A Convex Optimization Approach to Multiple Stopping: Pricing Chooser Caps and Swing Options

Shyam S Chandramouli*

(Working Version: All errors are mine)

Abstract

In this current work, we generalize the recent Pathwise Optimization approach of Desai et al. [11] to Multiple stopping problems. The approach also minimizes the dual bound as in Desai et al. [11] to find the best approximation architecture for the Multiple stopping problem. Though, we establish the convexity of the dual operator, in this setting as well, we cannot directly take advantage of this property because of the computational issues that arise due to the combinatorial nature of the problem. Hence, we deviate from the pure martingale dual approach to *marginal* dual approach of Meinshausen and Hambly [19] and solve each such optimal stopping problem in the framework of Desai et al. [11]. Though, this Pathwise Optimization approach as generalized to the Multiple stopping problem is computationally intensive, we highlight that it can produce superior dual and primal bounds in certain settings.

*IEOR Department, Columbia University, New York, NY; sc3102@columbia.edu

1 Introduction

Multiple optimal stopping problems have received considerable attention in recent years. They appear in the context of pricing swing options in electricity markets, e.g. Keppo [17], and exotic options such as chooser caps, e.g. Meinshausen and Hambly [19] in the fixed income derivatives markets. Standard approaches to solving these problems are based on dynamic programming but the curse of dimensionality implies that only small problems can be solved exactly, e.g. Jaillet, Ronn and Tompaidis [16]. As a result, considerable effort has been spent on obtaining good approximate solutions to multiple optimal stopping problems. These approximate solution techniques have been inspired by the earlier work on the pricing of American options.

In the case of American options, the cross-path regression approach of Longstaff and Schwartz [18] and Tsitsiklis and Van Roy [24] has been fundamental. They used simple linear regressions to obtain very good approximation to the value function. Moreover, the resulting approximate value function could then be used to obtain a good feasible exercise policy. The value of this policy therefore constituted a lower, i.e. primal, bound on the true value of the option. In order to determine just how good such a policy was, Haugh and Kogan [15] and Rogers [21] independently developed dual representations that could be used to generate upper, i.e. dual, bounds on the optimal value function of the optimal stopping problem. If the primal and dual bounds were sufficiently close to one another, then we could conclude that the sub-optimal exercise policy corresponding to the lower bound was close to optimal. Since these dual representations were developed, there have been other significant developments including, for example, Andersen and Broadie [2], Chen and Glasserman [10] and Desai, Farias and Moallemi [11]

Meinshausen and Hambly [19] were the first to apply these duality ideas to multiple optimal stopping problems by constructing a dual bound on the *marginal* value of the value function. That is, they computed a dual bound on $J_0^{*n+1} - J_0^{*n}$ where J_0^{*n} is the optimal value function at time $t = 0$ when there are a total of n exercise rights. A perceived weakness of their formulation is that their dual representation requires an optimization over both martingales and stopping times. In contrast, Schoenmakers [22] derived a pure martingale representation which only required an optimization over martingales.

Chandramouli and Haugh [9] develop a unified framework to derive these representations based on the information relaxation ideas of Brown et al. [7]. They use the ideal penalty structure of Brown et al. to derive the pure martingale representation of Schoenmakers [22] and with the additional theory on structural policies, they give a much simpler proof of the marginal representation of Meinshausen and Hambly [19]. They establish numerically the superiority of the pure martingale bound over the *marginal* dual bound and observe that the need of additional stopping time approximations in the Meinshausen and Hambly [19] bound can reduce the quality of the bound. Theoretically speaking, the question is still open about which of the two representations is superior in terms of producing better quality upper bounds?

Very recently Bender [4] extended the marginal dual representation to account for volume constraints. Aleksandrov and Hambly [1] also extended the marginal approach to volume constraints but they used a different solution approach to Bender and considered a more general payoff function whereby the payoff for exercising the ith option in a given period was a decreasing function of i. Other related work includes Bender [5] who extends these ideas to a continuous-time framework.

From an implementation perspective, the work of Schoenmakers [22], Meinshausen and Hambly [19], Bender [4] etc. all compute their value function approximations based on the generalization of the cross-path regression approach of Longstaff and Schwartz [18]. We skip the further details of this implementation here. The value function approximation typically obtained from this regression policy is used to construct the dual penalty functions in the upper bounds discussed above. The dual bound of Schoenmakers [22] is can be computed by solving the dual problem by backward induction. For further details on the induction refer to Bender and Schoenmakers [6]. A generalization of inner simulation procedure of Andersen and Broadie [2] might possibly produce tighter bounds as in the case of a optimal stopping problem. But this advantage is offset by severse computational requirements that might accompany with such an approach.

In the approach of Meinshausen and Hambly [19], a sequence of marginal dual bounds is computed, each dual problem in the sequence solved as a dual of a single optimal stopping problem. This approach further requires an adjustment to the penalties and approximation to the successive stopping times (which are fixed in a sequential way) between marginal computations. These stopping times are again obtained by a greedy policy based on the value function approximation. Refer to Meinshausen and Hambly [19] for more details on the implementation.

Recently, Desai, Farias and Moallemi [11] develop a convex optimization approach to obtain primal and dual bounds on the price of an American Option. Termed as the *pathwise optimization* procedure, the approach is a dual based compared to the existing methods. In more detail, the pathwise approach seeks to minimize the dual optimization problem directly and in the process finds the best approximation architecture to the value function[1]. This approximation architecture is then used to obtain the continuation value estimates and the corresponding basis function approximation for the same.

In this current work, we generalize the approach of Desai et al. to obtain valid primal and dual bounds on a multiple stopping problem. In the process, we highlight that the naive reformulation of Schoenmakers [22] dual representation as a convex optimization problem suffers from combinatorial issues which leads to computational difficulty. Hence, we work with the dual formulation of Meinshausen and Hambly [19] which breaks down the original problem into a sequence of sub problems using the structural properties of the optimal stopping times associated with it. We reformulate each of these sub problems as a pathwise optimization problem. For a reader familiar with the above literature, the intuition of such a reformulation follows immediately as each sub problem of Meinshausen and Hambly is a specific instance of the optimal stopping problem of Desai, Farias and Moallemi [11].

The remainder of this paper is organized as follows. After a brief review of the applications of Multiple stopping problems in Section 2, the mathematical description of the problem is discussed in Section 3. Next, we discuss the pathwise approach in Section 4. We conclude the work with related theory in Section 7 which establishes near optimal bounds produced by the current procedure. Finally, we conclude in Section 6 by performing numerical experiments to compare this new procedure with the existing techniques.

[1]Note: The methods of Longstaff and Schwartz [18] and the subsequent techniques approximate the continuation values whereas the pathwise technique is based on an approximation of the value function

2 Applications

- **Swing Options (Electricity Markets):** We trace it back to the work of Keppo [17] and Jaillet [16] the study of Swing options in the context of Electricity Markets. Typical contract that arise in this market is that the supplier delivers a fixed quantity of power for a fixed price for the contracted time period. In addition to this, the demander has the right to purchase additional quantity L (which is decided at the beginning of the contract) times during the time period. The demander would buy these extra amount at the strike price and not at the electricity spot market price (which would reward him with $(S(t) - K)^+$ if exercised at time t when the spot market price is $S(t)$ with a fixed strike price of K). This helps the demander hedge against sudden needs for additional power. Demander would now choose the best L time slots to exercise his right, which would maximize his profit and hence this boils down as a multiple stopping problem. This is the context of study as well in the recent paper of Meinshausen and Hambly [19], also with volume constraints in Bender [4] (which gives the demander the right to exercise more than once in the same time slot) and with refractive index in Bender [5] which constraints the time gap between successive exercise of rights.

