SINGLE THRESHOLD RULE FOR THE I.I.D. PROPHET INEQUALITY WITH FREEZE

GABRIEL BELOUZE[*]

supervisor
MARTIN HOEFER[†]

referent
CHIEN-CHUNG HUANG[‡]

Master I of Computer Science
March-July 2020

[*] École Normale Supérieure
[†] Köln Universität
[‡] École Normale Supérieure

ACKNOWLEDGEMENTS

I sincerely thank Chien-Chung Huang for guiding me through several internship subjects, and ultimately recommending me to Martin Hoefer, my internship's director.
I thank Martin for taking me in, helping me with advices and precise suggestions – even remotely – and overall for mentoring me.

CONTENTS

1 INTRODUCTION

Optimal stopping theory is concerned with finding the best policy that determines the time to take an action in a stochastic environment, in order to maximize some objective function. Historically, two canonical examples have been thoroughly studied : the secretary problem and the prophet inequality. Consider a seller with an item, and a finite sequence of bidders. Each bidder arrives with a "take-it-or-leave-it" offer to the seller, i.e. a price for the item.

secretary problem Secretary problem assumes no prior knowledge of the bidders and is concerned with choosing the one highest bid, with a strategy based only on the relative rank of the offers seen so far. It is known from Dynkin [5] that in this classical setting, the best strategy for n bidders consists in letting k_n bidders pass, then choosing the first bidder that offers more than all the previous ones. We have $\frac{k_n}{n} \to e^{-1}$ and $\mathbb{P}(\text{choosing the highest bid}) \to e^{-1}$. Variants of the secretary problem have been studied, where the number of bidders seen is stochastic [1], [13]. One particular variant of interest is when the total number of bidders is known, but only a subset of them are seen. An interpretation of this situation is that a freeze variable may stop the process of incoming bidders at a random time. This has been studied by Ester Samuel-Cahn [16], who showed results similar to the classical variant for uniform and geometric freeze.

prophet inequality In the prophet inequality setting, the seller has stochastic prior knowledge of the bidders, but the order of arrival is adversarial. The classical prophet inequality, from Krengel and Sucheston [11], states that she can obtain at least half the expected reward of a prophet who sells the same item but knows in advance the offers of each bidder – and hence chooses the highest. Interestingly, Samuel-Cahn showed that such approximation can be achieved when the seller follows a simple single threshold rule, choosing the first offer that comes above a fixed threshold [14].

Because of the wide range of applications, notably online auctions, prophet inequalities have gained a lot of interest and a number of variants have been thoroughly studied. This includes variants where the seller may choose more than one offer [2], and in particular where the set of accepted choices of bidders follows a matroid structure [10], or variants where all offers follow the same stochastic distribution [15], [8], [9], [3]. In every variant, the goal is to find a strategy that optimizes the approximation of the prophet's reward. Often however, the optimal strategy is difficult or too costly to implement, and it makes sense to study simpler strategies – albeit suboptimal.

My internship director Martin Hoefer was interested in how the introduction of a freeze in a prophet inequality setting changed the prophet's approximation ratio. Ester-Cahn has studied the setting where a freeze variable is present, and the seller has prior stochastic knowledge of the bidders, but tries only to maximize her probability of choosing the highest bid (instead of maximizing her expected reward) [17]. To our knowledge, no author have yet given a prophet's approximation ratio for a setting with freeze. The goal of the internship was hence to study what algorithms could be used in such situation, and to compute the derived ratio approximations.

To account for the freeze blocking access to part of the bids, we define a "fair" prophet who chooses the highest bid before the freeze. We also define a class of stopping rules, called "quantile rule", which, although suboptimal, are very easy to implement and allow for an easier performance analysis. In addition, finding the optimal quantile rule requires minimal knowledge of the freeze distribution. Such a rule consists in taking the first variable that comes with a value in the top q quantile of its distribution. Our main contribution is that for i.i.d. bidders, and a freeze variable F, the optimal quantile rule consists in choosing quantile $q = \frac{1}{\mathbb{E}[F]}$, which ensures an approximation ratio of the fair prophet of at least

$$\mathbb{E}[1-(1-q)^F]$$

.

In chapter 2, we present a more detailed state of the art that is closest to our setting, namely the classical prophet inequality, the i.i.d. prophet inequality and the works on freeze in optimal stopping theory. In chapter 3, we give a proof of our main result and derive some worst case adversarial behaviours for the buyers. Finally in chapter 4, some perspectives to further the work on freeze prophet inequalities are given.

2 STATE OF THE ART

2.1 THE CLASSICAL PROPHET INEQUALITY

A seller has one item to sell and looks to maximize her profit V. She sequentially gets offers from n independent buyers in a "take-it-or-leave-it" manner. In addition, the seller has prior stochastic knowledge of the buyers, that is if X_i is the i-th offer, she knows $\mathcal{D}_i$ the distribution from which X_i is sampled.

