A Prophet Inequality for the Moment-Knowledge Scenario

Stefan Lapointe
McGill University

Jiechen Zhang
McGill University

August 12, 2024

Abstract

Since the proof of the classic prophet inequality, many variations of the original game have been studied. We begin this paper by reviewing the classic prophet inequality, as well as a question about the moment-knowledge scenario asked by Jose Correa, Foncea, Hoeksma, Oosterwijk, and Vredeveld 2019. We then explore the question by proving a result for the 1-moment special case and showing how our technique can be generalized to the k-moment general case by using a family of probability distributions constructed to have certain desirable properties.

1 The Classic Problem

Given a sequence of n independent, non-negative random variables $X_1, X_2, \ldots, X_n$, each representing the value of an opportunity and drawn from a known distribution F_i for $i = 1, \ldots, n$, the decision-maker observes the realization of each X_i sequentially. The decision-maker must decide on the spot, for each X_i, whether to accept X_i and stop the process, or reject X_i and proceed to X_{i+1}, with the objective of maximizing the expected value of X_τ, where τ is the stopping time of the decision-maker. The classic prophet inequality asserts that there exists a stopping rule for the decision-maker that guarantees an expected payoff of at least $\frac{1}{2}\mathbb{E}[X^*]$, where $X^* = \max\{X_1, \ldots, X_n\}$, which represents the payoff of a prophet who knows the realization of each X_i at the start of the game.

Classic results

Theorem 1 (Krengel and Sucheston 1987, Samuel-Cahn 1984)**.** *A uniform threshold guarantees that the expected value of the gambler's choice,* $\mathbb{E}[\textit{Gambler}]$*, is at least* $\frac{1}{2}$ *of the expected maximum value,* $\mathbb{E}[\textit{Prophet}]$*. Furthermore, this is a tight bound.*

Tight Bound

The above bound was proven to be tight using a scenario consisting of just two random variables, X_1 and X_2, where $\epsilon > 0$, $X_1 = 1$, and $X_2 = (1/\epsilon)\,\text{bernoulli}(\epsilon)$.

- The prophet selects the maximum value of two outcomes,

$$X^* = \begin{cases} 1 & \text{with probability } 1-\epsilon, \\ \frac{1}{\epsilon} & \text{with probability } \epsilon. \end{cases}$$

Hence, the expected maximum value $E[\text{Prophet}]$ is calculated as:

$$E[\text{Prophet}] = 1 \cdot (1-\epsilon) + \frac{1}{\epsilon} \cdot \epsilon = 1 - \epsilon + 1 = 2 - \epsilon.$$

- The gambler can either accept the first outcome and get an expected value of 1, or reject it and accept the second outcome, which also has an expected value of 1. Therefore, the expected value for the gambler is:

$$E[\text{Gambler}] = 1.$$

- Taking $\epsilon \to 0$ shows that the bound is tight.

Moment-Knowledge Scenario

In the moment-knowledge scenario, given a sequence of n independent, non-negative random variables $X_1, X_2, \ldots, X_n$, the player knows only the first k moments of each distribution, unlike the classic scenario where the player knows each distribution in full. The notion of knowledge in this context is currently informal, but will soon be formalized.

In the final section of Jose Correa, Foncea, Hoeksma, Oosterwijk, and Vredeveld 2019, multiple open problems were proposed, including the problem of finding bounds for the moment-knowledge scenario. The remainder of this paper presents our work on that problem.

2 Preliminaries.

We begin by formalizing the notion of a gambler who knows only the values it has drawn and the first k moments of each random variable, inspired by the (k, n)-stopping rule of José Correa, Dütting, Fischer, and Schewior 2018.

For our purposes, a *stopping rule* $\mathbf{r}$ is a sequence $r_1, \ldots, r_n$ of functions $r_i : \mathbb{R}_+^i \to [0, 1]$ where $r_i(x_1, \ldots, x_i)$ represents the conditional probability of stopping on X_i given that $X_1 = x_1, \ldots, X_i = x_i$ in a prophet game with random variables $X_1, \ldots, X_n$. A k-*moment gambler* is a gambler whose choice of stopping rule depends only on the first k moments of each random variable.

