

 Open access • Journal Article • DOI:10.1145/3331033.3331039

Recent developments in prophet inequalities — Source link

José R. Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk ...+1 more authors

Institutions: University of Chile, Massachusetts Institute of Technology, University of Bremen, Max Planck Society ...+1 more institutions

Published on: 07 May 2019 - Sigecom Exchanges (Association for Computing Machinery)

Topics: Optimal stopping

Related papers:

- [An economic view of prophet inequalities](#)
- [Automated online mechanism design and prophet inequalities](#)
- [Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers](#)
- [Posted Price Mechanisms for a Random Stream of Customers](#)
- [Multi-parameter mechanism design and sequential posted pricing](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/recent-developments-in-prophet-inequalities-3o5132lgkl>

Recent Developments in Prophet Inequalities

JOSE CORREA

Universidad de Chile

and

PATRICIO FONCEA

Massachusetts Institute of Technology

and

RUBEN HOEKSMAS

Universität Bremen

and

TIM OOSTERWIJK

Max-Planck-Institut für Informatik

and

TJARK VREDEVELD

Maastricht University

The classic prophet inequality states that, when faced with a finite sequence of non-negative independent random variables, a gambler who knows their distribution and is allowed to stop the sequence at any time, can obtain, in expectation, at least half as much reward as a prophet who knows the values of each random variable and can choose the largest one. Following this classic theorem from the 70s, many results have been obtained for several related optimal stopping problems. Moreover, the recently uncovered connection between prophet inequalities and posted price mechanisms, has given the area a new surge. We survey some new developments and highlight some compelling open problems.

Categories and Subject Descriptors: K.4.4 [**Computers and Society**]: Electronic Commerce

General Terms: Algorithms, Design, Economics

Additional Key Words and Phrases: Optimal stopping; Prophet inequalities; Posted price mechanisms; Mechanism design

1. INTRODUCTION

Optimal stopping theory is concerned with choosing the right time to take a particular action, so as to maximize the expected reward. The famous prophet inequality is a key example of a result in optimal stopping. There, a gambler faces a finite sequence of non-negative independent random variables X_1, \dots, X_n with known distributions F_i from which iteratively a prize is drawn. After seeing a prize, the gambler can either accept the prize and the game ends, or reject the prize and the next prize is presented to her. The classic result of Krengel and Sucheston, and Garling [1977; 1978], states that the gambler can obtain at least half of the expected reward that a prophet can make who knows the realizations of the prizes

Authors' addresses: correa@uchile.cl, foncea@mit.edu, hoeksma@uni-bremen.de, toosterw@mpi-inf.mpg.de, t.vredeveld@maastrichtuniversity.nl

beforehand. That is, $\sup\{\mathbb{E}[X_t] : t \text{ stopping rule}\} \geq \frac{1}{2}\mathbb{E}\{\sup_{1 \leq i \leq n} X_i\}$. Moreover, Krengel and Sucheston also showed that this bound is best possible. Remarkably, Samuel-Cahn [1984] showed that the bound of $1/2$ can be obtained by a simple threshold rule, which stops as soon as a prize is above a fixed threshold. The existence of this $1/2$ -approximation is now known as the prophet inequality, and following the above results there was a significant line of work studying this optimality gap in different settings. We refer to the survey of Hill and Kertz [1992] for more results on prophet inequalities.

In the past decade, problems related to the prophet inequality regained interest. This was initiated by Hajiaghayi et al. [2007] and then Chawla et al. [2010], who observed a close connection between prophet inequalities and online mechanisms. In particular, Chawla et al. [2010] implicitly show that the problem of revenue maximization in posted price mechanisms can be reduced to that of finding stopping rules of a related optimal stopping problem. In a posted price mechanism, a seller holds a single item to sell to a set of customers, who have independent random valuations for the item and arrive sequentially. Upon arrival of a customer, the seller offers a price as a take-it-or-leave-it offer, and the customer either takes the item at that price or simply leaves it. The goal of the seller is to find the prices that maximize her expected revenue. Quite recently, Correa, Foncea, et al. [2019] proved that the reverse direction also holds, thus establishing an equivalence between designing posted price mechanisms and finding stopping rules for optimal stopping problems. The recent survey of Lucier [2017] is a good starting point to understand the economic perspective of prophet inequalities.

