

Trading Prophets with Initial Capital

Yossi Azar*

Niv Buchbinder[†]

Roie Levin[‡]

Or Vardi*

“The easiest way to become a millionaire is to be born one.” – Folklore

Abstract

Correa et al. [EC’ 2023] introduced the following trading prophets problem. A trader observes a sequence of stochastic prices for a stock, each drawn from a known distribution, and at each time must decide whether to buy or sell. Unfortunately, they observed that in this setting it is impossible to compete with a prophet who knows all future stock prices.

In this paper, we explore the trading prophets problem when we are given initial capital with which to start trading. We show that initial capital is enough to bypass the impossibility result and obtain a competitive ratio of 3 with respect to a prophet who knows all future prices (and who also starts with capital), and we show that this competitive ratio is best possible. We further study a more realistic model in which the trader must pay multiplicative and/or additive transaction costs for trading which model dynamics such as bid-ask spreads and broker fees.

1 Introduction

Is there an algorithm for making money in the stock market?

Building on the well-established foundations of optimal stopping theory and prophet inequalities [KS77, SC84], Correa et al. [CCD⁺23] formalized a version of this question which they termed the *trading prophets* problem. In their model, a trader observes a daily price for the stock or commodity, where the price $X_i \sim D_i$ is drawn from a known distribution D_i . Each day, the trader must decide whether to buy or sell, while holding at most one unit of the asset.¹ This constraint may arise from storage limitations, regulatory restrictions, or risk management policies that cap the trader’s position. Subject to these constraints, the trader’s objective is to maximize profit, defined as the total revenue from sales minus the total cost from purchases, and the benchmark to measure against is the performance of a “prophet” who has a perfect foresight of all future prices.

Correa et al. observe a simple example in which the competitive ratio of any algorithm is unbounded. The instance spans two days: on day one the price is 1, and on day two the price is $1/\epsilon$ with probability ϵ or 0 otherwise. The algorithm expects that buying on day one leads to profit $\epsilon \cdot 1/\epsilon - 1 = 0$, so it might as well never buy and sit the game out entirely. The prophet on the other hand, with perfect foresight, chooses to buy on day one only when profit is guaranteed on day two, and this grants him strictly positive expected returns of $\epsilon \cdot (1/\epsilon - 1) > 0$. As a result, Correa et al. focus on a relaxed (and arguably less realistic) setting in which the daily price distributions are either i.i.d. or presented in random order.

At first glance, the example above appears to close the door on the general version of the trading prophets problem. However, a closer look reveals that the impossibility of the instance relies on the buy-in cost for the algorithm being prohibitively expensive; in other words, the algorithm must start with no initial capital. One might argue that this is pathological, and that meaningful instances of the trading prophets problem typically span many time steps and enable the trader to amortize buy-in costs over time.

This observation raises the central question we explore in this work:

Do there exist competitive algorithms for trading prophets when both the algorithm and the prophet start with initial capital?

*Department of Computer Science, Tel-Aviv University, Tel-Aviv, Israel. Emails: azar@tauex.tau.ac.il, orvardi@mail.tau.ac.il.

[†]Department of Statistics and Operations Research, School of Mathematical Sciences, Tel Aviv University, Tel Aviv, Israel. Email: niv.buchbinder@gmail.com. The work of Niv Buchbinder is supported in part by the Israel Science Foundation (ISF) grant no. 3001/24, and the United States - Israel Binational Science Foundation (BSF) grant no. 2022418.

[‡]Department of Computer Science, Rutgers University, Piscataway, NJ 08854. Email: roie.levin@rutgers.edu.

¹More generally, the trader may be limited to a fixed number of units, which can be normalized to one for simplicity.

1.1 Our Results

Our first result is that, perhaps surprisingly, buy-in cost is the *only* barrier to achieving good competitive algorithms for the trading prophets problem. Specifically, if we allow the algorithm and the prophet to begin the game with one share of stock (or more generally, if we allow the algorithm to incur an additive loss on top of the multiplicative competitive ratio), the prophet's advantage is dramatically reduced, and the following simple algorithm is 3-competitive.

Algorithm 1: Buy Low Sell High

For all time steps $i = 0, 1, \dots, T$, observe the price $X_i \sim D_i$. Then,

- If the algorithm holds the stock, and $X_i > \mathbb{E}[X_{i+1}]$, sell the stock.
- If the algorithm does not hold the stock, and $X_i \leq \mathbb{E}[X_{i+1}]$, buy the stock.

Algorithm 1 has the added advantage that it is myopic, in that it does not require full foreknowledge of all future distributions (but only the expected stock price tomorrow). In fact it requires only a single bit of information about the future: whether the expected price of the stock tomorrow higher or lower than the today's price. This algorithm generalizes that of [RCV25] that was designed for the i.i.d./random order model (though theirs also works in a more general setting involving multiple stocks). We extend the observation in [RCV25] and show that this simple strategy is in fact optimal among all online algorithms for the adversarial order trading prophets problem. Our main result is:

Theorem 1.1. *Algorithm 1 is 3-competitive for the adversarial order trading prophets problem. Moreover, if $\mathbb{E}[X_i] = \mu$ for all $i = 1, \dots, T$, then the algorithm is 2-competitive. These competitive ratios are best possible.*

It is worth noting that the proof of this theorem relies solely on the initial stock granted to both the algorithm and the optimal solution, with no need for any further additive term in the competitive ratio.

Adding transaction costs or bid-ask spread. In reality, no trader can buy and sell the stock at the same price, and this can happen due to several reasons. Firstly, in trading platforms there is usually a *bid-ask spread*, i.e. an inherent difference between the cost of buying and selling the same commodity that arises naturally due to market dynamics.² Secondly, the trader often pays a *transaction cost* to the broker either as an added percentage of the price or as a fixed flat cost. To model these aspects, we study the *trading prophets problem with transaction costs* in which there are two additional parameters $\epsilon_\Pi, \epsilon_\Sigma \in [0, 1)$. At each time step i , the base price is still sampled as $X_i \sim D_i$, but the algorithm can now only buy the stock at a price of $(1 + \epsilon_\Pi)X_i + \epsilon_\Sigma$ or sell the stock at a price of $(1 - \epsilon_\Pi)X_i - \epsilon_\Sigma$. We refer to ϵ_Π and ϵ_Σ as the multiplicative and additive transaction costs respectively.