Lemma 2. *Let $\mathbf{r}$ be any stopping rule. Given two prophet games, one on $X_1, \ldots, X_n$ and the other on $Y_1, \ldots, Y_n$, where X_j and Y_j are discrete and have the same first k moments for all $1 \le j \le n$ and the same distribution for all $1 \le j \le i$ for some i, let τ be the stopping time of $\mathbf{r}$ on $X_1, \ldots, X_n$ and let π be the stopping time of $\mathbf{r}$ on $Y_1, \ldots, Y_n$. Then $\mathbf{P}(\tau = i) = \mathbf{P}(\pi = i)$.*

Remark. *This formalizes the notion that the behaviour of a k-moment gambler up to a point in time can only be affected by what it has seen up to that point in time, at least when the distributions are discrete.*

Proof of Lemma 2. For $x_1, \ldots, x_i$ respectively in the supports of $X_1, \ldots, X_i$ (which are also respectively the supports of $Y_1, \ldots, Y_i$), we have

$$\mathbf{P}(\tau = i | X_1 = x_1, \ldots, X_i = x_i) = \left[\prod_{j=1}^{i-1} (1 - r_j(x_1, \ldots, x_j)) \right] \cdot r_i(x_1, \ldots, x_i)$$

and similarly

$$\mathbf{P}(\tau = i | Y_1 = x_1, \ldots, Y_i = x_i) = \left[\prod_{j=1}^{i-1}(1 - r_j(x_1, \ldots, x_j))\right] \cdot r_i(x_1, \ldots, x_i)$$

Summing over all choices of $x_1, \ldots, x_j$, we obtain

$$\begin{aligned}
\mathbf{P}(\tau = i) = \sum_{(x_1,\ldots,x_j)} \mathbf{P}((X_1, \ldots, X_i) = (x_1, \ldots, x_i)) \cdot \left[\prod_{j=1}^{i-1}(1 - r_j(x_1, \ldots, x_j))\right] \cdot r_i(x_1, \ldots, x_i) \\
= \sum_{(x_1,\ldots,x_i)} \mathbf{P}((Y_1, \ldots, Y_i) = (x_1, \ldots, x_j)) \cdot \left[\prod_{j=1}^{i-1}(1 - r_j(x_1, \ldots, x_j))\right] \cdot r_i(x_1, \ldots, x_i) \\
= \mathbf{P}(\pi = i).
\end{aligned}$$

□

Lemma 3. *For any non-negative probability mass function f whose support has a maximum m, there exists n such that given n i.i.d. f random variables $X_1, \ldots, X_n$, we have*

$$\mathbf{E}\{\max_{1 \le i \le n}\{X_i\}\} \ge m/2.$$

Proof. Let $p = f(m)$. Then $0 \le (1 - p) < 1$, so there exists n such that $(1 - p)^n \le 1/2$. Thus, for n i.i.d. f random variables $X_1, \ldots, X_n$, we have

$$\begin{aligned}
\mathbf{E}\{\max_{1 \le i \le n}\{X_i\}\} &\ge m \cdot \mathbf{P}(\max_{1 \le i \le n}\{X_i\} \ge m) \\
&= m \cdot (1 - \mathbf{P}(\max_{1 \le i \le n}\{X_i\} < m)) \\
&= m \cdot (1 - \mathbf{P}(X_1 < m)^n) \\
&= m \cdot (1 - (1 - p)^n) \\
&\ge m \cdot (1 - 1/2) \\
&= m/2.
\end{aligned}$$

□

3 Special Case: 1-Moment Gambler

We begin by restricting our focus to the 1-moment gambler case. Its proof is more intuitive than that of the general case thanks to the use of familiar Bernoulli distributions, but it acts as a stepping stone to the general case.

