



# Trading Prophets

JOSÉ CORREA, Universidad de Chile, Chile

ANDRÉS CRISTI, Universidad de Chile, Chile

PAUL DÜTTING, Google Research, Switzerland

MOHAMMAD HAJIAGHAYI, University of Maryland, USA

JAN OLKOWSKI, University of Maryland, USA

KEVIN SCHEWIOR, University of Southern Denmark, Denmark

In this work we initiate the study of buy-and-sell prophet inequalities. We start by considering what is arguably the most fundamental setting. In this setting the online algorithm observes a sequence of prices one after the other. At each time step, the online algorithm can decide to buy and pay the current price if it does not hold the item already; or it can decide to sell and collect the current price as a reward if it holds the item.

We show that for i.i.d. prices a single-threshold online algorithm achieves at least  $1/2$  of the expected profit of the the optimal offline algorithm, and that this is the best possible ratio achievable by any online algorithm. For non-i.i.d. prices in random order, where prices are no longer independent, we give a single-threshold online algorithm that achieves at least a  $1/16$  fraction of the expected profit of the optimal offline algorithm. We also show that for this setting no online algorithm can yield a better than  $1/3$  approximation, and thus establish a formal separation from the i.i.d. case. On the other hand, we present a threshold-based online algorithm for this setting that yields a  $1/2 - o(1)$  approximation. For non-i.i.d. prices (and arbitrarily correlated distributions) no approximation is possible.

We use the results for these base cases to solve a variety of more complex settings. For instance, we show a  $1/2 - o(1)$  approximation for settings where prices are affiliated and the online algorithm has only access to a single sample. We also extend our upper and lower bounds for the single item case to  $k$  items, and thus in particular show that unlike in the classic prophet inequality problem with  $k$  items it is impossible to achieve  $1 - o(1)$  approximations. We also consider a budgeted version where the online algorithm can buy fractions of an item and re-invest gains, and show that it's possible to achieve constant-factor approximations to the growth rate of the optimal offline algorithm. For a setting with  $k$  types of items and streams of prices for each type of item, we show that for the unit-capacity case it is possible to achieve a  $\Omega(1/k)$  approximation, and that this is best possible.

CCS Concepts: • **Theory of computation** → **Algorithmic game theory and mechanism design; Online algorithms**; • **Applied computing** → **Electronic commerce**.

Additional Key Words and Phrases: optimal stopping, prophet inequalities, online trading

## ACM Reference Format:

José Correa, Andrés Cristi, Paul Dütting, Mohammad Hajiaghayi, Jan Olkowski, and Kevin Schewior. 2023. Trading Prophets. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC '23)*, July 9–12, 2023, London, United Kingdom. ACM, New York, NY, USA, 21 pages. <https://doi.org/10.1145/3580507.3597813>



This work is licensed under a Creative Commons Attribution International 4.0 License.

EC '23, July 9–12, 2023, London, United Kingdom

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0104-7/23/07.

<https://doi.org/10.1145/3580507.3597813>

## 1 INTRODUCTION

Consider a trader, let's call her Alice, buying and selling a certain good or commodity (e.g., paintings, used cars, barrels of oil, etc.). Suppose Alice has some limit on how many units she can store. Consider the situation where Alice faces a sequence of potential buyers/sellers that arrive online over time and approach her with offers, and wants to maximize her profit. Of course, she can only sell a unit of the good if she has it, and she can only buy a unit if she has enough place to store it. How should she proceed? Or consider Bob, who wants to invest \$1000 in Bitcoin. Suppose Bob seeks to buy and sell coins (which is possible in essentially arbitrary fractions) to make as much money as he can. He is willing to reinvest any gains but he doesn't want to put in extra money. So Bob's constraint is his current budget, and not his capacity. Suppose that on any given day, there is a single market price at which Bob can buy and sell coins. He can buy as many coins as he can afford, and he can sell some or all of his current shares. How should he proceed?

Of course, there is a whole range of situations like this, and different situations will justify different modelling assumptions [e.g. Charnes et al., 1966, Myerson and Satterthwaite, 1983, Osborne, 1959]. For example, in Alice's case, one could model buyer/seller valuations as independent draws from distinct distributions that arrive in random order. Another plausible assumption could be an affiliated-valuations model, where buyers/sellers share a random base value but their individual preferences are expressed by independent increments.

In this work, we initiate the study of this type of problems from a prophet-inequality perspective. That is, we assume that the trader has some prior (distributional) knowledge about the sequence of prices, and compare the expected performance of a given online algorithm to the expected performance of the best offline algorithm. A fundamental challenge of our buy-and-sell prophet-inequality problem and important departure from the classic prophet inequality problem is the mixed-sign objective. We start by considering what is arguably the most basic variant of a buy-and-sell problem, where the online algorithm can hold up to one indivisible item, she is not budget-constrained, and prices are either i.i.d. random variables, non-i.i.d. presented in random order, or non-i.i.d. presented in adversarial order. We then show how to use the results for these base cases to solve a variety of more complex settings, including settings with limited information about the distribution of prices, settings with additional forms of correlation among prices (such as the aforementioned affiliated-valuations model), settings with more than one item, settings with multiple types of items, settings with divisible items, and settings with budgets. We also observe that in the arbitrarily correlated case and in the case where prices form a martingale, no approximation is possible.

### 1.1 The Basic Trading Prophet Problem

At the heart of our work is the following basic *trading prophet* problem. A "gambler" (online algorithm) observes a sequence of  $n$  prices. The prices are generated by a stochastic process that is known to the agent, but the agent only observes the realized prices one-by-one in an online fashion. In the simplest and most fundamental version of the problem that we consider, the agent trades an indivisible good and can hold at most one copy of the good. In this case whenever the agent does not have the item, she can either buy it at the current price  $p_j$  (at a reward of  $-p_j$ ) or she can skip this price. Similarly, when she holds the item, she can either sell it at the current price  $p_j$  (collecting a reward of  $p_j$ ) or she can hold on to the item.

The agent can base her buy/sell decisions on knowledge about the stochastic process, the current state (e.g., whether she currently holds the item or not), and the history of prices. We seek to compare the expected reward that the agent can achieve this way to the expected reward of an all-knowing “prophet” (offline algorithm), who knows the entire sequence of prices in advance and can make optimal buy and sell decisions.

Our goal are worst-case approximation guarantees (“prophet inequalities”) that state that in the worst-case over all stochastic processes from a certain class, the expected reward of the online algorithm is at least an  $\alpha \in (0, 1]$  fraction of the expected reward of the optimal offline algorithm.

## 1.2 Results for the Basic Trading Prophet Problem

We show a number of tight and near-tight bounds for the fundamental version of the trading-prophet problem in which the agent can hold up to one item. Our results<sup>1</sup> for the base case are summarized in Table 1. We discuss a number of extensions, including additional results for correlated prices (beyond those discussed below) and budgets in Section 1.3.

*Results for I.I.D. Prices.* We start with the case where the prices  $p_j$  are i.i.d. draws from  $F$ . We show that the worst-case optimal online policy is a simple threshold policy and achieves a  $1/2$ -approximation. Note that we are assuming continuous distributions without loss of generality.

**Main Result 1.** If the  $p_i$  are i.i.d. draws from  $F$ , then buying whenever the price is below the median and selling otherwise yields a  $1/2$ -approximation; and no online policy can do better.

We formally prove this result in Section 3, but we give a high-level overview here: We argue that the optimal offline strategy buys whenever the price is a local minimum and sells at local maxima, and use this to show that the expected reward of the prophet is exactly  $(n - 1)/2$  times the expected absolute difference  $\mathbb{E}[|X_i - X_j|]$  between any two draws  $X_i, X_j, i \neq j$  from  $F$ , say between  $X_1$  and  $X_2$ . A simple application of the triangle inequality in combination with the fact that we are looking at i.i.d. random variables then leads to an upper bound of  $(n - 1) \cdot \mathbb{E}[|X_1 - T|]$  for any  $T$ . On the other hand, for  $T$  set to the median of  $F$ , we show the agent has the item at any given step with probability  $1/2$ . This means that for any intermediate step  $1 < i < n$ , with half probability whenever  $X_i < T$  the agent can buy and with half probability whenever  $X_i \geq T$  the agent can sell. So the expected reward from an intermediate step is  $1/2 \cdot (\mathbb{E}[X_1 \cdot \mathbb{1}_{X_1 \geq T}] - \mathbb{E}[X_1 \cdot \mathbb{1}_{X_1 < T}])$ . Now since  $T$  is the median,  $\mathbb{E}[T \cdot \mathbb{1}_{X_1 \geq T}] = \mathbb{E}[T \cdot \mathbb{1}_{X_1 < T}]$ . So we can add and subtract this term from the previous formula, and obtain that the expected reward from an intermediate step is  $1/2 \cdot \mathbb{E}[|X_1 - T|]$ . Together with a careful analysis of the boundary cases this shows that the expected reward achieved by the agent over all steps is exactly equal to  $(n - 1)/2 \cdot \mathbb{E}[|X_1 - T|]$ .

To establish that no online policy can achieve a better approximation guarantee than  $\frac{1}{2}$ , consider constructing an i.i.d. instance with  $n = 2$  random variables  $X_i \sim F$  for  $i \in \{1, 2\}$ , where  $F$  is supported on three values:  $0, \frac{1}{2}$ , and  $1$ . By appropriately selecting the probability distribution for these values, it is possible to ensure that the best online policy only buys in period 1 when it observes the lowest possible value. Namely, our distribution is such that the expected period 2 price is not larger than the middle value. In contrast, the prophet will occasionally buy when it observes the middle value in period 1 and when the period-2 price

<sup>1</sup>We would like to note that due to space limitations some proofs are omitted or provided in a shorter version. We refer the reader to a full version of the paper for more details.

	i.i.d. prices	non-i.i.d. prices	
		random order	worst order
upper bound	2 (Theorem 1)	16 (2 as $n \rightarrow \infty$ ) (Theorems 2 and 3)	–
lower bound	2 (Proposition 1)	3 (Proposition 2)	$\infty$ (simple example)

Table 1. The achievable approximation ratios in the basic versions of the trading-prophet problem.

is the highest possible value. This idea can also be generalized to  $n \geq 2$ , demonstrating that increasing  $n$  does not lead to better online policies.

