

Optimal Algorithms for k -Search with Application in Option Pricing

Julian Lorenz¹, Konstantinos Panagiotou², and Angelika Steger

Institute of Theoretical Computer Science, ETH Zurich, 8092 Zurich, Switzerland
{jlorenz,panagiok,stegeer}@inf.ethz.ch

Abstract. In the k -search problem, a player is searching for the k highest (respectively, lowest) prices in a sequence, which is revealed to her sequentially. At each quotation, the player has to decide *immediately* whether to accept the price or not. Using the competitive ratio as a performance measure, we give optimal deterministic and randomized algorithms for both the maximization and minimization problems, and discover that the problems behave substantially different in the worst-case. As an application of our results, we use these algorithms to price “lookback options”, a particular class of financial derivatives. We derive bounds for the price of these securities under a no-arbitrage assumption, and compare this to classical option pricing.

1 Introduction

1.1 k -Search Problem

We consider the following online search problem: a player wants to sell (respectively, buy) $k \geq 1$ units of an asset with the goal of maximizing her profit (minimizing her cost). At time points $i = 1, \dots, n$, the player is presented a price quotation p_i , and must *immediately* decide whether or not to sell (buy) one unit of the asset for that price. The player is required to complete the transaction by some point in time n . We ensure that by assuming that if at time $n - j$ she has still j units left to sell (respectively, buy), she is compelled to do so in the remaining j periods. We shall refer to the profit maximization version (selling k units) as k -max-search, and to the cost minimization version (purchasing k units) as k -min-search.

In this work, we shall make no modeling assumptions on the price path except that it has finite support, which is known to the player. That is, the prices are chosen from the real interval $\mathcal{I} = \{x \mid m \leq x \leq M\}$, where $0 < m < M$. We define the *fluctuation ratio* $\varphi = M/m$. Let $\mathcal{P} = \bigcup_{n \geq k} \mathcal{I}^n$ be the set of all price sequences of length at least k . Moreover, the length of the sequence is known to the player at the beginning of the game.

¹ Partially supported by UBS AG

² Partially supported by the SNF, grant number: 200021-107880/1

Sleator and Tarjan [1] proposed to evaluate the performance of online algorithms by using *competitive analysis*. In this model, an online algorithm ALG is compared with an offline optimum algorithm OPT (which knows all prices in advance), on the same price sequence. Here, the price sequence is chosen by an adversary out of the set \mathcal{P} of admissible sequences. Let $\text{ALG}(\sigma)$ and $\text{OPT}(\sigma)$ denote the objective values of ALG and OPT when executed on $\sigma \in \mathcal{P}$. The *competitive ratio* of ALG is defined for maximization problems as

$$\text{CR}(\text{ALG}) = \max \left\{ \frac{\text{OPT}(\sigma)}{\text{ALG}(\sigma)} \mid \sigma \in \mathcal{P} \right\} ,$$

and similarly, for minimization problems

$$\text{CR}(\text{ALG}) = \max \left\{ \frac{\text{ALG}(\sigma)}{\text{OPT}(\sigma)} \mid \sigma \in \mathcal{P} \right\} .$$

We say that ALG is *c-competitive* if it achieves a competitive ratio not larger than c . For randomized algorithms, we substitute the *expected* objective value $\mathbb{E}[\text{ALG}]$ for ALG in the definitions above.

Related Work. In 2001, El-Yaniv, Fiat, Karp and Turpin studied, among other problems, the case $k = 1$, i.e. *1-max-search*, and the closely related *one-way trading* problem [2] with the competitive ratio (defined above) as performance measure. In the latter, a player wants to exchange some initial wealth to some other asset, and is again given price quotations one-by-one. However, the player may exchange an *arbitrary fraction* of her wealth for each price. Hence, the k -max-search problem for general $k \geq 1$ can be understood as a natural bridge between the two problems considered in [2], with $k \rightarrow \infty$ corresponding to the one-way trading problem. This connection will be made more explicit later.

Several variants of search problems, which will be discussed below, have been extensively studied in operations research and mathematical economics. However, traditionally most of the work follows a *Bayesian* approach: optimal algorithms are developed under the assumption that the prices are generated by a known distribution. Naturally, such algorithms heavily depend on the underlying model.

Lippmann and McCall [3, 4] give an excellent survey on search problems with various assumptions on the price process. More specifically, they study the problem of job and employee search and the economics of uncertainty, which are two classical applications of series search problems. In [5], Rosenfield and Shapiro study the situation where the price follows a random process, but some of its parameters may be random variables with known prior distribution. Hence, the work in [5] tries to get rid of the assumption of the Bayesian search models that the underlying price process is *fully* known to the player. Ajtai, Megiddo and Waarts [6] study the classical *secretary* problem. Here, n objects from an ordered set are presented in random order, and the player has to accept k of them so that the final decision about each object is made only on the basis of its rank relative to the ones already seen. They consider the problems of maximizing the probability of accepting the best k objects, or minimizing the

expected sum of the ranks (or powers of ranks) of the accepted objects. In this context, Kleinberg designed in [7] an $(1 - \mathcal{O}(1/\sqrt{k}))$ -competitive algorithm for the problem of maximizing the sum of the k chosen elements.

Results & Discussion. In contrast to the Bayesian approaches, El-Yaniv et al. [2] circumvent almost all distributional assumptions by resorting to competitive analysis and the minimal assumption of a known finite price interval. In this paper we also follow this approach. The goal is to provide a generic search strategy that works with any price evolution, rather than to retrench to a specific stochastic price process. In many applications, where it is not clear how the generating price process should be modeled, this provides an attractive alternative to classical Bayesian search models. In fact, in the second part of the paper we give an application of k -max-search and k -min-search to *robust option pricing* in finance, where relaxing typically made assumptions on the (stochastic) price evolution to the minimal assumption of a *price interval* yields remarkably good bounds.

Before we proceed with stating our results, let us introduce some notation. For $\sigma \in \mathcal{P}$, $\sigma = (p_1, \dots, p_n)$, let $p_{\max}(\sigma) = \max_{1 \leq i \leq n} p_i$ denote the maximum price, and $p_{\min}(\sigma) = \min_{1 \leq i \leq n} p_i$ the minimum price. Let W denote *Lambert's W -function*, i.e., the inverse of $f(w) = w \exp(w)$. For brevity we shall write $f(x) \sim g(x)$, if $\lim_{x \rightarrow \infty} f(x)/g(x) = 1$. It is well-known that $W(x) \sim \ln x$.

Our results for deterministic k -max-search are summarized in Theorem 1.

Theorem 1. *Let $k \in \mathbb{N}$, $\varphi > 1$. There is a r^* -competitive deterministic algorithm for k -max-search, where $r^* = r^*(k, \varphi)$ is the unique solution of*

$$\frac{\varphi - 1}{r^* - 1} = \left(1 + \frac{r^*}{k}\right)^k, \quad (1)$$

and there exists no deterministic algorithm with smaller competitive ratio. Furthermore, we have

- (i) $r^*(k, \varphi) \sim \sqrt[k+1]{k^k \varphi}$ for fixed $k \geq 1$ and $\varphi \rightarrow \infty$,
- (ii) $r^*(k, \varphi) \sim 1 + W\left(\frac{\varphi-1}{e}\right)$ for fixed $\varphi > 1$ and $k \rightarrow \infty$.