This seemingly small change to the model makes the trading prophets problem significantly harder. In Appendix A we show that even in the special case when $\epsilon_\Pi = 0$ and there is only additive transaction cost, no myopic algorithm that at time i uses only information about D_{i+1} can be competitive.³ In fact, even in the i.i.d. case, the natural generalization of Algorithm 1 that buys the stock whenever $X_i \leq \mathbb{E}[X_{i+1}] - \epsilon$ and sells the stock whenever $X_i \geq \mathbb{E}[X_{i+1}] + \epsilon$ fails, and new ideas are needed. Nevertheless, we show that the following algorithm is 2-competitive for the i.i.d. setting.

Algorithm 2: Buy Below Sell Above

Define $z_H \geq z_L$ to be thresholds such that for all $i \in [T]$,^a

$$\Pr[X_i \geq z_H] = \Pr[X_i \leq z_L],$$

and

$$z_H(1 - \epsilon_\Pi) - \epsilon_\Sigma = z_L(1 + \epsilon_\Pi) + \epsilon_\Sigma.$$

For all time steps $i = 0, 1, \dots, T$, observe the price $X_i \sim D$. Then,

- If the algorithm holds the stock, and $X_i \geq z_H$, sell the stock.
- If the algorithm does not hold the stock, and $X_i \leq z_L$, buy the stock.

^aRecall that X_i are i.i.d. We show in Section 2 that such values z_H and z_L exist without loss of generality.

²See, e.g., https://en.wikipedia.org/wiki/Bid-ask_spread.

³We actually show the stronger fact that even a myopic algorithm with knowledge of the *realization* of X_{i+1} cannot be competitive.

Our algorithm is again myopic and only needs to know the two threshold values $z_L \leq z_H$ that depend on the distribution. We note that the median value M , i.e. the value for which $\Pr[X \geq M] = \Pr[X \leq M] = \frac{1}{2}$, always satisfies $z_L \leq M \leq z_H$ (but M need not be precisely in the middle of the range). When $\epsilon_{\Pi} = \epsilon_{\Sigma} = 0$ then $z_H = z_L = M$ and our algorithm reduces to the algorithm analyzed in [CCD⁺23] for the i.i.d. and random arrival models. Our second result is:

Theorem 1.2. *Algorithm 2 is 2-competitive for the trading prophets problem with transaction costs in the i.i.d. model.*

We leave as an intriguing open question whether there exists a constant competitive algorithm for the trading prophets problem with transaction costs in the adversarial order model.

1.2 Related Work

Online Trading. Correa et al. [CCD⁺23] initiated the study of the trading prophets problem. They obtained a tight 2-competitive algorithm for the i.i.d. model. They also obtained a 16-competitive algorithm for the non-i.i.d. random arrival model. Following their work Rajput, Chiplunkar and Vaish [RCV25] studied a generalization of the basic setting in which there are multiple stocks whose prices at a given time may be correlated. The algorithm can hold only a subset of the stocks, or more generally there are matroid constraints on the feasible subsets of stocks that it can hold. They designed algorithms for this more general setting in the i.i.d and random arrival models.

Very recently, Azar et al. [ABLV25] studied a significantly more general model in which the trader can buy and sell n types of goods in bundles. Unlike [CCD⁺23, RCV25], these authors no longer assume prices are stochastically generated, but instead changing in arbitrary adversarial ways; they also allow buyers and sellers to arrive at separate time steps, which is the same as allowing for unbounded bid-ask spreads. Because of the additional difficulty, Azar et al. only managed to obtain logarithmic competitive ratios and their results require resource augmentation, i.e. they compare their algorithm to an optimal solution that must trade at slightly worse prices. It is an interesting question to explore intermediate models that bridge the stochastic and fully adversarial settings.

Prophet Inequalities. The term “prophet inequality” was originally used to describe the sell-only analog of our setting. The first such result is due to Krengel and Sucheston [KS77] (see also [HK82] for an early survey). Other applications to computer science were given by Hajiaghayi et al. [HKS07] and Chawla et al. [CHMS10].

The literature on prophet inequalities is too vast to cover fairly here. We briefly mention some recent trends in this area: blends of prophet and secretary problems [EHL17, EHK24, CS21], single (or few) sample variants where the algorithm only gets sample access to the distributions D_i [AKW14, AKW19, CDF⁺22, GKL24], and versions that explore correlation structures beyond product distributions [RSC87, CL24, LPS24, CGLW21, GHKL24, DKP24, ISW23, MMR24].

2 Preliminaries

We begin by describing the trading prophets problem formally. A trader observes a sequence of T prices $X_1, X_2, \dots, X_T$ for a single stock, where $X_i \sim D_i$ is drawn from a non-negative distribution D_i that is known to the trader. After the price X_i is revealed, the trader must decide whether to buy or sell at that price, and the objective is to maximize profit. We assume the trader may hold at most one unit of stock at any point in time (see [CCD⁺23] for why this is without loss of generality). Define $\mu_i = \mathbb{E}[X_i]$ be the expected price at time i . For convenience, we assume there are extra distributions X_0 and X_{T+1} that are both 0 with probability 1: these allow the algorithm and the prophet to purchase a stock initially for a price of 0, and guarantee our algorithms are well defined at time T .

Using the notation $[i, j] := \{i, \dots, j\}$ for integers i and j , define $B \subseteq [0, T + 1]$ and $S \subseteq [1, T + 1]$ to be the subsets of times in which the algorithm buys and sells the stock respectively. Then the profit of the algorithm is $\text{ALG} := \sum_{i \in S} X_i - \sum_{i \in B} X_i$. When we write $\mathbb{E}[\text{ALG}]$, we take the expectation over the random variables $X_1, \dots, X_i$ (and any randomness of the algorithm).

