https://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Vol. 74, No. 1, January–February 2026, pp. 260–280
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Crosscutting Areas

Trading Prophets

José Correa,[a] Andrés Cristi,[b,*] Paul Dütting,[c] Mohammad Hajiaghayi,[d] Jan Olkowski,[d] Kevin Schewior[e,f]

[a] Departamento de Ingeniería Industrial, Universidad de Chile, Santiago 8370439, Chile; [b] College of Management of Technology, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland; [c] Google Research, 8002 Zurich, Switzerland; [d] Computer Science Department, University of Maryland, College Park, Maryland 20742; [e] Department of Mathematics and Computer Science, University of Cologne, 50969 Köln, Germany; [f] Department of Mathematics and Computer Science, University of Southern Denmark, 5230 Odense, Denmark

*Corresponding author

Contact: correa@uchile.cl, https://orcid.org/0000-0002-3012-7622 (JC); andres.cristi@epfl.ch, https://orcid.org/0000-0002-1227-2092 (AC); duetting@google.com, https://orcid.org/0000-0002-0635-6812 (PD); hajiagha@umd.edu, https://orcid.org/0000-0003-4842-0533 (MH); olkowski@umd.edu, https://orcid.org/0000-0001-9278-3935 (JO); k.schewior@uni-koeln.de, https://orcid.org/0000-0003-2236-0210 (KS)

Received: October 30, 2023
Revised: May 12, 2025; September 10, 2025
Accepted: September 15, 2025
Published Online in Articles in Advance: October 27, 2025

Area of Review: Market Analytics and Revenue Management

https://doi.org/10.1287/opre.2023.0593

Abstract. 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 identify settings in which a constant-factor buy-and-sell prophet inequality can be achieved. Interestingly, these settings are different from those in which a constant-factor standard prophet inequality can be achieved. In particular, no constant-factor buy-and-sell prophet inequality can be achieved in the case of arbitrary independent prices. In contrast, we show that in the case of independent prices arriving in random order a constant factor can be achieved. This in particular implies a constant-factor buy-and-sell prophet inequality for i.i.d. prices, for which we also show a tight ratio of $1/2$, using a single-threshold algorithm. We use the results for these base cases to solve a variety of more complex settings: affiliated prices, the single-sample setting, the case of k items and that of k item types, and a budgeted version where gains can be reinvested. We experimentally validate our results by assessing our algorithms' performance on a used car auction data set and synthetic data.

Funding: This work was supported by the Agencia Nacional de Investigación y Desarrollo [Grant FB210005], the National Science Foundation [Grants AF:Small 2114269 and AF:Small 2218678], the Danmarks Frie Forskningsfond [Grant DFF-0135-00018B], and the Defense Advanced Research Projects Agency [Grant QuICC].

Supplemental Material: All supplemental materials, including the code, data, and files required to reproduce the results, are available at https://doi.org/10.1287/opre.2023.0593.

Keywords: prophet inequality • online algorithms

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 that she 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 $1,000 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 does not want to put in extra money. Therefore, 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 modeling assumptions (Osborne 1959, Charnes et al. 1966, Myerson and Satterthwaite 1983). 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. Another option would be to model prices as following a random walk or Brownian motion.

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 compares 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. Our central goal is to identify conditions under which constant-factor prophet inequalities are achievable. Interestingly, these turn out to be distinct from those where this is possible in the standard prophet inequality problem.

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, independent random variables presented in random order, or independent random variables 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. 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 $n \geq 2$ (possibly negative) 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 first consider some fundamental settings in which the agent can hold up to one item and classify them according to whether a constant-factor buy-and-sell prophet inequality can be achieved. In contrast to the vanilla prophet inequality problem, we show that no approximation is possible when prices are independent but arrive in any order. This in particular implies that no approximation is possible with arbitrarily correlated prices. On the other hand, we show that independent prices that arrive in random order, which includes i.i.d. prices as an important special case, admit constant-factor approximations. We discuss a number of extensions, including extensions of our positive results to affiliated prices and settings with budgets in Section 1.3. Our results for the base case are summarized in Table 1.

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. Therefore, the expected reward from an intermediate step is $1/2 \cdot (\mathbb{E}[X_1 \cdot \mathbf{1}_{X_1 \geq T}] - \mathbb{E}[X_1 \cdot \mathbf{1}_{X_1 < T}])$. Now, because T is the median, $\mathbb{E}[T \cdot \mathbf{1}_{X_1 \geq T}] = \mathbb{E}[T \cdot \mathbf{1}_{X_1 < T}]$. Therefore, we can add and subtract this term from the previous formula and obtain that the expected reward from an intermediate step is

Table 1. Achievable Approximation Ratios in the Basic Versions of the Trading-Prophet Problem

		Independent prices	
	i.i.d. prices	Random order	Worst order
Upper bound	2 (Theorem 2)	16 (2 as $n \to \infty$) (Theorems 3 and 4)	–
Lower bound	2 (Proposition 1)	3 (Proposition 2)	∞ (simple example)

$1/2 \cdot \mathbb{E}[|X_1 - T|]$. 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 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 Independent Prices: Random Order. Now assume that prices p_j are first drawn independently from not necessarily 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 is still possible to achieve a $1/16$ approximation guarantee with a threshold policy, which sets a single nonadaptive 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 an approximation factor better than $1/3$ for $n = 2$. Interestingly, this impossibility applies even if the order of random variables, and not their realizations, is revealed to the algorithm before the sequence starts. We also present a threshold-based online algorithm that achieves a $(1/2 - c/n)$-approximation, for some constant c, which tends to $1/2$ as $n \to \infty$. Note that the latter result does not contradict the impossibility of $1/3$; in particular, in contrast to classic prophet inequalities, the competitive ratio is *not* invariant under adding deterministic variables of value 0.

Main Result 2. For independent 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 an approximation factor better than $1/3$.

Main Result 3. For independent prices $p_j \sim F_j$ presented in random order, setting the threshold T to the median of the mixture distribution yields a $(1/2 - c/n)$-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 σ is a uniform random permutation. Generalizing our approach for the i.i.d. case, we show that the expected reward of the prophet is $(n-1)/2 \cdot \mathbb{E}[|X_{\sigma(1)} - X_{\sigma(2)}|]$, whereas the expected reward of a threshold policy with threshold T is equal to

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

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 can be shown to buy only when $X_{\sigma(1)} = 0$ which leads to an expected profit of one. 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 Independent Prices: Adversarial Order. If prices are distributed, not necessarily 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 and therefore obtains an expected value of zero0. The prophet, however, buys X_1 precisely when $X_2 = 1/\varepsilon$ (which happens with probability ε), therefore obtaining an expected value of $\varepsilon \cdot (1/\varepsilon - 1)$. Note that this example can be extended to work also for any larger n by adding deterministic variables of value 1 in the beginning of the sequence.

(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 not necessarily 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. We in particular show how to extend our positive results to settings with affiliated prices and settings where the trader is subject to a budget constraint. We give an overview of our additional results 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. (2021b, 2022b). The same algorithm works and yields a $(1/2 - o(1))$-approximation when prices exhibit a form of correlation known as affiliation in economics; that is, 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. Therefore, 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 Reinvestment 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 is possible to achieve constant-factor approximations to the growth rate of the optimal offline policy.

Multiarmed 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 do not 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, that is, $p_j = p + \sum_{i=1}^{j} x_i$ 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. The resulting sequence of prices is a martingale. It thus follows from the well-known *optional stopping theorem* (Doob 1953) that, if one has to sell the bought item eventually, one cannot make profit. Thus, if $p \geq n$, in which case prices are nonnegative and not selling a bought item is not beneficial, the expected value of any online algorithm is nonpositive. Hence, unfortunately, 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. Experiments

We experimentally validate our results by assessing our algorithms' performance on a used car auction data set with prices from the United States and Canada, as well as synthetic data. In Section 6.1, we present our results for the used car data set. For this set of experiments, we applied our trading strategies to a fixed model and manufacturing year (in our case, "Kia Sorento 2015"). We observe that both our mean and median algorithms developed for the i.i.d. case perform close to their theoretical guarantees for the i.i.d. case, even though the empirical sequence of prices does not formally satisfy the i.i.d. assumption. The same holds for the algorithm for the unknown-distribution case, but the performance is noisier.

In Section 6.2, we explore a setting that is not i.i.d. and may serve as a good model for practical settings in which the algorithm performs well. Here, the prices are generated by adding noise to a (slow) random walk. We find that, the slower the random walk, the higher the profit that the algorithm extracts. In the limit, we obtain the i.i.d. setting and numerically recover a guarantee of close to $1/2$.

