The Fortieth AAAI Conference on Artificial Intelligence (AAAI-26)

Sequential Selling with Sunk Cost Bias

Yasushi Kawase[1,2], **Tomohiro Nakayoshi**[1]

[1]The University of Tokyo, Tokyo, Japan
[2]RIKEN Center for Advanced Intelligence Project, Tokyo, Japan
kawase@mist.i.u-tokyo.ac.jp

Abstract

We study a sequential selling problem in which an agent receives daily offers to sell a good, incurs a holding cost each day, and is subject to sunk cost bias—allowing past, irrecoverable costs to influence present decisions. We introduce a formal model parameterizing the degree of sunk cost bias and distinguish between three behavioral types: optimistic (who ignore future bias), naive (who assume their current bias persists), and sophisticated (who anticipate the evolution of their own bias). For each type, we characterize the optimal selling strategy and precisely quantify the worst-case gap in expected objective profit compared to an unbiased agent. Our results show that optimistic agents can suffer a quadratic loss in profit due to excessive waiting, naive agents perform identically to unbiased agents, and sophisticated agents limit their losses to a linear function of the time horizon. These findings clarify how different anticipations of sunk cost bias affect sequential decision-making and suggest targeted interventions to mitigate inefficiency.

1 Introduction

Imagine you possess a car and are considering selling it. Each day, its market price fluctuates stochastically, and you must decide whether to sell at the current price or wait for a potentially better one in the future. However, waiting incurs daily expenses such as parking fees, taxes, and maintenance. Although you may gain some benefit from keeping the car, these costs are assumed to outweigh those benefits. The challenge is to determine the optimal time to sell in order to maximize your overall utility.

This scenario is a classic example of the *optimal stopping problem*, which, in theory, can be solved using dynamic programming techniques (see, e.g., (Benomar, Baudry, and Perchet 2024)). However, in practice, individuals often deviate from the optimal strategy due to behavioral biases. One prominent bias is the *sunk cost bias*, where past, irrecoverable expenditures influence current decisions (Thaler 1980; Arkes and Blumer 1985; Arkes and Ayton 1999). For instance, after paying months of parking fees and maintenance for your unused car, you may hesitate to accept a low offer, thinking, "I have already invested so much, so I should wait for a better price." This reasoning can lead to repeatedly rejecting reasonable offers, accumulating further costs, and ultimately reducing your total utility. Motivated by this phenomenon, we study optimal stopping behavior in agents subject to sunk cost bias.

The goal of this research is to understand how the degree of sunk cost bias influences outcomes in sequential decision-making. We model the scenario as follows: an agent owns an indivisible good and seeks to maximize her profit over a time horizon of T days. Each day $t \in [T]$, its market price is X_t drawn from a known distribution F_t, and the agent must decide whether to sell at the current market price or wait for a potentially better one, incurring a daily holding cost $c > 0$. If holding the good also yields some benefit, then c represents the net cost, that is, the daily holding cost minus the daily gain. We refer to this as the *sequential selling problem*. To capture sunk cost bias, we introduce a parameter $\lambda \geq 0$ representing the agent's sensitivity to past, irrecoverable costs. By day t, after incurring a total holding cost $c \cdot (t-1)$, an agent with sunk cost bias perceives her utility from selling as the objective utility minus an additional subjective penalty proportional to the accumulated cost, i.e., $\lambda c \cdot (t-1)$.

Our model distinguishes between three behavioral types, depending on how agents anticipate and respond to their own sunk cost bias: *optimistic*, *naive*, and *sophisticated* agents. An *optimistic* agent applies sunk cost bias only to her current decision, assuming that her future choices will be unbiased. A *naive* agent is unaware that her bias may change over time and believes her current level of sunk cost bias will persist in the future. In contrast, a *sophisticated* agent fully anticipates how her sunk cost bias will evolve and plans her actions accordingly. We formalize the decision rules for each agent type and analyze how these differing anticipations of bias affect their selling strategies and overall performance.

Our analysis provides insights into how platforms in used goods markets can effectively intervene to improve outcomes. Moreover, our model is not limited to the context of selling a good; it can also be applied to a wide range of sequential decision-making scenarios involving sunk cost bias. Additionally, we only focus on the sunk cost bias and do not consider other behavioral biases such as present bias or reference point bias, since we aim to isolate the effects of sunk cost bias on decision-making. We will discuss these points in more detail in Section 5.

Optimistic	Naive	Sophisticated
$\lambda c \cdot (T-1)(T-2)/2$	0	$\lambda c \cdot (T-1)$

Table 1: The worst-case difference in objective profit between a biased agent and the unbiased agent.

Our Results

We rigorously analyze the performance of each agent type by quantifying the worst-case gap in expected objective profit relative to the unbiased agent, which serves as a benchmark. Specifically, we reveal the worst-case gaps for optimistic, naive, and sophisticated agents as functions of the time horizon T, daily cost c, and the bias parameter λ.

Our main findings are as follows. For optimistic agents, the worst-case profit loss compared to the unbiased agent grows quadratically with the time horizon T, specifically $\lambda c \cdot (T-1)(T-2)/2$ (Theorem 3.1). This demonstrates that optimistic agents can suffer substantial losses due to excessive waiting and the compounding effect of sunk cost bias. Naive agents, in contrast, achieve exactly the same expected profit as the unbiased agent for any value of the bias parameter λ (Theorem 3.4). This occurs because their decision rule coincides with that of the unbiased agent, as they do not anticipate changes in their own bias. Sophisticated agents, who fully anticipate the evolution of their own bias, limit their worst-case profit loss to a linear function of T, namely $\lambda c \cdot (T-1)$ (Theorem 3.5). To establish this, we show that the sophisticated agent behaves identically to the unbiased agent in a modified instance where the cost is scaled by $(1+\lambda)$. Our result illustrates that accounting for future bias can substantially mitigate inefficiency, though some loss remains inevitable.

These results are summarized in Table 1.

