Optimal Algorithms for k-Search with Application in Option Pricing

Julian Lorenz[1], Konstantinos Panagiotou[2], and Angelika Steger

Institute of Theoretical Computer Science, ETH Zurich, 8092 Zurich, Switzerland
{jlorenz,panagiok,steger}@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] This work was partially supported by UBS AG
[2] This work was 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 $\mathsf{ALG}(\sigma)$ and $\mathsf{OPT}(\sigma)$ 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

$$\mathrm{CR}(\mathsf{ALG}) = \max\left\{ \frac{\mathsf{OPT}(\sigma)}{\mathsf{ALG}(\sigma)} \;\middle|\; \sigma \in \mathcal{P} \right\} ,$$

and similarly, for minimization problems

$$\mathrm{CR}(\mathsf{ALG}) = \max\left\{ \frac{\mathsf{ALG}(\sigma)}{\mathsf{OPT}(\sigma)} \;\middle|\; \sigma \in \mathcal{P} \right\} .$$

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 $\mathbb{E}[\mathsf{ALG}]$ 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 $k \to \infty$ corresponding to the one-way trading problem. This connection will be made more explicit later.

Several variants of search problems 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 price sequence is generated by a known prior 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 that 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.

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 interesting 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 $\sigma \in \mathcal{P}, \sigma = (p_1, \ldots, p_n)$, let $p_{max}(\sigma) = \max_{1 \le i \le n} p_i$ denote the maximum price, and $p_{min}(\sigma) = \max_{1 \le i \le n} p_i$ 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 $\lim_{x \to \infty} f(x)/g(x) = 1$. It is well-known that $W(x) \sim \log x$.

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*

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

and there exists no deterministic algorithm with smaller competitive ratio. Furthermore, it holds

(i) For fixed $k \ge 1$ and $\varphi \to \infty$, we have $r^*(k, \varphi) \sim \sqrt[k+1]{k^k \varphi}$.
(ii) For fixed $\varphi > 1$ and $k \to \infty$, we have $r^*(k, \varphi) \sim 1 + W(\frac{\varphi-1}{e})$.

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 $1 + W(\frac{\varphi-1}{e})$ (for $n \to \infty$), which coincides with the ratio of our algorithm given by the theorem above for $k \to \infty$. 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 r^*-competitive deterministic algorithm for k-min-search, where $r^* = r^*(k, \varphi)$ is the unique solution of*

$$\frac{1 - 1/\varphi}{1 - 1/r^*} = \left(1 + \frac{1}{kr^*}\right)^k , \tag{2}$$

and there exists no deterministic algorithm with smaller competitive ratio. Furthermore, it holds

(i) For fixed $k \ge 1$ and $\varphi \to \infty$, we have $r^*(k, \varphi) \sim \sqrt{\frac{k+1}{2k}\varphi}$.
(ii) For fixed $\varphi > 1$ and $k \to \infty$, we have $r^*(k, \varphi) \sim (W(-\frac{\varphi-1}{e\varphi}) + 1)^{-1}$.

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 $\log \varphi$, 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 $\mathcal{O}(\log \varphi)$-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*

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

Furthermore, there is a $(\ln \varphi)/\ln 2$-competitive randomized algorithm.

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 $\log \varphi$, 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*

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

Again, the given lower bound is independent of k. More surprising, combined with Theorem 2, the theorem above states that for *all* $k \in \mathbb{N}$, randomization *does not* improve the performance (up to possibly a multiplicative constant) of algorithms for k-min-search, compared to deterministic algorithms. This is again a slightly unexpected 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. [7]). 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. 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 [8] 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. [9]), which then became the standard model for option pricing. While it certainly shows many features that fairly resemble reality, 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 [10], 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 [11–13]).

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 [14] 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 [15] propose non-probabilistic models for pricing interest rate securities in a framework of "worst-case scenarios". Korn [16] 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 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 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 works through the reservation prices in a strictly sequential manner. Note that RPP 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.