*Results for Non-I.I.D. Prices, Random Order.* Now assume that prices  $p_j$  are first drawn independently from non-identical distributions  $F_j$ , and that they are then presented to the algorithms in random order. Note that when prices are generated this way, then prices observed on different days are no longer independent of each other. We show that despite this, it’s still possible to achieve a  $1/16$  approximation guarantee with a threshold policy, which sets a single non-adaptive threshold. We also show that this case is strictly harder than the i.i.d. case by showing that no online algorithm (whether single threshold or not) can achieve a better than  $1/3$  approximation. Interestingly, this impossibility applies even if the order of random variables, not their realizations, is revealed to the algorithm before the sequence starts. We also present a threshold-based online algorithm that achieves a  $1/2 - \frac{c}{n}$ -approximation, for some constant  $c$ .

**Main Result 2.** For non-i.i.d. prices  $p_j \sim F_j$  presented in random order there exists a threshold  $T$  (different from the median of the mixture distribution) such that the corresponding threshold policy that buys below  $T$  and sells above  $T$  achieves a  $1/16$ -approximation, and no online policy can achieve a better than  $1/3$ -approximation.

**Main Result 3.** For non-i.i.d. prices  $p_j \sim F_j$  presented in random order, setting the threshold  $T$  to the median of the mixture distribution yields a  $(1/2 - \frac{c}{n})$ -approximation, for some constant  $c$ .

We give formal proofs in Section 4, but we again give some overview here. The core of our argument is a reduction to a two-period problem with two (random) random variables  $X_{\sigma(1)}$  and  $X_{\sigma(2)}$ , where  $\sigma$  is a uniform random permutation. Generalizing our approach for the i.i.d. case we show that the expected reward of the prophet is  $(n - 1)/2 \cdot \mathbb{E}[|X_{\sigma(1)} - X_{\sigma(2)}|]$ , while the expected reward of a threshold policy with threshold  $T$  is equal to

$$(n - 1)/2 \cdot \mathbb{E}[|X_{\sigma(1)} - X_{\sigma(2)}| \cdot \mathbb{1}_{T \in [\min(X_{\sigma(1)}, X_{\sigma(2)}), \max(X_{\sigma(1)}, X_{\sigma(2)})]}].$$

The main difficulty for relating the two quantities is that, unlike in the i.i.d. case where the comparison was via  $X_1, X_2$  and these were independent random variables, now we are looking at  $X_{\sigma(1)}, X_{\sigma(2)}$ , which are no longer independent. For the universal  $1/16$ -approximation we address this correlation by arguing that we can approximate the two quantities of interest up to constant factors with two independent random variables  $X_{a'}, X_{b'}$ . For this we consider all possible ways of splitting the random variables into two equal halves and drawing  $X_{a'}, X_{b'}$  from the two halves, and argue that there must be a split that leads to the required property.

For the asymptotic approximation guarantee of  $(1/2 - o(1))$  we formalize the intuitive claim that, the larger  $n$ , the less correlated the two random variables  $X_{\sigma(1)}, X_{\sigma(2)}$  should be.

To show that no online policy can obtain a better than  $1/3$ -approximation, we consider the following instance with  $n = 2$  random variables. Let  $M$  be a large constant. The first random variable is equal to  $M+2$  with probability  $M/(M+2)$  and zero otherwise. The second random variable is  $M$  with probability  $M/(M+2)$  and  $2M+2$  otherwise. With probability  $1/2$  we first see  $X_1$  and then  $X_2$  and with probability  $1/2$  we first see  $X_2$  and then  $X_1$ . The best online policy (found by backward induction) can be shown to buy only when  $X_{\sigma(1)} = 0$  which leads to an expected profit of  $1$ . The prophet in turn buys and sells except when  $X_{\sigma(1)} = 2M+2$  or  $X_{\sigma(2)} = 0$ , which leads to an expected value of  $3 - O(1/M)$ .

*Results for Non-I.I.D. Prices, Adversarial Order.* If prices are distributed non-identically and presented in arbitrary order, then no constant approximation factor is possible. Somewhat reminiscent of the classic lower-bound example for the standard prophet inequality, for some  $\varepsilon > 0$ , let  $n = 2$  and

$$X_1 = 1; X_2 = \begin{cases} \frac{1}{\varepsilon} & \text{w.p. } \varepsilon, \\ 0 & \text{w.p. } 1 - \varepsilon. \end{cases}$$

An optimal algorithm is indifferent between buying  $X_1$  and not doing so, therefore obtaining an expected value of  $0$ . The prophet, however, buys  $X_1$  precisely when  $X_2 = 1/\varepsilon$  (which happens with probability  $\varepsilon$ ), therefore obtaining an expected value of  $\varepsilon \cdot (1/\varepsilon - 1)$ . Note how the exact same example can be interpreted as a joint distribution over pairs of prices, namely  $(1, \frac{1}{\varepsilon})$  with probability  $\varepsilon$  and  $(1, 0)$  with probability  $1 - \varepsilon$ , and so it also shows that with arbitrarily correlated prices no approximation is possible.

*(Non-)Connection to Bilateral Trade.* Our analysis—especially the random-order case—reveals an interesting (non-)connection to bilateral trade. Namely, the two-period problem that we reduce to can be seen as the problem of maximizing *gains from trade* when the item is allocated to one of the two parties at random.

For the classic variant of this problem where the item is allocated to a fixed side of the market, classic work by Myerson and Satterthwaite [1983] showed that the “second best” (i.e., the gains from trade achievable with a truthful mechanism) is generally strictly smaller than the “first best” (theoretical optimum). McAfee [2008] showed that a fixed price mechanism achieves a  $1/2$  approximation when the two sides are identically distributed.

For general non-identical distributions it remained open whether a constant-factor approximation is possible. Only very recently, Deng et al. [2022] were able to resolve this question in an affirmative way. They showed that—unlike in our case where no constant-factor approximation is possible in the general case—the classic problem admits a  $1/8.23$  approximation.

### 1.3 Extensions and Applications

Just like the classic single-item prophet inequality [Krengel and Sucheston, 1977, 1978, Samuel-Cahn, 1984], the basic trading-prophet problem that we introduce in this paper can be used as a tool to solve a number of related problems, and it can be extended in different directions. This includes settings where prices rather than being independent are affiliated, or settings where the trader is subject to a budget constraint. We give an overview here, which is made precise in Section 5.

*Unknown Distribution and Affiliated Prices.* We show how to handle the version of the i.i.d. problem in which the distribution is not known. We give a simple trading strategy that achieves a  $(1/2 - o(1))$ -approximation. The idea is to use a random sample rather than the median as a threshold, and economize on the use of samples using a variation of the fresh-looking samples idea of Correa et al. [2019, 2021a]. The same algorithm works and yields a  $(1/2 - o(1))$ -approximation when prices exhibit a form of correlation known as affiliation in Economics, i.e., there is a distribution  $G$  and  $n$  distributions  $F_1, \dots, F_n$  and price  $p_j = x_j + y$  where  $x_j \sim F_j$  and  $y \sim G$ .

*More Than One Item.* We show that unlike in the single-item prophet inequality problem, where better approximation guarantees can be obtained when there are  $k$  copies of the good [Alaei, 2014], both the optimal offline algorithm and the optimal online algorithm for the trading-prophet problem satisfy an “all or nothing” property. So they either buy all  $k$  copies that they are allowed to buy, or they sell all of the copies that they currently hold. So a simple scaling argument implies that the upper and lower bounds that we identified for the single-item case also apply in the multi-item case.

*Budgeted Version with Fractional Purchase and Re-investment of Gains.* We show how our result can be applied in a budgeted-setting similar to that faced by a budget-constrained investor in the stock market. We assume that the agent starts with a budget of  $B_1 = 1$  and no goods  $S_1 = 0$ . In each step  $j$  the agent can either buy a fractional amount of the good (at most  $B_j/p_j$  units, where  $B_j$  is the budget in period  $j$ ) or sell any fraction of the amount  $S_j$  of the good that she currently possesses. We show through reduction that it’s possible to achieve constant-factor approximations to the growth rate of the optimal offline policy.

*Multi-Armed Bandit Version.* We present results for a variant of the basic trading-prophet problem in which there are  $k$  different types of good. We model this by running  $k$  parallel copies of our base model, so there are  $k$  streams of prices, one for each type of good. Both the agent and the prophet can buy any of the items at the current price if they currently don’t hold an item, or they can sell the item that they hold at the current price of that item. We use our results for the basic trading-prophet problem to show that it is possible to achieve  $O(k)$  approximations, and present an asymptotically tight  $\Omega(k)$  lower bound.

*Random Walks and Martingales.* A natural extension/variant of our model concerns situations where the prices form a balanced random walk, i.e.,  $p_j = p + \sum_{i=1}^j x_i$  where  $p$  is some base price and  $x_1, \dots, x_n$  are i.i.d. variables, each either  $-1$  or  $1$  with equal probability. Unfortunately, the resulting sequence of prices is a martingale, and it follows from the well-known *optional stopping theorem* that the expected value of any online algorithm is non-positive. Hence, no approximation ratio can be achieved by an online algorithm. Identifying a class of supermartingales that admits a meaningful approximation ratio is an interesting direction for future research.

#### 1.4 Further Related Work

To the best of our knowledge, ours is the first work to address the repeated buying and selling problem from a prophet inequality lens. In particular, the aspect that we allow the agent to buy and sell multiple times seems new.

On a technical level our work is related to the vast literature on prophet inequalities. This includes work on the classic single-item prophet inequality [Krengel and Sucheston, 1977, 1978, Samuel-Cahn, 1984], the i.i.d. case [Abolhassani et al., 2017, Correa et al., 2021b, Kleinberg and Kleinberg, 2018, Liu et al., 2021, Singla, 2018], and the prophet secretary

problem [Correa et al., 2021c, Ehsani et al., 2018, Esfandiari et al., 2017]. Our unknown distributions result is related to a recent stream of work that looks at prophet inequalities with samples [Caramanis et al., 2022, Correa et al., 2019, 2020, Rubinstein et al., 2020]. Ideas from the prophet inequalities literature have also been used to construct simple near-optimal mechanisms for two-sided markets [Braun and Kesselheim, 2021, Colini-Baldeschi et al., 2016, 2017a, Dütting et al., 2021].

Relevant pointers into the literature on two-sided markets include the aforementioned seminal work by Myerson and Satterthwaite [1983] and their celebrated impossibility result for the bilateral trade problem, as well as [Brustle et al., 2017, Colini-Baldeschi et al., 2017b, Deng et al., 2022, McAfee, 2008] who show constant-factor approximation guarantees for maximizing the “gains from trade”.