Due to this regained interest, new variants of the prophet inequality setting have been studied. These include problems when the gambler has to select more than one item [Ezra et al. 2018], when the selected set has to satisfy combinatorial constraints [Kleinberg and Weinberg 2012; Düetting et al. 2017], when the underlying distributions are unknown [Azar, Kleinberg, and Weinberg 2018], and when the random variables arrive in continuous time [Allaart 2007; Kleinberg and Kleinberg 2018], among many others. In this letter however we focus on three variants of the classic prophet inequality, namely, the *order selection* prophet inequality, the *prophet secretary* version, and the *i.i.d.* prophet inequality.

In the *order selection* prophet inequality, the gambler is allowed to select the order in which she examines the random variables, and therefore it is natural to expect that the constant $1/2$ of the classic prophet inequality can be improved. Indeed, Chawla et al. [2010] improve the bound to $1 - 1/e$. In recent work, the bound was further improved first by Azar, Chiplunkar, and Kaplan [2018] to $1 - 1/e + 1/400 \approx 0.6346$, then by Beyhaghi et al. [2018] to $1 - 1/e + 0.022 \approx 0.6543$, and further to $1 - 1/e + 1/30 \approx 0.6655$ by Correa, Saona, and Ziliotto [2019]. It is worth mentioning that the bounds of Azar, Chiplunkar, and Kaplan [2018] and Correa, Saona, and Ziliotto [2019] also apply to the prophet secretary version described next.

The *prophet secretary* problem is a nice combination of the prophet inequality and the secretary problem. In this version, the random variables are shown to the gambler in random order, as in the secretary problem. The problem was first studied by Esfandiari et al. [2015] who found a bound of $1 - 1/e$. Their algorithm defines a nonincreasing sequence of n thresholds that only depend on the expecta-

tion of the maximum of the random variables, and the gambler stops whenever the currently sampled value surpasses the threshold corresponding to the current time. A remarkable feature of this algorithm is that the decision to stop is anonymous in the sense that it does not depend on the index of the sampled random variable. Later, Correa et al. [2017] proved that the same factor of $1 - 1/e$ can be obtained with a personalized but time invariant sequence of thresholds. In recent work, Ehsani et al. [2018] show that the bound of $1 - 1/e$ can even be achieved using a single threshold (randomization may be needed to break ties in some situations). Shortly after the work of Ehsani et al., the mentioned works of Azar, Chiplunkar, and Kaplan [2018] and Correa, Saona, and Ziliotto [2019] beat the bound of $1 - 1/e$, although the tight bound is still unknown.

Finally, we mention the *i.i.d.* prophet inequality, where the random variables are identically distributed. Clearly, the optimal bound in this case is not smaller than in both previous cases. For this problem, Hill and Kertz [1982] provided the family of worst possible instances from which Kertz [1986] proved the largest possible bound one could expect is $1/\beta \approx 0.7451$ (β is defined later). Regarding algorithms, Hill and Kertz also proved a bound of $1 - 1/e$, which was improved by Abolhassani et al. [2017] to 0.7380. Finally, Correa et al. [2017] proved that $1/\beta$ is a tight value.

In the remainder of this letter, we describe in more detail some of the results that have been obtained for these three variants of the problem. For the sake of completeness, we start by providing an elegant proof of the classic prophet inequality due to Kleinberg and Weinberg [2012].

2. THE PROPHET INEQUALITY

There are a number of proofs of the classic prophet inequality. Besides the original one by Krengel and Sucheston, and Garling [1978], there is an inductive proof by Hill and Kertz [1981], a proof based on the analysis of a single threshold rule by Samuel-Cahn [1984], and a very short and elegant one that the authors recently learned from Jason Hartline. Here, we have decided to include the proof due to Kleinberg and Weinberg [2012], since it has an approximate stochastic dominance flavor that will be useful when discussing the prophet secretary problem.

