6. Online Companion

6.1. Pricing the Bermudan max call

In this section of our online companion we prove that the Bermudan max call, a popular test benchmark for OS algorithms (Andersen and Broadie (2004)), satisfies Assumptions 2 and 3 (for appropriate parameters). Therefore, the classical problem of pricing the Bermudan max call is an instance within our theoretical framework, and thus enjoys the guarantee of Theorem 4.

6.1.1. Basic setup Following Andersen (1999), we assume that the underlying stochastic asset process evolves according to the dynamics of a multi-dimensional geometric B.m., i.e. $dG^k_t = (r - \varrho)G^k_t dt + \sigma G^k_t dW^k_t$, for all $t \geq 0$ and $k = 1, ..., D$, with $G^k_0 = y_0$ for all k, where $\{W^k_t, t \geq 0, k = 1, ..., D\}$ are independent one-dimensional B.m. and r, ϱ, σ are constants capturing the interest rate, the dividend rate and the volatility respectively. We assume that the option lives in $[0, \mathcal{J}]$ for some horizon length $\mathcal{J}$,and that the exercising dates are restricted to $\mathcal{S} = \{t_0, ..., t_T\}$ where $t_i \triangleq \frac{i}{T}\mathcal{J}$ for $i \in \{0, \ldots, T\}$ (assumed to start from 0 in line with standard notation in this setting). We also assume that $\mathcal{J}, T$ are both positive integers, and T is divisible by $\mathcal{J}$. One should think of $\mathcal{J}$ as, for example, the number of months, and $\frac{T}{\mathcal{J}}$ the number of exercising opportunities within one month. The payoff is given by $g_t(G_t) = \big(\max(G^1_t, ..., G^D_t) - \kappa\big)^+$, with $\kappa > 0$ some positive constant and $x^+ = \max(x, 0)$ for any scalar x. The OS problem associated with Bermudan max call pricing is thus

$$\text{OPT} = \sup_{\tau \in \{0, ..., T\}} E\Big[e^{-rt_\tau} g_{t_\tau}(G_{t_\tau})\Big],$$

with τ a stopping time adapted to the natural filtration $\mathcal{F}$ generated by the asset process. Since here the information process is simply a D-dimensional geometric B.m. discretized on the mesh $\{t_i, i = 0, \ldots, T\}$, the existence of an efficient simulator (as required by Assumption 2) is easily verified.

6.1.2. Main result Our main result for Bermudan max call is summarized in the next corollary. Let $\Phi(\cdot)$ denote the standard normal c.d.f., and $\Phi^c(\cdot) \triangleq 1 - \Phi(\cdot)$.

COROLLARY 5. *Suppose $\varrho > 1.5\sigma^2$. Then Algorithm $\mathcal{A}$ (with properly chosen parameters and an additional truncation step) implements a randomized stopping time τ_ϵ s.t. $E\Big[e^{-rt_{\tau_\epsilon}} g_{t_{\tau_\epsilon}}(G_{t_{\tau_\epsilon}})\Big] \geq (1 - \epsilon)\text{OPT} - \epsilon\kappa$. In addition, at each time step, the decision of whether to stop (if one has not yet stopped) requires total computational time at most $(10TD + 4) \times \exp\Bigg(1.6 \times 10^9 \times \gamma_0^8 \times \epsilon^{-6}\Bigg) \times T^{3\times 10^4 \times \gamma_0^4 \times \epsilon^{-3}}$, where $\gamma_0 = 1 + \Bigg(\Phi^c\big(\frac{\varrho}{\sigma} + \frac{\sigma}{2}\big)\Bigg)^{-2} \big(\frac{e}{e-1}\big)^2 \big(1 + \frac{4\sigma^2 + \frac{4\sigma^4}{(\varrho + \sigma^2/2)^2}}{\rho - 1.5\sigma^2}\big)$.*

Let us briefly comment on the assumptions of Corollary 5. The requirement that $T, \mathcal{J}$ are integers with T divisible by $\mathcal{J}$ is a technical convenience, and could very likely be removed through a more involved analysis. Regarding the assumption $\varrho > 1.5\sigma^2$, let us point out that previous studies using Bermudan max call as their numerical benchmark (e.g. Andersen and Broadie (2004), Becker et al. (2021), Rogers (2015), Belomestny (2011)) all have parameter choices satisfying this condition. Indeed, a common default setting for these parameters in numerical experiments from the literature is $\sigma = 20\%$ and $\varrho = 0.1$, which satisfy our condition. Theoretically, the sign of $\varrho - 1.5\sigma^2$ determines whether the r.v. $\sup_{t\geq 0} e^{-rt}G^1_t$ has finite fourth moment, and thus our assumption can be interpreted as enforcing a finite fourth moment assumption on a critical quantity ensuring a strong stability of Monte-Carlo simulation approaches for this problem. Indeed, the definition of geometric B.m. (and parameters governing the drift and variance of G^i) implies that for $k \geq 1$, $E\big[\big(\sup_{t\geq 0} e^{-rt}G^1_t\big)^k\big] = E\big[\exp\Big(\sup_{t\geq 0}\big(-k(\varrho+\sigma^2/2)t + k\sigma W^1_t\big)\Big)\big]$. It is well known that the all-time supremum of a negatively drifted B.m is exponentially distributed, i.e. $\sup_{t\geq 0}\big(-k(\varrho+\sigma^2/2)t + k\sigma W^1_t\big) \sim \text{Exp}\left(\frac{2(\varrho+\sigma^2/2)}{k\sigma^2}\right)$. Combining, it follows from the basic properties of the moment generating function of an exponential distribution that for $k \geq 1$,

$$E\big[\big(\sup_{t\geq 0} e^{-rt}G^1_t\big)^k\big] < \infty \leftrightarrow \frac{2(\varrho+\sigma^2/2)}{k\sigma^2} > 1 \leftrightarrow \rho > \frac{k-1}{2}\sigma^2. \tag{14}$$

Hence $\rho > 1.5\sigma^2$ iff $E\big[\big(\sup_{t\geq 0} e^{-rt}G^1_t\big)^4\big] < \infty$.