Related stopping problems have been studied in the finance/applied probability literature. First, du Toit and Peskir [2007], Graversen et al. [2006] study the problem of stopping a Brownian motion, with the benchmark of minimizing the expected squared difference between the value at which the online algorithm stops and the maximum value in the sequence. Second, du Toit and Peskir [2009], for the same stochastic process, study the problem of minimizing the expected ratio between the maximum value in the sequence and the value at which the online algorithm stops. In addition to studying different benchmarks, these works also do not quantify the worst-case performance guarantee.

There is also a rich literature on online portfolio selection problems, including work that takes a worst-case competitive analysis approach [Borodin and El-Yaniv, 1998, Chapter 14]. We refer the reader to the survey of Li and Hoi [2014] for details.

## 2 MODEL

In our basic model we trade a single unit of an item. We are given  $n$  distributions  $F_1, \dots, F_n$  and, in random order, in each of  $n$  periods we are offered a price for the item, drawn from one of the distributions. More precisely, we have independent prices  $X_i \sim F_i$  for  $1 \leq i \leq n$ , and an independent uniformly random permutation  $\sigma : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ ; and on period  $i$  we get to observe the price  $X_{\sigma(i)}$ . In each period we have two possible states: We either have the item, or we do not. In each period  $i$ , immediately after observing the price  $X_{\sigma(i)}$  we must make a decision. If we have the item, we can sell it at price  $X_{\sigma(i)}$ , or pass. If we do not have the item, we can buy it at price  $X_{\sigma(i)}$ , or pass. In period 1 we do not have the item, and the state of period  $i + 1$  is determined by the decision of period  $i$ .

Our objective is to design a decision rule that maximizes the expectation of the profit we make, where the profit is the sum of the prices on periods in which we sell minus the sum of prices on periods in which we buy. For an algorithm ALG, slightly abusing notation, we also denote its (random) profit by ALG and its expected profit by  $\mathbb{E}(\text{ALG})$ .

We denote by OPT the optimal profit in hindsight. This is, OPT is the (random) maximum profit given the realization of the prices and the permutation  $\sigma$ , over all possible sequences of buy/sell decisions. We compare ourselves with  $\mathbb{E}(\text{OPT})$ , where the expectation is over the realizations of the prices and the permutation  $\sigma$ . We say a decision rule or algorithm ALG is an  $\alpha$ -approximation for a family of instances if

$$\mathbb{E}(\text{ALG}) \geq \alpha \cdot \mathbb{E}(\text{OPT})$$

for all instances in that family.

The following simple observation about OPT, which can be proved by doing local changes in the decisions of OPT, will be important in our analysis. It states that OPT buys in local minima and sells in local optima. See Figure 1 for an illustration.




Fig. 1. Example of decisions made by OPT, the optimal buy/sell decisions in hindsight.

**OBSERVATION 1.** To handle corner cases in a unified manner, define  $X_{\sigma(0)} := \infty$  and  $X_{\sigma(n+1)} := -\infty$ . Then OPT buys in period  $1 \leq i < n$  if and only if the price  $X_{\sigma(i)}$  is a local minimum, i.e.,  $X_{\sigma(i)} < X_{\sigma(i-1)}$  and  $X_{\sigma(i)} \leq X_{\sigma(i+1)}$ . Analogously, OPT sells in period  $1 < i \leq n$  if and only if the price  $X_{\sigma(i)}$  is a local maximum, i.e., if  $X_{\sigma(i)} \geq X_{\sigma(i-1)}$  and  $X_{\sigma(i)} > X_{\sigma(i+1)}$ .<sup>2</sup>

In our proofs we assume all distributions are absolutely continuous. In particular, this implies that for every  $i$  there always exists  $T_i \in \mathbb{R}$  (a median) such that  $\mathbb{P}(X_i < T_i) = \mathbb{P}(X_i \geq T_i) = 1/2$ . This assumption is w.l.o.g., as we can add to each price an auxiliary independent perturbation  $\varepsilon_i \sim \text{Uniform}[-\varepsilon, \varepsilon]$ , for an  $\varepsilon > 0$  such that  $n \cdot \varepsilon \ll \mathbb{E}(\text{OPT})$ .

### 3 I.I.D. PRICES

In this section we consider the case where all distributions are equal, so the prices are i.i.d. For this setting we remove the dependency on the permutation  $\sigma$ , and simply denote by  $X_i$  the price in period  $i$ . We present a simple algorithm that achieves a  $1/2$ -approximation, and show this factor is tight.

Our algorithm is a single-threshold rule, i.e., there is a threshold  $T \in \mathbb{R}$  such that our decision in each period  $1 \leq i < n$  is to sell whenever  $X_i \geq T$ , and to buy whenever  $X_i < T$ . In period  $n$  the algorithm sells if we have the item and passes if we do not, regardless of the price. We denote such an algorithm as  $\text{ALG}_T$ .

**THEOREM 1.** If the prices are i.i.d. and  $T$  is the median of the distribution, then  $\text{ALG}_T$  is a  $1/2$ -approximation, i.e.,  $\mathbb{E}(\text{ALG}_T) \geq \frac{1}{2}\mathbb{E}(\text{OPT})$ .

To prove this theorem, we first obtain in Lemma 1 an upper bound on  $\mathbb{E}(\text{OPT})$  in terms of  $T$ , and then show in Lemma 2 that the expected profit of the algorithm is in fact half of this upper bound when  $T$  is the median of the distribution.

**LEMMA 1.** If the prices are i.i.d., then for any  $T \in \mathbb{R}$ ,

$$\mathbb{E}(\text{OPT}) = \frac{n-1}{2} \cdot \mathbb{E}(|X_1 - X_2|) \leq (n-1) \cdot \mathbb{E}(|X_1 - T|).$$

**PROOF.** Let us denote by  $\text{OPT}_i$  the gain of OPT in period  $i$ , i.e., if OPT buys in period  $i$ , then  $\text{OPT}_i = -X_i$ , if OPT sells in period  $i$ , then  $\text{OPT}_i = X_i$  and  $\text{OPT}_i = 0$  otherwise. From Observation 1, we have the following.

$$\text{OPT}_i = \begin{cases} -X_1 \cdot \mathbb{1}_{X_1 \leq X_2} & \text{if } i = 1 \\ X_i \cdot \mathbb{1}_{X_i \geq X_{i-1}} - X_i \cdot \mathbb{1}_{X_i \leq X_{i+1}} & \text{if } 2 \leq i \leq n-1 \\ X_n \cdot \mathbb{1}_{X_n \geq X_{n-1}} & \text{if } i = n. \end{cases}$$

<sup>2</sup>The strict inequalities break ties in the border case where there is a set of consecutive periods with equal prices. In such a set OPT buys or sells at most once, and furthermore, it can do the operation in any period and the result is the same. With this particular choice, OPT will only buy in the first period of such a set, and sell only in the last one.



Now we calculate  $\mathbb{E}(\text{OPT})$  by adding up all these terms and taking expectation.

$$\begin{aligned}\mathbb{E}(\text{OPT}) &= \sum_{i=1}^n \mathbb{E}(\text{OPT}_i) = \sum_{i=2}^n \mathbb{E}(X_i \cdot \mathbb{1}_{X_i \geq X_{i-1}}) - \sum_{i=1}^{n-1} \mathbb{E}(X_i \cdot \mathbb{1}_{X_i \leq X_{i+1}}) \\ &= (n-1) \cdot \mathbb{E}\left((X_1 - X_2) \cdot \mathbb{1}_{X_1 \geq X_2}\right).\end{aligned}$$

The last line follows from the fact that prices are i.i.d. Finally, also from the fact that prices are i.i.d. and from the triangle inequality, we conclude that for any  $T \in \mathbb{R}$ ,

$$\mathbb{E}(\text{OPT}) = \frac{n-1}{2} \cdot \mathbb{E}(|X_1 - X_2|) \leq (n-1) \cdot \mathbb{E}(|X_1 - T|). \quad \square$$

LEMMA 2. *If the prices are i.i.d. and  $T$  is the median of the distribution, then*

$$\mathbb{E}(\text{ALG}_T) = \frac{n-1}{2} \cdot \mathbb{E}(|X_1 - T|).$$

PROOF. As in Lemma 1, we analyze the gains of  $\text{ALG}_T$  in period  $i$ , which we denote by  $\text{ALG}_T(i)$ . If  $\text{ALG}_T$  buys in period  $i$ , then  $\text{ALG}_T(i) = -X_i$ , if it sells, then  $\text{ALG}_T(i) = X_i$ , and  $\text{ALG}_T(i) = 0$  otherwise. By the definition of  $\text{ALG}_T$ , we have that

$$\text{ALG}_T(1) = -X_1 \cdot \mathbb{1}_{X_1 < T}.$$

Denote by  $p_i$  the probability that we have the item in period  $i$  (before making a decision). For  $2 \leq i \leq n-1$ , since the price of period  $i$  is independent of the prices in periods  $1, \dots, i-1$ , we have the following.

$$\mathbb{E}(\text{ALG}_T(i)) = \mathbb{E}(X_i \cdot \mathbb{1}_{X_i \geq T}) \cdot p_i - \mathbb{E}(X_i \cdot \mathbb{1}_{X_i < T}) \cdot (1 - p_i).$$

In the last period, the algorithm only sells whenever we have the item. Thus,

$$\mathbb{E}(\text{ALG}_T(n)) = \mathbb{E}(X_n) \cdot p_n = \left( \mathbb{E}(X_n \cdot \mathbb{1}_{X_n \geq T}) + \mathbb{E}(X_n \cdot \mathbb{1}_{X_n < T}) \right) \cdot p_n.$$

We show now that if  $T$  is the median of the distribution,  $p_i = 1/2$  for all  $i \geq 2$ . In fact, we have the following formula, which is obtained by conditioning on the event that we have (or not) the item in period  $i-1$ .

$$\begin{aligned}p_i &= \mathbb{P}(X_{i-1} < T) \cdot p_{i-1} + \mathbb{P}(X_{i-1} \geq T) \cdot (1 - p_{i-1}) \\ &= \mathbb{P}(X_{i-1} < T) = 1/2.\end{aligned}$$

