x(1+R))], \hat{g}(x) \right\}, \tag{53}$$

where $\hat{\mathcal{U}}$ is the discretized uncertainty set with step length ϵ. The performance can be analyzed by using the same techniques presented in Section 4, which gives the following corollary:

Corollary 7.2. *Let δ be an arbitrary constant. Consider using the multinomial tree approximation algorithm to find the American option's upper bound. When $\epsilon = c\delta^2/(L^2\tau^2)$ for some constant c, the algorithm gives an $\hat{g}_0(S_0)$ such that $g_0(S_0) - \delta S_0 \le \hat{g}_0(S_0) \le g_0(S_0)$.*

7.3 Continuous-time exercise right

One arguably unsatisfying feature in our model for the upper bound of American options is that the nature can only exercise the option at discrete rounds, *i.e.,* at the times when the trader can execute trade decisions. There is a natural formulation to relax this constraint, and it turns out that this relaxation does not change our existing model.

Definition 7.3 (American option hedging game with continuous-time exercise right.)**.** *We model the dynamics for an American option, with payoff $g(\cdot)$ and expiration T, by considering a τ-round game between the trader and the nature. The time length for each round is $\gamma \triangleq T/\tau$. Let $\{-\underline{\zeta}_i, \overline{\zeta}_i\}$ be the uncertainty parameters for the i-th round. In this game,*

- *the investor is only allowed to trade at the beginning of each round,* i.e., *at time $0, \gamma, 2\gamma, ..., T$.*
- *the adversary is allowed to exercise the option at any time. The adversary also decides the (continuous) trajectory of the price movement subject to the following constraints specified by the uncertainty parameters: let $t = i\gamma + \delta$, where $\delta < \gamma$; we require $-\underline{\zeta}_{i+1} \cdot \delta/\gamma \leq S_t/S_{i\gamma} - 1 \leq \overline{\zeta}_{i+1} \cdot \delta/\gamma$ for any $\delta < \gamma$.*

Intuitively, the price movement lies in the uncertainty set in the form of a "cone" that extends from time $i\gamma$ to $(i+1)\gamma$. By using a simple change of variable trick, we have the following observation:

Corollary 7.4. *In the model in Definition 7.3, the upper bound of an American option's price is the same as the upper bound from the ordinary hedging game introduced in Section 7.2.*

References

[AB09] S. Arora and B. Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009.

[AFW12] Jacob Abernethy, Rafael M. Frongillo, and Andre Wibisono. Minimax option pricing meets black-scholes in the limit. In *STOC*, pages 1029–1040, 2012.

[AK94] Kaushik Amin and Ajay Khanna. Convergence of american option values from discrete-to continuous-time financial models1. *Mathematical Finance*, 4(4):289–304, 1994.

[Ami93] K. I. Amin. Jump diffusion option valuation in discrete time. *Journal of Finance*, pages 1833–1863, 1993.

[BB12] Chaithanya Bandi and Dimitris Bertsimas. Tractable stochastic analysis in high dimensions via robust optimization. *Mathematical programming*, 134(1):23–70, 2012.

[Ber05a] P. Bernhard. A robust control approach to option pricing including transaction costs. *Annals of the ISDG*, 7:391–416, 2005.

[Ber05b] Pierre Bernhard. The robust control approach to option pricing and interval models: an overview. In Michèle Breton and Hatem Ben-Ameur, editors, *Numerical Methods in Finance*, pages 91–108. Springer, New York, 2005.

[BG04] Mark Broadie and Paul Glasserman. A stochastic mesh method for pricing high-dimensional american options. *Journal of Computational Finance*, 7:35–72, 2004.

[Bil09] Patrick Billingsley. *Convergence of Probability Measures*. Wiley Series in Probability and Statistics, 2009.

[BS73] Fischer Black and Myron Scholes. The Pricing of Options and Corporate Liabilities. *The Journal of Political Economy*, 81(3):637–654, 1973.

[Che10] Si Chen. Robust option pricing : An ϵ-arbitrage approach. Master's thesis, MIT, 2010.