- **Chooser Caps (Fixed Income Markets):** In similar spirit, traded are the Caplets/ Chooser Caps in Fixed Income Markets which the owner used to hedge interest rate risks. In this setting, say a chooser cap on LIBOR with L exercise rights, would give the owner the right to exercise L times with a payoff of $\delta(L_i - K)^+$ if exercised at time i. Here, L_i is the LIBOR at time i, K pre determined strike price and δ, a constant. The work of Meinshausen and Hambly [19] and Schoenmakers et al. [3] explore the modeling and simulation of Chooser Caps as a multiple stopping problem in greater detail.

- **Storage Valuation (Gas Storage):** Typical storage valuation problems, is to decide on quantity of gas or other commodity to order now and storing as inventory versus ordering at a later time period. Under reasonable modeling assumptions, the optimal storage valuation problem can be treated as a multiple stopping problem [13]

- **Portfolio Liquidation:** We again refer to Hambly et al. [13] for a discussion on the optimal liquidation problem as a multiple stopping problem. Further treatment and an numerical example of optimal liquidation can be seen in Schoenmakers et al. [6]. The interesting variant in liquidation problems is the price impact of a particular execution over the future executions and has to be treated as a coupled stochastic control problem.

3 Formulation

We now consider multiple optimal stopping problems which are defined as follows. There are T time periods and a probability space $(\Omega, \mathcal{F}, P)$ with an associated filtration, $\mathbb{F} := \{\mathcal{F}_t : t = 0, \ldots, T\}$, where $\mathcal{F}_T = \mathcal{F}$, $\mathcal{T} := \{0, 1, ...T\}$. α is the discount factor.m_j denotes the number of remaining exercise rights at time j. We have a positive adapted reward process, g_t, which satisfies $\sum_{t=1}^{T} \mathrm{E}[|g_t|] < \infty$. This reward process may be *exercised* $L \leq T$ times subject to the constraint that it may not be exercised more than once on any given date (This assumption can be relaxed for our problem, about which we discuss later). For notational convenience below, we will define

$\mathcal{F}_t := \mathcal{F}_T$ and $g_t := 0$ for all $t > T$. For the same reason we will also suppress the dependence of all stochastic processes on any state variables that may be driving the dynamics of the system. The goal then is to maximize the expected reward by optimally choosing the L exercise times. We have the following definition.

Definition 3.1 *For $0 \leq t \leq T$ and each L, we define $\mathcal{S}_t(L)$ to be the set of $\mathbb{F}$-stopping vectors $\tau_t := \{\tau^1, \ldots, \tau^L\}$ such that $t \leq \tau^1 < \tau^2 < \cdots < \tau^L \leq T$.*

For a given $\tau_t \in \mathcal{S}_t(L)$ we define the value function associated with τ_t to be

$$V_t^{\tau_t, L} \;:=\; \mathrm{E}_t \left[\sum_{l=1}^{L} g_{\tau^l} \right]. \tag{1}$$

The multiple stopping problem at time t may then be formulated as finding $\tau_t^* := \{\tau^{*1}, \ldots, \tau^{*L}\} \in \mathcal{S}_t(L)$ such that for $0 \leq t \leq T$,

$$J_t^{*L} \;:=\; \sup_{\tau_t \in \mathcal{S}_t(L)} V_t^{\tau_t, L} \;=\; \mathrm{E}_t \left[\sum_{l=1}^{L} g_{\tau^{*l}} \right] \tag{2}$$

where $\mathrm{E}_t[\cdot]$ denotes an expectation conditional on $\mathcal{F}_t$ and the sup in (2) is understood as the *essential supremum*. The process J_t^{*L} defined in (2) is known as the *Snell envelope* of g under L exercise rights. It is the optimal time t value function when there are L exercise rights remaining.

We refer the readers to Carmona and Touzi [8] for a detailed formulation of the multiple stopping problem as a dynamic programming problem. We know that the multiple optimal stopping problem can be reduced to L nested single exercise optimal stopping problems. This is achieved by noting that $J^{*0} = 0$, that J^{*1} is equal to the Snell envelope of g, and that for any $l > 1$, J^{*l} is the Snell envelope of the generalized cash flow process $g_t + \mathrm{E}_t \left[J_{t+1}^{*l-1} \right]$. Finally, we let J_t^* denote the optimal value function for the multiple stopping problem at time t with the understanding that $J_t^* = J_t^{*l}$ if there are l exercise opportunities remaining at that time.

Intuitively, the investor decides the first stopping time, τ^1 of the stopping vector τ_t in the following way: At a particular time i, the investor decides between the 2 possibilities, to take the cash flow g_i at time i and start a new stopping problem from time $i+1$ with $L-1$ contracts versus not exercise his right and have a stopping problem with L exercise rights from time $i+1$. The investor proceeds with this optimal decision at each time step.

With the transformation to generalized cash flows, we can essentially adapt the theory for single optimal stopping as discussed in Neveu [20]. Caramona and Touzi [8] present a detailed study of the same. We summarize below the optimal stopping times, starting from a time t

$$\begin{aligned}
\tau^1 &= \inf_{k \geq t} \{ J_k^{(L)} = g_k + E_k J_{k+1}^{(L-1)} \} \\
\tau^2 &= \inf_{k \geq \tau_1} \{ J_k^{(L-1)} = g_k + E_k J_{k+1}^{(L-2)} \} \\
&\;\;\;

If any of the above sets are empty, then all the stopping times from that step and onwards can be set to ∞. We can see from the description above that for the case of one exercise right, i.e. $L = 1$, the analysis boils down to the classical single stopping or American option [20].

We now define the Doob martingale of the Snell envelope process. Let $\mathcal{S}$ denote the space of real-valued measurable functions that are defined on the state space $\mathbb{R}^n$. We now define an operator Δ that maps $\mathcal{S}$ to the space of real-valued measurable functions on $\mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^m$ according to

$$(\Delta V_t)(x_t, x_{t-1}, u_{t-1}) \;:=\; V_t(x_t) \;-\; \mathbb{E}\left[V_t(x_t) \,|\, x_{t-1}, u_{t-1}\right]. \tag{4}$$

Loosely speaking, Δ is an operator on (approximate) value functions. Note in particular that $\mathbb{E}_0\left[(\Delta V_t)(x_t, x_{t-1}, u_{t-1})\right] = 0$ for all integrable V_t. Let $\mathcal{D}$ be the space of real-valued functions on $\mathbb{R}^n \times \mathcal{T}$ such that if $V \in \mathcal{D}$ then $V_t := V(\cdot, t)$ is measurable and $\mathbb{E}_0\left[|V_t(x_t)|\right] < \infty$ for all $t \in \mathcal{T}$

Definition 3.2 *Let M^{*l} be the Doob martingale of the Snell envelope for l exercise rights. That is,*

$$J_t^{*l} \;=\; J_0^{*l} \;+\; M_t^{*l} \;-\; A_t^{*l}, \quad t = 0, \ldots, T \tag{5}$$

where $M_0^{*l} = A_0^{*l} = 0$ and A^{*l} is predictable and non-decreasing for each $l = 0, \ldots, L$.