Putting all together, we have that

$$\begin{aligned}\mathbb{E}(\text{ALG}_T) &= \sum_{i=1}^n \mathbb{E}(\text{ALG}_T(i)) \\ &= -\mathbb{E}(X_1 \cdot \mathbb{1}_{X_1 < T}) + \frac{1}{2} \cdot \sum_{i=2}^{n-1} \mathbb{E}(X_i \cdot \mathbb{1}_{X_i \geq T}) - \frac{1}{2} \cdot \sum_{i=2}^{n-1} \mathbb{E}(X_i \cdot \mathbb{1}_{X_i < T}) \\ &\quad + \frac{1}{2} \cdot \left( \mathbb{E}(X_n \cdot \mathbb{1}_{X_n \geq T}) + \mathbb{E}(X_n \cdot \mathbb{1}_{X_n < T}) \right) \\ &= \frac{n-1}{2} \cdot \left( \mathbb{E}(X_1 \cdot \mathbb{1}_{X_1 \geq T}) - \mathbb{E}(X_1 \cdot \mathbb{1}_{X_1 < T}) \right).\end{aligned}$$

Here, the last line comes from the fact that the prices are i.i.d. Finally, notice that  $\mathbb{P}(X_1 \geq T) = \mathbb{P}(X_1 < T) = 1/2$ , so  $\mathbb{E}(T \cdot \mathbb{1}_{X_1 \geq T}) = \mathbb{E}(T \cdot \mathbb{1}_{X_1 < T})$ , and therefore, by adding and subtracting

this quantity in the last line, we conclude that

$$\mathbb{E}(\text{ALG}_T) = \frac{n-1}{2} \cdot \mathbb{E}(|X_1 - T|). \quad \square$$

**PROPOSITION 1.** *For any  $n \geq 2$  and  $\varepsilon > 0$ , there is an instance with  $n$  i.i.d. prices such that for the optimal algorithm ALG,*

$$\mathbb{E}(\text{ALG}) \leq \left(\frac{1}{2} + \varepsilon\right) \cdot \mathbb{E}(\text{OPT}).$$

#### 4 NON-I.I.D. PRICES, RANDOM ORDER

In this section we consider the case where the prices  $X_1, \dots, X_n$  are independent draws from not necessarily identical distributions  $F_1, \dots, F_n$ ; and these prices are presented to us in random order. We show three results, a  $1/16$  approximation by a threshold policy, an impossibility showing that no online policy can achieve a better than  $1/3$  approximation, and an asymptotic  $1/2 - o(1)$  approximation by a threshold policy as  $n \rightarrow \infty$ . The asymptotic  $1/2 - o(1)$  approximation is obtained by setting the threshold to the median of the mixture distribution. The  $1/16$  approximation requires a different threshold. We describe how we choose the threshold for the  $1/16$  approximation in Section 4.2, and provide a hard instance for the median of the mixture distribution for small  $n$  in the full version of the paper.

**THEOREM 2.** *If the prices are presented in order  $X_{\sigma(1)}, \dots, X_{\sigma(n)}$ , where  $\sigma : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$  is a uniform random permutation, then there exist a threshold  $T$  such that*

$$\mathbb{E}(\text{ALG}_T) \geq \frac{1}{16} \cdot \mathbb{E}(\text{OPT}).$$

Moreover, the threshold  $T$  can be computed in polynomial time.

**PROPOSITION 2.** *For every  $\varepsilon > 0$ , there is an instance such that if the prices arrive in uniform random order, then for the optimal algorithm ALG,*

$$\mathbb{E}(\text{ALG}) \leq \left(\frac{1}{3} + \varepsilon\right) \cdot \mathbb{E}(\text{OPT}).$$

**THEOREM 3.** *If the prices are presented in order  $X_{\sigma(1)}, \dots, X_{\sigma(n)}$ , where  $\sigma : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$  is a uniform random permutation, then there exist a threshold  $T$  and a constant  $C$ , independent of these distributions, such that*

$$\left(1 + \frac{C}{n}\right) \cdot 2 \geq \frac{\mathbb{E}(\text{OPT})}{\mathbb{E}(\text{ALG}_T)}.$$

We provide the proofs of Theorem 2 and Theorem 3 in Section 4.1, Section 4.2, and Section 4.3; the proof of Proposition 2 in the full version.

##### 4.1 Reduction to Two Periods with Correlated Random Variables

We start by reducing the problem of showing an approximation guarantee for the random order model with  $n$  periods to a two period problem with correlated random variables.

**LEMMA 3.** *If the prices are presented in uniformly random order  $\sigma$ , then*

$$\mathbb{E}(\text{OPT}) = \frac{n-1}{2} \cdot \mathbb{E}(|X_{\sigma(1)} - X_{\sigma(2)}|).$$

PROOF. For a fixed order  $\sigma$  and realizations  $X_{\sigma(1)}, \dots, X_{\sigma(n)}$ , we can argue as in the beginning of the proof of Lemma 1 and use Observation 1 to obtain

$$\text{OPT} = \sum_{i=2}^n X_{\sigma(i)} \cdot \mathbb{1}_{X_{\sigma(i)} \geq X_{\sigma(i-1)}} - \sum_{i=1}^{n-1} X_{\sigma(i)} \cdot \mathbb{1}_{X_{\sigma(i)} \leq X_{\sigma(i+1)}} = \sum_{i=1}^{n-1} (X_{\sigma(i+1)} - X_{\sigma(i)}) \cdot \mathbb{1}_{X_{\sigma(i+1)} \geq X_{\sigma(i)}}.$$

Then, taking expectation over the possible orders  $\sigma$  and realizations  $X_{\sigma(1)}, \dots, X_{\sigma(n)}$ ,

$$\begin{aligned} \mathbb{E}(\text{OPT}) &= \sum_{i=1}^{n-1} \mathbb{E}((X_{\sigma(i+1)} - X_{\sigma(i)}) \cdot \mathbb{1}_{X_{\sigma(i+1)} \geq X_{\sigma(i)}}) \\ &= (n-1) \cdot \mathbb{E}((X_{\sigma(1)} - X_{\sigma(2)}) \cdot \mathbb{1}_{X_{\sigma(1)} \geq X_{\sigma(2)}}) = \frac{n-1}{2} \cdot \mathbb{E}(|X_{\sigma(1)} - X_{\sigma(2)}|), \end{aligned}$$

where the first equality holds by linearity of expectation and the second and third equality hold because  $\sigma$  is a uniformly random order.  $\square$

LEMMA 4. *If the prices are presented in uniformly random order  $\sigma$ , then for any threshold  $T \in \mathbb{R}$  we have that*

$$\mathbb{E}(\text{ALG}_T) = \frac{n-1}{2} \cdot \mathbb{E}(|X_{\sigma(1)} - X_{\sigma(2)}| \cdot \mathbb{1}_{T \in [\min(X_{\sigma(1)}, X_{\sigma(2)}), \max(X_{\sigma(1)}, X_{\sigma(2)})]}).$$

PROOF. Let  $\text{ALG}_T(i)$  denote the gain (or the loss) of  $\text{ALG}_T$  in period  $i$ . By the definition of  $\text{ALG}_T$ , it can only buy in period  $i = 1$ . Thus,  $\text{ALG}_T(1) = -X_{\sigma(1)} \cdot \mathbb{1}_{X_{\sigma(1)} < T}$ .

Consider now any period  $2 \leq i \leq n-1$ . In such period, the algorithm buys whenever  $X_{\sigma(i)} < T$  and it does not have the item. The latter event is equivalent to the event  $X_{\sigma(i-1)} \geq T$ . Indeed, assuming that  $X_{\sigma(i-1)} \geq T$  either  $\text{ALG}_T$  in period  $i-1$  does not have the item, but in this case it will not buy it in this period; or it does have the item, but then it will certainly sell the item in period  $i-1$ . In both cases, the algorithm does not have the item in period  $i$ . On the other hand, if  $X_{\sigma(i-1)} < T$  then either the algorithm already has the item in period  $i-1$  or it will buy it in that period, so in this case it will certainly hold the item in period  $i$ .

Selling an item is similar. In any period  $2 \leq i \leq n-1$  the algorithm sells the item iff  $X_{\sigma(i)} \geq T$  and it has the item. Analogously, the algorithm has the item in period  $i$  iff  $X_{\sigma(i-1)} < T$ . Therefore, we can express the gains of  $\text{ALG}_T$  in period  $2 \leq i \leq n-1$  by

$$\text{ALG}_T(i) = -X_{\sigma(i)} \cdot \mathbb{1}_{X_{\sigma(i)} < T} \cdot \mathbb{1}_{X_{\sigma(i-1)} \geq T} + X_{\sigma(i)} \cdot \mathbb{1}_{X_{\sigma(i)} \geq T} \cdot \mathbb{1}_{X_{\sigma(i-1)} < T}.$$

In the last period, the algorithm never buys and always sells if it has the item. Thus, from the same reasons as above we have that  $\text{ALG}_T(n) = X_{\sigma(n)} \cdot \mathbb{1}_{X_{\sigma(n-1)} < T}$ . Putting all together we obtain

$$\begin{aligned} \mathbb{E}(\text{ALG}_T) &= \sum_{i=1}^n \mathbb{E}(\text{ALG}_T(i)) \\ &= -\mathbb{E}(X_{\sigma(1)} \cdot \mathbb{1}_{X_{\sigma(1)} < T}) + \sum_{i=2}^{n-1} \mathbb{E}(X_{\sigma(i)} \cdot \mathbb{1}_{X_{\sigma(i)} \geq T} \cdot \mathbb{1}_{X_{\sigma(i-1)} < T}) \\ &\quad - \sum_{i=2}^{n-1} \mathbb{E}(X_{\sigma(i)} \cdot \mathbb{1}_{X_{\sigma(i)} < T} \cdot \mathbb{1}_{X_{\sigma(i-1)} \geq T}) + \mathbb{E}(X_{\sigma(n)} \cdot \mathbb{1}_{X_{\sigma(n-1)} < T}) \\ &\stackrel{(*)}{=} (n-2) \cdot \mathbb{E}(X_{\sigma(1)} \cdot \mathbb{1}_{X_{\sigma(1)} \geq T} \cdot \mathbb{1}_{X_{\sigma(2)} < T} - X_{\sigma(2)} \cdot \mathbb{1}_{X_{\sigma(2)} < T} \cdot \mathbb{1}_{X_{\sigma(1)} \geq T}) \\ &\quad + \mathbb{E}(X_{\sigma(1)} \cdot \mathbb{1}_{X_{\sigma(2)} < T}) - \mathbb{E}(X_{\sigma(1)} \cdot \mathbb{1}_{X_{\sigma(1)} < T}), \end{aligned}$$

where (\*) follows from linearity of expectation and the fact that  $\sigma$  is a uniformly random permutation. By manipulating the last two terms of the above sum, we obtain