Consider a finite collection of nonnegative and independent random variables X_1, \dots, X_n with known distributions F_1, \dots, F_n , respectively. We will look at a very simple (and anonymous) stopping rule for the gambler, namely, pick a threshold T and accept the first value above T . Let r be the index of the first sampled value above the threshold. Thus, X_r is the reward of the gambler while $\hat{X} := \max_{1 \leq i \leq n} X_i$ is that of the prophet. We show that $\mathbb{E}(X_r) \geq \frac{1}{2} \mathbb{E}(\hat{X})$.

The proof starts by defining $p = \mathbb{P}(\hat{X} > T)$ and noting that

$$\mathbb{P}(X_r > x) \geq \begin{cases} p & x \leq T, \\ (1 - p)\mathbb{P}(\hat{X} > x) & x > T. \end{cases}$$

Indeed, the first case is trivial since the gambler stops only when seeing a value of at least T and the probability that this value exists equals p . On the other hand,

for $x > T$, by conditioning on the stopping time and using the union bound, we get

$$\begin{aligned} \mathbb{P}(X_r > x) &= \sum_{i=1}^n \mathbb{P}(X_i > x) \prod_{j < i} \mathbb{P}(X_j \leq T) \\ &\geq (1-p) \sum_{i=1}^n \mathbb{P}(X_i > x) \geq (1-p) \mathbb{P}(\hat{X} > x). \end{aligned}$$

The latter inequality allows to lower bound the reward of the gambler as follows:¹

$$\begin{aligned} \mathbb{E}(X_r) &= \int_0^\infty \mathbb{P}(X_r > x) dx \geq \int_0^T p dx + (1-p) \int_T^\infty \mathbb{P}(\hat{X} > x) dx \\ &\geq pT + (1-p) (\mathbb{E}(\hat{X}) - T). \end{aligned}$$

To conclude, observe that by either picking $T = \mathbb{E}(\hat{X})/2$ or by picking T so that $p = 1/2$, the reward of the gambler is at least half that of the prophet.

The fact that the inequality is best possible follows easily from considering just two random variables: X_1 is deterministic and always equals 1, while X_2 is $1/\varepsilon$ with probability ε and 0 with probability $1 - \varepsilon$. Clearly, the best the gambler can get (under any stopping rule) in expectation is 1, while the prophet gets $\varepsilon \cdot (1/\varepsilon) + 1 - \varepsilon = 2 - \varepsilon$. The ratio thus goes to $1/2$ as ε goes to 0.

3. ORDER SELECTION

In this section, we discuss the *order selection prophet inequality*, in which the gambler is allowed to determine the order in which she examines the random variables. In the setting of posted price mechanisms, Chawla et al. [2010] showed that a threshold rule achieves a bound of $1 - 1/e$. We present this rule in the setting of prophet inequalities, where, for simplicity reasons, we assume that the distributions are continuous and strictly increasing.

First, we denote the probability that X_i attains the maximum value by $q_i = \mathbb{P}(X_i = \hat{X})$, we let $\tau_i = F_i^{-1}(1 - q_i)$ be the threshold that X_i exceeds with probability exactly q_i , and we let $b_i = \mathbb{E}(X_i | X_i > F_i^{-1}(1 - q_i))$ denote the expected value of X_i conditioned on the event that X_i is at least τ_i . Then, we sort the random variables so that $b_1 \geq b_2 \geq \dots \geq b_n$ and accept the first variable X_i that exceeds its threshold τ_i . Letting r again be the index of the first sampled value above its threshold, the expected value obtained by this rule is equal to

$$\mathbb{E}(X_r) = \sum_i \left(\prod_{j=1}^{i-1} (1 - q_j) \cdot q_i b_i \right) = \sum_i c_i q_i b_i,$$

where $c_i = \prod_{j=1}^{i-1} (1 - q_j)$ is the probability that the sample of variable X_i is considered. This probability is decreasing, i.e., $c_1 \geq c_2 \geq \dots \geq c_n$. As $\sum_i q_i = 1$,² we have that $\sum_i c_i b_i q_i \geq (\sum_i c_i q_i)(\sum_i b_i q_i)$. Hence, we can bound the expected obtained value by $\mathbb{E}(X_r) \geq (\sum_i b_i q_i)(\sum_i c_i q_i)$.

