CS 580: Algorithmic Game Theory
Lecture 12: Prophet Inequality, Variants and Extensions

Vasilis Livanos*

1 Recap: Prophet Inequality

In the last lecture, we saw that when we want to sell a single item and we have n buyers with valuations drawn from (potentially) different distributions $D_1, \ldots, D_n$, then the problem of maximizing revenue can be reduced to the design of a prophet inequality. We saw that there exists a $1/2$-competitive prophet inequality for the single-item setting, and that this factor $1/2$ is tight, i.e. we cannot do better for every instance.

The standard approach to obtain this factor is to set a single threshold τ and accept the first value $X_i \geq \tau$. Appropriate choices for this threshold include the median of the distribution of $\max_i X_i$ [11] as well as $1/2\, \mathbb{E}[\max_i X_i]$ [9].

Today, we will explore some generalizations of the prophet inequality setting; in particular, when we want to select more than one item, when the items arrive in random order or when the distributions $D_1, \ldots, D_n$ are all the same distribution D. In all such cases, the optimal prophet inequality can no longer be obtained by a single threshold. Instead, one needs to set a different threshold τ_i for each random variable X_i.

2 Extensions to Multiple Items

2.1 Background: Chernoff bounds

In the following section, we will need a result that provides bounds on the probability that a random variable is far from its expectation. Such results are called *concentration bounds*. You have probably seen such bounds before; simple example include Markov's and Chebyshev's inequalities. In this case, we will need a stronger result, due to Chernoff.

Theorem 1. *Let $Y_1, \ldots, Y_n$ be real-valued independent random variables such that $Y_i \in \{0, 1\}$ for all i. If $S_n = \sum_{i=1}^n Y_i$, then, for all $t > 0$*

$$\Pr[(1-t)\,\mathbb{E}[\,|S_n|\,] \leq |S_n| \leq (1+t)\,\mathbb{E}[\,|S_n|\,] \leq e^{-\frac{\mathbb{E}[\,|S_n|\,]\cdot t^2}{2+t}} + e^{-\frac{\mathbb{E}[\,|S_n|\,]\cdot t^2}{2}}.$$

2.2 Cardinality constraint k

One can generalize the classical prophet inequality setting in several different ways. A natural generalization of the single item setting is the setting where, instead of one, we can select up to k random variables, where $k \geq 1$ is a fixed constant. For linear objectives, if we select a set $S \subseteq \boldsymbol{X}$, where $|S| \leq k$, the value we obtain is

$$\mathbb{E}\left[\sum_{i \in S} X_i\right].$$

Of course, the prophet, who sees all realizations at once and thus can pick the best set of k random variables, receives value

$$\mathbb{E}\left[\max_{S:|S| \leq k} \sum_{i \in S} X_i\right].$$

* University of Illinois Urbana-Champaign, Urbana IL 61801, USA `livanos3@illinois.edu`

For this setting, Hajiaghayi, Kleinberg and Sandholm [7] showed that one can use a natural generalization of the single-item algorithm for this setting as well. Specifically, we can set a threshold T such that the expected number of random variables who have realizations above T are $k-\delta$ for some appropriately-chosen $\delta > 0$. Then, we can use standard concentration bounds, in this case a Chernoff bound, to show that the actual number of random variables with realizations above T will be between $k-2\delta$ and k, with high probability, thus achieving a $1 - \mathrm{O}\left(\sqrt{\frac{\log k}{k}}\right)$-competitive ratio.