$$\begin{aligned} & \mathbb{E}(X_{\sigma(1)} \cdot \mathbb{1}_{X_{\sigma(2)} < T}) - \mathbb{E}(X_{\sigma(1)} \cdot \mathbb{1}_{X_{\sigma(1)} < T}) \\ &= \mathbb{E}(X_{\sigma(1)} \cdot \mathbb{1}_{X_{\sigma(2)} < T} \cdot \mathbb{1}_{X_{\sigma(1)} < T} + X_{\sigma(1)} \cdot \mathbb{1}_{X_{\sigma(2)} < T} \cdot \mathbb{1}_{X_{\sigma(1)} \geq T}) \\ &\quad - \mathbb{E}(X_{\sigma(1)} \cdot \mathbb{1}_{X_{\sigma(1)} < T} \cdot \mathbb{1}_{X_{\sigma(2)} \geq T} + X_{\sigma(1)} \cdot \mathbb{1}_{X_{\sigma(1)} < T} \cdot \mathbb{1}_{X_{\sigma(2)} < T}) \\ &= \mathbb{E}(X_{\sigma(1)} \cdot \mathbb{1}_{X_{\sigma(2)} < T} \cdot \mathbb{1}_{X_{\sigma(1)} \geq T}) - \mathbb{E}(X_{\sigma(1)} \cdot \mathbb{1}_{X_{\sigma(1)} < T} \cdot \mathbb{1}_{X_{\sigma(2)} \geq T}) \\ &= \mathbb{E}(X_{\sigma(1)} \cdot \mathbb{1}_{X_{\sigma(1)} \geq T} \cdot \mathbb{1}_{X_{\sigma(2)} < T} - X_{\sigma(2)} \cdot \mathbb{1}_{X_{\sigma(2)} < T} \cdot \mathbb{1}_{X_{\sigma(1)} \geq T}). \end{aligned}$$

Substituting this back into the formula for  $\mathbb{E}(\text{ALG}_T)$  yields

$$\begin{aligned} \mathbb{E}(\text{ALG}_T) &= (n-1) \cdot \mathbb{E}(X_{\sigma(1)} \cdot \mathbb{1}_{X_{\sigma(1)} \geq T} \cdot \mathbb{1}_{X_{\sigma(2)} < T} - X_{\sigma(2)} \cdot \mathbb{1}_{X_{\sigma(2)} < T} \cdot \mathbb{1}_{X_{\sigma(1)} \geq T}) \\ &= (n-1) \cdot \mathbb{E}((X_{\sigma(1)} - X_{\sigma(2)}) \cdot \mathbb{1}_{X_{\sigma(1)} \geq T} \cdot \mathbb{1}_{X_{\sigma(2)} < T}) \\ &= \frac{n-1}{2} \cdot \mathbb{E}(|X_{\sigma(1)} - X_{\sigma(2)}| \cdot \mathbb{1}_{T \in [\min(X_{\sigma(1)}, X_{\sigma(2)}), \max(X_{\sigma(1)}, X_{\sigma(2)})]}), \end{aligned}$$

where the last lines follows from the fact that  $\sigma$  is a uniformly random permutation.  $\square$

#### 4.2 Proof of Theorem 2

Lemma 3 and Lemma 4 imply that, in order to show Theorem 2, it suffices to show that there exists a threshold  $T \in \mathbb{R}$  such that

$$\mathbb{E}(|X_{\sigma(1)} - X_{\sigma(2)}| \cdot \mathbb{1}_{T \in [\min(X_{\sigma(1)}, X_{\sigma(2)}), \max(X_{\sigma(1)}, X_{\sigma(2)})]}) \geq \frac{1}{16} \cdot \mathbb{E}(|X_{\sigma(1)} - X_{\sigma(2)}|). \quad (1)$$

The main difficulty in showing Inequality (1) is in the fact that the random variables  $X_{\sigma(1)}$  and  $X_{\sigma(2)}$  are *not* independent. Indeed, if they were independent then the inequality would be implied by the following key lemma.

**LEMMA 5.** *Let  $X_1, X_2$  be two independent prices with distributions  $F_1, F_2$ . Then, there exists a threshold  $T \in \mathbb{R}$  such that*

$$\mathbb{E}(|X_1 - X_2| \cdot \mathbb{1}_{T \in [\min(X_1, X_2), \max(X_1, X_2)]}) \geq \frac{1}{4} \cdot \mathbb{E}(|X_1 - X_2|) \quad (2)$$

In the appendix, we show that Inequality (2) can be fulfilled by setting  $T$  to at least one of  $M_1$  and  $M_2$ , the medians of  $F_1$  and  $F_2$ , respectively. Note that this is implied by

$$\begin{aligned} & 2 \left( \mathbb{E}(|X_1 - X_2| \cdot \mathbb{1}_{M_1 \in [\min(X_1, X_2), \max(X_1, X_2)]}) + \mathbb{E}(|X_1 - X_2| \cdot \mathbb{1}_{M_2 \in [\min(X_1, X_2), \max(X_1, X_2)]}) \right) \\ & \geq \mathbb{E}(|X_1 - X_2|). \end{aligned}$$

A way to view the proof is that we charge any elementary event, say,  $X_1 = x_1$  and  $X_2 = x_2$ , to another elementary event  $X_1 = x'_1, X_2 = x'_2$  such that

- (i)  $|x_1 - x_2| \leq |x'_1 - x'_2|$ ,
- (ii)  $M_1 \in [\min(x'_1, x'_2), \max(x'_1, x'_2)]$  or  $M_2 \in [\min(x'_1, x'_2), \max(x'_1, x'_2)]$ ,
- (iii) no elementary event is charged to more than two times.

First note that we can charge each event  $X_1 = x_1, X_2 = x_2$  satisfying  $M_1 \in [\min(x_1, x_2), \max(x_1, x_2)]$  or  $M_2 \in [\min(x_1, x_2), \max(x_1, x_2)]$  to itself. For the other events, we distinguish different cases visualized in Figure 2 and charge in such a way that each realization of the former type is only charged to one additional time. In the appendix, we also show that the analysis is tight for this way of setting the threshold.

Using the lemma, we now proceed to the proof of the theorem.




Fig. 2. The different cases distinguished in the proof of Lemma 5. Realizations of  $X_1$  and  $X_2$  are shown in red and blue, respectively. The realizations that are charged to are depicted in lighter colors. While the specific values depend on the actual distributions, what is unaffected is the order of the realizations with respect to the median of the corresponding distribution as well as among each other.

PROOF OF THEOREM 2. The goal of this proof is constructing two *independent* random variables such that if they are put to Inequality (1) they, with a loss of only a constant factor, estimate the right-hand side from above and the left-hand side from below.

Let set  $H_1 \subseteq \{1, \dots, n\}$  be a set chosen uniformly at random from all subset of size  $\lceil \frac{n}{2} \rceil$ . Let  $a, b$  be random elements chosen uniformly from  $H_1$  and  $\{1, \dots, n\} \setminus H_1$ . For any two  $i \neq j, i, j \in \{1, \dots, n\}$  we have that

$$\mathbb{P}(a = i, b = j) = \mathbb{P}(\sigma(1) = i, \sigma(2) = j) = \frac{1}{n(n-1)},$$

which gives us that

$$\mathbb{E}(|X_a - X_b|) = \mathbb{E}(|X_{\sigma(1)} - X_{\sigma(2)}|).$$

Since choice of each subset of size  $\lceil \frac{n}{2} \rceil$  of set  $\{1, \dots, n\}$  is equally likely and happens with probability  $1/\binom{n}{\lceil n/2 \rceil}$ , there must exist a set  $S$  such that

$$\mathbb{E}(|X_{a'} - X_{b'}|) \geq \mathbb{E}(|X_{\sigma(1)} - X_{\sigma(2)}|), \tag{3}$$

where  $a'$  is a uniformly random element from  $S$  while  $b'$  is a uniformly random element from  $\{1, \dots, n\} \setminus S$ . Note here that  $X_{a'}$  and  $X_{b'}$  are independent random variables since the sets of random variables  $\{X_i | i \in S\}$  and  $\{X_i | i \in \{1, \dots, n\} \setminus S\}$  are pair-wise independent. Therefore, we can apply Lemma 5 and get a threshold  $T \in \mathbb{R}$  such that

$$\mathbb{E}(|X_{a'} - X_{b'}| \cdot \mathbb{1}_{T \in [\min(X_{a'}, X_{b'}), \max(X_{a'}, X_{b'})]}) \geq \frac{1}{4} \mathbb{E}(|X_{a'} - X_{b'}|).$$

This together with Inequality (3) yields

$$\mathbb{E}(|X_{a'} - X_{b'}| \cdot \mathbb{1}_{T \in [\min(X_{a'}, X_{b'}), \max(X_{a'}, X_{b'})]}) \geq \frac{1}{4} \mathbb{E}(|X_{\sigma(1)} - X_{\sigma(2)}|). \tag{4}$$

To the end, observe that

$$4 \cdot \mathbb{E}(|X_{\sigma(1)} - X_{\sigma(2)}| \cdot \mathbb{1}_{T \in [\min(X_{\sigma(1)}, X_{\sigma(2)}), \max(X_{\sigma(1)}, X_{\sigma(2)})]})$$

$$\begin{aligned}
 &= \sum_{i,j \in \{1, \dots, n\}, i \neq j} \frac{4}{n(n-1)} \mathbb{E}(|X_i - X_j| \cdot \mathbb{1}_{T \in [\min(X_i, X_j), \max(X_i, X_j)]}) \\
 &\geq \sum_{i \in S, j \in \{1, \dots, n\} \setminus S} \frac{1}{\lceil n/2 \rceil} \cdot \frac{1}{n - \lceil n/2 \rceil} \mathbb{E}(|X_i - X_j| \cdot \mathbb{1}_{T \in [\min(X_i, X_j), \max(X_i, X_j)]}) \\
 &= \mathbb{E}(|X_{a'} - X_{b'}| \cdot \mathbb{1}_{T \in [\min(X_{a'}, X_{b'}), \max(X_{a'}, X_{b'})]}),
 \end{aligned}$$

where the inequality above follows from the following reasons. First, all random variables on both sides of the inequality are non-negative. Second, on the left-hand side the sum iterates over all possible distinct pairs  $i, j$  while on the right-hand side the sum iterates only over the pairs for which  $i$  belongs to  $S$  and  $j$  belongs to the complement of  $S$ . Third, a simple case analysis assures that  $\frac{4}{n(n-1)} = \frac{2}{n} \cdot \frac{2}{n-1} \geq \frac{1}{\lceil n/2 \rceil} \cdot \frac{1}{n - \lceil n/2 \rceil}$  for all integer  $n$ . This together with Inequality (4) proves the theorem.  $\square$

### 4.3 Proof of Theorem 3

Let us now turn to proving Theorem 3. We will be analyzing the fraction

$$\frac{\mathbb{E}(|X_1 - X_2|)}{\mathbb{E}(|X_1 - X_2| \cdot \mathbb{1}_{T \in [\min(X_1, X_2), \max(X_1, X_2)]})},$$

where  $T$  is the median of the mixture random variable  $\frac{1}{n}X_1 + \dots + \frac{1}{n}X_n$ . We will relate the nominator and denominator of the above fraction to two independent random variables of the *same* distribution which makes the key difference compared to the previous approach.

