UNIVERSIDAD DE CHILE
FACULTAD DE CIENCIAS FÍSICAS Y MATEMÁTICAS
DEPARTAMENTO DE INGENIERÍA MATEMÁTICA

PROPHET INEQUALITY
THROUGH SCHUR-CONVEXITY AND OPTIMAL CONTROL

TESIS PARA OPTAR AL GRADO DE
MAGÍSTER EN CIENCIAS DE LA INGENIERÍA,
MENCIÓN MATEMÁTICAS APLICADAS

MEMORIA PARA OPTAR AL TÍTULO DE
INGENIERO CIVIL MATEMÁTICO

RAIMUNDO JULIÁN SAONA URMENETA

PROFESOR GUÍA:
JOSÉ CORREA HAEUSSLER

MIEMBROS DE LA COMISIÓN:
JAIME SAN MARTÍN ARISTEGUI
BRUNO ZILIOTTO

Este trabajo ha sido parcialmente financiado por
CONICYT-Chile, ECOS-CONICYT, Google
y CMM-CONICYT PIA AFB170001.

SANTIAGO DE CHILE
2019

RESUMEN DE LA MEMORIA PARA OPTAR AL TÍTULO DE
INGENIERO CIVIL MATEMÁTICO Y AL GRADO DE
MAGÍSTER EN CIENCIAS DE LA INGENIERÍA, MENCIÓN MATEMÁTICAS APLICADAS
POR: RAIMUNDO JULIÁN SAONA URMENETA
FECHA: 2019
PROF. GUÍA: JOSÉ CORREA HAEUSSLER

PROPHET INEQUALITY
THROUGH SCHUR-CONVEXITY AND OPTIMAL CONTROL

En el clásico problema de tiempo de parada óptimo conocido como Desigualdad del profeta realizaciones de variables positivas e independientes son descubiertas secuencialmente. Una jugadora que conoce las distribuciones, pero no puede ver en el futuro, debe decidir cuándo parar y tomar la última variable revelada. Su objetivo es maximizar la esperanza de lo obtenido y su rendimiento está dado por el peor caso del cociente entre la esperanza de que obtiene y la esperanza de lo que obtendría un profeta (que puede ver en el futuro y así siempre elegir el máximo). En los setenta, Krengel y Sucheston, y Garling, [20] determinaron que el rendimiento de una jugadora puede ser una constante y que $1/2$ es la mejor constante. En la última década, la desigualdad del profeta ha resurgido como un problema importante dada su conexión con "Posted Price Mechanisms", una teoría usada en ventas en línea. Una variante de particular interés es "Prophet Secretary", donde la única diferencia es que las relaciones son descubiertas en orden aleatorio. Para esta variante, varios algoritmos logran un rendimiento de $1 - 1/e \approx 0.63$ y recientemente Azar et al. [2] mejoraron este resultado. En cuanto a cotas superiores, se sabe que una jugadora no puede hacerlo mejor que 0.745, en el límite sobre el tamaño de la instancia.

En esta tesis se deriva una forma de analizar estrategias que dependen sólo del tiempo: dada una instancia, se calcula una secuencia decreciente de exigencias que son usadas para decidir si parar o no. La jugadora tomará el primer valor que supere la exigencia correspondiente al momento en que fue descubierta. Específicamente, se considera una clase robusta de estrategias que denominamos "blind strategies". Constituyen una generalización de fijar una sola exigencia para todo el proceso y consisten en fijar una función, independiente de la instancia, que determina cómo calcular las exigencias una vez la instancia es conocida. El resultado principal es que la jugadora logra un rendimiento de al menos 0.669, superando el estado del arte (Azar et al. [2]) tanto para "Prophet Secretary" como para la variante en la que la jugadora tiene la libertad de escoger el orden en que descubre las variables (Beyhaghi et al [3]). El análisis se reduce a estudiar la distribución del tiempo de parada inducido por estas estrategias, a través de la teoría de Schur-convexidad. También, se demuestra que este tipo de estrategias no pueden lograr más que 0.675, a través de calcular el rendimiento óptimo de la jugadora contra dos instancias particulares, resolviendo un problema de control óptimo.

Finalmente, se demuestra que el conjunto más amplio de estrategias no adaptativas no pueden lograr más de $\sqrt{3} - 1 \approx 0.73$, cota que también mejora el estado del arte en cotas superiores para estrategias simples (Azar et al [2]). Se considera una estrategia como no adaptativa si al decisión de parar depende del valor, la identidad y el tiempo en que fue descubierta la variable, pero no toma en cuenta la identidad de las variables anteriores.

RESUMEN DE LA MEMORIA PARA OPTAR AL TÍTULO DE
INGENIERO CIVIL MATEMÁTICO Y AL GRADO DE
MAGÍSTER EN CIENCIAS DE LA INGENIERÍA, MENCIÓN MATEMÁTICAS APLICADAS
POR: RAIMUNDO JULIÁN SAONA URMENETA
FECHA: 2019
PROF. GUÍA: JOSÉ CORREA HAEUSSLER

PROPHET INEQUALITY
THROUGH SCHUR-CONVEXITY AND OPTIMAL CONTROL

In the classic prophet inequality, a problem in optimal stopping theory, samples from independent random variables (possibly differently distributed) arrive online. A gambler that knows the distributions, but cannot see the future, must decide at each point in time whether to stop and pick the current sample or to continue and lose that sample forever. The goal of the gambler is to maximize the expected value of what she picks and the performance measure is the worst case ratio between the expected value the gambler gets and what a prophet, that sees all the realizations in advance, gets. In the late seventies, Krengel and Sucheston, and Garling [20], established that this worst case ratio is a constant and that $1/2$ is the best possible such constant. In the last decade the theory of prophet inequalities has resurged as an important problem due to its connections to posted price mechanisms, frequently used in online sales. A particularly interesting variant is the so-called Prophet Secretary problem, in which the only difference is that the samples arrive in a uniformly random order. For this variant several algorithms are known to achieve a constant of $1 - 1/e$ and very recently this barrier was slightly improved by Azar et al. [2], while the best known upper bound is approximately 0.745.

In this thesis we derive a way of analyzing time-threshold strategies that basically sets a nonincreasing sequence of thresholds to be applied while discovering values. The gambler will thus stop the first time a sample surpasses the corresponding threshold. Specifically we consider a class of very robust strategies that we call blind quantile strategies. These constitute a clever generalization of single threshold strategies and consist in fixing a function which is used to define a sequence of thresholds once the instance is revealed. Our main result shows that these strategies can achieve a constant of 0.669 in the Prophet Secretary problem, improving upon the best known result of Azar et al. [2], and even that of Beyhaghi et al. [3] that works in the case the gambler can select the order of the samples. The crux of the analysis is a very precise analysis of the underlying stopping time distribution for the gambler's strategy that is inspired by the theory of Schur convex functions. We further prove that our family of blind strategies cannot lead to a constant better than 0.675, solving an optimal control problem based in the achievable performance in two carefully chosen instances.

Finally we prove that no nonadaptive algorithm for the gambler can achieve a constant better than 0.732, which also improves upon a recent result of Azar et al. [2]. Here, 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.

To the perfection of reality and the simplicity of Mathematics.

Acknowledgments

This thesis would not have been possible without the support, guidance and fellowship of José Correa and Bruno Ziliotto. This work is the result of our shared contributions and I am therefore profoundly grateful to them. I would also like to thank my family, for supporting me during all my studies and giving me all the freedom I have been able to enjoy.

The theoretical contribution exposed in this thesis is mostly contained in a paper ([8]) by Bruno, José and myself.

This work has been partially supported by CONICYT-Chile through grant PII 20150140, ECOS-CONICYT thorough grant C15E03, CMM-CONICYT through grant PIA AFB170001 and Google Research for Latin America Award.

Contents

Chapter 1

Introduction

When to be satisfied with the present and stop searching in the future something that might be better? This is a quite philosophical question, but any mathematician would understand that this thesis is about an optimal stopping problem, and therefore the natural questions are:

1. What do I know? (information available at the moment I take a decision)

2. How does this work? (dynamic one is subject to when deciding)

3. What is a good choice? (function to be maximized)

The first and most classical model, proposed in the 1960 partly as a recreational problem, is the "Secretary Problem" and consists in the following:

1. There is a fixed and known number of options (say numbers) available and, if faced with two options, you are able to determine which is better.

2. One by one, the options are shown to you following a random order. After discovering the first option, either you take it or drop it for ever. If you drop an option, the next is shown to you and so the process continues until you pick an option.

3. You only win if you pick the best option (the largest number).

This simple setting started a whole field of research for online decisions and a more modern survey is available in [12, 13]. Another model proposed in 1977 by Krengel and Sucheston is called "Prophet Inequality", where the decision maker know more in advance about the options. Prophet inequalities, in contrast with the Secretary problem, consist in the following:

1. There is a fixed number of positive numbers available. You know that each of these numbers were drawn from probability distributions fixed and known in advance. Moreover, you the order in which these realizations are going to be shown to you.

2. One by one, these realizations of random variables are shown to you. After discovering

the first option, either you take it or drop it for ever. If you drop an option, the next is shown to you and so the process continues until you pick an option.

3. You only care about the expectation that you get, and you compare yourself against a prophet who can see in the future and always chooses the maximum of all the realizations.

The starting point of this thesis is the work of Hajiaghayi et al. [15] and that of Chawla et al. [5] who established an interesting connection between (revenue maximizing) PPMs and prophet inequalities. Posted price mechanisms (PPM) constitute a widely used way of selling items to strategic consumers. Indeed, in the last years online sales have been moving from an auction format, to posted price formats [10], and the basic reason for this trend switch seems to be that PPMs are much simpler than optimal auctions, yet efficient enough. Furthermore, several companies have started to apply price discrimination to sell their products. Under this policy, companies set different prices for different consumers based on purchase history or other factors that may affect their willingness to pay. For example, the online data provider Lexis-Nexis sells to virtually every user at a different price. In 2012, Orbitz online travel agency found that people who use Mac computers spent as much as 30% more on hotels, so it started to show them different, and sometimes costlier, travel options than those shown to Windows visitors. Furthermore, similar pricing strategies are used to determine the (personalized) reserve prices for Google ads and have significant potential impact in the pricing of cloud services.

The way these mechanisms work is as follows. Suppose a seller has an item to sell. Consumers arrive one at a time and the seller proposes to each consumer a take-it-or-leave-it offer. The first customer accepting the offer pays that price and takes the item. These type of mechanisms are flexible and adapt well to different scenarios. Furthermore, their simplicity and the fact that strategic behaviour vanishes make them quite suitable for many applications [5]. Of course, PPMs are suboptimal and therefore the study of what their approximation guarantees – where the benchmark is that given by the optimal Myerson's auction [22] – has been an extremely active area in the last decade, in particular in the computer science community.

Another explanation of the "Prophet Inequality" setting is the following: a gambler is faced to a sequence of random variables and has to pick a stopping time so that the expected value he gets is as close as possible to the expectation of the maximum of all random variables, interpreted as what a prophet, who knows the realizations in advance, could get. Chawla et al. implicitly show that any prophet type inequality can be turned into a PPM with the same approximation guarantee. This is obtained by noting that a PPM for revenue maximization can be seen as (threshold) stopping rule for the gambler but on the virtual values, and later identify these thresholds with prices. Very recently Correa et al. [6] show that also the converse holds proving that a posted price mechanism can be turned into a prophet type inequality with the same approximation guarantee. As a consequence of these developments, most work in the field concentrated on prophet inequalities and then applied the obtained results to sequential posted price mechanisms.

1.1 Prophet-Inequalities

For fixed $n > 1$, let $V_1, \ldots, V_n$ be non-negative, independent random variables and T_n their set of stopping rules. A classic result of Krengel and Sucheston, and Garling [20, 19] asserts that the quotient between $\sup\{\mathbb{E}(V_T) : T \in T_n\}$ and $\mathbb{E}(\max\{V_1, \ldots, V_n\})$ must be greater or equal to 1/2, and that 1/2 is the best possible bound. The interpretation of this result says that a *gambler*, who only knows the distribution of the random variables and that looks at them sequentially, can select a stopping rule that guarantees her half of the value that a *prophet*, who knows all the realizations could get. The study of this type of inequalities, known as *prophet inequalities*, was initiated by Gilbert and Mosteller [14] and attracted a lot of attention in the eighties [16, 17, 18, 24]. In particular Samuel-Cahn [24] noted that rather than looking at the set of all stopping rules one can (quite naturally) only look at threshold stopping rules in which the decision to stop depends on whether the value of the currently observed random variable is above a certain threshold (and possibly on the rest of the history).

Example. Consider two options V_1 and V_2. V_1 is the constant 1 and V_2 takes the value $1/\varepsilon$ with probability ε and zero otherwise. In the first step, the gambler is faced with the value 1. The expectation of the future (V_2) is also 1 and therefore, the optimal mechanism is indifferent between picking it or not. Therefore, the expectation of the gambler is 1. In the other hand, the expectation of the maximum is $2 - 1/\varepsilon$, so, in the limit, we get that the gambler can not get more than half the expectation of the maximum, which is what the prophet gets.

Although the situation for the standard prophet inequality just described is well understood, there are variants of the problem, which are particularly relevant given the connection to PPMs, for which the situation is very different. In what follows we describe three such variants that are connected to each other. The second one constitutes the main focus of this thesis.

- *Order selection.* In this version the gambler is allowed to select the order in which she examines the random variables. For this version [5] improved the bound of 1/2 (of the standard prophet inequality) to $1 - 1/\mathrm{e} \approx 0.6321$. This bound remained the best known for quite some time until Azar et al. [2] improved it to $1 - 1/\mathrm{e} + 1/400 \approx 0.6346$. Interestingly, the bound of Azar et al. actually applies to the random order case described below. Very recently Beyhaghi et al. [3], use order selection to further improve the bound to $1 - 1/\mathrm{e} + 0.022 \approx 0.6541$.

  Example. Consider the same instance $V_1 \equiv 1$ and $V_2 = 1/\varepsilon B(\varepsilon)$. If the order of discovering where V_2 first and V_1 second, the gambler could actually always pick the maximum: if V_2 is $1/\varepsilon$ she takes it, if not, she drops it and take $V_1 = 1$. Now the quotient between her expectation and that of the maximum is simply 1.

- *Prophet secretary (or random order).* In this version the random variables are shown to the gambler in random order, as in the secretary problem. This version was first studied by Esfandiari et al. [11] who found a bound of $1 - 1/\mathrm{e}$. Their algorithm defines a nonincreasing sequence of n thresholds $\tau_1, \ldots, \tau_n$ that only depend on the expectation of the maximum of the $V_i's$ and on n. The gambler at time-step i stops if the value of

V_{σ_i} (the variable shown at step i) surpasses τ_i. Later, Correa et al. [7] proved that the same factor of $1 - 1/e$ can be obtained with a personalized and nonadaptive sequence of thresholds, that is thresholds $\tau_1, \ldots, \tau_n$ such that whenever variable V_i is shown the gambler stops if its value is above τ_i, not taking into account previously seen variables. In recent work, Ehsani et al. [9] show that the bound of $1 - 1/e$ can be achieved using a single threshold (having to randomize to break ties in some situations). This result appears to be surprising since without the ability of breaking ties at random $1/2$ is the best possible constant and this insight turns out to be the starting point of this work. Shortly after the work of Ehsani et al., Azar et al. [2] improved it to $1 - 1/e + 1/400 \approx 0.6346$ through an algorithm that relies on some subtle case analysis.

Example. Consider the instance $V_1 \equiv 1$ and $V_2 = 1/\varepsilon B(\varepsilon)$. If the order of discovering where V_2 first and V_1 second, we know the gambler gets the maximum. If the order is $\sigma = (1, 2)$, we know the gambler gets 1 and the expected maximum is almost 2. Therefore, in the limit, the quotient of interest is $3/4$.

- *IID Prophet inequality.* Finally, we mention the case when the random variables are identically distributed. Here, the constant $1/2$ can also be improved. Hill and Kertz [16] provided a family of "bad" instances from which Kertz [17] proved the largest possible bound one could expect is $1/\beta \approx 0.7451$, where β is the unique solution to $\int_0^1 \frac{1}{y(1-\ln(y))+(\beta-1)} dy = 1$. Quite surprisingly, this upper bound is still the best known upper bound for the two variants above. Regarding algorithms Hill and Kertz also proved a bound of $1 - 1/e$ which was improved by Abolhassan et al. [1] to 0.7380. Finally Correa et al. [7] proved that $1/\beta \approx 0.7451$ is a tight value. To this end they exhibit a quantile strategy for the gambler in which some quantiles $q_1 < \ldots < q_n$, that only depend on n (and not on the distribution), are precomputed and then translated into thresholds so that if the gambler gets to step i, she will stop with probability q_i.