We leave as an open question whether this assumption can be relaxed to e.g. a finite second moment assumption on the same quantity, which would translate into an assumption that $\varrho > .5\sigma^2$. Here a key challenge is the fact that our results require a uniform bound for the s.c.v. over all values of T and D, which (using the natural representation of geometric B.m. and relaxing the relevant suprema over discrete time points to their continuous-time analogues) seems to require such a uniform bound on the s.c.v. of $Z_{D,T} \triangleq \max_{i=1,\ldots,D} \exp\Big(\sup_{t\in[0,T]}\big(-(\varrho+\sigma^2/2)t + \sigma W^i_t\big)\Big)$. This question is complicated by at least three subtleties : (1) the s.c.v. does not enjoy some of the natural monotonicities that any given moment does in isolation (e.g. increasing T clearly increases any given moment of $Z_{D,T}$, but it is unclear whether it increases the s.c.v. of $Z_{D,T}$); (2) the tail of $Z_{D,T}$ behaves very differently depending on the relative scaling of D and T; and (3) although for any one value of T (even the choice $T = \infty$) one can use extreme value theory to (asymptotically) control the s.c.v. of $Z_{D,T}$, the subtle setting in which D and T may scale jointly seems to require a very fine-grained analysis beyond the scope of our results.

6.1.3. Proof of Corollary 5 Let us first state a technical lemma which follows from e.g. the results of Gordon (1941), Baricz (2008), and Gasull and Utzet (2014).

LEMMA 13. *For all $x > 0$, one has*

$$\frac{1}{\sqrt{2\pi}}\frac{x}{1+x^2}e^{-\frac{x^2}{2}} \le \Phi^c(x) \le \frac{1}{\sqrt{2\pi}}\frac{1}{x}e^{-\frac{x^2}{2}}.$$

Also, $\Phi^c(x) \le \sqrt{\frac{2}{\pi}}e^{-\frac{x^2}{2}}(x+\sqrt{x^2+\frac{8}{\pi}})^{-1}$. Furthermore, let $R(x,a) \triangleq \frac{\Phi^c(x+a)}{\Phi^c(x)}$. Then $R(x,a)$ is monotone decreasing in a for each fixed $x \in \mathbb{R}$, and monotone decreasing in x for each fixed $a \ge 0$.

As noted before, the main part of the proof of Corollary 5 is to verify Assumptions 2 and 3, especially Assumption 3 whose proof will be non-trivial. Let $\mathcal{M} \triangleq \max_{i=1,\ldots,D}\max_{t\in\{t_0,\ldots,t_T\}} e^{-rt}(G^i_t - \kappa)^+ + \kappa$.

LEMMA 14 **(Bermudan max call satisfies Assumption 3)**. *For any Bermudan max call instance satisfying $\varrho > 1.5\sigma^2$,*

$$\frac{E[\mathcal{M}^2]}{E[\mathcal{M}]^2} \le 1 + \left(\Phi^c\big(\frac{\varrho}{\sigma}+\frac{\sigma}{2}\big)\right)^{-2}\left(\frac{e}{e-1}\right)^2\big(1+\frac{4\sigma^2+\frac{4\sigma^4}{(\varrho+\sigma^2/2)^2}}{\rho-1.5\sigma^2}\big).$$

Proof Note that $e^{-rt}(G^i_t-\kappa)^+ + \kappa \ge e^{-rt}(G^i_t-\kappa) + e^{-rt}\kappa = e^{-rt}G^i_t$. It follows that

$$\begin{aligned}
\frac{E[\mathcal{M}^2]}{E[\mathcal{M}]^2} &= 1 + \frac{\text{Var}(\mathcal{M})}{E[\mathcal{M}]^2}, &(15)\\
&\le 1 + \frac{\text{Var}\left(\max_{i=1,\ldots,D}\max_{t\in\{t_0,\ldots,t_T\}} e^{-rt}(G^i_t-\kappa)^+ + \kappa\right)}{E\left[\max_{i=1,\ldots,D}\max_{t\in\{t_0,\ldots,t_T\}} e^{-rt}G^i_t\right]^2}\\
&= 1 + \frac{\text{Var}\left(\max_{i=1,\ldots,D}\max_{t\in\{t_0,\ldots,t_T\}} e^{-rt}(G^i_t-\kappa)^+\right)}{E\left[\max_{i=1,\ldots,D}\max_{t\in\{t_0,\ldots,t_T\}} e^{-rt}G^i_t\right]^2}\\
&\le 1 + \frac{E\left[\left(\max_{i=1,\ldots,D}\max_{t\in\{t_0,\ldots,t_T\}} e^{-rt}G^i_t\right)^2\right]}{E\left[\max_{i=1,\ldots,D}\max_{t\in\{t_0,\ldots,t_T\}} e^{-rt}G^i_t\right]^2}. &(16)
\end{aligned}$$