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 *with reservation prices given by*

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

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

Proof. For $0 \le j \le k$, let $\mathcal{P}_j \subseteq \mathcal{P}$ be the sets of price sequences for which RPP 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

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

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

$$\frac{\mathsf{OPT}(\sigma)}{\mathsf{RPP}(\sigma)} \le \frac{k p_{\max}(\sigma)}{\mathsf{RPP}(\sigma)} \ \le \ \frac{k p_{j+1}^*}{\mathsf{RPP}(\sigma_j)} \ . \tag{6}$$

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

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

But then we have for $\varepsilon \to 0$,

$$\forall\, 0 \le j \le k : \quad \frac{k p_{j+1}^*}{\mathsf{RPP}(\sigma_j)} = \frac{k p_{j+1}^*}{\sum_{i=1}^{j} p_i^* + (k - j)m} = r^* \ .$$

Thus, from (6) the r^*-competitiveness of RPP follows immediately. ☐

The competitive ratio (1) is based on the worst-case assumption. Thus RPP tends to be overly pessimistic on non-worst-case sequences. In practical situations more "lenient" sequences might be presented. In Appendix B, we describe an algorithm that always performs as well as RPP when given the worst-case sequences, but will strictly improve on non worst-case sequences.

By considering the choice of reservation prices in Lemma 1, we see that in fact no deterministic algorithm will be able to do better than RPP 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.*

By combining Lemma 1 and Lemma 2 we immediately obtain the first part of Theorem 1. The proofs of Lemma 2 and the statements about the asymptotic behavior of the competitive ratio are deferred to Appendix A.

Similarly, we can construct a reservation price policy RPP for k-min-search. Naturally, RPP 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 $r^* = r^*(k, \varphi)$ be defined as in* (2). *Then the reservation price policy* RPP *with reservation prices $p_1^* > \cdots > p_k^*$,*

$$p_i^* = M\left[1 - \left(1 - \frac{1}{r^*}\right)\left(1 + \frac{1}{kr^*}\right)^{i-1}\right] \;, \tag{7}$$

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

Again, no deterministic algorithm can do better than RPP in Lemma 3.

Lemma 4. *Let $k \geq 1$, $\varphi > 1$. Then $r^*(k, \varphi)$ given by* (7) *is the lowest possible competitive ratio that a deterministic k-min-search algorithm can achieve.*

The proofs of Lemma 3 and 4 and the asymptotic behavior of r^* in Theorem 2 can be found in Appendix A.

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 $\mathsf{RPP}(j)$ be the reservation price policy with reservation $m2^j$, and define EXPO to be a uniform probability mixture over $\{\mathsf{RPP}(j)\}_{j=0}^{\ell-1}$.

Lemma 5 (Levin, see [2]). *Algorithm* EXPO *is $\mathcal{O}(\log \varphi)$-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* [17].

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,*

$$\mathrm{CR}(\mathsf{RALG}) = \max_{\sigma\in\mathcal{S}} \frac{\mathsf{OPT}(\sigma)}{\mathbb{E}[\mathsf{RALG}(\sigma)]} \geq \left(\max_{\mathsf{ALG}\in\mathcal{A}} \mathbb{E}\left[\frac{\mathsf{ALG}(\mathrm{S})}{\mathsf{OPT}(\mathrm{S})}\right]\right)^{-1} . \tag{8}$$

The proof of Theorem 5 is given in Appendix C.1 for completeness. In words, Yao's principle says 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 (8) gives a lower bound for *arbitrary* chosen input distributions. However, only for well-chosen y's we will obtain strong lower bounds.

We first 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.*

The proof is in Appendix C.2. The next lemma yields the desired lower bound.