The classical prophet inequality states that the seller has a strategy that guarantees at least half the expected profit compared to that of a prophet with complete foresight on the offers. That is,

$$\sup\{\mathbb{E}_T[X_1, \ldots, X_n] \mid T \text{ stopping rule}\} \geqslant \frac{1}{2}\mathbb{E}[X_1 \vee \ldots \vee X_n]$$

where $A \vee B$ is the maximum of A and B, and a stopping rule is a strategy for the seller based only on the offers seen so far. Formally, $T : \mathbb{R}_+^n \to [1, n]$ is a stopping rule if

$$T(x_1, \ldots, x_n) = i \implies \forall y_{i+1}, \ldots, y_n,\ T(x_1, \ldots, x_i, y_{i+1}, \ldots, y_n) = i$$

We used the natural notation $\mathbb{E}_T[X_1, \ldots, X_n] \coloneqq \mathbb{E}[X_{T(X_1,\ldots,X_n)}]$.

2.1.1 *Optimal algorithm*

Let us first describe the idea for the seller's optimal strategy, which is achieved by *backward induction*. Suppose the seller has rejected the first $n-1$ offers. Then she should always take the last one, giving her $\mathbb{E}[X_n]$ as an expected reward. Now the seller should accept the $n-1$-th buyer's offer only when it outbids her expected reward if she rejects it, which is $\mathbb{E}[X_n]$; this means that having reached the $n-1$-th step, optimal behaviour gives her an expected reward of $\mathbb{E}[X_{n-1} \vee \mathbb{E}[X_n]]$. Inductively, she can thus define n thresholds $\tau_1, \ldots, \tau_n$ such that the optimal strategy is to accept the first X_i higher than τ_i.

Now let V_i be the expected reward for the seller following an optimal policy, given that she has reached step i. From what we described, we have

$$\begin{aligned} V_n &= \mathbb{E}[X_n] \\ V_i &= \mathbb{E}[X_i \vee V_{i+1}] \text{ for } i = 1, \ldots, n-1 \end{aligned}$$

Notice that these correspond to the thresholds mention earlier, namely $\tau_i = V_{i+1}$). Indeed the thresholds should act as "indifference values", where the seller has no preference between receiving τ_i or moving to step $i+1$. We see that the thresholds are decreasing. This is fairly intuitive, as the more we approach from the end, the more we should be lenient to accept smaller values.

The *value* of the optimal policy is $\mathsf{ALG}(X_1, \dots, X_n) \coloneqq V_1$. This is to be compared with the prophet's expected reward $\mathsf{OPT}(X_1, \dots, X_n) \coloneqq \mathbb{E}[X_1 \vee \dots \vee X_n]$.

THEOREM

The optimal policy achieves a $\frac{1}{2}$-approximation of the prophet, i.e.

$$\mathsf{ALG}(X_1, \dots, X_n) \geqslant \frac{1}{2}\mathsf{OPT}(X_1, \dots, X_n)$$

Proof. Let $b_i \coloneqq (X_i - V_{i+1}) \vee 0$ be the "bonus" the seller gets from the i-th buyer. Recursively, we have that $V_i = \mathbb{E}[\sum_{j \geqslant i} b_j] \leqslant \mathsf{ALG}(X_1, \dots, X_n)$ and hence $\mathsf{ALG}(X_1, \dots, X_n) = \mathbb{E}[b_1 + \dots + b_n]$. Now note that $X_i \leqslant b_i + V_{i+1}$. We can write

$$\begin{aligned}
\mathsf{OPT}(X_1, \dots, X_n) &\leqslant \mathbb{E}[\max_i (b_i + V_{i+1})] \\
&\leqslant \mathbb{E}[\max_i (b_i + \mathsf{ALG}(X_1, \dots, X_n))] \\
&= \mathsf{ALG}(X_1, \dots, X_n) + \mathbb{E}[\max_i b_i] \\
&\leqslant \mathsf{ALG}(X_1, \dots, X_n) + \mathbb{E}[\sum_i b_i] \\
&= 2\mathsf{ALG}(X_1, \dots, X_n)
\end{aligned}$$

□

This short proof comes from Bo Waggoner[19], although the inequality is known from Krengel and Sucheston [11].

In fact the approximation ratio $\frac{1}{2}$ is tight. To see this, take the case $n = 2$, and define $X_1 = 1$ almost surely, and $X_2 = 1/p$ with probability p, 0 otherwise, for some $p > 0$. The seller's expected reward is 1 regardless of her strategy, while the prophet's reward is $2 - p$. When p approaches 0, the approximation ratio comes arbitrarily close to $1/2$.

In many variants of the prophet inequality setting, the optimal policy can easily be inductively defined as above, although its performance analysis is rarely straightforward. But the optimal strategy suffers from many drawbacks (typical from optimal solution for Bellman equations). It is fairly complicated (n numbers are necessary to describe it), it can be expensive to compute, it requires good knowledge of the distributions. It is also sensible to small changes of the distributions (like changing the order of arrival, or changing the i-th distribution). For this reason, simpler algorithms are often developped. They are easier to describe, to implement and do not necessarily achieve a worst approximation. Moreover, they might be helpful to form a better understanding of the problem.

2.1.2 *A single threshold algorithm achieves a $1/2$ approximation ratio*

A single threshold algorithm consists in choosing a threshold τ and accepting the first X_i such that $X_i \geqslant \tau$. A high threshold ensures that the selected value is large, but risks rejecting every offer. A low threshold accepts an offer with high probability, but does not guarantee that the accepted offer is good. In the classical setting, the optimal approximation ratio can actually be achieved with a single threshold policy, by choosing half of $\mathrm{OPT}(X_1, \ldots, X_n)$ (Samuel-Cahn [14]). In fact, there is a range of thresholds that all guarantee a half approximation of the prophet (see for instance [19]).