Example. Consider the instance given by the uniform distribution in $[0, 1]$. Once the first variable is revealed, the expectation of the future is $1/2$, therefore, the optimal algorithms picks V_1 only if its value is above $1/2$ (if not, it picks the second option). Thus, the expectation of the gambler is given by $5/8$. In the other hand, the expectation of the maximum is $3/4$, leading to a quotient of $5/6$.

A striking fact about the three problems above is that, in terms of attainable performance, no separation between them is known. For instance, it is perfectly plausible that the best constant achievable in all three cases is that of the i.i.d. case, while it is also possible that all three versions admit different optimal constants. To be more specific, the upper bound of approximately 0.7451 for the achievable performance in the i.i.d. case is a natural bound for the other two variants. Moreover, it is the only upper bound for the optimal algorithm in both cases. In terms of lower bounds, the i.i.d. case is "solved" and there is an algorithm that achieves this upper bound. In the case of Order Selection, the algorithm of Beyhaghi et al. [3] achieves a performance of approximately 0.6541. For Prophet Secretary, the proposal of Azar et al. [2] rises up to approximately 0.6346.

1.2 Contribution

In this thesis we study the prophet secretary problem and propose improved algorithms for it. In particular our work improves upon the recent work of Ehsani et al. [9], Azar et al. [2], and Beyhaghi et al. [3] by providing an algorithm that guarantees the gambler a fraction of $0.669 \approx 1 - 1/\mathrm{e} + 1/27$ in the Prophet secretary setting (and therefore in the Order selection case too). Our main contribution however is not the actual numerical improvement but rather the way in which this is obtained.

From a conceptual viewpoint we introduce a class of algorithms which we call *blind quantile* strategies, that are very robust in nature. This type of algorithm is a clever generalization of the single threshold algorithm of Ehsani et al. to a multi-threshold setting. In their algorithm Ehsani et al. first compute a threshold τ such that $\mathbb{P}(\max\{V_1, \ldots, V_n\} \leq \tau) = 1/\mathrm{e}$ and then use this τ as a single threshold strategy, so that the gambler stops the first time in which the observed value surpasses τ. They observe that this strategy only works for random variables with continuous distributions, however they also note that by allowing randomization the strategy can be extended to general random variables. Rather than fixing a single probability of acceptance we fix a function $\alpha : [0, 1] \to [0, 1]$ which is used to define a sequence of thresholds in the following way. Given an instance with n continuous distributions we draw uniformly and independently n random values in $[0, 1]$, and reorder them as $u_{[1]} < \cdots < u_{[n]}$. Then we set thresholds $\tau_1, \ldots, \tau_n$ such that $\mathbb{P}(\max\{V_1, \ldots, V_n\} \leq \tau_\mathrm{i}) = \alpha(u_{[\mathrm{i}]})$ and the gambler stops at time i if $V_{\sigma_\mathrm{i}} > \tau_\mathrm{i}$. Note that if the function α is nonincreasing this will lead to a nonincreasing sequence of thresholds.

The idea of blind quantile strategies comes from the i.i.d. case mentioned above. In that setting the strategies are indeed best possible as shown by Correa et al. [7]. What makes blind quantile strategies attractive is that although decisions are time dependent, this dependence lies completely in the choice of the function α, which is independent of the instance. This independence significantly simplifies the analysis of multi-threshold strategies. Again, when facing discontinuous distributions we also require randomization for the results to hold.

From a technical standpoint our analysis introduces the use of Schur convexity [23] in the prophet inequality setting. We start our analysis by revisiting the single threshold strategy of Ehsani et al., which corresponds to a constant blind quantile strategy $\alpha(\cdot) = 1/\mathrm{e}$. We exhibit a new analysis of this strategy showing stochastic dominance type result. Indeed we prove that the probability that the gambler gets a value of more than t is at least that of the maximum being more than t, rescaled by a factor $1 - 1/\mathrm{e}$. This result uses Schur convexity to deduce that if there is a value above the threshold τ, then it is chosen by the gambler with probability at least $1 - 1/\mathrm{e}$. Then we extend this analysis to deal with more general functions α which require precise bounds on the distribution of the stopping time corresponding to a function α. These bounds make use of results of Esfandiari et al. [11] and Azar et al. [2] and of newly derived bounds that follow from the core ideas in Schur convexity theory.

Again in this more general setting we find an appropriate stochastic dominance type bound on the probability that the gambler obtains at least a certain amount with respect to the probability that the prophet obtains the same amount. Interestingly we manage to make the bound solely dependent on the blind strategy by basically controlling the implied stopping time distribution (patience of the gambler). Then optimizing over blind strategies leads to the improved bound of 0.669. Through the thesis we show two lower bounds on the performance of a blind strategy, the second more involved than the first one. In the first case, there is a natural way to optimize over the choice of α solving an ordinary differential equation, leading to a guarantee of 0.665. In the second case, using a refined analysis, we derive the stated bound of 0.669. Although it may seem that our general approach still leaves some room for improvement, we prove that blind strategies cannot lead to a factor better than 0.675. This bound is obtained by taking two carefully chosen instances and proving that no blind strategy can perform well in both, through Optimal Control theory.

Finally, we prove a quite general upper bound on the performance of any nonadaptive strategy, namely an algorithm whose decision to stop at any step can be precomputed and does not take into account the history that has been observed. In particular, in a nonadaptive strategy the decision to stop can depend on the distributions of the instance, on the index of the random variable being sampled, on the value sampled, and on the current time, but not on the history that has been observed. Our result here is to find an instance (which is not i.i.d.) in which no nonadaptive strategy can perform better than $\sqrt{3}-1 \approx 0.732$. This improves upon the best possible bound known of 0.745 which corresponds to the i.i.d. case and was proved by Hill and Kertz [16]. Furthermore it improves and generalizes a recent bound of $11/15 \approx 0.733$ of Azar et al. [2] for the more restricted class of *Deterministic distribution-insensitive algorithms*.

1.3 Preliminaries and notation

Given nonnegative independent random variables $V_1, \ldots, V_n$ and a random permutation $\sigma : [n] \to [n]$[1], in the Prophet Secretary problem a gambler is presented realizations of the random variables, along with their identities, in the order given by σ, i.e., at time j she sees the realization of V_{σ_j}, and also σ_j. The goal of the gambler is to find a stopping time T such that $\mathbb{E}(V_{\sigma_T})$ is as large as possible. In particular we want to find the largest possible constant c such that

$$\sup\{\mathbb{E}(V_{\sigma_T}) : T \in \mathcal{T}_n\} \geq c \cdot \mathbb{E}(\max\{V_1, \ldots, V_n\}),$$

where $\mathcal{T}_n$ is the set of stopping rules.

Throughout this thesis we denote by $F_1, \ldots, F_n$ the underlying distributions of $V_1, \ldots, V_n$, which we assume to be continuous. All our results apply unchanged to the case of general distributions by introducing random tie-breaking rules (this is made precise in Section 8). To see why random tie-breaking rules are actually needed, consider the single threshold strategy of Ehsani et al. [9]. Recall that they compute a threshold τ such that $\mathbb{P}(\max\{V_1, \ldots, V_n\} \leq \tau) = 1/\mathrm{e}$ and then use this τ as a single threshold strategy, which, by allowing random tie-breaking, leads to a performance of $1 - 1/\mathrm{e}$. However, if random tie-breaking is not allowed, a single threshold strategy cannot achieve a constant better than 1/2. Indeed, consider the instance with $n-1$ deterministic random variables equal to 1 and one random variable giving n with probability $1/n$ and zero otherwise. Now, for a fixed threshold $\tau < 1$ the gambler

[1]Here $[n]$ denotes the set $\{1, \ldots, n\}$

gets n with probability $1/n^2$ and 1 otherwise so that she gets approximately 1, whereas if $\tau \geq 1$ the gambler gets n with probability $1/n$, leading to an expected value of 1. Noting that the expectation of the maximum in this instance equals 2, we conclude that fixed thresholds cannot achieve a guarantee better than 1/2. However, as Ehsani et al. note, if in this instance the gambler accepts a deterministic random variable with probability $1/n$, then her expected value approaches $2(1 - 1/\mathrm{e})$. In Section 8 we extend this idea for the more general multi-threshold strategies.

The main type of stopping rules we deal with in this thesis uses a nonincreasing threshold approach. This is a quite natural idea, since Esfandiari et al.[11] already use such an approach to derive a guarantee of $1 - 1/\mathrm{e}$. Interestingly, the analysis of multi-threshold strategies becomes rather difficult when trying to go beyond this bound. This is evident from the fact that the more recent results take a different approach. In this thesis we use a rather restrictive class of multi-threshold strategies that we call *blind quantile* strategies. These are simply given by a nonincreasing function $\alpha : [0, 1] \to [0, 1]$ which is turned into an algorithm as follows: given an instance $F_1, ..., F_n$ of continuous distributions, we independently draw $u_1, ..., u_n$ from a uniform distribution on $[0, 1]$ and find thresholds $\tau_1, \ldots, \tau_n$ such that

$$\mathbb{P}(\max_{\mathrm{i} \in [n]} \{V_\mathrm{i}\} \leq \tau_\mathrm{i}) = \alpha(u_{[\mathrm{i}]}) ,$$

where $u_{[\mathrm{i}]}$ is the i-th order statistic of $u_1, \ldots, u_n$. Then the algorithm for the gambler stops at the first time in which a value exceeds the corresponding threshold, in other words the gambler applies the following.

Algorithm 1 Time Threshold Algorithm ($TTA_{\tau_1,\ldots,\tau_n}$)

for i = 1, ..., n **do**
    if $V_{\sigma_\mathrm{i}} > \tau_\mathrm{i}$ **then**
        Take V_{σ_i}
    end if
end for

Note that a blind strategy is uniquely determined by the choice of function α, independent of the actual distributions or even size of the instance. This justifies that we may simply talk about strategy α. Our goal is thus to find a *good* function α such that the previous algorithm performs well against any instance.

For a blind strategy α and an instance $F_1, \ldots, F_n$, we will be interested in the underlying stopping time T which is the random variable defined as $T := \inf\{j \in [n] : V_{\sigma_j} > \tau_j\}$, where the $\tau_1, \ldots, \tau_n$ are the corresponding thresholds. In particular the reward of the gambler is $V_{\sigma_T} \mathbb{1}_{T<\infty}$, which we simply denote by V_{σ_T}.

Chapter 2

Schur-convexity

A general presentation of the theory of Schur-convex functions is available in the third chapter of [21]. Here we will present only what is necessary and useful for the understanding of this work. The objective we have in mind is to solve the following problem:

$$(P)\begin{cases} opt_x & \phi(x) \\ s.t. & a\mathbb{1} \leq x \leq b\mathbb{1} \\ & \sum_{\mathrm{i}\in[n]} x_{\mathrm{i}} = \beta\,, \end{cases}$$

where *opt* stands for either maximize or minimize, $a, b, \beta \in \overline{\mathbb{R}}$ and $\phi : \mathbb{R}^n \to \mathbb{R}$ is a *smooth* function and also symmetric with respect to permutations of the input.

Solving this kind of problems will allow us to prove inequalities which are used to derive lower bounds in the performance of blind strategies.

2.1 Definitions

For $a, b, \beta \in \overline{\mathbb{R}}$, define the set $I_{a,b,\beta} := \{x \in \mathbb{R}^n : a\mathbb{1} \leq x \leq b\mathbb{1} \quad , \quad \sum_{\mathrm{i}\in[n]} x_{\mathrm{i}} = \beta\}$, the domain we are interested in. We can define a partial order $\preceq$ over this set by the following construction. Take $S : \mathbb{R}^n \to \mathbb{R}^n$ defined by $S(x)_{\mathrm{i}} := \max\{\sum_{\mathrm{i}\in I} x_{\mathrm{i}} : I \subseteq [n], |I| = \mathrm{i}\}$, then for $x, y, \in \mathbb{R}^n$

$$x \preceq y \quad \Leftrightarrow \quad \forall \mathrm{i} \in [n] \quad S(x)_{\mathrm{i}} \leq S(y)_{\mathrm{i}}\,,$$

in other words, $x \preceq y$ if and only if all its partial sums are less or equal to the corresponding partial sums of y, when these are ordered respect to the usual order $\leq$.

With this relation we can define a Schur-convex function as a $\preceq$-preserving function, ie: $\phi : I_{a,b,\beta} \to \mathbb{R}$ is Schur-convex if for all $x, y \in I_{a,b,\beta}$ we have that

$$x \preceq y \quad \Rightarrow \quad \phi(x) \leq \phi(y)\,.$$

The other concept we need to define is that of "symmetric respect to permutations of the input". A set $I \subseteq \mathbb{R}^n$ is said to be permutation-symmetric if for all $x \in I$ and all permutations

$\pi : [n] \to [n]$, $\pi(x) \in I$. Similarly, a function $\phi : I \to \mathbb{R}$ is said to be permutation-symmetric if I is a permutation-symmetric set and for all $x \in I$ and all permutations $\pi : [n] \to [n]$, $\phi(x) = \phi(\pi(x))$. With this concept, we can notice that if $\phi : I_{a,b,\beta} \to \mathbb{R}$ is permutation-symmetric, then it is sufficient to verify Schur-convexity in the set $I_{a,b,\beta} \cap \{x \in \mathbb{R}^n : x_1 \geq x_2 \geq \ldots \geq x_n\}$.

Example. Consider $I_{0,1,1/2} \subseteq \mathbb{R}^2$ and $\phi(x) = -x_1 x_2$. In this domain, we can simplify the analysis since $x_2 = 1/2 - x_1$ and therefore $\phi(x) = -x_1(1/2 - x_1) = x_1^2 - x_1/2$, which is clearly Schur-convex.

An important remark to make is that, even when $\preceq$ is only a partial order relation, for any $a, b, \beta \in \mathbb{R}$, there exists a unique $\preceq$-minimum (x_m) and $\preceq$-maximum (x_M) in $I_{a,b,\beta}$ (provided it is non-empty) given by

$$x_m = \frac{\beta}{n}\mathbb{1} \quad , \quad x_M = (b, b, \ldots, b, a + \varepsilon, a, a, \ldots, a)\,,$$

where $\varepsilon \geq 0$ and x_M concentrates as much positivity as possible in the first coordinates. This is important because if one needs to solve the optimization problem (P), proving Schur-convexity of ϕ would immediately tell us that the value of the problem is either $\phi(x_m)$ or $\phi(x_M)$, depending on the symbol *opt*.

2.2 Sufficient conditions

Given $\phi \in \mathcal{C}^1(\overline{I_{a,b,\beta}}; \mathbb{R})$ (continuous in $I_{a,b,\beta}$ and continuously differentiable in its interior) a permutation symmetric function, for it to be Schur-convex it is sufficient to verify the following condition, known as the Schur-Ostrowski condition [23],

$$\forall x \in \mathrm{i}nt(I_{a,b,\beta}) \quad (x_1 - x_2)\left[\partial_{x_1}\phi(x) - \partial_{x_2}\phi(x)\right] \geq 0\,,$$

which simply states that ϕ grows in the correct direction when making a differential change for all points in the interior of the domain.

Example. Consider $I_{0,1,1/2}$ and $\phi(x) = -\prod_{\mathrm{i}\in[n]} x_{\mathrm{i}}$. Note that $[\partial_{x_1}\phi(x) - \partial_{x_2}\phi(x)] = (x_1 - x_2)[x_3 x_4 \ldots x_n]$, then $\forall x \in \mathrm{i}nt(I_{0,1,1/2}) \quad (x_1 - x_2)\left[\partial_{x_1}\phi(x) - \partial_{x_2}\phi(x)\right] = (x_1 - x_2)^2[x_3 x_4 \ldots x_n] \geq 0$, and therefore ϕ is Schur-convex.