Note also that since A^{*l} is predictable we can apply the Δ operator of (4) to J^{*l} and obtain

$$\Delta J_t^{*l} \;=\; M_t^{*l} \;-\; M_{t-1}^{*l}. \tag{6}$$

We now consider an alternative approach due to Meinshausen and Hambly [19]. They focus on characterizing the marginal value of having one additional exercise right. We therefore need the following definition.

Definition 3.3 *For every $\tau \in \mathcal{S}_t(n)$ and $1 \leq n \leq L$ the marginal value is defined as*

$$\Delta V_t^{\tau,n} \;:=\; V_t^{\tau,n} \;-\; V_t^{\tau,n-1}. \tag{7}$$

Note that the definition of Δ operator in (7) is different to its definition in (4). We can distinguish between the two definitions by noting that that argument of the Δ operator in (7) must have a superscript, n, denoting the number of exercise opportunities remaining. We use $\Delta J_t^{*,n}$ to denote the marginal value of the optimal value function so that, for example, $\Delta J_t^{*,1}$ is simply the time t value of the usual optimal stopping problem with just one exercise opportunity. This follows since $J_0^{*0} = 0$. The following result is intuitively obvious but nonetheless it requires some effort[2] to prove it.

Proposition 1 *Let $\tau_t^* = \{\tau^{*1}, \ldots, \tau^{*n-1}\} \in \mathcal{S}_t(n-1)$ be an optimal policy at time t when there are $n-1$ exercise rights remaining. Then there exists an optimal policy, $\sigma_t^* \in \mathcal{S}_t(n)$, for the time t problem when there are n exercise rights remaining such that $\tau_t^* \subset \sigma_t^*$.*

[2]See Proposition 5.1 in Bender [4] which in turns uses Proposition 5.2 in Meinshausen and Hambly [19].

3.1 Martingale Duality

Now that we have discussed the two different representations of the multiple stopping problem, in this section, we derive the dual representations of these problems. Intuitively, the dual problems give the decision maker additional information about the future process of the system and penalizes him (with a negative reward) for having this additional information. Since, the additional information can let the decision maker make better decision than the original setting, with a zero penalty the decision maker earns a higher value in the new setting.

Brown et al. [7] establish a methodology for deriving good penalties for a general dynamic programming problem. They show that there exists an ideal penalty for these decision problems and when the decision maker is penalized accordingly, the value of both the primal and dual problems are the same. In order that an ideal penalty is not available, any other feasible penalty[3] would give a upper bound on the original primal problem. For a reader, familiar with the work of Brown et al. [7], we work here with the full information relaxation.

Hence, we work with these dual problems to derive efficient upper bounds on the multiple stopping problem. Below, we give the representation of the Martingale dual problem which we later use in our approach.

We define for each $t \in \mathcal{T}$, the *pure martingale duality upper bound operator*, $F_t : \mathcal{D} \mapsto \mathcal{S}$ according to,

$$(F_t J)(x) = \mathrm{E}_t \left[\max_{t \le j_1 < \ldots j_L \le T} \sum_{k=1}^{L} \alpha^{\mathrm{j}_k - t} g(x_{j_k}) - \sum_{p=t+1}^{j_L} \alpha^{p-t} \left\{ J_p^{(m_p)}(x_p) - \mathrm{E}_{p-1}[J_p^{(m_p)}(x_p)] \right\} \right] \tag{8}$$

and the *marginal martingale duality upper bound operator* as,

$$(F_t \Delta V_t^{\tau,n})(x) = \mathrm{E}_t \left[\max_{j_n \notin \{\tau_1, \tau_2 \ldots \tau_{n-1}\}} g(j_n) - \sum_{j=t+1}^{j_n} \{ V_j^{\tau,m_j} - V_j^{\tau,m_j-1} - \mathrm{E}_{j-1}[V_j^{\tau,m_j} - V_j^{\tau,m_j-1}] \} \right] \tag{9}$$

The dual bound operator when operated on any feasible J, would return an upper bound value on the corresponding value-to-go function and return the optimal value function when operated on J^* (Verbatim for ΔV). This is summarized in the following lemma:

Lemma 1 ***(Martingale Duality)***

(Weak Duality) *For any $J \in \mathcal{P}$ and all $x \in \mathcal{S}$ and $t \in \mathcal{T}$, $J_t^*(x) \le F_t J(x)$ and in addition for each $n \in \{1, 2, \ldots L\}$ and $\tau \in S_t(L)$, $\Delta V_t^{*,n}(x) \le (F_t \Delta V_t^{\tau,n})(x)$*

(Strong Duality) *For all $x \in \mathcal{S}$ and $t \in \mathcal{T}$, $J_t^*(x) = F_t J^*(x)$ and in addition for each $n \in \{1, 2, \ldots L\}$, $\Delta V_t^{*,n}(x) = (F_t \Delta V_t^{*,n})(x)$*

[3] Refer to Brown et al. [7] for a detailed discussion on feasible penalties and ideal penalty

Proof: Refer to Chandramouli and Haugh [9]

The result when viewed as an optimization problem: for each $t \in \mathcal{T}$ and $x \in S$,

$$J_t^*(x) = \inf_{J \in D} F_t J(x) \tag{10}$$

and for each n,

$$\Delta V_t^{*,n}(x) = \inf_{\Delta V \in D, \tau \in S_t(n)} (F_t \Delta V_t^{\tau,n})(x) \tag{11}$$

The optimization approach discussed next intends to find the best J (ΔV) by minimizing the RHS of 10 (11).

4 Convex Optimization Approach

Following the approach of Desai et al. [11], we seek a parametric approximation for the value-to-go function of the multiple stopping problem. A brief review of such approximation dynamic programming (ADP) approaches can be found in Bertsekas [12] and Topagolu [23]. We wish to highlight here that most of the recent pricing methods in literature starting from Longstaff and Schwartz [18], Haugh and Kogan [15], Andersen and Broadie [2], Rogers [21] all use the approximate dynamic programming framework. The optimal value function, we posit, as a set of linear basis functions and try to find the best set of regression coefficients that would suit the approximation. The performance of the approximation should improve with the number of chosen orthogonal basis functions (as in the limit, these basis functions span the Hilbert Space that contains the original optimal value function).

The idea to develop dual bounds and primal policies based on the optimization of the dual problem stems from Desai et al. [11]. Unlike the aforementioned earlier papers which focus on the primal problem, Desai et al. center their attention on the dual problem. The numerical experiments in their work suggest that such an approach can obtain superior primal and dual bounds as compared to the approach that uses the method of Longstaff and Schwartz (LS) [18] in conjunction with Haugh and Kogan(HK). But the pathwise approach typically takes a lot of computational work than LS and HK together. The extra work arrives in setting up and solving the LP reformulation of the basis function approximation of the dual convex problem.