¹Recall that $\mathbb{E}(\hat{X}) = \int_0^\infty \mathbb{P}(\hat{X} > x) dx \leq T + \int_T^\infty \mathbb{P}(\hat{X} > x) dx$

²Due to the assumption that the distributions are continuous and strictly increasing, the probability that two or more random variables simultaneously attain the maximum is 0.

Since $(X_i|X_i = \hat{X})$ is stochastically dominated by $(X_i|X_i > \tau_i)$, we have that $\mathbb{E}(\hat{X}) = \sum_i q_i \mathbb{E}(X_i|X_i = \hat{X}) \leq \sum_i b_i q_i$. On the other hand, we have that $\sum c_i q_i = 1 - \prod_j (1 - q_j) \geq 1 - 1/e$. Therefore, combining these inequalities we obtain that $\mathbb{E}(X_r) \geq (1 - 1/e)\mathbb{E}(\hat{X})$, which shows the result.

Recently, the bound of $1 - 1/e$ has been improved by Beyhaghi et al. [2018]. Like Chawla et al. [2010], Beyhaghi et al. [2018] consider the setting of posted price mechanisms. Whereas Chawla et al. [2010] compute for each customer a threshold based on the probability of winning in an optimal mechanism, Beyhaghi et al. [2018] takes the best of two threshold strategies. The first one computes a threshold for each customer based on samples from the other customers, in such a way that the probability of obtaining the good equals the probability of winning in an optimal mechanism. The second strategy, on the other hand, takes a uniform threshold for each customer. The idea behind such a best-of-two strategy is that for the instances in which the sampling strategy performs bad, a uniform threshold strategy actually performs well. To analyze the performance of their strategy, the authors use a factor-revealing LP to obtain a bound of approximately 0.6543.

4. PROPHET SECRETARY

We now discuss the prophet secretary problem, which is a natural combination of the prophet inequality problem and the well-known secretary problem. In this setting, we assume that the random variables X_i , $i = 1, \dots, n$, arrive in an order σ , chosen uniformly at random among all permutations of $[n] = \{1, \dots, n\}$. A generic threshold algorithm for this problem can be described as a set of thresholds τ_i^j , with $i, j \in [n]$, where the gambler stops the sequence at time j if sampling random variable X_i (i.e., if $\sigma(i) = j$) and the sampled value is at least τ_i^j . We say that the threshold algorithm is anonymous (as opposed to personalized) when the τ_i^j 's do not depend on i , that is, $\tau_i^j = \tau_k^j$ for all $i, j, k \in [n]$, while on the other hand we say that the threshold algorithm is time invariant if τ_i^j 's do not depend on j , that is, $\tau_i^j = \tau_i^k$ for all $i, j, k \in [n]$. Furthermore, when all τ_i^j 's are equal we say that the algorithm is a single threshold algorithm. Finally, if the thresholds of an algorithm can depend on the index of the random variable being sampled, on the value sampled, and on the time, but not on the history that has been observed, the algorithm is called nonadaptive.

This problem was introduced by Esfandiari et al. [2015], who started by showing that a single threshold cannot achieve a better bound than $\frac{1}{2}$. Then they exhibit an anonymous threshold rule that yields a bound of $1 - 1/e \approx 0.6321$. The first fact follows from the following simple example. Consider n random variables that arrive in random order, $n - 1$ of them are deterministic and always give 1, while the other gives n with probability $1/n$ and zero with probability $1 - 1/n$. Clearly, $\mathbb{E}(\hat{X}) = n \cdot 1/n + 1 \cdot (1 - 1/n) \approx 2$. Now fix a threshold T . If $T \leq 1$, the gambler (who uses T as single threshold) gets n with probability $1/n^2$, and 1 otherwise, so she gets approximately 1. If on the contrary $T > 1$ the gambler gets n with probability $1/n$, so in total she gets 1.