2.2 THE I.I.D. PROPHET INEQUALITY

An important variant of the prophet inequality is the *i.i.d. prophet inequality*, which results from the added hypothesis that $X_1, \ldots, X_n$ are identically distributed, say with the same distribution as that of X. For this variant, the classical 0.5 approximation ratio can be significantly improved upon to $1/\beta \approx 0.7451$ where β is the unique solution to $\int_0^1 (y - y\ln(y) + \beta - 1)^{-1} \, dy = 1$.

SINGLE THRESHOLD POLICY A well known trick in mechanism design is to select a threshold τ such that $\mathbb{P}(X \geqslant \tau) = \frac{1}{n}$. That is, we only select from the top $\frac{1}{n}$ quantile of the distribution. This policy guarantees a $1 - e^{-1}$ approximation ratio.

Indeed let $u(x)$ be the expected value of X given that it takes value in its top x quantile. Consider the i-th variable X_i, which is selected with probability $\frac{1}{n}$ by the prophet. It can at most contribute $\frac{1}{n}u(\frac{1}{n})$ to $\mathrm{OPT}(X)$. Hence overall, $\mathrm{OPT}(X)$ is bounded by $u(\frac{1}{n})$. Now note that the seller accepts an offer with probability $1 - (1 - \frac{1}{n})^n \geqslant 1 - e^{-1}$. It follows that $\mathrm{ALG}(X) = (1 - (1 - \frac{1}{n})^n) u(\frac{1}{n}) \geqslant (1 - e^{-1})\mathrm{OPT}(X)$ (see [7], theorem 1.5 for more details).

OPTIMAL APPROXIMATION Hill and Kertz [8] provided a family of worst instances for X, and the sequence (a_n) of optimal approximation ratios, which come as solutions of a family of recursively defined –and fairly complicated– functions. However, they were not able to compute the constant that this sequence implies, i.e. its lower bound.

Kertz later proved that $\frac{1}{\beta}$ is an upper bound of the optimal ratio (β is defined above). In 2017 Correa and al. [3] provided an algorithm that achieve a $1/\beta$ approximation, therefore proving that $\frac{1}{\beta}$ is tight. Their idea is very basic : the less buyers remain, the less exigent we should be about the incoming offers and their algorithm thus defines n increasing quantiles $q_1, \ldots, q_n$ and accepts the first X_i that takes value in the top q_i quantile of the distribution. The way these quantiles are determined and the performance analysis are however quite tricky, and exceed the purpose of this report. Singla ([18], chapter 10) reaches the same bound by relaxing the buyer's arrival time (which is classically discrete, with value taken in $[\![1, n]\!]$) to the continuous range $\mathbb{R}_+$. Upon arrival of a buyer, his algorithm simulates a Poisson arrival process to obtain a time of arrival, and turns this time into a threshold with a well chosen function.

The ratio $\frac{1}{\beta}$ will be a benchmark to compare our results in a freeze setting.

2.3 OPTIMAL STOPPING UNDER FREEZE

The concept of freeze was introduced for the secretary problem by Samuel-Cahn [16]. She then provided a general theorem that states that a freeze in a prophet inequality setting is equivalent to a discounted reward [17]. More precisely, she proves that the seller's optimal reward is the same if there is a freeze F, or if there is no freeze, but the benefits from accepting the i-th offer is discounted by a factor of $q_i := \mathbb{P}(F \geqslant i)$. Actually, the result shown is a little stronger : she proved that the equivalence holds whether the seller gets $\phi_i(X_1, \ldots, X_i)$ in a freeze setting, or $q_i \phi_i(X_1, \ldots, X_i)$ without freeze, for any ϕ_i.

This theorem is then used to solve the problem where the seller has stochastic knowledge of the buyers (as in the prophet inequality) and tries to maximize her probability of choosing the highest (as in the secretary problem), for uniform, geometric, and Poisson freezes. This is not relevant for our problem.

3 OUR CONTRIBUTION

For clarity, we repeat here the notation introduced in chapter 2. Let $X, X_1, \ldots, X_n$ be i.i.d random positive bounded variables, with density μ_X. X_i represents the offer of the i-th bidder. Let also $\mathcal{F}_X(x) = \mathbb{P}(X \leqslant x)$ be the cumulative distribution of X. For simplicity, we will assume that $\mathcal{F}_X$ is continuous and strictly increasing, but all the results easily carry over for all $\mathcal{F}_X$.

Let F be the freeze variable, taking values in $1, \ldots, n$. $F = i$ means that the bidding process will stop after the seller receives the offer X_i. Let V be the reward of the seller (it is a random variable).