[CHW11] Qi Cheng, Joshua E. Hill, and Daqing Wan. Counting value sets: Algorithm and complexity. *CoRR*, abs/1111.1224, 2011.

[CRR79] John C. Cox, Stephen A. Ross, and Mark Rubinstein. Option pricing: A simplified approach. *Journal of Financial Economics*, 7(3):229–263, September 1979.

[CT12] R. Cont and P. Tankov. *Financial Modelling with Jump Processes, Second Edition*. Chapman & Hall/Crc Financial Mathematics Series. CRC PressINC, 2012.

[DKM06] Peter DeMarzo, Ilan Kremer, and Yishay Mansour. Online trading algorithms and robust option pricing. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, STOC '06, pages 477–486. ACM, 2006.

[Duf05] D. Duffie. *Dynamic Asset Pricing Theory*. New Age International (P) Limited, 2005.

[Dur10] R. Durrett. *Probability: Theory and Examples*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2010.

[FIA12] FIA. http://www.futuresindustry.org/volume-.asp. *Trading Volume Statistics*, 2012.

[FS06] W. H. Fleming and H. M. Soner. *Controlled Markov Processes and Viscosity Solutions*. Springer, 2006.

[GM11] Eyal Gofer and Yishay Mansour. Regret minimization algorithms for pricing lookback options. In *Proceedings of the 22nd international conference on Algorithmic learning theory*, ALT'11, pages 234–248. Springer-Verlag, 2011.

[He90] Hua He. Convergence from discrete-to continuous-time contingent claims prices. *Review of Financial Studies*, 3(4):523–546, 1990.

[Hes93] Steven L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6:327–343, 1993.

[HF11] Alexander Schied Hans Follmer. *Stochastic Finance: An Introduction in Discrete Time*. Walter de Gruyter, 2011.

[Hul09] J. Hull. *Options, Futures, & Other Derivatives*. Prentice Hall finance series. Prentice Hall, 2009.

[KD01] H.J. Kushner and P. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Applications of Mathematics Series. Springer, 2001.

[Kol11] Vassili Kolokoltsov. Game theoretic analysis of incomplete markets: emergence of probabilities, nonlinear and fractional black-scholes equations. *CoRR*, abs/1105.3053, 2011.

[Kou02] Steven G Kou. A jump-diffusion model for option pricing. *Management science*, 48(8):1086–1101, 2002.

[Lue97] D.G. Luenberger. *Optimization by Vector Space Methods*. Series in Decision and Control. Wiley, 1997.

[Mer73] Robert C. Merton. Theory of Rational Option Pricing. *Bell Journal of Economics and Management Science*, 4(1):141–183, 1973.

[Mer76] Robert C Merton. Option pricing when underlying stock returns are discontinuous. *Journal of financial economics*, 3(1):125–144, 1976.

[MMS89] Dilip B Madan, Frank Milne, and Hersh Shefrin. The multinomial option pricing model and its brownian and poisson limits. *Review of Financial Studies*, 2(2):251–265, 1989.

[REK00] Richard Rouge and Nicole El Karoui. Pricing via utility maximization and entropy. *Mathematical Finance*, 10(2):259–276, 2000.

[RES05] Berend Roorda, Jacob Engwerda, and Johannes M. Schumacher. Performance of hedging strategies in interval models. *Kybernetika*, 41(5):[575]–592, 2005.

[Ser80] R. J. Serfling. *Approximation Theorems of Mathematical Statistics*. John Wily & Sons, Inc., 1980.

[Shr04] Steven Shreve. *Stochastic Calculus for Finance II: Continuous-Time Models*. Springer Finance, 2004.

[Ste01] M. J. Steele. *Stochastic calculus and financial applications*. Springer, 2001.

A Hardness results

A.1 Computational lower bound: $\#P$-hard result for non-uniform uncertainty sets

This section proves that the problem of exactly computing the price upper bound is $\#P$-hard for non-uniform uncertainty sets, even when the payoff function is convex.