The bound of $1 - 1/e$ can also be achieved by a personalized but time invariant threshold algorithm [Correa et al. 2017]. The cornerstone of the analysis is the so-called Bernoulli Selection Lemma, which may be of independent interest. The

result states that given a set of non-homogeneous independent Bernoulli random variables with associated prizes, there is a subset of variables so that the expected average prize of the successes within the subset is at least a factor $1 - 1/e$ of the expected maximum prize over all random variables.

Surprisingly, Ehsani et al. [2018] show that the same bound of $1 - 1/e$ can even be obtained with a single threshold. This may appear to contradict the upper bound of $\frac{1}{2}$ of Esfandiari et al. [2015], however, the subtle issue is that they allow to break ties at random. In particular, if in the previous example the gambler sets $T = 1$ and if there is a tie, she stops with probability $1/n$. We can see that the probability that the gambler gets a nonzero value is $1 - (1 - 1/n)^n \approx 1 - 1/e$. Also, the expected reward of the gambler given that she got a nonzero value can be computed as $n \cdot 1/n + 1 \cdot (1 - 1/n) \approx 2$. Therefore, the expected reward of the gambler is approximately $2(1 - 1/e) \approx (1 - 1/e)\mathbb{E}(\hat{X})$. In general, if one can show a bound for a single threshold algorithm in the case of continuous distributions, it is straightforward to obtain the same bound for general distributions with a single threshold, but allowing to break ties at random.

To wrap-up this section, we provide a proof sketch of this single threshold result of Ehsani et al. [2018] for continuous distributions. In their algorithm, the single threshold T is simply defined as the value such that $\mathbb{P}(\hat{X} < T) = 1/e$. The proof we present here is due to Correa, Saona, and Ziliotto [2019] and uses the idea of approximate stochastic dominance. Actually, they showed a stronger result, namely that $\mathbb{P}(X_r > x) \geq (1 - 1/e)\mathbb{P}(\hat{X} > x)$ (where r is the random time at which the algorithm for the gambler stops). The key insight of the proof is to show that if $X_i > T$, then the gambler will stop with that random variable with probability at least $1 - 1/e$, i.e., $\mathbb{P}(r = \sigma(i) | X_i > T) \geq 1 - 1/e$. The intuition comes from the fact that it is unlikely that many other X_j 's are also above the threshold, though a formal proof of this statement follows from the theory of Schur-convex functions. With this result at hand, the proof follows quite simply by conditioning on the stopping time and using the union bound:

$$\begin{aligned} \mathbb{P}(X_r > x) &= \sum_{i=1}^n \mathbb{P}(X_i > x | r = \sigma(i)) \mathbb{P}(r = \sigma(i)) \\ &\geq \sum_{i=1}^n \mathbb{P}(X_i > x | r = \sigma(i)) \left(1 - \frac{1}{e}\right) \mathbb{P}(X_i > T) \\ &= \left(1 - \frac{1}{e}\right) \sum_{i=1}^n \frac{\mathbb{P}(X_i > x)}{\mathbb{P}(X_i > T)} \mathbb{P}(X_i > T) \geq \left(1 - \frac{1}{e}\right) \mathbb{P}(\hat{X} > x). \end{aligned}$$

Here we assumed that $x \geq T$, for $x < T$ the statement is actually trivial.

It is worth noting that the bound of $1 - 1/e$ is best possible for single threshold algorithms (the example above is actually tight), however, it can be further improved for more general algorithms. The first to break the barrier were Azar, Chiplunkar, and Kaplan [2018] who improved the bound to $1 - 1/e + 1/400 \approx 0.6346$, through an algorithm that relies on subtle case distinctions. Quite recently, Correa, Saona, and Ziliotto [2019] further improved the bound to 0.668, following ideas similar to those in the previous proof. They also prove an upper bound of $\sqrt{3} - 1 \approx 0.7321$ on any nonadaptive algorithm.