Theorem 2. *Consider the prophet inequality setting in which one can select up to k random variables for some fixed constant $k > 0$. Then, there exists a value τ such that an algorithm that selects the first k values that exceed τ receives value V such that, with probability at least $1 - \frac{2}{k}$,*

$$V \geq \left(1 - \sqrt{\frac{8 \ln k}{k}}\right) \cdot \max_{S : |S| \leq k} \sum_{i \in S} X_i.$$

Proof. Select a threshold T such that the expected number of values $\geq T$ are $k - \delta$ for some appropriately chosen δ to be decided later. For fixed realizations, let $S_T = \{i \in [n] \mid X_i \geq T\}$, where $[n] = \{1, 2, \ldots, n\}$. In other words, we want to select T such that

$$\mathbb{E}\left[\,|S_T|\,\right] = k - \delta.$$

Since the realizations of the X_i's are independent, for an appropriately chosen δ, one can show that the number of realizations that are at least T are between $k - 2\delta$ and k, with high probability. To do this, we will use Theorem 1. Let Y_i be a random variable that is 1 if $X_i \geq T$ and 0 otherwise. Then, for $t = \frac{\delta}{k-\delta}$, we have that

$$(1-t)\,\mathbb{E}[\,|S_T|\,] = \left(1 - \frac{\delta}{k-\delta}\right)(k-\delta) = k - 2\delta,$$

$$(1+t)\,\mathbb{E}[\,|S_T|\,] = \left(1 + \frac{\delta}{k-\delta}\right)(k-\delta) = k,$$

and thus

$$\begin{aligned}
\Pr[k - 2\delta \leq |S_T| \leq k] &\leq e^{-\frac{(k-\delta)\cdot\left(\frac{\delta}{k-\delta}\right)^2}{2+\frac{\delta}{k-\delta}}} + e^{-\frac{(k-\delta)\cdot\left(\frac{\delta}{k-\delta}\right)^2}{2}} \\
&\leq \frac{2}{k} + \frac{1}{k} \\
&= \frac{3}{k}.
\end{aligned}$$

Thus, we have that $k - 2\delta \leq |S_T| \leq k$ with high probability.

Now, we have

$$\sum_{i \in S_T} X_i = \sum_{i \in S_T} T + (X_i - T) = T \cdot |S_T| + \sum_{i \in S_T} (X_i - T).$$

Since $|S_T| \geq k - 2\delta$, our revenue is at least $(k - 2\delta)T$.

Let S^* be the optimal set selected by the prophet. Then

$$OPT = \sum_{i \in S^*} X_i \leq \sum_{i \in S^*} T + (X_i - T) \leq kT + \sum_{i=1}^{n} (X_i - T),$$

Since $|S_T| \leq k$, we accepted every value that was at least T. Thus

$$\begin{aligned}
\sum_{i \in S_T} (X_i - T) = \sum_{i=1}^{n} (X_i - T) \geq OPT - kT &\geq \frac{k - 2\delta}{k}(OPT - kT) \\
&= \left(1 - \frac{2\delta}{k}\right) OPT - (k - 2\delta)T.
\end{aligned}$$

Now, for $\delta = \sqrt{2k \log k}$, we get

$$\sum_{i \in S_T} X_i \geq \left(1 - \frac{2\delta}{k}\right) OPT = \left(1 - \sqrt{\frac{8 \log k}{k}}\right) OPT.$$

Later, Alaei [1] was able to strengthen Theorem 2 and obtain the following result. His algorithm selects a threshold τ_i for each X_i adaptively, based on the realizations of $X_1, \ldots, X_{i-1}$ seen so far, instead of a fixed threshold.

Theorem 3. *Consider the prophet inequality setting in which one can select up to k random variables for some fixed constant $k > 0$. Then, there exists an algorithm that selects a set $S \subseteq \boldsymbol{X}$ such that $|S| \leq k$ and*

$$\mathbb{E}\left[\sum_{i \in S} X_i\right] \geq \left(1 - \frac{1}{\sqrt{k+3}}\right) \mathbb{E}\left[\max_{S:|S| \leq k} \sum_{i \in S} X_i\right].$$

Hajiaghayi, Kleinberg and Sandholm [7] also gave a lower bound for this setting, which (asymptotically) matches the upper bound of Theorem 3. Subsequently, Jiang, Ma and Zhang [8] solved the cardinality case completely, providing tight competitive ratios for every k instead of just asymptotically optimal. Their idea is more complicated and involves formulating the problem as an LP and showing optimality via the LP's dual.