We reduce our problem from the counting subset sum problem. Let $A = \{a_1, ..., a_\tau\}$ be a set of positive integers. The counting problem here is to count the number of subsets $T \subseteq A$ such that the sum of integers in T equals to b. This problem is known to be $\#P$-hard [CHW11].

We use a Cook-Reduction (see *e.g.,* Chapter 17 in [AB09]) and our reduction proceeds as follows. Given a counting subset sum instance $\{A, b\}$, we construct the price upper bound for a τ-round European option hedging game such that the uncertainty set $\mathcal{U}_i$ at the i-th round is $[-\underline{\zeta}_i, \overline{\zeta}_i]$, where $\underline{\zeta}_i = \overline{\zeta}_i = \frac{e^{a_i}-1}{e^{a_i}+1}$. The payoff function is the ordinary call option with strike K, *i.e.,* $g(x) = (x-K)^+$. Since the payoff is convex, the optimal risk-neutral probability measure only has uniform probability mass on $\{-\underline{\zeta}_i, \overline{\zeta}_i\}$ at each round, *i.e.,* the price moves up by a factor $(1+\overline{\zeta}_i)$ with probability $\frac{1}{2}$ and moves down by a factor $(1-\underline{\zeta}_i)$ with probability $\frac{1}{2}$ for each round.

Let P^* be this optimal probability measure. We next build a natural coupling between sampling a subset $T \subseteq A$ and moving a price trajectory in the binomial tree: an element $a_i \in T$ if and only if the price trajectory moves up at the i-th round. Under this probability measure, the total number of subsets T in which the elements sum up to b is

$$2^\tau \Pr_{T \leftarrow P^*}[\sum_{a_i \in T} a_i = b].$$

We next express the probability $\Pr_{T\leftarrow P^*}[\sum_{a_i\in T} a_i = b]$ in terms of option prices. When T is sampled from P^*, the coupled trajectory's final price is

$$S_0 \prod_{a_i\in T}(1+\bar{\zeta}_i)\prod_{a_i\notin T}(1-\underline{\zeta}_i) = S_0\prod_{a_i\in A}\frac{2}{e^{a_i}+1}\prod_{a_i\in T}e^{a_i} = S_0\prod_{a_i\in A}\frac{2}{e^{a_i}+1}\exp\left(\sum_{a_i\in T}a_i\right).$$

Let $\mathcal{C} = \prod_{a_i\in A}\frac{2}{e^{a_i}+1}$. We can see from the above equation that

$$\Pr_{T\leftarrow P^*}[\sum_{a_i\in T}a_i = b] = \Pr_{P^*}[S_\tau = S_0\mathcal{C}\exp(b)].$$

Let us consider options with three different strike prices $K_1 = S_0\mathcal{C}\exp(b-1)$, $K_2 = S_0\mathcal{C}\exp(b)$, and $K_3 = S_0\mathcal{C}\exp(b+1)$. Let their corresponding prices be V_1, V_2, and V_3. We can compute V_1 as follows:

$$V_1 = \int(S_\tau - K)^+dP^* = \sum_{j\geq b}\Pr[\sum_{a_i\in T}a_i = j](S_0\mathcal{C}\exp(j) - K_1).$$

Similarly, we have

$$V_2 = \sum_{j\geq b+1}\Pr[\sum_{a_i\in T}a_i = j](S_0\mathcal{C}\exp(j) - K_2),$$

and

$$V_3 = \sum_{j\geq b+2}\Pr[\sum_{a_i\in T}a_i = j](S_0\mathcal{C}\exp(j) - K_3).$$

From the above equalities, we have

$$\begin{aligned} V_1 - V_2 &= \Pr[\textstyle\sum_{a_i\in T}a_i \geq b]S_0\mathcal{C}(\exp(b) - \exp(b-1)) \\ V_2 - V_3 &= \Pr[\textstyle\sum_{a_i\in T}a_i \geq b+1]S_0\mathcal{C}(\exp(b+1) - \exp(b)). \end{aligned}$$

Therefore, we have

$$\Pr[\sum_{a_i\in T}a_i = b] = \frac{V_1 - V_2}{S_0\mathcal{C}(\exp(b) - \exp(b-1))} - \frac{V_2 - V_3}{S_0\mathcal{C}(\exp(b+1) - \exp(b))},$$

which completes our reduction.