2.3 Example

Consider $\beta > 0$ and $\phi : I_{0,\infty,\beta} \subseteq \mathbb{R}_+^n \longrightarrow \mathbb{R}$ given by

$$\phi(x) := \sum_{S\subseteq[n]} \frac{1}{|S|+1} \prod_{j\in S} \mathrm{e}^{x_j} - 1\,.$$

Trying to verify the Schur-Ostrowski condition, straightforward calculations yield

$$\begin{aligned}
\partial_{x_1}\phi(x) &= \sum_{\substack{S\subseteq[n-1]\\ S\ni 1}} \frac{1}{|S|+1} \quad \mathrm{e}^{x_1}\prod_{\substack{j\in S\\ j\neq 1}} \mathrm{e}^{x_j} - 1 \\
&= \frac{\mathrm{e}^{x_1}}{\mathrm{e}^{x_1}-1}\left(\sum_{\substack{S\subseteq[n-1]\\ S\ni 1,2}} \frac{1}{|S|+1} \quad \prod_{j\in S} \mathrm{e}^{x_j} - 1\right) \\
&+ \frac{\mathrm{e}^{x_1}}{(\mathrm{e}^{x_1}-1)(\mathrm{e}^{x_2}-1)}\left(\sum_{\substack{S\subseteq[n-1]\\ S\ni 1,2}} \frac{1}{|S|} \quad \prod_{j\in S} \mathrm{e}^{x_j} - 1\right) \\
&=: \frac{\mathrm{e}^{x_1}}{\mathrm{e}^{x_1}-1}A + \frac{\mathrm{e}^{x_1}}{(\mathrm{e}^{x_1}-1)(\mathrm{e}^{x_2}-1)}B
\end{aligned}$$

and, by symmetry, $\partial_{x_2}\phi(x) = \frac{\mathrm{e}^{x_2}}{\mathrm{e}^{x_2}-1}A + \frac{\mathrm{e}^{x_2}}{(\mathrm{e}^{x_2}-1)(\mathrm{e}^{x_1}-1)}B$. Then,

$$\begin{aligned}
[\partial_{x_1}\phi(x) - \partial_{x_2}\phi(x)] &= A\left[\frac{\mathrm{e}^{x_1}}{\mathrm{e}^{x_1}-1} - \frac{\mathrm{e}^{x_2}}{\mathrm{e}^{x_2}-1}\right] + B\left[\frac{\mathrm{e}^{x_1}}{(\mathrm{e}^{x_1}-1)(\mathrm{e}^{x_2}-1)} - \frac{\mathrm{e}^{x_2}}{(\mathrm{e}^{x_2}-1)(\mathrm{e}^{x_1}-1)}\right] \\
&= (B-A)\left[\frac{\mathrm{e}^{x_1}-\mathrm{e}^{x_2}}{(\mathrm{e}^{x_2}-1)(\mathrm{e}^{x_1}-1)}\right].
\end{aligned}$$

Finally, we consider $x \in int(I_{0,\infty,\beta})$ and note that $\mathrm{e}^{x_j}-1>0$ and $B>A$. Then, $(x_1-x_2)[\partial_{x_1}\phi(x)-\partial_{x_2}\phi(x)] \geq 0$ if and only if $(x_1-x_2)(\mathrm{e}^{x_1}-\mathrm{e}^{x_2}) \geq 0$, which holds by monotonicity of the exponential function. With this, the Schur-Ostrowski condition holds and ϕ is Schur-convex.

Chapter 3

Optimal Control

A general presentation of the theory of Optimal Control is available in the book "Optimization—Theory and Applications: Problems with Ordinary Differential Equations" [4]. Here, we will focus only in what is known as a "Mayer Problem", which consists in the following minimization problem:

$$(P)\begin{cases} min_\alpha & h(t_1, x(t_1), t_2, x(t_2)) & & \\ s.t. & \frac{\mathrm{d}x}{\mathrm{d}t}(t) & = f(t, x(t), \alpha(t)) & t \in [t_1, t_2] \\ & (t, x(t)) & \in A & t \in [t_1, t_2] \\ & \alpha(t) & \in U & t \in [t_1, t_2] \\ & (t_1, x(t_1), t_2, x(t_2)) & \in B\,. & \end{cases}$$

In this setting, we think of a dynamic described by the controlled differential equation given by $\dot{x}(t) = f(t, x(t), \alpha(t))$, where α is the control. The rest of the constrains concern: controlling the admissible trajectories by $(t, x(t)) \in A$, stating admissible controls by restricting its values to a fixed set by $\alpha(t) \in U$ and initial and/or final conditions given by $(t_1, x(t_1), t_2, x(t_2)) \in B$.

Our intention is to consider a blind strategy given by some function $\alpha : [0, 1] \to [0, 1]$, then define the sets A, U, B and the function f such that the solution of the controlled differential equation allows us to compute an upper bound on the performance of the blind strategy α. This way, we will be able to minimize this upper bound over all possible blind strategies (solving the Optimal Control problem) and constrain the achievable performance for this kind of algorithms.

The general approach to solve a Mayer problem, an optimization problem in infinite dimension, is first to prove the existence of an optimal solution satisfying some easy-to-verify conditions. Then, one is able to deduce four necessary conditions on the optimal solution. These are used to characterize, not always uniquely, the solutions and reduce the space in which we search the optimal control.

3.1 Existence of an appropriate solution

As we said, we need the existence of an appropriate solution. Formally, there must exists an optimal solution (x^*, α^*) that not only is an admissible pair of the Mayer problem and achieves the minimum, but also satisfy the following conditions:

(a) $\{(t, x^*(t)) : t \in [t_1, t_2]\} \subseteq int(A)$.

(b) α^* is bounded.

(c) $(t_1, x^*(t_1), t_2, x^*(t_2)) \in B$ and B possesses a tangent linear variety, $B'(t_1, x^*(t_1), t_2, x^*(t_2))$, whose vectors will be denoted $(\mathrm{d}t_1, \mathrm{d}x_1, \mathrm{d}t_2, \mathrm{d}x_2)$.

(d) h possesses a differential $\mathrm{d}h$ at $(t_1, x^*(t_1), t_2, x^*(t_2))$.

As we said earlier, in our particular application these conditions are trivially satisfied.

3.2 Necessary conditions for optimality

To state the necessary conditions of optimality, we must first define two functions:

$$H(t, x, u, \lambda) = \sum_{j \in [n]} \lambda_j f_j(t, x, u)$$

and

$$M(t, x, \lambda) = \inf_{u \in U} H(t, x, u, \lambda)\,.$$

Then, there are four necessary conditions to be satisfied by an optimal pair (x^*, u^*).

(P1) $\exists \lambda = (\lambda_1, \ldots, \lambda_n)$ an absolutely continuous vector function such that for $j \in [n]$

$$\frac{\mathrm{d}\lambda_j}{\mathrm{d}t} = -\frac{\mathrm{d}H}{\mathrm{d}x_j}(t, x^*(t), u^*(t), \lambda(t))\,,$$

for $t \in [t_1, t_2]$ (a.e.). Moreover, if $\mathrm{d}h \neq 0$, then $\lambda(t)$ is never zero in $[t_1, t_2]$.

(P2) For $t \in [t_1, t_2]$ (a.e.),

$$M(t, x^*, \lambda) = H(t, x^*, u^*, \lambda)\,.$$

(P3) $M(t) := M(t, x^*(t), \lambda(t))$ is absolutely continuous in $[t_1, t_2]$ and

$$\frac{\mathrm{d}M}{\mathrm{d}t}(t) = \frac{\mathrm{d}H}{\mathrm{d}t}(t, x^*(t), u^*(t), \lambda(t))\,.$$

(P4) $\exists \lambda_0 \in [0, \infty)$ such that $\forall (\mathrm{d}t_1, \mathrm{d}x_1, \mathrm{d}t_2, \mathrm{d}x_2) \in B'(t_1, x^*(t_1), t_2, x^*(t_2))$,

$$\lambda_0 \mathrm{d}h + \left(M(t)\mathrm{d}t - \sum_{j \in [n]} \lambda_j(t) \mathrm{d}x_j \right)_1^2 = 0\,.$$

3.3 Example

Consider the following problem:

$$(P)\begin{cases} min_\alpha \quad h(0,x(0),1,x(1)) & = -x_1(1) \\ s.t. \quad \frac{dx}{dt}(t) & = f(t,x(t),\alpha(t)) = \begin{pmatrix} 1-\alpha(t) \\ \exp[x_3(t)] \\ \ln(\alpha(t)) \end{pmatrix} & t\in[0,1] \\ (t,x(t)) & \in A = [0,1]\times\mathbb{R}^3 & t\in[0,1] \\ \alpha(t) & \in U = [0,1] & t\in[0,1] \\ (t_1,x(t_1),t_2,x(t_2)) & \in B = \{(0,(0,0,0)',1,(r,r,s)') : u,v\in\mathbb{R}\}\,. \end{cases}$$

This problem arises in the proof of our upper bound in section 7.1, and the idea is to characterize as much as possible the optimal solution.

The set B fixes the initial and final time to zero and one respectively, the initial condition to be the null vector and imposes a restriction over the final condition: $x_1(1) = x_2(1)$. Simply from these conditions it is not clear if there exists some α such that the controlled differential equation has a solution. But, for now, assume the existence of an appropriate optimal solution as stated in section 3.1 and we will focus our attention in applying the necessary conditions for optimality.

We should first compute $B'(\mathrm{e}[x^*]), H$ and M.

$$\begin{aligned} B'(\mathrm{e}[x^*]) &= \{(0,(0,0,0)',0,(\mathrm{d}r,\mathrm{d}r,\mathrm{d}s)') : \mathrm{d}r,\mathrm{d}s\in\mathbb{R}\} \\ H(t,x,\alpha,\lambda) &= \lambda_1(1-\alpha)+\lambda_2\mathrm{e}^{x_3}+\lambda_3\ln\alpha \\ M(t,x,\lambda) &= \begin{cases} H(t,x,\alpha=0,\lambda) = -\infty & ; \quad \lambda_3>0 \\ H(t,x,\alpha=1,\lambda) = \lambda_2\mathrm{e}^{x_3} & ; \quad \lambda_3\le 0, \lambda_1\ge\lambda_3 \\ H(t,x,\alpha=\frac{\lambda_3}{\lambda_1},\lambda) = \lambda_1+\lambda_2\mathrm{e}^{x_3}-\lambda_3[1+\ln\frac{\lambda_1}{\lambda_3}] & ; \quad \lambda_3\le 0, \lambda_1<\lambda_3 \end{cases} \end{aligned}$$

Then, by condition $(P1)$,

$$\dot{\lambda}_1(t) = 0 \quad , \quad \dot{\lambda}_2(t) = 0 \quad , \quad \dot{\lambda}_3(t) = -\lambda_2(t)\mathrm{e}^{x_3(t)}\,,$$

so λ_1 and λ_2 are constants. Also notice that $h = -x_1(1)$ and therefore $\mathrm{d}h = -\mathrm{d}x_1(1) = -\mathrm{d}r \in \mathbb{R}$.

By $(P4)$, $\exists\lambda_0 \ge 0$ such that for all $\mathrm{d}r, \mathrm{d}s \in \mathbb{R}$

$$-\lambda_0\mathrm{d}r - \lambda_1\mathrm{d}r - \lambda_2\mathrm{d}r - \lambda_3(1)\mathrm{d}s = 0\,.$$

Therefore, $\lambda_3(1) = 0$ and $\lambda_0+\lambda_1+\lambda_2 = 0$. Since $\dot{\lambda}_3(t) = -\lambda_2(t)\mathrm{e}^{x_3(t)}$, we have that

$$\lambda_3(t) = \lambda_2\int_t^1 \mathrm{e}^{x_3(s)}\mathrm{d}s\,.$$

By $(P3)$, $M(t) = M(t,x^*(t),\lambda(t))$ is an absolutely continuous function, therefore $\lambda_3(t)\le 0$ for all $t\in[0,1]$. Then, by the previous formula, $\lambda_2 \le 0$, moreover, we can do this inequality strict. Assume $\lambda_2 = 0$, then $\lambda_3 = 0$ and we have that $0 \leq \lambda_0 = -\lambda_1 - \lambda_2 = -\lambda_1$, therefore, $\lambda_1 \leq 0$, but since $\mathrm{d}h \neq 0$, $\lambda_1 < 0 = \lambda_3$. Then, $\alpha = \lambda_3/\lambda_1 = 0$, which is a contradiction. We can conclude that for all $t \in [0,1]$

$$\lambda_2, \lambda_3(t) < 0\,.$$

In a similar way we can deduce that $\lambda_0 > 0$. Assume $\lambda = 0$, then $\lambda_1 = -\lambda_2 > 0$. Since $\lambda_3(t) < 0 < \lambda_1$, we have that for all $t \in [0,1]$, $\alpha = 1$, which is not an optimal solution. Therefore, and w.l.o.g.

$$\lambda_0 = 1\,.$$

Notice that $\lambda_3(\cdot)$ is an increasing function, then $\{t \in [0,1) : \lambda_3(t) \leq \lambda_1\}$ is an interval, possibly empty. This implies that in a first stage $\alpha = 1$ and then, as λ_3 increases, $\alpha = \lambda_3/\lambda_1$. Since $\alpha \equiv 1$ is not an optimal solution, we can deduce that $\lambda_1 < 0$, because $\lambda_3(\cdot)$ is always negative and must eventually over pass it. To simplify notation, consider $\lambda_2 := -\lambda$, then $\lambda_1 = 1 - \lambda$. Also, take $\bar{t} \in (0,1)$ such that $\{t \in (0,1) : \lambda_3(t) \leq \lambda_1\} = (0,\bar{t}]$. Then, we can parametrize the space of all optimal functions α by $\lambda, \bar{t} \in [0,1)$ in the following way:

$$\alpha(t) = \begin{cases} 1 & ; 0 \leq t < \bar{t} \\ \frac{\lambda}{1-\lambda} \int\limits_t^1 \mathrm{e}^{\int_0^s \ln \alpha^*(r)\mathrm{d}r} \mathrm{d}s & ; \bar{t} \leq t \leq 1\,. \end{cases}$$

Note that we replaced $x_3(s) = \int_0^s \ln \alpha(r)\mathrm{d}r$ since $x_3(0) = 0$ and $\dot{x}_3(t) = \ln \alpha(t)$, for all $t \in (0,1)$.

Finally, since for $s \geq \bar{t}$, $\int_0^s \ln \alpha^*(r)\mathrm{d}r = \int_{\bar{t}}^s \ln \alpha^*(r)\mathrm{d}r$, therefore we can further describe α by $K > 0$ and $\bar{t} \in (0,1)$ as follows:

$$\alpha_{K,\bar{t}}(t) = \begin{cases} 1 & ; 0 \leq t < \bar{t} \\ \beta_{K,\bar{t}}(t) & ; \bar{t} \leq t \leq 1\,, \end{cases}$$

where $\beta_{K,\bar{t}}$ is the solution of the following ODE:

$$\begin{cases} \dot{\beta}(t) = -K \exp\left[\int\limits_{\bar{t}}^t \ln \beta(s)\mathrm{d}s\right] & t \in (\bar{t},1) \\ \beta(1) = 0\,. \end{cases}$$

With this, we have been able to reduce the space of infinite dimension to a two dimensional space where we can find the optimal solution of the Mayer problem.

Chapter 4

Single Threshold

As a warm-up exercise, we illustrate the main ideas in this thesis by providing an alternative proof of a recent result by Eshani et al. [9]. Consider the blind strategy given by $\alpha \equiv p$, where $p \in [0,1]$ is a fixed number (taking $p = 1/e$ gives exactly the single threshold algorithm of Eshani et al.).

Teorema 4.1 *Given $\alpha \equiv p$, for all $t \geq 0$,*

$$\mathbb{P}(V_{\sigma_T} > t) \geq \min\left\{1-p, \frac{1-p}{-\ln p}\right\} \mathbb{P}(\max_{i \in [n]}\{V_i\} > t)\,.$$

Proof. Recall that given an instance $V_1, \ldots, V_n$, the blind strategy α first computes τ such that $\mathbb{P}(\max_{i\in[n]}\{V_i\} \leq \tau) = p$ and then uses $TTA_{\tau_1=\tau,\ldots,\tau_n=\tau}$, which simply stops the first time a value above τ is observed.

Note that for $t \leq \tau$, we have that

$$\mathbb{P}(V_{\sigma_T} > t) = \mathbb{P}(\max_{i\in[n]}\{V_i\} > \tau) \qquad = 1-p\,,$$

because whenever we pick something, that value is above τ.