Let OPT denote the maximum profit achievable offline with full information about the realizations of $X_1, \dots, X_T$ and, currently, with no transaction costs. As observed in [CCD⁺23], this optimal solution always buys at local minima (time steps for which $X_i \leq \max(X_{i-1}, X_{i+1})$) and sells at local maxima

(time step for which $X_i \geq \max(X_{i-1}, X_{i+1})$), and since successive difference terms telescope, we can write $\text{OPT} = \sum_{i=1}^T (X_i - X_{i-1})_+$, where $(X_i - X_{i-1})_+ \triangleq \max\{0, (X_i - X_{i-1})\}$.

Observation 2.1. *For the trading prophets problem, $\mathbb{E}[\text{OPT}] = \sum_{i=1}^T \mathbb{E}[(X_i - X_{i-1})_+]$.*

In the more general version of the trading prophets problem with transaction costs, we are given two additional parameters $\epsilon_\Pi, \epsilon_\Sigma \in [0, 1]$. At each time step the algorithm, as well as the offline prophet, can buy the stock at a price of $(1 + \epsilon_\Pi)X_i + \epsilon_\Sigma$ and sell the stock at a price of $(1 - \epsilon_\Pi)X_i - \epsilon_\Sigma$. We observe that in this case it is no longer true that $\mathbb{E}[\text{OPT}] = \sum_{i=1}^T \mathbb{E}[(X_i - X_{i-1})_+]$, because the difference terms no longer telescope.

We say that an algorithm is α -competitive if

$$(2.1) \quad \mathbb{E}[\text{OPT}] \leq \alpha \cdot \mathbb{E}[\text{ALG}] + c,$$

where c is a constant independent of the length of the sequence T . We remark that allowing such an additive constant is common in online computation and competitive analysis.

To implement our algorithm, we require the following fact.

Fact 2.2. *For any real valued, nonnegative random variable X with continuous cdf, and any $\epsilon_\Pi, \epsilon_\Sigma \in [0, 1]$, there exist thresholds z_L and z_H such that simultaneously*

- (i) $\Pr[X \geq z_H] = \Pr[X \leq z_L]$, and
- (ii) $z_H(1 - \epsilon_\Pi) - \epsilon_\Sigma = z_L(1 + \epsilon_\Pi) + \epsilon_\Sigma$.

Proof. Given a nonnegative real number h , define $\ell(h) = \frac{h(1 - \epsilon_\Pi) - 2\epsilon_\Sigma}{1 + \epsilon_\Pi}$, such that $z_H = h$ and $z_L = \ell(h)$ satisfy (ii). Then define the function

$$f(h) = \Pr(X \geq h) - \Pr(X \leq \ell(h)).$$

When $h = 0$, we have $f(h) = 1$, and when $h \rightarrow \infty$, we have $f(h) = -1$, so by the intermediate value theorem, there exists a value h^* such that $z_H = h^*$ and $z_L = \ell(h^*)$ satisfy both (i) and (ii). $\square$

As observed in [CCD⁺23], this assumption is without loss of generality, as we can add to each price an auxiliary independent uniform perturbation $\delta_i \in [-\delta, \delta]$ and taking δ arbitrary small. We remark that after adding this perturbation, our algorithm may become randomized.

3 Trading Prophets in the Adversarial Model

In this section we analyze Algorithm 1 for the trading prophets problem, proving Theorem 1.1.

The upper bound on the competitive ratio follows directly by the following two lemmas.

Lemma 3.1. *Algorithm 1 satisfies $\mathbb{E}[\text{ALG}] = \sum_{i=1}^T \mathbb{E}[(\mu_i - X_{i-1})_+]$.*

Lemma 3.2. *It holds that $\mathbb{E}[\text{OPT}] \leq 3 \cdot \sum_{i=1}^T \mathbb{E}[(\mu_i - X_{i-1})_+]$. Moreover, if $\mathbb{E}[X_i] = \mu$ for all $i \in [1, T]$, then $\mathbb{E}[\text{OPT}] \leq 2 \cdot \sum_{i=1}^T \mathbb{E}[(\mu - X_{i-1})_+]$*

Proof of Lemma 3.1. For any $i \in [0, T]$ define $p_i \triangleq \Pr(X_i \leq \mu_{i+1})$, $q_i \triangleq \Pr(X_i > \mu_{i+1}) = 1 - p_i$. Let $\mathcal{A}_i$ be the event that the algorithm holds the stock at time step i (before the price x_i is revealed). We observe that the probability that the algorithm holds the stock at time step $i \in [1, T]$ is exactly p_{i-1} . Indeed, if $X_{i-1} \leq \mu_i$ (which happens with probability p_{i-1}), then the algorithm does not sell the stock (if the algorithm has the stock in inventory), and buys the stock (if the algorithm does not have the stock in inventory). Similarly, if $X_{i-1} > \mu_i$ the algorithm sells the stock (if the algorithm has the stock in inventory) and does not buy the stock (if the algorithm does not have the stock in inventory). By the law of total expectation,

$$\begin{aligned} \mathbb{E}[\text{ALG}] &= \sum_{i=1}^T \left(\Pr[\mathcal{A}_{i-1}] \cdot q_i \cdot \mathbb{E}[X_i \mid X_i > \mu_{i+1}] - \Pr[\neg \mathcal{A}_{i-1}] \cdot p_i \cdot \mathbb{E}[X_i \mid X_i \leq \mu_{i+1}] \right) \\ &= \sum_{i=1}^T \left(p_{i-1} \cdot q_i \cdot \mathbb{E}[X_i \mid X_i > \mu_{i+1}] - (1 - p_{i-1}) \cdot p_i \cdot \mathbb{E}[X_i \mid X_i \leq \mu_{i+1}] \right) \\ &= \sum_{i=1}^T \left(p_{i-1} \cdot \left(q_i \cdot \mathbb{E}[X_i \mid X_i > \mu_{i+1}] + p_i \cdot \mathbb{E}[X_i \mid X_i \leq \mu_{i+1}] \right) - p_i \cdot \mathbb{E}[X_i \mid X_i \leq \mu_{i+1}] \right) \end{aligned}$$

$$\begin{aligned}
&= \sum_{i=1}^T \left(p_{i-1} \cdot \mathbb{E}[X_i] - p_i \cdot \mathbb{E}[X_i \mid X_i \leq \mu_{i+1}] \right) \\
(3.1) \quad &= \sum_{i=1}^T p_{i-1} \left(\mu_i - \mathbb{E}[X_{i-1} \mid X_{i-1} \leq \mu_i] \right) \\
&= \sum_{i=1}^T \mathbb{E}[(\mu_i - X_{i-1})_+].
\end{aligned}$$