Observe that there is a trade off between the quality of bound and the computational time requirements. This is further supported by the dual bound obtained by Andersen and Broadie [2], which is superior than both Desai et al. and Haugh and Kogan but needs much more computational time than both these methods. This is because the approach of Andersen and Broadie [2] requires the need of additional inner simulation paths at each exercise step.

The numerical experiments of Desai et al. compare the bounds of various methods for a fixed set of basis functions. From the literature, we can see that, with a good choice of basis functions the LS-HK procedure can produce much tighter bounds than indicated by Desai et al.. In the absence of a theory for the right choice of basis functions, we believe pathwise procedure might still be helpful in certain settings when fine tuning of basis functions becomes absolutely necessary. Hence, in the spirit of Desai et al., we restrict the attention to the subspace spanned by the chosen basis functions i.e. functions of the form

$$(\Phi r^{(n)})_t(x) := \sum_{i=1}^{K} r_i^{(n)} \phi_i(x,t) \quad \forall x \in S, t \in T \tag{12}$$

where $r^{(n)} \in R^K$ is the weight vector when n exercise rights are remaining. In the current approximation architecture we represent the value function for each n by a linear combination of basis functions. As an alternative, one can try approximations where basis functions can depend on n.

The general idea of basis function architecture is to obtain a low dimensional approximation to the original value function. We wish to find a good architecture where in, instead of optimizing over the entire function space, when we optimize only over this restricted space (by basis functions), we are able to obtain reasonably good solutions.

Using the approximation in 12, we can rewrite the optimization problem 10 as,

$$\inf_r F_0 \Phi r^{(L)}(x) \tag{13}$$

where

$$(F_t(\Phi r^{(L)}))(x) = \mathrm{E}_t \left[\max_{t \le j_1 \ldots j_L \le T} \sum_{k=1}^{L} \alpha^{\mathrm{J}_k - t} g(x_{j_k}) - \sum_{p=t+1}^{j_L} \alpha^{p-t} \Delta(\Phi r^{(m_p)})_p(x_p, x_{p-1}) \right] \tag{14}$$

This optimization problem is the general multiple stopping *Pathwise Optimization* (PO) problem. Solving this optimization problem, provides an approximation architecture to the multiple stopping problem. This architecture can then be used to construct a feasible primal policy. We denote the solution to the optimization problem 13 by $r_{PO} := \{r_{PO}^{(1)}, .., r_{PO}^{(L)}\}$, where $r_{PO}^{(n)}$ is the solution when n exercise rights are left.

We make the following observation, en route to solving the problem:

Lemma 2 *For every $t \in T$, $x \in S$, the function $r \mapsto F_t \Phi r^{(L)}(x)$ is convex in r*

Proof: Let us consider the term inside the expectation. The L stopping times can be chosen in $\binom{n}{L}$ ways where $(n = T - t)$. So, the inner term is a maximum of $\binom{n}{L}$ functions of r. Each such function is affine in r (sum of affine functions which is again affine). Now, maximum of affine functions is convex and non negative linear combination of convex functions is convex. So, the convexity of the upper bound operator follows.

Before prescribing an approach to solve the convex optimization problem in 13, we describe a method to generate a lower bound policy (primal bound) once the approximation architecture is available. The reason we describe this earlier will become clear in the coming sections where this procedure would be used as a sub routine while solving the marginal dual problem.

4.1 Primal Bound (Lower) : Cross-Path Regression

We here describe a policy construction based on estimation of continuation values in similar spirit of Longstaff and Schwartz [18]. Then, a greedy execution with respect to this policy can be evaluated

to find the primal bounds. Once we have the solution to the convex optimization problem 13 ($r_{PO}^{(n)}$), then construct the upper bound estimates on various continuation values, $C_t^{*,n}(.)$. By weak duality for multiple stopping 1, we have,

$$C_t^{*,n}(x_t) \leq \mathrm{E}_t \left[\max_{t+1 \leq j_1 \ldots j_n \leq T} \sum_{k=1}^{n} g(x_{j_1}) - \sum_{p=t+2}^{n} \Delta(\Phi r_{PO}^{(m_p)})_p(x_p, x_{p-1}) \right] \tag{15}$$

Define the point estimate for each sample path for each n by,

$$c_t^{(i),n} := \max_{t+1 \leq j_1 \ldots j_n \leq T} \sum_{k=1}^{n} g(x_{j_1}) - \sum_{p=t+2}^{n} \left\{ (\Phi r_{PO}^{(m_p)})_p(x_p^i) - \mathrm{E}_{p-1} \left[(\Phi r_{PO}^{m})_p(x_p^i) \right] \right\} \tag{16}$$

We immediately note that the continuation values for each sample path i can be recursively represented as a dynamic programming problem and hence solved via backward induction. Refer to details in Algorithm 4.1 below. Hence, the computational requirements for finding a lower bound policy is same as that of the Longstaff and Schwartz regression procedure.

The first step is to choose a set of basis functions $\psi_1(X) \ldots \psi_m(X)$. These basis functions define a *linear architecture* that will be used to approximate the continuation (Q-value) functions. In particular, we will approximate $Q_t^{(k)}(X_t)$ with

$$\tilde{Q}_t^{(k)}(X_t) = r_t^{(k,1)} \psi_1(X_t) + \cdots + r_t^{(k,m)} \psi_m(X_t) \tag{17}$$

where $Q_t^{(k)}(X_t)$ is the continuation value and $r_t^{(k)} := (r_t^{(k,1)} \ldots r_t^{(k,m)})$ is a vector of time t parameters when the number of exercise tights left is k. These parameters are estimated by regression in the following way. In total, NT regression computations are needed.

Now, a greedy policy, π, with respect to the regressed weights $r_t^{(n)}$ is defined as follows: (This policy hinges on the structure of the optimal stopping times discussed in section 3)

$$\pi^n(X_t) = \begin{cases} \textit{Continue}, & \text{if } g_t(X_t) < \tilde{Q}_{t+1}^{(n)} - \tilde{Q}_{t+1}^{(n-1)} \\ \textit{Stop}, & \text{otherwise} \end{cases}$$

$\pi^n(.)$ is a greedy exercise policy defined for each n. Starting from $n = L \to 1$, an approximation to the L optimal stopping times are found and hence the corresponding primal lower bound.