We also analyze the case where the daily prices are identically distributed, and the time horizon T is large (Section 4). In this setting, the unbiased and naive agents both adopt a stationary threshold policy, selling as soon as the offer exceeds a fixed value. The sophisticated agent, anticipating future bias, uses a lower threshold to avoid accumulating holding costs. In contrast, the optimistic agent's threshold increases over time, causing her to wait excessively and incur greater losses. This highlights how different anticipations of sunk cost bias fundamentally alter the timing of selling and long-run outcomes.

These findings highlight that naive and sophisticated agents, despite exhibiting sunk cost bias, do not experience significant profit losses compared to the unbiased agent. In contrast, optimistic agents can incur substantial losses due to their tendency to overestimate future opportunities and postpone selling in an attempt to recover sunk costs. This behavior leads to excessive waiting and the accumulation of holding costs.

Our analysis suggests that interventions targeting optimistic or sophisticated agents can meaningfully improve outcomes. For example, introducing early-bird incentives, time-limited offers, or imposing escalating fees can motivate optimistic agents to sell earlier, thereby reducing inefficiency. For sophisticated agents, providing subsidies for holding costs may enable them to wait for higher offers and achieve better sales outcomes. Conversely, from the perspective of the party collecting parking fees (such as a facility owner), our results suggest that increasing the fee can actually increase revenue when dealing with optimistic agents, not only because the daily income per user rises, but also because such agents tend to hold onto the item longer and thus pay fees over a longer period.

Related Work

Behavioral economics has gained significant attention in recent years, particularly in the context of decision-making under uncertainty. Traditional economic models often assume that agents are rational and make decisions based solely on expected utility. However, empirical evidence suggests that individuals frequently deviate from these assumptions, leading to suboptimal choices. Sunk cost bias is one such deviation, where individuals consider past, irrecoverable costs when making current decisions. The sunk cost fallacy can be observed in various situations, such as continuing to invest in a failing project, holding onto a losing stock, or remaining in a bad relationship. Thaler (1980) introduced the concept of sunk cost bias and provided an explanation through the lens of prospect theory (Kahneman and Tversky 1979). Arkes and Blumer (1985) conducted a series of experiments demonstrating that individuals often allow past costs to influence their current decisions, even when it is irrational to do so. Arkes and Ayton (1999) provided a comprehensive review of the sunk cost bias literature, highlighting its prevalence and implications in various domains.

The optimal stopping problem is a classical topic in probability theory and decision science, with applications in areas such as finance, economics, and operations research. A fundamental setting is the *prophet inequality* problem, where an agent must decide whether to accept or reject a sequence of offers over time (Krengel and Sucheston 1977; Samuel-Cahn 1984; Kleinberg and Weinberg 2012; Correa et al. 2019; Lucier 2017; Harada, Kawase, and Sumita 2025). Specifically, it is well known that the optimal strategy for this problem is guaranteed to achieve at least half of the expected maximum value of the offers (Krengel and Sucheston 1977). Recently, (Kleinberg, Kleinberg, and Oren 2022) studied the prophet inequality with reference point bias, where agents are influenced by their past experiences and expectations. (Cai, Gardner, and Weinberg 2023) extended this work to consider agents with both reference point bias and present bias, providing a more comprehensive framework for understanding decision-making under uncertainty. Our model can be viewed as a variant of the prophet inequality problem, where the agent's preferences are influenced by sunk cost bias.

Kleinberg and Oren (2018) introduced a graph-theoretic model of tasks and goals for time-inconsistent agents. This model has been extended in various directions (Kleinberg, Oren, and Raghavan 2016, 2017; Kleinberg et al. 2021; Kleinberg, Kleinberg, and Oren 2022; Akagi, Marumo, and Kurashima 2024; Akagi, Kim, and Kurashima 2025; Tang et al. 2017; Albers and Kraft 2019; Gravin and Immorlica

2016). The sophisticated agent was introduced by Kleinberg, Oren, and Raghavan (2016), building on a formalization from O'Donoghue and Rabin (2008). Although there has been limited research specifically focused on sunk cost bias, Kleinberg et al. (2021) is the most closely related to our work. They analyzed the case where an agent traverses a path from a starting node s to a target node t in a directed graph, obtaining a reward R upon reaching the target. Each node is associated with a cost for moving forward, and there is a probability distribution on its outgoing edges, determining the next node. At each step, the agent must choose whether to move forward or stop. They showed that the gap between the expected reward of the optimal agent and that of the optimistic agent can be exponentially large in the number of nodes. They also showed that the gap between the expected reward of the optimal agent and that of the sophisticated agent is at most $\lambda/(1+\lambda) \cdot R$. A key difference between our results and those of Kleinberg et al. (2021) lies in the underlying structure of cost and profit. In their model, the cost of each step can be set to an arbitrary value, but the profit is fixed. In contrast, in our model, the cost is fixed and the profit is stochastic, which leads to different agent behaviors. As a result of this structural difference, our analysis yields distinct conclusions regarding the impact of sunk cost bias on agent behavior and performance.

2 Model

For a positive integer n, let $[n] := \{1, 2, \ldots, n\}$. The sequential selling problem is defined as follows. Consider a situation where an agent wishes to sell a good over T days. On each day $t \in [T]$, the agent decides whether to sell the good at a market price X_t, which is sampled from a known non-negative distribution F_t. We assume that the expected value of X_t is finite for each $t \in [T]$. Holding the item incurs a daily cost $c > 0$.[1] An instance of the problem is specified by the tuple $(c, T, (F_t)_{t \in [T]})$.

We model an agent exhibiting *sunk cost bias*, where perceived utility depends on accumulated costs. The agent is characterized by a bias parameter λ (≥ 0), representing sensitivity to sunk cost. We refer to an agent with bias parameter λ as λ-biased. An agent with no bias ($\lambda = 0$) is called an *unbiased agent*.

On each day $t \in [T-1]$, the agent chooses to either

1. Sell the good with value X_t and terminate, or
2. Pay cost c to continue to day $t+1$.

On day T, the good must be sold at value X_T.