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

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

Proof. Let $b > 1$ and $\ell = \log_b \varphi$. We define a finite approximation of $\mathcal{I}$ by $\mathcal{I}_b = \{mb^i \mid i = 0\ldots\ell\}$, 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

$$\sigma_i = mb^0, \ldots, mb^i, m, \ldots, m . \tag{9}$$

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 $\mathsf{ALG} \in \mathcal{A}_b$. 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 $0 \leq \chi(\mathsf{ALG}) \leq \ell-1$ be such that ALG accepts the $(\chi(\mathsf{ALG})+1)$-th price, i.e. $mb^{\chi(\mathsf{ALG})}$, in $\sigma_{\ell-1}$. ALG will then earn $m\,b^{\chi(\mathsf{ALG})}$ on all σ_j with $j \geq \chi(\mathsf{ALG})$, and m on all σ_j with $j < \chi(\mathsf{ALG})$. To shorten notation, we write χ instead of $\chi(\mathsf{ALG})$ in the following. Thus, we have

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

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,

$$\max_{\mathsf{ALG} \in \mathcal{A}_b} \mathbb{E}\left[\frac{\mathsf{ALG}}{\mathsf{OPT}}\right] \leq \frac{1}{\ell}\left(1 - \frac{1}{\sqrt{\varphi}}\right)\frac{2b}{b-1} \leq \frac{1}{\ln\varphi} \cdot \frac{2b\ln b}{b-1} \ . \tag{10}$$

Let Υ_b be the set of all randomized algorithms for 1-max-search with possible price sequences $\mathcal{P}_b$. By Lemma 6, each $\mathsf{RALG}_b \in \Upsilon_b$ 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 $\mathsf{RALG}_b \in \Upsilon_b$, we have

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

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

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

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

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

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 $\{0, \ldots, \ell-1\}$, 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* $\mathrm{CR}(\mathsf{EXPO}_k) = \ln\varphi/\ln 2$ *as $\varphi \to \infty$.*

The proofs of Lemma 8 and Lemma 9 can be found in Appendix C.4.

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 Appendix C.4).

Proof (Theorem 4). We only give the proof for $k = 1$. Let $\mathcal{S} = \{\sigma_1, \sigma_2\}$ with

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

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 $y(\sigma)$ must be either ALG_1 or ALG_2. Since

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

the desired lower bound follows from the min-cost version of Yao's principle, see Theorem 7 in Appendix C.4. The proof for general $k \geq 1$ is straightforward by repeating the prices $m\sqrt{\varphi}$ and m in σ_1 and σ_2 in blocks of k times. □

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. At each price quotation S_t, only one unit of stock (or one lot of constant size) can be traded. In reality, certainly the amount of shares tradeable varies, but we can model this by multiple quotations at the same price. We assume that a lookback option is on a fixed number $k \geq 1$ of shares. Recall that the holder of a lookback call has the right to buy shares from the option writer for the price $S_{min} = \min\{S_t \,|\, 0 \leq t \leq T\}$. Neglecting stock price appreciation, upwards and downwards movement is equally probably. Consequently, we 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 6. *Assume $(S_t)_{0 \leq t \leq T}$ is a (S_0, φ) stock price path. Let $r^*(k, \varphi)$ be given by* (2)*, and let*

$$V^*_{Call}(k, S_0, \varphi) = kS_0(r^*(k, \varphi) - 1)/\sqrt{\varphi} \ . \tag{11}$$

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, applying RPP for k-min-search with reservation prices as given by (7). Let H be the total sum of money spent for purchasing k units of stock. By Lemma 3 we have $H \leq kr^*(k, \varphi)S_{min}$. At time T the option holder executes her option, buying k shares from the option writer for kS_{min} in cash. Thus, after everything has been settled, we have $C_T = V - H + kS_{min} \geq V + kS_{min}(1 - r^*(k, \varphi))$. Because of $S_{min} \geq S_0/\sqrt{\varphi}$ and $V > V^*_{\text{Call}}(k, S_0, \varphi)$, 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. □

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 $V \leq V^*_{Call}(k, S_0, \varphi)$, with $V^*_{Call}(k, S_0, \varphi)$ as defined in* (11).