Equation (3.1) holds since $\mathbb{E}[X_0 \mid X_0 \leq \mu_1] = \mathbb{E}[X_T \mid X_T \leq \mu_{T+1}] = 0$ (as $X_0 = X_{T+1} = 0$). $\square$

We finish with the bound on OPT.

Proof of Lemma 3.2. We first observe that for any $i \in [1, T]$ we have,

$$\begin{aligned}
&\mathbb{E}[(X_i - X_{i-1})_+] = \mathbb{E}[(X_i - \mu_{i+1}) + (\mu_{i+1} - \mu_i) + (\mu_i - X_{i-1})]_+ \\
(3.2) \quad &\leq \mathbb{E}[(X_i - \mu_{i+1})_+] + (\mu_{i+1} - \mu_i)_+ + \mathbb{E}[(\mu_i - X_{i-1})_+] \\
(3.3) \quad &= (\mu_i - \mu_{i+1} + \mathbb{E}[(\mu_{i+1} - X_i)_+]) + (\mu_{i+1} - \mu_i)_+ + \mathbb{E}[(\mu_i - X_{i-1})_+] \\
(3.4) \quad &\leq (\mu_i - \mu_{i+1} + \mathbb{E}[(\mu_{i+1} - X_i)_+]) + \mathbb{E}[(\mu_{i+1} - X_i)_+] + \mathbb{E}[(\mu_i - X_{i-1})_+] \\
&= \mu_i - \mu_{i+1} + 2 \cdot \mathbb{E}[(\mu_{i+1} - X_i)_+] + \mathbb{E}[(\mu_i - X_{i-1})_+].
\end{aligned}$$

Inequality (3.2) follows as $(a + b + c)_+ \leq a_+ + b_+ + c_+$ (and by linearity of expectation). Equality (3.3) follows as $\mathbb{E}[(X_i - \mu_{i+1})_+] = \mu_i - \mu_{i+1} + \mathbb{E}[(\mu_{i+1} - X_i)_+]$ because

$$\mathbb{E}[(X_i - \mu_{i+1})_+] + \mu_{i+1} - \mu_i = \mathbb{E}[\max\{X_i - \mu_{i+1}, 0\} + \mu_{i+1} - X_i] = \mathbb{E}[(\mu_{i+1} - X_i)_+].$$

Inequality (3.4) follows as $(\mu_{i+1} - \mu_i)_+ \leq \mathbb{E}[(\mu_{i+1} - X_i)_+]$ since

$$(\mu_{i+1} - \mu_i)_+ = \max\{0, \mathbb{E}[\mu_{i+1} - X_i]\} \leq \mathbb{E}[\max\{0, \mu_{i+1} - X_i\}] = \mathbb{E}[(\mu_{i+1} - X_i)_+].$$

By Observation 2.1 summing over all $T \in [1, T]$ and using that $X_0 = X_{T+1} = 0$, we get

$$\begin{aligned}
\mathbb{E}[\text{OPT}] &= \sum_{i=1}^T \mathbb{E}[(X_i - X_{i-1})_+] \\
&\leq \sum_{i=1}^T \left(\mu_i - \mu_{i+1} + 2 \cdot \mathbb{E}[(\mu_{i+1} - X_i)_+] + \mathbb{E}[(\mu_i - X_{i-1})_+] \right) \\
&= \sum_{i=1}^T (\mu_i - \mu_{i+1}) + 2 \sum_{i=1}^T \mathbb{E}[(\mu_{i+1} - X_i)_+] + \sum_{i=1}^T \mathbb{E}[(\mu_i - X_{i-1})_+] \\
&= (\mu_1 - \mu_{T+1}) + 3 \cdot \sum_{i=2}^T \mathbb{E}[(\mu_i - X_{i-1})_+] + \mathbb{E}[(\mu_1 - X_0)_+] + 2 \cdot \mathbb{E}[(\mu_{T+1} - X_T)_+] \\
&= 2 \cdot \mathbb{E}[(\mu_1 - X_0)_+] + 3 \cdot \sum_{i=2}^T \mathbb{E}[(\mu_i - X_{i-1})_+] \leq 3 \cdot \sum_{i=1}^T \mathbb{E}[(\mu_i - X_{i-1})_+].
\end{aligned}$$

Hence, $\mathbb{E}[\text{OPT}] \leq 3 \cdot \sum_{i=1}^T \mathbb{E}[(\mu_i - X_{i-1})_+]$.

Finally, in the case that $\mu_i = \mu$ for all $i \in [1, T]$, observe that we can remove one of the $\mathbb{E}[(\mu_{i+1} - X_i)_+]$ terms from the inequality and obtain $\mathbb{E}[\text{OPT}] \leq 2 \cdot \sum_{i=1}^T \mathbb{E}[(\mu - X_{i-1})_+]$. This concludes the proof of Lemma 3.2. $\square$

3.1 A Tight Lower Bound

In this section we prove the second part of Theorem 1.1,

Proposition 3.3. *For any $\epsilon < 1$, there is a large enough T , and a sequence of distributions $X_1, \dots, X_T$ such that for any online algorithm for the trading prophets problem it holds that*

$$\mathbb{E}[\text{OPT}] \geq (3 - O(\epsilon)) \cdot \mathbb{E}[\text{ALG}]$$

Similarly, there is a sequence of i.i.d. distributions such that

$$\mathbb{E}[\text{OPT}] \geq (2 - O(\epsilon)) \cdot \mathbb{E}[\text{ALG}]$$

We first prove the following lemma.