Now, for $t > \tau$,

$$\begin{aligned}
\mathbb{P}(V_{\sigma_T} > t) &= \sum_{i\in[n]} \mathbb{P}(V_i > t | \sigma_T = i)\mathbb{P}(\sigma_T = i) \\
&= \sum_{i\in[n]} \frac{\mathbb{P}(V_i > t)}{\mathbb{P}(V_i > \tau)}\mathbb{P}(\sigma_T = i) && \text{independence} \\
&= \sum_{i\in[n]} \mathbb{P}(V_i > t)\mathbb{P}(\sigma_T = i | V_i > \tau) \\
&\geq \left(\frac{1-p}{-\ln p}\right) \sum_{i\in[n]} \mathbb{P}(V_i > t) && \text{Lemma 4.3} \\
&\geq \left(\frac{1-p}{-\ln p}\right) \mathbb{P}(\max_{i\in[n]}\{V_i\} > t) && \text{union bound}\,.
\end{aligned}$$

Then, taking the minimum of both cases we have the result for any $t \geq 0$. □

For a nonnegative random variable V we have that $\mathbb{E}(V) = \int_0^\infty \mathbb{P}(V > t)\mathrm{d}t$. Thus, an immediate consequence of Theorem 4.1 is a result of Eshani et al. [9].

Corolario 4.2 ([9]) *Take* $\alpha \equiv \frac{1}{\mathrm{e}}$, *then* $\mathbb{E}(V_{\sigma_T}) \geq \left(1 - \frac{1}{\mathrm{e}}\right)\mathbb{E}(\max_{\mathrm{i}\in[n]}\{V_\mathrm{i}\})$.

To complete the previous proof, we prove the following lemma.

Lema 4.3 *Consider* $V_1, \ldots, V_n$ *independent random variables and* σ *an independent random uniform permutation of* $[n]$. *Let* T *be the stopping time of* $TTA_{\tau_1=\tau,\ldots,\tau_n=\tau}$ *and* $p = \mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_\mathrm{i}\} \leq \tau)$, *then for all* $\mathrm{i} \in [n]$ *such that* $\mathbb{P}(V_\mathrm{i} > \tau) > 0$ *we have that*

$$\mathbb{P}(\sigma_T = \mathrm{i} | V_\mathrm{i} > \tau) \geq \frac{1-p}{-\ln p}.$$

Proof. Fix $\mathrm{i} \in [n]$ and denote the distribution of V_j by F_j, then

$$\begin{aligned}
\mathbb{P}(\sigma_T = \mathrm{i} | V_\mathrm{i} > \tau) &= \sum_{S\subseteq[n]\setminus\{\mathrm{i}\}} \mathbb{P}(\sigma_T = \mathrm{i}, S \cup \{\mathrm{i}\} = \{j : V_j > \tau\} | V_\mathrm{i} > \tau) \\
&= \sum_{S\subseteq[n]\setminus\{\mathrm{i}\}} \frac{1}{|S|+1} \prod_{j\in S} 1 - F_j(\tau) \prod_{j\in[n]\setminus(S\cup\{\mathrm{i}\})} F_j(\tau) \qquad \text{independence} \\
&= \prod_{j\in[n]\setminus\{\mathrm{i}\}} F_j(\tau) \sum_{S\subseteq[n]\setminus\{\mathrm{i}\}} \frac{1}{|S|+1} \prod_{j\in S} \frac{1-F_j(\tau)}{F_j(\tau)} \\
&= \frac{p}{F_\mathrm{i}(\tau)} \sum_{S\subseteq[n]\setminus\{\mathrm{i}\}} \frac{1}{|S|+1} \prod_{j\in S} \frac{1-F_j(\tau)}{F_j(\tau)} \qquad \text{def. of } \tau\,.
\end{aligned}$$

Forget about i for an instant and consider variables $(y_j)_{j\in[n]\setminus\{\mathrm{i}\}}$ such that $\mathrm{e}^{-y_j} = F_j(\tau)$. Then, define $\phi : \mathbb{R}_+^{n-1} \longrightarrow \mathbb{R}$ by

$$\phi(y) := \sum_{S\subseteq[n-1]} \frac{1}{|S|+1} \prod_{j\in S} \frac{1-\mathrm{e}^{-y_j}}{\mathrm{e}^{-y_j}} = \sum_{S\subseteq[n-1]} \frac{1}{|S|+1} \prod_{j\in S} \mathrm{e}^{y_j} - 1$$

and $\beta := -\ln p + \ln F_\mathrm{i}(\tau)$. Thus, we have that

$$\mathbb{P}(\sigma_T = \mathrm{i} | V_\mathrm{i} > \tau) \geq \frac{p}{F_\mathrm{i}(\tau)} \min\left\{\phi(y) : \sum_{j\in[n-1]} y_j = \beta, y_j \geq 0\right\}.$$

As seen in section 2.3, ϕ is Schur-convex and therefore the minimum is achieved with the constant vector $y = \beta/(n-1)\mathbb{1}$. Consequently, for fixed $F_\mathrm{i}(\tau)$, and under the constraint that $\prod_{j\in[n]\setminus\{\mathrm{i}\}} F_j(\tau) = p/F_\mathrm{i}(\tau)$, the quantity $\mathbb{P}(\sigma_T = \mathrm{i}|V_\mathrm{i} > \tau)$ is minimal when, for all $j \neq \mathrm{i}$, $F_j(\tau) = (p/F_\mathrm{i}(\tau))^{\frac{1}{n-1}}$.

It follows that, since σ and V_{i} are independent,

$$\begin{aligned}
\mathbb{P}(\sigma_T = \mathrm{i} | V_{\mathrm{i}} > \tau) &= \frac{1}{n} \sum_{j=1}^{n} \mathbb{P}(\sigma_T = \mathrm{i} | V_{\mathrm{i}} > \tau, \sigma_j = \mathrm{i}) \\
&\geq \frac{1}{n} \sum_{j=0}^{n-1} \left(\frac{p}{F_{\mathrm{i}}(\tau)} \right)^{\frac{j}{n-1}} && \text{random order} \\
&\geq \frac{1}{n} \sum_{j=0}^{n-1} p^{\frac{j}{n-1}} && F_{\mathrm{i}}(\tau) \in [0,1] \\
&= \frac{1}{n} \frac{1 - p^{\frac{n}{n-1}}}{1 - p^{\frac{1}{n-1}}} .
\end{aligned}$$

Now we note that the left hand side does not depend on n: we can add some dummy variables $(V_{n+1}, V_{n+2}, \ldots \equiv 0)$ and the probability does not change. Therefore, taking limit on $n \longrightarrow \infty$ we get

$$\mathbb{P}(\sigma_T = \mathrm{i} | V_{\mathrm{i}} > \tau) \geq \frac{1-p}{-\ln p} .$$

□

In this Lemma, we have seen how Schur-convexity allowed us to prove an inequality which was used to deduce a lower bound over the performance of certain blind strategy. The same methodology is used when proving a lower bound for general blind strategies, using the concepts of Schur-convexity to prove inequalities which are used in the analysis of lower bounds for their performance.

In terms of upper bounds, it is known that constant thresholds can not obtain a better guarantee than $1 - 1/\mathrm{e}$. This analysis is already done by Eshani et al. [9] studying a specific instance and optimizing over all possible thresholds. This upper bound can be proved to hold for a more general class of algorithms in the work of Correa et al. [7].

Chapter 5

Beating $1 - \frac{1}{\mathrm{e}}$

Since a single threshold (or a constant blind strategy) can already achieve a performance of $1 - 1/\mathrm{e}$, would it not be natural for two thresholds, applying the second after some time, to perform better? This is the key question that led to the definition of blind strategies, which consist in a particular way of defining thresholds to use along the process of discovering values. In this chapter we improve upon the bound of $1 - 1/\mathrm{e}$ by presenting first the answer to the previous question and then by using the idea of blind strategies previously described. Notice that a single threshold is a deterministic algorithm and blind strategies are stochastic, therefore, both approaches might have their own interest.

5.1 Two-thresholds

In this section we present one way of analyzing the use of two thresholds, using one for some fixed portion of the game and then switching to the other. The analysis is easy to generalize to a fixed amount of thresholds and blind strategies are simply the result of taking an infinite number of thresholds. To be more precise, consider $\lambda \in (0,1)$ and tw thresholds τ_1, τ_2, then we are thinking in applying the TTA with the threshold τ_1 until we get to time $\lfloor \lambda n \rfloor + 1$, when we start using the threshold τ_2.

By using only Lemma 5.8 (which has nothing to do with Schur-convexity), we can get the following result:

Teorema 5.1 *Given $\tau_1 \geq \tau_2$ and defining T as the stopping time corresponding to the strategy $TFTA_{\tau_1=\tau_1,\dots,\tau_{\lambda n}=\tau_1,\tau_{\lambda n+1}=\tau_2,\dots,\tau_n=\tau_2}$, we have that*

$$\mathbb{P}(V_{\sigma_T} > t) \geq \begin{cases} \mathbb{P}(T < \infty) & ; \quad t \leq \tau_2 \\ \frac{\mathbb{P}(T \leq \lambda n)}{1-\beta} \, \mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\} > t) + \left[\frac{1}{n}\sum_{k>\lambda n}^{n} \mathbb{P}(T > k)\right] \sum_{\mathrm{i}\in[n]} \mathbb{P}(V_{\mathrm{i}} > t) & ; \quad \tau_2 < t \leq \tau_1 \\ \left[\frac{1}{n}\sum_{k=1}^{n} \mathbb{P}(T > k)\right] \sum_{\mathrm{i}\in[n]} \mathbb{P}(V_{\mathrm{i}} > t) & ; \quad \tau_1 < t\,, \end{cases}$$

where $\beta := \mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\} \leq \tau_2)$.

Proof. The proof consists in analyzing each of the relevant intervals and applying Lemma 5.8. For $t \leq \tau_2$,

$$\mathbb{P}(V_{\sigma_T} > t) = \mathbb{P}(T < \infty)\,.$$

For $\tau_2 < t \leq \tau_1$,

$$\begin{aligned}\mathbb{P}(V_{\sigma_T} > t) &= \mathbb{P}(V_{\sigma_T} > t, T \leq \lambda n) + \mathbb{P}(V_{\sigma_T} > t, T > \lambda n)\\ &= \mathbb{P}(T \leq \lambda n) + \mathbb{P}(V_{\sigma_T} > t, T > \lambda n)\\ &\geq \frac{\mathbb{P}(T \leq \lambda n)}{1-\beta}\mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\} > t) + \mathbb{P}(V_{\sigma_T} > t, T > \lambda n)\,.\end{aligned}$$

Note that, for $t \geq \tau_2$,

$$\begin{aligned}\mathbb{P}(V_{\sigma_T} > t, T > \lambda n) &= \sum_{\mathrm{i}\in[n]} \mathbb{P}(V_{\mathrm{i}} > t, \sigma_T = \mathrm{i}, T > \lambda n)\\ &= \sum_{\mathrm{i}\in[n]} \sum_{k>\lambda n}^{n} \mathbb{P}(V_{\mathrm{i}} > t, \sigma_k = \mathrm{i}, T \geq k)\\ &= \sum_{\mathrm{i}\in[n]} \mathbb{P}(V_{\mathrm{i}} > t) \sum_{k>\lambda n}^{n} \mathbb{P}(\sigma_T = \mathrm{i}, T \geq k) && \text{independence}\\ &\geq \left[\frac{1}{n}\sum_{k>\lambda n}^{n} \mathbb{P}(T > k)\right] \sum_{\mathrm{i}\in[n]} \mathbb{P}(V_{\mathrm{i}} > t) && \text{Lemma 5.8}\,.\end{aligned}$$

Finally, for $t > \tau_1$,

$$\begin{aligned}\mathbb{P}(V_{\sigma_T} > t) &= \sum_{\mathrm{i}\in[n]} \mathbb{P}(V_{\mathrm{i}} > t, \sigma_T = \mathrm{i})\\ &= \sum_{\mathrm{i}\in[n]} \sum_{k=1}^{n} \mathbb{P}(V_{\mathrm{i}} > t, \sigma_k = \mathrm{i}, T \geq k)\\ &= \sum_{\mathrm{i}\in[n]} \mathbb{P}(V_{\mathrm{i}} > t) \sum_{k=1}^{n} \mathbb{P}(\sigma_T = \mathrm{i}, T \geq k) && \text{independence}\\ &\geq \left[\frac{1}{n}\sum_{k=1}^{n} \mathbb{P}(T > k)\right] \sum_{\mathrm{i}\in[n]} \mathbb{P}(V_{\mathrm{i}} > t)\,,\end{aligned}$$

so we get the stated result. □

Now, we introduce Lemma 5.7, which is proved using the ideas of Schur-convexity to get the following result:

Corolario 5.2 *Given $\tau_1 \geq \tau_2$ such that $\mathbb{P}(V_1, ..., V_n \leq \tau_1) = \alpha, \mathbb{P}(V_1, ..., V_n \leq \tau_2) = \beta$ and defining T as the stopping time corresponding to $TFTA_{\tau_1=\tau_1,...,\tau_{\lambda n}=\tau_1,\tau_{\lambda n+1}=\tau_2,...,\tau_n=\tau_2}$, we have that for $\lambda \in \{\frac{k}{n} : k \in [n]\}$*

$$\mathbb{E}(V_{\sigma_T}) \geq factor(\alpha, \beta, \lambda)\mathbb{E}(\max_{i\in[n]}\{V_i\}),$$

where

$$factor(\alpha, \beta, \lambda) := \min\{1 - \lambda\alpha - (1-\lambda)\beta, \lambda\frac{1-\alpha}{1-\beta} + \alpha^\lambda\frac{1-\beta^{(1-\lambda)}}{-\ln\beta}, \frac{1-\alpha^\lambda}{-\ln\alpha} + \alpha^\lambda\frac{1-\beta^{(1-\lambda)}}{-\ln\beta}\}.$$

In particular, $\exists\alpha, \beta, \lambda$ such that

$$\mathbb{E}(V_{\sigma_T}) \geq 0.6494 \quad \mathbb{E}(\max_{i\in[n]}\{V_i\}).$$

The proof consists in replacing the terms concerning the distribution of T in Theorem 5.1 in terms of α and β using Lemma 5.7.

Proof. Notice that, by Lemma 5.7

$$\mathbb{P}(T < \infty) = 1 - \mathbb{P}(T > n) \geq 1 - \lambda\alpha - (1-\lambda)\beta$$

and

$$\mathbb{P}(T \leq \lambda n) = 1 - \mathbb{P}(T > \lambda n) \geq 1 - \lambda\alpha - (1-\lambda) = \lambda(1-\alpha)$$

Also, we have that

$$\frac{1}{n}\sum_{k>\lambda n}^{n}\mathbb{P}(T > k) \geq \alpha^\lambda\frac{1}{n}\sum_{k>\lambda n}^{n}\beta^{\frac{k-\lambda n}{n}} = \alpha^\lambda\frac{\beta^{\frac{1}{n}}}{n}\frac{1-\beta^{(1-\lambda)}}{1-\beta^{\frac{1}{n}}} \xrightarrow[n\to\infty]{} \alpha^\lambda\frac{1-\beta^{(1-\lambda)}}{-\ln\beta}$$

and similarly,

$$\begin{aligned}\frac{1}{n}\sum_{k=1}^{n}\mathbb{P}(T > k) &= \frac{1}{n}\sum_{k=1}^{\lambda n}\mathbb{P}(T > k) + \frac{1}{n}\sum_{k>\lambda n}^{n}\mathbb{P}(T > k) \\ &\geq \frac{1}{n}\sum_{k=1}^{\lambda n}\alpha^{\frac{k}{n}} + \alpha^\lambda\frac{\beta^{\frac{1}{n}}}{n}\frac{1-\beta^{(1-\lambda)}}{1-\beta^{\frac{1}{n}}} \\ &= \frac{\alpha^{\frac{1}{n}}}{n}\frac{1-\alpha^\lambda}{1-\alpha^{\frac{1}{n}}} + \alpha^\lambda\frac{\beta^{\frac{1}{n}}}{n}\frac{1-\beta^{(1-\lambda)}}{1-\beta^{\frac{1}{n}}} \\ &\xrightarrow[n\to\infty]{} \frac{1-\alpha^\lambda}{-\ln\alpha} + \alpha^\lambda\frac{1-\beta^{(1-\lambda)}}{-\ln\beta}.\end{aligned}$$

Integrating over t, we get the result. For the numerical performance, notice that taking $\alpha = 0.4126, \beta = 0.2185, \lambda = 0.6793$, we get

$$factor(\alpha, \beta, \lambda) = 0.6496.$$

□

5.2 Blind strategies

In this section we present how to use the idea of blind strategies previously described. To this end, we find an appropriate stochastic dominance type bound on the probability that the gambler obtains at least a certain amount with respect to the probability that the prophet obtains the same amount. This bound is presented in Theorem 5.6, and improved in Lemma 6.1, where we manage to make the bound solely dependent on the blind strategy. Then optimizing over blind strategies leads to the improved bound of 0.665, and 0.669 respectively. The analysis of these bounds are separated for the sake of clarity and the second bound is presented in the following chapter since it is more technical. As for the rest of the thesis we assume for simplicity that $F_1, F_2, ..., F_n$ are continuous (see Section 8 for an explanation on how to extend the results to the discontinuous case). In summary we prove the following result.