Consider a random vector  $(X_1, \dots, X_n)$  and its independent copy with the same distribution  $(X'_1, \dots, X'_n)$ . Let  $Y, Y'$  be two independent random numbers from 1 to  $n$ . Then, the variables  $X_Y$  and  $X'_{Y'}$  are also independent and have the same distribution.

We define

$$S_1 := \mathbb{E}(|X_{\sigma(1)} - X_{\sigma(2)}|), \text{ and } S_2 := \mathbb{E}(|X_{\sigma(1)} - X_{\sigma(2)}| \cdot \mathbb{1}_{T \in [\min(X_{\sigma(1)}, X_{\sigma(2)}), \max(X_{\sigma(1)}, X_{\sigma(2)})]}),$$

as well as

$$S'_1 := \mathbb{E}(|X_Y - X'_{Y'}|), \text{ and } S'_2 := \mathbb{E}(|X_Y - X'_{Y'}| \cdot \mathbb{1}_{T \in [\min(X_Y, X'_{Y'}), \max(X_Y, X'_{Y'})]}).$$

Since  $T$  is the median of  $X_Y$  and  $X'_{Y'}$ , thus by Lemma 1 we have that  $2 \geq S'_1/S'_2$ . The following lemma relates  $S'_2$  to  $S_2$ .

**LEMMA 6.** *If  $T$  is the median of distribution  $X_Y$ , then there exists a constant  $C > 0$ , independent from  $X_Y$ , such that the following inequality holds:*

$$\left(1 + \frac{C}{n}\right) \cdot S_2 \geq \frac{n}{n-1} \cdot S'_2.$$

**PROOF.** Throughout this proof, for any real  $A, B$ , we will use  $[A, B]$  to denote a continuous interval  $[\min(A, B), \max(A, B)]$ . We have that

$$\frac{n}{n-1} \cdot S'_2 = \frac{n}{n-1} \cdot \mathbb{E}(|X_Y - X'_{Y'}| \cdot \mathbb{1}_{T \in [X_Y, X'_{Y'}]}) = \frac{1}{(n-1)n} \cdot \sum_{i,i' \in [n]} \frac{1}{n^2} \cdot \mathbb{E}(|X_i - X'_{i'}| \cdot \mathbb{1}_{T \in [X_i, X'_{i'}]}).$$

On the other hand it holds that

$$\begin{aligned}
 \left(1 + \frac{C}{n}\right) \cdot S_2 &= \left(1 + \frac{C}{n}\right) \cdot \mathbb{E}(|X_{\sigma(1)} - X_{\sigma(2)}| \cdot \mathbb{1}_{T \in [X_{\sigma(1)}, X_{\sigma(2)}]}) \\
 &= \left(1 + \frac{C}{n}\right) \cdot \sum_{i \neq j \in [n]} \frac{1}{(n-1)n} \cdot \mathbb{E}(|X_i - X_j| \cdot \mathbb{1}_{T \in [X_i, X_j]}).
 \end{aligned}$$

Since  $X_i$  and  $X'_j$  are independent if  $i \neq j$  it makes sense to subtract  $\sum_{i \neq j \in [n]} \frac{1}{(n-1)n} \cdot \mathbb{E}(|X_i - X_j| \cdot \mathbb{1}_{T \in [X_i, X_j]})$  from both sides of the original inequality, which reduces our task to proving

$$\frac{C}{(n-1)n^2} \cdot \sum_{i \neq j, i, j \in [n]} \mathbb{E}(|X_i - X'_j| \cdot \mathbb{1}_{T \in [X_i, X'_j]}) \geq \frac{1}{n(n-1)} \cdot \sum_{i \in [n]} \mathbb{E}(|X_i - X'_i| \cdot \mathbb{1}_{T \in [X_i, X'_i]}).$$

Let us now fix  $i \in [n]$ . We will show that

$$\frac{C}{n} \cdot \sum_{j \in [n], j \neq i} \mathbb{E}(|X_i - X'_j| \cdot \mathbb{1}_{T \in [X_i, X'_j]}) \geq \mathbb{E}(|X_i - X'_i| \cdot \mathbb{1}_{T \in [X_i, X'_i]}), \quad (5)$$

which, if summed over all choices of  $i \in [n]$ , will prove the lemma. To do so, let  $A$  be the set of the indices  $k$  for which  $\mathbb{P}(X'_k > T) \geq 1/4$ . Observe that  $|A| \geq n/4$ . If not, then we have  $\mathbb{P}(X'_Y > T) < \mathbb{P}(Y \in A) + \frac{1}{4} \cdot \mathbb{P}(Y \notin A) < \frac{1}{4} + \frac{1}{4} = \frac{1}{2}$ , which contradicts with the choice of  $T$ . An analogous argument shows that set  $B := \left\{ k \mid \mathbb{P}(X'_k \leq T) \geq \frac{1}{4} \right\}$  has size at least  $n/4$ . Consider now the sum  $\sum_{j \in [n], j \neq i} \mathbb{E}(|X_i - X'_j| \cdot \mathbb{1}_{T \in [X_i, X'_j]})$ . We have that

$$\begin{aligned} & 2 \sum_{j \in [n], j \neq i} \mathbb{E}(|X_i - X'_j| \cdot \mathbb{1}_{T \in [X_i, X'_j]}) \\ & \geq \sum_{a \in A-i} \mathbb{E}(|X_i - X'_a| \cdot \mathbb{1}_{X_i \leq T} \cdot \mathbb{1}_{X'_a > T}) + \sum_{b \in B-i} \mathbb{E}(|X_i - X'_b| \cdot \mathbb{1}_{X_i > T} \cdot \mathbb{1}_{X'_b \leq T}) \\ & = \sum_{a \in A-i} \mathbb{E}((X'_a - X_i) \cdot \mathbb{1}_{X_i \leq T} \cdot \mathbb{1}_{X'_a > T}) + \sum_{b \in B-i} \mathbb{E}((X_i - X'_b) \cdot \mathbb{1}_{X_i > T} \cdot \mathbb{1}_{X'_b \leq T}) \\ & \geq \sum_{a \in A-i} \mathbb{E}((T - X_i) \cdot \mathbb{1}_{X_i \leq T} \cdot \mathbb{1}_{X'_a > T}) + \sum_{b \in B-i} \mathbb{E}((X_i - T) \cdot \mathbb{1}_{X_i > T} \cdot \mathbb{1}_{X'_b \leq T}), \end{aligned} \quad (6)$$

where the last inequality follows from the fact that  $X_a > T$  and  $X_b \leq T$ . Since for every  $a \in A-i$  variable  $X'_a$  is independent from  $X_i$  and, by the choice of  $A$ , we have  $\mathbb{P}(X'_a > T) \geq 1/4$ , it holds that

$$\sum_{a \in A-i} \mathbb{E}((T - X_i) \cdot \mathbb{1}_{X_i \leq T} \cdot \mathbb{1}_{X'_a > T}) \geq \frac{1}{4} \cdot \left( \frac{n}{4} - 1 \right) \cdot \mathbb{E}((T - X_i) \cdot \mathbb{1}_{X_i \leq T}).$$

The last inequality follows from the fact that  $|A| \geq \frac{n}{4}$ . By a symmetric reasoning for set  $B$ , we see that

$$\sum_{b \in B-i} \mathbb{E}((X_i - T) \cdot \mathbb{1}_{X_i > T} \cdot \mathbb{1}_{X'_b \leq T}) \geq \frac{1}{4} \cdot \left( \frac{n}{4} - 1 \right) \cdot \mathbb{E}((X_i - T) \cdot \mathbb{1}_{X_i > T}).$$

The two above inequalities combined with Inequality (6) gives us that

$$\sum_{j \in [n], j \neq i} 2\mathbb{E}(|X_i - X'_j| \cdot \mathbb{1}_{T \in [X_i, X'_j]}) \geq \left( \frac{n}{16} - \frac{1}{4} \right) \cdot (\mathbb{E}((T - X_i) \cdot \mathbb{1}_{X_i \leq T}) + \mathbb{E}((X_i - T) \cdot \mathbb{1}_{X_i > T})) \quad (7)$$

On the other hand, the left-hand side of (Inequality (5)) can be rewritten as

$$\begin{aligned} & \mathbb{E}(|X_i - X'_i| \cdot \mathbb{1}_{T \in [X_i, X'_i]}) \\ & = \mathbb{E}((X_i - X'_i) \cdot \mathbb{1}_{X_i \geq T} \cdot \mathbb{1}_{X'_i < T}) + \mathbb{E}((X'_i - X_i) \cdot \mathbb{1}_{X_i < T} \cdot \mathbb{1}_{X'_i \geq T}) = 2\mathbb{E}((X_i - X'_i) \cdot \mathbb{1}_{X_i \geq T} \cdot \mathbb{1}_{X'_i < T}) \end{aligned}$$

since  $X_i$  and  $X'_i$  are independent random variables with the same distribution. By linearity of expectation we get

$$2\mathbb{E}((X_i - X'_i) \mathbb{1}_{X_i \geq T} \cdot \mathbb{1}_{X'_i < T}) \leq 2(\mathbb{E}((T - X'_i) \mathbb{1}_{X_i \geq T} \cdot \mathbb{1}_{X'_i < T}) + \mathbb{E}((X_i - T) \cdot \mathbb{1}_{X_i \geq T} \cdot \mathbb{1}_{X'_i < T}))$$

$$\leq 2(\mathbb{E}((T - X'_i) \cdot \mathbb{1}_{X'_i < T}) + \mathbb{E}((X_i - T) \cdot \mathbb{1}_{X_i \geq T})),$$

where the last inequality follows from the independence of  $X_i$  and  $X'_i$ . This inequality, combined with Inequality (7) leads to

$$\begin{aligned} \sum_{j \in [n], j \neq i} 2\mathbb{E}(|X_i - X'_j| \cdot \mathbb{1}_{T \in [X_i, X'_j]}) &\geq \left(\frac{n}{16} - \frac{1}{4}\right) \cdot (\mathbb{E}((T - X_i) \cdot \mathbb{1}_{X_i \leq T}) + \mathbb{E}((X_i - T) \cdot \mathbb{1}_{X_i > T})) \\ &\geq \frac{1}{2} \cdot \left(\frac{n}{16} - \frac{1}{4}\right) \cdot \mathbb{E}(|X_i - X'_i| \cdot \mathbb{1}_{T \in [X_i, X'_i]}), \end{aligned}$$

which is equivalent to

$$\frac{64}{n} \cdot \mathbb{E}(|X_i - X'_j| \cdot \mathbb{1}_{T \in [X_i, X'_j]}) \geq \mathbb{E}(|X_i - X'_i| \cdot \mathbb{1}_{T \in [X_i, X'_i]}).$$

This proves the claimed Inequality (5) with constant  $C := 64$ , and therefore the lemma follows.  $\square$

We are now ready to finish the proof of Theorem 3.