Hence it suffices to upper bound (16). As (16) takes the same value for all values of y_0 (due to the multiplicative nature of geometric brownian motion), we may w.l.o.g. assume $y_0 = 1$. We now bound (16) (where the maximum is over a discrete set of times in $[0,\mathcal{J}]$) in terms of the analogous ratio in which the suprema are over all times in $[0,\mathcal{J}]$. Trivially, we may bound the numerator as follows:

$$E\bigg[\Big(\max_{i=1,...,D}\max_{t\in\{t_0,...,t_T\}} e^{-rt}G_t^i\Big)^2\bigg] \le E\bigg[\Big(\max_{i=1,...,D}\sup_{t\in[0,\mathcal{J}]} e^{-rt}G_t^i\Big)^2\bigg], \tag{17}$$

where we replace the grids $\{t_0,...,t_T\}$ by the interval $[0,\mathcal{J}]$. We now provide a similar lower bound on the denominator. Indeed, we show that

$$\frac{E\big[\max_{i=1,...,D}\max_{t\in\{t_0,...,t_T\}} e^{-rt}G_t^i\big]}{E\big[\max_{i=1,...,D}\sup_{t\in[0,\mathcal{J}]} e^{-rt}G_t^i\big]} \ \ge\ \Phi^c\big(\frac{\varrho}{\sigma}+\frac{\sigma}{2}\big). \tag{18}$$

We argue as follows. Let $\max_{t\in\{t_0,...,t_T\}} e^{-rt}G_t^i \triangleq M_i'$ and $\sup_{t\in[0,\mathcal{J}]} e^{-rt}G_t^i \triangleq M_i$. Then $\{(M_i',M_i), i=1,...,D\}$ are $i.i.d.$ random vectors, where w.p.1 $M_i\ge M_i'$ for all i. Here we note that M_i, M_i' have no connection to the constants M_1, M_2 used in the previous section. By definition, process $\{e^{-rt}G_t^i, t\in[0,\mathcal{J}]\}$ has the explicit form $\Big\{\exp\Big((-\varrho-\frac{\sigma^2}{2})t+\sigma W_t^i\Big), t\in[0,\mathcal{J}]\Big\}$ namely a geometric B.m. with negative drift. Here $\{W_t^i, t\in[0,\mathcal{J}]\}$ are independent standard B.m. To prove bound 18, we first show M_i' is near M_i in the following precise sense:

$$P(M_i'\ge z)\ge\Phi^c\Big(\frac{\varrho}{\sigma}+\frac{\sigma}{2}\Big)P(M_i\ge z) \text{ for all } z\ge 0 \text{ and all } i. \tag{19}$$

In particular, let $X^i=\{X_t^i, t\ge 0\}$ denote the stochastic process $\{\exp\big((-\varrho-\frac{\sigma^2}{2})t+\sigma W_t^i\big), t\ge 0\}$, where we note that $X_0^i=1$ w.p.1. Fix $z>1$ and define $\tau_z=\inf\{t\le\mathcal{J}: X_t^i=z\}$ (with $\tau_z=\mathcal{J}$ if no such time exists in $[0,\mathcal{J}]$). Recall that $\mathcal{S}=\{t_0,...,t_T\}$ where $t_i\triangleq\frac{i}{T}\mathcal{J}$ for $i\in\{0,\dots,T\}$. On $\{\tau_z<\mathcal{J}\}$, let $t^+=\min\{s\in\mathcal{S}:\ s\ge\tau_z\}$. Then $t^+-\tau_z\in(0,\frac{\mathcal{J}}{T}]\subseteq(0,1]$, and

$$\frac{X_{t^+}^i}{z}=\exp\big((-\varrho-\tfrac{\sigma^2}{2})(t^+-\tau_z)+\sigma(W_{t^+}^i-W_{\tau_z}^i)\big).$$

By the strong Markov property, $W_{t^+}^i-W_{\tau_z}^i\sim\mathcal{N}(0,t^+-\tau_z)$ independent of $\mathcal{F}_{\tau_z}$, where $N(a,b)$ denotes a normally distributed r.v. with mean a and variance b. Hence conditional on the event $\{\tau_z<\mathcal{J}\}$,

$$E\big[I(X_{t^+}^i\ge z)\,|\,\mathcal{F}_{\tau_z}\big]=P\big(\sigma N(0,t^+-\tau_z)>(\rho+\frac{\sigma^2}{2})(t^+-\tau_z)\big)=\Phi^c\Big((\tfrac{\varrho}{\sigma}+\tfrac{\sigma}{2})\sqrt{t^+-\tau_z}\Big)\ge\Phi^c\Big(\tfrac{\varrho}{\sigma}+\tfrac{\sigma}{2}\Big).$$

Noting that $P(\{\tau_z<\mathcal{J}\})=P(M_i\ge z)$, it follows that $P(M_i'\ge z)\ge\Phi^c\Big(\frac{\varrho}{\sigma}+\frac{\sigma}{2}\Big)\,P(M_i\ge z)$, which proves (19).