3.1 THE FAIR PROPHET

The usual prophet inequality compares the seller's reward with $OPT_{\text{no freeze}}(X) = \mathbb{E}[X_1 \vee \ldots \vee X_n]$, that is the reward of a seller with complete foresight. In a freeze setting, not every bidder will be seen, and thus $OPT_{\text{no freeze}}(X)$ does not represent anymore the reward for an optimal auction. Instead, we naturally compare the seller with a *fair* prophet who chooses the highest bid before the freeze, that is

$$OPT(X) = \mathbb{E}[X_1 \vee \ldots \vee X_F]$$

In fact, the seller cannot hope to ensure a constant approximation ratio of the unfair prophet. Indeed consider $X = 1$ with probability p, and $X = 0$ otherwise, and $F = 1$ almost surely. The optimal strategy for the seller is to always accept the first bid. She gets p in expectation. The unfair prophet on the other hand always gets 1 unless all X_i's equal 0, hence $OPT_{\text{no freeze}}(X) = 1 - (1-p)^n$. Letting p go to 0, the approximation ratio approaches $\frac{1}{n}$, which can be rendered arbitrarily small. Note that in fact $\frac{1}{n}$ is tight, as

$$ALG \geqslant \mathbb{E}[X_1] = \frac{1}{n}\mathbb{E}[\sum X_i] \geqslant \frac{1}{n}\mathbb{E}[X_1 \vee \ldots \vee X_n] = \frac{1}{n}OPT_{\text{no freeze}}(X)$$

BOUNDED APPROXIMATION HYPOTHESIS Our hope is that there exists some constant α such that for any n, F and X, the seller can ensure an α approximation of $OPT(X)$.

3.2 QUANTILE RULE

Recall that in the no freeze, i.i.d. bidders setting, the tight approximation ratio is ≈ 0.7451. The algorithms achieving this bound are clever, and their performance analysis non trivial. Our approach is instead motivated by the single $\frac{1}{n}$ quantile that achieves a $1 - e^{-1} \approx 0.632$ ratio, which is a fairly good approximation.

We thus restrict our analysis to a subset of stopping rules called *single quantile rules*. Our hope is that one can readily obtain a good bound for the optimal approximation ratio by optimizing over the set of single quantile rules.

> DEFINITION 1: THRESHOLD RULE
>
> A τ-**threshold rule** for the seller consists in taking the first value higher than τ.

> DEFINITION 2: SINGLE QUANTILE RULE
>
> Let $0 \leqslant q \leqslant 1$ be a **quantile**. The q-**quantile rule** is the τ-threshold rule where τ is chosen such that $\mathbb{P}(X \geqslant \tau) = q$, that is $\tau = \mathcal{F}_X^{-1}(1 - q)$. We note $\mathrm{ALG}_q(X)$ the corresponding reward expectation.

Note that for a large q, the seller almost always accepts an offer, but with little guarantee that this offer is good. On the other hand, if q is low, the accepted offers are always good but often the item is not sold. The difficulty lies in finding the optimal trade-off between accepting often enough and accepting high valued offers.

BREAKING TIES For non continuous distribution and some τ-threshold rule, it may happen that an offer comes exactly at the threshold, i.e. $X_i = \tau$ for some i. It is intuitive to see that all results shown for continuous distribution will carry over non-continuous ones, *if we allow the seller to break ties at random*. In fact, not breaking ties can give strictly worse approximations (see Correa and al. [4] section 4 for an example of such difference in a variant called prophet secretary).

3.3 MAIN RESULT

Theorem 1 states what the optimal quantile is for a given freeze distribution, and gives the associated approximation ratio. The proof is later detailed in section 3.4.

> THEOREM 1
>
> Let F be a freeze variable. The optimal single quantile rule chooses the quantile $q^* = \frac{1}{\mathbb{E}[F]}$. Then the following inequality holds
>
> $$\mathrm{ALG}_{q^*}(X) \geqslant \alpha_F \cdot \mathrm{OPT}(X)$$
>
> where $\alpha_F = \mathbb{E}[1 - (1 - \frac{1}{\mathbb{E}[F]})^F]$.

An important limitation of the result is that the quantile is optimized *across all possible buyer's distributions*. Against one given distribution, there might be a better specific quantile.

Note that choosing the optimal single quantile rule requires minimal knowledge of the freeze, namely its expectation. We also note that for the special case $F = n$ a.s., which is the classical no freeze setting, we do find that $\frac{1}{n}$ is the optimal quantile, and that the associated approximation ratio is $1 - (1 - \frac{1}{n})^n$ which tends to $1 - e^{-1}$ as n grows.

For uniform freeze over $[1, n]$, the approximation ratio obtained by the optimal single quantile rule approaches $\frac{1+e^{-2}}{2} \approx 0.567$ as n grows.

3.4 PROOF OF THEOREM 1

First note that because X is bounded, we can assume without loss of generality that it takes value in $[0, 1]$.

For a freeze variable F, we note

$$a_F(i) = \mathbb{P}(F = i)$$

LEMMA 1

We can express

$$\mathbb{E}(X) = 1 - \int_0^1 \mathcal{F}_X(x)dx$$

Proof.

$$\begin{aligned} \mathbb{E}(X) &= \int_0^1 x\mu_X(x)dx \\ &= [x\mathcal{F}_X]_0^1 - \int_0^1 \mathcal{F}_X(x)dx \qquad \textit{(by part integration)} \\ &= 1 - \int_0^1 \mathcal{F}_X(x)dx \end{aligned}$$

□

Let q be a quantile, and τ the associated threshold for X. We can express $\mathrm{OPT}(X)$ and $\mathrm{ALG}_q(X)$ in terms of $\mathcal{F}_X$.