Teorema 5.3 *There exists a nonincreasing function $\alpha : [0,1] \to [0,1]$ such that*

$$\mathbb{E}(V_{\sigma_T}) \geq 0.669\ \mathbb{E}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\}),$$

where T is the stopping time of the blind strategy α.

To do the analysis we first need to note that a blind strategy can be interpreted as the limit, as the size of the instance goes to infinity, of strategies that do not use randomization.

Definición 5.4 *Consider a nonincreasing function $\alpha : [0,1] \to [0,1]$. The deterministic blind strategy given by α is the strategy that applies $TTA_{\tau_1,\ldots,\tau_n}$ to the sequence of thresholds $\tau_1, \ldots, \tau_n$ defined by the following conditions:*

$$\forall j \in [n], \quad \mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\} \leq \tau_j) = \alpha\left(\frac{j}{n}\right).$$

To turn a deterministic blind strategy into a blind strategy consider an instance $F_1, \ldots, F_n$ and add to this instance m deterministic random variables equal to zero, $F_{n+\mathrm{i}} = \mathbb{1}_{[0,\infty)}$ for $\mathrm{i} = 1, \ldots, m$, so that the new instance becomes $F_1, \ldots, F_n, F_{n+1}, \ldots, F_{n+m}$. Denoting by T_m the stopping time given by the deterministic blind strategy applied to this instance, we have that

$$\lim_{m\to\infty} \mathbb{E}(V_{\sigma_{T_m}}) = \mathbb{E}(V_{\sigma_T}).$$

Indeed, recalling the definition of blind strategies in Section 1.3, it is easy to see that a deterministic blind strategy applied to instance $F_1, \ldots, F_{n+m}$ is approximately a blind strategy applied to instance $F_1, \ldots, F_n$, where the random variables $u_1, \ldots, u_n$ are drawn from $U(\{\frac{1}{n+m}, \ldots, 1\})$, rather than from $U(0,1)$. Thus, by taking the limit as $m \to \infty$ the claim follows.

The conclusion of this remark is that in order to analyze the performance of blind strategies it is sufficient to study the performance of deterministic blind strategies and then take the limit as n grows to infinity.

We are now ready to start analyzing deterministic blind strategies.

Lema 5.5 *Given an instance $F_1, F_2, ..., F_n$ and nonincreasing thresholds $\infty = \tau_0 \geq \tau_1 \geq \cdots \geq \tau_n \geq \tau_{n+1} = -\infty$, it holds that, for $j \in [n+1]$ and $t \in [\tau_j, \tau_{j-1})$,*

$$\mathbb{P}(V_{\sigma_T} > t) = \mathbb{P}(T \leq j-1) + \sum_{\mathrm{i} \in [n]} \mathbb{P}(V_{\mathrm{i}} > t) \left(\sum_{k > j-1}^{n} \frac{\mathbb{P}(T \geq k | \sigma_k = \mathrm{i})}{n} \right),$$

where T is the stopping time given by $TTA_{\tau_1,\ldots,\tau_n}$.

Proof. Notice that, since thresholds are nonincreasing,

$$\mathbb{P}(V_{\sigma_T} > t) = \mathbb{P}(T \leq j-1) + \mathbb{P}(V_{\sigma_T} > t, T \geq j)$$

and therefore, we simply must analyze the second term.

$$\begin{aligned}
\mathbb{P}(V_{\sigma_T} > t, T \geq j) &= \sum_{\mathrm{i} \in [n]} \mathbb{P}(V_{\mathrm{i}} > t, \sigma_T = \mathrm{i}, T \geq j) \\
&= \sum_{\mathrm{i} \in [n]} \sum_{k=j}^{n} \mathbb{P}(V_{\mathrm{i}} > t, \sigma_k = \mathrm{i}, T = k) \\
&= \sum_{\mathrm{i} \in [n]} \sum_{k=j}^{n} \mathbb{P}(V_{\mathrm{i}} > t, \sigma_k = \mathrm{i}, T \geq k) \\
&= \sum_{\mathrm{i} \in [n]} \mathbb{P}(V_{\mathrm{i}} > t) \left(\sum_{k=j}^{n} \mathbb{P}(\sigma_k = \mathrm{i}, T \geq k) \right) \\
&= \sum_{\mathrm{i} \in [n]} \mathbb{P}(V_{\mathrm{i}} > t) \left(\frac{1}{n} \sum_{k=j}^{n} \mathbb{P}(T \geq k | \sigma_k = \mathrm{i}) \right).
\end{aligned}$$

□

Note that for a nonincreasing blind strategy α, the previous lemma holds since the thresholds used are also nonincreasing. In the rest of this section we will present a not so technical way to derive a function α with a guarantee of 0.665, relying on two lemmata shown in section 5.3. The first of these, Lemma 5.7, gives precise bounds for the distribution of the stopping time of T of $TTA_{\tau_1,\ldots,\tau_n}$, namely:

$$\frac{1}{n} \sum_{j \in [k]} 1 - \alpha_j \leq \mathbb{P}(T \leq k) \leq 1 - \left(\prod_{l=1}^{k} \alpha_l \right)^{\frac{1}{n}}.$$

The second result, Lemma 5.8, (inspired by a result of Esfandiari et al. [11]) bounds the distribution of T conditional on the k-th element in the random permutation.

$$\mathbb{P}(T \geq k | \sigma_k = \mathrm{i}) \geq \frac{\mathbb{P}(T > k)}{1 - \frac{k}{n} + \frac{1}{n} \sum_{l \in [k]} \mathbb{P}(V_{\mathrm{i}} \leq \tau_l)}.$$

Notice that the denominator is positive and less than one. In this section we deduce a lower bound for the performance of blind strategies by using the simpler inequality $\mathbb{P}(T \geq k | \sigma_k =$

i) $\geq \mathbb{P}(T > k)$. This leads us to use the union bound $\sum_{\mathrm{i}\in[n]} \mathbb{P}(V_{\mathrm{i}} > t) \geq \mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\} > t)$ and then we conclude the existence of a blind strategy with an overall performance of at least 0.665. This approach does not take full advantage of Lemma 5.8 and the more involved analysis is done in the next section, where we also replace the union bound by a more refined bound given by Lemma 6.3.

Teorema 5.6 *Let $\alpha : [0,1] \to [0,1]$ be nonincreasing, and let T be the* deterministic blind strategy *stopping time. For every instance $F_1, \ldots, F_n$ and $t > 0$,*

$$\mathbb{P}(V_{\sigma_T} > t) \geq \min_{j\in[n+1]} \{f_j(\alpha)\} \, \mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\} > t)\,,$$

where, for all $j \in [n+1]$, taking $\alpha(\frac{n+1}{n}) = 0$,

$$f_j(\alpha) = \sum_{k=1}^{j-1} \frac{1-\alpha\left(\frac{k}{n}\right)}{n\left(1-\alpha\left(\frac{j}{n}\right)\right)} + \frac{1}{n}\sum_{k=j}^{n}\left(\prod_{l=1}^{k}\alpha\left(\frac{l}{n}\right)\right)^{\frac{1}{n}}.$$

Proof. Since α is nonincreasing, we can apply Lemma 5.5. Fix $j \in [n+1]$ and $t \in [\tau_j, \tau_{j-1})$, by Lemma 5.7 we have that

$$\begin{aligned}\mathbb{P}(T \leq j-1) &\geq \sum_{k=1}^{j-1}\frac{1-\alpha\left(\frac{k}{n}\right)}{n}\\ &\geq \sum_{k=1}^{j-1}\frac{1-\alpha\left(\frac{k}{n}\right)}{n\left(1-\alpha\left(\frac{j}{n}\right)\right)}\mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\} > t)\,,\end{aligned}$$

since $t \in [\tau_j, \tau_{j-1})$ implies that $1-\alpha(j/n) = \mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\} > \tau_j) \geq \mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\} > t)$.

On the other hand, by Lemma 5.8, Lemma 5.7 and the union bound,

$$\begin{aligned}\sum_{\mathrm{i}\in[n]}\mathbb{P}(V_{\mathrm{i}} > t)\left(\frac{1}{n}\sum_{k=j}^{n}\mathbb{P}(T \geq k|\sigma_k = \mathrm{i})\right) &\geq \left(\frac{1}{n}\sum_{k=j}^{n}\mathbb{P}(T > k)\right)\sum_{\mathrm{i}\in[n]}\mathbb{P}(V_{\mathrm{i}} > t)\\ &\geq \left(\frac{1}{n}\sum_{k=j}^{n}\left(\prod_{l=1}^{k}\alpha\left(\frac{l}{n}\right)\right)^{\frac{1}{n}}\right)\sum_{\mathrm{i}\in[n]}\mathbb{P}(V_{\mathrm{i}} > t)\\ &\geq \frac{1}{n}\sum_{k=j}^{n}\left(\prod_{l=1}^{k}\alpha\left(\frac{l}{n}\right)\right)^{\frac{1}{n}}\mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\} > t)\,.\end{aligned}$$

Both bounds together prove the theorem. □

Thus, for every n, we get a lower bound on the performance of a deterministic blind strategy α, that only depends on $\alpha(\frac{1}{n}), \ldots, \alpha(\frac{n}{n})$. As we explained before, we only care about the performance of this strategy when n tends to $+\infty$. Assuming that α is continuous, a standard Riemann sum analysis shows that

$$\lim_{n\to\infty}\min_{j\in[n+1]}\{f_j(\alpha)\} = \min\left\{\int_0^1 1-\alpha(y)\mathrm{d}y\,,\ \inf_{x\in[0,1]}\int_0^x \frac{1-\alpha(y)}{1-\alpha(x)}\mathrm{d}y + \int_x^1 \mathrm{e}^{\int_0^y \ln\alpha(w)\mathrm{d}w}\mathrm{d}y\right\}. \tag{5.1}$$

Thus, in order to prove Theorem 5.3, we would like to find a blind strategy α maximizing the latter expression. As this is a nontrivial optimal control problem we aim at finding a function α such that the above expression is larger than 0.665. As we already said, the 0.669 bound is proved in the next section.

Remark. Consider α being constant equal to $1/\mathrm{e}$. Then the above quantity is equal to $1 - 1/\mathrm{e}$. Thus, we recover the one-threshold result in Corollary 4.2. Furthermore, if for instance we take $\alpha(x) = 0.53 - 0.38x$ the guarantee of the strategy (given by expression (5.1)) is greater than 0.657. This gives an explicit α that beats significantly $1 - 1/\mathrm{e}$.

To maximize over expression (5.1), we resort to a numerical approximation. Note that if α is such that $\alpha(1) = 0$ and $x \mapsto \int_0^x \frac{1-\alpha(y)}{1-\alpha(x)} \mathrm{d}y + \int_x^1 \mathrm{e}^{\int_0^y \ln \alpha(w) \mathrm{d}w} \mathrm{d}y$ is a constant, then this constant is a lower bound for the infimum in (5.1).

Consequently, we solve the following integro-differential equation:

$$\begin{cases} \frac{\mathrm{d}}{\mathrm{d}x} \left(\int_0^x \frac{1-\alpha(y)}{1-\alpha(x)} \mathrm{d}y + \int_x^1 \mathrm{e}^{\int_0^y \ln \alpha(w) \mathrm{d}w} \mathrm{d}y \right) = 0 & ; x \in (0,1) \\ \alpha(1) = 0. \end{cases}$$

To this end we consider a change of variables leading to the following second order ODE:

$$\begin{cases} (u'(x))^2 K(x,u) - u''(x)u(x) = 0 & ; x \in (0,1) \\ u'(1) = 1 \\ u(0) = 0, \end{cases}$$

where

$$u(x) := \int_0^x 1 - \alpha(x) \mathrm{d}x \quad \text{and} \quad K(x,u) := 1 - \exp\left(\int_0^x \ln(1 - u'(t)) \mathrm{d}t \right).$$

We approximately solved this equation by taking an initial guess u_0 and defining u_{n+1} as the solution to $(u'(x))^2 K(x, u_n) - u''(x)u(x) = 0$. To be more precise, the initial guess u_0 was the result of maximizing over α $\min_{j \in [n+1]} \{f_j(\alpha)\}$, given in Theorem 5.6, for $n = 23$. Then, we iterated the process eleven times and obtained an α with $\alpha(1) = 0$ and such that the function $x \mapsto \int_0^x \frac{1-\alpha(y)}{1-\alpha(x)} \mathrm{d}y + \int_x^1 \exp\left(\int_0^y \ln \alpha(w) \mathrm{d}w \right) \mathrm{d}y$ varies between 0.6653 and 0.6720. Even if we did not find an exact solution for the ODE, its performance is given by computing expression (5.1), which gives the claimed factor of 0.665.

5.3 Inequalities

We now present the two lemmata used in the previous sections. The first establishes bounds on the distribution of the stopping time T. Both the lower and upper bounds are sharp in the sense that they are achieved by different instances: the lower bound corresponds to the case where there is only one non-zero variable and the upper bound corresponds to the case where all distributions are equal.

Lema 5.7 *Consider* $\alpha_1, \ldots, \alpha_n \in [0,1]$ *fixed. For every instance* $F_1, \ldots, F_n$ *consider* $\tau_1, \ldots, \tau_n$ *the sequence of thresholds such that*

$$\mathbb{P}(\max_{\mathrm{i} \in [n]}\{V_{\mathrm{i}}\} \leq \tau_{\mathrm{i}}) = \alpha_{\mathrm{i}}.$$

Denoting T *the stopping time of* $TTA_{\tau_1,\ldots,\tau_n}$*, we have that* $\forall k \in [n]$

$$\frac{1}{n}\sum_{j \in [k]} 1 - \alpha_j \leq \mathbb{P}(T \leq k) \leq 1 - \left(\prod_{l=1}^{k} \alpha_l\right)^{\frac{1}{n}}.$$

Proof. The proof consists in highlighting the role of F_1 and F_2 in $\mathbb{P}(T > k)$ and using the symmetry that the random order σ induces. We focus on three different cases according to σ:

1. $\sigma^{-1}(1) \leq k \veebar \sigma^{-1}(2) \leq k$, i.e. : only one of the variables V_1 and V_2 shows before time k.
2. $\sigma^{-1}(1) \leq k \wedge \sigma^{-1}(2) \leq k$, i.e. : both V_1 and V_2 show before time k.
3. $\sigma^{-1}(1) > k \wedge \sigma^{-1}(2) > k$, i.e. : neither V_1 nor V_2 shows before time k.

To express this formally, denote

$$\begin{aligned}\Sigma(k) &:= \{\sigma,\ \text{ordered subset of } [n] \text{ with size } k\} \\ \Sigma_{-1,-2}(k) &:= \{\sigma,\ \text{ordered subset of } [n] \setminus \{1,2\} \text{ with size } k\},\end{aligned}$$

then if $\sigma \in \Sigma(k)$, we have that either

1. $\exists p \in [k], \exists \mathrm{i} \in \{1,2\} \quad$ s.t. $\sigma_p = \mathrm{i}$ and
$$(\sigma_j)_{j \in [k]\setminus\{p\}} \in \Sigma_{-1,-2}(k-1).$$
2. $\exists p < q \in [k] \quad$ s.t. $\{\sigma_p, \sigma_q\} = \{1,2\}$ and
$$(\sigma_j)_{j \in [k]\setminus\{p,q\}} \in \Sigma_{-1,-2}(k-2).$$
3. $\sigma \in \Sigma_{-1,-2}(k)$.