1.5. Further Related Work

To the best of our knowledge, ours is the first work to address a repeated buying and selling problem through a prophet-inequality lens. We focus on the natural objective of maximizing the gambler's profit from a sequence of buy and sell operations.

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, Kleinberg and Kleinberg 2018, Singla 2018, Correa et al. 2021c, Liu et al. 2021), and the prophet secretary problem (Esfandiari et al. 2017, Ehsani et al. 2018, Correa et al. 2021a). Our unknown distributions result is related to a recent stream of work that looks at prophet inequalities with samples (Correa et al. 2020, 2022b; Rubinstein et al. 2020; Caramanis et al. 2022).

Ideas from the prophet inequalities literature have also been used to construct simple near-optimal mechanisms for two-sided markets (Colini-Baldeschi et al. 2016, 2017b; Braun and Kesselheim 2021; 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 McAfee 2008, Brustle et al. 2017, Colini-Baldeschi et al. 2017a, and Deng et al. 2022, who show constant-factor approximation guarantees for maximizing the "gains from trade."

Recent work by Ekbatani et al. (2024) studies a prophet inequality problem, which can also be understood as modeling a situation with multiple buy and sell operations. However, the exact model and objective are incomparable. Namely, in their model, which is inspired by the "buyback model" in auctions, the gambler faces a sequence of stochastic rewards from which it can choose one. Although the gambler has to decide which reward to keep in an online fashion, the gambler is allowed to return a previously accepted reward x at a cost $f \cdot x$, where $f \geq 0$ is some fixed parameter. The gambler's utility is the reward it eventually holds minus all cancellation costs.

Additional stopping problems motivated by financial applications have been studied in the finance/applied probability literature. First, Graversen et al. (2006) and du Toit and Peskir (2007) 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 $F_1, \ldots, F_n$ on $\mathbb{R}$ and, in random order, in each of $n \geq 2$ 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 $\sigma : \{1, \ldots, n\} \to \{1, \ldots, n\}$; 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. That is, OPT is the (random) maximum profit given the realization of the prices and the permutation σ, 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 σ. We say a decision rule or algorithm ALG is an α-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 maxima. See Figure 1 for an illustration.

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, that is, $X_{\sigma(i)} < X_{\sigma(i-1)}$ 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, that is, if $X_{\sigma(i)} \geq X_{\sigma(i-1)}$ and $X_{\sigma(i)} > X_{\sigma(i+1)}$. Here, 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 choice, OPT will buy in the first period and sell in the last.

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 without loss of generality, 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 σ and simply denote by X_i the price in period i. We start our discussion by describing an optimal algorithm, that is, an algorithm maximizing $\mathbb{E}(\text{ALG})$. Note that such an algorithm can be obtained by a straightforward backward induction. There is, however, a simpler description. The algorithm we will describe is a single-threshold algorithm; that is, 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 ALG_T.

Theorem 1. *If the prices are i.i.d. and $T = \mathbb{E}(X_1)$ is the mean of the distribution, then ALG_T is an optimal algorithm, that is, $\mathbb{E}(\text{ALG}_T) \geq \mathbb{E}(\text{ALG})$ for any algorithm ALG.*

Proof. For any period $i \in \{1, \ldots, n\}$, denote by $V_1(i)$ ($V_0(i)$) the optimal expected total profit that can be generated starting from period i with having the item (not having the item, respectively). For convenience, let $V_1(n+1) = V_0(n+1) = 0$. For $i \geq 2$, consider two cases:

- The algorithm has the item in the beginning of period $i-1$. The expected loss after period $i-1$ from selling the item is, by linearity of expectation, $V_1(i) - V_0(i)$. Therefore, the algorithm should sell the item at price X_{i-1} if $X_{i-1} > V_1(i) - V_0(i)$ and should not sell it if $X_{i-1} < V_1(i) - V_0(i)$, with indifference at equality.
- The algorithm does not have the item in the beginning of period $i-1$. The expected profit after period $i-1$ from buying the item is, by linearity of expectation, $V_1(i) - V_0(i)$. Therefore, the algorithm should buy the item at price X_{i-1} if $X_{i-1} < V_1(i) - V_0(i)$ and should not buy it if $X_{i-1} > V_1(i) - V_0(i)$, with indifference at equality.

Therefore, setting the threshold $V_1(i) - V_0(i)$ in period $i-1$ is optimal. Finally, note that, by definition of $V_1(i)$ and $V_0(i)$,

$$V_1(i) = \mathbb{E}(\max\{X_i + V_0(i+1), V_1(i+1)\})$$

and

$$\begin{aligned} V_0(i) &= \mathbb{E}(\max\{V_1(i+1) - X_i, V_0(i+1)\}) \\ &= \mathbb{E}(\max\{V_1(i+1), V_0(i+1) + X_i\}) - \mathbb{E}(X_i). \end{aligned}$$

This implies that, indeed, $V_1(i) - V_0(i) = \mathbb{E}(X_i) = \mathbb{E}(X_1)$. □

Figure 1. (Color online) Example of Decisions Made by OPT: Optimal Buy/Sell Decisions in Hindsight

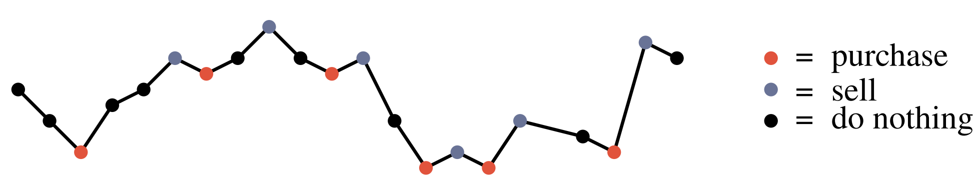

It seems, however, difficult to directly analyze this algorithm with respect to OPT. Therefore, we consider another simple algorithm that achieves a 1/2-approximation and show this factor is tight. This algorithm is again a single-threshold algorithm.

Theorem 2. *If the prices are i.i.d. and T is the median of the distribution, then ALG_T is a 1/2-approximation, that is,* $\mathbb{E}(\mathrm{ALG}_T) \geq (1/2) \cdot \mathbb{E}(\mathrm{OPT})$.

To prove this theorem, we first obtain in Lemma 1 an upper bound on $\mathbb{E}(\mathrm{OPT})$ 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 $T \in \mathbb{R}$,*

$$\begin{aligned}\mathbb{E}(\mathrm{OPT}) &= \frac{n-1}{2} \cdot \mathbb{E}(|X_1 - X_2|)\\ &\leq (n-1) \cdot \mathbb{E}(|X_1 - T|).\end{aligned}$$

Proof. Let us denote by OPT_i the gain of OPT in period i; that is, if OPT buys in period i, then $\mathrm{OPT}_i = -X_i$, and if OPT sells in period i, then $\mathrm{OPT}_i = X_i$ and $\mathrm{OPT}_i = 0$ otherwise. From Observation 1, we have the following:

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

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

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

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 $T \in \mathbb{R}$,

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

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

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

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

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

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

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

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

$$\begin{aligned}\mathbb{E}(\mathrm{ALG}_T(n)) &= \mathbb{E}(X_n) \cdot p_n\\ &= (\mathbb{E}(X_n \cdot \mathbf{1}_{X_n \geq T}) + \mathbb{E}(X_n \cdot \mathbf{1}_{X_n < T})) \cdot p_n.\end{aligned}$$

We show now that if T is the median of the distribution, $p_i = 1/2$ for all $i \geq 2$. Indeed, notice that, conditioned on having the item in period $i-1$ or conditioned on not having it, we have the item in period i if and only if X_{i-1} is below T:

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

Putting all together, we have that

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

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 $\mathbb{E}(T \cdot \mathbf{1}_{X_1 \geq T}) = \mathbb{E}(T \cdot \mathbf{1}_{X_1 < T})$, and therefore, by adding and subtracting this quantity in the last line, we conclude that

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

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:*

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

Proof. Fix n, and we construct a parameterized instance with n i.i.d. prices and obtain the claimed bound as the parameter $\varepsilon \in [0,1]$ tends to zero.