The objective profit from selling at day t, as evaluated from day t, is X_t. However, due to sunk cost bias, the agent's *perceived utility* from selling at day t, as evaluated from day t, is $X_t - \lambda c \cdot (t-1)$, since the total cost incurred by day t is $c \cdot (t-1)$.

The biased agent makes selling decisions based on perceived utility. We analyze three types of agents: *optimistic*, *naive*, and *sophisticated*. Although their perceived utility on day t is always $X_t - \lambda c \cdot (t-1)$, their expectations about future decisions differ.

Let $t, t' \in [T]$ with $t' \geq t$. For the unbiased agent, the predicted utility at day t for selling at day t' is:

$$u^{(\mathrm{u})}_{t,t'}(X_{t'}) := X_{t'} - c \cdot (t'-t).$$

The perceived utility at day t for selling at day t' is defined as follows for each agent type. The optimistic agent expects that future decisions will be unbiased, resulting in perceived utility:

$$u^{(\mathrm{o})}_{t,t'}(X_{t'}) := \begin{cases} X_{t'} - \lambda c \cdot (t-1) & \text{if } t' = t, \\ X_{t'} - c \cdot (t'-t) & \text{if } t' > t. \end{cases}$$

The naive agent regards the bias present today as persisting into the future, leading to perceived utility:

$$u^{(\mathrm{n})}_{t,t'}(X_{t'}) := X_{t'} - c \cdot (t'-t) - \lambda c \cdot (t-1).$$

The sophisticated agent anticipates that the bias will continue to accumulate in the future, resulting in perceived utility:

$$u^{(\mathrm{s})}_{t,t'}(X_{t'}) := X_{t'} - c \cdot (t'-t) - \lambda c \cdot (t'-1).$$

If the agent type is unspecified, we denote the superscript symbol by a variable $a \in \{\mathrm{u}, \mathrm{o}, \mathrm{n}, \mathrm{s}\}$ for **u**nbiased, **o**ptimistic, **n**aive, and **s**ophisticated agents, respectively.

On each day, each agent makes an optimal decision according to her perceived utility. For a type-a agent and $t, t' \in [T]$ with $t \leq t'$, let $\pi^{(a)}_{t,t'}$ denote the expected perceived utility of the agent from day t given that she does not sell the good before day t'.

We first observe the unbiased agent. At day T, the agent must sell the good, and her expected utility is $\pi^{(\mathrm{u})}_{T,T} = \mathbb{E}[X_T]$. For $t \in [T-1]$, the agent compares the utility $u^{(\mathrm{u})}_{t,t}(X_t) = X_t$ of selling at day t with the utility $\pi^{(\mathrm{u})}_{t,t+1} = \pi^{(\mathrm{u})}_{t+1,t+1} - c$ of not selling at day t. Thus, the agent sells at day t if and only if $X_t \geq \pi^{(\mathrm{u})}_{t+1,t+1}$ and the expected utility is

$$\pi^{(\mathrm{u})}_{t,t} = \mathbb{E}\Big[\max\Big\{X_t,\, \pi^{(\mathrm{u})}_{t+1,t+1} - c\Big\}\Big].$$

The overall expected utility of the unbiased agent is $\pi^{(\mathrm{u})}_{1,1}$, which can be computed by recursively solving the above equation for $t = T-1, T-2, \ldots, 1$. It is worth mentioning that the expected utility of the unbiased agent, $\pi_{1,1}$, is not equal to the expected optimal profit that could be achieved if all the realizations of the prices were known in advance, i.e., $\mathbb{E}\big[\max_{t \in [T]}(X_t - c \cdot (t-1))\big]$.

Next, we examine the behavior of a biased agent of type $a \in \{\mathrm{o}, \mathrm{n}, \mathrm{s}\}$. Since the agent must sell the good at day T, her perceived utility at that time is defined as

$$\pi^{(a)}_{t,T} = \mathbb{E}\Big[u^{(a)}_{t,T}(X_T)\Big].$$

[1] If we wish to handle possibly varying daily costs c_t, we can model this by setting c to a sufficiently small unit and, for each day t, inserting $c_t/c - 1$ consecutive days with identically zero distributions between days.

For $t \le t' < T$, the agent at day t considers that the good will be sold at day t' if it is not sold before day t' and $u^{(a)}_{t,t'}(X_{t'}) \ge \pi^{(a)}_{t,t'+1}$.[2] Thus, we have

$$\pi^{(a)}_{t,t'} = \mathbb{E}\Big[\max\Big\{u^{(a)}_{t,t'}(X_{t'}),\, \pi^{(a)}_{t,t'+1}\Big\}\Big].$$

At day $t \in [T-1]$, the agent sells the good if and only if $u^{(a)}_{t,t}(X_t) \ge \pi^{(a)}_{t,t+1}$. Since $u^{(a)}_{t,t}(X_t) = X_t - \lambda c \cdot (t-1)$, the agent sells the good at day t if and only if the observed value X_t is at least the threshold $\pi^{(a)}_{t,t+1} + \lambda c \cdot (t-1)$.

Proposition 2.1. *For any agent type $a \in \{\mathrm{o}, \mathrm{n}, \mathrm{s}\}$, the agent at day t sells the good if and only if the observed value X_t is at least the threshold*

$$\theta^{(a)}_t := \pi^{(a)}_{t,t+1} + \lambda c \cdot (t-1).$$

This value can be computed by performing $T - t$ integrations.

Let us see the differences among the agents with a concrete example.

Example 2.2. *Consider the instance with $T = 12$, $c = 1$, and $\lambda = 1$, where each offer X_t is drawn from the uniform distribution* $\mathrm{Uniform}[0, 8]$*. Table 2 shows the sales thresholds $\theta^{(a)}_t$ for optimistic, naive, and sophisticated agents. By simple computation, the resulting expected objective profits at the beginning are approximately* 3.120 *(optimistic),* 5.000 *(naive), and* 4.757 *(sophisticated).*

From this example, we observe that the optimistic agent's threshold increases almost linearly over time, leading to excessive waiting and accumulated costs. Notably, from day 6 onward, the optimistic agent sets thresholds so high that they exceed the maximum possible offer, making it impossible to sell on those days. This behavior reflects a tendency to continually postpone the decision to sell, hoping for a better opportunity to recover sunk costs, even though such opportunities may never arise. The naive agent uses an almost constant threshold (≈ 4.0), which is the same as the unbiased agent's threshold, as we will see in Theorem 3.4. The sophisticated agent adopts a lower threshold (≈ 2.3) and sells early to reduce holding costs.