PROOF OF THEOREM 3. Recall that our goal is to prove

$$\left(1 + \frac{C}{n - C}\right) \cdot 2 \geq \frac{S_1}{S_2}. \tag{8}$$

To do so, first observe that

$$\begin{aligned} S_1 &= \mathbb{E}(|X_{\sigma(1)} - X_{\sigma(2)}|) = \frac{n}{n - 1} \cdot \mathbb{E}(|X_Y - X'_{Y'}|) - \frac{1}{n(n - 1)} \cdot \sum_{i \in [n]} \mathbb{E}(|X_i - X'_i|) \\ &= \frac{n}{n - 1} \cdot S'_1 - \frac{1}{n(n - 1)} \cdot \sum_{i \in [n]} \mathbb{E}(|X_i - X'_i|), \end{aligned}$$

which leads to

$$\frac{n}{n - 1} \cdot S'_1 = S_1 + \frac{1}{n(n - 1)} \cdot \sum_{i \in [n]} \mathbb{E}(|X_i - X'_i|).$$

As observed before, we have that  $2 \geq \frac{S'_1}{S'_2}$  which combined with the above equality gets that

$$2 \geq \frac{S_1 + \frac{1}{n(n-1)} \cdot \sum_{i \in [n]} \mathbb{E}(|X_i - X'_i|)}{S'_2} \geq \frac{S_1}{S'_2}, \tag{9}$$

where in the last inequality we used the fact that  $\sum_{i \in [n]} \mathbb{E}(|X_i - X'_i|) \geq 0$ . By Lemma 6 we get that there exists a constant  $C$  such that  $\left(1 + \frac{C}{n}\right) \cdot S_2 \geq \frac{n}{n-1} \cdot S'_2 \geq S'_2$ , which, when plugged into the denominator of the right-hand side of Inequality (9), implies  $2 \cdot \left(1 + \frac{C}{n}\right) \cdot S_2 \geq S_1$ . This completes the proof of the theorem.  $\square$

## 5 GENERALIZATIONS

In the following subsections, we consider four generalizations of our basic trading-prophet problem: An unknown-distribution version with sample access, a version with more than one item, a budgeted version with re-investments of gains, and a multi-armed-bandit version.



### 5.1 Unknown Distribution and Affiliated Prices

We consider in this section a variant of the model studied in Section 3 in which the prices are i.i.d., but we are not given the distribution beforehand. We prove the result for i.i.d. prices, but notice that the exact same argument applies to affiliated prices where  $p_j = x_j + y$  and  $y \sim G$  and  $x_j \sim F$  for all  $j$  independently. The reason is that once  $y$  is fixed the increments are i.i.d., and both the online algorithm and the optimal offline algorithm buy and sell the same number of times so that  $y$  cancels out.

**THEOREM 4.** *If prices are i.i.d. or affiliated, but we do not know the distribution, there is an algorithm ALG such that  $\mathbb{E}(\text{ALG}) \geq \frac{1}{2} \cdot \frac{n-2}{n-1} \cdot \mathbb{E}(\text{OPT})$ .*

To prove this result we first analyze the case where, instead of knowing the distribution, we have access to  $n - 1$  independent samples  $S_1, S_2, \dots, S_{n-1}$ . For this case, consider the following algorithm we denote by  $\text{ALG}^s$ . In period  $i < n$ , if we do not have the item, we buy if  $X_i < S_i$ ; if we have the item, we sell if  $X_i \geq S_i$ . In period  $n$  we simply sell the item if we still have it. For the case of affiliated prices, we require that the samples are also correlated, i.e., that  $S_j = y + s_j$  and  $X_j = y + x_j$ , where  $y \sim G$  and  $s_j \sim F, x_j \sim F$  for all  $j$ .

**LEMMA 7.**  $\mathbb{E}(\text{ALG}^s) \geq \frac{n-1}{4} \cdot \mathbb{E}(|X_1 - X_2|)$ .

The proof of this lemma can be found in the full version and is very similar to the proof of Lemma 2. The key observation is that, as in the median-threshold algorithm, at any given period  $i \geq 2$  the probability that we have the item is exactly  $1/2$ . Combining this lemma with Lemma 1 we have that  $\text{ALG}^s$  is actually a  $1/2$  approximation. From  $\text{ALG}^s$  we derive  $\text{ALG}^{s*}$ , an algorithm that uses a single sample  $S_1$  of the distribution. For  $2 \leq i \leq n - 1$ , we define  $S'_i$  by selecting a uniformly random element from  $\{S_1, X_1, \dots, X_{i-1}\}$ . Now,  $\text{ALG}^{s*}$  behaves like  $\text{ALG}^s$ , taking  $S_1, S'_2, \dots, S'_{n-1}$  as the set of samples.

**LEMMA 8.**  $\mathbb{E}(\text{ALG}^{s*}) = \mathbb{E}(\text{ALG}^s)$ .

The main observation needed for the proof of this lemma is that, when analyzing the expected gains,  $S_1, S'_2, \dots, S'_{n-1}$  essentially behave as independent samples.

Now, Theorem 4 is a direct application of Lemma 8. If we do not know the distribution, we can skip the first period and use  $\text{ALG}^{s*}$  in  $X_2, \dots, X_n$ , treating  $X_1$  as a sample. This yields an expected profit of  $\frac{n-2}{4} \cdot \mathbb{E}(|X_1 - X_2|)$ , which by Lemma 1 is exactly  $\frac{1}{2} \cdot \frac{n-2}{n-1} \cdot \mathbb{E}(\text{OPT})$ .

### 5.2 More Than One Item

In this section we analyze a variant of the main model where we are allowed to store up to  $k$  copies of the item and the prices are independent and are presented in uniformly random order. It is clear that, in hindsight, the optimal strategy always buy  $k$  items or sell all  $k$  items. We prove the optimal online strategy also presents this behavior, which implies that this variant is equivalent to the single item case.

**LEMMA 9.** *If we are allowed to store  $k$  items at any time, then the optimal online strategy always trades with all  $k$  items, i.e. in a single period either the algorithm buys  $k$  items or it sells all  $k$  items.*

The proof of this lemma involves writing down the value function in a given step of the backward induction argument, and showing it is linear in the number of items we buy or sell. This implies that the optimal decision is always an extreme point: We should always buy or sell 0 or  $k$  items.

### 5.3 Budgeted Version with Fractional Purchase and Re-investment of Gains

A natural extension of our model is to allow, as in a stock market, the transaction of any fraction of the item; and at a given period, only limit ourselves by the budget we have in that period. More precisely, in this section we consider the following model. We are given a sequence of positive prices  $X_{\sigma(1)}, \dots, X_{\sigma(n)}$  generated in the same way as in our basic model. In period  $i$  we have a state given by a pair  $(S_i, B_i)$ , where  $S_i \geq 0$  is the (possibly fractional) number of stocks we hold in period  $i$ , and  $B_i \geq 0$  is our budget in period  $i$ . We start with  $(S_1, B_1) = (0, 1)$ , meaning that we start with 0 stocks and a budget of 1 unit of money. In period  $i$  we can buy (or sell) any number of stocks  $s \in [-S_i, B_i/X_{\sigma(i)}]$ , which determines the state in the next period as  $(S_{i+1}, B_{i+1}) = (S_i + s, B_i - s \cdot X_{\sigma(i)})$ . Our objective is to maximize the money we have at the end of the process, i.e.,  $B_{n+1}$ .

Denote by  $\text{OPT}_F$  the gains of the prophet in this model, i.e., the maximum possible profit in hindsight. It is easy to see that the prophet either sells all stocks or spends all money, i.e.,  $s \in \{-S_i, B_i/X_{\sigma(i)}\}$ . Thus, if  $I$  denotes the set of periods where the prophet buys, and  $J$  the periods where the prophet sells, we have that  $\text{OPT}_F = \prod_{j \in J} X_{\sigma(j)} / \prod_{i \in I} X_{\sigma(i)}$ . This formula immediately implies that, in this model, the prophet behaves as the prophet of the basic model with prices  $X'_{\sigma(i)} = \log(X_{\sigma(i)})$ . In turn, this implies that the prophet has an expected profit that grows exponentially in  $n$ .

**OBSERVATION 2.** *By Jensen's inequality and Lemma 3, we have that*

$$\mathbb{E}(\text{OPT}_F) \geq \exp(\mathbb{E}(\log \text{OPT}_F)) = \exp\left(\frac{n-1}{2} \cdot \mathbb{E}(|\log X_{\sigma(1)} - \log X_{\sigma(2)}|)\right).$$

It is clear from this that there is no hope for a constant approximation against  $\mathbb{E}(\text{OPT}_F)$ . Instead, we can approximate the expected growth rate of  $\text{OPT}_F$ , i.e.  $\mathbb{E}(\log \text{OPT}_F)$ . Indeed, by considering prices  $X'_{\sigma(i)} = \log(X_{\sigma(i)})$ , we can apply our algorithms from Sections 3 and 4 to obtain approximately optimal expected growth rates. In this model, for  $T \in \mathbb{R}_+$  we denote by  $\text{ALG}_T$  the algorithm that sells all stocks if  $X_{\sigma(i)} \geq T$  or  $i = n$ , and spends the whole budget when  $X_{\sigma(i)} < T$  and  $i < n$ .