The algorithm in the theorem above is given explicitly in Section 2. Interestingly, the optimal competitive deterministic algorithm for the one-way trading problem studied in [2] has competitive ratio exactly $1 + W\left(\frac{\varphi-1}{e}\right)$ (for $n \rightarrow \infty$), which coincides with the ratio of our algorithm given by the theorem above for $k \rightarrow \infty$. Hence, k -max-search can indeed be understood as a natural bridge between the 1-max-search problem and the one-way trading problem.

For deterministic k -min-search we obtain the following statement.

Theorem 2. *Let $k \in \mathbb{N}$, $\varphi > 1$. There is a s^* -competitive deterministic algorithm for k -min-search, where $s^* = s^*(k, \varphi)$ is the unique solution of*

$$\frac{1 - 1/\varphi}{1 - 1/s^*} = \left(1 + \frac{1}{ks^*}\right)^k, \quad (2)$$

and there exists no deterministic algorithm with smaller competitive ratio. Furthermore, we have

- (i) $s^*(k, \varphi) \sim \sqrt{\frac{k+1}{2k}}\varphi$ for fixed $k \geq 1$ and $\varphi \rightarrow \infty$,
(ii) $s^*(k, \varphi) \sim (W(-\frac{\varphi-1}{e\varphi}) + 1)^{-1}$ for fixed $\varphi > 1$ and $k \rightarrow \infty$.

The algorithm in the theorem above is also given explicitly in Section 2. Surprisingly, although one might think that k -max-search and k -min-search should behave similarly with respect to competitive analysis, Theorem 2 states that this is in fact *not* the case. Indeed, according to Theorems 1 and 2, for large φ , the best algorithm for k -max-search achieves a competitive ratio of roughly $k \sqrt[k]{\varphi}$, while the best algorithm for k -min-search is at best $\sqrt{\varphi/2}$ -competitive. Similarly, when k is large, the competitive ratio of a best algorithm for k -max-search behaves like $\ln \varphi$, in contrast to k -min-search, where a straightforward analysis (i.e. series expansion of the W function around its pole) shows that the best algorithm achieves a ratio of $\Theta(\sqrt{\varphi})$. Hence, algorithms for k -min-search perform in the worst-case rather poorly compared to algorithms for k -max-search.

Furthermore, we investigate the performance of *randomized* algorithms for the problems in question. In [2] the authors gave a $\mathcal{O}(\ln \varphi)$ -competitive randomized algorithm for 1-max-search, but did not provide a lower bound.

Theorem 3. *Let $k \in \mathbb{N}$, $\varphi > 1$. For every randomized k -max-search algorithm RALG we have*

$$\text{CR}(\text{RALG}) \geq (\ln \varphi)/2 . \quad (3)$$

Furthermore, there is a $2 \ln \varphi$ -competitive randomized algorithm for $\varphi > 3$.

Note that the lower bound above is independent of k , i.e., randomized algorithms cannot improve their performance when k increases. In contrast to that, by considering Theorem 1, as k grows the performance of the best *deterministic* algorithm improves, and approaches $\ln \varphi$, which is only a multiplicative factor away from the best ratio that a randomized algorithm can achieve.

Our next result is about randomized algorithms for k -min-search.

Theorem 4. *Let $k \in \mathbb{N}$, $\varphi > 1$. For every randomized k -min-search algorithm RALG we have*

$$\text{CR}(\text{RALG}) \geq (1 + \sqrt{\varphi})/2 . \quad (4)$$

Again, the lower bound is independent of k . Furthermore, combined with Theorem 2, the theorem above states that for *all* $k \in \mathbb{N}$, randomization *does not* improve the performance (up to a multiplicative constant) of algorithms for k -min-search, compared to deterministic algorithms. This is again a difference between k -max-search and k -min-search.

1.2 Application to Robust Valuation of Lookback Options

In the second part of the paper we will use competitive k -search algorithms to derive upper bounds for the price of *lookback options*, a particular class of financial derivatives (see e.g. [8]). An option is a contract whereby the option holder has the right (but not obligation) to exercise a feature of the option contract on or before an exercise date, delivered by the other party – the writer

of the option. Thus, an option is in general an asymmetric contract. Since the option gives the buyer a right, it will have a price that the buyer has to pay to the option writer.

The most basic type of options are European options on a stock. They give the holder the right to buy (respectively, sell) the stock on a prespecified date T (expiry date) for a prespecified price K . Besides these standard and well-understood types, there is also a plethora of options with more complex features. One type are so called *lookback options*. A lookback *call* allows the holder to buy the underlying stock at time T from the option writer at the historical minimum price observed over $[0, T]$, and a lookback *put* to sell at the historical maximum.

A fundamental question is to determine the value of an option at time $t < T$. Black and Scholes [9] studied European call and put options on non-dividend paying stocks in a seminal paper. The key argument in their derivation is a *no arbitrage condition*. Loosely speaking, an arbitrage is a zero-risk, zero net investment strategy that still generates profit. If such an opportunity came about, market participants would immediately start exploiting it, pushing prices until the arbitrage opportunity ceases to exist. Black and Scholes essentially give a dynamic trading strategy in the underlying stock by which an option writer can risklessly hedge an option position. Thus, the no arbitrage condition implies that the cost of the trading strategy must equal the price of the option to date.

In the model of Black and Scholes trading is possible continuously in time and in arbitrarily small portions of shares. Moreover, a central underlying assumption is that the stock price follows a geometric Brownian motion (see e.g. [10]), which then became the standard model for option pricing. While it certainly shows many features that resemble reality fairly, the behavior of stock prices in practice is not fully consistent with this assumption. For instance, the distribution observed for the returns of stock price processes are non-Gaussian and typically heavy-tailed [11], leading to underestimation of extreme price movements. Furthermore, in practice trading is discrete, price paths include price jumps and stock price volatility is not constant. As a response, numerous modifications of the original Black-Scholes setting have been proposed, examining different stochastic processes for the stock price (for instance [12–14]).

In light of the persistent difficulties of finding and formulating the “right” model for the stock price dynamic, there have also been a number of attempts to price financial instruments by *relaxing the Black-Scholes assumptions* instead. The idea is to provide *robust* bounds that work with (almost) any evolution of the stock price rather than focusing on a specific formulation of the stochastic process. In this fashion, DeMarzo, Kremer and Mansour [15] derive both upper and lower bounds for option prices in a model of bounded quadratic variation, using competitive online trading algorithms. In the mathematical finance community, Epstein and Wilmott [16] propose non-probabilistic models for pricing interest rate securities in a framework of “worst-case scenarios”. Korn [17] combines the random walk assumption with a worst-case analysis to tackle optimal asset allocation under the threat of a crash.

In this spirit, using the deterministic k -search algorithms from Section 2 we derive in Section 4 upper bounds for the price of lookback calls and puts,

under the assumption of *bounded stock price paths* and non-existence of arbitrage opportunities. Interestingly, the resulting bounds are remarkably good, showing similar qualitative properties and quantitative values as pricing in the standard Black-Scholes model. Note that the assumption of a bounded stock price is indeed very minimal, since without *any* assumption about the magnitude of the stock price fluctuation in fact *no* upper bounds for the option price apply.