We next prove (18).

$$E\big[\max_{i=1,...,D}\max_{t\in\{t_0,...,t_T\}} e^{-rt}G_t^i\big]\ =\ E\big[\max_{i=1,...,D} M_i'\big]\ =\ \int_{z=0}^{\infty} E\Big[I\Big(\max_{i=1,...,D} M_i'\ge z\Big)\Big]dz$$

$$= \int_{z=0}^{\infty}\sum_{j=1}^{D} E\Big[I\Big(\max_{i=1,...,D} M_i' \geq z\Big) \times I\Big(M_1,...,M_{j-1} < z, M_j \geq z\Big)\Big]dz$$

since $\max_{i=1,...,D} M_i' \geq z \rightarrow \max_{i=1,...,D} M_i \geq z$, implying some M_j is the first to exceed z

$$\geq \int_{z=0}^{\infty}\sum_{j=1}^{D} E\Big[I\Big(M_j' \geq z\Big) \times I\Big(M_1,...,M_{j-1} < z, M_j \geq z\Big)\Big]dz$$

$$= \int_{z=0}^{\infty}\sum_{j=1}^{D} E\Big[I\Big(M_j' \geq z\Big) \times I\Big(M_j \geq z\Big)\Big] \times E\Big[I\Big(M_1,...,M_{j-1} < z\Big)\Big]dz$$

by independence

$$= \sum_{j=1}^{D}\int_{z=0}^{\infty} E\Big[I\Big(M_j' \geq z\Big) \times I\Big(M_j \geq z\Big)\Big] \times E\Big[I\Big(M_1,...,M_{j-1} < z\Big)\Big]dz$$

$$= \sum_{j=1}^{D}\int_{z=0}^{\infty} E\Big[I\Big(M_j' \geq z\Big)\Big] \times E\Big[I\Big(M_1,...,M_{j-1} < z\Big)\Big]dz$$

since $M_j' \geq z$ implies $M_j \geq z$

$$\geq \sum_{j=1}^{D}\int_{z=0}^{\infty} \Phi^c\Big(\frac{\varrho}{\sigma}+\frac{\sigma}{2}\Big)E\Big[I\Big(M_j \geq z\Big)\Big] \times E\Big[I\Big(M_1,...,M_{j-1} < z\Big)\Big]dz \qquad \text{by (19)}$$

$$= \Phi^c\Big(\frac{\varrho}{\sigma}+\frac{\sigma}{2}\Big)\int_{z=0}^{\infty}\sum_{j=1}^{D} E\Big[I\Big(M_1,...,M_{j-1} < z, M_j \geq z\Big)\Big]dz$$

$$= \Phi^c\Big(\frac{\varrho}{\sigma}+\frac{\sigma}{2}\Big)\int_{z=0}^{\infty} E\Big[I\Big(\max_{1\leq i\leq D} M_i \geq z\Big)\Big]dz$$

$$= \Phi^c\Big(\frac{\varrho}{\sigma}+\frac{\sigma}{2}\Big)E[\max_{1\leq i\leq D} M_i].$$

Combining the above completes the proof of (18).

Recall again that $M_i = \sup_{t\in[0,\mathcal{J}]} e^{-rt}G_t^i = \sup_{t\in[0,\mathcal{J}]} \exp\big(-(\varrho+\sigma^2/2)t+\sigma W_t^i\big)$, and let us define $M^* \triangleq \max_{i=1,..,D} M_i$. Then combining (16) with both the upper bound (17) and the lower bound (18), we have that

$$\frac{E[\mathcal{M}^2]}{E[\mathcal{M}]^2} \leq 1 + \bigg(\Phi^c\Big(\frac{\varrho}{\sigma}+\frac{\sigma}{2}\Big)\bigg)^{-2}\frac{E[(M^*)^2]}{E[M^*]^2}. \tag{20}$$

In other words, it suffices to upper bound the s.c.v of $\max_{i=1,...,D}\sup_{t\in[0,\mathcal{J}]} e^{-rt}G_t^i$ (i.e. $\frac{E[(M^*)^2]}{E[M^*]^2}$) where the inner maximum is taken over the continuous interval $[0,\mathcal{J}]$ and which is easier to handle due to its continuous-time nature. We consider two cases (where the first case is essentially trivial).

Case I: $D=1$. In this case, we prove that $E[(M^*)^2]/E[M^*]^2 \le \frac{2\varrho+\sigma^2}{2\varrho-\sigma^2}$. Note that $E[M^*]^2 \ge 1$ trivially by considering the relevant supremum evaluated at $t=0$. We also have $E\left[\left(M^*\right)^2\right] = E\left[\sup_{t\in[0,\mathcal{J}]} \exp\left(-2(\varrho+\sigma^2/2)t + 2\sigma W_t\right)\right]$, which is upper bounded by $E\left[\sup_{t\in[0,\infty)} \exp\left(-2(\varrho+\sigma^2/2)t + 2\sigma W_t\right)\right]$. Combining with the the same argument used to prove (14), which solved for this expectation in closed form, we conclude that when $D=1$, $E[(M^*)^2]/E[M^*]^2 \le \frac{2\varrho+\sigma^2}{2\varrho-\sigma^2}$ as desired.

Case 2: $D \ge 2$. In this case, we prove that

$$\frac{E[(M^*)^2]}{E[M^*]^2} \le \left(\frac{e}{e-1}\right)^2 \left(1 + \frac{4\sigma^2 + \frac{4\sigma^4}{(\varrho+\sigma^2/2)^2}}{\rho - 1.5\sigma^2}\right). \tag{21}$$

Let $y^* \triangleq \sup\{y : P(M_i \ge y)\} = \frac{1}{D}$ be the $(1-\frac{1}{D})$-quantile of the distribution of each (identically distributed) M_i. By previous reasoning, $y^* \ge 1$. We next prove M^* is well concentrated around y^* in an appropriate multiplicative sense. More precisely, we prove that for all $\Delta \ge 0$,

$$P(M^* \ge y^* e^{\Delta}) \;\le\; 2\left(1 + \frac{\sigma^2}{(\varrho+\sigma^2/2)^2}\right) e^{-\frac{\varrho+\sigma^2/2}{\sigma^2}\Delta}; \tag{22}$$

$$P(M^* \ge y^*) \;\ge\; 1 - \frac{1}{e}. \tag{23}$$

We first prove (23). It follows from straightforward reasoning about quantiles and the maximum of i.i.d. r.v.s that $P(M^* \ge y^*) = 1 - \left(1 - P(M_1 > y^*)\right)^D = 1 - (1-\frac{1}{D})^D$. Combining with a straightforward calculus exercise then completes the proof of (23).