A.2 Information theoretic lower bound: additive dependencies on S_0

We now present a lower bound under the oracle model to justify the necessity of our algorithm's additive dependency on the stock's initial price S_0 and the Lipschitz parameter L, even for a single-round model. Specifically, we shall show that in a single-round hedging game, if the number of queries to the oracle is b, then there exists two payoff functions $g(x)$ and $h(x)$ such that:

1. Both functions are monotonic, L-Lipschitz, and have the same values at the queried points.
2. There exists an S_0 such that the difference between the price upper bounds at S_0 for the functions is $\Theta(LS_0/b)$.

In other words, so long as the number of queries is only polynomial in τ, there will be an additive error that is linear in S_0 and L.

We now explain our construction. Let u be the size of the uncertainty set. Let the query points be $q_1 < q_2 < ... < q_i < 0 < ... < q_b$, where $q_b - q_1 \leq u$. For expositional purpose, let us assume there are at least two queried points on the negative axis. Also, let us assume $\mathcal{U}$ is symmetric, *i.e.,* $\mathcal{U} = [-\frac{u}{2}, \frac{u}{2}]$. Both assumptions can easily be relaxed.

By using an averaging argument, we see that there exists i and j such that ($q_i < 0$ and $q_i - q_{i-1} \geq \frac{u}{2b}$) and ($q_j > 0$ and $q_{j+1} - q_j \geq \frac{u}{2b}$). Now we define $g(x)$ and $h(x)$ as follows:

Definition of $g(x)$: Let C be a sufficiently large number, *e.g.,* $C = (1 + 100u)S_0$.

$$g(x) = \begin{cases} \frac{Lx}{2} & x \leq C \\ Lx & x > C. \end{cases} \tag{54}$$

Definition of $h(x)$: Let $I_1 = [(1 + q_{i-1})S_0, (1 + q_i)S_0]$ and $I_2 = [(1 + q_j)S_0, (1 + q_{j+1})S_0]$.

$$h(x) = \begin{cases} g(x) & x \notin I_1, I_2. \\ Lx & x \in [(1 + q_{i-1})S_0, (1 + \frac{q_{i-1}+q_i}{2})S_0] \\ L(1 + q_i)S_0/2 & x \in [(1 + \frac{q_{i-1}+q_i}{2})S_0, (1 + q_i)S_0] \\ Lx & x \in [(1 + q_j)S_0, (1 + \frac{q_j+q_{j+1}}{2})S_0] \\ L(1 + q_{j+1})S_0/2 & x \in [(1 + \frac{q_j+q_{j+1}}{2})S_0, (1 + q_{j+1})S_0]. \end{cases} \tag{55}$$

We can see that the price upper bound for $g(S_0)$ in this single-round model is $LS_0/2$. For $h(x)$, we can see that in the dual characterization of its optimal solution, the corresponding risk-neutral measure has probability masses only at $r_1 \triangleq \frac{q_i+q_{i-1}}{2}$ and $r_2 \triangleq \frac{q_{j+1}+q_j}{2}$. Thus, the price upper bound for $h(S_0)$ is at least $S_0(\frac{L}{2} + \frac{uL}{2b})$, which completes our argument.

B Reduction to Black-Scholes model from controlled diffusion process

This section elaborates on Section 5 and provides an alternative proof for the continuous-time convergence of our price upper bound to the Black-Scholes model when the payoff function is convex. We shall derive a heuristic partial differential equation (PDE) that characterizes the solution for the controlled diffusion in Theorem 5.1. Then, under additional convexity assumption, we will demonstrate that our PDE is rigorously defined and coincides with the PDE for the Black-Scholes model.