2 Deterministic Search

Let us consider the following *reservation price policy* RPP_{\max} for k -max-search. Prior to the start of the game, we choose *reservation prices* p_i^* ($i = 1 \dots k$). As the prices are sequentially revealed, RPP_{\max} accepts the first price that is at least p_1^* and sells one unit. It then waits for the first price that is at least p_2^* , and subsequently continues with all reservation prices. RPP_{\max} works through the reservation prices in a strictly sequential manner. Note that RPP_{\max} may be forced to sell at the last prices of the sequence, which may be lower than the remaining reservations, to meet the constraint of completing the sale.

The proof of the lemma below generalizes ideas used for *1-max-search* in [2].

Lemma 1. *Let $k \in \mathbb{N}$, $\varphi > 1$. Let $r^* = r^*(k, \varphi)$ be defined as in (1). Then the reservation price policy RPP_{\max} with reservation prices given by*

$$p_i^* = m \left[1 + (r^* - 1) \left(1 + \frac{r^*}{k} \right)^{i-1} \right], \quad (5)$$

satisfies $kp_{\max}(\sigma) \leq r^ \cdot \text{RPP}_{\max}(\sigma)$ for all $\sigma \in \mathcal{P}$. In particular, RPP_{\max} is a r^* -competitive algorithm for the k -max-search problem.*

Proof. For $0 \leq j \leq k$, let $\mathcal{P}_j \subseteq \mathcal{P}$ be the sets of price sequences for which RPP_{\max} accepts *exactly* j prices, excluding the forced sale at the end. Then \mathcal{P} is the disjoint union of the \mathcal{P}_j 's. To shorten notation, let us write $p_{k+1}^* = M$. Let $\varepsilon > 0$ be fixed and define the price sequences

$$\forall 0 \leq i \leq k : \quad \sigma_i = p_1^*, p_2^*, \dots, p_i^*, \underbrace{p_{i+1}^* - \varepsilon, \dots, p_{i+1}^* - \varepsilon}_k, \underbrace{m, m, \dots, m}_k.$$

Observe that as $\varepsilon \rightarrow 0$, each σ_j is a sequence yielding the worst-case ratio in \mathcal{P}_j , in the sense that for all $\sigma \in \mathcal{P}_j$

$$\frac{\text{OPT}(\sigma)}{\text{RPP}_{\max}(\sigma)} \leq \frac{kp_{\max}(\sigma)}{\text{RPP}_{\max}(\sigma)} \leq \frac{kp_{j+1}^*}{\text{RPP}_{\max}(\sigma_j)}. \quad (6)$$

Thus, to prove the statement we show that for $0 \leq j \leq k$ it holds $kp_{j+1}^* \leq r^* \cdot \text{RPP}_{\max}(\sigma_j)$. A straightforward calculation shows that for all $0 \leq j \leq k$

$$\sum_{i=1}^j p_i^* = m \left[j + k(1 - 1/r^*) \left((1 + r^*/k)^j - 1 \right) \right].$$

But then we have for $\varepsilon \rightarrow 0$, the competitive ratio is arbitrarily close to

$$\forall 0 \leq j \leq k: \quad \frac{kp_{j+1}^*}{\text{RPP}_{\max}(\sigma_j)} = \frac{kp_{j+1}^*}{\sum_{i=1}^j p_i^* + (k-j)m} = r^* .$$

Thus, from (6) the r^* -competitiveness of RPP_{\max} follows immediately. \square

Remark 1. While the proof above shows that the reservations prices in (5) are in fact the optimal choice, let us also briefly give an intuition on how to construct them. First, note that we have to choose the p_i^* 's such that

$$\frac{kp_1^*}{km} \stackrel{!}{=} \frac{kp_2^*}{p_1^* + (k-1)m} \stackrel{!}{=} \dots \stackrel{!}{=} \frac{kM}{\sum_{i=1}^k p_i^*} \stackrel{!}{=} r^* . \quad (7)$$

(The nominator is the objective value of OPT on σ_i as $\varepsilon \rightarrow 0$, whereas the denominator is the value of RPP_{\max} on the same sequence.) For $0 \leq i \leq k$, let $r_i = p_i^*/p_1^*$. By comparing adjacent terms in (7), it is easy to see that r_i satisfies the simple recurrence

$$r_i = r_{i-1}(1 + p_1^*/(km)) - 1/k, \quad \text{and} \quad r_1 = 1 ,$$

and standard methods readily yield a closed formula for p_i^* in terms of p_1^* . Furthermore, using (7) we obtain the explicit expression for p_1^* .

From the choice of reservation prices in Lemma 1, we see that in fact no deterministic algorithm will be able to do better than RPP_{\max} in the worst-case.

Lemma 2. *Let $k \geq 1$, $\varphi > 1$. Then $r^*(k, \varphi)$ given by (1) is the lowest possible competitive ratio that a deterministic k -max-search algorithm can achieve.*

Proof. Let ALG be any deterministic algorithm. We shall show that ALG cannot achieve a ratio lower than $r^*(k, \varphi)$. Let p_1^*, \dots, p_k^* be the price sequence defined by (5). We start by presenting p_1^* to ALG , at most k times or until ALG accepts it. If ALG never accepts p_1^* , we drop the price to m for the remainder, and ALG achieves a competitive ratio of $p_1^*/m = r^*(k, \varphi)$. If ALG accepts p_1^* , we continue the price sequence by presenting p_2^* to ALG at most k times. Again, if ALG never accepts p_2^* before we presented it k times, we drop to m for the remainder and ALG achieves a ratio not lower than $kp_2^*/(p_1^* + (k-1)m) = r^*(k, \varphi)$. We continue in that fashion by presenting each p_i^* at most k times (or until ALG accepts it). Whenever ALG doesn't accept a p_i^* after presenting it k times, we drop the price to m . If ALG subsequentially accepts all p_1^*, \dots, p_k^* , we conclude with k times M . In any case, ALG achieves only a ratio of at most $r^*(k, \varphi)$. \square

With the above preparations we are ready to prove Theorem 1.

Proof (of Theorem 1). The first statement follows directly from Lemma 1 and Lemma 2. To show (i), first observe that for $k \geq 1$ fixed, $r^* = r^*(\varphi)$ must satisfy $r^* \rightarrow \infty$ as $\varphi \rightarrow \infty$, and r^* is an increasing function of φ . Let $r_+ = k^{\frac{k}{k+1}} k^{+1}\sqrt{\varphi}$. Then, for $\varphi \rightarrow \infty$, we have

$$(r_+ - 1) \left(1 + \frac{r_+}{k}\right)^k = (1 + o(1)) \left(k^{\frac{k}{k+1}} k^{+1}\sqrt{\varphi} \cdot \left(k^{-\frac{1}{k+1}} k^{+1}\sqrt{\varphi}\right)^k\right) = (1 + o(1))\varphi .$$

Furthermore, let $\varepsilon > 0$ and set $r_- = (1 - \varepsilon)k^{\frac{k}{k+1}} \sqrt[k+1]{\varphi}$. A similar calculation as above shows that for sufficiently large φ we have

$$(r_- - 1) \left(1 + \frac{r_-}{k}\right)^k \geq (1 - 3k\varepsilon) \varphi .$$

Thus, $r = (1 + o(1))k^{\frac{k}{k+1}} \sqrt[k+1]{\varphi}$ indeed satisfies (1) for $\varphi \rightarrow \infty$. For the proof of (ii), note that for $k \rightarrow \infty$ and φ fixed, equation (1) becomes $(\varphi - 1)/(r^* - 1) = e^{r^*}$, and thus

$$(\varphi - 1)/e = (r^* - 1) e^{r^* - 1} .$$

The claim follows from the definition of the W -function. \square

Similarly, we can construct a reservation price policy RPP_{\min} for k -min-search. Naturally, RPP_{\min} is modified such that it accepts the first price *lower* than the current reservation price.