Lemma 3.4. *The expected profit of any online algorithm for the trading prophets problem is at most $\sum_{i=1}^T \mathbb{E}[(\mu_i - X_{i-1})_+]$.*

Together with Lemma 3.1, this implies that the expected profit of Algorithm 1 is at least as large as the expected profit of any other online algorithm. We note that the proof of Lemma 3.4 generalizes that of [RCV25, Lemma 2.2] to the adversarial model.

Proof of Lemma 3.4. Consider an algorithm that buys the stock at times $b_1, \dots, b_k$ and sells the stock at times $s_1, \dots, s_k$ such that $b_1 \leq s_1 \leq b_2 \leq s_2 \leq \dots \leq b_k \leq s_k$. We can equivalently rewrite the algorithm as one that, for every $i \in [k]$, buys at every time step in the interval $[b_i, s_i - 1]$ and sells at every time step in the interval $[b_i + 1, s_i]$. Hence, we may assume that every algorithm always sells one time step after buying, no matter the selling price. The price at time $i + 1$ is independent of the price at time i and of the decision of the algorithm at time i . Therefore, if the realizations of the prices are $x_1, \dots, x_T$, and p_i is the probability that the algorithm buys the stock at time i given the realization x_i , then the algorithm's profit is

$$\mathbb{E}[\text{ALG}] = \sum_{i=1}^T p_{i-1} (\mathbb{E}[X_i] - x_{i-1}).$$

This expression is maximized by the algorithm that buys at step i if and only if $x_i \leq \mathbb{E}[X_{i+1}]$, and in this case the expected profit is $\sum_{i=1}^T \mathbb{E}[(\mu_i - X_{i-1})_+]$. $\square$

We are now ready to prove the proposition, which is a slight adaptation of the examples in [CCD⁺23, Propositions 1 and 2].

Proof of Proposition 3.3. For the first statement, fix $\epsilon < 1$ and consider the following sequence of distributions

$$\begin{aligned} X_0 &= 0 \\ X_{2i-1} &= \begin{cases} 1/\epsilon & \text{w.p. } (1 - \epsilon), \\ 0 & \text{w.p. } \epsilon. \end{cases} \quad \forall i \in [T/2] \\ X_{2i} &= \begin{cases} 1/\epsilon - 1 & \text{w.p. } (1 - \epsilon), \\ 2/\epsilon - 1 & \text{w.p. } \epsilon. \end{cases} \end{aligned}$$

These random variables are non-negative and have expectations $\mu_{2i-1} = 1/\epsilon - 1$ and $\mu_{2i} = 1/\epsilon$. Simple calculations yield

$$\begin{aligned} \mathbb{E}[(\mu_{2i+1} - X_{2i})_+] &= \mathbb{E}[(1/\epsilon - 1 - X_{2i})_+] = 0, \\ \mathbb{E}[(\mu_{2i} - X_{2i-1})_+] &= \mathbb{E}[(1/\epsilon - X_{2i-1})_+] = \epsilon \cdot 1/\epsilon = 1, \\ \mathbb{E}[(X_{2i} - X_{2i-1})_+] &= \epsilon \cdot (1 - \epsilon)(1/\epsilon - 1) + \epsilon \cdot \epsilon(2/\epsilon - 1) + (1 - \epsilon) \cdot \epsilon(1/\epsilon - 1) \\ &= 2(1 - \epsilon)^2 + \epsilon(2 - \epsilon) \geq 2 - 4\epsilon, \\ \mathbb{E}[(X_{2i-1} - X_{2i})_+] &= (1 - \epsilon) \cdot (1 - \epsilon) \cdot 1 \geq 1 - 2\epsilon, \\ \mathbb{E}[(X_1 - X_0)_+] &= \mathbb{E}[(\mu_1 - X_0)_+] = 1/\epsilon - 1. \end{aligned}$$

By Observation 2.1 and Lemma 3.4, for small enough $\epsilon > 0$,

$$\begin{aligned} \mathbb{E}[\text{OPT}] &= \sum_{i=1}^T \mathbb{E}[(X_i - X_{i-1})_+] \geq \frac{T}{2} (3 - 6\epsilon) + (1/\epsilon - 1) - (1 - 2\epsilon) \geq \frac{T}{2} (3 - 6\epsilon), \\ \mathbb{E}[\text{ALG}] &\leq \sum_{i=1}^T \mathbb{E}[(\mu_i - X_{i-1})_+] = T/2 + 1/\epsilon - 1 \leq T/2 + 1/\epsilon. \end{aligned}$$

Taking T sufficiently large concludes the proof of the first statement.

For the second statement, consider the following distribution (as in [CCD⁺23]): $X_0 = 0$ and

$$X_i = \begin{cases} 1 & \text{w.p. } \frac{\epsilon}{2}, \\ \frac{1}{2} & \text{w.p. } 1 - \epsilon, \\ 0 & \text{w.p. } \frac{\epsilon}{2}. \end{cases} \quad \forall i \in [T]$$

This time $\mu = \frac{1}{2}$ and

$$\begin{aligned} \mathbb{E}[\text{OPT}] &= \sum_{i=1}^T \mathbb{E}[(X_i - X_{i-1})_+] = \frac{1}{2} + T \cdot \left(\frac{\epsilon}{2}(1 - \epsilon) \cdot \frac{1}{2} + \left(\frac{\epsilon}{2}\right)^2 \cdot 1 + (1 - \epsilon) \cdot \frac{\epsilon}{2} \cdot \frac{1}{2} \right) \geq T \cdot \left(\frac{\epsilon}{2} - \frac{1}{4}\epsilon^2 \right), \\ \mathbb{E}[\text{ALG}] &\leq \sum_{i=1}^T \mathbb{E}[(\mu - X_{i-1})_+] = \frac{1}{2} + T \cdot \frac{\epsilon}{4}. \end{aligned}$$

Once again taking T sufficiently large concludes the proof. $\square$

4 Trading Prophets with Transaction Fees in the IID Model

In this section we analyze Algorithm 2, proving Theorem 1.2. First, define

- $p \triangleq \Pr[X \leq z_L] = \Pr[X \geq z_H]$,
- $v \triangleq z_H(1 - \epsilon_{\Pi}) - \epsilon_{\Sigma} = z_L(1 + \epsilon_{\Pi}) + \epsilon_{\Sigma}$,
- $v_L \triangleq \mathbb{E}[X \mid X \leq z_L]$,
- $v_H \triangleq \mathbb{E}[X \mid X \geq z_H]$.