We now prove (22). As M_i is the maximum of a geometric B.m. with drift, its c.d.f. has an explicit formula (see e.g. Boukai (1990)). Namely, for all i and $y > 0$,

$$P(M_i > y) = e^{-2\frac{\varrho+\sigma^2/2}{\sigma^2}\log y}\Phi^c\left(\frac{\log y - (\varrho+\sigma^2/2)\mathcal{J}}{\sigma\sqrt{\mathcal{J}}}\right) + \Phi^c\left(\frac{\log y + (\varrho+\sigma^2/2)\mathcal{J}}{\sigma\sqrt{\mathcal{J}}}\right). \tag{24}$$

Let us define $u^+ \triangleq \frac{\log y^* + (\varrho+\sigma^2/2)\mathcal{J}}{\sigma\sqrt{\mathcal{J}}}$, $u^- \triangleq \frac{\log y^* - (\varrho+\sigma^2/2)\mathcal{J}}{\sigma\sqrt{\mathcal{J}}}$, and $\nu \triangleq \frac{\varrho+\sigma^2/2}{\sigma^2}$. Then (24) is easily seen to imply that

$$\begin{aligned}
\frac{P(M_i > y^* e^{\Delta})}{P(M_i > y^*)} &= \frac{e^{-2\nu(\log y^* + \Delta)}\Phi^c\left(u^- + \frac{\Delta}{\sigma\sqrt{\mathcal{J}}}\right) + \Phi^c\left(u^+ + \frac{\Delta}{\sigma\sqrt{\mathcal{J}}}\right)}{e^{-2\nu\log y^*}\Phi^c\left(u^-\right) + \Phi^c\left(u^+\right)} \\
&\le \frac{e^{-2\nu(\log y^* + \Delta)}\Phi^c\left(u^- + \frac{\Delta}{\sigma\sqrt{\mathcal{J}}}\right)}{e^{-2\nu\log y^*}\Phi^c\left(u^-\right)} \qquad (25)\\
&\quad + \frac{\Phi^c\left(u^+ + \frac{\Delta}{\sigma\sqrt{\mathcal{J}}}\right)}{\Phi^c\left(u^+\right)}. \qquad (26)
\end{aligned}$$

We now bound (25) and (26) separately, beginning with (25). By monotonicity of $\Phi^c(\cdot)$, (25) is at most

$$\frac{e^{-2\nu(\log y^*+\Delta)}\Phi^c\left(u^-\right)}{e^{-2\nu\log y^*}\Phi^c\left(u^-\right)} = e^{-2\nu\Delta}.$$

We now bound (26), and begin by proving a bound for $\frac{\Phi^c(u+s)}{\Phi^c(u)}$ for general $u, s \geq 0$. In particular, by Lemma 13, for all $u \geq 0, s \geq 0$,

$$\begin{aligned}
\frac{\Phi^c(u+s)}{\Phi^c(u)} &\leq \frac{\frac{1}{u+s}\exp\left(-\frac{1}{2}(u+s)^2\right)}{\frac{u}{1+u^2}\exp(-\frac{1}{2}u^2)} \\
&= \frac{1+u^2}{u(u+s)}\exp\left(-us-\frac{1}{2}s^2\right) \\
&\leq \frac{1+u^2}{u^2}\exp(-us) = (1+\frac{1}{u^2})\exp(-us).
\end{aligned}$$

Plugging in u^+, and using $u^+ = \frac{\log y^*+(\varrho+\sigma^2/2)\mathcal{J}}{\sigma\sqrt{\mathcal{J}}} \geq \frac{(\varrho+\sigma^2/2)\sqrt{\mathcal{J}}}{\sigma}$ (since $\log y^* \geq 0$), we find that (26) is at most

$$\begin{aligned}
&\left(1+\frac{1}{(u^+)^2}\right)e^{-u^+\frac{\Delta}{\sigma\sqrt{\mathcal{J}}}} \\
\leq &\left(1+\frac{\sigma^2}{(\varrho+\sigma^2/2)^2\mathcal{J}}\right)e^{-u^+\frac{\Delta}{\sigma\sqrt{\mathcal{J}}}} \\
\leq &\left(1+\frac{\sigma^2}{(\varrho+\sigma^2/2)^2}\right)e^{-\frac{(\varrho+\sigma^2/2)\sqrt{\mathcal{J}}}{\sigma}\times\frac{\Delta}{\sigma\sqrt{\mathcal{J}}}} \\
= &\left(1+\frac{\sigma^2}{(\varrho+\sigma^2/2)^2}\right)e^{-\nu\Delta}.
\end{aligned}$$

Combining our bounds for (25) and (26), we conclude that

$$\frac{P(M_i > y^*e^{\Delta})}{P(M_i > y^*)} \leq 2\left(1+\frac{\sigma^2}{(\varrho+\sigma^2/2)^2}\right)e^{-\nu\Delta}. \tag{27}$$

We now complete the proof of (22). By a straightforward union bound, $P(M^* \geq y^*e^{\Delta}) \leq D \times P(M_1 \geq y^*e^{\Delta})$. Combining with (27) and the fact that by definition $P(M_1 > y^*) = \frac{1}{D}$ then completes the proof of (22).