3 Performance Gap Analysis

In this section, we characterize the worst-case gap in expected profit between each biased agent type and the unbiased agent.

Optimistic Agents

We begin by showing that the worst-case profit loss for an optimistic λ-biased agent, relative to the unbiased agent, grows quadratically with the time horizon T.

Theorem 3.1. *For any $\lambda \ge 0$, the worst-case difference in expected profit between an optimistic λ-biased agent and the unbiased agent is $\lambda c \cdot (T-1)(T-2)/2$.*

[2] If $u^{(a)}_{t,t'}(X_{t'}) = \pi^{(a)}_{t,t'+1}$, the agent is indifferent between selling and waiting. For simplicity, we assume that the agent sells the good at day t' in this case. Although this assumption may affect the analysis in minor ways, the same results can still be obtained.

We first establish the upper bound on this profit gap.

Lemma 3.2. *For any instance $(c, T, (F_t)_{t\in[T]})$, the difference in expected objective profit between an optimistic λ-biased agent and the unbiased agent is at most $\lambda c \cdot (T-1)(T-2)/2$.*

Proof. Let $\pi^{(\mathrm{u})}_t$ and $\pi^{(\mathrm{o})}_t$ denote the expected objective profit of the unbiased and optimistic agents at day t, respectively, conditioned on not having sold before day t. Note that $\pi^{(\mathrm{u})}_t = \pi^{(\mathrm{u})}_{t,t}$, but $\pi^{(\mathrm{o})}_t$ may differ from $\pi^{(\mathrm{o})}_{t,t}$ due to the bias.

We claim that for all $t \in [T]$,

$$\pi^{(\mathrm{o})}_t \ge \pi^{(\mathrm{u})}_t - \lambda c \cdot (T-t)(T+t-3)/2.$$

We prove this by backward induction on t. If this claim holds, then at $t = 1$ the difference in objective profits is bounded by

$$\pi^{(\mathrm{u})}_1 - \pi^{(\mathrm{o})}_1 \le \lambda c \cdot (T-1)(T-2)/2.$$

Base case: For $t = T$, both the optimistic and unbiased agents are required to sell the good, so their expected objective profits coincide: $\pi^{(\mathrm{o})}_T = \pi^{(\mathrm{u})}_T = \mathbb{E}[X_T]$. This establishes the base case.

Inductive step: Assume that for some $t \in [T-1]$, the inequality holds at $t+1$, i.e.,

$$\pi^{(\mathrm{o})}_{t+1} \ge \pi^{(\mathrm{u})}_{t+1} - \lambda c \cdot (T+t-2)(T-t+1)/2.$$

Since the optimistic agent has no bias for future decisions, we have $\pi^{(\mathrm{o})}_{t,t+1} = \pi^{(\mathrm{u})}_{t+1} - c$ for all $t \in [T-1]$. Thus, we have

$$\begin{aligned}
\pi^{(\mathrm{o})}_t &= \mathbb{E}\Big[X_t \mathbb{1}_{u^{(\mathrm{o})}_{t,t}(X_t) \ge \pi^{(\mathrm{o})}_{t,t+1}}\Big] \\
&\quad + \mathbb{E}\Big[(\pi^{(\mathrm{o})}_{t+1} - c)\mathbb{1}_{u^{(\mathrm{o})}_{t,t}(X_t) < \pi^{(\mathrm{o})}_{t,t+1}}\Big] \\
&= \mathbb{E}\Big[X_t \mathbb{1}_{u^{(\mathrm{o})}_{t,t}(X_t) \ge \pi^{(\mathrm{u})}_{t+1} - c}\Big] \\
&\quad + \mathbb{E}\Big[(\pi^{(\mathrm{o})}_{t+1} - c)\mathbb{1}_{u^{(\mathrm{o})}_{t,t}(X_t) < \pi^{(\mathrm{u})}_{t+1} - c}\Big] \\
&\ge \mathbb{E}\Big[X_t \mathbb{1}_{u^{(\mathrm{o})}_{t,t}(X_t) \ge \pi^{(\mathrm{u})}_{t+1} - c}\Big] \\
&\quad + \mathbb{E}\Big[\big(\pi^{(\mathrm{u})}_{t+1} - \lambda c\,\tfrac{(T+t-2)(T-t+1)}{2} - c\big)\mathbb{1}_{u^{(\mathrm{o})}_{t,t}(X_t) < \pi^{(\mathrm{u})}_{t+1} - c}\Big] \\
&\ge \mathbb{E}\Big[X_t \mathbb{1}_{u^{(\mathrm{o})}_{t,t}(X_t) \ge \pi^{(\mathrm{u})}_{t+1} - c}\Big] \\
&\quad + \mathbb{E}\Big[(\pi^{(\mathrm{u})}_{t+1} - c)\mathbb{1}_{u^{(\mathrm{o})}_{t,t}(X_t) < \pi^{(\mathrm{u})}_{t+1} - c}\Big] \\
&\quad - \lambda c\,\tfrac{(T+t-2)(T-t+1)}{2} \\
&= \mathbb{E}\Big[\max\{X_t, \pi^{(\mathrm{u})}_{t+1} - c\}\Big] \\
&\quad - \mathbb{E}\Big[\big(X_t - (\pi^{(\mathrm{u})}_{t+1} - c)\big)\mathbb{1}_{X_t \ge \pi^{(\mathrm{u})}_{t+1} - c > u^{(\mathrm{o})}_{t,t}(X_t)}\Big] \\
&\quad - \lambda c\,\tfrac{(T+t-1)(T-t)}{2} \\
&\ge \pi^{(\mathrm{u})}_t - \lambda c(t-1) - \lambda c\,\tfrac{(T+t-1)(T-t)}{2} \\
&= \pi^{(\mathrm{u})}_t - \lambda c \cdot \tfrac{(T+t-2)(T-t+1)}{2},
\end{aligned}$$