Note that $z_L \leq v \leq z_H$. We prove the theorem by upper bounding OPT and lower bounding ALG.

Lemma 4.1. *Algorithm 2 satisfies $\mathbb{E}[\text{ALG}] \geq \frac{1}{2} \cdot pT \cdot [v_H(1 - \epsilon_{\Pi}) - v_L(1 + \epsilon_{\Pi}) - 2\epsilon_{\Sigma}]$.*

Lemma 4.2. *It holds that $\mathbb{E}[\text{OPT}] \leq v + pT \cdot [v_H(1 - \epsilon_{\Pi}) - v_L(1 + \epsilon_{\Pi}) - 2\epsilon_{\Sigma}]$.*

Note, this implies that Algorithm 2 is 2-competitive with additive cost v . We start by bounding OPT.

Proof of Lemma 4.2. To upper bound $\mathbb{E}[\text{OPT}]$, we begin by giving extra power to the optimal solution, which can only increase $\mathbb{E}[\text{OPT}]$. In addition to the option of buying the stock at price $X_i(1 + \epsilon_{\Pi}) + \epsilon_{\Sigma}$ and selling the stock at a price of $X_i(1 - \epsilon_{\Pi}) - \epsilon_{\Sigma}$, we allow the optimal offline solution to (a) immediately sell stock purchased in the same round (as well as the initial inventory stock) for a price of v , and (b) to buy stock (just before the current round) at a price of v .

With this change, the strategy of the optimal solution that sees all realizations X_i is simple. At any time $i = 1, \dots, T$, if $X_i \leq z_L$, buy the stock and sell it immediately at a price of v . At any time when $X_i \geq z_H$, buy the stock for a price of v and sell it immediately.

To see that this is the optimal strategy, first observe that we can assume any algorithm with this extra power only ever buys stock that it sells immediately in the same round for price v , and only ever sells stock that it bought just prior to the round at price v . To this end, consider any solution that buys in round i and sells in round $j > i$. Then we can modify this solution to sell the stock in round i for a price of v , and buy the stock just before round j at a price of v , without changing the overall profit.

Next, observe that for round i , the profit obtainable from buying and then immediately selling at price v is $v - (X_i(1 + \epsilon_{\Pi}) + \epsilon_{\Sigma})$, which is nonnegative if and only if $X_i \leq z_L$. Hence, the solution should perform this transaction if and only if $X_i \leq z_L$, otherwise the overall profit can be improved by following this rule. Similarly, the profit obtainable from buying at price v and then immediately selling is $X_i(1 - \epsilon_{\Pi}) - \epsilon_{\Sigma} - v$, which is nonnegative if and only if $X_i \geq z_H$. Once again, the solution should perform this transaction if and only if $X_i \geq z_H$, otherwise profit can be improved by following the rule.

To conclude, assuming the optimal solution uses the strategy outlined above, the total expected profit is

$$\begin{aligned} \mathbb{E}[\text{OPT}] &\leq v + \sum_{i=1}^T \Pr[X \leq z_L] \cdot \mathbb{E}[v - (X_i(1 + \epsilon_{\Pi}) + \epsilon_{\Sigma}) \mid X \leq z_L] \\ &\quad + \Pr[X \geq z_H] \cdot \mathbb{E}[X_i(1 - \epsilon_{\Pi}) - \epsilon_{\Sigma} - v \mid X \geq z_H] \end{aligned}$$

$$\begin{aligned}
&= v + \sum_{i=1}^T p \cdot [v - (v_L(1 + \epsilon_{\Pi}) + \epsilon_{\Sigma})] + p \cdot [v_H(1 - \epsilon_{\Pi}) - \epsilon_{\Sigma} - v] \\
&= v + pT \cdot [v_H(1 - \epsilon_{\Pi}) - v_L(1 + \epsilon_{\Pi}) - 2\epsilon_{\Sigma}],
\end{aligned}$$

where the extra v comes from selling the initial stock. $\square$

We end with the bound on ALG's performance.

Proof of Lemma 4.1. Let B_i and S_i be indicators that the algorithm buys and sells the stock respectively at time step i . As X_i is independent of the event that the algorithm holds the stock at time step i , we have $\mathbb{E}[X_i \mid S_i = 1] = \mathbb{E}[X_i \mid X_i \geq z_H] = v_H$ and $\mathbb{E}[X_i \mid B_i = 1] = \mathbb{E}[X_i \mid X_i \leq z_L] = v_L$. Moreover, for any time step $i \in [T]$, we have that $\mathbb{E}[B_i + S_i] = p$. This is true because conditioned on holding the stock, the algorithm sells with probability p , and conditioned on not holding the stock, the algorithm buys with probability p . Defining $B = \sum_{i=1}^T B_i$ to be the total number of purchases and $S = \sum_{i=1}^T S_i$ to be the total number of sales, we get $\mathbb{E}[B + S] = pT$. Since the algorithm initially owns the stock, we have $S - 1 \leq B \leq S$, which implies that $\mathbb{E}[S] \geq \frac{pT}{2}$ and $\mathbb{E}[B] \leq \frac{pT}{2}$. Hence,