To begin, let us write down a Hamilton-Jacobi-Bellman (HJB) equation informally using (35) and (36). Define $G(t, x) = \max_\xi \mathrm{E}_{(t,x)}[g(S(T))]$, where $\mathrm{E}_{(t,x)}$ denotes the expectation conditional on $S(t) = x$ and using the optimal control (and hence $G(0, S(0))$ is as defined in (35)). Assuming for the moment that $G \in \mathcal{C}^2$, we can heuristically write

$$\begin{aligned} G(t, x) &= \max_\xi \mathrm{E}_{(t,x)}[G(t + \delta, S(t + \delta))] \\ &= \max_\xi \mathrm{E}_{(t,x)}[G(t, x) + G_t(t, x)\delta + G_x(t, x)xR^\delta + \frac{1}{2}G_{xx}x^2(R^\delta)^2 + \cdots] \text{ (by Taylor's series)} \\ &= G(t, x) + G_t(t, x)\delta + \frac{1}{2}G_{xx}x^2 \max_\xi \mathrm{E}_{(t,x)}[(R^\delta)^2] + \cdots \text{ (since } \mathrm{E}_{(t,x)}[R^\delta] = 0) \end{aligned}$$

where $R^\delta = S(t + \delta)/S(t) - 1$, $G_x(t, x) = \frac{\partial}{\partial x}G(t, x)$, $G_{xx}(t, x) = \frac{\partial^2}{\partial x^2}G(t, x)$ *etc.*

Since $U = [0, \underline{\zeta}\overline{\zeta}]$ and $\mathrm{E}_{(t,x)}[(R^\delta)^2] = \sigma(u)^2\delta$, one can attempt to establish that $G(t, x)$ is the solution of the following PDE:

$$G_t(t, x) + \frac{1}{2} \max_{u \in U} \sigma(u)^2 x^2 G_{xx}(t, x) = 0 \tag{56}$$

with the boundary condition $G(T,x) = g(x)$. In general, the solution of this PDE may not exist in the classical sense, and there is no guarantee to coincide with the optimal solution to the control problem in (35) [FS06]. The following theorem, nevertheless, presents a verification of the PDE's solution as the control problem's optimum, under a priori smoothness condition on G:

Theorem B.1. *Suppose $G^* \in \mathcal{C}^2$ is a solution for (56) (hence implying that the payoff function g must be in $\mathcal{C}^2$), and that there exists an optimal $u^*(t,x)$ as the optimal solution to the* $\max$ *in* (56)*. Then $G^*(0,x)$ is the optimal solution to the control problem in (35), and $u^*(t,x)$ is the optimal control in (36).*

Proof. The proof follows by a standard application of Ito's lemma. For any $t < t'$ and x, we have

$$
\begin{aligned}
&\mathrm{E}^*_{(t,x)}[G^*(t', S(t'))] - G^*(t,x) \\
&= \mathrm{E}^*_{(t,x)}\left[\int_t^{t'} G^*_t(t,S(t))dt + \int_t^{t'} G^*_x(t,S(t))dS_t + \int_t^{t'} \frac{1}{2}G^*_{xx}(t,S(t))\sigma^2(u^*(t,S(t)))dt\right] \\
&= \mathrm{E}^*_{(t,x)}\left[\int_t^{t'} \left(G^*_t(t,S(t)) + \frac{1}{2}G^*_{xx}(t,S(t))\sigma^2(u^*(t,S(t)))\right) dt\right] \quad \text{(since } S(t) \text{ is a martingale)} \\
&= 0
\end{aligned}
$$

by (56), where $\mathrm{E}^*_{(t,x)}$ denotes the expectation taken when the control is u^*. On the other hand, any control $u(t,x)$ must satisfy

$$
G_t(t,x) + \frac{1}{2}\sigma(u(t,x))^2 x^2 G_{xx}(t,x) \leq 0 \tag{57}
$$

by the definition of (56). Hence the same argument leads to

$$
\begin{aligned}
&\mathrm{E}^u_{(t,x)}[G(t', S(t'))] - G(t,x) \\
&= \mathrm{E}^u_{(t,x)}\left[\int_t^{t'} G_t(t,S(t))dt + \int_t^{t'} G_x(t,S(t))dS(t) + \int_t^{t'} \frac{1}{2}G_{xx}(t,S(t))\sigma^2(u(t,S(t)))dt\right] \\
&\leq 0
\end{aligned}
$$