Lemma 3. *Let $k \in \mathbb{N}$, $\varphi > 1$. Let $s^* = s^*(k, \varphi)$ be defined as in (2). Then the reservation price policy RPP_{\min} with reservation prices $p_1^* > \dots > p_k^*$,*

$$p_i^* = M \left[1 - \left(1 - \frac{1}{s^*}\right) \left(1 + \frac{1}{ks^*}\right)^{i-1} \right] , \quad (8)$$

satisfies $\text{RPP}_{\min}(\sigma) \leq s^(k, \varphi) \cdot k p_{\min}(\sigma)$, and is a $s^*(k, \varphi)$ -competitive deterministic algorithm for k -min-search.*

Proof. The proof is analogous to the proof of Lemma 1. Again, for $0 \leq j \leq k$, let $\mathcal{P}_j \subseteq \mathcal{P}$ be the sets of price sequences for which RPP_{\min} accepts *exactly* j prices, excluding the forced sale at the end. To shorten notation, define $p_{k+1}^* = m$. Let $\varepsilon > 0$ be fixed and define the price sequences

$$\sigma_i = p_1^*, p_2^*, \dots, p_i^*, \underbrace{p_{i+1}^* + \varepsilon, \dots, p_{i+1}^* + \varepsilon}_k, \underbrace{M, \dots, M}_k , \quad \text{for } 0 \leq i \leq k .$$

As $\varepsilon \rightarrow 0$, each σ_j is a sequence yielding the worst-case ratio in \mathcal{P}_j , in the sense that for all $\sigma \in \mathcal{P}_j$, $0 \leq j \leq k$,

$$\frac{\text{RPP}_{\min}(\sigma)}{\text{OPT}(\sigma)} \leq \frac{\text{RPP}_{\min}(\sigma)}{k p_{\min}(\sigma)} \leq \frac{\text{RPP}_{\min}(\sigma_j)}{k p_{j+1}^*} . \quad (9)$$

Straightforward calculation shows that for $\varepsilon \rightarrow 0$

$$\forall 0 \leq j \leq k : \quad \frac{\text{RPP}_{\min}(\sigma_j)}{k p_{j+1}^*} = \frac{\sum_{i=1}^j p_i^* + (k-j)M}{k p_{j+1}^*} = s^* ,$$

and hence

$$\forall \sigma \in \mathcal{P} : \quad \frac{\text{RPP}_{\min}(\sigma)}{k p_{\min}(\sigma)} \leq s^* .$$

Since $\text{OPT}(\sigma) \geq k p_{\min}(\sigma)$ for all $\sigma \in \mathcal{P}$, this also implies that RPP_{\min} is s^* -competitive. \square

Again, no deterministic algorithm can do better than RPP_{\min} in Lemma 3.

Lemma 4. *Let $k \geq 1$, $\varphi > 1$. Then $s^*(k, \varphi)$ given by (2) is the lowest possible competitive ratio that a deterministic k -min-search algorithm can achieve.*

The proof of Lemma 4 is identical to the proof of Lemma 2, except that the adversary now uses the prices defined by (8) and the roles of m and M interchanged.

Using Lemma 3 and Lemma 4 we can now prove Theorem 2.

Proof (of Theorem 2). The first part of the Theorem follows directly from Lemma 3 and Lemma 4. To show (i), first observe that for $k \geq 1$ fixed, $s^* = s^*(\varphi)$ must satisfy $s^* \rightarrow \infty$ as $\varphi \rightarrow \infty$, and s^* is an increasing function of φ . With this assumption we can expand the right-hand side of (2) with the binomial theorem to obtain

$$\frac{1 - 1/\varphi}{1 - 1/s^*} = 1 + \frac{1}{s^*} + \frac{k-1}{2k(s^*)^2} + \Theta((s^*)^{-3}) \implies \frac{1}{\varphi} = \frac{k+1}{2k(s^*)^2} + \Theta((s^*)^{-3}) .$$

By solving this equation for s^* we obtain (i). For the proof of (ii), first observe that for $\varphi \geq 1$ fixed, $s^* = s^*(k)$ must satisfy $s^*(k) \leq C$, for some constant C which may depend on φ . Indeed, if $s^*(k) \rightarrow \infty$ with $k \rightarrow \infty$, then by taking limits on both sides of (2) yields

$$1 - \frac{1}{\varphi} = \lim_{k \rightarrow \infty} \left(1 + \frac{1}{ks^*(k)} \right)^k = 1 ,$$

which is a contradiction. Thus, $s^* = \Theta(1)$ and we obtain from (2) by taking limits

$$\frac{1 - 1/\varphi}{1 - 1/s^*} = \lim_{k \rightarrow \infty} \left(1 + \frac{1}{ks^*} \right)^k = e^{1/s^*} ,$$

and (ii) follows immediately by the definition of the W -function. □

3 Randomized Search

3.1 Lower Bound for Randomized k -max-search

We consider $k = 1$ first. The optimal deterministic online algorithm achieves a competitive ratio of $r^*(1, \varphi) = \sqrt{\varphi}$. As shown in [2], randomization can dramatically improve this. Assume for simplicity that $\varphi = 2^\ell$ for some integer ℓ . For $0 \leq j < \ell$ let $\text{RPP}_{\max}(j)$ be the reservation price policy with reservation $m2^j$, and define EXPO to be a uniform probability mixture over $\{\text{RPP}_{\max}(j)\}_{j=0}^{\ell-1}$.

Lemma 5 (Levin, see [2]). *Algorithm EXPO is $\mathcal{O}(\ln \varphi)$ -competitive.*

We shall prove that EXPO is in fact the optimal randomized online algorithm for 1-max-search. We will use the following version of *Yao's principle* [18].

Theorem 5 (Yao’s principle). *For an online maximization problem denote by \mathcal{S} the set of possible input sequences, and by \mathcal{A} the set of deterministic algorithms, and assume that \mathcal{S} and \mathcal{A} are finite. Fix any probability distribution $y(\sigma)$ on \mathcal{S} , and let S be a random sequence according to this distribution. Let RALG be any mixed strategy, given by a probability distribution on \mathcal{A} . Then,*

$$\text{CR}(\text{RALG}) = \max_{\sigma \in \mathcal{S}} \frac{\text{OPT}(\sigma)}{\mathbb{E}[\text{RALG}(\sigma)]} \geq \left(\max_{\text{ALG} \in \mathcal{A}} \mathbb{E} \left[\frac{\text{ALG}(S)}{\text{OPT}(S)} \right] \right)^{-1}. \quad (10)$$

Note that the first expectation is taken with respect to the randomization of the algorithm RALG, whereas the second expectation is taken with respect to the input distribution $y(\sigma)$.

