DS 574 Algorithmic Mechanism Design	**Lecture #15**
Fall 2022	Prof. Kira Goldner

Prophet Inequalities

You're at a casino gambling, and are offered to play the following game. Items will arrive one-by-one. As an item arrives, you see its value. You may only take a single item, and once you take an item, the game ends. A priori, you know the *distribution* of each item. At some point there will be a red item with the red distribution of values, and at some point there will be a blue item with the blue distribution of values, and so forth. However, you do not know the order of items (it is adversarial), and you do not know the exact values of the items (they are drawn from their specific distributions). Your goal is to come up with an algorithm that competes with the *prophet* who is all knowing, so knows the realization of values and the arrival order.

That is, n items will arrive in adversarial order. Item i (which is a label, not necessarily the order) has value v_i drawn from known distribution F_i. Your goal is to determine an algorithm $\textsc{Alg}$ such that the value you get from gambling competes with the prophet who always gets $\max_i v_i$. However, your competition is over the randomness of the values that are drawn, so you only have to compete with $\textsc{opt} = \mathbb{E}_{\mathbf{v}}[\max_i v_i]$.

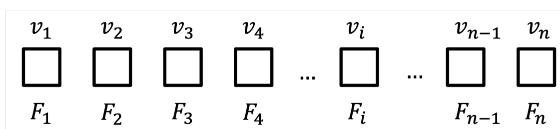

Figure 1: The prophet inequality problem.

To summarize:

- Goal: Pick one item; maximize its value.
- Gambler knows distribution for each item.
- Order is adversarial.
- Inspect each item online (see v_i) and irrevocably decide whether to take or pass forever.
- Compete with $\textsc{opt} = \mathbb{E}_{\mathbf{v}}[\max_i v_i]$.

The Prophet Inequality problem was posed by Samuel-Cahn '84 [6], with the original solution and analysis that we'll see by Krengel Sucheston '78 [4] and Garling. It was brought to Algorithmic Mechanism Design by Hajiaghayi Kleinberg Sandholm '07 [1], and a new analysis for this case was developed by Kleinberg Weinberg '12 [2, 3].

Prove the following.

Theorem 1. *There is a threshold algorithm* $\textsc{Alg}$ *such that when the gambler takes an item if and only if its value is above* T*,* $\textsc{Alg} \geq \frac{1}{2}\textsc{opt}$*.*

Determine what threshold T to use and prove this statement using the following steps:

1. Divide what the algorithm yields from an item (in expectation) into exactly the threshold and the surplus above the threshold.

2. Lower bound your surplus term.

3. Set your threshold in order to combine like-terms and have $\textsc{opt}$ pop out.

Note: Can you find two different thresholds that give this same approximation?

Proof. We consider two different ways to set the threshold, starting a proof of what our algorithm obtains using the framework above. Let p denote the probability that *some* (at least one) $v_i \geq T$ for $i \in [n]$.

We will set the threshold T such that either (1) $T = \frac{1}{2}\mathbb{E}[\max_i v_i]$, or (2) such that $p = \frac{1}{2}$.

$$\begin{aligned}
\textsc{Alg} &= \Pr[\text{any above}]T + \sum_i \Pr[\text{all } j < i \text{ below}] \cdot \mathbb{E}[(v_i - T)^+] \\
&\geq pT + (1-p)\sum_i \mathbb{E}[(v_i - T)^+] \\
&\geq pT + (1-p)\mathbb{E}[\sum_i (v_i - T)^+] \\
&\geq pT + (1-p)(\mathbb{E}[\max_i v_i] - T) \qquad (*)
\end{aligned}$$

$$\text{using (1)} \qquad \begin{aligned}
(*) &\geq p\left(\frac{1}{2}\mathbb{E}[\max_i v_i]\right) + (1-p)\left(\frac{1}{2}\mathbb{E}[\max_i v_i]\right) \\
&= \frac{1}{2}\textsc{opt}.
\end{aligned}$$

$$\text{using (2)} \qquad \begin{aligned}
(*) &= \frac{1}{2}T + \frac{1}{2}\left(\mathbb{E}[\max_i v_i] - T\right) \\
&= \frac{1}{2}\textsc{opt}.
\end{aligned}$$

Hence a threshold algorithm set using (1) or (2) produces a $\frac{1}{2}$-approximation to the prophet (a $\frac{1}{2}$-competitive ratio). □

Exercise: You could see this as a mechanism for a buyer to maximize social welfare. Could you design a mechanism to maximize revenue using the prophet inequality?

[Hint: Use virtual values.]

See Roughgarden Twenty Lectures (364A) Lecture 6 Section 3 for a formal treatment on how to do this.

I'm not aware of any textbooks on the subject, but here is a list of resources on prophet inequalities and the breadth of work in more recent research:

- 2017 Survey "An Economic View of Prophet Inequalities" by Brendan Lucier [5]: https://sigecom.org/exchanges/volume_16/1/LUCIER.pdf
- 2016 Simons Bootcamp Talks by Matt Weinberg
  - Part I: https://www.youtube.com/watch?v=NwF4Xr0-6Rc
  - Part II: https://www.youtube.com/watch?v=E19TWolvn8I
- EC 2021 Tutorial on Prophet Inequalities by Michal Feldman, Thomas Kesselheim, and Sahil Singla
  - Website (slides and reading list): http://www.thomas-kesselheim.de/tutorial-prophet-inequalities/
  - Part 1: https://www.youtube.com/watch?v=qbHd0g9RkCg
  - Part 2: https://www.youtube.com/watch?v=l20KP5IIgcQ
  - Part 3: https://www.youtube.com/watch?v=lyOUcYfNEiA

References

[1] Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *Proceedings of the 22nd National Conference on Artificial Intelligence - Volume 1*, AAAI'07, page 58?65. AAAI Press, 2007.

[2] Robert Kleinberg and S. Matthew Weinberg. Matroid prophet inequalities. In Howard J. Karloff and Toniann Pitassi, editors, *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 123–136. ACM, 2012.

[3] Robert Kleinberg and S. Matthew Weinberg. Matroid prophet inequalities and applications to multi-dimensional mechanism design. *Games Econ. Behav.*, 113:97–115, 2019.

[4] Ulrich Krengel and Louis Sucheston. Semiamarts and finite values. *Bulletin of the American Mathematical Society*, 83(4):745–747, 1977.

[5] Brendan Lucier. An economic view of prophet inequalities. *SIGecom Exch.*, 16(1):24?47, September 2017.

[6] Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent non-negative random variables. *The Annals of Probability*, pages 1213–1216, 1984.