LEMMA 2

We have

$$\mathrm{OPT}(X) = 1 - \sum_{r=1}^{n} a_F(r) \int_0^1 \mathcal{F}_X^r$$

Proof. This follows immediately from

$$\mathrm{OPT}(X) = \sum_{r=1}^{n} \mathbb{P}(F = r)\mathbb{E}(X_1 \vee X_2 \vee \ldots \vee X_r)$$

noting that $\mathcal{F}_{X_1 \vee \ldots \vee X_r} = \mathcal{F}_X^r$. □

LEMMA 3
We have

$$\mathrm{ALG}_q(X) = [\tau + \frac{1-\tau}{q} - \frac{1}{q} \int_\tau^1 \mathcal{F}_X] \times [1 - \sum_{r=1}^{n} a_F(r)(1-q)^r]$$

Proof. For a given realisation r of F, the probability that any value is accepted is $1-(1-q)^r$. The expected reward, given that a box is accepted, is

$$\begin{aligned}
\mathbb{E}(V \mid V > 0) &= \mathbb{E}(X \mid X \geqslant \tau) \\
&= \frac{1}{q} \int_\tau^1 x\mu_X dx \\
&= \frac{1}{q} \left([x\mathcal{F}_X]_\tau^1 - \int_\tau^1 \mathcal{F}_X \right) \\
&= \tau + \frac{1-\tau}{q} - \int_\tau^1 \frac{\mathcal{F}_X}{q}
\end{aligned}$$

Summing $\mathrm{ALG}_q(X) = \sum_{r=1}^{n} \mathbb{E}(V \mid V > 0, F = r)\mathbb{P}(V > 0 \mid F = r)\mathbb{P}(F = r)$ concludes, noting that V is independent of F given that $V > 0$. □

Finding the best quantile approximation ratio thus leads to finding

$$\begin{aligned}
\alpha &= \sup_{0<q<1} [1 - \sum_{r=1}^{n} a_F(r)(1-q)^r] \cdot \inf_X \frac{\tau + \frac{1-\tau}{q} - \frac{1}{q}\int_\tau^1 \mathcal{F}_X}{1 - \sum_{r=1}^{n} a_F(r) \int_0^1 \mathcal{F}_X^r} \\
&= \sup_{0<q<1} [1 - \sum_{r=1}^{n} a_F(r)(1-q)^r] \cdot \inf_{0<\tau<1} \inf_{\substack{X \\ \mathbb{P}(X \geqslant \tau) = q}} \frac{\tau + \frac{1-\tau}{q} - \frac{1}{q}\int_\tau^1 \mathcal{F}_X}{1 - \sum_{r=1}^{n} a_F(r) \int_0^1 \mathcal{F}_X^r}
\end{aligned}$$

Let $\phi(q,\tau) = \inf_{\substack{X \\ \mathbb{P}(X \geqslant \tau) = q}} \frac{\tau + \frac{1-\tau}{q} - \frac{1}{q}\int_\tau^1 \mathcal{F}_X}{1 - \sum_{r=1}^{n} a_F(r) \int_0^1 \mathcal{F}_X^r}$. Now how should $\mathcal{F}_X$ be chosen to reach the infimum $\phi(q,\tau)$?

LEMMA 4: FIRST RESTRICTION ON $\mathcal{F}$

Let $\mathfrak{F}$ be the set of increasing functions over $[0,1]$ that verify $\mathcal{F}(0) = 1-q$ and $\mathcal{F}(1) = 1$. We can constrain $\mathcal{F}_X$ to be arbitrarily close to 0 over $[0,\tau[$, so that

$$\phi(q,\tau) = \inf_{\mathcal{F}\in\mathfrak{F}} \frac{\tau + \frac{1-\tau}{q}(1-\int_0^1 \mathcal{F})}{1-(1-\tau)\sum_{r=1}^{n} a_F(r)\int_0^1 \mathcal{F}^r}$$

Proof. Note that the value of $\mathcal{F}_X$ over $[0,\tau[$ can be arbitrarily changed without changing $\mathrm{ALG}_q(X)$. Hence, since we should minimize the denominator, which increases with $\int_0^1 \mathcal{F}_X$, we should have $\mathcal{F}_X$ be close to 0 over $[0,\tau[$. This shows that

$$\phi(q,\tau) = \inf_{\substack{X \\ \mathbb{P}(X\geqslant\tau)=q}} \frac{\tau + \frac{1-\tau}{q} - \frac{1}{q}\int_\tau^1 \mathcal{F}_X}{1-\sum_{r=1}^{n} a_F(r)\int_\tau^1 \mathcal{F}_X^r}$$

Substituting x in the integrals, so that they are taken over $[0,1]$, yields the lemma's result. □

The interpretation of lemma 4 is that before τ, the worst case scenario is when the weight of μ_X is condensed at τ^-.

LEMMA 5: SECOND RESTRICTION ON $\mathcal{F}$

In the equation of lemma 4, $\mathcal{F}$ can be considered almost constant, meaning that

$$\phi(q,\tau) = \inf_{1-q<a<1} \left(\tau + \frac{(1-\tau)(1-a)}{q}\right) \Big/ \left(1-(1-\tau)\sum_{r=1}^{n} a_F(r)a^r\right)$$

Proof. Let $a = \int_0^1 \mathcal{F}$. By mean power theorem ([12]), we have that

$$\begin{aligned}\sum_{r=1}^{n} a_F(r)\int_0^1 \mathcal{F}^r &\geqslant \sum_{r=1}^{n} a_F(r)(\int_0^1 \mathcal{F})^r \\ &= \sum_{r=1}^{n} a_F(r)a^r\end{aligned}$$

And equality can be achieved when $\mathcal{F}$ is constant.