The reader is referred to standard textbooks for a proof (e.g. Chapter 6 and 8 in [19]). In words, Yao’s principle states that we obtain a lower bound on the competitive ratio of the best randomized algorithm by calculating the performance of the best deterministic algorithm for a chosen probability distribution of input sequences. Note that (10) gives a lower bound for *arbitrary* chosen input distributions. However, only for well-chosen y ’s we will obtain strong lower bounds.

We first need to establish the following lemma on the representation of an arbitrary randomized algorithm for k -search.

Lemma 6. *Let RALG be a randomized algorithm for the k -max-search problem. Then RALG can be represented by a probability distribution on the set of all deterministic algorithms for the k -max-search problem.*

Proof. The proof of the statement is along the lines of the proof of Theorem 1 in [2]. Here we only sketch the proof idea. Using game-theoretic terminology, RALG may be either a mixed strategy (a distribution on deterministic algorithms, from which one is randomly chosen prior to the start of the game) or a behavioral strategy (where an independent random choice may be made at each point during the game). As we have perfect recall in k -search (player has no memory restrictions), k -search is a linear game. For linear games, every behavioral strategy has an equivalent mixed algorithm. Thus, we can always assume that RALG is a mixed strategy given by a probability distribution on the set of all deterministic algorithms. \square

The next lemma yields the desired lower bound.

Lemma 7. *Let $\varphi > 1$. Every randomized 1-max-search algorithm RALG satisfies*

$$\text{CR}(\text{RALG}) \geq (\ln \varphi)/2 .$$

Proof. Let $b > 1$ and $\ell = \log_b \varphi$. We define a finite approximation of \mathcal{I} by $\mathcal{I}_b = \{mb^i \mid i = 0 \dots \ell\}$, and let $\mathcal{P}_b = \bigcup_{n \geq k} \mathcal{I}_b^n$. We consider the 1-max-search problem on \mathcal{P}_b . As \mathcal{P}_b is finite, also the set of deterministic algorithms \mathcal{A}_b is finite. For $0 \leq i \leq \ell - 1$, define sequences of length ℓ by

$$\sigma_i = mb^0, \dots, mb^i, m, \dots, m . \quad (11)$$

Let $\mathcal{S}_b = \{\sigma_i \mid 0 \leq i \leq \ell - 1\}$ and define the probability distribution y on \mathcal{P}_b by

$$y(\sigma) = \begin{cases} 1/\ell & \text{for } \sigma \in \mathcal{S}_b, \\ 0 & \text{otherwise.} \end{cases}$$

Let $\text{ALG} \in \mathcal{A}_b$. Note that for all $1 \leq i \leq \ell$, the first i prices of the sequences σ_j with $j \geq i - 1$ coincide, and ALG cannot distinguish them up to time i . As ALG is deterministic, it follows that if ALG accepts the i -th price in $\sigma_{\ell-1}$, it will accept the i -th price in all σ_j with $j \geq i - 1$. Thus, for every ALG , let $0 \leq \chi(\text{ALG}) \leq \ell - 1$ be such that ALG accepts the $(\chi(\text{ALG}) + 1)$ -th price, i.e. $mb^{\chi(\text{ALG})}$, in $\sigma_{\ell-1}$. ALG will then earn $mb^{\chi(\text{ALG})}$ on all σ_j with $j \geq \chi(\text{ALG})$, and m on all σ_j with $j < \chi(\text{ALG})$. To shorten notation, we write χ instead of $\chi(\text{ALG})$ in the following. Thus, we have

$$\mathbb{E} \left[\frac{\text{ALG}}{\text{OPT}} \right] = \frac{1}{\ell} \left[\sum_{j=0}^{\chi-1} \frac{m}{mb^j} + \sum_{j=\chi}^{\ell-1} \frac{mb^{\chi}}{mb^j} \right] = \frac{1}{\ell} \left[\frac{1-b^{-\chi}}{1-b^{-1}} + \frac{1-b^{-(\ell-\chi)}}{1-b^{-1}} \right],$$

where the expectation $\mathbb{E}[\cdot]$ is with respect to the probability distribution $y(\sigma)$. If we consider the above term as a function of χ , then it is easily verified that it attains its maximum at $\chi = \ell/2$. Thus,

$$\max_{\text{ALG} \in \mathcal{A}_b} \mathbb{E} \left[\frac{\text{ALG}}{\text{OPT}} \right] \leq \frac{1}{\ell} \left(1 - \frac{1}{\sqrt{\varphi}} \right) \frac{2b}{b-1} \leq \frac{1}{\ln \varphi} \cdot \frac{2b \ln b}{b-1}. \quad (12)$$

Let \mathcal{Y}_b be the set of all randomized algorithms for 1-max-search with possible price sequences \mathcal{P}_b . By Lemma 6, each $\text{RALG}_b \in \mathcal{Y}_b$ may be given as a probability distribution on \mathcal{A}_b . Since \mathcal{A}_b and \mathcal{S}_b are both finite, we can apply Theorem 5. Thus, for all $b > 1$ and all $\text{RALG}_b \in \mathcal{Y}_b$, we have

$$\text{CR}(\text{RALG}_b) \geq \left(\max_{\text{ALG} \in \mathcal{A}_b} \mathbb{E} \left[\frac{\text{ALG}}{\text{OPT}} \right] \right)^{-1} \geq \ln \varphi \frac{b-1}{2b \ln b}.$$

Let \mathcal{Y} be the set of all randomized algorithms for 1-max-search on \mathcal{P} . Since for $b \rightarrow 1$, we have $\mathcal{A}_b \rightarrow \mathcal{A}$, $\mathcal{Y}_b \rightarrow \mathcal{Y}$ and $(b-1)/(2b \ln b) \rightarrow \frac{1}{2}$, the proof is completed. \square

In fact, Lemma 7 can be generalized to arbitrary $k \geq 1$.

Lemma 8. *Let $k \in \mathbb{N}$, $b > 1$ and $\varphi > 1$. Let RALG be any randomized algorithm for k -max-search. Then, we have*

$$\text{CR}(\text{RALG}) \geq (\ln \varphi)/2.$$

Proof. Let $1 < b < \varphi$ and $\ell = \log_b \varphi$. We define \mathcal{P}_b and \mathcal{A}_b as in the proof of Lemma 7. For $0 \leq i \leq \ell - 1$, define

$$\sigma_i = \underbrace{mb^0, \dots, mb^0}_k, \dots, \underbrace{mb^i, \dots, mb^i}_k, \underbrace{m, \dots, m}_{k(\ell-i-1)}. \quad (13)$$

Let $\mathcal{S}_b = \bigcup_{1 \leq i \leq \ell} \sigma_i$ and define the probability distribution y on \mathcal{P}_b by

$$y(\sigma) = \begin{cases} 1/\ell & \text{for } \sigma \in \mathcal{S}_b, \\ 0 & \text{otherwise.} \end{cases}$$