3 Benefits of Random Order: Prophet Secretary

One may wonder whether the constraint of requiring our prophet inequality results to hold for any arrival order of the random variables, and thus also for the worst-case order, is too stringent. In practice, a grocer usually expects their customers to arrive in a uniformly random order, and this is a standard assumption in many other applications in which real-life noise affects the buyers' choices on when to arrive and renders the arrival order essentially random.

As it turns out, if we consider the standard single-item prophet inequality setting but we assume that the algorithm sees the random variables in a uniformly random order, then we can do better than $1/2$! The resulting setting is called the *prophet secretary* setting and the following result is due to Esfandiari et al [6].

Theorem 4. *Consider the prophet inequality setting with random arrival order, i.e. the prophet secretary setting. There exists an algorithm for this setting that selects a value V such that*

$$\mathbb{E}\left[V\right] \geq \left(1 - \frac{1}{e}\right) \mathbb{E}[\max_i X_i].$$

The surprising result here is that this is not tight! In fact, there has been significant progress recently on the prophet secretary problem [3,4] and the state-of-the-art competitive ratio is 0.669 (note that $1 - 1/e \approx 0.632$), due to Correa et al [4], via a multiple-threshold strategy. In the same paper, they also show an upper bound; no algorithm can achieve a competitive ratio better than $\sqrt{3} - 1 \approx 0.732$. Figuring out the tight constant in the prophet secretary setting is a very interesting open problem.

For the cardinality case of selecting at most k values under a random arrival order, very recently, Arnosti and Ma [2] gave a very nice result which completely resolves this setting. They showed that if one sets a *single threshold* T such that, on expectation, you have $k \cdot \gamma_k$ realizations above T, where $\gamma_k = 1 - e^{-k} \frac{k^k}{k!}$, then one obtains a γ_k-competitive ratio with respect to the prophet's value, and this is tight for every k among single-threshold strategies! Notice that for $k = 1$, we retrieve the known $1 - 1/e$ ratio.

Asymptotically, this displays the same behaviour as in the adversarial case since, by Stirling's approximation, we have

$$\gamma_k = 1 - e^{-k} \frac{k^k}{k!} \approx 1 - \frac{1}{\sqrt{2\pi k}},$$

thus showing that, for large enough k, the arrival order of the random variables does not matter all that much.

4 I.I.D. Prophet Inequality

Yet another variant of the prophet inequality is the *i.i.d. prophet inequality*. In this setting, we assume that all distributions are actually the same distribution D, i.e. $D_1 = D_2 = \cdots = D_n = D$. Of course, this is a stronger condition than the arrival order being random since, in the i.i.d. setting, there is no concept of arrival order, since all draws are the same. Therefore, the "arrival order" in this instance is essentially a random order, as it is equivalent to drawing n values from D, randomly permuting them and then revealing them one by one to the algorithm.

For this setting, there has been significant work only in the single-item setting and, here, we encounter something peculiar; the optimal algorithm is simple and in fact had been known for a while, yet its analysis had proven difficult. Finally, Correa et al [5] showed that the optimal algorithm yields a 0.745-competitive ratio with respect to the prophet's value.

We have already alluded that the optimal threshold algorithm for this instance sets distinct thresholds τ_i for each X_i. However, how should one think about these thresholds? What is the correct approach to reason about their value? Well, clearly, if we have reached X_n, then we should accept its value no matter what, since it is the last random variable. Therefore, the optimal threshold for X_n is $\tau_n = 0$. What about the rest?

If one ponders this question for a while, they might quickly arrive at the following realization: Since there exists only one distribution, the optimal algorithm should accept the realization of X_i if and only if it is higher than what the optimal algorithm would expect to get from the remaining random variables, i.e. $\tau_i = \mathbb{E}[ALG(X_{i+1}, \ldots, X_n)]$. In fact, this intuition turns out to be exactly correct! Let $G(i)$ denote $\mathbb{E}[ALG(X_i, \ldots, X_n)]$, for brevity, where ALG is the optimal threshold algorithm.