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.

4.2 Comparison to Pricing in Black-Scholes Model

Goldman, Sosin and Gatto [18] 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

$$V_{\text{Call}}^{\text{BS}}(S_0, T, \sigma) = S_0\big(2\,\Phi(\sigma\sqrt{T}/2) - 1\big) \ . \tag{12}$$

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 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 (12), we choose the fluctuation rate $\varphi = \varphi(T)$ such that the expected trading range $\big[\mathbb{E}(\min_{0\le t\le T} S_t),\, \mathbb{E}(\max_{0\le t\le T} S_t)\big]$ of a geometric Brownian motion starting at S_0 with volatility σ is $\big[\varphi^{-1/2} S_0, \varphi^{1/2} S_0\big]$. Figure 1 shows the results for $\sigma = 0.2$, $S_0 = 20$ and $k = 10$.

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. Hull, J.C.: Options, Futures, and Other Derivatives. Prentice Hall (2002)

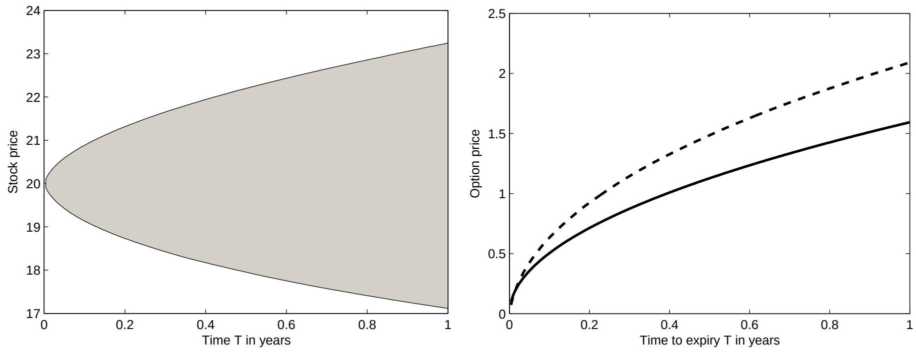

Fig. 1. The left plot shows the expected trading range of a geometric Brownian motion with volatility $\sigma = 0.2$ and $S(0) = 20$. The right plot shows the price of a lookback call with maturity T in the Black-Scholes model (solid line) and the bound V^*_{Call} (dashed line), with $k = 100$ and $\varphi(T)$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

8. Black, F., Scholes, M.S.: The pricing of options and corporate liabilities. Journal of Political Economy **81**(3) (1973) 637–54
9. Shreve, S.E.: Stochastic calculus for finance. II. Springer Finance. Springer-Verlag, New York (2004) Continuous-time models.
10. Cont, R.: Empirical properties of asset returns: stylized facts and statistical issues. Quantitative Finance **1** (2001) 223–236
11. Merton, R.C.: Option pricing when underlying stock returns are discontinuous. Journal of Financial Economics **3**(1-2) (1976) 125–144
12. Cont, R., Tankov, P.: Financial Modelling with Jump Processes. CRC Press (2004)
13. 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
14. 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
15. Epstein, D., Wilmott, P.: A new model for interest rates. International Journal of Theoretical and Applied Finance **1**(2) (1998) 195–226
16. Korn, R.: Worst-case scenario investment for insurers. Insurance Math. Econom. **36**(1) (2005) 1–11
17. Yao, A.C.C.: Probabilistic computations: toward a unified measure of complexity. (In: 18th Symp. on Foundations of Comp. Sci. (1977)) 222–227
18. 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
19. Borodin, A., El-Yaniv, R.: Online computation and competitive analysis. Cambridge University Press, New York (1998)
20. Loomis, L.H.: On a theorem of von Neumann. Proc. Nat. Acad. Sci. U.S.A. **32** (1946) 213–215

Appendix

A Proofs from Section 2

A.1 Proofs for k-max-search

Proof (of Lemma 2). Let ALG be any deterministic algorithm. We shall show that ALG cannot achieve a ratio lower than $r^*(k, \varphi)$. Let $p_1^*, \ldots, p_k^*$ 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 have 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 $p_1^*, \ldots, p_k^*$, we conclude with k times M. In any case, ALG achieves only a ratio of at most $r^*(k, \varphi)$. □

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

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

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 $\varphi \to \infty$.