This is the key decomposition we use to show the inequality. Note that

$$
\begin{aligned}
&\mathbb{P}(T > k; F_1, \dots, F_n) \\
&= \frac{1}{|\Sigma(k)|} \sum_{\sigma \in \Sigma(k)} \prod_{\mathrm{i} \in [k]} F_{\sigma_\mathrm{i}}(\tau_\mathrm{i}) \\
&= \frac{(n-k)!}{n!} \left(\sum_{\sigma \in \Sigma_{-1,-2}(k)} \prod_{\mathrm{i} \in [k]} F_{\sigma_\mathrm{i}}(\tau_\mathrm{i}) \right. \\
&+ \sum_{\substack{\sigma \in \Sigma_{-1,-2}(k-1) \\ p \in [k]}} \prod_{\mathrm{i}=1}^{p-1} F_{\sigma_\mathrm{i}}(\tau_\mathrm{i}) \left[F_1(\tau_p) + F_2(\tau_p) \right] \prod_{\mathrm{i}=p}^{k-1} F_{\sigma_\mathrm{i}}(\tau_{\mathrm{i}+1}) \\
&+ \left. \sum_{\substack{\sigma \in \Sigma_{-1,-2}(k-2) \\ p < q \in [k]}} \prod_{\mathrm{i}=1}^{p-1} F_{\sigma_\mathrm{i}}(\tau_\mathrm{i}) \left[F_1(\tau_p) F_2(\tau_q) + F_2(\tau_p) F_1(\tau_q) \right] \prod_{\mathrm{i}=p}^{q-1} F_{\sigma_\mathrm{i}}(\tau_{\mathrm{i}+1}) \prod_{\mathrm{i}=q}^{k-1} F_{\sigma_\mathrm{i}}(\tau_{\mathrm{i}+2}) \right) .
\end{aligned}
$$

To simplify the notation, let us define

$$
\begin{aligned}
A(F_1, F_2) &:= \sum_{\substack{\sigma \in \Sigma_{-1,-2}(k-1) \\ p \in [k]}} \left[F_1(\tau_p) + F_2(\tau_p) \right] \prod_{\mathrm{i} \in [k-1]} F_{\sigma_\mathrm{i}}(\tau_{\mathrm{i} + \mathbb{1}_{\mathrm{i} \geq p}}) \\
B(F_1, F_2) &:= \sum_{\substack{\sigma \in \Sigma_{-1,-2}(k-2) \\ p < q \in [k]}} \left[F_1(\tau_p) F_2(\tau_q) + F_2(\tau_p) F_1(\tau_q) \right] \prod_{\mathrm{i} \in [k-1]} F_{\sigma_\mathrm{i}}(\tau_{\mathrm{i} + \mathbb{1}_{\mathrm{i} \geq p} + \mathbb{1}_{\mathrm{i} \geq q}}) \\
C &:= \sum_{\sigma \in \Sigma_{-1,-2}(k)} \prod_{\mathrm{i} \in [k]} F_{\sigma_\mathrm{i}}(\tau_\mathrm{i}) .
\end{aligned}
$$

Then,

$$
\mathbb{P}(T > k; F_1, \dots, F_n) = \frac{(n-k)!}{n!} \left[A(F_1, F_2) + B(F_1, F_2) + C \right] .
$$

Let's show that both A and B change in the correct direction when we change F_1 and F_2, by $F_1 F_2$ and $\mathbb{1}_{\mathbb{R}_+}$, or $\sqrt{F_1 F_2}$ and $\sqrt{F_1 F_2}$, respectively. For this, note that $\forall p \in [k]$

$$
1 + F_1(\tau_p) F_2(\tau_p) \geq F_1(\tau_p) + F_2(\tau_p) \geq 2 \sqrt{F_1(\tau_p) F_2(\tau_p)} ,
$$

and $\forall p < q \in [k]$

$$
F_1(\tau_p) F_2(\tau_p) + F_2(\tau_q) F_1(\tau_q) \geq F_1(\tau_p) F_2(\tau_q) + F_2(\tau_p) F_1(\tau_q) \geq 2 \sqrt{F_1(\tau_p) F_2(\tau_p) F_1(\tau_q) F_2(\tau_q)} .
$$

Then,

$$
A(F_1 F_2, \mathbb{1}_{\mathbb{R}_+}) \geq A(F_1, F_2) \geq A(\sqrt{F_1 F_2}, \sqrt{F_1 F_2}) ,
$$

and

$$
B(F_1 F_2, \mathbb{1}_{\mathbb{R}_+}) \geq B(F_1, F_2) \geq B(\sqrt{F_1 F_2}, \sqrt{F_1 F_2}) .
$$

We can conclude the lower bound by applying the inequality n times and noticing that

$$\mathbb{P}\left(T \le k; \prod_{\mathrm{i}\in[n]} F_{\mathrm{i}}, \mathbb{1}_{\mathbb{R}_+}, \ldots, \mathbb{1}_{\mathbb{R}_+}\right) = \frac{1}{n}\sum_{j\in[k]} 1-\alpha_j\,.$$

The upper bound follows from applying the inequality infinitely many times and noticing that

$$\mathbb{P}\left(T \le k; \prod_{\mathrm{i}\in[n]} F_{\mathrm{i}}^{\frac{1}{n}}, \ldots, \prod_{\mathrm{i}\in[n]} F_{\mathrm{i}}^{\frac{1}{n}}\right) = 1-\prod_{j\in[k]} \alpha_j^{\frac{1}{n}}\,.$$

□

Remember that in the proof of Lemma 4.3 we solved the following optimization problem: $\min\left\{\phi(y);\ s.t. \sum_{j\in[n-1]} y_j = \beta\right\}$. The value of this problem was obtained by noticing that ϕ is Schur convex. This time we considered the problem

$$\begin{cases} \text{opt } \mathbb{P}(T > k; F_1, \ldots, F_n) \\ s.t. \prod_{\mathrm{i}\in[n]} F_{\mathrm{i}} = F \text{ and } F_{\mathrm{i}} \text{ is a distribution.} \end{cases}$$

where "opt" is a symbol in $\{\min, \max\}$. This problem is harder since it involved optimizing over functions rather than real numbers. Trying to apply Schur convexity theory again, one could see $\mathbb{P}(T > k; F_1, \ldots, F_n)$ as a function of the distributions evaluated at each threshold, that is, as a function of the vector $(F_1(\tau_1), \ldots, F_1(\tau_n), \ldots, F_n(\tau_1), \ldots, F_n(\tau_n))$. Unfortunately doing this results in a domain which is no longer symmetric and moreover the constraint of the product being constant results in n different constraints, making it hard to apply the theory.

However, note that the previous lemma shows that $\mathbb{P}(T > k; F_1, \ldots, F_n)$ is nearly log-Schur-convex: it increases when the components of the argument get more concentrated in some coordinate. Nevertheless, the behavior of $\mathbb{P}(T > k)$ is not always monotone along the curve $\lambda \in [0,1] \mapsto (F_1F_2^{\lambda}, F_2^{1-\lambda}, \ldots, F_n)$, a property that would be satisfied by a log-Schur-convex function if $F_1, \ldots, F_n$ were numbers. In spite of the latter, there is a step by step way to go from $(F_1, F_2, \ldots, F_n)$ to $(F_1F_2 \ldots F_n, \mathbb{1}_{\mathbb{R}_+}, \ldots, \mathbb{1}_{\mathbb{R}_+})$ that exhibits a monotonic behavior, while maintaining the product. This property could be called *weak* log-Schur-convexity and it is enough to solve the optimization problem. The same can be said about the points $(F_1, F_2, \ldots, F_n)$ and $(\sqrt[n]{F_1F_2\ldots F_n}, \ldots, \sqrt[n]{F_1F_2\ldots F_n})$.

Lema 5.8 *Given $V_1, \ldots, V_n$ independent random variables and $\tau_1 \ge \ldots \ge \tau_n$ a sequence of nonincreasing thresholds, we denote by T the stopping time of the $TTA_{\tau_1,\ldots,\tau_n}$. Then, $\forall \mathrm{i}, k \in [n]$*

$$\mathbb{P}(T \ge k | \sigma_k = \mathrm{i}) \ge \frac{\mathbb{P}(T > k)}{1-\frac{k}{n}+\frac{1}{n}\sum\limits_{l\in[k]} \mathbb{P}(V_{\mathrm{i}} \le \tau_l)}\,.$$

Proof. Inspired by the proof given by Esfandiari et al. [11], fix $\mathrm{i}, k \in [n]$ and define

$$\Sigma_{-\mathrm{i}}(k) := \{\sigma,\ \text{ordered subset of } [n]\setminus\{\mathrm{i}\} \text{ with size } k\}\,.$$

Then, conditioning on the value of $\sigma^{-1}(\mathrm{i})$,

$$\mathbb{P}(T > k) = \frac{1}{n}\sum_{l\in[k]}\mathbb{P}(T > k|\sigma_l = \mathrm{i}) + \frac{1}{n}\sum_{l=k+1}^{n}\mathbb{P}(T > k|\sigma_l = \mathrm{i})\,. \tag{5.2}$$

Note that for all $l = k+1,\dots,n$,

$$\mathbb{P}(T > k|\sigma_l = \mathrm{i}) = \mathbb{P}(T > k|\sigma_{k+1} = \mathrm{i}) \leq \mathbb{P}(T \geq k|\sigma_k = \mathrm{i})$$

and given $l \in [k]$, since the thresholds are nonincreasing,

$$\begin{aligned}\mathbb{P}(T > k|\sigma_l = \mathrm{i}) &= \frac{\mathbb{P}(V_\mathrm{i} \leq \tau_l)}{|\Sigma_{-\mathrm{i}}(k-1)|}\sum_{\sigma\in\Sigma_{-\mathrm{i}}(k-1)}\prod_{j\in[k-1]}F_{\sigma_j}(\tau_{j+\mathbb{1}_{j\geq l}})\\ &\leq \frac{\mathbb{P}(V_\mathrm{i} \leq \tau_l)}{|\Sigma_{-\mathrm{i}}(k-1)|}\sum_{\sigma\in\Sigma_{-\mathrm{i}}(k-1)}\prod_{j\in[k-1]}F_{\sigma_j}(\tau_j)\\ &= \mathbb{P}(V_\mathrm{i} \leq \tau_l)\mathbb{P}(T \geq k|\sigma_k = \mathrm{i})\,.\end{aligned}$$

Plugging both inequalities back into equation (5.2) we get the result. □

Chapter 6

Improved analysis

To prove Theorem 5.3 (i.e., to improve upon the bound proved in the last section) we turn to a particular type of blind strategies where $\alpha = \alpha_{\alpha_1,\ldots,\alpha_m}$ is given by

$$\alpha_{\alpha_1,\ldots,\alpha_m}(x) = \sum_{j\in[m]} \alpha_j \mathbb{1}_{\left[\frac{j-1}{m},\frac{j}{m}\right)}(x)\,,$$

in other words, piece-wise constant functions.

The idea now is to use the property that, for any instance of size larger than m, this blind strategy uses only m thresholds. In addition, observe that Lemma 5.8 was used in a loose way, so to take more advantage of it we present Lemma 6.3 at the end of this section.

Define for $m \geq 1$

$$g_{m,p}(k) = \begin{cases} \frac{1}{1-\frac{k}{m}(1-p)} & ; k \leq \frac{m}{2} \\ \frac{2}{1+p} & ; k > \frac{m}{2}\,. \end{cases}$$

This function is used in the next lemma. Notice that it is a nondecreasing function in k that is always greater than 1. We can prove the following statement, which has the same flavor as Theorem 5.6.

Lema 6.1 *Let $\alpha = \alpha_{\alpha_1,\ldots,\alpha_m}$ be a nonincreasing function where $\alpha_m > 0$, and let T be the blind strategy stopping time. Then,*

$$\frac{\mathbb{E}(V_{\sigma_T})}{\mathbb{E}(\max\limits_{\mathrm{i}\in[n]}\{V_\mathrm{i}\})} \geq \min_{j\in[m+1]}\{f_j(\alpha_1,\ldots,\alpha_m)\}\,,$$

where is given by

$$f_j(\alpha_1,\ldots,\alpha_m) = \begin{cases} \sum\limits_{k=1}^{m}\left(\prod\limits_{l\in[k-1]}\alpha_l\right)^{\frac{1}{m}}\left(\frac{1-\alpha_k^{\frac{1}{m}}}{-\ln\alpha_k}\right) & ; j = 1 \\ \frac{1}{m}\sum\limits_{k\in[m]} 1-\alpha_k & ; j = m+1 \\ \sum\limits_{k\in[j-1]}\frac{1-\alpha_k}{m(1-\alpha_j)} + \sum\limits_{k=j}^{m}\left(\prod\limits_{l\in[k-1]}\alpha_l\right)^{\frac{1}{m}} g_{m,\alpha_1}(k-1)\left(\frac{1-\alpha_k^{\frac{1}{m}}}{-\ln\alpha_k}\right) & ; 2 \leq j \leq m\,. \end{cases}$$

Proof. As done in section 5, we analyze the performance of the corresponding deterministic blind strategy with an instance of size n and we only care about the performance guarantee of α as n grows to ∞. Consider an instance $F_1, \ldots, F_{Nm}$, and take $j \in [m+1]$ and $t \in [\tau_j, \tau_{j-1})$, where $\tau_0 = \infty$, $\tau_{m+1} = 0$ and $\alpha_{m+1} = 0$. In the same spirit as in Lemma 5.5, we have that

$$\begin{aligned}\mathbb{P}(V_{\sigma_T} > t) &= \mathbb{P}(T \leq (j-1)N) + \mathbb{P}(V_{\sigma_T} > t, T > (j-1)N) \\ &= \mathbb{P}(T \leq (j-1)N) + \sum_{\mathrm{i}\in[n]} \mathbb{P}(V_\mathrm{i} > t) \left(\sum_{k=(j-1)N+1}^{Nm} \frac{\mathbb{P}(T \geq k | \sigma_k = \mathrm{i})}{Nm} \right) .\end{aligned}$$

One key point is that since the same thresholds are used N times, we can deduce better bounds. For the first term, as before, we have that

$$\mathbb{P}(T \leq (j-1)N) \geq \frac{\mathbb{P}(T \leq (j-1)N)}{1-\alpha_j} \mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_\mathrm{i}\} > t) .$$

Then, for $j = m+1$ (i.e. : $t \in [0, \tau_m)$), using Lemma 5.7, we have that

$$\mathbb{P}(V_{\sigma_T} > t) \geq \frac{1}{m} \sum_{k\in[m]} 1 - \alpha_k ,$$

which concludes the case $j = m+1$, since $\alpha_{m+1} = 0$.

Now, for $j \in \{2, \ldots, m\}$, we must show the lower bound for

$$\mathbb{P}(T \leq (j-1)N) + \sum_{\mathrm{i}\in[n]} \mathbb{P}(V_\mathrm{i} > t) \sum_{k=(j-1)N+1}^{Nm} \frac{\mathbb{P}(T \geq k | \sigma_k = \mathrm{i})}{Nm} .$$

For the first term we use again Lemma 5.7 and deduce that

$$\begin{aligned}\mathbb{P}(T \leq (j-1)N) &\geq \sum_{k\in[(j-1)N]} \frac{1-\alpha(k/Nm)}{Nm} \\ &= \sum_{k\in[j-1]} \frac{1-\alpha_k}{m} \\ &\geq \sum_{k\in[j-1]} \frac{1-\alpha_k}{m(1-\alpha_j)} \mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_\mathrm{i}\} > t) .\end{aligned}$$

Noticing that α is nonincreasing, the corresponding thresholds are nonincreasing and we can use both Lemma 5.8 and Lemma 6.3 in the following way. First, for every $\mathrm{i}, k \in [Nm]$,

$$\mathbb{P}(T \geq k | \sigma_k = \mathrm{i}) \geq \frac{\mathbb{P}(T > k)}{1 - \frac{1}{Nm} \sum_{l\in[k]} \mathbb{P}(V_\mathrm{i} > \tau_l)} .$$

Then, interchanging the order of the sums,

$$\sum_{\mathrm{i}\in[n]} \mathbb{P}(V_\mathrm{i} > t) \sum_{k=(j-1)N+1}^{Nm} \frac{\mathbb{P}(T \geq k | \sigma_k = \mathrm{i})}{Nm} \geq \sum_{k=(j-1)N+1}^{Nm} \mathbb{P}(T > k) \sum_{\mathrm{i}\in[n]} \frac{\mathbb{P}(V_\mathrm{i} > t)}{Nm - \sum_{l\in[k]} \mathbb{P}(V_\mathrm{i} > \tau_l)} .$$

Now, by Lemma 6.3, for $k \leq Nm/2$,

$$\sum_{i\in[n]} \frac{\mathbb{P}(V_i > t)}{Nm - \sum_{l\in[k]} \mathbb{P}(V_i > \tau_l)} \geq \frac{\mathbb{P}(\max_{i\in[n]}\{V_i\} > t)}{Nm - k\mathbb{P}(\max_{i\in[n]}\{V_i\} > \tau_1)}$$

$$= \frac{g_{Nm,\alpha_1}(k)}{Nm}\mathbb{P}(\max_{i\in[n]}\{V_i\} > t)\,,$$

and for $k > Nm/2$,

$$\sum_{i\in[n]} \frac{\mathbb{P}(V_i > t)}{Nm - \sum_{l\in[k]} \mathbb{P}(V_i > \tau_l)} \geq \frac{\mathbb{P}(\max_{i\in[n]}\{V_i\} > t)}{Nm - Nm/2\mathbb{P}(\max_{i\in[n]}\{V_i\} > \tau_1)}$$

$$= \frac{g_{Nm,\alpha_1}(k)}{Nm}\mathbb{P}(\max_{i\in[n]}\{V_i\} > t)\,,$$

which is an improvement over using the union bound (as in the previous section). All in all, we have proven the following bound.