5. THE I.I.D. PROPHET INEQUALITY

Finally, we discuss the situation in which the random variables X_1, \dots, X_n are independent and identically distributed according to F . For this variant, Hill and Kertz [1982] characterized the worst instances as a function of n , however, they could not determine the constant that these instances imply. Kertz [1986] proved that these instances yield an upper bound of $1/\beta \approx 0.7451$, where β is the unique value solving $\int_0^1 \frac{1}{y(1-\ln(y))+(\beta-1)} dy = 1$. On the other hand, [Hill and Kertz 1982] also proved a lower bound of $1 - 1/e$, which was recently improved to 0.7380 (for large enough n) by [Abolhassani et al. 2017]. Finally, Correa et al. [2017] prove that $1/\beta$ is the tight value, thus resolving the i.i.d. prophet inequality.

The idea behind the Correa et al. [2017] algorithm is to use a *quantile* stopping rule. Rather than directly constructing a threshold rule, they find a collection of acceptance probabilities $q_1 < \dots < q_n$ which are independent of the instance (and only depend on n). Then, these q_i 's are naturally turned into thresholds by defining $\tau_i = F^{-1}(1 - q_i)$, so that the algorithm stops at time i if the sampled value X_i beats τ_i . Again, we will denote the random time at which the algorithm for the gambler stops by r .

The key step of the algorithm is therefore the choice of q_1, \dots, q_n . This choice is driven by two simple observations. First, note that if we use quantile q at any given point in time, the expected reward is given by $R(q) = \mathbb{E}(X|X > F^{-1}(1 - q))q = \int_0^q F^{-1}(1 - \theta)d\theta$, and this allows to express the expected reward of the prophet as $\mathbb{E}(\hat{X}) = n \int_0^1 (n - 1)(1 - q)^{n-2} R(q) dq$. Second, note that $\psi(q) = (n - 1)(1 - q)^{n-2}$ integrates to 1 in the interval $[0, 1]$, so that the reward of the prophet is the expectation of the function $nR(q)$, when q is distributed according to $\Psi(q) = 1 - (1 - q)^{n-1}$.³ It is thus natural to choose the q_i 's as follows. Partition the interval $A = [0, 1]$ into n intervals $A_i = [\varepsilon_{i-1}, \varepsilon_i]$, and draw q_i at random from the interval A_i according to probability density function $\frac{\psi(q)}{\gamma_i}$, where γ_i is the normalization equal to $\gamma_i = \int_{q \in A_i} \psi(q) dq$. With this choice the reward of the gambler can be expressed as

$$\begin{aligned} \mathbb{E}(X_r) &= \sum_{i=1}^n \mathbb{E}(R(q_i)) \prod_{j=1}^{i-1} \mathbb{E}(1 - q_j) \\ &= \sum_{i=1}^n \int_{\varepsilon_{i-1}}^{\varepsilon_i} (n - 1)(1 - q)^{n-2} R(q) dq \frac{\prod_{j=1}^{i-1} \int_{\varepsilon_{j-1}}^{\varepsilon_j} \psi(q)(1 - q) dq}{\prod_{j=1}^i \gamma_j} \\ &= \sum_{i=1}^n \rho_i \int_{\varepsilon_{i-1}}^{\varepsilon_i} (n - 1)(1 - q)^{n-2} R(q) dq, \end{aligned}$$

where $\rho_1 = \frac{1}{\gamma_1}$ and $\rho_{i+1} = \frac{\rho_i}{\gamma_{i+1}} \int_{\varepsilon_{i-1}}^{\varepsilon_i} \psi(q)(1 - q) dq$ for $i = 1, \dots, n - 1$. If we finally choose the ε_i 's so that $\rho_1 = \dots = \rho_n$ we have that $\mathbb{E}(X_r) = \frac{1}{n\gamma_1} \mathbb{E}(\hat{X})$.

The final part of the proof is to show that this choice of ε_i implies that $n\gamma_1 \leq \beta$.