Agent type	1	2	3	4	5	6	7	8	9	10	11	12
Optimistic	3.999	4.998	5.997	6.994	7.988	8.975	9.950	10.90	11.79	12.56	13.00	0.000
Naive	3.999	3.998	3.997	3.994	3.988	3.975	3.950	3.899	3.793	3.562	3.000	0.000
Sophisticated	2.343	2.343	2.343	2.343	2.343	2.342	2.341	2.335	2.316	2.250	2.000	0.000

Table 2: Sales thresholds $\theta_t^{(a)}$ at day t for each agent type.

where the first inequality holds by the inductive hypothesis and the last inequality holds because $x \geq \pi_{t+1}^{(u)} - c > u_{t,t}^{(o)}(x) = x - \lambda c \cdot (t-1)$ implies $\lambda c \cdot (t-1) > x - (\pi_{t+1} - c)$. This completes the induction and the proof. □

Next, we establish the matching lower bound.

Lemma 3.3. *Fix $T \geq 3$, $c > 0$, $\lambda > 0$, and any $\varepsilon \in (0, \lambda c)$. Then, there exists an instance of the sequential selling problem $(c, T, (F_t)_{t\in[T]})$ such that the difference in expected objective profit between an optimistic λ-biased agent and the unbiased agent is at least $\lambda c \cdot (T-1)(T-2)/2 - \varepsilon$.*

Proof. For each day $t \in [T]$, define the deterministic price

$$x_t := c\left((t-1) - \tfrac{\lambda(t-1)(t-2)}{2}\right) + \tfrac{t-1}{T-2} \cdot \varepsilon,$$

and let F_t be the point-mass distribution at x_t.

Since prices are deterministic, the unbiased agent will sell at the day maximizing objective profit. For $t \in [T]$,

$$u_{1,t}^{(u)}(x_t) = x_t - c \cdot (t-1) = -\lambda c \cdot \tfrac{(t-1)(t-2)}{2} + \tfrac{t-1}{T-2} \cdot \varepsilon.$$

This is maximized at $t = 2$, yielding the profit of $\pi_1^{(u)} = u_{1,2}^{(u)}(x_2) = \varepsilon/(T-2)$.

Now consider the objective profit of the optimistic agent. At each day $t \in [T-1]$, her perceived utility from selling immediately is

$$\begin{aligned} u_{t,t}^{(o)}(x_t) &= x_t - \lambda c \cdot (t-1) \\ &= c \cdot (t-1) - \lambda c \tfrac{(t-1)(t-2)}{2} + \tfrac{t-1}{T-2} \cdot \varepsilon. \end{aligned}$$

Her perceived utility from waiting and selling at day $t+1$ is

$$\begin{aligned} u_{t,t+1}^{(o)}(x_{t+1}) &= x_{t+1} - c \\ &= c \cdot (t-1) - \lambda c \tfrac{t(t-1)}{2} + \tfrac{t}{T-2} \cdot \varepsilon. \end{aligned}$$

It is straightforward to verify that $u_{t,t+1}^{(o)}(x_{t+1}) > u_{t,t}^{(o)}(x_t)$. As

$$\begin{aligned} \pi_{t,t+1}^{(o)} &= \max\{u_{t,t+1}^{(o)}(x_{t+1}),\, \pi_{t,t+1}^{(o)}\} \\ &\geq u_{t,t+1}^{(o)}(x_{t+1}) > u_{t,t}^{(o)}(x_t), \end{aligned}$$

the optimistic agent always prefers to wait rather than sell. Therefore, she postpones selling until the final day T. Hence, her objective profit is

$$x_T - c \cdot (T-1) = -\lambda c \tfrac{(T-1)(T-2)}{2} + \tfrac{T-1}{T-2} \cdot \varepsilon.$$

Therefore, the gap in objective profit is

$$\tfrac{\varepsilon}{T-2} - \left(-\lambda c \tfrac{(T-1)(T-2)}{2} + \tfrac{T-1}{T-2}\varepsilon\right) = \lambda c \tfrac{(T-1)(T-2)}{2} - \varepsilon.$$

This completes the proof. □

Note that the loss for the optimistic agent is bounded by a quadratic function because the cumulative cost only grows linearly with time. If we allow the daily cost c to vary over time, the loss can grow even faster. Indeed, we can show that the loss can be exponential in the time horizon T, similar to the result by Kleinberg et al. (2021).

Naive Agents

We now show that a naive λ-biased agent achieves exactly the same expected objective profit as the unbiased agent, regardless of the value of λ.

Theorem 3.4. *For any $\lambda \geq 0$, a naive λ-biased agent follows exactly the same policy as the unbiased agent and achieves the same expected objective profit.*

Proof. At day $t \in [T]$, the naive agent's perceived utility for selling at day $t' \geq t$ is $u_{t,t'}^{(n)}(X_{t'}) = u_{t,t'}^{(u)}(X_{t'}) - \lambda c \cdot (t-1)$. Thus, the expected perceived utility satisfies $\pi_{t,t'}^{(n)} = \pi_{t,t'}^{(u)} - \lambda c \cdot (t-1)$, and in particular, $\pi_{t,t+1}^{(n)} = \pi_{t,t+1}^{(u)} - \lambda c \cdot (t-1)$. The naive agent sells at day t if and only if $u_{t,t}^{(n)}(X_t) \geq \pi_{t,t+1}^{(n)}$, which is equivalent to $X_t \geq \pi_{t,t+1}^{(u)}$. This is exactly the same threshold as the unbiased agent. Therefore, the naive agent's selling rule and expected objective profit coincide with those of the unbiased agent. □

Therefore, in our model, the naive agent always follows the same optimal policy as the unbiased agent and achieves identical expected profit, regardless of the bias parameter λ. This result arises from the specific structure of sunk cost bias in our model, where the bias cancels out in the naive agent's decision process. In contrast, for present bias, naive agents suffer significant losses compared to unbiased agents (Kleinberg, Oren, and Raghavan 2016).