With (22) and (23) in hand, we are able to bound $E[M^*]$ and $E[(M^*)^2]$ and complete the proof of (21). On the one hand, we have $E[M^*] \geq E[y^*I(M^* \geq y^*)] \geq \frac{e-1}{e}y^*$ by (23). On the other hand, by (22),

$$\begin{aligned}
E[(M^*)^2] &= \int_0^{\infty} P\Big((M^*)^2 > t\Big)\mathrm{d}t \\
&= \int_0^{(y^*)^2} P\Big((M^*)^2 > t\Big)\mathrm{d}t + \int_{(y^*)^2}^{\infty} P\Big((M^*)^2 > t\Big)\mathrm{d}t
\end{aligned}$$

$$
\begin{aligned}
&\le (y^*)^2 + \int_{(y^*)^2}^{\infty} P\left(M^* > e^{\log(\sqrt{\frac{t}{(y^*)^2}})} y^*\right) \mathrm{d}t \\
&\le (y^*)^2 + 2\left(1 + \frac{\sigma^2}{(\varrho+\sigma^2/2)^2}\right) \int_{(y^*)^2}^{\infty} e^{-\frac{\varrho+\sigma^2/2}{\sigma^2}\log(\sqrt{\frac{t}{(y^*)^2}})} dt \\
&= (y^*)^2 + 2\left(1 + \frac{\sigma^2}{(\varrho+\sigma^2/2)^2}\right) (y^*)^{\frac{\varrho+\sigma^2/2}{\sigma^2}} \int_{(y^*)^2}^{\infty} t^{-\frac{1}{2}\frac{\varrho+\sigma^2/2}{\sigma^2}} dt \\
&= (y^*)^2 + 2\left(1 + \frac{\sigma^2}{(\varrho+\sigma^2/2)^2}\right) (y^*)^{\frac{\varrho+\sigma^2/2}{\sigma^2}} \frac{(y^*)^{2-\frac{\varrho+\sigma^2/2}{\sigma^2}}}{\frac{1}{2}\frac{\varrho+\sigma^2/2}{\sigma^2} - 1} \\
&= (y^*)^2 + 2(y^*)^2 \frac{2\sigma^2 + \frac{2\sigma^4}{(\varrho+\sigma^2/2)^2}}{\rho - 1.5\sigma^2} \\
&= (y^*)^2 \big(1 + \frac{4\sigma^2 + \frac{4\sigma^4}{(\varrho+\sigma^2/2)^2}}{\rho - 1.5\sigma^2}\big)
\end{aligned}
$$

Combining the above completes the proof of (21).

We now complete the proof of Lemma 14. Combining the two cases $D = 1$ and $D \ge 2$, with the fact that (1) $\frac{2\rho+\sigma^2}{2\rho-\sigma^2} = 1 + \frac{2\sigma^2}{2\rho-\sigma^2}$, (2) $\rho > 1.5\sigma^2$ implies $\frac{2\sigma^2}{2\rho-\sigma^2} \le 1$, and (3) $(\frac{e}{e-1})^2 \ge 2$, we conclude that for all $D \ge 1$,

$$
\frac{E[(M^*)^2]}{E[M^*]^2} \le \left(\frac{e}{e-1}\right)^2 \big(1 + \frac{4\sigma^2 + \frac{4\sigma^4}{(\varrho+\sigma^2/2)^2}}{\rho - 1.5\sigma^2}\big).
$$

Combining with (20) then completes the proof. $Q.E.D.$

LEMMA 15 **(Bermudan max call satisfies Assumption 2)**. *For any Bermudan max call instance, the runtime to generate one (conditioned) trajectory and to compute the associated stopping rewards along the trajectory is bounded by* $10TD$ *(for all* t *and* $\gamma_{[t]}$*), under our computational model. Namely,* $C + G \le 10TD$*.*

Proof $Y_{[T]}$ is distributed as D independent discretized geometric B.m. Thus conditional on any t and trajectory $\gamma_{[t]}$, the remainder of the trajectory will be (conditionally) distributed as D independent geometric B.m. (sampled discretely at times $t+1, \ldots, T$), initialized at the values consistent with $\gamma_{[t]}$ (i.e. the conditioned sample path at time t and considering the Markovian structure of geometric B.m.). Thus these trajectories can be simulated in a straightfoward iterative manner by generating $(T-t) \times D$ independent standard normals, and iteratively scaling, exponentiating, and multiplying with the D respective processes to construct the remaining trajectories. As the reward functions are simple functions of the maximum of these processes (at any given time), they too can be evaluated (for any given $Y_{[T]}$) efficiently in our computational model. It may be easily verified that these straightforward steps can be implemented in time bounded by $10TD$ in our computational model, and we omit the details. Q.E.D.

With Lemmas 14 and 15 proven, we complete the proof of Corollary 5.