³Already at this point we can use the concavity of $R(q)$ and Jensen's inequality to deduce that $\mathbb{E}(\hat{X}) \leq nR(1/n)$, since $1/n$ is the expectation of a random variable with distribution Ψ . Thus by setting $q_i = 1/n$, for all i , the gambler gets a reward of $\sum_{i=0}^{n-1} R(1/n)(1 - 1/n)^i \geq (1 - 1/e)\mathbb{E}(\hat{X})$.

Although straightforward, the analysis is intricate; the imposed condition $\rho_1 = \dots = \rho_n$ translates into a recursion for ε_i 's which turns out to be quite hard to solve explicitly. However it can be analyzed tightly by approximating it through a differential equation, leading to the bound $n\gamma_1 \leq \beta$. We refer the reader to the work of Correa et al. [2017] for the details.

We remark once again that the choice of q_i 's only depend on the choice of ε_i 's, which in turn only depend on n (note that ψ does not depend on F) and thus the algorithm just described is *blind*, in that the quantiles do not depend on the particular instance at hand. One may imagine that a similar approach would work for the more general prophet secretary problem. Indeed, Correa, Saona, and Ziliotto [2019] take this path and show that *blind quantile strategies* lead to a bound of 0.668 for prophet secretary, however they also show an upper bound of 0.675 on the performance of any blind quantile strategy. Therefore, to get closer to 0.7451 in prophet secretary one needs to go beyond the approach for the i.i.d. case.

6. CONCLUSIONS AND FUTURE DIRECTIONS

In this letter, we considered several variants of the prophet inequality problem. In the classic setting the obtained bounds, both for general and for i.i.d. distributions, have been shown to be tight. However, for the order selection variant and the prophet secretary problem, there remains a gap between the proved lower and upper bounds. A remarkable open problem in the area is to find some separation between the optimal constant for the three variants discussed in this letter, or rather to show that they all coincide. Obviously, the optimal constant for the i.i.d. prophet inequality is larger than or equal to that for order selection, which in turn is larger than or equal to that for prophet secretary. Interestingly, the upper bound of 0.7451 obtained by Hill and Kertz [1982] is still the smallest upper bound known for all three problems. In particular, an interesting question is to determine the potential benefits of using adaptivity. In this context, a nonadaptive algorithm is an algorithm whose decision to stop can depend on the index of the random variable being sampled, on the value sampled, and on the time, but not on the history that has been observed. Clearly, adaptivity does not help in the i.i.d. prophet inequality, and moreover, Hill [1983] shows that it neither does in order selection prophet inequality. However, adaptivity does help in the prophet secretary problem. For this latter problem it is known that no nonadaptive algorithm for the gambler can achieve a constant better than 0.7320 [Correa, Saona, and Ziliotto 2019]. However, no upper bound better than that of the i.i.d. case is known for adaptive algorithms.

Another interesting research direction is to investigate how much knowledge of the distributions is required. The proof described in Section 2 uses full knowledge of the distributions, while Azar, Kleinberg, and Weinberg [2018] only require one sample of each distribution. On the other hand, Azar, Chiplunkar, and Kaplan [2018] only need to know $\mathbb{E}(\hat{X})$. This raises the question what bounds can be obtained when knowing only the first k moments of the maximum, or the first k moments of the random variables with or without any knowledge of the moments of the maximum. Finally, there are many extensions to combinatorial settings in which more than one random variable can be selected.