Theorem 4. *For any real $\epsilon > 0$ and any 1-moment gambler, there exists a prophet game $Y_1, \ldots, Y_m$ such that*

$$\mathbf{E}\{Y_\pi\} \le \epsilon \mathbf{E}\{\max_i\{Y_i\}\},$$

where π is the stopping time of that 1-moment gambler on that prophet game.

Proof. Let $N = \lceil 4/\epsilon \rceil$. We must now show that for any 1-moment gambler, there is a prophet game where the expected payoff for the prophet is at least $N/4$ times as high as that of the gambler. We begin by constructing an intermediate prophet game from which the desired prophet game will be constructed.

For $1 \le i \le N$, let f_i be the $N^i\,\text{bernoulli}(1/N^i)$ probability mass function and let m_i be a number such that the expected maximum of m_i i.i.d. f_i random variables is at least $N^i/2$, which is guaranteed to exist by Lemma 3. Now for $0 \le i \le N$ let $M_i = \sum_{j=1}^{i} m_i$. Consider the prophet game corresponding to a sequence of independent random variables $X_1, \ldots, X_{M_N}$ where for all $1 \le i \le N$ and $M_{i-1} < j \le M_i$, X_j is a random variable with probability mass function f_i. Here is how $X_1, \ldots, X_{M_N}$ looks:

$$X_1 \ldots, X_{M_1}, \ldots, \underbrace{X_{M_{i-1}+1}, \ldots, X_{M_i}}_{m_i \text{ random variables with pmf } f_i \text{, together having expected maximum } \ge N^i/2}, \ldots, X_{M_{N-1}+1}, \ldots, X_{M_N}$$

Now let g be any 1-moment gambler, let $\mathbf{r}$ be the stopping rule that g chooses for $X_1, \ldots, X_{M_N}$, let τ be the stopping time of $\mathbf{r}$ on $X_1, \ldots, X_{M_N}$, and let

$$p_i = \mathbf{P}(M_{i-1} < \tau \le M_i)$$

for all $1 \le i \le N$. Then there exists $1 \le n \le N$ such that $p_n \le 1/N$.

Now let $Y_1, \ldots, Y_{M_N}$ be a new sequence of independent random variables where $Y_1, \ldots, Y_{M_n}$ are respectively the same as $X_1, \ldots, X_{M_n}$ while $Y_{M_n+1}, \ldots, Y_{M_N}$ have distribution bernoulli(1). By definition of 1-moment gambler, g must choose the stopping rule $\mathbf{r}$ for $Y_1, \ldots, Y_{M_N}$. Let π be the stopping time of $\mathbf{r}$ on $Y_1, \ldots, Y_{M_N}$. By Lemma 2,

$$\mathbf{P}(M_{n-1} < \pi \le M_n) = \mathbf{P}(M_{n-1} < \tau \le M_n) \le 1/N.$$

Now, because every random variable that isn't among $X_{M_{n-1}+1}, \ldots, X_{M_n}$ can take on value at most N^{n-1}, we have

$$\mathbf{E}\{X_\pi\} \le \frac{1}{N} N^n + \frac{N-1}{N} N^{n-1} = 2N^{n-1}.$$

At the same time,

$$\mathbf{E}\{\max_i \{X_i\}\} \ge \frac{N^n}{2},$$

by choice of m_n using Lemma 3. As a result,

$$\mathbf{E}\{\max_i \{X_i\}\} \ge \frac{N}{4} \mathbf{E}\{X_\pi\},$$

which implies that

$$\mathbf{E}\{X_\pi\} \le \epsilon \mathbf{E}\{\max_i \{X_i\}\}$$

by definition of N. □

Remark. *Varying* x *in* $x\,\text{bernoulli}(1/x)$ *provides a way of controlling how the expected maximum grows when chaining copies of this distribution together. In the above proof, we have taken advantage of the gambler's inability to predict this growth.*

4 Existence of the Distributions Needed to Generalize.

The following lemma essentially asserts the existence of a family of distributions where the maximum of the support can be chosen to be arbitrarily high without affecting the first k moments.