Lemma 1. *The optimal thresholds for the i.i.d. prophet inequality setting with random variables $X_1, X_2, \ldots, X_n$ are: $\tau_n = 0$, and for every $1 \le i < n$,*

$$\tau_i = G(i+1).$$

5 Prophet Inequalities for Cost Minimization

The problem's difficulty changes dramatically in the case where one attempts to minimize cost instead of maximizing reward. The corresponding *cost prophet inequalities* [10] find applications in procurement auctions, in which there exists one buyer attempting to purchase items from several sellers. In this setting, there exist no algorithms that guarantee any bounded approximation to the minimum cost, even for the simple case of two random variables. Constant-factor cost prophet inequalities are so far only known for special cases, like in the i.i.d. setting with the distribution being an MHR[1] distribution [10].

References

1. Alaei, S.: Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. SIAM Journal on Computing **43**(2), 930–972 (2014). https://doi.org/10.1137/120878422, https://doi.org/10.1137/120878422
2. Arnosti, N., Ma, W.: Tight guarantees for static threshold policies in the prophet secretary problem. In: Pennock, D.M., Segal, I., Seuken, S. (eds.) EC '22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 - 15, 2022. p. 242. ACM (2022). https://doi.org/10.1145/3490486.3538377, https://doi.org/10.1145/3490486.3538377
3. Azar, Y., Chiplunkar, A., Kaplan, H.: Prophet secretary: Surpassing the 1-1/e barrier. CoRR **abs/1711.01834** (2017), http://arxiv.org/abs/1711.01834
4. Correa, J., Saona, R., Ziliotto, B.: Prophet secretary through blind strategies. Mathematical Programming (08 2020). https://doi.org/10.1007/s10107-020-01544-8

[1] A distribution is said to be a *Monotone Hazard Rate (MHR)* distribution if its hazard rate function, $h(x) = \frac{f(x)}{1-F(x)}$ is monotonically increasing. The observant reader will recognize this function appearing in the definition of virtual welfare. Indeed, the virtual welfare function is $\phi(v) = v - \frac{1-F(v)}{f(v)} = v - \frac{1}{h(v)}$, and thus the optimal reserve price in the Vickrey auction, $\phi^{-1}(0)$, can also be thought of as the price p such that $p \cdot h(p) = 1$.

5. Correa, J.R., Foncea, P., Hoeksma, R., Oosterwijk, T., Vredeveld, T.: Posted price mechanisms and optimal threshold strategies for random arrivals. Math. Oper. Res. **46**(4), 1452–1478 (2021). https://doi.org/10.1287/moor.2020.1105, https://doi.org/10.1287/moor.2020.1105
6. Esfandiari, H., Hajiaghayi, M., Liaghat, V., Monemizadeh, M.: Prophet secretary. CoRR **abs/1507.01155** (2015), http://arxiv.org/abs/1507.01155
7. Hajiaghayi, M.T., Kleinberg, R., Sandholm, T.: Automated online mechanism design and prophet inequalities. In: Proceedings of the 22Nd National Conference on Artificial Intelligence - Volume 1. pp. 58–65. AAAI'07, AAAI Press (2007), http://dl.acm.org/citation.cfm?id=1619645.1619656
8. Jiang, J., Ma, W., Zhang, J.: Tight guarantees for multi-unit prophet inequalities and online stochastic knapsack. In: Naor, J.S., Buchbinder, N. (eds.) Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022. pp. 1221–1246. SIAM (2022). https://doi.org/10.1137/1.9781611977073.51, https://doi.org/10.1137/1.9781611977073.51
9. Kleinberg, R., Weinberg, S.M.: Matroid prophet inequalities. In: Proceedings of the forty-fourth annual ACM symposium on Theory of computing. pp. 123–136 (2012)
10. Livanos, V., Mehta, R.: Prophet inequalities for cost minimization (2022). https://doi.org/10.48550/ARXIV.2209.07988, https://arxiv.org/abs/2209.07988
11. Samuel-Cahn, E.: Comparison of threshold stop rules and maximum for independent nonnegative random variables. The Annals of Probability **12**(4), 1213–1216 (1984), http://www.jstor.org/stable/2243359