Similarly as in the proof of Lemma 7, we characterize every algorithm $\text{ALG} \in \mathcal{A}_b$ by a vector $(\chi_i)_{1 \leq i \leq k}$ where mb^{χ_i} is the price for which ALG sells the i -th unit on $\sigma_{\ell-1}$. By construction, we have $\chi_1 \leq \dots \leq \chi_k$. (Recall that $\sigma_{\ell-1}$ is the sequence that is increasing until the very end.) Note that for all $1 \leq i \leq \ell$, the sequences $\{\sigma_j \mid j \geq i-1\}$ are not distinguishable up to time ik , since the first ik prices of those sequences are identical. Let $0 \leq j \leq \ell-1$ and $t = \max\{i \mid \chi_i \leq j\}$. When presented σ_j , ALG accepts all prices mb^{χ_i} for which $\chi_i \leq j$. Hence, we have $\text{OPT}(\sigma_j) = kmb^j$ and $\text{ALG}(\sigma_j) = (k-t)m + \sum_{s=1}^t mb^{\chi_s}$, i.e. ALG can successfully sell for its first t reservation prices. To abbreviate notation, let $\chi_0 = 0$ and $\chi_{k+1} = \ell$, and define $\delta_t = \chi_{t+1} - \chi_t$. Taking expectation with respect to $y(\sigma)$, we have

$$\begin{aligned} \mathbb{E} \left[\frac{\text{ALG}}{\text{OPT}} \right] &= \frac{1}{\ell} \sum_{t=0}^k \sum_{j=\chi_t}^{\chi_{t+1}-1} \frac{(k-t)m + \sum_{s=1}^t mb^{\chi_s}}{kmb^j} \\ &= \frac{1}{\ell} \sum_{t=0}^k \frac{(k-t + \sum_{s=1}^t b^{\chi_s}) \sum_{j=0}^{\delta_t-1} b^{-j}}{kb^{\chi_t}} \\ &= \frac{1}{\ell} \sum_{t=0}^k \frac{(k-t + \sum_{s=1}^t b^{\chi_s})(1-b^{-\delta_t})}{kb^{\chi_t}(1-b^{-1})}. \end{aligned}$$

Straightforward yet tedious algebra simplifies this expression to

$$\mathbb{E} \left[\frac{\text{ALG}}{\text{OPT}} \right] = \frac{\sum_{t=1}^k 1 - b^{-\chi_t} + \sum_{t=1}^k 1 - b^{-(\ell-\chi_t)}}{\ell k(1-b^{-1})},$$

and the maximum over $\{\chi_1, \dots, \chi_k\}$ is attained at $\chi_1 = \dots = \chi_k = \ell/2$. Thus, defining $\chi = \ell/2$ we have

$$\max_{\text{ALG} \in \mathcal{A}_b} \mathbb{E} \left[\frac{\text{ALG}}{\text{OPT}} \right] \leq \frac{1}{\ell} \left[\frac{1-b^{-\chi}}{1-b^{-1}} + \frac{1-b^{-(\ell-\chi)}}{1-b^{-1}} \right] = \frac{2b}{\ell(b-1)} \left(1 - \frac{1}{\sqrt{\varphi}} \right),$$

which is exactly (12) in the proof of Lemma 7. Thus, we can argue as in the remainder of the proof of Lemma 7, and let again $b \rightarrow 1$ to conclude that $\text{CR}(\text{RALG}) \geq (\ln \varphi)/2$ for all randomized algorithms RALG for k -max-search. \square

Giving an optimal randomized algorithm for k -max-search is straightforward. For $1 < b < \varphi$ and $\ell = \log_b \varphi$, EXPO_k chooses j uniformly at random from $\{0, \dots, \ell-1\}$, and sets all its k reservation prices to mb^j .

Lemma 9. *Let $k \in \mathbb{N}$. EXPO_k is an asymptotically optimal randomized algorithm for the k -max-search problem with $\text{CR}(\text{EXPO}_k) = \ln \varphi \cdot \frac{(b-1)}{\ln b} \frac{\varphi}{\varphi-1}$.*

Proof. We want to determine

$$\text{CR}(\text{EXPO}_k) = \max_{\sigma \in \mathcal{P}} R(\sigma), \quad \text{where} \quad R(\sigma) = \frac{\text{OPT}(\sigma)}{\mathbb{E}[\text{EXPO}_k(\sigma)]} . \quad (14)$$

Obviously, a sequence σ that maximizes R is non-decreasing, since rearranging σ does not change the objective value of OPT , but may increase the objective value of EXPO_k . Let

$$\hat{\sigma} = \underbrace{m, \dots, m}_k, \underbrace{mb^1, \dots, mb^1}_k, \underbrace{mb^2, \dots, mb^2}_k, \dots, \underbrace{mb^{\ell-1}, \dots, mb^{\ell-1}}_k, \underbrace{M, \dots, M}_k .$$

We will now show that for every non-decreasing sequence $\sigma = p_1, p_2, \dots, p_N$ it holds

$$R(\sigma) \leq R(\sigma \circ \hat{\sigma}) , \quad (15)$$

where $\sigma \circ \hat{\sigma}$ is the concatenation of σ and $\hat{\sigma}$. Let P be the reservation price of EXPO_k , and let $\text{EXPO}_k^j(\sigma) = \mathbb{E}[\text{EXPO}_k(\sigma) \mid P = mb^j]$. To see the first inequality we shall show that for all $0 \leq j < \ell$,

$$\frac{\text{EXPO}_k^j(\sigma \circ \hat{\sigma})}{\text{OPT}(\sigma \circ \hat{\sigma})} \leq \frac{\text{EXPO}_k^j(\sigma)}{\text{OPT}(\sigma)} , \quad (16)$$

which yields

$$\frac{1}{R(\sigma)} = \mathbb{E} \left[\frac{\text{EXPO}_k(\sigma)}{\text{OPT}(\sigma)} \right] \geq \mathbb{E} \left[\frac{\text{EXPO}_k(\sigma \circ \hat{\sigma})}{\text{OPT}(\sigma \circ \hat{\sigma})} \right] = \frac{1}{R(\sigma \circ \hat{\sigma})} ,$$

i.e., $R(\sigma) \leq R(\sigma \circ \hat{\sigma})$. To see (16), note that if p_l is the first price accepted by EXPO_k^j in σ , then EXPO_k^j will also accept $p_{l+1}, \dots, p_{l+k-1}$. This follows from the property of σ being non-decreasing and from the fact that all reservation prices of EXPO_k^j are identical. Now we distinguish two cases: either $l = N - k + 1$ (i.e. EXPO_k^j accepts the last k prices in σ , possibly forced by the constraint to finish the sale by the end of the sequence σ) or $l < N - k + 1$ (i.e. EXPO_k^j can successfully sell all k units for prices of at least its reservation price mb^j). In the first case, $\text{EXPO}_k^j(\sigma) = \text{OPT}(\sigma)$ and (16) follows trivially, since we always have $\text{EXPO}_k^j(\sigma \circ \hat{\sigma}) \leq \text{OPT}(\sigma \circ \hat{\sigma})$. In the second case, $\text{OPT}(\sigma \circ \hat{\sigma}) = kM \geq \text{OPT}(\sigma)$ and $\text{EXPO}_k^j(\sigma \circ \hat{\sigma}) = \text{EXPO}_k^j(\sigma)$, since EXPO_k^j already accepted k prices already before the end of σ was reached, and it cannot accept any prices in the second part of $\sigma \circ \hat{\sigma}$. Hence, (16) also holds in this case. This shows (15).