Algorithm 1 (4.1) Policy for Lower Bound

for $t = N-1 \to 1$ **do**
  if $t = N-1$ **then**
    $c_t^{(i),n} = g(x_T^i) \quad \forall n = 1, 2, ...L$
  else
    for $n = 1 \to L$ **do**
      $c_{exec} = g(x_{t+1}^{(i)}) + c_{t+1}^{(i),n-1} - \left\{ (\Phi r_{PO}^{(m_p)})_{t+2}(x_{t+2}^{(i)}) - \mathrm{E}_{t+1}[(\Phi r_{PO}^{(m_p)})_{t+2}(x_{t+2}^{i})] \right\}$
      $c_{cont} = c_{t+1}^{(i),n} - \left\{ (\Phi r_{PO}^{(m_p)})_{t+2}(x_{t+2}^{(i)}) - \mathrm{E}_{t+1}[(\Phi r_{PO}^{(m_p)})_{t+2}(x_{t+2}^{(i)})] \right\}$
      $c_t^{(i),n} = \max\{c_{exec}, c_{cont}\}$
    end for
  end if
  Compute the weight vectors $r_t^{(n)}$ via the regression
  $r_t^{(n)} \in \mathrm{argmin}_k \sum_{i=1}^{S} (\psi_t r_t^{(n)}(x_t^{(i)}) - c_t^{(i),n})^2$
end for
Use the greedy algorithm of Section 4.1 to find the stopping times and the lower bound.

4.2 Solving the Convex Problem by Sampling

In order to solve the convex optimization problem by a stochastic gradient method we have to solve a dynamic program at each step. This arises because of the combinatorial nature of the problem which allows for only a fixed number of exercise rights over a discrete set.

One practical way is to reformulate as an LP and solve using any state-of-art LP solvers. To implement the same, we observe that the expectation in the objective of the dual optimization problem can be replaced by its sample average i.e.

$$(\tilde{F}_t(\Phi r^{(L)}))(x) = \sum_{i=1}^{S} \left[\max_{t \le j_1 \dots j_L \le T} \sum_{k=1}^{L} \alpha^{\mathrm{j}_k - t} g(x_{j_k}^{(i)}) - \sum_{p=t+1}^{j_L} \alpha^{p-t} \Delta(\Phi r^{(m_p)})_p(x_p^{(i)}, x_{p-1}^{(i)}) \right] \tag{18}$$

where S is the number of monte carlo paths chosen to approximation the expectation in the dual problem. By the strong law of large numbers, the sample average $(\tilde{F}_t(\Phi r^{(L)}))(x) \to (F_t(\Phi r^{(L)}))(x)$ as $S \to \infty$.

The expectations in the penalty terms are estimated by inner sampling in Desai et al. In this work, we find a closed form expression to evaluate the same. Hence, we avoid inner sampling. Following the ideas of Desai et. al, the sampling variant of the optimization problem 13 can be reformulated as 4.2,

The number of constraints in LP 4.2 ($\binom{T}{L}$) increases exponentially with the number of exercise rights L. Since, Desai et al. [11] formulate the pathwise LP for an American Option ($L = 1$), the size of the LP is reasonable to handle. But in typical commodity trading problems, the size of the swing is very large, say $L = 50 - 100$ [19, 4] hence 4.2 is impossible to solve on a PC to obtain good bounds. Therefore, we try to exploit the structural properties of the multiple stopping problem (see Section 3) to reduce it to a sequence of reasonable size LP's. One immediate approach is to set the marginal dual problems in the pathwise optimization framework. Such an approximation is

Algorithm 2 (4.2) LP approach to pathwise optimization

Function LP

$\min_{r,u} \quad \frac{1}{S}\sum_{i=1}^{S} u_i$

subject to

$u_i + \sum_{p=t+1}^{j_L} \alpha^{p-t} \left\{ (\Phi r^{(m_p)})_p(x_p^{(i)}) - \mathrm{E}_{p-1}[(\Phi r^{(m_p)})_p(x_p^{(i)})\right\}] \geq \sum_{k=1}^{L} \alpha^{j_k - t} g(x_{j_k}^{(i)})$

$\forall i, j^L \in \mathcal{C}$

where $\mathcal{C} = \{j^L : t \leq j_1 < j_2 < ....j_L \leq T\}$

end function

interesting, since each of the marginal dual problems is a optimal stopping problem, the pathwise approach for which is described in Desai et al. [11]. Hence, we explore this approach more in this section.

Now, on approximation by basis functions, the marginal dual bound operator (defined in Equation 9) takes the following form,

$$F_t \Phi r^{(n)}(x) = \mathrm{E}_t \left[\max_{j_n \notin \{\tau_1, \tau_2 .... \tau_{n-1}\}} g(j_n) - \sum_{j=t+1}^{j_n} \Delta(\Phi r^{(m_j)})_p(x_j, x_{j-1}) \right] \tag{19}$$

The marginal dual optimization problem (11) for each n reduces to,

$$\inf_{r^{(n)}} F_0 \Phi r^{(n)}(x) \tag{20}$$

With a slight abuse of notation, we use the same operator for pure dual as well as the marginal optimization problem. Henceforth, the discussion is restricted to the marginal optimization problem.

Lemma 3 *For every $t \in T$, $x \in S$, $n \in 1, 2...L$, the function $r \mapsto F_t \Phi r^{(n)}(x)$ is convex in r*

Proof: Each marginal pathwise optimization problem is a optimal stopping problem. Hence, convexity of the marginal dual operator follows for the same reasons as in Desai et al. [11].

We have already noted in section 3, in the approach of Meinshausen and Hambly [19] the optimal structure of stopping times can be exploited to obtain the marginal dual problems. The value functions of these marginal problems can be summed up to obtain the upper bound on the option price (see Meinshausen and Hambly). As already noted, the marginal approach also requires an approximation to the optimal stopping times in addition to the value function approximation. In order to solve the nth marginal dual problem, we need the approximation of earlier $(n-1)$ optimal stopping times as discussed in section 3.

Since, Meinshausen and Hambly [19] is a primal based approach, the value function approximations are already computed before solving the marginal dual problems. Hence, the necessary stopping times can be computed using a greedy policy based on these value function approximations. Whereas, our algorithm is a dual based technique, hence, we follow the steps discussed next to obtain the stopping time approximations.

Solving the $(n-1)^{th}$ marginal dual value problem outputs the optimal values of the variables, $r_{PO}^{(n-1)}$. A control policy when $n-1$ exercise rights are remaining can be constructed using the approximation of continuation value discussed in previous section 4.1. This is a cross path regression approach similar in spirit to the Longstaff and Schwartz [18]. More specifically, we run the $(n-1)^{th}$ iteration of the loop discussed in section 4.1. As it is clear from the discussion there, the $(n-1)^{th}$ iteration (which produces the value function approximations when $n-1$ exercise rights are remaining) depends only on the continuation values for marginal problems $k = 1, 2...(n-2)$ which are available by now. Once, this control policy is generated, we use a greedy policy to just fix the $(n-1)^{th}$ approximate stopping time. Then, continue solving the nth marginal pathwise optimization problem.

Since, after solving the kth dual marginal problem we are computing an approximation architecture with k exercise rights remaining. when we are complete solving all the marginal pathwise optimization problems, we have the complete architecture ($r^{(n)}$ for $n \in \{1, 2..L\}$). Also, since we are running the regression procedure of section 4.1 at the end of each marginal problem we have a control policy at the end of the procedure. This policy can be then used to obtain a lower bound by generating a new set of sample paths.

In short, we use the pathwise optimization approach of Desai et al. [11] in conjunction with the marginal dual value approach of Meinshausen and Hambly [19] and solve a sequence of LP's to obtain the upper bounds on the corresponding marginal dual value.