Fix $\varepsilon \in [0,1]$ and consider

$$X_i = \begin{cases} 1 & \text{w.p. } \frac{\varepsilon}{2} \\ \frac{1}{2} & \text{w.p. } 1-\varepsilon \\ 0 & \text{w.p. } \frac{\varepsilon}{2} \end{cases} \quad \text{for } i \in \{1,2,\dots,n\}.$$

To calculate the expected profit of OPT, we use the characterization described in Lemma 1. Here, it gives

$$\begin{aligned}\mathbb{E}(\mathrm{OPT}) &= \frac{n-1}{2}\cdot \mathbb{E}(|X_1 - X_2|) \\ &= \frac{n-1}{2}\cdot\left(4\cdot\frac{1}{2}\cdot(1-\varepsilon)\cdot\frac{\varepsilon}{2} + 2\cdot 1\cdot\left(\frac{\varepsilon}{2}\right)^2\right) \\ &= \frac{n-1}{2}\cdot\left(\varepsilon - \frac{\varepsilon^2}{2}\right).\end{aligned}$$

Consider any algorithm ALG. We can bound its expected profit as follows. First, selling is always optimal in the last period. Because the expected price in this period is $1/2$, we can replace it with a deterministic price of $1/2$, which by linearity of expectation does not change the profit of ALG. Second, any optimal algorithm performs buy/sell operations only when the prices belong to one of the following pairs: $(1/2,1/2)$, $(0,1/2)$, $(0,1)$, or $(1/2,1)$. In all these cases, we can distribute the profit from the buy/sell operations into two parts. The algorithm gains 0 when it trades at a value of $1/2$, and it gains $1/2$ each time it trades at 0 (buy) or 1 (sell). Now, consider period t, and let $q(t)$ represent the probability that ALG holds an item before observing the price in this period. Because the price at period t is independent of $q(t)$, using the above reasoning, for $t \le n-1$, we can write that

$$\begin{aligned}&\mathbb{E}(\text{profit of ALG at period } t) \\ &\le q(t)\frac{1}{2}\varepsilon + (1-q(t))\frac{1}{2}\varepsilon = \frac{1}{4}\varepsilon.\end{aligned}$$

By the linearity of expectation, we conclude that

$$\mathbb{E}(\mathrm{ALG}) \le \frac{(n-1)}{4}\cdot\varepsilon.$$

It follows that

$$\frac{\mathbb{E}(\mathrm{ALG})}{\mathbb{E}(\mathrm{OPT})} \le \frac{\frac{n-1}{4}\cdot\varepsilon}{\frac{n-1}{2}\left(\varepsilon - \frac{\varepsilon^2}{2}\right)}.$$

Sending ε to zero, we derive

$$\lim_{\varepsilon\to 0}\frac{\mathbb{E}(\mathrm{ALG})}{\mathbb{E}(\mathrm{OPT})} \le \lim_{\varepsilon\to 0}\frac{\frac{n-1}{4}\cdot\varepsilon}{\frac{n-1}{2}\left(\varepsilon - \frac{\varepsilon^2}{2}\right)} = \frac{1}{2}. \quad \square$$

4. Independent 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 $n\to\infty$. 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 online appendix.

Theorem 3. *If the prices are presented in order $X_{\sigma(1)},\dots,X_{\sigma(n)}$, where $\sigma:\{1,\dots,n\}\to\{1,\dots,n\}$ is a uniform random permutation, then there exist a threshold T such that*

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

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

Theorem 4. *If the prices are presented in order $X_{\sigma(1)},\dots,X_{\sigma(n)}$, where $\sigma:\{1,\dots,n\}\to\{1,\dots,n\}$ is a uniform random permutation, then there exist a threshold T and a constant C, independent of these distributions, such that*

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

We provide the proofs of Theorems 3 and 4 in Sections 4.1, 4.2, and 4.3.

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,*

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

Proof. Consider a large constant $M > 0$. We define an instance with $n = 2$. Take X_1 such that $X_1 = M+2$ w.p. $M/(M+2)$, and $X_1 = 0$ w.p. $2/(M+2)$. Take X_2 such that $X_2 = M$ w.p. $M/(M+2)$ and $X_2 = 2M+2$ w.p. $2/(M+2)$.

Notice that $\mathbb{E}(X_1) = M$ and $\mathbb{E}(X_2) = M+2$. The expectation of the optimal algorithm can be written as

$$\begin{aligned}&\mathbb{E}(\mathrm{ALG}) \\ &= \frac{1}{2}\cdot\mathbb{E}([\mathbb{E}(X_2) - X_1]_+) + \frac{1}{2}\cdot\mathbb{E}([\mathbb{E}(X_1) - X_2]_+) \\ &= \frac{1}{2}\cdot(M+2)\cdot\frac{2}{M+2} + 0 \quad = 1.\end{aligned}$$

Let us calculate $\mathbb{E}(\mathrm{OPT})$. There are four possible sequences of prices where OPT buys and sells: $(M+2, 2M+2)$, $(0,M)$, $(0,2M+2)$, $(M,M+2)$. In all other cases, OPT does nothing. Therefore,

$$\begin{aligned}
&\mathbb{E}(\mathrm{OPT})\\
&=\frac{1}{2}\cdot\Bigg(M\cdot\frac{M}{M+2}\cdot\frac{2}{M+2}+M\cdot\frac{2}{M+2}\cdot\frac{M}{M+2}\\
&\quad+(2M+2)\cdot\left(\frac{2}{M+2}\right)^2+2\cdot\left(\frac{M}{M+2}\right)^2\Bigg)\\
&=3+O(1/M).
\end{aligned}$$

This completes the proof. □

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. The first lemma is analogous to Lemma 1 from the i.i.d. case.

Lemma 3. *If the prices are presented in uniformly random order σ, then*

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

Proof. For a fixed order σ 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

$$\begin{aligned}
\mathrm{OPT}&=\sum_{i=2}^{n}X_{\sigma(i)}\cdot\mathbf{1}_{X_{\sigma(i)}\ge X_{\sigma(i-1)}}\\
&\quad-\sum_{i=1}^{n-1}X_{\sigma(i)}\cdot\mathbf{1}_{X_{\sigma(i)}\le X_{\sigma(i+1)}}\\
&=\sum_{i=1}^{n-1}(X_{\sigma(i+1)}-X_{\sigma(i)})\cdot\mathbf{1}_{X_{\sigma(i+1)}\ge X_{\sigma(i)}}.
\end{aligned}$$

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

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

where the first equality holds by linearity of expectation and the second and third equality hold because σ is a uniformly random order. □

The second lemma, generalizing Lemma 2 from the i.i.d. case, is more complicated than that lemma because it involves the indicator variable of T being between the two variables from the two-variable problem we are reducing to.

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

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

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

Consider now any period $2\le i\le n-1$. In such a period, the algorithm buys whenever $X_{\sigma(i)}<T$ and it does not have the item. The latter event is equivalent to the event $X_{\sigma(i-1)}\ge T$. Indeed, assuming that $X_{\sigma(i-1)}\ge T$, either 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 $2\le i\le n-1$, the algorithm sells the item iff $X_{\sigma(i)}\ge T$, 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 ALG_T in period $2\le i\le n-1$ by

$$\begin{aligned}
\mathrm{ALG}_T(i)=&-X_{\sigma(i)}\cdot\mathbf{1}_{X_{\sigma(i)}<T}\cdot\mathbf{1}_{X_{\sigma(i-1)}\ge T}\\
&+X_{\sigma(i)}\cdot\mathbf{1}_{X_{\sigma(i)}\ge T}\cdot\mathbf{1}_{X_{\sigma(i-1)}<T}.
\end{aligned}$$

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 $\mathrm{ALG}_T(n)=X_{\sigma(n)}\cdot\mathbf{1}_{X_{\sigma(n-1)}<T}$. Putting all together, we obtain

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

where the last equality follows from linearity of expectation and the fact that σ is a uniformly random permutation. By manipulating the last two terms of the above sum, we obtain

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

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

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

where the last lines follows from the fact that σ is a uniformly random permutation. □

4.2. Proof of Theorem 3

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

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

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*

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

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

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

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 $M_2 \in [\min(x_1', x_2'), \max(x_1', x_2')]$, and

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. Because of the presence of long calculations in the formal proof, we refer the reader to the online appendix for details. In there, 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.