$$\begin{aligned}
\mathbb{E}[\text{ALG}] &= \sum_{i=1}^T \Pr[S_i = 1] \cdot \mathbb{E}[X_i(1 - \epsilon_{\Pi}) - \epsilon_{\Sigma} \mid S_i = 1] - \Pr[B_i = 1] \cdot \mathbb{E}[X_i(1 + \epsilon_{\Pi}) + \epsilon_{\Sigma} \mid B_i = 1] \\
&= \sum_{i=1}^T \mathbb{E}[S_i] \cdot [v_H(1 - \epsilon_{\Pi}) - \epsilon_{\Sigma}] - \mathbb{E}[B_i] \cdot [v_L(1 + \epsilon_{\Pi}) + \epsilon_{\Sigma}] \\
&= \mathbb{E}[S] \cdot [v_H(1 - \epsilon_{\Pi}) - \epsilon_{\Sigma}] - \mathbb{E}[B] \cdot [v_L(1 + \epsilon_{\Pi}) + \epsilon_{\Sigma}] \\
&\geq \frac{pT}{2} \cdot [v_H(1 - \epsilon_{\Pi}) - v_L(1 + \epsilon_{\Pi}) - 2\epsilon_{\Sigma}].
\end{aligned}$$

$\square$

5 Conclusion

In this work, we show that allowing an additive term in the competitive ratio, or simply permitting the algorithm to start with an initial stock, bypasses the impossibility theorem of prior work and enables competitive algorithms for the trading prophets problem. In particular, we present an optimal 3-competitive algorithm for this original problem as well as a 2-competitive algorithm for a harder version with transaction costs but i.i.d. prices.

This raises a natural open question: can we design a competitive algorithm for the trading prophets problem with transaction costs in the adversarial (i.e. non-i.i.d.) setting? Another interesting direction is to generalize our results to the multiple commodity setting studied by [RCV25], or to investigate intermediate versions that bridge between ours and the fully adversarial (and significantly harder) model of [ABLV25].

References

- [ABLV25] Yossi Azar, Niv Buchbinder, Roie Levin, and Or Vardi. Competitive bundle trading. *arXiv*, <https://arxiv.org/abs/2507.23047>, 2025. $\rightarrow$ cited on page 3, 8
- [AKW14] Pablo Daniel Azar, Robert Kleinberg, and S. Matthew Weinberg. Prophet inequalities with limited information. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA*, pages 1358–1377, 2014. $\rightarrow$ cited on page 3
- [AKW19] Pablo Daniel Azar, Robert Kleinberg, and S. Matthew Weinberg. Prior independent mechanisms via prophet inequalities with limited information. *Games Econ. Behav.*, 118:511–532, 2019. $\rightarrow$ cited on page 3
- [CCD⁺23] José Correa, Andrés Cristi, Paul Duetting, MohammadTaghi Hajiaghayi, Jan Olkowski, and Kevin Schewior. Trading prophets. In *Proceedings of the 24th ACM Conference on Economics and Computation, EC 2023, London, United Kingdom*, pages 490–510. ACM, 2023. $\rightarrow$ cited on page 1, 3, 4, 6, 7

- [CDF⁺22] Constantine Caramanis, Paul Dütting, Matthew Faw, Federico Fusco, Philip Lazos, Stefano Leonardi, Orestis Papadigenopoulos, Emmanouil Pountourakis, and Rebecca Reiffenhäuser. Single-sample prophet inequalities via greedy-ordered selection. In *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA*, pages 1298–1325. SIAM, 2022. → cited on page 3
- [CGLW21] Ioannis Caragiannis, Nick Gravin, Pinyan Lu, and Zihe Wang. Relaxing the independence assumption in sequential posted pricing, prophet inequality, and random bipartite matching. In *Web and Internet Economics - 17th International Conference, WINE 2021, Potsdam, Germany*, volume 13112, pages 131–148, 2021. → cited on page 3
- [CHMS10] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA*, pages 311–320, 2010. → cited on page 3
- [CL24] Chandra Chekuri and Vasilis Livanos. On submodular prophet inequalities and correlation gap. *Theor. Comput. Sci.*, 1019:114814, 2024. → cited on page 3
- [CSZ21] José Correa, Raimundo Saona, and Bruno Ziliotto. Prophet secretary through blind strategies. *Math. Program.*, 190(1):483–521, 2021. → cited on page 3
- [DKP24] Shaddin Dughmi, Yusuf Hakan Kalayci, and Neel Patel. Limitations of stochastic selection problems with pairwise independent priors. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada*, pages 479–490. ACM, 2024. → cited on page 3
- [EHKS24] Soheil Ehsani, MohammadTaghi Hajiaghayi, Thomas Kesselheim, and Sahil Singla. Prophet secretary for combinatorial auctions and matroids. *SIAM J. Comput.*, 53(6):1641–1662, 2024. → cited on page 3
- [EHL17] Hossein Esfandiari, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Morteza Monemizadeh. Prophet secretary. *SIAM J. Discret. Math.*, 31(3):1685–1701, 2017. → cited on page 3
- [GHKL24] Anupam Gupta, Jinqiao Hu, Gregory Kehne, and Roie Levin. Pairwise-independent contention resolution. In *Integer Programming and Combinatorial Optimization - 25th International Conference, IPCO 2024, Wrocław, Poland*, volume 14679 of *Lecture Notes in Computer Science*, pages 196–209. Springer, 2024. → cited on page 3
- [GKL24] Anupam Gupta, Gregory Kehne, and Roie Levin. Set covering with our eyes wide shut. In *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms, SODA 2024, Alexandria, VA, USA*, pages 4530–4553, 2024. → cited on page 3
- [HK82] Theodore P. Hill and Robert P. Kertz. Comparisons of stop rule and supremum expectations of i.i.d. random variables. *Annals of Probability*, 10:336–345, 1982. → cited on page 3
- [HKS07] Mohammad Taghi Hajiaghayi, Robert D. Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence, Vancouver, British Columbia, Canada*, pages 58–65, 2007. → cited on page 3
- [ISW23] Nicole Immorlica, Sahil Singla, and Bo Waggoner. Prophet inequalities with linear correlations and augmentations. *ACM Trans. Economics and Comput.*, 11(3-4):1–29, 2023. → cited on page 3
- [KS77] Ulrich Krengel and Louis Sucheston. Semiamarts and finite values. *Bulletin of the American Mathematical Society*, 83:745–747, 1977. → cited on page 1, 3
- [LPS24] Vasilis Livanos, Kalen Patton, and Sahil Singla. Improved mechanisms and prophet inequalities for graphical dependencies. In *Proceedings of the 25th ACM Conference on Economics and Computation, EC 2024, New Haven, CT, USA*, pages 782–805. ACM, 2024. → cited on page 3
- [MMR24] Simon Mauras, Divyarthi Mohan, and Rebecca Reiffenhäuser. Optimal stopping with interdependent values. In *Proceedings of the 25th ACM Conference on Economics and Computation, EC 2024, New Haven, CT, USA*, pages 246–265. ACM, 2024. → cited on page 3