We summarize the complete algorithm next, (see 4.2)

Algorithm 3 Modified LP approach to pathwise optimization

for $n = 1 \rightarrow L$ **do**
  Solve LP (n)
  Run StopTime (n) (%Finds the approximation to the nth stopping time)
end for

Function LP(n)
$\min_{r,u} \ \frac{1}{S}\sum_{i=1}^{S} u_i$
subject to
  $u_i + \sum_{j=t+1}^{j_n}((\Phi r^{*,(m_j)})_j - (\Phi r^{*,(m_j-1)})_j - \mathrm{E}_{j-1}[(\Phi r^{*(m_j)})_j - (\Phi r^{*(m_j-1)})_j]) \geq \alpha^{j_n-t} g(x_{j_n})$
  $\forall i, j_n \notin \{\tau_1, \tau_2, ...\tau_{n-1}\}$
end function

Function StopTime (n)
  Solve for $\psi_t^n \quad \forall t = \{1, 2, ..T\}$ as indicated in section 4.1
  Fix an approximation of nth stopping time τ_n based on a greedy policy.
end function

From an implementation perspective, to compute the inner expectations we do not use inner sampling but we obtain it rather as a closed form analytic expression. We discuss it more in the numerical experiments section 6.

For $n = 1$, our method is equivalent to the approach of Desai et al. [11].

4.3 Unbiased Upper Bound

The upper bound obtained from solving the marginal dual problem 20 is a biased estimate. This is due to Jensen's inequality (see Glasserman [14]) and also follow the discussion in Desai et al. Hence, an unbiased upper bound is computed by generating a new set of monte carlo paths. Now, again there are 2 approaches to compute the upper bound Schoenmakers vs. Meinshausen-Hambly. Chandramouli and Haugh [9] elucidate with a numerical example that the Schoenmakers pure dual upper bound is superior because there is only a value function approximation whereas Meinshausen and Hambly has a additional stopping time approximation.

Hence, to computer the unbiased upper bound, we compute the following quantity:

$$\frac{1}{S}\sum_{i=1}^{S}\left[\max_{t\leq j_1\ldots j_L\leq T}\sum_{k=1}^{L}\alpha^{j_k-t}g(x_{j_k}^{(i)})-\sum_{p=t+1}^{j_L}\alpha^{p-t}\Delta(\Phi r^{(m_p)})_p(x_p^{(i)},x_{p-1}^{(i)})\right] \tag{21}$$

which is the sample average approximation (SAA) of the pure dual bound operator 8. On each sample path, the maximization over L stopping times can be computed by backward induction as discussed earlier (further details see Schoenmakers et al. [6])

By the Strong Law of Large Numbers (SLLN), this will converge to $F_t\Phi r(x)$ as $S\to\infty$.

5 Other Extensions

5.1 Volume Constraints

In practice swing options can often be exercised more than once per time period. It is therefore necessary to consider more general volume constraints. In recent work Bender [4] extended the dual characterization of Meinshausen and Hambly [19] to the case where the option could be exercised at most k_t times in period t, but still maintaining the constraint of at most L exercise rights. Aleksandrov and Hambly [1] then extended the result of Bender to the case where $g_t^{(i)}\geq g_t^{(i+1)}$ for $i=1,\ldots,k_t$ and $t=0,\ldots,T$ where $g_t^{(i)}$ is the payoff when the option is exercised for the ith time at time t. That is while Bender assumed that each option exercise at time t had identical payoffs, Aleksandrov and Hambly allow for the possibility that the ith option payoff at time t is a non-increasing function of i. We define

$$\begin{aligned} N_t^m(\tau^1,\ldots,\tau^m) &:= \#\{j\,:\,\tau^j=t\}\\ \bar{N}_t^m(\tau^1,\ldots,\tau^m) &:= \#\{j\,:\,\tau^j\leq t\}\end{aligned}$$

so that $N_t^m(\tau^1,\ldots,\tau^m)$ is the number of stopping times taking the value t and $\bar{N}_t^m(\tau^1,\ldots,\tau^m)$ is the number of stopping times taking the value at most t. Let $\mathcal{C}_V$ be the space of vectors in $\mathcal{T}^L$ that satisfy the capacity constraints so that

$$\mathcal{C}_V\ :=\ \{j^{(L)}=\{j_1\leq j_2\leq\cdots\leq j_L\}\in\mathcal{T}^L{:}N_u^L(j_1,\ldots,j_L)\leq k_u\ \text{ for } u=0,1\ldots,T\}.$$

The marginal dual representations were obtained initially by Bender [4]. Chandramouli and Haugh [9] present a shorter proof of the same result using ideas from information relaxation. Once,

we have the marginal dual value problems, then the pathwise approach of preceding section can be used to obtain a upper bound on the multiple stopping problem with volume constraints. At a high level, the only part that changes significantly in the entire algorithm is the step that fixes the stopping time at every iteration. When computing the stopping time, the greedy policy is allowed to exercise more than once at particular time slot, hence the policy needs to be designed to execute in such a way that it does not violate the volume constraints. More details on the policy can be found in Bender [4].

5.2 Refractive Index

In practice, contracts are often written in such a way that if an option right is exercised at time t then the holder of the option is forbidden from exercising another right until a further δ_t time units have elapsed. In practice these constraints are often referred to as refractive index constraints. The problems that we considered in earlier sections all have $\delta_t = 1$ and so they are special cases of the problem considered here. We now define $\mathcal{C}_R$ to be the space of vectors in $\mathcal{T}^L$ that satisfy the refractive index constraints so that

$$\mathcal{C}_R \ := \ \{j^{(L)} = \{j_1 \leq j_2 \leq \cdots \leq j_L\} \in \mathcal{T}^L : j_{k+1} - j_k \geq \delta_{j_k} \ \text{ for } k =, 1 \ldots, L-1\}.$$

Our model cannot handle multiple stopping problems for the same reasons the Meinshausen and Hambly approach fails to do so. This is because the proposition 1 regarding the stopping times does not hold any longer. Hence, we cannot exploit such structural property to develop marginal bounds. Hence, the only approach to this problem is to develop pure martingale dual bounds. But as we saw earlier the pure martingale bound gives rise to computational difficulties in our formulation. Hence, this is a limitation of our approach. For these problems, Longstaff and Schwartz approach used along with Schoenmaker's pure dual bounds gives the required bounds.

6 Numerical Experiments

6.1 Model: Geometric Brownian Motion

Assume Black-Scholes framwork and risk neutral asset pricing, the stock price at time t (S_t) follows a Geometric Brownian Motion (GBM) given by,

$$dS_t = rS_t dt + \sigma S_t dW_t \tag{22}$$

where r is the risk free interest rate, σ is the volatility and $W_t \sim \mathrm{N}(0, 1)$. The stock price process is generated by the SDE 22. The model parameters are set to $\sigma = 0.2$, $r = 0.06$ and $S_0 = 36$. The strike price, K, for each put option was set to 40 and we assumed that each swing option had $T = 1000$ exercise opportunities with the number of exercise rights varying from $L = 1$ to $L = 25$.