$$\sum_{i\in[n]} \mathbb{P}(V_i > t) \sum_{k>(j-1)N}^{Nm} \frac{\mathbb{P}(T \geq k|\sigma_k = i)}{Nm} \geq \left[\sum_{k=(j-1)N+1}^{Nm} \mathbb{P}(T > k)\frac{g_{Nm,\alpha_1}(k)}{Nm}\right] \mathbb{P}(\max_{i\in[n]}\{V_i\} > t)\,.$$

Moreover,

$$\begin{aligned}
\sum_{k=(j-1)N+1}^{Nm} \mathbb{P}(T > k)\frac{g_{Nm,\alpha_1}(k)}{Nm} &= \sum_{l=j}^{m}\sum_{k=1}^{N} \mathbb{P}(T > (l-1)N + k)\frac{g_{Nm,\alpha_1}((l-1)N + k)}{Nm} \\
&\geq \sum_{l=j}^{m}\left(\prod_{l'=1}^{l-1}\alpha_{l'}\right)^{\frac{1}{m}} \sum_{k=1}^{N}\left(\alpha_l^{\frac{1}{Nm}}\right)^k \frac{g_{Nm,\alpha_1}((l-1)N)}{Nm} \\
&= \sum_{l=j}^{m}\left(\prod_{l'=1}^{l-1}\alpha_{l'}\right)^{\frac{1}{m}} g_{m,\alpha_1}(l-1)\frac{\alpha_l^{\frac{1}{Nm}}}{Nm}\frac{1-\alpha_l^{\frac{1}{m}}}{1-\alpha_l^{\frac{1}{Nm}}} \\
&\xrightarrow[N\to\infty]{} \sum_{k=j}^{m}\left(\prod_{l\in[k-1]}\alpha_l\right)^{\frac{1}{m}} g_{m,\alpha_1}(k-1)\left(\frac{1-\alpha_k^{\frac{1}{m}}}{-\ln\alpha_k}\right)\,.
\end{aligned}$$

Putting these two inequalities together, we can conclude the case $j \in \{2, \ldots, m\}$.

Lastly, for $j = 1$ (i.e. : $t \in [\tau_1, \infty)$) we do as before, in Section 5, and use Lemma 5.8, Lemma 5.7 and the union bound to derive the following

$$\begin{aligned}
\mathbb{P}(V_{\sigma_T} > t) &= \frac{1}{n}\sum_{i\in[n]} \mathbb{P}(V_i > t)\sum_{k=1}^{Nm} \mathbb{P}(T \geq k|\sigma_k = i) \\
&\geq \left[\frac{1}{n}\sum_{k=1}^{Nm} \mathbb{P}(T > k)\right]\mathbb{P}(\max_{i\in[n]}\{V_i\} > t) \\
&\geq \left[\sum_{k=1}^{m}\left(\prod_{l\in[k-1]}\alpha_l\right)^{\frac{1}{m}}\left(\frac{1-\alpha_k^{\frac{1}{m}}}{-\ln\alpha_k}\right)\right]\mathbb{P}(\max_{i\in[n]}\{V_i\} > t)\,,
\end{aligned}$$

where the last inequality is only valid in the limit as $N \to \infty$. $\square$

With the previous lemma we can easily establish the improved guarantee. The idea is quite simple, we take the right-hand-side of the expression in Lemma 6.1 and optimize over the choice of $\alpha_1, \dots, \alpha_m$. We do this optimization numerically and find a particular collection $\alpha_1, \dots, \alpha_m$ such that the guarantee evaluates to 0.669, as stated in the following corollary. We must note however that there might be other choices leading to slightly improved guarantees.

Corolario 6.2 *There exists* $1 \geq \alpha_1 \geq \ldots \geq \alpha_m \geq 0$ *such that*

$$\frac{\mathbb{E}(V_{\sigma_T})}{\mathbb{E}(\max_{i\in[n]}\{V_i\})} \geq 0.66975\,,$$

where T *is the stopping time corresponding to the blind strategy* $\alpha = \alpha_{\alpha_1,\dots,\alpha_m}$.

In particular, taking $m = 30$ was enough to derive this result. The missing part on the previous analysis is Lemma 6.3, which we now prove for completeness.

Lema 6.3 *Given* $V_1, ..., V_n$ *independent random variables and* $\tau_1 \geq \max\{\tau_2, \tau_3, ..., \tau_n\}$ *a sequence of thresholds. Then, for any* $t < \tau_1$ *and* $k \leq \frac{n}{2}$,

$$\sum_{i\in[n]} \frac{\mathbb{P}(V_i > t)}{1 - \frac{1}{n}\sum_{l\in[k]}\mathbb{P}(V_i > \tau_l)} \geq \frac{\mathbb{P}(\max_{i\in[n]}\{V_i\} > t)}{1 - \frac{k}{n}\mathbb{P}(\max_{i\in[n]}\{V_i\} > \tau_1)}\,.$$

Proof. Define $\lambda := \frac{k}{n} \in [0, \frac{1}{2}]$ and notice that

$$\begin{aligned}\sum_{i\in[n]} \frac{\mathbb{P}(V_i > t)}{1 - \frac{1}{n}\sum_{l\in[k]}\mathbb{P}(V_i > \tau_l)} &= \sum_{i\in[n]} \frac{\mathbb{P}(V_i > t)}{\frac{n-k}{n} + \frac{1}{n}\sum_{l\in[k]}\mathbb{P}(V_i \leq \tau_l)}\\ &\geq \sum_{i\in[n]} \frac{\mathbb{P}(V_i > t)}{1 - \lambda + \lambda\mathbb{P}(V_i \leq \tau_1)}\\ &= \sum_{i\in[n]} \frac{1 - F_i(t)}{1 - \lambda + \lambda F_i(\tau_1)}\\ &=: C(t; \lambda, F_1, \dots, F_n)\,.\end{aligned}$$

Therefore, it is sufficient to prove that

$$\begin{aligned}\frac{1 - F_1(t)}{1 - \lambda + \lambda F_1(\tau_1)} &+ \frac{1 - F_2(t)}{1 - \lambda + \lambda F_2(\tau_1)}\\ &\geq \frac{1 - F_1(t)F_2(t)}{1 - \lambda + \lambda F_1(\tau_1)F_2(\tau_1)}\,,\end{aligned} \tag{6.1}$$

since we can iterate this argument n times to deduce the result. To prove this inequality, define the following variables

$$\beta := F_1(\tau_1)F_2(\tau_1), \qquad \gamma := F_1(t)F_2(t),$$
$$x := F_1(\tau_1), \qquad y := F_1(t)\,.$$

We can now simply solve the following optimization problem:

$$
(P)\begin{cases} \min_{x,y} & f_\lambda(x,y) := \frac{1-y}{1-\lambda+\lambda x} + \frac{1-\frac{\gamma}{y}}{1-\lambda+\lambda\frac{\beta}{x}} \\ s.t. & \beta \le x \le 1 \\ & \beta \le \frac{\beta}{x} \le 1 \\ & \gamma \le y \le x \\ & \gamma \le \frac{\gamma}{y} \le \frac{\beta}{x}. \end{cases}
$$

Notice that the function f_λ defined by

$$
f_\lambda(x,y) = \frac{1}{1-\lambda+\lambda x} - \left(\frac{1}{1-\lambda+\lambda x}\right) y + \frac{1}{1-\lambda+\lambda\frac{\beta}{x}} - \left(\frac{\gamma}{1-\lambda+\lambda\frac{\beta}{x}}\right)\frac{1}{y},
$$

is concave in y. Rearranging the inequalities, the problem reduces to

$$
(P)\begin{cases} \min_{x} & f_\lambda(x) := \frac{1-x}{1-\lambda+\lambda x} + \frac{1-\frac{\gamma}{x}}{1-\lambda+\lambda\frac{\beta}{x}} \\ s.t. & \beta \le x \le 1 \end{cases}
$$

Then, all we have to show is that, if $\lambda \in [0, \frac{1}{2}]$, $x = 1$ is the minimum, since this would imply inequality (6.1). The following are three sufficient conditions for $x = 1$ to be the minimum:

1. $f_\lambda(\beta) \ge f_\lambda(1)$.

2. $x = 1$ is local minimum.

3. There exists at most one critical point in the interval $[\beta, 1]$.

All these conditions are true for $\lambda \in [0, 1/2]$, but involve tedious, but straight-forward, computations that are skipped. The first condition is a simple calculation. The second condition can be proved by checking that $\partial_x f_\lambda(1) \le 0$, which is true for $\lambda \le 1/2$. The third condition follows by noticing that the critical points of $f_\lambda(\cdot)$ are the solution to a polynomial equation of degree two, so there are at most two such points. Moreover, for $\lambda \in [0, 1/2]$, one of them must be negative, given it is a real number.

With these three conditions we can prove inequality (6.1) and therefore the lemma is proved iterating this inequality n times. □

Chapter 7

Upper bounds

7.1 A 0.675 upper bound for blind strategies

We first see an upper bound for the performance of blind strategies, which is very close to the lower bound shown in the previous section. To this end we consider two instances and show that no blind strategy can guarantee better than 0.675 for both instances.

The first instance consists simply in a single random variable which is nearly deterministic, given by $V_1 \sim U(1-\varepsilon, 1+\varepsilon)$. The second instance has n i.i.d. random variables defined by (and we take $n \to \infty$):

$$V_{\mathrm{i}} \sim \begin{cases} 1/\varepsilon & \text{w.p.} \quad \varepsilon \\ U(0,\varepsilon) & \text{w.p.} \quad 1-\varepsilon\,. \end{cases}$$

Combining these two instances one can show the following result.

Lema 7.1 *Let T be the stopping time corresponding to the blind strategy given by α. Then*

$$\sup_{\alpha} \inf_{n;F_1,\dots,F_n} \frac{\mathbb{E}(V_{\sigma_T})}{\mathbb{E}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\})} \leq \sup_{\alpha}\ \min\left\{1-\int_0^1 \alpha(s)\mathrm{d}s \quad , \quad \int_0^1 \mathrm{e}^{\int_0^s \ln\alpha(w)\mathrm{d}w}\mathrm{d}s\right\}.$$

With this result we need to compute the quantity on the right-hand-side of the previous lemma to obtain an upper bound on the performance guarantee of any blind strategy. This is done using optimal control theory. The basic procedure consists first in proving that the supremum right-hand-side in Lemma 7.1 is attained, then we have Mayer's optimal control problem for which the necessary optimality conditions can be expressed as an integro-differential equation. We conclude by solving this equation numerically, and thus have the following result.

Corolario 7.2 *Let T be the stopping time corresponding to the blind strategy given by α. Then*

$$\sup_{\alpha} \inf_{n;F_1,\dots,F_n} \frac{\mathbb{E}(V_{\sigma_T})}{\mathbb{E}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\})} \leq 0.675\,.$$

PROOF OF LEMMA 7.1. The first instance is $V_1 = U(1-\varepsilon, 1+\varepsilon)$, with $\varepsilon > 0$. Notice that, $\tau = 1 - \varepsilon + \alpha(u)2\varepsilon$, where $u \sim U(0,1)$ and, by direct computation, we have that

$$\begin{aligned}
\mathbb{E}(V_{\sigma_T}) &= \int_0^1 \mathbb{E}(V_{\sigma_T} | u = s)\mathrm{d}s \\
&= \int_0^1 \mathbb{E}(V_1 | \tau = 1 - \varepsilon + \alpha(s)2\varepsilon)\mathrm{d}s \\
&= \int_0^1 (1 + \varepsilon\alpha(s))\,(1 - \alpha(s))\,\mathrm{d}s \\
&\xrightarrow[\varepsilon\to 0]{} \int_0^1 (1 - \alpha(s))\mathrm{d}s\,.
\end{aligned}$$

The second instance has n i.i.d. random variables defined by:

$$V_{\mathrm{i}} \sim \begin{cases} \frac{1}{\varepsilon} & w.p. \quad \varepsilon \\ U(0,\varepsilon) & w.p. \quad 1-\varepsilon\,. \end{cases}$$

Moreover, for ε small enough,

$$\mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\} \le t) = \begin{cases} 0 & ; t < 0 \\ \left(\frac{1-\varepsilon}{\varepsilon}\right)^n t^n & ; 0 \le t < \varepsilon \\ (1-\varepsilon)^n & ; \varepsilon \le t < \frac{1}{\varepsilon} \\ 1 & ; \frac{1}{\varepsilon} \le t \end{cases}$$

Notice that we can assume $\alpha(x) < 1$, for $x > 0$, since there is no gain in rejecting anything at any point in time. Then, with probability $(1-\varepsilon)^{n^2}$, for all $\mathrm{i} \in [n]$, $u_{\mathrm{i}} < (1-\varepsilon)^n$ and therefore, $\tau_{\mathrm{i}} = \sqrt[n]{\alpha(u_{[\mathrm{i}]})}\frac{\varepsilon}{1-\varepsilon}$ and we have that $\mathbb{P}(V_{\mathrm{i}} \le \tau_{\mathrm{i}}) = \sqrt[n]{\alpha(u_{[\mathrm{i}]})}$. By direct computation,

$$\begin{aligned}
\lim_{\varepsilon\to 0}\mathbb{E}(V_{\sigma_T}|u) &= 1 + \sqrt[n]{\alpha(u_{[1]})} + \sqrt[n]{\alpha(u_{[1]})\alpha(u_{[2]})} + \ldots + \sqrt[n]{\alpha(u_{[1]})...\alpha(u_{[n-1]})} + o(\varepsilon) \\
&= \sum_{\mathrm{i}=0}^{n-1}\prod_{j=1}^{\mathrm{i}}\sqrt[n]{\alpha(u_{[j]})} + o(\varepsilon)\,.
\end{aligned}$$

In addition, we have $\mathbb{E}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\}) = \frac{1}{\varepsilon}(1-(1-\varepsilon)^n) \xrightarrow[\varepsilon\to 0]{} n$. Then,

$$\lim_{n\to\infty}\lim_{\varepsilon\to 0}\frac{\mathbb{E}(V_{\sigma_T})}{\mathbb{E}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\})} = \lim_{n\to\infty}\mathbb{E}_u\left[\frac{1}{n}\sum_{\mathrm{i}=0}^{n-1}\prod_{j=1}^{\mathrm{i}}\sqrt[n]{\alpha(u_{[j]})}\right] = \int_0^1 \mathrm{e}^{\int_0^s \ln\alpha(w)\mathrm{d}w}\mathrm{d}s\,.$$

□

SKETCH OF PROOF OF COROLLARY 7.2. We first show that the supremum given by Lemma 7.1 is attained at a certain α^*. To this end we note that, without loss of generality, we can consider the supremum over nonincreasing functions α, by a simple exchange of mass argument. Then, we note that the set of nonincreasing functions from $[0,1]$ to itself is compact for the $||\cdot||_\infty$ and the functional being optimized is continuous for that metric. Then we deduce the existence of α^*, and furthermore it satisfies

$$1-\int_0^1 \alpha^*(t)\mathrm{d}t = \int_0^1 \exp\left[\int_0^t \ln \alpha^*(w)\mathrm{d}w\right]\mathrm{d}t\,. \tag{7.1}$$

Therefore, α^* is the solution of the following optimal control problem:

$$(P)\left\{\begin{array}{ll} \min\limits_{\alpha} & -x_1(1) = -\int\limits_0^1 1-\alpha(t)\mathrm{d}t \\ s.t.: & \dot{x}(t) = \begin{pmatrix} 1-\alpha(t) \\ \exp[x_3(t)] \\ \ln\alpha(t) \end{pmatrix} \\ & x(0) = (0,0,0)' \\ & (t,x(t)) \in [0,1]\times\mathbb{R}^3 \\ & \alpha(t)\in[0,1] \\ & x_1(1) = x_2(1)\,. \end{array}\right.$$

Note that the objective is to maximize $x_1(1) = 1-\int_0^1 \alpha(s)\mathrm{d}s$ by choosing the right α. The dynamic is defined so that the restriction $x_1(1) = x_2(1)$ represents the condition (7.1) and the auxiliary component x_3 is needed to express this in standard form.