For the proof of (ii), note that for $k \to \infty$ and φ fixed, equation (1) becomes $(\varphi - 1)/(r^* - 1) = e^{r^*}$, and thus

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

The claim follows from the definition of the W-function. □

A.2 Proofs for k-min-search

Proof (Lemma 3). 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 accepts *exactly* j prices, excluding the forced sale at the end. Let $\varepsilon > 0$ be fixed and define the price sequences

$$\sigma_0 = \underbrace{p_1^* + \varepsilon, \ldots, p_1^* + \varepsilon}_{k}, \underbrace{M, M, \ldots, M}_{k} \ ,$$

$$\vdots$$

$$\sigma_{k-1} = p_1^*, p_2^*, \ldots, p_{k-1}^*, \underbrace{p_k^* + \varepsilon, \ldots, p_k^* + \varepsilon}_{k}, M \ ,$$

$$\sigma_k = p_1^*, p_2^*, \ldots, p_k^*, \underbrace{m, \ldots, m}_{k} \ .$$

To shorten notation, define $p_{k+1}^* = m$. As $\varepsilon \to 0$, each σ_j is a sequence yielding the worst-case ratio in $\mathcal{P}_j$, in the sense that for all $\sigma \in \mathcal{P}_j$, $0 \leq j \leq k$,

$$\frac{\mathsf{RPP}(\sigma)}{\mathsf{OPT}(\sigma)} \leq \frac{\mathsf{RPP}(\sigma)}{kp_{\min}(\sigma)} \leq \frac{\mathsf{RPP}(\sigma_j)}{kp_{j+1}^*} \ . \tag{13}$$

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

$$\forall\, 0 \leq j \leq k: \quad \frac{\mathsf{RPP}(\sigma_j)}{kp_{j+1}^*} = \frac{\sum_{i=1}^{j} p_i^* + (k-j)M}{kp_{j+1}^*} = r^* \ ,$$

and hence

$$\forall\, \sigma \in \mathcal{P}: \quad \frac{\mathsf{RPP}(\sigma)}{kp_{\min}(\sigma)} \leq r^* \ .$$

Since $\mathsf{OPT}(\sigma) \geq kp_{\min}(\sigma)$ for all $\sigma \in \mathcal{P}$, this also implies that RPP is r^*-competitive. □

Proof (Lemma 4). The proof is similar to the proof of Lemma 2. Let ALG be any deterministic algorithm for k-min-search. We shall show that ALG cannot achieve a ratio lower than $r^*(k, \varphi)$. Let $p_1^*, \ldots, p_k^*$ be the price sequence defined by (7). We start by presenting p_1^* to ALG, at most k times or until ALG accepts it. If ALG never accepts p_1^*, we set the price to M for the remainder, and ALG achieves a competitive ratio of $M/p_1^* = r^*(k, \varphi)$. 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 have presented it k times, we raise to M for the remainder and ALG achieves a ratio not lower than $(p_1^* + (k-1)M)/(kp_2^*) = r^*(k, \varphi)$. 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 raise the price to M for the remainder. If ALG subsequentially accepts all $p_1^*, \ldots, p_k^*$, we conclude with k times m. In any case, ALG achieves only a ratio of at most $r^*(k, \varphi)$. □

Proof (of Theorem 2). The first part follows directly from Lemma 3 and Lemma 4. To show (i), first observe that for $k \geq 1$ fixed, $r^* = r^*(\varphi)$ must satisfy $r^* \to \infty$

as $\varphi \to \infty$, and r^* is an increasing function of φ. With this assumption we can expand the right-hand side of (2) with the binomial theorem to obtain

$$\frac{1-1/\varphi}{1-1/r^*} = 1 + \frac{1}{r^*} + \frac{k-1}{2k\,(r^*)^2} + \Theta\left((r^*)^{-3}\right) \implies \frac{1}{\varphi} = \frac{k+1}{2k\,(r^*)^2} + \Theta\left((r^*)^{-3}\right) \quad .$$

By solving this equation for r^*, we obtain (i).