Sophisticated Agents

Finally, we show that the worst-case difference in expected profit between a sophisticated λ-biased agent and the unbiased agent grows linearly with the time horizon T.

Theorem 3.5. *For any $\lambda \geq 0$, the worst-case difference in expected profit between a sophisticated λ-biased agent and the unbiased agent is $\lambda c \cdot (T-1)$.*

We first prove the upper bound of the profit difference. The key observation is that the sophisticated agent behaves exactly as the unbiased agent in a modified instance where the cost is scaled by $(1+\lambda)$.

Lemma 3.6. *For any instance $(c, T, (F_t)_{t\in[T]})$, the difference in expected objective profit between a sophisticated λ-biased agent and the unbiased agent is at most $\lambda c \cdot (T-1)$.*

Proof. Let τ denote the random variable of the day that the unbiased agent sells the good in $(c, T, (F_t)_{t\in[T]})$, and let τ' denote the random variable of the day that the unbiased agent sells the good in $((1+\lambda)c, T, (F_t)_{t\in[T]})$. The expected objective profits of the unbiased agent for $(c, T, (F_t)_{t\in[T]})$ and $((1+\lambda)c, T, (F_t)_{t\in[T]})$ are $\mathbb{E}[X_\tau - c\cdot(\tau-1)]$ and $\mathbb{E}[X_{\tau'} - (1+\lambda)c\cdot(\tau'-1)]$, respectively.

By definition, the sophisticated λ-biased agent acting on instance $(c, T, (F_t)_{t\in[T]})$ behaves exactly as the unbiased agent for $((1+\lambda)c, T, (F_t)_{t\in[T]})$, i.e., sells the good at day τ'. Therefore, the sophisticated agent achieves profit $\mathbb{E}[X_{\tau'} - c\cdot(\tau'-1)]$ for $(c, T, (F_t)_{t\in[T]})$.

In addition, we have

$$\begin{aligned}&\mathbb{E}[X_\tau - c(\tau-1)]\\&\quad= \mathbb{E}[X_\tau - (1+\lambda)c\cdot(\tau-1)] + \lambda c\cdot\mathbb{E}[\tau-1]\\&\quad\leq \mathbb{E}[X_{\tau'} - (1+\lambda)c\cdot(\tau'-1)] + \lambda c\cdot\mathbb{E}[\tau-1].\end{aligned}$$

where the inequality holds because τ' is the optimal policy for $((1+\lambda)c, T, (F_t)_{t\in[T]})$.

Hence, we have

$$\begin{aligned}&\mathbb{E}[X_\tau - c\cdot(\tau-1)] - \mathbb{E}[X_{\tau'} - c\cdot(\tau'-1)]\\&\quad\leq \mathbb{E}[X_{\tau'} - (1+\lambda)c\cdot(\tau'-1)] + \lambda c\cdot\mathbb{E}[\tau-1]\\&\qquad - \mathbb{E}[X_{\tau'} - c\cdot(\tau'-1)]\\&\quad= -\lambda c\cdot\mathbb{E}[\tau'-1] + \lambda c\cdot\mathbb{E}[\tau-1]\\&\quad= \lambda c\cdot\mathbb{E}[\tau-\tau'] \leq \lambda c\cdot(T-1),\end{aligned}$$

where the last inequality holds since $\tau' \geq 1$ and $\tau \leq T$.

Therefore, the difference in expected objective profit is at most $\lambda c\cdot(T-1)$. □

We next prove the lower bound of the profit difference.

Lemma 3.7. *Fix $T \geq 2$, $\lambda \geq 0$, $c > 0$, and any $\varepsilon \in (0, \lambda c)$. Then, there exists an instance of the sequential selling problem $(c, T, (F_t)_{t\in[T]})$ such that the difference in expected objective profit between a sophisticated λ-biased agent and the unbiased agent is at least $\lambda c(T-1) - \varepsilon$.*

Proof. For each day $t \in [T]$, define the deterministic price

$$x_t := \begin{cases}0 & \text{if } t = 1,\\(1+\lambda)c\cdot(t-1) - \varepsilon & \text{otherwise,}\end{cases}$$

and let F_t be the point-mass distribution at x_t.

The objective profit of selling at day $t \in [T]$ is

$$\begin{aligned}u^{(\mathrm{u})}_{1,t}(x_t) &= x_t - c\cdot(t-1)\\&= \begin{cases}0 & \text{if } t = 1,\\\lambda c\cdot(t-1) - \varepsilon & \text{otherwise.}\end{cases}\end{aligned}$$

Hence, the unbiased agent waits until day T and achieves $\pi^{(\mathrm{u})}_{1,1} = \lambda c\cdot(T-1) - \varepsilon$.

Now, consider the sophisticated agent at day 1. The perceived utility from selling at day 1 is

$$u^{(\mathrm{s})}_{1,1}(x_1) = x_1 - (1+\lambda)c\cdot 0 = 0.$$

Moreover, the perceived utility at day 1 for selling at a future day $t' \in \{2, 3, \ldots, T\}$ is

$$u^{(\mathrm{s})}_{1,t'}(x_{t'}) = x_{t'} - c\cdot(t'-1) - \lambda c\cdot(t'-1) = -\varepsilon.$$

This implies that the expected perceived utility of not selling at day 1 is $\pi^{(\mathrm{s})}_{1,2} = -\varepsilon < 0$. Hence, the sophisticated agent sells the good at day 1 and her objective profit is 0.