Now observe that for any non-decreasing σ we have

$$\mathbb{E}[\text{EXPO}_k(\sigma \circ \hat{\sigma})] \geq \mathbb{E}[\text{EXPO}_k(\hat{\sigma})] ,$$

since for every j algorithm EXPO_k^j accepts k prices in $\sigma \circ \hat{\sigma}$ that are at least mb^j , but in $\hat{\sigma}$ it accepts k times exactly its reservation price mb^j . Combined with the fact that $\text{OPT}(\sigma \circ \hat{\sigma}) = \text{OPT}(\hat{\sigma}) = kM$, this yields

$$R(\sigma \circ \hat{\sigma}) \leq R(\hat{\sigma}) .$$

With (15), this implies that $\hat{\sigma}$ is a worst-case sequence for (14). Therefore, we have

$$\text{CR}(\text{EXPO}_k) = \text{R}(\hat{\sigma}) = \frac{kM}{\frac{1}{\ell} \sum_{j=0}^{\ell-1} kmb^j} = \ell \frac{\varphi(b-1)}{\varphi-1} = \ln \varphi \frac{\varphi}{\varphi-1} \cdot \frac{(b-1)}{\ln b},$$

since $M = \varphi m$ and $b^\ell = \varphi$. \square

For $\varphi > 3$ and $b < 3/2$, we have $\frac{\varphi}{\varphi-1} \frac{(b-1)}{\ln b} < 2$, and hence combining Lemma 8 and 9 we immediately obtain Theorem 3.

3.2 Randomized k -Min-Search

The proof of the lower bound for k -min-search, Theorem 4, uses an analogous version of Yao's principle (see for instance Theorem 8.5 in [19]).

Theorem 6 (Yao's principle for cost minimization problems). *For an online cost minimization problem Π , let the set of possible input sequences \mathcal{S} and the set of deterministic algorithms \mathcal{A} be finite, and given by $\mathcal{S} = \{\sigma_1, \dots, \sigma_n\}$ and $\mathcal{A} = \{\text{ALG}_1, \dots, \text{ALG}_m\}$. Fix any probability distribution $y(\sigma)$ on \mathcal{S} . Let RALG be any mixed strategy, given by a probability distribution $x(a)$ on \mathcal{A} . Then,*

$$\text{CR}(\text{RALG}) = \max_{\sigma \in \mathcal{S}} \frac{\mathbb{E}[\text{RALG}(\sigma)]}{\text{OPT}(\sigma)} \geq \min_{\text{ALG} \in \mathcal{A}} \mathbb{E} \left[\frac{\text{ALG}}{\text{OPT}} \right].$$

We are now ready to prove Theorem 4.

Proof (Theorem 4). We shall consider first the case $k = 1$. Let $\mathcal{S} = \{\sigma_1, \sigma_2\}$ with

$$\sigma_1 = m\sqrt{\varphi}, M, \dots, M \quad \text{and} \quad \sigma_2 = m\sqrt{\varphi}, m, M, \dots, M,$$

and let $y(\sigma)$ be the uniform distribution on \mathcal{S} . For $i \in \{1, 2\}$, let ALG_i be the reservation price policy with reservation prices $p_1^* = m\sqrt{\varphi}$ and $p_2^* = m$, respectively. Obviously, the best deterministic algorithm against the randomized input given by the distribution $y(\sigma)$ behaves either like ALG_1 or ALG_2 . Since

$$\mathbb{E} \left[\frac{\text{ALG}_i}{\text{OPT}} \right] = (1 + \sqrt{\varphi})/2, \quad i \in \{1, 2\},$$

the desired lower bound follows from the min-cost version of Yao's principle. For general $k \geq 1$, we repeat the prices $m\sqrt{\varphi}$ and m in σ_1 and σ_2 k times each. Observe that in that case we can partition the set of all deterministic algorithms into $k+1$ equivalence classes, according to the number price quotations accepted from the first k prices $m\sqrt{\varphi}, \dots, m\sqrt{\varphi}$, as σ_1 and σ_2 are not distinguishable until the $(k+1)$ th price. Suppose ALG accepts j times the price $m\sqrt{\varphi}$. Then we have

$$\mathbb{E} \left[\frac{\text{ALG}}{\text{OPT}} \right] = \frac{1}{2} \left(\frac{j m \sqrt{\varphi} + (k-j) M}{k m \sqrt{\varphi}} + \frac{j m \sqrt{\varphi} + (k-j) m}{k m} \right) = (1 + \sqrt{\varphi})/2$$

for all $0 \leq j \leq k$, and the lower bound follows again from Yao's principle. \square

4 Robust Valuation of Lookback Options

In this section, we use the deterministic k -search algorithms from Section 2 to derive upper bounds for the price of lookback options under the assumption of *bounded stock price paths* and non-existence of arbitrage opportunities. We consider a discrete-time model of trading. For simplicity we assume that the interest rate is zero. The price of the stock at time $t \in \{0, 1, \dots, T\}$ is given by S_t , with S_0 being the price when seller and buyer enter the option contract.

Recall that the holder of a lookback call has the right to buy shares from the option writer for the price $S_{min} = \min\{S_t \mid 0 \leq t \leq T\}$. We shall assume that the lookback call is on $k \geq 1$ units of the underlying stock.

Note that since $S_{min} \leq S_T$, the option holder is never worse off executing the option at the expiry date T (and then immediately selling the shares for S_T) rather than to forgo his option. Hence, a lookback call option can always be considered as executed at expiry. This is in contrast to a European call option, where the option holder is not interested in executing his option if the price S_T at expiry is below the pre-specified strike price K .

Neglecting stock price appreciation, upwards and downwards movement of the stock price is equally likely. Consequently, we will assume a symmetric trading range $[\varphi^{-1/2}S_0, \varphi^{1/2}S_0]$ with $\varphi > 1$. We refer to a price path that satisfies $S_t \in [\varphi^{-1/2}S_0, \varphi^{1/2}S_0]$ for all $1 \leq t \leq T$ as a (S_0, φ) price path.

4.1 Upper Bounds for the Price of Lookback Options

Theorem 7. *Assume $(S_t)_{0 \leq t \leq T}$ is a (S_0, φ) stock price path. Let $s^*(k, \varphi)$ be given by (2), and let*

$$V_{Call}^*(S_0, \varphi) = S_0(s^*(k, \varphi) - 1)/\sqrt{\varphi} . \quad (17)$$

Let V be the option premium paid at time $t = 0$ for a lookback call option on k shares expiring at time T . Suppose we have $V > V_{Call}^(k, S_0, \varphi)$. Then there exists an arbitrage opportunity for the option writer, i.e., there is a zero-net-investment strategy which yields a profit for all (S_0, φ) stock price paths.*

Proof. In the following, let C_t denote the money in the option writer's cash account at time t . At time $t = 0$, the option writer receives V from the option buyer, and we have $C_0 = V$. The option writer then successively buys k shares, one-by-one, applying RPP_{min} for k -min-search with reservation prices as given by (8). Let H be the total sum of money spent for purchasing k units of stock. By Lemma 3 we have $H \leq ks^*(k, \varphi)S_{min}$. At time T the option holder purchases k shares from the option writer for kS_{min} in cash. As noted above, a lookback call option can always be considered executed at the expiration time T ; if the option holder does not execute his option, the option writer simply sells the k shares again on the market for $kS_T \geq kS_{min}$.