Now with $\int_0^1 \mathcal{F}$ fixed, $\sum_{r=1}^{n} a_F(r)\int_0^1 \mathcal{F}^r$ should be minimized to reach the infinimum $\phi(q,\tau)$, so $\mathcal{F}$ should be chosen almost constant. Since $\mathcal{F}(0)$ must be equal to $1-q$, and $\mathcal{F}(1)$ must be equal to 1, $\int_0^1 \mathcal{F}$ can be chosen to be equal to any a where $1-q < a < 1$. This yields the result. □

The interpretation of lemma 5 is that after τ, the worst case scenario is when the weight of μ_X is focused at the extremal points, i.e. τ^+ and 1

Now let $\psi(q,\tau,a) = \left(\tau + \frac{(1-\tau)(1-a)}{q}\right) \ / \ (1-(1-\tau)\sum_{r=1}^{n} a_F(r)a^r)$. We aim to find $\inf_{a,\tau}\psi(q,a,\tau)$. Note that $\tau \mapsto \psi(q,\tau,a)$ is a homographic function of τ. Its infinimum should thus be reached at $\tau = 0$ or $\tau = 1$.

Before finishing the analysis, we introduce a useful function and list some properties that it satisfies.

> DEFINITION 3
>
> Define over $]0,1]$
>
> $$W : x \longmapsto \frac{1 - \sum a_F(r)(1-x)^r}{x}$$
>
> and $W(0) = \mathbb{E}(F)$

> PROPOSITION 1
>
> W is continuous, decreasing, convex.

> PROPOSITION 2
>
> $x \mapsto xW(x)$ is continuous, increasing, concave.

Demonstrations of these proposition can be found in Appendix A.

Now we have

$$\psi(q,0,a) = \frac{1}{q \cdot W(1-a)}$$
$$\psi(q,1,a) = 1$$

And because W is decreasing

$$\inf_{\tau,a}\psi(q,\tau,a) = \min(1, \frac{1}{q\mathbb{E}(F)})$$

We can finally compute

$$\begin{aligned}\alpha &= \sup_{0<q<1} [1 - \sum_{r=1}^{n} a_F(r)(1-q)^r] \min(1, \frac{1}{q\mathbb{E}(F)}) \\ &= \sup_{0<q<1} \min(\frac{W(q)}{\mathbb{E}(F)}, qW(q))\end{aligned}$$

$q \mapsto \frac{W(q)}{\mathbb{E}(F)}$ decreases from 1 to 0, and $qW(q)$ increases from 0 to 1. The sup of the min of the two is reached when they are equal. This happens for $q = \frac{1}{\mathbb{E}(F)}$. This yields

$$\alpha = 1 - \sum_{r=1}^{n} a_F(r)(1 - \frac{1}{\mathbb{E}(F)})^r$$

Or more consisely

$$\alpha = \mathbb{E}[1 - (1 - \frac{1}{\mathbb{E}(F)})^F]$$

3.5 ADVERSARIAL STRATEGY AGAINST A SINGLE QUANTILE RULE SELLER

Looking at the proof allows us to understand what distribution for X produces the worst case scenario. Notice that lemmas 4 and 5 show that X should push its weight to the extremal points for a given constraint $\mathcal{F}_X(1-q) = \tau$, i.e. splitting its weight between τ and 1.

Indeed we can show that two-point distributions suffice to achieve the worst case scenario for a seller following a single quantile policy. The proof is similar to that of section 3.4, although quite easier.

Fix some q-quantile strategy. For some $a \geqslant 0$, define

$$X = \begin{cases} a & \text{w.p. } 1-p \\ 1 & \text{w.p. } p \end{cases}$$

We are going to show that the worst case scenario is achieved for some a and p.

Now for such non continuous distribution, it might be conceptually simpler to think of the quantile rule as a single τ threshold rule. Here without loss of generality, we can assume that the seller chooses either the threshold a or 1, as well as some $0 \leqslant s \leqslant 1$ to break ties, i.e. when a value comes equal to τ, she accepts the offer with probability s. Notice that to be consistent with the quantile, we have

- if $p \leqslant q$, $\tau = a$ and $q = p + (1-p)s$
- if $p > q$, $\tau = 1$ and $q = s$

Let us compute $ALG_q(X)$ and $OPT(X)$

$$\begin{aligned}
OPT(X) &= \sum_{r=1}^{n} a_F(r)\mathbb{E}[X_1 \vee \ldots \vee X_r] \\
&= \sum_{r=1}^{n} a_F(r)(1-(1-a)(1-p)^r) \\
&= 1-(1-a)\sum_{r=1}^{n} a_F(r)(1-p)^r \\
ALG(X,s) &= \sum_{r=1}^{n} a_F(r)\mathbb{P}(\text{an offer is accepted} \mid F = r)\mathbb{E}[V \mid V > 0] \\
&= \sum_{r=1}^{n} a_F(r)(1-(1-p)^r(1-s)^r)\mathbb{E}[V \mid V > 0] \\
&= \mathbb{E}[V \mid V > 0] \times (1 - \sum_{r=1}^{n} a_F(r)(1-q)^r)
\end{aligned}$$

where

$$\mathbb{E}[V \mid V > 0] = \begin{cases} \frac{p+a(1-p)s}{q} & \text{if } p \leqslant q \\ 1 & \text{otherwise} \end{cases}$$