REFERENCES

- ABOLHASSANI, M., EHSANI, S., ESFANDIARI, H., HAJIAGHAYI, M., KLEINBERG, R., AND LUCIER, B. 2017. Beating $1-1/e$ for ordered prophets. In *Proceedings of 49th Annual ACM Symposium on the Theory of Computing*. STOC 2017.
- ALLAART, P. C. 2007. Prophet inequalities for i.i.d. random variables with random arrival times. *Sequential Analysis* 26, 4, 403–413.
- AZAR, P. D., KLEINBERG, R., AND WEINBERG, S. M. 2018. Prior independent mechanisms via prophet inequalities with limited information. *Games and Economic Behavior*.
- AZAR, Y., CHIPLUNKAR, A., AND KAPLAN, H. 2018. Prophet secretary: Surpassing the $1-1/e$ barrier. In *Proceedings of the 19th ACM Conference on Economics and Computation*. EC '18.
- BEYHAGHI, H., GOLREZAEI, N., LEME, R. P., PAL, M., AND SIVAN, B. 2018. Improved approximations for free-order prophets and second-price auctions. *CoRR abs/1807.03435*.
- CHAWLA, S., HARTLINE, J., MALEC, D., AND SIVAN, B. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*. STOC 2010.
- CORREA, J., FONCEA, P., HOEKSMAN, R., OOSTERWIJK, T., AND VREDEVELD, T. 2017. Posted price mechanisms for a random stream of customers. In *Proceedings of the 18th ACM Conference on Economics and Computation*. EC 2017.
- CORREA, J., FONCEA, P., PIZARRO, D., AND VERDUGO, V. 2019. From pricing to prophets, and back! *Operations Research Letters* 47, 1, 25–29.
- CORREA, J., SAONA, R., AND ZILLOTTO, B. 2019. Prophet secretary through blind strategies. In *Proceedings of the thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '19.
- DÜETTING, P., FELDMAN, M., KESSELHEIM, T., AND LUCIER, B. 2017. Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs. In *Proceedings of the 58th IEEE Annual Symposium on Foundations of Computer Science*. FOCS '17.
- EHSANI, S., HAJIAGHAYI, M., KESSELHEIM, T., AND SINGLA, S. 2018. Prophet secretary for combinatorial auctions and matroids. In *Proceedings of the 29th annual ACM-SIAM symposium on Discrete algorithms*. SODA 2018.
- ESFANDIARI, H., HAJIAGHAYI, M., LIAGHAT, V., AND MONEMIZADEH, M. 2015. Prophet secretary. In *The 23rd Annual European Symposium on Algorithms*. ESA 2015.
- EZRA, T., FELDMAN, M., AND NEHAMA, I. 2018. Prophets and secretaries with overbooking. In *Proceedings of the 19th ACM Conference on Economics and Computation*. EC '18.
- HAJIAGHAYI, M., KLEINBERG, R., AND SANDHOLM, T. 2007. Automated online mechanism design and prophet inequalities. In *Proceedings of the 22nd Conference on Artificial Intelligence*. AAAI 2007.
- HILL, T. 1983. Prophet inequalities and order selection in optimal stopping problems. *Proceedings of the American Mathematical Society* 88, 1, 131–137.
- HILL, T. AND KERTZ, R. 1982. Comparisons of stop rule and supremum expectations of i.i.d. random variables. *The Annals of Probability* 10, 2, 336–345.
- HILL, T. AND KERTZ, R. 1992. A survey of prophet inequalities in optimal stopping theory. *Contemporary Mathematics* 125, 191–207.
- HILL, T. P. AND KERTZ, R. P. 1981. Ratio comparisons of supremum and stop rule expectations. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* 56, 2 (Jun), 283–285.
- KERTZ, R. 1986. Stop rule and supremum expectations of i.i.d. random variables: A complete comparison by conjugate duality. *Journal of Multivariate Analysis* 19, 88–112.
- KLEINBERG, J. M. AND KLEINBERG, R. 2018. Delegated search approximates efficient search. In *Proceedings of the 19th ACM Conference on Economics and Computation*. EC '18.
- KLEINBERG, R. AND WEINBERG, S. M. 2012. Matroid prophet inequalities. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*. STOC '12.
- KRENGEL, U. AND SUCHESTON, L. 1977. Semiamarts and finite values. *Bull. Amer. Math. Soc.* 83, 745–747.

- KRENGEL, U. AND SUCHESTON, L. 1978. On semiamarts, amarts, and processes with finite value. *Adv. in Probability* 4, 197–266.
- LUCIER, B. 2017. An economic view of prophet inequalities. *ACM SIGecom Exchanges* 16, 1, 24–47.
- SAMUEL-CAHN, E. 1984. Comparisons of threshold stop rule and maximum for independent nonnegative random variables. *The Annals of Probability* 12, 4, 1213–1216.