After everything has been settled, we have

$$C_T = V - H + kS_{min} \geq V + kS_{min}(1 - s^*(k, \varphi)) .$$

Because of $S_{min} \geq S_0/\sqrt{\varphi}$ and $V > V_{\text{Call}}^*(S_0, \varphi)$, we conclude that $C_T > 0$ for all possible (S_0, φ) stock price paths. Hence, this is indeed a zero net investment profit for the option writer on all (S_0, φ) stock price paths. \square

Under the no-arbitrage assumption, we immediately obtain an upper bound for the value of a lookback call option.

Corollary 1. *Under the no-arbitrage assumption, we have $V \leq V_{\text{Call}}^*(k, S_0, \varphi)$, with $V_{\text{Call}}^*(k, S_0, \varphi)$ as defined in (17).*

Using Lemma 1 and similar no-arbitrage arguments, also an upper bound for the price of a lookback put option can be established.

Note that it is not possible to derive non-trivial *lower bounds* for lookback options in the bounded stock price model, as a (S_0, φ) -price path may stay at S_0 throughout, making the option mature worthless for the holder. To derive lower bounds, there must be a promised fluctuation of the stock. In the classical Black-Scholes pricing model, this is given by the volatility of the Brownian motion.

We shall remark that in practice there is certainly no trading range in which the stock price stays with certainty; what we rather can give are trading ranges in which the price stays with sufficiently high probability. V_{Call}^* is then to be understood as a bound for the option price up to a certain residual risk. Note that whereas the Black-Scholes-type price (18), which shall be given in the next section, has no such residual risk within the Black-Scholes model, it does certainly have significant *model risk* due to the fact that the Black-Scholes assumptions might turn out to be incorrect in the first place.

4.2 Comparison to Pricing in Black-Scholes Model

Goldman, Sosin and Gatto [20] give closed form solutions for the value of lookback puts and calls in the Black-Scholes setting. Let σ be the volatility of the stock price, modeled by a geometric Brownian motion, $S(t) = S_0 \exp(-\sigma^2 t/2 + \sigma B(t))$, where $B(t)$ is a standard Brownian motion. Let $\Phi(x)$ denote the cumulative distribution function of the standard normal distribution. Then, for zero interest rate, at time $t = 0$ the value of a lookback call on one share of stock, expiring at time T , is given by

$$V_{\text{Call}}^{\text{BS}}(S_0, T, \sigma) = S_0(2\Phi(\sigma\sqrt{T}/2) - 1) . \quad (18)$$

The hedging strategy is a certain roll-over replication strategy of a series of European call options. Everytime the stock price hits a new all-time low, the hedger has to “roll-over” her position in the call to one with a new strike. Interestingly, this kind of behavior to act only when a new all-time low is reached resembles the behavior of RPP_{\min} for k -min-search.

For a numerical comparison of the bound $V_{\text{Call}}^*(k, S_0, \varphi, T)$ with the Black-Scholes-type pricing formula (18), we choose the fluctuation rate $\varphi = \varphi(T)$ such that the expected trading range $[\mathbb{E}(\min_{0 \leq t \leq T} S_t), \mathbb{E}(\max_{0 \leq t \leq T} S_t)]$ of a geometric Brownian motion starting at S_0 with volatility σ is $[\varphi^{-1/2}S_0, \varphi^{1/2}S_0]$.

Figure 1 shows the results for $\sigma = 0.2$, $S_0 = 20$ and $k = 10$. As can be seen from the graph, the bound V_{Call}^* is qualitatively and quantitatively similar to the Black-Scholes-type pricing (18). However, it is important to note that the two pricing formulas rely on two different models. In the Black-Scholes model, (18) is the correct price for a lookback option. On the other hand, the key advantage of our price bound are its weak modeling assumptions on the price dynamics, and hence the price bound holds even in situations where the Black-Scholes assumptions might break down. Certainly, both concepts have strengths and weaknesses, and a good analyst consults, compares and combines both.




Fig. 1. The left plot shows the expected trading range of a geometric Brownian motion with volatility $\sigma = 0.2$ and $S(0) = 20$. The right plot shows the price of a lookback call with maturity T in the Black-Scholes model (solid line) and the bound V_{Call}^* (dashed line), with $k = 100$ and $\varphi(T)$ chosen to match the expected trading range.

References

1. Sleator, D.D., Tarjan, R.E.: Amortized efficiency of list update and paging rules. *Comm. ACM* **28**(2) (1985) 202–208
2. El-Yaniv, R., Fiat, A., Karp, R.M., Turpin, G.: Optimal search and one-way trading online algorithms. *Algorithmica* **30**(1) (2001) 101–139
3. Lippmann, S.A., McCall, J.J.: The economics of job search: a survey. *Economic Inquiry* **XIV** (1976) 155–189
4. Lippmann, S.A., McCall, J.J.: The economics of uncertainty: selected topics and probabilistic methods. *Handbook of mathematical economics* **1** (1981) 211–284
5. Rosenfield, D.B., Shapiro, R.D.: Optimal adaptive price search. *Journal of Economic Theory* **25**(1) (1981) 1–20
6. Ajtai, M., Megiddo, N., Waarts, O.: Improved algorithms and analysis for secretary problems and generalizations. *SIAM Journal on Disc. Math.* **14**(1) (2001) 1–27
7. Kleinberg, R.D.: A multiple-choice secretary algorithm with applications to online auctions. In: *SODA, SIAM* (2005) 630–631
8. Hull, J.C.: *Options, Futures, and Other Derivatives*. Prentice Hall (2002)

9. Black, F., Scholes, M.S.: The pricing of options and corporate liabilities. *Journal of Political Economy* **81**(3) (1973) 637–54
10. Shreve, S.E.: *Stochastic calculus for finance. II.* Springer Finance. Springer-Verlag, New York (2004) Continuous-time models.
11. Cont, R.: Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance* **1** (2001) 223–236
12. Merton, R.C.: Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* **3**(1-2) (1976) 125–144
13. Cont, R., Tankov, P.: *Financial Modelling with Jump Processes.* CRC Press (2004)
14. Heston, S.L.: A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Rev. of Fin. Stud.* **6**(2) (1993) 327–343
15. DeMarzo, P., Kremer, I., Mansour, Y.: Online trading algorithms and robust option pricing. In: *Proc. of the ACM Symp. on Theory of Comp., STOC.* (2006) 477–486
16. Epstein, D., Wilmott, P.: A new model for interest rates. *International Journal of Theoretical and Applied Finance* **1**(2) (1998) 195–226
17. Korn, R.: Worst-case scenario investment for insurers. *Insurance Math. Econom.* **36**(1) (2005) 1–11
18. Yao, A.C.C.: Probabilistic computations: toward a unified measure of complexity. In: *18th Symp. on Foundations of Comp. Sci.* (1977). (1977) 222–227
19. Borodin, A., El-Yaniv, R.: *Online computation and competitive analysis.* Cambridge University Press, New York (1998)
20. Goldman, M. B., Sosin, H. B., Gatto, M. A.: Path dependent options: "buy at the low, sell at the high". *The Journal of Finance* **34**(5) (1979) 1111–1127