Notice that in either case, $\frac{ALG_q(X)}{OPT(X)}$ is a homographic function of a. Hence we only look at the cases $a = 0$ and $a = 1$. In fact, we can only look at the case $a = 0$, as $a = 1$ is equivalent to $p = 1$ and $a = 0$. Hence

$$\begin{aligned} \frac{ALG_q(X)}{OPT(X)} &= \min(1, \frac{p}{q}) \times \frac{1 - \sum_{r=1}^{n} a_F(r)(1-q)^r}{1 - \sum_{r=1}^{n} a_F(r)(1-p)^r} \\ &= \min(\frac{W_n(q)}{W_n(p)}, \frac{qW_n(q)}{pW_n(p)}) \end{aligned}$$

Depending on q, the worst case scenario happens when $p = 0$ or $p = 1$, which gives the ratio

$$\min(\frac{W_n(q)}{\mathbb{E}[F]}, qW_n(q))$$

which is lower or equal than α_F.

3.6 single threshold rule does not validate the bounded approximation hypothesis

Unfortunately, Theorem 1 does not allow us to validate the bounded approximation hypothesis. Indeed, the family α_F is not lower bounded by a strictly positive constant. Indeed consider the freeze

$$F = \begin{cases} 1 & \text{w.p. } 1-t \\ n & \text{w.p. } t \end{cases}$$

for some large n and small t. It can render α_F arbitrarily small.

4 FURTHER PERSPECTIVES

4.1 FURTHER WORK ON THE PROPHET INEQUALITY WITH FREEZE

I.I.D. SETTING The work done here on the i.i.d. prophet inequality with freeze still leaves a lot of room to work. Can we validate or invalidate the bounded approximation hypothesis ? More importantly, how much can we improve the approximation ratio if we let the thresholds vary as we advance ? Can we improve the approximation ratio by simply letting the quantile depend on the distribution ? What is the optimal approximation ratio for some specific freeze families (e.g. the uniform freeze over $[1, n]$, or the geometric freeze over $[1, n]$) ? We made some unsuccessful attempts to generalize Singla's idea to relax the buyer's time of arrival (see [18]), in order to derive the optimal approximation radio for uniform freeze. We know from our work that the optimal uniform freeze approximation ratio is lower bounded by $\frac{1+e^{-2}}{2} \approx 0.567$, and it is upper bounded by the average of the optimal ratio in the no freeze setting (this would be achieved if the seller had foresight on the freeze), which in this case tends to $\frac{1}{\beta} \approx 0.745$ as n grows.

GENERAL CASE We did not look at the general case where the buyers are not required to follow the same distribution. Although this is the variant closest to the classical prophet inequality, there exists as far as we know no work in the litterature about it. Here we only provide an example that shows that no constant approximation ratio can be expected in all generality. We use the fact that knowing when the freeze is exactly occurring is considerably advantageous to the prophet. The idea is to have each X_i dwarf the one before, so that the seller has to guess exactly when the freeze is occurring in order to do achieve a good approximation.

Let $n \geqslant 2$, $x > 1$, and let $X_i = x^i$ almost surely. Let the freeze F verify for $1 \leqslant i \leqslant n$,

$$\mathbb{P}(F \geqslant i \mid F \geqslant i-1) = \frac{1}{x}$$

Hence $\mathbb{P}(F = i) = \frac{1-1/x}{x^{i-1}}$. Now let V_i be the optimal expected reward of the seller, conditioned on her reaching step X_i after having rejected the first $i-1$. The optimal expected reward is achieved by following the optimal strategy that is determined by backward induction. We have

$$V_n = \mathbb{E}(X_n) = x^n$$

$$V_{n-1} = \mathbb{E}(\max(X_{n-1}, \mathbb{P}(F \geqslant n \mid F \geqslant n-1)V_n)) = x^{n-1}$$
$$\vdots$$
$$V_i = x^i$$

And hence the optimal strategy gives the seller an expected reward of $V_1 = x$. On the other hand, the prophet's expected reward is

$$\begin{aligned}\sum_{i=1}^{n-1} \mathbb{P}(F = i)x^i + \mathbb{P}(F \geqslant n)x^n &= (1 - \frac{1}{x}) \sum_{i=1}^{n-1} \frac{1}{x^{i-1}} x^i + \frac{1}{x^n} x^n \\ &= (1 - \frac{1}{x})(n-1)x + 1 \\ &= (n-1)x\end{aligned}$$

The approximation ratio is $\frac{1}{n-1}$

4.2 PROPHET SECRETARY WITH FREEZE

Some other variants were briefly investigated – albeit unsuccessfully. In particular, the prophet secretary with freeze looks to be a good "mid-point" between the i.i.d case and the general case. In prophet secretary, the seller still tries to maximize her expected reward (just like prophet inequality), but every order of arrival of the buyers is equally likely (just like the secretary problem). It is known that without freeze, the optimal approximation ratio is $1 - e^{-1}$ and can be achieved with a single threshold algorithm [6]. Investigating the effect of a freeze variable appears interesting to us. In particular, approximation ratios found in this setting would carry as a lower bound in the i.i.d. case.