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

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

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

$$\frac{1-1/\varphi}{1-1/r^*} = \lim_{k\to\infty}\left(1 + \frac{1}{kr^*}\right)^k = e^{1/r^*} \quad ,$$

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

B Adapting against a Lenient Adversary for k-max-search

If RPP is presented favorable prices in the beginning, it will achieve a competitive ratio strictly better than the worst-case ratio, even if it leaves the remaining $k-\ell$ reservations unchanged. However, it is even better off adjusting them.

Lemma 10. *Suppose that an algorithm for the k-max-search problem has already sold $\ell < k$ units and has captured D units of money. Let $N = \frac{D+(k-\ell)m}{km}$. Then every optimal online algorithm for the k-max-search problem has competitive ratio at least $\hat{r}^* = \hat{r}^*(k, \varphi, D)$, given by*

$$\varphi = 1 + (\hat{r}^* N - 1)\left(1 + \frac{\hat{r}^*}{k}\right)^{k-\ell} \quad . \tag{14}$$

An optimal algorithm is RPP with $k-\ell$ reservation prices $\hat{p}^*_{\ell+1}, \ldots, \hat{p}^*_k$ given by

$$\hat{p}^*_i = m\left[1 + (\hat{r}^* N - 1)\left(1 + \frac{\hat{r}^*}{k}\right)^{i-\ell-1}\right] \quad . \tag{15}$$

Proof. The proof goes along the same lines as the proof of Lemma 1. We have

$$\forall\, \ell \leq j \leq k-1: \;\; \frac{k\hat{p}^*_{j+1}}{D + \sum_{i=\ell+1}^{j} \hat{p}^*_i + (k-j)m} \leq r^*, \text{ and } \frac{kM}{D + \sum_{i=\ell+1}^{k} \hat{p}^*_i} \leq r^*,$$

and it is again straightforward to check that when choosing the reservation prices by (15), each of these equations is satisfied with equality. □

That is, against an adversary that deviates from the worst-case, the algorithm will deviate from the reservation prices stated in Lemma 2. It is easy to see that it will in fact *lower* them. Against the worst-case sequence, it has to stick to the original reservation prices and achieves the ratio given by (1).

C Proofs from Section 3

C.1 Proof of Yao's principle

Competitive analysis of an online problem can be viewed as the analysis of a two-person zero-sum game. The reader is referred to standard textbooks for further details (e.g. chapter 6 and 8 in [19]). A finite two-person zero-sum game in strategic form can be represented by a payoff matrix $G = (h_{ij}) \in \mathbb{R}^{m,n}$. The row index i enumerates the pure strategies for the first player, and the column index j enumerates the pure strategies of the second player. If the first player chooses i and the second player chooses j, the payoff for the first player is h_{ij}, and since the game is zero-sum, the payoff for the second player is $-h_{ij}$. *Mixed strategies* are probability distributions on the set of pure strategies, written as $x = (x_1, x_2, \ldots, x_m)$ and $y = (y_1, y_2, \ldots, y_n)$ for players 1 and 2, respectively. We define

$$H(x,y) = \sum_{i=1}^{m} \sum_{j=1}^{n} h_{ij}\; x_i y_j \ , \tag{16}$$

the expected payoff for player 1 under x and y. The game with $H(x,y)$ as payoff function is referred to as the mixed extension of G.