Proof of Theorem 3. 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 $H_1 \subseteq \{1, \ldots, n\}$ be a set chosen uniformly at random from all subsets of size $\lceil n/2 \rceil$. Let a, b be random elements chosen uniformly from H_1 and $\{1, \ldots, n\} \setminus H_1$. For any two $i \neq j, i, j \in \{1, \ldots, n\}$, 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

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

Because choice of each subset of size $\lceil n/2 \rceil$ of set $\{1, \ldots, n\}$ 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, whereas b' is a uniformly random element from $\{1, \ldots, n\} \setminus S$. Note here that $X_{a'}$ and $X_{b'}$ are independent random variables because the sets of random variables $\{X_i | i \in S\}$ and $\{X_i | i \in \{1, \ldots, n\} \setminus S\}$ are pair-wise independent. Therefore, we can apply Lemma 5 and

Figure 2. (Color online) Different Cases Distinguished in the Proof of Lemma 5

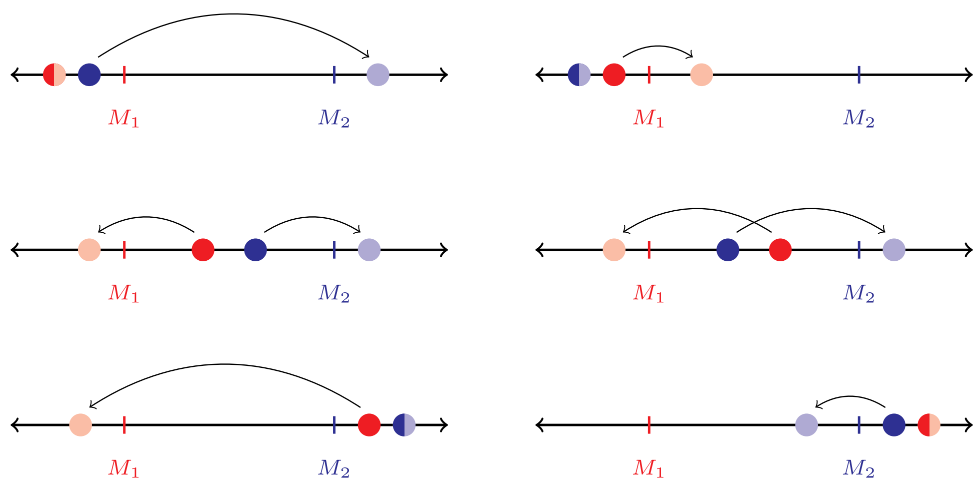

Notes. Realizations of X_1 and X_2 are shown in darker colors, respectively. The realizations that are charged to are depicted in lighter colors. Although the specific values depend on the actual distributions, what is unaffected is the order of the realizations with respect to the median of the corresponding distribution, as well as among each other.

get a threshold $T \in \mathbb{R}$ such that

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

This together with Inequality (3) yields

$$\begin{aligned}&\mathbb{E}(|X_{a'} - X_{b'}| \cdot \mathbf{1}_{T\in[\min(X_{a'},X_{b'}),\max(X_{a'},X_{b'})]})\\&\quad\geq \frac{1}{4}\mathbb{E}(|X_{\sigma(1)} - X_{\sigma(2)}|). \qquad (4)\end{aligned}$$

Finally, observe that

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

where the inequality above follows from the following reasons. First, all random variables on both sides of the inequality are nonnegative. Second, on the left-hand side, the sum iterates over all possible distinct pairs i,j, whereas 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 $4/(n(n\text{-}1)) = (2/n)\cdot(2/(n\text{-}1)) \geq (1/\lceil n/2\rceil)\cdot(1/(n-\lceil n/2\rceil))$ for all integer n. This, together with Inequality (4), proves the theorem. □

4.3. Proof of Theorem 4

Let us now turn to proving Theorem 4. We will be analyzing the fraction

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

where T is the median of the mixture random variable, i.e., the median of X_Y where Y is a uniformly random index in $[n]$. 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 with 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 one to n. Then, the variables X_Y and $X'_{Y'}$ are also independent and have the same distribution.

We define

$$\begin{aligned}&S_1 := \mathbb{E}(|X_{\sigma(1)} - X_{\sigma(2)}|),\\&S_2 := \mathbb{E}(|X_{\sigma(1)} - X_{\sigma(2)}| \cdot \mathbf{1}_{T\in[\min(X_{\sigma(1)},X_{\sigma(2)}),\max(X_{\sigma(1)},X_{\sigma(2)})]}),\end{aligned}$$

as well as

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

Because 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

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

On the other hand, it holds that

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

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

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

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

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

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 $B := \{k \mid \mathbb{P}(X'_k \leq T) \geq \frac{1}{4}\}$ has size at least $n/4$. Consider now the sum $\sum_{j\in[n], j\neq i}\mathbb{E}(|X_i - X'_j|\cdot \mathbf{1}_{T\in[X_i, X'_j]}$. We have that

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

where the last inequality follows from the fact that $X_a > T$ and $X_b \leq T$. Because 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

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

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

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

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

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

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

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

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

$$\begin{aligned}
&2\mathbb{E}((X_i - X_i')\mathbf{1}_{X_i \ge T} \cdot \mathbf{1}_{X_i' < T}) \\
&\quad \le 2(\mathbb{E}((T - X_i')\mathbf{1}_{X_i \ge T} \cdot \mathbf{1}_{X_i' < T}) \\
&\qquad + \mathbb{E}((X_i - T) \cdot \mathbf{1}_{X_i \ge T} \cdot \mathbf{1}_{X_i' < T})) \\
&\quad \le 2(\mathbb{E}((T - X_i') \cdot \mathbf{1}_{X_i' < T}) + \mathbb{E}((X_i - T) \cdot \mathbf{1}_{X_i \ge T})),
\end{aligned}$$

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

$$\begin{aligned}
&\sum_{j \in [n], j \ne i} 2\mathbb{E}(|X_i - X_j'| \cdot \mathbf{1}_{T \in [X_i, X_j']}) \\
&\ge \left(\frac{n}{16} - \frac{1}{4}\right) \cdot (\mathbb{E}((T - X_i) \cdot \mathbf{1}_{X_i \le T}) + \mathbb{E}((X_i - T) \cdot \mathbf{1}_{X_i > T})) \\
&\ge \frac{1}{2} \cdot \left(\frac{n}{16} - \frac{1}{4}\right) \cdot \mathbb{E}(|X_i - X_i'| \cdot \mathbf{1}_{T \in [X_i, X_i']}),
\end{aligned}$$

which is equivalent to

$$\frac{64}{n} \cdot \mathbb{E}(|X_i - X_j'| \cdot \mathbf{1}_{T \in [X_i, X_j']}) \ge \mathbb{E}(|X_i - X_i'| \cdot \mathbf{1}_{T \in [X_i, X_i']}).$$

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

We are now ready to finish the proof of Theorem 4.

Proof of Theorem 4. Recall that our goal is to prove

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

To do so, first observe that

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

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 $2 \ge \frac{S_1'}{S_2'}$, which combined with the above equality, gets

$$2 \ge \frac{S_1 + \frac{1}{n(n-1)} \cdot \sum_{i \in [n]} \mathbb{E}(|X_i - X_i'|)}{S_2'} \ge \frac{S_1}{S_2'}, \tag{9}$$

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

5. Generalizations

In the following sections, 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 reinvestments of gains, and a multiarmed 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 $p_j = x_j + y$, $y \sim G$, 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 5. *If prices are i.i.d. or affiliated, but we do not know the distribution, there is an algorithm ALG such that*

$$\mathbb{E}(\text{ALG}) \ge \frac{1}{2} \cdot \frac{n-2}{n-1} \cdot \mathbb{E}(\text{OPT}).$$

A natural approach to prove such a result would be to use the empirical median of past prices. However, proving a guarantee for that algorithm would require analyzing the distribution of the empirical median (similar to the analysis of Correa et al. (2022b) and Rubinstein et al. (2020)). If M_t denotes the median of the first t prices, we would need to analyze how the terms $\mathbb{E}([X_{t+1} - M_t]_+)$ and $\mathbb{E}([M_t - X_{t+1}]_+)$ compare with $\mathbb{E}([X_1 - X_2]_+)$. Instead of this, we use as a threshold an independently drawn sample from the distribution of prices, which results in a straightforward proof, and then we show how to implement this only using past prices.

Assume we have access to $n-1$ independent samples $S_1, S_2, \dots, S_{n-1}$. For this case, consider the following algorithm, which we denote by 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 $X_i \ge S_i$. 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, that is, 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. *We have that* $\mathbb{E}(\text{ALG}^s) \ge \frac{n-1}{4} \cdot \mathbb{E}(|X_1 - X_2|)$.