A

APPENDIX

W_n is continuous

$$\frac{1-\sum_{r=1}^{n} a_F(r)(1-x)^r}{x} \underset{0}{=} \frac{1}{x} \times (1 - \sum_{r=1}^{n} a_F(r)(1-rx) + o(x))$$
$$\underset{0}{=} \frac{1}{x} \times (1 - 1 + x\sum_{r=1}^{n} a_F(r)r) + o(1)$$
$$\underset{0}{\rightarrow} \mathbb{E}[F]$$

and thus W_n is continuous in 0.

W_n is convex, decreasing Let $\xi_r : x \mapsto \frac{1-(1-x)^r}{x}$. We have $W_n = \sum_{r=1}^{n} a_F(r)\xi_r$. Now we have

$$\xi_r(x) = \sum_{k=0}^{r-1} (1-x)^k$$

where every $(1-x)^k$ is convex decreasing. Hence every ξ_r is convex decreasing as well, and W_n is convex decreasing.

$xW_n(x)$ is continuous, increasing, concave Each $-(1-x)^r$ are increasing, concave functions of x, hence xW_n is increasing concave. It is obviously continuous.

[1] A. R. Abdel-Hamid, J. A. Bather, and G. B. Trustrum. "The Secretary Problem with an Unknown Number of Candidates". In: *Journal of Applied Probability* 19.3 (1982), pp. 619–630. ISSN: 00219002. URL: http://www.jstor.org/stable/3213519.

[2] S. Alaei. "Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers". In: (2011), pp. 512–521.

[3] José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. "Posted Price Mechanisms for a Random Stream of Customers". In: EC '17 (2017), 169–186. URL: https://doi.org/10.1145/3033274.3085137.

[4] Jose Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. "Recent Developments in Prophet Inequalities". In: 17.1 (May 2019), 61–70. URL: https://doi.org/10.1145/3331033.3331039.

[5] E. B. Dynkin. "Optimal choice of the stopping moment of a Markov process". In: *Dokl. Akad. Nauk SSSR* 150 (2 1963), pp. 238–240.

[6] Soheil Ehsani, MohammadTaghi Hajiaghayi, Thomas Kesselheim, and Sahil Singla. "Prophet Secretary for Combinatorial Auctions and Matroids". In: (2018), 700–714.

[7] Jason D. Hartline. *Mechanism Design and Approximation*. 2016. URL: http://jasonhartline.com/MDnA/MDnA-ch1to8.pdf.

[8] T. P. Hill and Robert P. Kertz. "Comparisons of Stop Rule and Supremum Expectations of I.I.D. Random Variables". In: *Ann. Probab.* 10.2 (May 1982), pp. 336–345. DOI: 10.1214/aop/1176993861. URL: https://doi.org/10.1214/aop/1176993861.

[9] Robert P Kertz. "Stop Rule and Supremum Expectations of i.i.d. Random Variables: A Complete Comparison by Conjugate Duality". In: *J. Multivar. Anal.* 19.1 (1986), 88–112. URL: https://doi.org/10.1016/0047-259X(86)90095-3.

[10] Robert Kleinberg and Seth Matthew Weinberg. "Matroid Prophet Inequalities". In: STOC '12 (2012), 123–136.

[11] U. Krengel and L. Sucheston. "On semiamarts, amarts, and processes with finite value". In: *Probability on Banach Spaces* 4 (1978), pp. 197–266.

[12] Pevcaric, J., Pejkovic, T. "On an integral inequality." eng. In: *JIPAM. Journal of Inequalities in Pure & Applied Mathematics [electronic only]* 5.2 (2004), Paper No. 47, 6 p., electronic only–Paper No. 47, 6 p., electronic only. URL: http://eudml.org/doc/124838.

[13] E. L. Presman and I. M. Sonin. "The best choice problem for a random number of objects". In: *Theory Probab. Appl.* 17 (1972), pp. 657–668.

[14] Ester Samuel-Cahn. "Comparison of Threshold Stop Rules and Maximum for Independent Nonnegative Random Variables". In: *Annals of Probability* (1984), pp. 1213–1216.

[15] Ester Samuel-Cahn. "A Difference Prophet Inequality for Bounded I.I.D. Variables, with Cost for Observations". In: *Ann. Probab.* 20.3 (July 1992), pp. 1222–1228.

[16] Ester Samuel-Cahn. "The best-choice secretary problem with random freeze on jobs". In: *Stochastic Processes and their Applications* 55 (Feb. 1995), pp. 315–327.

[17] Ester Samuel-Cahn. "Optimal Stopping With Random Horizon With Application to the Full-Information Best-Choice Problem With Random Freeze". In: *Journal of the American Statistical Association* 91.433 (1996), pp. 357–364. URL: http://www.jstor.org/stable/2291415.

[18] Sahil Singla. "Combinatorial Optimization Under Uncertainty : Probing and Stopping-Time Algorithms". In: (2018).

[19] Bo Waggoner. *The Tiger's Stripes*. 2018. URL: https://www.bowaggoner.com/blog/2018/08-25-prophet-inequalities/index.html.

COLOPHON

This document was typeset using the typographical look-and-feel `classicthesis` developed by André Miede. The style was inspired by Robert Bringhurst's seminal book on typography "*The Elements of Typographic Style*". `classicthesis` is available for both LaTeX and LyX:

https://bitbucket.org/amiede/classicthesis/

Final Version as of May 8, 2023 (`classicthesis` version 0.1).