Lower and upper bounds were generated by the pathwise approach described earlier and compared with the lower bounds computed using a simple generalization of the cross-path regression approach of Longstaff and Schwartz [18] and Tsitsiklis and Van Roy [24]. Details of this generalization are given in Meinshausen and Hambly [19]. The benchmark for the upper bounds is the pure martingale bound of Schoenmakers [22]

Basis Functions

We use the following 4 basis functions for the computation:

$$\phi_1(x) = 1 \quad \phi_2(x) = x \quad \phi_3(x) = x^2 \quad \phi_4(x) = x^3 \tag{23}$$

These are the first 4 linear polynomials of x, We discuss the nice property of these basis functions for Gaussian models like GBM and AR(1) discussed earlier. Recall that the value function approximations, are the linear combinations of these basis functions. Since the past work of Haugh and Kogan [15], Andersen and Broadie [2],Meinshausen and Hambly [19], Benders [4] is based on approximation to the continuation values, the inner expectations were approximated once again by sampling. Inner Expectations arise in the penalty term in the form $(E_t(V_{t+1}))$. But in the current work, we approximate the value function directly instead of the continuation values. Hence, it is possible to obtain nice analytic expressions for the inner expectations and avoid sampling. This way, we would have the exact value rather than close estimated values which should improve accuracy and speed up the computation drastically.

$$\begin{aligned}
\mathrm{E}_t\phi_1(X_{t+1}) &= 1 \\
\mathrm{E}_t\phi_2(X_{t+1}) &= X_t * (1+\alpha dt) \\
\mathrm{E}_t\phi_3(X_{t+1}) &= X_t^2 * (1+\alpha dt)^2 + \sigma^2 dt \\
\mathrm{E}_t\phi_4(X_{t+1}) &= X_t^3 * (1+\alpha dt)^3 + X_t^3 * 3 * (1+\alpha dt) * \sigma^2 dt
\end{aligned} \tag{24}$$

Results

Pathwise approach produces better lower and upper bounds than the benchmark methods. We see that as the number of exercise rights L increases, the improvement produced by the pathwise method is not much. This might be due to that the increase in inaccuracies due to the stopping time approximations. Also, our numerical experiments worked on a very few number of paths ($S = 3000$) as compared to the experiments on Desai et al. (They use around $30,000$ sample paths). Hence, these are just the initial tests of the approach. One needs to experiment with different basis functions and more number of samples for better estimation. I believe, with increasing number of monte carlo paths, one might get better bounds consistently even for large values of L. But one might incur heavy cost of computation i.e. because of the setting of the LP at each step.

Exercise Rights	LS-LB	PO-LB	PO-UB	Schoenmakers (UB)
1	4.40	4.47	4.74	4.85
2	8.81	8.98	9.46	9.81
3	13.20	13.44	14.16	14.66
4	17.61	17.91	18.84	19.47
5	21.94	22.32	23.50	24.21
6	26.32	26.73	28.12	28.91
7	30.69	31.13	32.73	33.56
8	35.00	35.45	37.32	38.17
9	39.27	39.75	41.88	42.75
10	43.59	44.02	46.43	47.30
12	52.05	52.58	55.48	56.25
15	64.77	65.19	68.86	69.42
18	77.24	77.62	81.03	82.27
21	89.67	89.95	94.01	94.80
24	101.98	102.07	106.79	107.03
25	106.03	106.05	111.00	111.02

Table 1: Comparing the Primal/Dual Bounds of Pathwise Approach with Longstaff-Schwartz (LB)/Schoenmakers(DB)

6.2 Model: Auto Regressive Process

We now present the results of a simple numerical experiment where we compare the performance of the Meinshausen and Hambly [19] dual bound with the pure martingale dual bound of Schoenmakers [22]. We consider the problem setup of Meinshausen-Hambly where the underlying security price follows an AR(1) type model. In particular, the underlying price, S_t, on day t satisfies

$$\log S_{t+1} \;=\; (1-k)\,(\log S_t - \mu) \;+\; \mu \;+\; \sigma W_t$$

where $W_t \sim \mathrm{N}(0,1)$ and the model parameters are set to $\sigma = 0.5$, $k = 0.9$, $\mu = 0$ and $S_1 = 1$. The strike price, K, for each option was set to zero and we assumed that each swing option had $T = 1000$ exercise opportunities with the number of exercise rights varying from $L = 1$ to $L = 10$. Lower bounds can be computed using a simple generalization of the cross-path regression approach of Longstaff and Schwartz [18] and Tsitsiklis and Van Roy [24]. Details of this generalization are given in Meinshausen and Hambly [19].

Basis Functions

We use the following 4 basis functions for the computation:

$$\phi_1(x) = 1 \quad \phi_2(x) = x \quad \phi_3(x) = x^2 \quad \phi_4(x) = x^3 \tag{25}$$

These are the first 4 linear polynomials of x, We discuss the nice property of these basis functions for Gaussian models like GBM and AR(1) discussed earlier. Recall that the value function approximations, are the linear combinations of these basis functions. Since the past work of Haugh

and Kogan [15], Andersen and Broadie [2],Meinshausen and Hambly [19], Benders [4] is based on approximation to the continuation values, the inner expectations were approximated once again by sampling. Inner Expectations arise in the penalty term in the form $(E_t(V_{t+1}))$. But in the current work, we approximate the value function directly instead of the continuation values. Hence, it is possible to obtain nice analytic expressions for the inner expectations and avoid sampling. This way, we would have the exact value rather than close estimated values which should improve accuracy and speed up the computation drastically.

$$\begin{aligned} \mathrm{E}_t\phi_1(X_{t+1}) &= 1 \\ \mathrm{E}_t\phi_2(X_{t+1}) &= c_1 * exp(0.1 * log(X_t)) \\ \mathrm{E}_t\phi_3(X_{t+1}) &= c_2 * exp(0.2 * log(X_t)) \\ \mathrm{E}_t\phi_4(X_{t+1}) &= c_3 * exp(0.3 * log(X_t)) \end{aligned} \tag{26}$$

where $c_1 = \exp\{\frac{0.5^2}{2}\}$, $c_2 = \exp\{\frac{1}{2}\}$, $c_3 = \exp\{\frac{1.5^2}{2}\}$

Results

For this model, where the underlying security process is auto regressive, pathwise approach produces lower and upper bounds very close to the existing methods. Existing methods seems to outperform the pathwise method by a small accuracy. Nevertheless, this approach can be seen as a good alternative if there are setting in which the existing methods fail to produce good confidence bounds.

Exercise Rights	LS-LB	PO-LB	PO-UB	Schoenmakers (UB)
1	3.27	3.24	3.37	3.27
2	6.10	6.06	6.26	6.13
3	8.71	8.66	8.89	8.74
4	11.15	11.09	11.36	11.19
5	13.46	13.37	13.70	13.50
6	15.66	15.57	15.95	15.71
7	17.77	17.67	18.20	17.83
8	19.80	19.70	20.36	19.87
9	21.80	21.68	22.40	21.88
10	23.72	23.59	24.50	23.85

Table 2: Comparing the Primal/Dual Bounds of Pathwise Approach with Longstaff-Schwartz (LB)/Schoenmakers(DB)

7 Theory

We conclude our work with a guarantee on the upper bound generated by the pathwise technique.