This problem was studied in section 3.3, leading to identify α^* with $\alpha_{K,\bar{t}}$ defined by

$$\alpha_{K,\bar{t}}(t) = \left\{\begin{array}{ll} 1 & ;0\le t<\bar{t} \\ \beta_{K,\bar{t}}(t) & ;\bar{t}\le t\le 1, \end{array}\right.$$

where $K>0$ and $\bar{t}\in[0,1)$ and $\beta_{K,\bar{t}}$ is the solution of the following integro-differential equation

$$\left\{\begin{array}{ll} \dot{\beta}(t) = -K\exp\left[\int\limits_{\bar{t}}^t \ln\beta(s)\mathrm{d}s\right] & t\in(\bar{t},1) \\ \beta(1) = 0\,. \end{array}\right. \tag{7.2}$$

We can further restrict the space of search for K and $\bar{t}$ using the fact that there is a blind strategy with performance 0.665. First notice that

$$0.665 \le \int_0^1 1-\alpha^*(t)\mathrm{d}t \le 1-\bar{t}\,,$$

therefore $\bar{t}\in(0,0.335)$. In the other hand, consider $K,\bar{t}$ such that $\alpha_{K,\bar{t}} = \alpha^*$, then

$$\begin{aligned}\beta_{K,\bar{t}}(\bar{t}) &= K\int_{\bar{t}}^1 \exp\left[\int\limits_{\bar{t}}^t \ln\beta(s)\mathrm{d}s\right]\mathrm{d}t \\ &= K\int_{\bar{t}}^1 \exp\left[\int\limits_0^t \ln\alpha^*(s)\mathrm{d}s\right]\mathrm{d}t\,.\end{aligned}$$

Using this equality, we can notice that

$$
\begin{aligned}
0.665 &\leq \int_0^1 \exp\left[\int_0^t \ln \alpha^*(w)\mathrm{d}w\right]\mathrm{d}t \\
&= \bar{t} + \int_{\bar{t}}^1 \exp\left[\int_{\bar{t}}^{t} \ln \alpha^*(s)\mathrm{d}s\right]\mathrm{d}t \\
&= \bar{t} + \frac{\beta_{K,\bar{t}}(\bar{t})}{K} \\
&\leq 0.335 + \frac{1}{K}\,,
\end{aligned}
$$

therefore, $K \leq 10/3$.

To solve numerically equation 7.2, consider the change of variables $g(t) = \int_{\bar{t}}^{t} \ln \beta(s)\mathrm{d}s$, so that the equation becomes the second order ODE

$$
\begin{cases}
\mathrm{e}^{\dot{g}(t)}\ddot{g}(t) = -K\mathrm{e}^{g(t)} \quad ; \quad t \in (\bar{t}, 1) \\
g(\bar{t}) = 0 \\
\dot{g}(\bar{t}) = \ln \beta(\bar{t})\,.
\end{cases}
$$

Because $\exp(\cdot)$ is continuous and locally Lipschitz, this is a well-posed Cauchy problem with a unique local solution. The initial condition $\dot{g}(\bar{t}) = \ln \beta(\bar{t})$ turns out to be simply a replacement for $\dot{g}(1) = -\infty$ in the sense that we search for the solutions g such that $g(\bar{t}) = 0$ and exploits at time 1. This seemingly numerical difficulty is well treated using solvers for stiff ODE such as *ode15s* of *Matlab*.

Then, we numerically compute (7.2) to determine that

$$
\begin{aligned}
&\sup_{\alpha} \min\left\{1 - \int_0^1 \alpha(s)\mathrm{d}s, \int_0^1 \mathrm{e}^{\int_0^s \ln \alpha(w)\mathrm{d}w}\mathrm{d}s\right\} \\
&\quad = \sup_{\substack{K\in[0,10/3] \\ \bar{t}\in[0,0.335)}} \min\left\{1 - \int_0^1 \alpha_{K,\bar{t}}(s)\mathrm{d}s, \int_0^1 \mathrm{e}^{\int_0^s \ln \alpha_{K,\bar{t}}(w)\mathrm{d}w}\mathrm{d}s\right\} \\
&\quad \leq 0.675\,.
\end{aligned}
$$

Finally we note that if (7.2) has no solution, this simply means that α^* does not corresponds to $\alpha_{K,\bar{t}}$ and thus it is not taken into account in the previous supremum. □

7.2 A 0.732 upper bound for nonadaptive strategies

Recall that a nonadaptive strategy 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. Our bound improves upon the lower bound of 0.745 which holds for the i.i.d. case [16].

Surprisingly, our bound comes from analyzing the following simple instance. Take $a \in [0,1]$ and consider n random variables whose values are distributed as:

$$\begin{aligned} &V_1 \sim \begin{cases} n & w.p. \quad \frac{1}{n} \\ 0 & w.p. \quad 1-\frac{1}{n} \end{cases} \\ &V_2 \equiv a \\ &V_3 = \dots V_n \equiv 0\,. \end{aligned}$$

Clearly any reasonable algorithm would always accept a value of n and never accept a value of 0. Therefore the only decision a nonadaptive algorithm has to make is that of whether accepting a value of a or not. As we are looking at nonadaptive algorithms this decision can depend only on the time the value a is sampled. Furthermore, the optimal decision follows the following simple rule: "pick the value a as long as the expectation of the future is less or equal to a". By this principle, taking only the time i in which the value a is discovered, this value should be picked if and only if

$$a \geq \mathbb{E}\left(\max\{V_{\sigma_{\mathrm{i}+1},\dots,V_{\sigma_n}}\} | V_{\sigma_1} = \dots = V_{\sigma_{\mathrm{i}-1}} = 0, V_{\sigma_\mathrm{i}} = a\right) = \phi(\mathrm{i})\,.$$

The function $\phi(\cdot)$ is strictly decreasing and $\phi(1) = 1$, $\phi(n) = 0$. Therefore, there is an index i^* from which the optimal algorithm starts to pick the value a if seen. Define $\lambda = \mathrm{i}^*/n$, then applying the optimal algorithm the gambler obtains the following

$$\begin{aligned} \mathbb{E}(V_{\sigma_T}) &= \frac{1}{n} \sum_{\mathrm{i} \in [n]} \mathbb{E}(V_{\sigma_T} | \sigma_\mathrm{i} = 2) \\ &= \frac{1}{n}\left[\mathrm{i}^* - 1 + \frac{1}{n}\sum_{\mathrm{i}=\mathrm{i}^*}^{n} \sum_{j \in [n]\setminus\{\mathrm{i}\}} \mathbb{E}(V_{\sigma_T} | \sigma_\mathrm{i} = 2, \sigma_j = 1)\right] \\ &= \frac{1}{n}\left[\mathrm{i}^* - 1 + \frac{1}{n-1}\sum_{\mathrm{i}=\mathrm{i}^*}^{n} (1+a)(\mathrm{i}-1) + a(n-\mathrm{i})\right] \\ &= \frac{1}{n}\left[\mathrm{i}^* - 1 + \frac{1}{n-1}\left(\frac{n(n+1)}{2} - \frac{\mathrm{i}^*(\mathrm{i}^*+1)}{2} + (n-\mathrm{i}^*)(an-1-a)\right)\right] \\ &= \lambda + \frac{1}{2} - \frac{\lambda^2}{2} + (1-\lambda)a + O\left(\frac{1}{n}\right) \\ &\leq 1 + \frac{a^*}{2} + O\left(\frac{1}{n}\right)\,, \end{aligned}$$

where the last inequality comes from optimizing over $\lambda \in [0,1]$.

On the other hand $\mathbb{E}(\max_{\mathrm{i}\in[n]}\{V_\mathrm{i}\}) = 1 + a - a/n$ which leads to conclude that, for nonadaptive algorithms

$$\frac{\mathbb{E}(V_{\sigma_T})}{\mathbb{E}(\max\limits_{\mathrm{i}\in[n]}\{V_\mathrm{i}\})} \leq \frac{1 + \frac{a^2}{2} + O\left(\frac{1}{n}\right)}{1 + a + O\left(\frac{1}{n}\right)}\,.$$

By minimizing the right-hand-side of the previous expression for $a \in [0,1]$, we obtain that the optimal value is attained at $a = \sqrt{3} - 1$, and therefore

$$\frac{\mathbb{E}(V_{\sigma_T})}{\mathbb{E}(\max\limits_{\mathrm{i}\in[n]}\{V_\mathrm{i}\})} \leq \frac{6 - 2\sqrt{3} + O\left(\frac{1}{n}\right)}{2\sqrt{3} + O\left(\frac{1}{n}\right)} \xrightarrow[n\to\infty]{} \sqrt{3} - 1 \approx 0.732\,.$$

Chapter 8

Dealing with discontinuous distributions

In this section we explain how to use a *blind strategy* in instances where the distributions $F_1, \ldots, F_n$ are not necessarily continuous. Recall that in the definition of blind strategies in Section 1.3, we need the existence of τ_{i} such that $\mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\} \leq \tau_{\mathrm{i}}) = \alpha(u_{[\mathrm{i}]})$. So, what happens if such thresholds $\tau_1, ..., \tau_n$ do not exist? For the purpose of studying the prophet inequality, the performance of a strategy defined over instances with continuous distributions is always extendable to discontinuous ones allowing stochastic tie breaking. In this case, we can explicitly define the strategy that α induces over discontinuous instances. The resulting strategy no longer depends on the distribution of the maximum only.

The procedure to compute the tie breaking is quite natural:

1. Approximate the instance.

2. Study the strategy induced by α in the approximated instance.

3. Replicate what would happen in the original instance, allowing tie breaking.

Given a realization of uniform random variables $u_1, \ldots, u_n$, assume that τ_{i} does not exist, in other words, for some $\mathrm{i} \in [n]$, there is a $\tau \in \mathbb{R}$ such that

$$\lim_{\varepsilon\to 0} \mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\} \leq \tau - \varepsilon) < \alpha(u_{[\mathrm{i}]}) < \mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\} \leq \tau)\,.$$

The stochastic tie breaking consists in accepting the value τ with some probability, say p_{i}. This acceptance rate depends on the whole instance, not only on the distribution of the maximum, and on the identity of the revealed variable. To compute these acceptance rates we use the following procedure. For $\varepsilon > 0$, consider the following approximated instance

$$F_{\mathrm{i}}^{\varepsilon}(t) = \begin{cases} F_{\mathrm{i}}(t) \\ F_{\mathrm{i}}(\tau-\varepsilon) + \frac{t-\tau+\varepsilon}{\varepsilon}\left(F(\tau) - F_{\mathrm{i}}(\tau-\varepsilon)\right) \end{cases},$$

for $t \notin [\tau-\varepsilon, \tau]$ in the first case and $t \in [\tau-\varepsilon, \tau]$ in the second case. This instance has a continuous distribution of the maximum in $[\tau-\varepsilon, \tau]$ and we are able to find τ^{ε}, the corresponding threshold for the approximated instance, such that

$$\mathbb{P}(\max_{\mathrm{i}\in[n]}\{V_{\mathrm{i}}\} \leq \tau^{\varepsilon}) = \alpha(u_{[\mathrm{i}]})\,.$$

Then, we compute, for $j \in [n]$, $\beta_j := \lim_{\varepsilon \to 0} F_j^\varepsilon(\tau^\varepsilon)$. To finish, we define, for j such that $\mathbb{P}(V_j = \tau) > 0$,

$$p_{\mathrm{i}}(j; F_1, ..., F_n) := \frac{F_j(\tau) - \beta_j}{\mathbb{P}(V_j = \tau)}$$

and $p_{\mathrm{i}} = 0$ otherwise. In other words, $p_{\mathrm{i}}(j; F_1, ..., F_n)$ corresponds to the portion of the probability that variable j is equal to τ that should be considered as a value of τ^+ and so be accepted. This will induce that, faced with V_j at time i, the gambler accepts its realization with probability $1 - \beta_j$. To be more precise, we use the following procedure.

Algorithm 2 Stochastic TTA

for $\mathrm{i} = 1, ..., n$ **do**
  if $V_{\sigma_{\mathrm{i}}} > \tau_{\mathrm{i}}$ **then**
    Take $V_{\sigma_{\mathrm{i}}}$
  else if $V_{\sigma_{\mathrm{i}}} = \tau_{\mathrm{i}}$ **then**
    Take $V_{\sigma_{\mathrm{i}}}$ with probability $p_{\mathrm{i}}(\sigma_{\mathrm{i}}; F_1, ..., F_n)$
  end if
end for

With this procedure, it is easy to see that all results extend to general instances.

Bibliography

[1] Melika Abolhassani, Soheil Ehsani, Hossein Esfandiari, MohammadTaghi HajiAghayi, Robert Kleinberg, and Brendan Lucier. Beating 1-1/e for ordered prophets. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, pages 61–71, New York, NY, USA, 2017. ACM.

[2] Yossi Azar, Ashish Chiplunkar, and Haim Kaplan. Prophet secretary: Surpassing the 1-1/e barrier. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC '18, pages 303–318, New York, NY, USA, 2018. ACM.

[3] Hedyeh Beyhaghi, Negin Golrezaei, Renato Paes Leme, Martin Pal, and Balasubramanian Sivan. Improved approximations for free-order prophets and second-price auctions. *CoRR*, abs/1807.03435, 2018.

[4] L. Cesari. *Optimization—Theory and Applications: Problems with Ordinary Differential Equations*. Stochastic Modelling and Applied Probability. Springer New York, 2012.

[5] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the Forty-second ACM Symposium on Theory of Computing*, STOC '10, pages 311–320, New York, NY, USA, 2010. ACM.

[6] José Correa, Patricio Foncea, Dana Pizarro, and Victor Verdugo. From pricing to prophets, and back! *Operations Research Letters*, 47(1):25 – 29, 2019.

[7] José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Posted price mechanisms for a random stream of customers. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, pages 169–186, New York, NY, USA, 2017. ACM.

[8] Jose Correa, Raimundo Saona, and Bruno Ziliotto. Prophet secretary through blind strategies. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1946–1961, 2019.

[9] Soheil Ehsani, MohammadTaghi Hajiaghayi, Thomas Kesselheim, and Sahil Singla. Prophet secretary for combinatorial auctions and matroids. *CoRR*, abs/1710.11213, 2017.

[10] Liran Einav, Chiara Farronato, Jonathan Levin, and Neel Sundaresan. Auctions versus posted prices in online markets. *Journal of Political Economy*, 126(1):178 – 215, 2018.

[11] Hossein Esfandiari, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Morteza Monemizadeh. Prophet secretary. *CoRR*, abs/1507.01155, 2015.

[12] Thomas S. Ferguson. Who solved the secretary problem? *Statist. Sci.*, 4(3):282–289, 08 1989.

[13] P. R. Freeman. The secretary problem and its extensions: A review. *International Statistical Review / Revue Internationale de Statistique*, 51(2):189–206, 1983.

[14] John P. Gilbert and Frederick Mosteller. Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, 61(313):35–73, 1966.

[15] 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, pages 58–65. AAAI Press, 2007.

[16] T. P. Hill and Robert P. Kertz. Comparisons of stop rule and supremum expectations of i.i.d. random variables. *The Annals of Probability*, 10(2):336–345, 1982.

[17] Robert P Kertz. Stop rule and supremum expectations of i.i.d. random variables: A complete comparison by conjugate duality. *J. Multivar. Anal.*, 19(1):88–112, June 1986.

[18] Robert Kleinberg and Seth Matthew Weinberg. Matroid prophet inequalities. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, STOC '12, pages 123–136, New York, NY, USA, 2012. ACM.

[19] Ulrich Krengel and Louis Sucheston. *Adv. Probab. Related Topics.*

[20] Ulrich Krengel and Louis Sucheston. Semiamarts and finite values. *Bull. Amer. Math. Soc.*, 83(4):745–747, 07 1977.

[21] Albert W. Marshall, Ingram Olkin, and Barry C. Arnold. *Inequalities: Theory of Majorization and its Applications*, volume 143. Springer, second edition, 2011.

[22] Roger B. Myerson. Optimal auction design. *Math. Oper. Res.*, 6(1):58–73, February 1981.

[23] J.E. Peajcariaac, Proschan Frank, and Y.L. Tong. *Convex Functions, Partial Orderings, and Statistical Applications.* Mathematics in Science and Engineering. Elsevier Science, 1992.

[24] Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *Ann. Probab.*, 12(4):1213–1216, 11 1984.