Lemma 5. *For any integer $k \geq 1$, there exists a family of discrete, non-negative distributions $\{D_k(\theta) : \theta \geq 1\}$ such that for all $a \geq 1$, $D_k(a)$ and $D_k(1)$ have the same first k moments and a is the maximum of the support of $D_k(a)$.*

Remark. *For $k = 1$, $x\,\mathrm{bernoulli}(1/x)$ satisfies the properties required of $D_k(x)$. These are in fact the only properties required for the proof of Theorem 4 to work, and so we will later be able to use the existence of D_k to generalize Theorem 4.*

Proof of Lemma 5. Let $D_k(1)$ be the uniform distribution over $\{0/k, 1/k, \ldots, k/k\}$. Now let $a > 1$. It will suffice to find a satisfactory definition of $D_k(a)$. Consider the function $F : \mathbb{R} \times \mathbb{R}^k \to \mathbb{R}^k$ defined

$$F(x, \mathbf{y}) = \begin{bmatrix} a \\ a^2 \\ \vdots \\ a^k \end{bmatrix} x + M\mathbf{y}$$

where

$$M = \begin{bmatrix} (1/k)^1 & (2/k)^1 & \ldots & (k/k)^1 \\ (1/k)^2 & (2/k)^2 & \ldots & (k/k)^2 \\ \vdots & \vdots & \ddots & \vdots \\ (1/k)^k & (2/k)^k & \ldots & (k/k)^k \end{bmatrix}.$$

For suitable values of x and $\mathbf{y}$, $F(x, \mathbf{y})$ is the vector of the first k moments of a distribution on $0/k, 1/k, \ldots, k/k$, and a with weight x on a and weight y_i on i/k for all $1 \leq i \leq k$ (with the remaining weight being put on 0). Let $x_0 = 0$ and $\mathbf{y}_0 = [1/(k+1), \ldots, 1/(k+1)] \in R^k$ so that x_0 and $\mathbf{y}_0$ are the weights of $D_k(1)$ on a and $1/k, 1/k, \ldots, k/k$. Define a new function $G : \mathbb{R} \times \mathbb{R}^k \to \mathbb{R}^k$ by

$$G(x, \mathbf{y}) = F(x, \mathbf{y}) - F(x_0, \mathbf{y}_0).$$

The Jacobian $\mathbf{J}F_{\mathbf{y}}(x_0, \mathbf{y}_0)$ is M, which is a Vandermonde matrix, and hence inevitable. As a consequence of the implicit function theorem, there exist a neighbourhood U of x_0 and a continuous function $g : U \to \mathbb{R}^k$ such that $g(x_0) = \mathbf{y}_0$ and $G(x, g(x)) = \mathbf{0}$ for all $x \in U$.

There is also a neighbourhood V around $\mathbf{y}_0$ where each element's components are non-negative and sum to at most $1 - \frac{1}{2k+2}$, since the components of $\mathbf{y}_0$ are strictly positive and sum to $1 - \frac{1}{k+1}$. By continuity of g, $g^{-1}(V)$ is a neighbourhood of x_0.

Now pick $0 < p \leq \frac{1}{2k+2}$ in the intersection of U and $g^{-1}(V)$, which is a neighbourhood of $x_0 = 0$. Then p and the components of $g(p)$ are all non-negative and sum to at most 1, meaning there is a discrete, non-negative distribution $D_k(a)$, the maximum of whose support is a, defined by putting a weight of p on a, a weight of $g(p)_i$ on i/k for all $1 \leq i \leq k$, and the remaining weight on 0. By definition of g, $F(p, g(p)) - F(x_0, \mathbf{y}_0) = \mathbf{0}$, that is, $D_k(a)$ has the same first k moments as $D_k(1)$. □