Proof. Denote by $\text{ALG}^s(i)$ the gains of the algorithm in period i. Note first that in every period $i \ge 2$, the event that we have the item is independent of X_i and S_i and has probability exactly $1/2$. This is because this event is exactly the event $\{X_{i-1} < S_{i-1}\}$. Therefore,

$$\mathbb{E}(\mathrm{ALG}^s(i)) = \begin{cases} \mathbb{E}(-X_1 \cdot \mathbf{1}_{X_1 < S_1}) & i = 1 \\ \frac{1}{2} \cdot \mathbb{E}(X_n) & i = n \\ \frac{1}{2} \cdot (\mathbb{E}(X_i \cdot \mathbf{1}_{X_i \ge S_i}) + \mathbb{E}(-X_i \cdot \mathbf{1}_{X_i < S_i})) & \text{else.} \end{cases}$$

Because $X_1, \dots, X_n, S_1, \dots S_{n-1}$ are i.i.d. realizations of the same distribution, we conclude that

$$\begin{aligned}\mathbb{E}(\mathrm{ALG}^s) &= \sum_{i=1}^{n} \mathbb{E}(\mathrm{ALG}^s(i)) \\ &= \frac{n-1}{2} \cdot \mathbb{E}((X_1 - X_2) \cdot \mathbf{1}_{X_1 \ge X_2}) \\ &= \frac{n-1}{4} \cdot \mathbb{E}(|X_1 - X_2|).\end{aligned}$$

For the case of affiliated random variables, that is, if $X_j = x_j + y$ and $S_j = s_j + y$ for $y \sim G$ and $x_j \sim s_j \sim F$, notice that the events $\{X_{i-1} < S_{i-1}\}$ and $\{x_{i-1} < s_{i-1}\}$ are equivalent. This means that the event that we have the item in period i is independent of X_i and S_i, so all previous equations still hold. □

Combining this lemma with Lemma 1, we have that ALG^s is actually a 1/2 approximation. From ALG^s, we derive ALG^{s*}, an algorithm that uses a single sample S_1 of the distribution. For $2 \le i \le n-1$, we define S_i' by selecting a uniformly random element from $\{S_1, X_1, \dots, X_{i-1}\}$. Now, ALG^{s*} behaves like ALG^s, taking $S_1, S_2', \dots, S_{n-1}'$ as the set of samples.

Lemma 8. *We have that* $\mathbb{E}(\mathrm{ALG}^{s*}) = \mathbb{E}(\mathrm{ALG}^s)$.

Proof. We simply prove that the expected gain of the two algorithms is the same in each period, that is, $\mathbb{E}(\mathrm{ALG}^{s*}(i)) = \mathbb{E}(\mathrm{ALG}^s(i))$. This is easy to see for period $i = 1$. For $i > 1$, note first that S_i' and X_i are i.i.d. random variables. This is because $\{S_1, X_1, \dots, X_{i-1}\}$ is independent of X_i.

The critical step is to show that the event that we have the item in period i is independent of the pair (X_i, S_i'). Recall that this is equivalent to the event $\{X_{i-1} < S_{i-1}'\}$, so conditioning on this event might affect the distribution of S_i'. This in fact is not the case, because it refers to the relative order within the set $\{S_1, X_1, \dots, X_{i-1}\}$, which is independent of the value of a uniformly random element of the same set (given that they are i.i.d. random variables). Notice that if we now add the same number $y \sim G$ to all random variables, the relative orders do not change, so the argument also applies to affiliated prices. □

Now, Theorem 5 is a direct application of Lemma 8. If we do not know the distribution, we can skip the first period and use 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 $(1/2) \cdot ((n-2)/(n-1)) \cdot \mathbb{E}(\mathrm{OPT})$.

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. For a formal proof of this result, we refer the reader to the online appendix.

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

5.3. Budgeted Version with Fractional Purchase and Reinvestment 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 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 $S_i \ge 0$ is the (possibly fractional) number of stocks we hold in period i, and $B_i \ge 0$ is our budget in period i. We start with $(S_1, B_1) = (0, 1)$, meaning that we start with zero stocks and a budget of one 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, that is, B_{n+1}.

Denote by OPT_F the gains of the prophet in this model, that is, the maximum possible profit in hindsight. It is easy to see that the prophet either sells all stocks or spends all money, that is, $s \in \{-S_i, B_i/X_{\sigma(i)}\}$. Thus, if I denotes the set of periods where the prophet buys, and J denotes the periods where the prophet sells, we have that $\mathrm{OPT}_F = \prod_{j \in J} X_{\sigma(j)} / \prod_{i \in I} X_{\sigma(i)}$. 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

$$\begin{aligned}\mathbb{E}(\mathrm{OPT}_F) &\ge \exp(\mathbb{E}(\log \mathrm{OPT}_F)) \\ &= \exp\left(\frac{n-1}{2} \cdot \mathbb{E}(|\log X_{\sigma(1)} - \log X_{\sigma(2)}|)\right).\end{aligned}$$

It is clear from this that there is no hope for a constant approximation against $\mathbb{E}(\mathrm{OPT}_F)$. Instead, we can approximate the expected growth rate of OPT_F, that is, $\mathbb{E}(\log \mathrm{OPT}_F)$. 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 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* $\mathbb{E}(\log \mathrm{ALG}_T) \geq \frac{1}{2}\mathbb{E}(\log \mathrm{OPT}_F)$.

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

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* $\mathbb{E}(\log \mathrm{ALG}_T) \geq \frac{1}{2\cdot(1+C/n)}\mathbb{E}(\log \mathrm{OPT}_F)$.

5.4. Multiarmed 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 the jth price of the ith item. On period I, we observe k different prices $X^1_{\sigma^1(i)}, \dots, X^k_{\sigma^k(i)}$ where $\sigma^j : \{1, \dots, n\} \to \{1, \dots, n\}$ 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 $X^j_{\sigma^j(i)}$, 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 $X^j_{\sigma^j(i)}$. 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 (this makes no difference in the original model because it would be equivalent to keep the item until the next period). 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 Theorems 3 and 4 with $\frac{1}{k}$ multiplicative loss, against the prophet, The key observation involves understanding the strategy of the prophet in the multiarmed bandit version.

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

$$\mathbb{E}(\mathrm{OPT}) = (n-1)\cdot \mathbb{E}\left(\max_{i\in[k]} [X^i_{\sigma^i(1)} - X^i_{\sigma^i(2)}]_+\right),$$

where $[\cdot]_+ := \max\{\cdot, 0\}$.

Proof. Consider a difference in prices between period i and $i+1$ for $1 \leq i \leq n-1$. On one hand, the prophet can gain at most $\max_{j\in[k]}((X^j_{\sigma^j(i+1)} - X^j_{\sigma^j(i)})\mathbf{1}_{X^j_{\sigma^j(i+1)} > X^j_{\sigma^j(i)}})$, because he can possess only one item during this period. On the other hand, he can gain exactly this value, because performing this transaction has no effects on next periods as he starts the next period without any item. Thus, we have that

$$\mathbb{E}(\mathrm{OPT}) = \mathbb{E}\left(\sum_{i=1}^{n-1} \max_{j\in[k]} [X^j_{\sigma^j(i+1)} - X^j_{\sigma^j(i)}]_+\right).$$

Using linearity of the expectation and the fact that the permutations σ^j are independent and uniformly random, we conclude the formula in the statement of the lemma. □

This result allows us to immediately extend Theorems 3 and 4.

Theorem 6. *There exist algorithms* ALG_1 *and* ALG_2 *that achieve the following approximation for k-armed bandit version of trading prophets:*

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

Proof. Let ALG_1 be an algorithm that first picks one of the k items uniformly and randomly and then applies the strategy from Theorem 3 to prices of the chosen item:

$$\mathbb{E}(\mathrm{ALG}_1) = \sum_{i=1}^{k} \frac{1}{k}\cdot \mathbb{E}(\mathrm{ALG}_1^i),$$

where $\mathbb{E}(\mathrm{ALG}_1^i)$ denotes the expected gain of algorithm given by Theorem 3 applied to prices of the ith item. By Lemma 4, we have that

$$\mathbb{E}(\mathrm{ALG}_1^i) \geq \frac{1}{16}\cdot \mathbb{E}(\mathrm{OPT}^i), \tag{10}$$

where OPT^i denotes the expectation of the optimal strategy in hindsight applied to prices of the ith item. By Lemma 3, we get

$$\begin{aligned}\frac{1}{16}\cdot \mathbb{E}(\mathrm{OPT}^i) &= \frac{1}{16}\frac{n-1}{2}\cdot \mathbb{E}(|X^i_{\sigma^i(1)} - X^i_{\sigma^i(2)}|)\\ &= \frac{n-1}{16}\cdot \mathbb{E}([X^i_{\sigma^i(1)} - X^i_{\sigma^i(2)}]_+).\end{aligned}$$

The above inequality plugged into Inequality (10) yields

$$\begin{aligned}\mathbb{E}(\mathrm{ALG}_1) &= \sum_{i=1}^{k}\frac{1}{k}\cdot\mathbb{E}(\mathrm{ALG}_1^i)\\ &\geq \frac{1}{16k}\cdot\sum_{i=1}^{k}(n-1)\cdot\mathbb{E}([X^i_{\sigma^i(1)}-X^i_{\sigma^i(2)}]_+).\end{aligned}$$

Observe that all random variables inside expectations are nonnegative. Thus,

$$\begin{aligned}&\sum_{i=1}^{k}(n-1)\cdot\mathbb{E}([X^i_{\sigma^i(1)}-X^i_{\sigma^i(2)}]_+)\\ &\quad\geq (n-1)\cdot\mathbb{E}\left(\max_{i\in[k]}[X^i_{\sigma^i(1)}-X^i_{\sigma^i(2)}]_+\right)\\ &\quad= \mathbb{E}(\mathrm{OPT}),\end{aligned}$$

where the last equality holds by Lemma 10. Therefore, the first part of the theorem is proven.