Proof of Corollary 5 By applying a constant shift κ to the stopping reward, namely $\max_{i=1,\ldots,D} e^{-rt}(G_t^i-\kappa)^+ + \kappa$, the optimal value increases by κ and the optimal stopping time remains unchanged. It is thus equivalent for us to solve a stopping problem associated with this new payoff. Let $\text{OPT}^s = \text{OPT} + \kappa$ denote the optimal value under the new payoff. The squared CV of the path-wise maximum of this new payoff is $\frac{E[M^2]}{E[M]^2}$ where $M = \max_{t\in\{t_0,\ldots,t_T\}} \max_{i=1,\ldots,D} e^{-rt}(G_t^i - \kappa)^+ + \kappa$. By Lemma 14, it is bounded by a constant $\gamma_0 = 1 + \left(\Phi^c\left(\frac{\varrho}{\sigma}+\frac{\sigma}{2}\right)\right)^{-2}\left(\frac{e}{e-1}\right)^2\left(1+\frac{4\sigma^2+\frac{4\sigma^4}{(\varrho+\sigma^2/2)^2}}{\rho-1.5\sigma^2}\right)$ satisfying Assumption 3. By Lemma 15, the unit sampling and evaluation cost is bounded by $C+G\le 10TD$, satisfying Assumption 2. Recall that here $Z_t = \max_{i=1,\ldots,D} e^{-rt}(G_t^i - \kappa)^+$ and $\text{OPT} = \sup_{\tau\in\mathcal{T}} E[Z_\tau]$. By Theorem 4, Algorithm $\mathcal{A}$ (with properly chosen parameters and an additional truncation step) implements a randomized stopping time τ_ϵ s.t. $\frac{E[Z_{\tau_\epsilon}]+\kappa}{\text{OPT}^s} \ge 1-\epsilon$, implying $E[Z_{\tau_\epsilon}]+\kappa \ge (1-\epsilon)(\text{OPT}+\kappa)$, itself implying $E[Z_{\tau_\epsilon}] \ge (1-\epsilon)\text{OPT} - \epsilon\kappa$. Combining the above completes the proof. $Q.E.D.$

6.2. Numerical Experiments

In this section, we test our theory numerically on simple problems of stopping *i.i.d.* r.v.s., largely to illustrate the practical limitations of our approach. As outside of this setting we are not easily able to simulate even $\mathbf{Z}^4$ (not to mention $\mathbf{Z}^k$ for $k = O(\frac{1}{\epsilon})$) on a standard laptop (in spite of its theoretical polynomial complexity), we do not pursue an in-depth numerical exploration of our algorithm's performance on the complex numerical options pricing benchmarks from e.g. Becker et al. (2021), as in general those authors are interested in obtaining highly accurate solutions. For the simple i.i.d. problems we consider, we can benchmark the performance of our expansion against the true optimal stopping value, since the DP recursion is tractable.

More specifically, consider the problem $\sup_{\tau\in\mathcal{T}} E[Z_\tau]$ where $Z_1,\ldots,Z_T$ are *i.i.d.* r.v.s distributed according to F, and the filtration is the natural one generated by $\{Z_t, t=1,\ldots,T\}$. We experiment with two types of distributions: $F = \mathsf{U}[0,1]$ (i.e. uniformly distributed on $[0,1]$) and $F = \mathsf{Exp}(1)$ (i.e. exponentially distributed with mean 1). Note that the choice of exponential distribution is beyond the scope of the approximation theory established in Theorem 2 under Assumption 1 (as the reward is unbounded), but falls within the scope of our theory for the setting of unbounded rewards (i.e. Assumption 3) since the coefficient of variation of the maximum of i.i.d. exponentially distributed r.v.s is well-known to be uniformly bounded (independent of T). As our expansion for OPT is at the core of our theory, we focus on numerically evaluating the accuracy of our truncated expansion in this setting. We use Monte-Carlo simulation to estimate our expansion $E_k = \sum_{i=1}^{k} L_i$, where we recall $L_i = E[\max_{t=1}^{T} Z_t^i]$ and $Z_t^{i+1} = Z_t^i - E[\max_{j=1}^{T} Z_j^i|\mathcal{F}_t]$; $Z_t^1 = Z_t$ for $t = 1,\ldots,T$. For both types of distributions, we compute E_k for $k = 1,2,3,4$ using naive nested Monte Carlo simulation. To speed up the simulation, we leverage our knowledge of F, which directly yields the analytical form of $\mathbf{Z}^2$. More precisely, it follows from a straightforward calculation that when $F = \mathsf{U}[0,1]$, we have $Z_t^2 = Z_t - \frac{T-t+\left(\max_{j\le t} Z_j\right)^{T-t+1}}{T-t+1}$; when $F = \mathsf{Exp}(1)$, we have $Z_t^2 = Z_t - \max_{j\le t} Z_j - \sum_{j=1}^{T-t} \frac{T-t}{j}(-1)^{j+1} e^{-j\max_{j\le t} Z_j}$. Building on these expressions, we then compute the third and the fourth term in our expansion with nested Monte Carlo simulation, which is feasible on a standard laptop (as the recursive depth is reduced by one since we have an analytical formula for $\mathbf{Z}^2$). More precisely, we implement the following algorithm.