**COROLLARY 1.** *If the prices are i.i.d. and  $T$  is the median of the distribution, then  $\mathbb{E}(\log \text{ALG}_T) \geq \frac{1}{2}\mathbb{E}(\log \text{OPT}_F)$ .*

**COROLLARY 2.** *If the prices are independent and presented in uniformly random order, then there exists a threshold  $T$  such that  $\mathbb{E}(\log \text{ALG}_T) \geq \frac{1}{16}\mathbb{E}(\log \text{OPT}_F)$ .*

**COROLLARY 3.** *There is a constant  $C$  so that if prices are independent and presented in uniformly random order, then there exists a threshold  $T$  such that  $\mathbb{E}(\log \text{ALG}_T) \geq \frac{1}{2 \cdot (1+C/n)}\mathbb{E}(\log \text{OPT}_F)$ .*

### 5.4 Multi-armed-bandit version

Consider the following model we call  $k$ -armed bandit trading prophet. We are able to trade  $k$  different kinds of items, each with a sequence of prices generated independently as in the basic model, but we can hold at most one item, across all kinds. More precisely, let us denote  $F_1^i, \dots, F_n^i$  the initial  $n$  distributions of an item  $1 \leq i \leq k$  which are all independent. Let  $X_j^i \sim F_j^i$  for  $1 \leq j \leq n$  denote  $j$ -th price of  $i$ -th item. On period  $i$  we observe  $k$  different prices  $X_{\sigma^1(i)}^1, \dots, X_{\sigma^k(i)}^k$  where  $\sigma^j : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$  for  $1 \leq j \leq k$  are  $k$  permutations chosen independently from the uniform distribution over all permutations. Like in the basic model, in each period we have two possibilities: Either we have an item  $j$  for  $1 \leq j \leq k$ , and in this case we can sell this item with the price  $X_{\sigma^j(i)}^j$ ; or we do not have *any* item, and in

this case we can buy any item  $j$  for  $1 \leq j \leq k$  for the price  $X_{\sigma^j(i)}^j$ . On top of that, we assume that if in a given period we have an item, then we can sell it and within the same period buy a different one.<sup>3</sup> We start at period 1 with no item and, as before, we want to design a decision rule that maximizes the expected profit obtained after  $n$  periods.

Consider now an algorithm that in the beginning of an execution chooses one item uniformly at random, and next trades only on this item applying the threshold decision rule from Section 4. In the following, we will show that this algorithm carries the results stated in Theorem 2 and Theorem 3 with  $\frac{1}{k}$  multiplicative loss, against the prophet. The key observation involves understanding the strategy of the prophet in the multi-armed bandit version.

**LEMMA 10.** *The expected gain of the optimal strategy in hindsight for the  $k$ -armed bandit trading-prophet problem is*

$$\mathbb{E}(\text{OPT}) = (n - 1) \cdot \mathbb{E} \left( \max_{i \in [k]} \left[ X_{\sigma^i(1)}^i - X_{\sigma^i(2)}^i \right]_+ \right), \text{ where } [\cdot]_+ := \max\{\cdot, 0\}.$$

To prove this lemma, we simply take the expected gain of OPT between each pair of consecutive periods and notice that this expected gain is equal across all these pairs. A formal proof is deferred to the full version. This result allows us to immediately extend Theorems 2 and 3.

**THEOREM 5.** *There exist algorithms  $\text{ALG}_1$  and  $\text{ALG}_2$  that achieve the following approximation for  $k$ -armed bandit version of trading prophets:*

$$\mathbb{E}(\text{ALG}_1) \geq \frac{1}{16k} \cdot \mathbb{E}(\text{OPT}); \quad \mathbb{E}(\text{ALG}_2) \geq \left( \frac{1}{2k} - o(1) \right) \cdot \mathbb{E}(\text{OPT}).$$

The natural question is whether there is any algorithm with an approximation factor  $o(1/k)$ . Not surprisingly, the answer is negative even we restrict to i.i.d. random variables  $X^i$ .

**LEMMA 11.** *For  $k \geq 2$ , there is an instance of the  $k$ -armed bandit trading-prophet problem with i.i.d. prices  $X_j^i$  such that for any algorithm ALG it holds that*

$$\mathbb{E}(\text{ALG}) \leq \frac{1}{(1 - e^{-1/2})k} \cdot \mathbb{E}(\text{OPT}).$$

In our instance, there are only two periods. Each price is independently  $k$  with probability  $1/k$  and 0 with probability  $(1 - 1/k)$ . Here, OPT knows exactly which item will have a high price in the second period, while ALG can only get a high price with probability  $1/k$ .

## ACKNOWLEDGMENTS

This work was initiated during the Fall 2021 Virtual Chair Semester on Prophet Inequalities [Correa et al., 2022]. We would like to thank the organizers and the participants of the semester for their feedback. Kevin Schewior was in part supported by Independent Research Fund Denmark, Natural Sciences under Grant No. DFF-0135-00018B. MohammadTaghi Hajiaghayi and Jan Olkowski are partially supported by DARPA QuICC, NSF AF:Small 2218678, and NSF AF:Small 2114269.

<sup>3</sup>Note that this assumption makes no difference in the original model, since it would be equivalent to simply keeping the item until the next period.

## REFERENCES

- Melika Abolhassani, Soheil Ehsani, Hossein Esfandiari, MohammadTaghi Hajiaghayi, Robert D. Kleinberg, and Brendan Lucier. Beating  $1 - 1/e$  for ordered prophets. In *ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 61–71, 2017.
- Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014.
- Allan Borodin and Ran El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.
- Alexander Braun and Thomas Kesselheim. Truthful mechanisms for two-sided markets via prophet inequalities. In *ACM Conference on Economics and Computation (EC)*, pages 202–203, 2021.
- Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. Approximating gains from trade in two-sided markets via simple mechanisms. In *ACM Conference on Economics and Computation (EC)*, pages 589–590, 2017.
- Constantine Caramanis, Paul Dütting, Matthew Faw, Federico Fusco, Philip Lazos, Stefano Leonardi, Orestis Papadigenopoulos, Emmanouil Pountourakis, and Rebecca Reiffenhäuser. Single-sample prophet inequalities via greedy-ordered selection. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1298–1325, 2022.
- Abraham Charnes, Jacques Drèze, and Merton Miller. Decision and horizon rules for stochastic planning problems: A linear example. *Econometrica*, 34:307–330, 1966.
- Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Approximately efficient double auctions with strong budget balance. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1424–1443, 2016.
- Riccardo Colini-Baldeschi, Paul W. Goldberg, Bart de Keijzer, Stefano Leonardi, Tim Roughgarden, and Stefano Turchetta. Approximately efficient two-sided combinatorial auctions. In *ACM Conference on Economics and Computation (EC)*, pages 591–608, 2017a.
- Riccardo Colini-Baldeschi, Paul W. Goldberg, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Fixed price approximability of the optimal gain from trade. In *International Conference on Web and Internet Economics (WINE)*, pages 146–160, 2017b.
- José R. Correa, Paul Dütting, Felix A. Fischer, and Kevin Schewior. Prophet inequalities for I.I.D. random variables from an unknown distribution. In *ACM Conference on Economics and Computation (EC)*, pages 3–17, 2019.
- José R. Correa, Andrés Cristi, Boris Epstein, and José A. Soto. The two-sided game of googol and sample-based prophet inequalities. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2066–2081, 2020.
- José R. Correa, Paul Dütting, Felix A. Fischer, Kevin Schewior, and Bruno Ziliotto. Unknown I.I.D. prophets: Better bounds, streaming algorithms and a new impossibility. In *Innovations in Theoretical Computer Science Conference (ITCS)*, pages 86:1–86:1, 2021a.
- José R. Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Posted price mechanisms and optimal threshold strategies for random arrivals. *Mathematics of Operations Research*, 46(4):1452–1478, 2021b.
- José R. Correa, Raimundo Saona, and Bruno Ziliotto. Prophet secretary through blind strategies. *Mathematical Programming*, 190(1):483–521, 2021c.
- Jose R. Correa, Jason Hartline, and Nicole Immorlica. A semester virtual institute. <https://cacm.acm.org/blogs/blog-cacm/258538-a-semester-virtual-institute/fulltext>, 2022.
- Yuan Deng, Jieming Mao, Balasubramanian Sivan, and Kangning Wang. Approximately efficient bilateral trade. In *ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 718–721, 2022.
- Jacques du Toit and Goran Peskir. The trap of complacency in predicting the maximum. *The Annals of Probability*, 35:340–365, 2007.
- Jacques du Toit and Goran Peskir. Selling a stock at the ultimate maximum. *The Annals of Applied Probability*, 19: 983–1014, 2009.
- Paul Dütting, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. Efficient two-sided markets with limited information. In *ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 1452–1465, 2021.
- Soheil Ehsani, MohammadTaghi Hajiaghayi, Thomas Kesselheim, and Sahil Singla. Prophet secretary for combinatorial auctions and matroids. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 700–714, 2018.
- Hossein Esfandiari, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Morteza Monemizadeh. Prophet secretary. *SIAM Journal on Discrete Mathematics*, 31(3):1685–1701, 2017.
- Svend E. Graversen, Goran Peskir, and Albert N. Shiryaev. Stopping brownian motion without anticipation as close as possible to its ultimate maximum. *Theory of Probability and its Applications*, 45:41–50, 2006.
- Jon M. Kleinberg and Robert Kleinberg. Delegated search approximates efficient search. In *ACM Conference on Economics and Computation (EC)*, pages 287–302, 2018.

- Ulrich Krengel and Louis Sucheston. Semiamarts and finite values. *Bulletin of the American Mathematical Society*, 83: 745–747, 1977.
- Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. *Advances in Probability and Related Topics*, 4:197–266, 1978.
- Bin Li and Steven C. Hoi. Online portfolio selection: A survey. *ACM Computing Surveys*, 46:1–33, 2014.
- Allen Liu, Renato Paes Leme, Martin Pál, Jon Schneider, and Balasubramanian Sivan. Variable decomposition for prophet inequalities and optimal ordering. In *ACM Conference on Economics and Computation (EC)*, page 692, 2021.
- R. Preston McAfee. The gains from trade under fixed price mechanisms. *Applied Economics Research Bulletin*, 1:1–10, 2008.
- Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:265–281, 1983.
- Matthew F. M. Osborne. Brownian motion in the stock market. *Operations Research*, 7(2):145–173, 1959.
- Aviad Rubinfeld, Jack Z. Wang, and S. Matthew Weinberg. Optimal single-choice prophet inequalities from samples. In *Innovations in Theoretical Computer Science Conference (ITCS)*, pages 60:1–60:10, 2020.
- Esther Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *Annals of Probability*, 12:1213–1216, 1984.
- Sahil Singla. *Combinatorial Optimization Under Uncertainty: Probing and Stopping-Time Algorithms*. PhD thesis, Carnegie Mellon University, 2018.