where $\mathrm{E}^*_{(t,x)}$ denotes the expectation taken when the control is u. Hence $\mathrm{E}^*_{(t,x)}[G^*_t(t', S(t'))] = G^*_t(t,x) \geq \mathrm{E}^u_{(t,x)}[G(t', S(t'))]$ for any u. Take $t = 0$ and $x = S_0$, we obtain the result. □

From Theorem B.1, we can obtain the Black-Scholes price when the payoff is convex. Observe that (56) can in fact be written as

$$
G_t(t,x) + \frac{1}{2}\underline{\zeta}\overline{\zeta}x^2 G_{xx}(t,x) I(G_{xx}(t,x) \geq 0) = 0 \tag{58}
$$

with boundary condition $G(T,x) = g(x)$. Suppose $G \in \mathcal{C}^2$ is convex in x and so $G_{xx}(t,\cdot) \geq 0$, then the equation (58) reduces to

$$
G_t(t,x) + \frac{1}{2}\underline{\zeta}\overline{\zeta}x^2 G_{xx}(t,x) = 0
$$

which is the ordinary Black-Scholes PDE with zero risk-free rate, whose solution is known to be convex and lies in $\mathcal{C}^2$ when the payoff g is convex and in $\mathcal{C}^2$ [Ste01]. Hence by Theorem B.1 it is the solution to the control problem in (35), which is the limit of our hedging game model.

Lastly, suppose that $G_{xx}(t,\cdot)$ is concave, and so $G_{xx}(t,x) < 0$. The equation (58) then reduces to $G_t(t,x) = 0$, which implies that the solution is $g(x)$, constant over $t \in [0,T]$. If the payoff g is concave, we know again by Theorem B.1 that $g(x)$ is the optimal value of (35). This also characterizes the limit in Section 3.2 when the payoff is concave.

C Missing calculation

Proof of (29)

$$
\begin{aligned}
\frac{\hat{r}^{(2)}}{\hat{r}^{(2)} - \hat{r}^{(1)}} &= \frac{r^{(2)}}{r^{(2)} - r^{(1)}} + \frac{\hat{r}^{(1)} r^{(2)} - \hat{r}^{(2)} r^{(1)}}{\big(\hat{r}^{(2)} - \hat{r}^{(1)}\big)\big(r^{(2)} - r^{(1)}\big)} \\
&= \frac{r^{(2)}}{r^{(2)} - r^{(1)}} + \frac{r^{(2)}\big(\hat{r}^{(1)} - r^{(1)}\big) - r^{(1)}\big(\hat{r}^{(2)} - r^{(2)}\big)}{\big(\hat{r}^{(2)} - \hat{r}^{(1)}\big)\big(r^{(2)} - r^{(1)}\big)} \\
&\geq \frac{r^{(2)}}{r^{(2)} - r^{(1)}} - \frac{r^{(2)}\epsilon}{\big(\hat{r}^{(2)} - \hat{r}^{(1)}\big)\big(r^{(2)} - r^{(1)}\big)} \quad \text{(since } -\hat{r}^{(1)}(\hat{r}^{(2)} - r^{(2)}) \geq 0 \text{ and } |\hat{r}^{(1)} - r^{(1)}| < \epsilon\text{)} \\
&\geq \frac{r^{(2)}}{r^{(2)} - r^{(1)}} - \frac{r^{(2)}\epsilon}{\big(r^{(2)} - r^{(1)}\big)^2} \quad \text{(by construction we have } \hat{r}^{(2)} - \hat{r}^{(1)} \geq r^{(2)} - r^{(1)}\text{)} \\
&\geq \frac{r^{(2)}}{r^{(2)} - r^{(1)}} \left(1 - \frac{\epsilon}{\eta}\right)
\end{aligned}
$$