To prove the second part, it suffices to take an algorithm that settles the threshold with respect to Theorem 4 instead of Theorem 3. Because Lemmas 3 and 10 hold irrespective of the threshold chosen by the algorithm, other parts of the above reasoning are valid and prove that ALG_2 has an approximation factor of $1/(2k)-o(1)$. □

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

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

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

Proof. Consider a random variable X that takes zero with probability $1-1/k$ and k with probability $1/k$. Assume that $n=2$ and each random variable X^i_j for $1\leq i\leq k$, $1\leq j\leq n$ is drawn independently according to the distribution X. Because there are only two periods, there can be at most one transaction. Obviously, the prophet trades only if the price on the first period is zero and the price of the same item on the second period is k. For a fixed item, the probability of such prices realization is $(1-1/k)/k$. Because we have k different items whose prices drawn independently, we see that the probability the sequence of prices $(0,k)$ does not happen in any of these price realizations are

$$\begin{aligned}\left(1-\left(1-\frac{1}{k}\right)\frac{1}{k}\right)^k &= \left(1-\frac{1}{k}+\frac{1}{k^2}\right)^k\\ &\leq \left(1-\frac{1}{2k}\right)^k \leq \sqrt{\frac{1}{e}}.\end{aligned}$$

Therefore, we have

$$\mathbb{E}(\mathrm{OPT}) \geq (1-e^{-1/2})\cdot k.$$

On the other hand, because any algorithm can make only one transaction, it always buys an item that have price 0 in the first period. Because prices presented in the second period are independent from choices of the algorithm, it does not matter which item the algorithm selects among these with price 0. Therefore, we have that

$$\mathbb{E}(\mathrm{ALG}) \leq k\cdot\frac{1}{k}=1,$$

which, if compared with $\mathbb{E}(\mathrm{OPT})$, proves the lemma. □

6. Experiments

In this section, we present some numerical experiments to validate our results, using a data set containing used car auction prices from the United States and Canada.[1] We then numerically investigate how robust our results are to dropping the i.i.d. condition.

6.1. Used Cars Data Set

We present the results for auction prices of a specific car model in the data set, but other models give similar results. We report results for a sequence of 936 prices from transactions between December 2014 and June 2015 for the car model Kia Sorento 2015 (Figure 3).

To evaluate our mean-and median-threshold algorithms, we took a sliding window of w consecutive prices and computed the algorithms' performances and that of the prophet on each time window. We tested the algorithms using the empirical mean and median within the window, which is impossible in practice but should be closer to the "true" mean and median, and

Figure 3. (Color online) Sequence of Auction Prices for Kia Sorento 2015

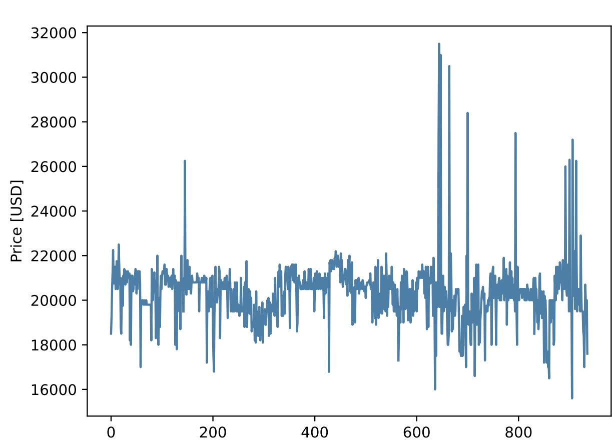

using the empirical mean and median of all past prices (before the start of the window).

That is, assuming n data points, we looked at the windows comprising the price points at indices $\{j, j+1, \ldots, j+w-1\}$ for $1 \le j \le n-w+1$. We refer to the prices at indices $\{j, j+1, \ldots, j+w-1\}$ as the jth time window. Then, to evaluate the algorithm's performance on the jth time window, in the first version, we used the mean and median on the jth time window itself, whereas in the second version, we used the mean and median on $\{1, \ldots, j-1\}$.

Figure 4 shows the results for a window size of $w = 300$. We observe that, for both approaches of computing the empirical mean and median, the algorithms have a performance close to the $1/2$-approximation predicted by our theoretical guarantee for i.i.d. prices. Interestingly, this is also true for the variant that uses past price points for early time windows, where the mean and median are based only on a single (or handful) of past price points.

Figure 5 gives the relative performance of the median algorithm for varying window sizes of $w = 50, w = 100, w = 200$, and $w = 300$. On average over the starting point of the window, we observe a similar performance guarantee for all window sizes, but smaller window sizes seem to lead to more variance in the performance.

We also tested the sample-based algorithm for the unknown-distribution case from Section 5.1. We plot the results in Figure 6, which consist of applying the algorithm in each window, using only observed prices within the window. The obtained performance is much noisier than that of the mean and median algorithms but is also close to the theoretical guarantee.

Figure 4. (Color online) Performance of the Mean and Median Algorithms on Kia Sorento 2015 Prices

Note. Each point in the plots is the total profit within a window of 300 consecutive prices.

Figure 5. (Color online) Performance of the Empirical-Median Algorithm (for Prices Before the Window) with Varying Time-Window Sizes

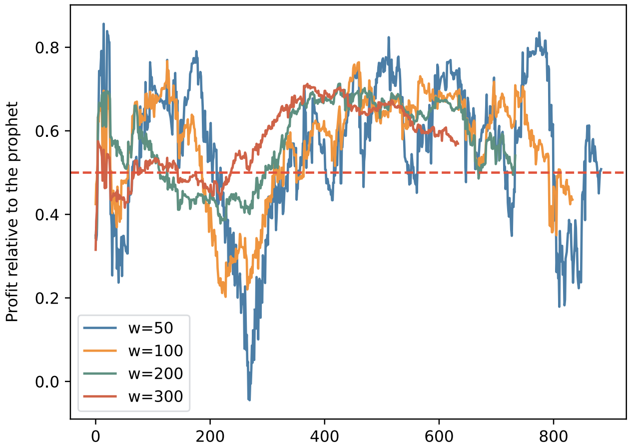

Note. The graphs for the longer time windows end earlier because the value of the x axis is the starting point of the window.

From Figure 3, we can see that the prices in our experiments with used car prices are clearly not i.i.d., so our experimental findings point to a certain robustness of our theoretical guarantees. Another simplification of our model, compared with the actual application, is that we do not model transaction costs. However, because the profit of our algorithms grows linearly with a coefficient comparable to the average difference between two consecutive prices, if the transaction cost is small compared with that average difference, then the transaction cost will not significantly affect the results.

6.2. Robustness on Synthetic Data

As observed earlier, no algorithm can extract a positive profit from prices that follow a martingale. To test the robustness of our algorithm to non-i.i.d. data, we apply it on a sequence given by a martingale plus i.i.d. noise, an interpolation between our i.i.d. setting and a worst-case setting. This is a reasonable model of markets for goods with high friction, where prices follow a certain underlying trend but are significantly affected by the valuation for the good of each particular buyer and seller and generalize the model of affiliated prices.

We generated 200 sequences of prices over a horizon of 800 periods given by a random walk plus noise. We used i.i.d. increments for the random walk, drawn from a normal distribution with mean 0 and standard deviation 0.15, and i.i.d. noise with mean 0 and standard deviation 0.3. We set the starting price to be 100 in all sequences (Figure 7(a)).