- [RCV25] Surbhi Rajput, Ashish Chiplunkar, and Rohit Vaish. Trading prophets: How to trade multiple stocks optimally. In *2025 Symposium on Simplicity in Algorithms, SOSA 2025, New Orleans, LA, USA*, pages 238–252. SIAM, 2025. → cited on page 2, 3, 6, 8
- [RSC87] Yosef Rinott and Ester Samuel-Cahn. Comparisons of optimal stopping values and prophet inequalities for negatively dependent random variables. *Ann. Statist.*, 15(4):1482–1490, 1987. → cited on page 3
- [SC84] Ester Samuel-Cahn. Comparison of Threshold Stop Rules and Maximum for Independent Nonnegative Random Variables. *The Annals of Probability*, 12(4):1213 – 1216, 1984. → cited on page 1

A Barriers for the Non-i.i.d. Model with Transaction Costs

In this section we explore the trading prophets problem with fixed additive transaction costs $\epsilon > 0$. We show that several natural approaches fail.

Theorem A.1. *Consider the trading prophets problem with fixed additive transaction costs of $\epsilon > 0$. Then, any online algorithm whose decisions at time step t are only based on the prices in the next k time steps for any fixed k cannot have a bounded competitive ratio. This holds even if at time step t , the algorithm is given the realizations of the prices in time steps $t + 1, \dots, t + k$.*

Proof. We give a deterministic sequence of prices such that the profit of an optimal solution that knows the entire sequence is unbounded, while an online algorithm that bases its decisions at time step t only on the (deterministic) prices at time steps $t + 1, \dots, t + k$ cannot make any non-negative profit. We prove the theorem for $k = 1$. For larger, k just duplicate each price k consecutive times. We assume that the transaction cost is additive ϵ , but the proof holds also for multiplicative transaction cost.

The price sequence is composed of any number of phases. A phase may start at an arbitrary time step t , but for clarity, assume that the first time step in the phase is at time step $t = 1$. The price at $t = 1$ is $x_1 = 1$ and the algorithm is given the price at time step $t = 2$ which is $x_2 = 1 + \epsilon$. We maintain the invariant that if the online algorithm holds the stock at the end of time step $t = 1$, then the price paid by the algorithm for the stock is at least 1. If the algorithm buys the stock at $t = 1$ or holds the stock after at the end of time step $t = 1$, then we reveal two prices $x_3 = 1 - \epsilon$ and $x_4 = 1 + \frac{3}{2}\epsilon$. In this case, the optimal solution can buy the stock at $t = 3$ and sell it at $t = 4$, making profit of $x_4 - x_3 - 2\epsilon = \frac{\epsilon}{2}$. The online algorithm that purchased the stock for a price of 1 cannot make any profit. It can, in principle, sell the stock making a negative profit, but as the phase ends and the next price is 1, there is no reason to do so.

If, on the other hand, at the end of time step $t = 1$ the algorithm does not hold the stock, then the price at $t = 3$ is $x_3 = 1 + \frac{5}{2}\epsilon$ and the phase ends. The optimal solution can buy the stock at time step $t = 1$ and sell it at time step $t = 3$, making a profit of $x_3 - x_1 - 2\epsilon = \frac{\epsilon}{2}$. If the algorithm buys the stock at time step $t = 2$ and sells it at time $t = 3$ it makes a negative profit. In any case, as the phase ends and the next price is again 1 the algorithm has no reason to buy the stock at time $t = 2$ at a price $1 + \epsilon$. This sequence of phases can be repeated any number of times. $\square$

Finally, we show that even for the i.i.d model with fixed transaction costs, a natural generalization of Algorithm 1 fails.

Observation A.2. *Consider the trading prophets problem with fixed additive transaction costs of $\epsilon > 0$, and consider the algorithm that buys whenever the stock is not in inventory and $x_i \leq \mu_{i+1} - \epsilon$ and sells whenever the stock is in inventory and $x_i \geq \mu_{i+1} + \epsilon$. Then this algorithm is not competitive, even in the i.i.d setting.*

Proof. Consider the following distribution of prices.

$$X_i = \begin{cases} 1 + 2\epsilon & \text{w.p. } \frac{1}{5}, \\ 1 - \frac{\epsilon}{2} & \text{w.p. } \frac{4}{5}. \end{cases} \quad \forall i \in [T]$$

For this distribution, $\mu = \frac{1}{5}(1 + 2\epsilon) + \frac{4}{5}(1 - \frac{\epsilon}{2}) = 1$. Hence the algorithm never buys and can make a profit of at most $1 + 2\epsilon - \epsilon$ by selling the initial inventory. On the other hand,

$$\mathbb{E}[\text{OPT}] \geq \sum_{i=1}^{T-1} \Pr[X_i = 1 - \epsilon/2, X_{i+1} = 1 + 2\epsilon] \cdot ((1 + 2\epsilon - \epsilon) - (1 - \frac{\epsilon}{2} + \epsilon))$$

$$\geq (T-1) \cdot \frac{4}{25} \cdot \left((1+2\epsilon-\epsilon) - \left(1-\frac{\epsilon}{2}+\epsilon\right) \right) = (T-1) \cdot \frac{2\epsilon}{25}.$$

Thus, for a large enough T , the algorithm is not competitive. □