Yao's principle relies on the following corollary of the famous von Neumann Minimax Theorem, known as Loomis' lemma [20].

Lemma 11 (Loomis' lemma). *Let H be the payoff matrix of a finite two-person zero-sum game. Let $e_i = (\delta_{ij})_j$ be the i-th unit vector, representing the i-th pure strategy. For all mixed strategies x of player 1 and all mixed strategies y of player 2, we have*

$$\max_i H(e_i, y) \geq \min_j H(x, e_j) \ . \tag{17}$$

Now we are ready to prove Yao's principle.

Proof (of Theorem 5). We follow the proof as presented in [19]. Define the two-person zero-sum game by the payoff matrix

$$h_{ij} = \frac{\mathsf{ALG}_i(\sigma_j)}{\mathsf{OPT}(\sigma_j)} \qquad \text{for } 1 \leq i \leq m \text{ and } 1 \leq j \leq n \ , \tag{18}$$

where the online player is player 1 (who chooses i) and is the maximizer. Let $\mathbb{E}_y$ denote the expectation with respect to the probability distribution $y(\sigma)$ on $\mathcal{S}$. From Loomis' inequality (17), we have

$$\max_{1\leq i\leq m} \mathbb{E}_y \left[\frac{\mathsf{ALG}_i}{\mathsf{OPT}}\right] \geq \min_{1\leq j\leq n} \mathbb{E}\left[\frac{\mathsf{RALG}(\sigma_j)}{\mathsf{OPT}(\sigma_j)}\right] = \left(\max_{1\leq j\leq n} \frac{\mathsf{OPT}(\sigma_j)}{\mathbb{E}[\mathsf{RALG}(\sigma_j)]}\right)^{-1} \ ,$$

where $\mathbb{E}[\cdot]$ without subscript is with respect to the randomization of the algorithm RALG, given by the distribution $x(a)$ on $\mathcal{A}$. Thus, we have

$$\max_{1\leq i\leq m} \mathbb{E}_y \left[\frac{\mathsf{ALG}_i}{\mathsf{OPT}}\right] \geq \frac{1}{\mathrm{CR}(\mathsf{RALG})} \ .$$

□

For cost minimzation problems, the following analogous theorem can be established. The reader is referred to [19] for a proof.

Theorem 7 (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 $\mathcal{A} = \{\mathsf{ALG}_1, \dots, \mathsf{ALG}_m\}$. 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,*

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

C.2 Proof of Lemma 6

Proof (of Lemma 6). 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. □

C.3 Proof of Lemma 8

Proof (of Lemma 8). 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

$$\sigma_i = \underbrace{mb^0, \dots, mb^0}_{k}, \dots, \underbrace{mb^i, \dots, mb^i}_{k}, \underbrace{m, \dots, m}_{k(\ell-i-1)} . \tag{19}$$

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 $\mathsf{ALG} \in \mathcal{A}_b$ 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 $\sigma_{\ell-1}$, which is increasing until the very end. Note that for all $1 \leq i \leq \ell$, the first ik prices of the sequences $\{\sigma_j \mid j \geq i-1\}$ are not distinguishable up to time ik. Therefore, when presented σ_j, ALG accepts all prices mb^{χ_i} for which $\chi_i \leq j$. To abbreviate notation, let $\chi_0 = 0$ and $\chi_{k+1} = \ell$, and define $\delta_t = \chi_{t+1} - \chi_t$. Observe that on σ_j with $\chi_t \leq j \leq \chi_{t+1} - 1$, we have $\mathsf{OPT}(\sigma_j) = kmb^j$ and