Remark. *We can think of $D(a)$ intuitively as being constructed by taking $D(1)$, transferring a "sufficiently small" portion of the weight from 0 to a, and then correcting the first k moments using the k degrees of freedom afforded by the weights on $1/k, 2/k, \ldots, k/k$.*

5 General Case: k-Moment Gambler

Theorem 6. *For any real $\epsilon > 0$, any integer $k \geq 1$, and any k-moment gambler, there exists a prophet game $Y_1, \ldots, Y_m$ such that*

$$\mathbf{E}\{Y_\pi\} \leq \epsilon \mathbf{E}\{\max_i \{Y_i\}\},$$

where π is the stopping time of that k-moment gambler on that prophet game.

Proof. The proof is similar to that of Theorem 4 except that the 1-moment gambler is replaced by a k-moment gambler and $x\,\text{bernoulli}(1/x)$ is replaced by $D_k(x)$ from Lemma 5 for all x. □

6 Future Work

One critical aspect of future work could be distribution-specific analyses, where the primary focus would be to determine whether the result regarding the moment-knowledge scenario holds consistently across different families of distributions. This investigation can specifically address the differences between heavy-tailed and light-tailed distributions, analyzing how these characteristics impact the decision-making process in prophet inequalities. Additionally, the research can be extended to scenarios where the random variables are not independent in order to understand how inter-variable correlations affect the maximum achievable ratio of the player's expected gain to the expected maximum value.

Another possible direction lies in algorithmic development. Here, the emphasis would be on designing algorithms capable of computing stopping rules in real time, using only the known moments of the distributions. Such algorithms would be crucial for applications requiring quick decision-making under uncertainty. Moreover, the robustness of these algorithms could be rigorously tested against various scenarios where moment estimation errors may occur, thereby evaluating their effectiveness and practicality in diverse settings.

Finally, the problem remains of determining the tight upper bound for a k-moment gambler when the number of boxes is restricted to N, or equivalently, the minimum number of boxes required to force a given upper bound. For example, the construction in Theorem 4 requires

$$m = \sum_{n=1}^{12} \left\lceil \frac{\ln(1/2)}{\ln\left(\frac{12^n - 1}{12^n}\right)} \right\rceil = 6,742,003,513,416$$

boxes $Y_1, \ldots, Y_m$ just to force an upper bound of

$$\mathbf{E}\{Y_\pi\} \leq \frac{1}{3} \mathbf{E}\{\max_i \{Y_i\}\}$$

for the stopping time π of a given 1-moment gambler on $Y_1, \ldots, Y_m$, and we believe that dramatic improvement is likely to be possible.

7 Acknowledgements

This paper was written as a part of the Winter 2024 Directed Reading Program at McGill University. We would like to thank Caelan Atamanchuk for being an excellent mentor, Prof. Luc Devroye for many valuable conversations and the Directed Reading Program organizers for bringing this wonderful opportunity to McGill University.

References

Correa, Jose, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld (May 2019). "Recent developments in prophet inequalities". In: *SIGecom Exch.* 17.1, pp. 61–70. DOI: 10.1145/3331033.3331039. URL: https://doi.org/10.1145/3331033.3331039.

Correa, José, Paul Dütting, Felix Fischer, and Kevin Schewior (2018). "Prophet Inequalities for I.I.D. Random Variables from an Unknown Distribution". In: *arXiv preprint arXiv:1811.06114*. URL: https://arxiv.org/abs/1811.06114.

Krengel, Ulrich and Louis Sucheston (1987). "Prophet Compared to Gambler: An Inequality for Transforms of Processes". In: *The Annals of Probability* 15.4, pp. 1593–1599. URL: http://www.jstor.org/stable/2244022.

Samuel-Cahn, Ester (1984). "Comparison of Threshold Stop Rules and Maximum for Independent Nonnegative Random Variables". In: *The Annals of Probability* 12.4, pp. 1213–1216. URL: http://www.jstor.org/stable/2243359.