Therefore, the gap in objective profit between the sophisticated agent and the unbiased agent is

$$\lambda c\cdot(T-1) - \varepsilon - 0 = \lambda c\cdot(T-1) - \varepsilon,$$

as claimed. □

4 Identical Distributions

In this section, we discuss the case where the price on each day follows the same distribution, i.e., $F_t = F$ for all $t \in [T]$, and T goes to infinity. Let

$$\theta^{(a)}_t := \pi^{(a)}_{t,t+1} + \lambda c\cdot(t-1)$$

denote the sales threshold at day t for agent type $a \in \{\mathrm{u}, \mathrm{o}, \mathrm{n}, \mathrm{s}\}$, i.e., the minimum value of X_t at which the agent will sell. We assume in this section that the expected value of the price is at least the daily cost, i.e., $\mathbb{E}_{X\sim F}[X] \geq c$.

For the unbiased agent, the threshold is constant over time because there is no difference between the days, as the distributions are identical and there is no sunk cost effect. Let $\theta^{(\mathrm{u})}$ denote the threshold for the unbiased agent. The expected profit from each day is also the same; let $\pi^{(\mathrm{u})}$ denote the expected profit of the unbiased agent. Then, the threshold $\theta^{(\mathrm{u})}$ should be set to $\pi^{(\mathrm{u})} - c$, since choosing to wait yields an expected profit of $\pi^{(\mathrm{u})} - c$. This gives us the equation

$$\mathbb{E}_{X\sim F}[\max\{X,\ \theta^{(\mathrm{u})}\}] = \pi^{(\mathrm{u})} = \theta^{(\mathrm{u})} + c.$$

By Theorem 3.4, we know that the naive agent behaves the same as the unbiased agent. Moreover, by Theorem 3.5, a λ-biased sophisticated agent acts as the unbiased agent for $((1+\lambda)c, T, (F_t)_{t\in[T]})$. Thus, we have the following theorem.

Theorem 4.1. *Let $f(\theta) = \mathbb{E}_{X\sim F}[\max\{X - \theta, 0\}]$. For any λ-biased naive agent, the sales threshold at each day is a constant $\theta^{(\mathrm{n})} \in f^{-1}(c)$ if $c \leq f(0)$ and $\theta^{(\mathrm{n})} = 0$ otherwise. For any λ-biased sophisticated agent, the sales threshold at each day is a constant $\theta^{(\mathrm{s})} \in f^{-1}((1+\lambda)c)$.*

From this theorem, we see that the thresholds for naive and sophisticated agents are monotone non-increasing with respect to both c and λ. In other words, as the cost c or the bias λ increases, naive and sophisticated agents will try to sell the good earlier.

Example 4.2. *For the case where the distributions are independent and identically* $\mathrm{Uniform}[0, 8]$ *and $c = 1$, we have*

$$\begin{aligned}c &= \mathbb{E}\Big[\max\{X - \theta^{(\mathrm{n})},\ 0\}\Big]\\&= \tfrac{1}{2}\cdot(1 - \tfrac{\theta^{(\mathrm{n})}}{8})(8 - \theta^{(\mathrm{n})}),\end{aligned}$$

and the solution is $\theta^{(\mathrm{u})} = 4$. In addition, for a λ-biased sophisticated agent, we have

$$(1+\lambda)c = \mathbb{E}\Big[\max\{X - \theta^{(\mathrm{s})}, 0\}\Big] = \tfrac{1}{2} \cdot (1 - \tfrac{\theta^{(\mathrm{s})}}{8})(8 - \theta^{(\mathrm{s})}),$$

and the solution is $\theta^{(\mathrm{s})} = 8 - 4\sqrt{1+\lambda}$. Specifically, for $\lambda = 1$, we have $\theta^{(\mathrm{s})} = 8 - 4\sqrt{2} \approx 2.343$. These results align with the thresholds in Example 2.2.

For the optimistic agent, the sales threshold on the first day coincides with that of the unbiased agent, since there is no accumulated sunk cost and the agent assumes future decisions will be unbiased. However, on day t, the threshold increases linearly due to accumulation of sunk costs.

Theorem 4.3. *For any λ-biased optimistic agent, the sales threshold at day t is*

$$\theta_t^{(\mathrm{o})} = \theta^{(\mathrm{u})} + \lambda c \cdot (t-1),$$

where $\theta^{(\mathrm{u})}$ is the threshold for the unbiased agent.

This means that as time progresses, the optimistic agent requires increasingly higher offers to sell, attempting to "recover" all previously incurred costs, even though these costs are irrecoverable. In particular, as the daily cost c increases, the threshold rises more rapidly, which quickly reduces the probability that the agent will sell. Interestingly, increasing the daily cost can actually improve the objective profit for the optimistic agent. The following example illustrates this phenomenon.

Example 4.4. *Consider the case where the distributions are independent and identically* $\mathrm{Uniform}[0,8]$ *and $\lambda = 1$.*

If $c = 1$, we have $\theta_t^{(\mathrm{o})} = 4 + (t-1)$ by $\theta^{(\mathrm{u})} = 4$ (see Example 4.2). The agent never sells the good with probability

$$\tfrac{4}{8} \cdot \tfrac{5}{8} \cdot \tfrac{6}{8} \cdot \tfrac{7}{8} = \tfrac{105}{512} \approx 0.2051.$$

This implies that, with probability approximately 0.2051, the agent will never sell the good and will continue to pay a daily cost of 1 forever, resulting in an expected objective profit of negative infinity.

On the other hand, if $c = 4$, we have $\theta_1^{(\mathrm{o})} = 0$ since $\theta^{(\mathrm{u})} = 0$ by $f(0) = \mathbb{E}[X] = 4$. The agent will sell the good on day $t = 1$ with probability 1. The expected objective profit is $\mathbb{E}[X_1] = 4$.

Thus, increasing daily cost c can sometimes improve her expected objective profit. Similarly, monotonicity of the expected objective profit with respect to the strength of the bias λ does not necessarily hold.

We remark that such a counterintuitive behavior never occurs for naive or sophisticated agents.