We apply a median algorithm that at each period sets as a threshold the median of the last five prices. As one would expect, this algorithm is sensitive to the length of the window over which we take the median, because prices that are further away in the past deviate too much from the current value of the underlying random walk. As we observe in Figure 7(b), the algorithm performs surprisingly well and is very close to a $1/2$ fraction of the offline optimum.

We then analyze what happens if we subsample the data, meaning that we only observe one in every s prices, for s between one and five. This models a situation where we cannot trade as fast as necessary if we want to take advantage of the i.i.d. noise. Mathematically, increasing s has the effect of increasing the

Figure 6. (Color online) Performance of the Sample-Based Algorithm on Kia Sorento 2015 Prices

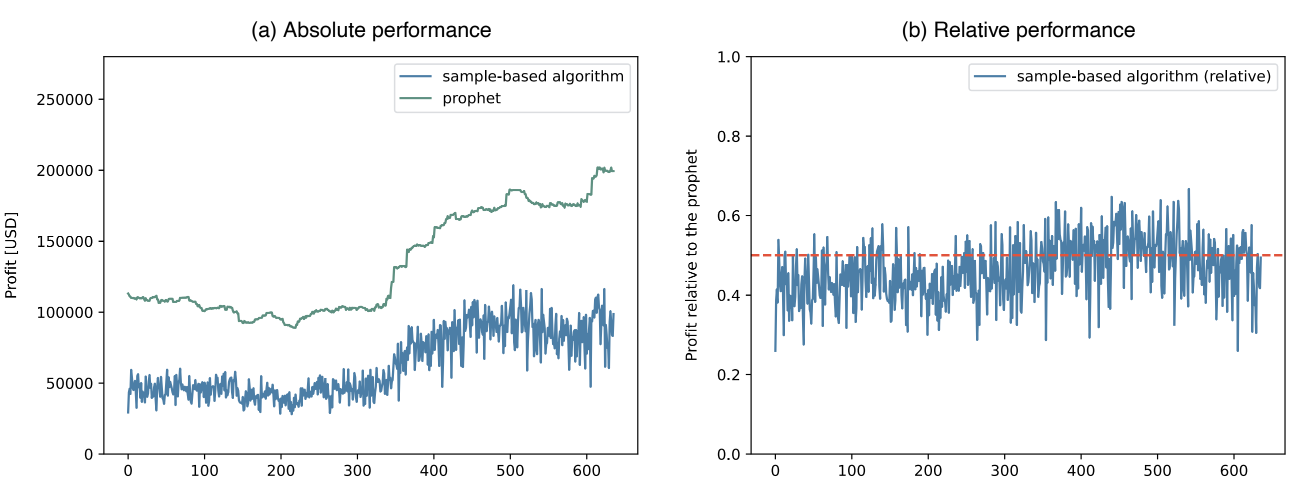

Note. Each point in the plots is the total profit within a window of 300 consecutive prices.

Operations Research, 2026, vol. 74, no. 1, pp. 260–280, © 2025 INFORMS

Figure 7. (Color online) Experiment with Synthetic Data with Gaussian i.i.d. Increments Plus Gaussian i.i.d. Noise

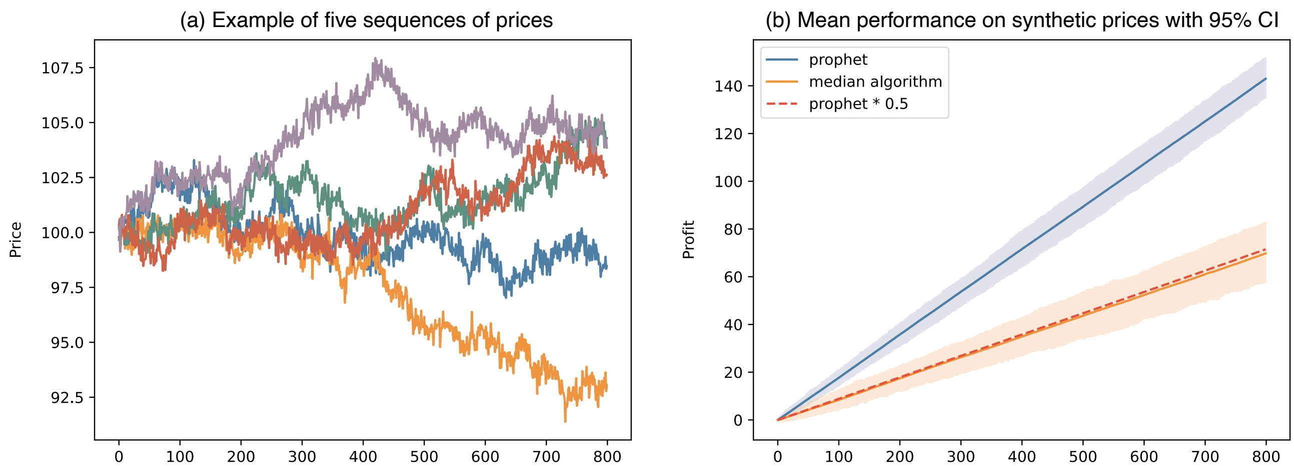

variance of the increments of the underlying random walk while keeping the size of the noise equal, thus reducing its relative size. Although the profit of both the prophet and the algorithm decreases (Figure 8(a)), as we observe in Figure 8(b), the performance of the online algorithm decreases faster and therefore its relative performance also decreases.

7. Future Work

An important avenue for future work is to identify additional practical settings in which large approximation ratios are achievable. The main priority would be to drop the i.i.d. assumption, which is too restrictive for many practical settings. We are interested in developing techniques for beyond i.i.d. prices and understanding how far one can relax the independence assumption while keeping good approximation ratios. This appears to be a technically challenging question, but our observations in the experiments with synthetic data suggest that there indeed are non-i.i.d. settings where one can obtain nontrivial guarantees.

Within the i.i.d. and the random order models, an interesting direction is to understand the case where we can hold up to $k > 1$ items, but we can trade only one item in each period. Our techniques do not extend to this version, but we believe similar guarantees should be possible. A second important direction is to investigate what is possible when we have transaction costs of considerable size compared with the average per-period profit of the offline optimum.

Figure 8. (Color online) Median Algorithm on Synthetic Data with Varying Subsampling Parameter

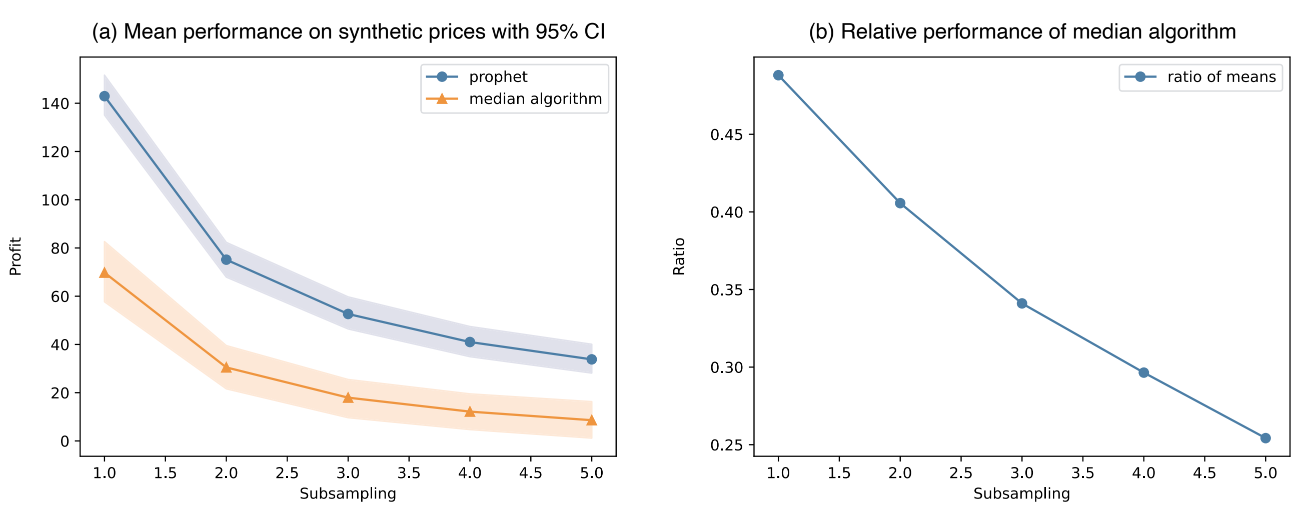

Note. A subsampling parameter of s indicates that we observe the price in one in every s periods.