Theorem 1 *For any policy $\pi = \{\tau^1, \tau^2, , ...\tau^{n-1}\}$ and value function V, the bias at time t with n exercise opportunities remaining $B_t(n, V)$ of the estimated marginal value denoted by*

$$B_t(n, V) := F_t \Delta V^{\tau,n} - \Delta V^{*,n} \tag{27}$$

is bounded above by

$$B_0^{\pi}(n, V) \leq D_{\pi}^{n-1} + 2\sqrt{(4D_V^2 + \sigma_k^2)T} \tag{28}$$

where

$$\begin{aligned} D_V &:= \sup_{m=1,2,..n, x \in S, 0 \leq t \leq T} |\Delta V_t^{*,m}(x) - \Delta V_t^m(x)| && (29) \\ D_{\pi}^n &:= V_0^{*,n} - V_0^{\pi,n} && (30) \\ \sigma_k^2 &:= \sup_{m=1,2,..n, x \in S, 0 \leq t \leq T} E_{t-1}\left[(E_{t-1}^m - E_{t-1}[\Delta V_t^m])^2\right] && (31) \end{aligned}$$

Proof: Refer to Meinshausen and Hambly [19]

The upper bound on the problem with n exercise rights as obtained by the pathwise optimization technique, is given by,

$$F_0 \Phi r_{PO}(x) := \sum_{i=1}^{n} F_0 \Phi r_{PO}^{(i)}(x) \tag{32}$$

Theorem 2 *Suppose that r_{PO} is an optimal solution of (20). Then,*

$$||F_0 \Phi r_{PO}(x) - V^*||_1 \leq \min_r \sum_{i=1}^{n} B_t(i, \Phi r) \tag{33}$$

where π is a policy that defines the stopping times. In our method, π is the greedy policy.

Proof:

$$\begin{aligned} ||F_0 \Phi r_{PO}(x) - V^*||_1 &:= ||\sum_{i=1}^{n} (F_0 \Phi r_{PO}^{(i)}(x) - \Delta V^{*,i}(x))||_1 && (34) \\ &\leq \sum_{i=1}^{n} ||(F_0 \Phi r_{PO}^{(i)}(x) - \Delta V^{*,i}(x))||_1 && (35) \\ &\leq \sum_{i=1}^{n} ||(F_0 \Phi r^{(i)}(x) - \Delta V^{*,i}(x))||_1 && (36) \\ &:= \sum_{i=1}^{n} B_t(i, \Phi r) && (37) \end{aligned}$$

Minimizing over all possible r establishes the theorem ∎

8 Conclusions

We generalize the pathwise optimization approach of Desai et al. to solve high dimensional multiple stopping problems. We show that this method outperforms existing methods when the underlying process is a Geometric Brownian Motion (GBM) and produces close confidence bounds when the underlying process is Auto Regressive.

Since, the pathwise method works on value function approximations, we can avoid inner sampling in settings when the inner expectations can be found analytically. But this is not the case in earlier methods since, they work on continuation value approximations and inner sampling becomes necessary.

The downside of the pathwise approach is that it solves a lot of LP's for a given instance of multiple stopping. Though the state of art solvers (For example, Gurobi which we use here) is extremely fast in solving the LP, it takes a lot of time to set up these LP's since one needs to adjust the penalties and find approximate stopping times between successive computation of marginal bound. This might make the pathwise method less attractive. Comparatively, the pathwise technique might need lot more sample paths for accurate estimation than the existing methods.

Acknowledgments I thank the discussions with Prof. Martin Haugh. The idea for this work came from our meetings while working on our earlier paper. I am also thankful to Prof Garud Iyengar and Rodrigo Carrasco who have been kind enough to offer their computational facility which helped running the large scale problem instances.

References

[1] N. Aleksandrov and BM Hambly. A dual approach to multiple exercise option problems under constraints. *Mathematical Methods of Operations Research*, 71(3):503–533, 2010.

[2] L. Andersen and M. Broadie. Primal-dual simulation algorithm for pricing multidimensional american options. *Management Science*, pages 1222–1234, 2004.

[3] S. Balder, A. Mahayni, and J. Schoenmakers. Primal–dual linear monte carlo algorithm for multiple stopping–an application to flexible caps. Technical report, WIAS preprint 1666, 2011.

[4] C. Bender. Dual pricing of multi-exercise options under volume constraints. *Finance and Stochastics*, pages 1–26, 2008.

[5] C. Bender. Primal and dual pricing of multiple exercise options in continuous time. *SIAM Journal of Financial Mathematics*, 2:562–586, 2011.

[6] C. Bender, J. Schoenmakers, and J. Zhang. Dual representations for general multiple stopping problems. Technical report, November, 2011.

[7] D.B. Brown, J.E. Smith, and P. Sun. Information relaxations and duality in stochastic dynamic programs. *Operations research*, 58(4):785–801, 2010.

[8] R. Carmona and N. Touzi. Optimal multiple stopping and valuation of swing options. *Mathematical Finance*, 18(2):239–268, 2008.

[9] Shyam Chandramouli and M. Haugh. A unified approach to multiple stopping and duality. 2011.

[10] N. Chen and P. Glasserman. Additive and multiplicative duals for american option pricing. *Finance and Stochastics*, 11(2):153–179, 2007.

[11] V.V. Desai, V.F. Farias, and C.C. Moallemi. Pathwise optimization for optimal stopping problems. Technical report, MIT, Sloan School of Management, 2010.

[12] Bertsekas D.P. Dynamic programming and optimal control. 1995.

[13] L. Gergely Gyurko, B. Hambly, and J. Hendrik Witte. Monte carlo methods via a dual approach for some discrete time stochastic control problems. 2011.

[14] P. Glasserman. *Monte Carlo methods in financial engineering*, volume 53. Springer verlag, 2004.

[15] M.B. Haugh and L. Kogan. Pricing american options: a duality approach. *Operations Research*, pages 258–270, 2004.

[16] P. Jaillet, E.I. Ronn, and S. Tompaidis. Valuation of commodity-based swing options. *Management science*, pages 909–921, 2004.

[17] J. Keppo. Pricing of electricity swing options. *Journal of Derivatives*, 11:26–43, 2004.

[18] F.A. Longstaff and E.S. Schwartz. Valuing american options by simulation: A simple least-squares approach. *Review of Financial Studies*, 14(1):113, 2001.

[19] N. Meinshausen and B.M. Hambly. Monte carlo methods for the valuation of multiple-exercise options. *Mathematical Finance*, 14(4):557–583, 2004.

[20] J. Neveu. *Discrete-parameter martingales*, volume 10. Elsevier, 1975.

[21] L.C.G. Rogers. Monte carlo valuation of american options. *Mathematical Finance*, 12(3):271–286, 2002.

[22] J. Schoenmakers. A pure martingale dual for multiple stopping. *Finance and Stochastics*, pages 1–16, 2010.

[23] H. Topaloglu. Computation and dynamic programming. 2010.

[24] J.N. Tsitsiklis and B. Van Roy. Regression methods for pricing complex american-style options. *Neural Networks, IEEE Transactions on*, 12(4):694–703, 2001.