For each fixed T and each distribution, we repeat the following for 20 runs (to obtain 20 estimates, so we may evaluate the empirical standard deviation). In each run, we simulate each of L_1, L_2, L_3, L_4. More precisely, in each run, we simulate $N_1 = 50000$ sample paths, which we use in common across our simulations of L_1, L_2, L_3, L_4. First, we estimate L_1, L_2 by averaging over the exact evaluations (via the above closed-form expressions) of the appropriate maximum (of the elements of either $\mathbf{Z}^1$ or $\mathbf{Z}^2$) on all N_1 sample paths. We then use nested simulation to compute L_3, L_4. To compute L_3, we first uniformly at random select $N_3 = 1000$ out of the $N_1 = 50,000$ sample paths (without replacement), to reduce the computational burden by a factor of 50. On each of these N_3 sample paths, and for $t = 1,\ldots,T$, we perform conditional nested simulation, and draw $M_3 = 100$ independent sample paths conditioned on the partial history of the given path up to time t. We then approximate $Z_t^3 = Z_t^2 - E[\max_{j=1}^{T} Z_j^2|\mathcal{F}_t]$ using naive sample average approximation with these conditionally generated sample paths (combined with our closed-form expression for Z_t^2). We note that the total number of conditional sample paths generated equals $1000 \times 100 \times T$ (along each of the 1000 sample paths, for each of the T time periods, 100 conditional paths are generated). This computation gives us approximate values for $Z_t^3, t = 1,\ldots,T$ along each of the $N_3 = 1000$ sample paths. We then take the maximum of these Z_t^3 estimates along each of these N_3 sample paths and average over the paths to get an estimate for L_3. To estimate L_4, we set parameter $M_4 = 100$, and apply the same logic recursively (with one greater level of recursive depth). Namely, for each of the $N_3 = 1000$ sample paths, for each of the T time periods, for each of the 100 conditional paths P previously generated, we generate an additional 100 conditional paths at each time $s = 1,\ldots,T$ along each of these conditional paths P. We then approximate $Z_t^4 = Z_t^3 - E[\max_{j=1}^{T} Z_j^3|\mathcal{F}_t]$ using sample average approximation, completely analogous to how we computed estimates of Z_t^3. We note that here the total number of conditional sample paths generated equals $1000 \times 100^2 \times T^2$. Using the above methodology, we get an approximate value for each of L_1, L_2, L_3, L_4 (and hence E_1, E_2, E_3, E_4) in each of the 20 repetitions. Our final outputs are the sample average as well as the sample standard deviation computed from the 20 repetitions.

For both distributions we vary the horizon $T \in \{2, 3, 5, 10\}$ and report the difference between $E_k, k = 1, 2, 3, 4$ and the true optimal stopping value. We focus on modest values of T primarily for two reasons : (1) for larger values of T the i.i.d. problem in the uniform setting becomes somewhat degenerate, with the optimal value concentrating near 1, making comparisons less interesting; (2) the relevant nested conditional expectations are computationally easier for small T. On a 2023 Macbook Pro with 12 core CPU, it takes roughly 6 seconds to finish each of the 20 runs of the algorithm.

Table 1 (2) summarizes the results for $F = \mathsf{U}[0,1](\mathsf{Exp}(1))$. First, we observe that good approximations to the optimal value can be obtained from our theory using a number of samples far less than that proven in our Theorem 3 and Corollary 3, at least on well-behaved instances (i.e. our formal theoretical analyses are very conservative in some cases). However, it is also clear from our numerical findings that our expansion theory is not immediately practical. As shown in Tables 1 and 2, in many cases the performance gap when truncating after 4 terms remains larger than 1% even for simple OS instances of stopping *i.i.d.* sequences, suggesting that one would typically need nested conditional expectations of depth greater than four to reach a high-accuracy solution in general. However, for problems with relatively long horizon and/or high-dimension, evaluating nested conditional expectations with more than four levels of nesting can be computationally challenging ((Syed and Wang 2023)), as the complexity grows exponentially with the depth. Indeed, most existing heuristic methods focus on exploiting only unconditional sample paths, or at most one level of nesting conditional expectations, rather than leveraging deeper levels of nesting. We note that recently there has been considerable progress in understanding how techniques such as multi-level Monte Carlo can be used to speed up nested simulation ((Zhou et al. 2021)), and we leave the exploration of connections between our work and those approaches as a direction for future research.

In summary, these numerical experiments provide evidence that although our theoretical bounds may be very conservative, the immediate contribution of our work is theoretical rather than practical, especially in settings where high-accuracy solutions are desired. Our theoretical results established a worst-case guarantee on the depth of nesting required to obtain an ϵ-optimal solution for any fixed ϵ, and somewhat surprisingly established that this depth is independent of T. However, even the proven depth required (scaling as $O(\frac{1}{\epsilon})$) is impractical using current techniques. We leave the question of how to leverage these theoretical insights to design practical algorithms to future research.

		Performance Gap: (Mean (S.E.))			
Horizon	**OPT**	$E_1 - \textbf{OPT}$	$E_2 - \textbf{OPT}$	$E_3 - \textbf{OPT}$	$E_4 - \textbf{OPT}$
$T=2$	0.6250	+0.0416 (0.0010)	+0.0142 (0.0012)	+0.0067 (0.0013)	+0.0037 (0.0012)
$T=3$	0.6953	+0.0548 (0.0008)	+0.0217 (0.0010)	+0.0109 (0.0012)	+0.0064 (0.0013)
$T=5$	0.7751	+0.0583 (0.0015)	+0.0276 (0.0018)	+0.0159 (0.0023)	+0.0103 (0.0025)
$T=10$	0.8611	+0.0482 (0.0008)	+0.0272 (0.0011)	+0.0175 (0.0016)	+0.0122 (0.0019)

Table 1 Expansion performance: the Uniform case.

		Performance Gap: (Mean (S.E.))			
Horizon	**OPT**	$E_1 - \textbf{OPT}$	$E_2 - \textbf{OPT}$	$E_3 - \textbf{OPT}$	$E_4 - \textbf{OPT}$
$T=2$	1.3679	+0.1314 (0.0039)	+0.0290 (0.0041)	+0.0105 (0.0041)	+0.0044 (0.0040)
$T=3$	1.6225	+0.2108 (0.0036)	+0.0518 (0.0040)	+0.0192 (0.0039)	+0.0090 (0.0041)
$T=5$	1.9820	+0.3029 (0.0043)	+0.0858 (0.0048)	+0.0353 (0.0052)	+0.0185 (0.0055)
$T=10$	2.5290	+0.3984 (0.0037)	+0.1286 (0.0041)	+0.0558 (0.0055)	+0.0284 (0.0059)

Table 2 Expansion performance: the Exponential case.