Acknowledgments

This work was initiated during the Fall 2021 Virtual Chair Semester on Prophet Inequalities (Correa et al. 2022b). The authors thank the organizers and participants of the semester for feedback. An extended abstract of this work appeared in the *Proceedings of the 24th ACM Conference on Economics and Computation*.

Endnote

[1] The data set is publicly available at https://www.kaggle.com/datasets/tunguz/used-car-auction-prices.

References

Abolhassani M, Ehsani S, Esfandiari H, Hajiaghayi M, Kleinberg RD, Lucier B (2017) Beating $1 - 1/e$ for ordered prophets. *Proc. ACM SIGACT Sympos. Theory Comput.* (ACM, New York), 61–71.

Alaei S (2014) Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM J. Comput.* 43(2):930–972.

Borodin A, El-Yaniv R (1998) *Online Computation and Competitive Analysis* (Cambridge University Press, Cambridge, UK).

Braun A, Kesselheim T (2021) Truthful mechanisms for two-sided markets via prophet inequalities. *Proc. ACM Conf. Econom. Comput.* (ACM, New York), 202–203.

Brustle J, Cai Y, Wu F, Zhao M (2017) Approximating gains from trade in two-sided markets via simple mechanisms. *Proc. ACM Conf. Econom. Comput.* (ACM, New York), 589–590.

Caramanis C, Dütting P, Faw M, Fusco F, Lazos P, Leonardi S, Papadigenopoulos O, et al. (2022) Single-sample prophet inequalities via greedy-ordered selection. *Proc. ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 1298–1325.

Charnes A, Drèze J, Miller M (1966) Decision and horizon rules for stochastic planning problems: A linear example. *Econometrica* 34(2):307–330.

Colini-Baldeschi R, de Keijzer B, Leonardi S, Turchetta S (2016) Approximately efficient double auctions with strong budget balance. *Proc. ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 1424–1443.

Colini-Baldeschi R, Goldberg PW, de Keijzer B, Leonardi S, Turchetta S (2017a) Fixed price approximability of the optimal gain from trade. *Proc. Internat. Conf. Web Internet Econom.* (Springer, Heidelberg, Berlin), 146–160.

Colini-Baldeschi R, Goldberg PW, de Keijzer B, Leonardi S, Roughgarden T, Turchetta S (2017b) Approximately efficient two-sided combinatorial auctions. *Proc. ACM Conf. Econom. Comput.* (ACM, New York), 591–608.

Correa JR, Hartline J, Immorlica N (2022a) A semester virtual institute. Accessed October 10, 2025, https://cacm.acm.org/blogs/blog-cacm/258538-a-semester-virtual-institute/fulltext.

Correa JR, Saona R, Ziliotto B (2021a) Prophet secretary through blind strategies. *Math. Programming* 190(1):483–521.

Correa JR, Cristi A, Epstein B, Soto JA (2020) The two-sided game of googol and sample-based prophet inequalities. *Proc. ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 2066–2081.

Correa J, Dütting P, Fischer FA, Schewior K (2022b) Prophet inequalities for independent and identically distributed random variables from an unknown distribution. *Math. Oper. Res.* 47(2):1287–1309.

Correa JR, Dütting P, Fischer FA, Schewior K, Ziliotto B (2021b) Unknown I.I.D. prophets: Better bounds, streaming algorithms and a new impossibility. *Proc. Innovations Theoretical Comput. Sci. Conf.* (Schloss Dagstuhl - Leibniz Zentrum für Informatik, Wadern, Germany), 86:1–86:1.

Correa JR, Foncea P, Hoeksma R, Oosterwijk T, Vredeveld T (2021c) Posted price mechanisms and optimal threshold strategies for random arrivals. *Math. Oper. Res.* 46(4):1452–1478.

Deng Y, Mao J, Sivan B, Wang K (2022) Approximately efficient bilateral trade. *Proc. ACM SIGACT Sympos. Theory Comput.* (ACM, New York), 718–721.

Doob JL (1953) *Stochastic Processes* (Wiley, Hoboken, NJ).

Du Toit J, Peskir G (2007) The trap of complacency in predicting the maximum. *Ann. Probability* 35(1):340–365.

Du Toit J, Peskir G (2009) Selling a stock at the ultimate maximum. *Ann. Appl. Probability* 19(3):983–1014.

Dütting P, Fusco F, Lazos P, Leonardi S, Reiffenhäuser R (2021) Efficient two-sided markets with limited information. *Proc. ACM SIGACT Sympos. Theory Comput.* (ACM, New York), 1452–1465.

Ehsani S, Hajiaghayi M, Kesselheim T, Singla S (2018) Prophet secretary for combinatorial auctions and matroids. *Proc. ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 700–714.

Ekbatani F, Niazadeh R, Nuti P, Vondrák J (2024) Prophet inequalities with cancellation costs. *Proc. ACM Sympos. Theory Comput.* (ACM, New York), 1247–1258.

Esfandiari H, Hajiaghayi M, Liaghat V, Monemizadeh M (2017) Prophet secretary. *SIAM J. Discrete Math.* 31(3):1685–1701.

Graversen SE, Peskir G, Shiryaev AN (2006) Stopping Brownian motion without anticipation as close as possible to its ultimate maximum. *Theory Probability Appl.* 45(1):41–50.

Kleinberg JM, Kleinberg R (2018) Delegated search approximates efficient search. *Proc. ACM Conf. Econom. Comput.* (ACM, New York), 287–302.

Krengel U, Sucheston L (1977) Semiamarts and finite values. *Bull. Amer. Math. Soc.* 83(4):745–747.

Krengel U, Sucheston L (1978) On semiamarts, amarts, and processes with finite value. *Adv. Probability Related Topics* 4:197–266.

Li B, Hoi SC (2014) Online portfolio selection: A survey. *ACM Comput. Survey* 46(3):1–33.

Liu A, Leme RP, Pál M, Schneider J, Sivan B (2021) Variable decomposition for prophet inequalities and optimal ordering. *Proc. ACM Conf. Econom. Comput.* (ACM, New York), 692.

McAfee RP (2008) The gains from trade under fixed price mechanisms. *Appl. Econom. Res. Bull.* 1:1–10.

Myerson RB, Satterthwaite MA (1983) Efficient mechanisms for bilateral trading. *J. Econom. Theory* 29(2):265–281.

Osborne MFM (1959) Brownian motion in the stock market. *Oper. Res.* 7(2):145–173.

Rubinstein A, Wang JZ, Weinberg SM (2020) Optimal single-choice prophet inequalities from samples. *Proc. Innovations Theoretical Comput. Sci. Conf.* (Schloss Dagstuhl - Leibniz Zentrum für Informatik, Wadern, Germany), 60:1–60:10.

Samuel-Cahn E (1984) Comparison of threshold stop rules and maximum for independent nonnegative random variables. *Ann. Probability* 12(4):1213–1216.

Singla S (2018) Combinatorial optimization under uncertainty: Probing and stopping-time algorithms. PhD thesis, Carnegie Mellon University, Pittsburgh, PA.

José Correa is a professor in the Department of Industrial Engineering at Universidad de Chile. His research interests are in game theory and mechanism design, mostly from an algorithmic viewpoint. Recently, he has worked on online allocation problems when users are forward looking.

Andrés Cristi is an assistant professor in the College of Management of Technology at the École Polytechnique Fédérale de Lausanne. His work studies how incentives shape the behavior of agents participating in a market and how the underlying combinatorial structure affects the efficiency of the potential outcomes.

Paul Dütting is a research scientist at Google Research Zurich. His research interests lie at the intersection of economics and computation, and in recent years, he has focused on the study of pricing in online platforms, as well as the efficient design of contracts.

Mohammad Hajiaghayi is the Jack and Rita G. Minker Professor of Computer Science at the University of Maryland at College Park. He is a Guggenheim fellow, and a fellow of the ACM, AAAS, IEEE, and EATCS. His research focuses on the design and analysis of graph algorithms, distributed computing, and algorithmic game theory.

Jan Olkowski is a PhD student at the University of Maryland, advised by Mohammad Hajiaghayi. His research interests are in algorithmic game theory, distributed computing, and consensus algorithms in communication networks.

Kevin Schewior is an assistant professor in the Department of Computer Science and Mathematics at University of Cologne and on leave from his position as an associate professor in the Department of Computer Science and Mathematics at the University of Southern Denmark. In his work, he studies algorithms for combinatorial optimization with a special focus on uncertainty and economic aspects.