$\mathsf{ALG}(\sigma_j) = (k-t)m + \sum_{s=1}^{t} mb^{\chi_s}$, i.e. ALG can successfully sell for its first t reservation prices. Taking expectation with respect to $y(\sigma)$, we have

$$\mathbb{E}\left[\frac{\mathsf{ALG}}{\mathsf{OPT}}\right] = \frac{1}{\ell}\sum_{t=0}^{k}\sum_{j=\chi_t}^{\chi_{t+1}-1}\frac{(k-t)m + \sum_{s=1}^{t} mb^{\chi_s}}{kmb^j}$$

$$= \frac{1}{\ell}\sum_{t=0}^{k}\frac{\left(k-t+\sum_{s=1}^{t} b^{\chi_s}\right)\sum_{j=0}^{\delta_t-1} b^{-j}}{kb^{\chi_t}}$$

$$= \frac{1}{\ell}\sum_{t=0}^{k}\frac{\left(k-t+\sum_{s=1}^{t} b^{\chi_s}\right)(1-b^{-\delta_t})}{kb^{\chi_t}(1-b^{-1})} \ .$$

Straightforward yet tedious algebra simplifies this expression to

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

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

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

which is exactly (10) in the proof of Lemma 7. Thus, we can argue as in the remainder of the proof of Lemma 7, and let again $b \to 1$ to conclude that $\mathrm{CR}(\mathsf{RALG}) \ge (\ln\varphi)/2$ for all randomized algorithms RALG for k-max-search.

□

C.4 Proof of Lemma 9

Proof (of Lemma 9). We want to determine

$$\mathrm{CR}(\mathsf{EXPO}_k) = \max_{\sigma\in\mathcal{P}} \frac{\mathsf{OPT}(\sigma)}{\mathbb{E}[\mathsf{EXPO}_k(\sigma)]} \ . \tag{20}$$

The maximum can only be attained for a non-decreasing σ, since a rearrangement to non-decreasing order can only increase $\mathsf{OPT}(\sigma)/\mathbb{E}[\mathsf{EXPO}_k(\sigma)]$. Obviously, the maximum price of the worst case sequence $\hat\sigma$ will be be arbitrary close to, but below the next possible reservation price, and appear at least k times. Furthermore, the ratio will only increase if always before a p_i is quoted, all possible reservation prices strictly smaller than p_i have already been triggered. That is, EXPO_k will never receive more than k times the reservation price mb^j it chose; note that it can receive less than that, if the reservation price is not met and EXPO_k is forced to sell for m at the end of the sequence. Let $p_{max} = mb^s - \varepsilon$ be the maximum of $\hat\sigma$. As noted, we can choose ε arbitrary close to zero, and therefore $\mathsf{OPT}(\hat\sigma) = kmb^s$, and

$$\mathbb{E}[\mathsf{EXPO}_k(\hat\sigma)] = \frac{1}{\ell}\Big(\sum_{j=0}^{s-1} kmb^j + (\ell-s)km\Big) \ .$$

Thus, we can recast (20) as

$$\mathrm{CR}(\mathsf{EXPO}_k) = \max_{s=1..\ell} \frac{\ell k m b^s}{\sum_{j=0}^{s-1} k m b^j + (\ell - s) k m} \ ,$$

and the maximum is attained for $s = \ell$. That is, the worst case may be given by

$$\hat{\sigma} = \underbrace{m, \ldots, m}_{k}, \underbrace{mb^1, \ldots, mb^1}_{k}, \underbrace{mb^2, \ldots, mb^2}_{k}, \ldots, \underbrace{mb^{\ell-1}, \ldots, mb^{\ell-1}}_{k}, \underbrace{M, \ldots, M}_{k} \ .$$

Therefore, we have

$$\mathrm{CR}(\mathsf{EXPO}_k) = \ell \frac{\varphi\,(b-1)}{\varphi - 1} = \ln \varphi \, \frac{\varphi}{\varphi - 1} \cdot \frac{(b-1)}{\ln b} \ .$$

For $\varphi \to \infty$ and $b = 2$, we get $\mathrm{CR}(\mathsf{EXPO}_k) = \varphi \ln \varphi / ((\varphi - 1) \ln 2) \to \ln \varphi / \ln 2$.
$\square$