5 Discussion

In this paper, we introduced a model of sequential selling with sunk cost bias and analyzed the performance of optimistic, naive, and sophisticated agents. We showed that the worst-case difference in profit between an optimistic agent and the unbiased agent is quadratic in T, while for a sophisticated agent it is linear in T. For naive agents, the bias cancels out in their decision process, so they achieve exactly the same expected profit as the unbiased agent, regardless of the bias parameter. Our analysis revealed that optimistic agents tend to be overly optimistic about future opportunities and postpone selling, leading to accumulated holding costs, whereas sophisticated agents anticipate their own future bias and sell earlier to avoid additional losses. Additionally, we provided the sales thresholds for each agent type when the distributions are identical across all days, showing that naive and sophisticated agents have constant thresholds, while optimistic agents have linearly increasing thresholds.

From the perspective of a used goods marketplace platform, our results suggest ways to increase profits by adjusting fees or storage charges to influence seller behavior. For example, optimistic sellers tend to delay selling and accumulate higher costs, so introducing a sales subsidy or bonus for selling on the first day, or increasing daily charges, can encourage earlier sales, thereby increasing transaction volume and fee revenue. For sophisticated sellers, who tend to sell earlier to avoid accumulating losses, reducing the daily holding cost can encourage them to wait longer for better offers. By lowering the daily cost, the platform enables sophisticated agents to be more patient, potentially achieving higher sales prices and improving both seller satisfaction and overall market efficiency. Thus, adjusting the fee structure according to the seller's bias can help the platform achieve more desirable outcomes.

Although we have focused on the context of selling a good, our model applies to a variety of other sequential decision-making scenarios involving sunk cost bias. For example, the same framework can be used to analyze quitting decisions in job search, where an agent must decide whether to accept a current job offer or continue searching while incurring search costs. Another application is in project management, where a manager decides whether to abandon or continue a project after investing resources, with ongoing costs influencing perceived utility. The model also extends to investment decisions, such as when to exit a financial position or terminate a research and development effort, where past expenditures may irrationally affect future choices. These examples illustrate the broad relevance of our analysis to economic, organizational, and personal decision-making processes affected by sunk cost bias.

There are several promising directions for future research. One direction is to incorporate other behavioral biases (e.g., present bias or loss aversion) into the model. When multiple biases coexist, the overall loss can naturally become larger. However, the interaction between different biases may lead to nontrivial effects on efficiency. Analyzing the combined impact and possible synergies among multiple behavioral biases is an important direction for future work. Another direction is to analyze settings with multiple agents, where strategic interactions and market dynamics may amplify or mitigate the effects of sunk cost bias. Finally, validating the model with real-world data from online marketplaces or field experiments could yield further insights into the practical impact of sunk cost bias and intervention effectiveness.

Acknowledgments

This work was partially supported by JST ERATO Grant Number JPMJER2301, JST PRESTO Grant Number JPMJPR2122, JSPS KAKENHI Grant Number JP20K19739, and Value Exchange Engineering, a joint research project between Mercari R4D Lab and RIISE (Research Institute for an Inclusive Society through Engineering).

References

Akagi, Y.; Kim, H.; and Kurashima, T. 2025. A Continuous-Time Tractable Model for Present-Biased Agents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 13510–13519.

Akagi, Y.; Marumo, N.; and Kurashima, T. 2024. Analytically tractable models for decision making under present bias. *Proc. Conf. AAAI Artif. Intell.*, 38(9): 9441–9450.

Albers, S.; and Kraft, D. 2019. Motivating Time-Inconsistent Agents: A Computational Approach. *Theory of Computing Systems*, 63(3): 466–487.

Arkes, H. R.; and Ayton, P. 1999. The sunk cost and concorde effects: Are humans less rational than lower animals? *Psychological Bulletin*, 125(5): 591–600.

Arkes, H. R.; and Blumer, C. 1985. The psychology of sunk cost. *Organizational Behavior and Human Decision Processes*, 35(1): 124–140.

Benomar, Z.; Baudry, D.; and Perchet, V. 2024. Lookback Prophet Inequalities. In *Advances in Neural Information Processing Systems*, volume 37.

Cai, L.; Gardner, J.; and Weinberg, S. M. 2023. Optimal Stopping with Multi-dimensional Comparative Loss Aversion. In *Proceedings of the International Conference on Web and Internet Economics*, 95–112. Springer.

Correa, J.; Foncea, P.; Hoeksma, R.; Oosterwijk, T.; and Vredeveld, T. 2019. Recent developments in prophet inequalities. *ACM SIGecom Exchanges*, 17(1): 61–70.

Gravin, D.; and Immorlica, N. 2016. Procrastination with Variable Present Bias. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, 361.

Harada, T.; Kawase, Y.; and Sumita, H. 2025. Online Algorithms for Repeated Optimal Stopping: Achieving Both Competitive Ratio and Regret Bounds. *arXiv preprint arXiv:2511.04484*.

Kahneman, D.; and Tversky, A. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2): 263–291.

Kleinberg, J.; Kleinberg, R.; and Oren, S. 2022. Optimal stopping with behaviorally biased agents: The role of loss aversion and changing reference points. *Games Econ. Behav.*, 133: 282–299.

Kleinberg, J.; and Oren, S. 2018. Time-inconsistent planning: a computational problem in behavioral economics. *Commun. ACM*, 61(3): 99–107.

Kleinberg, J.; Oren, S.; and Raghavan, M. 2016. Planning problems for sophisticated agents with present bias. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. New York, NY, USA: ACM.

Kleinberg, J.; Oren, S.; and Raghavan, M. 2017. Planning with Multiple Biases. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. New York, NY, USA: ACM.

Kleinberg, J.; Oren, S.; Raghavan, M.; and Sklar, N. 2021. Stochastic Model for Sunk Cost Bias. In *Proceedings of the 37th Conference on Uncertainty in Artificial Intelligence*, volume 161, 1279–1288. PMLR.

Kleinberg, R.; and Weinberg, S. M. 2012. Matroid prophet inequalities. *Proceedings of the 44th Symposium on Theory of Computing (STOC)*, 123–136.

Krengel, U.; and Sucheston, L. 1977. Semiamarts and finite values. *Bulletin of the American Mathematical Society*, 83(4): 745–747.

Lucier, B. 2017. An economic view of prophet inequalities. *ACM SIGecom Exchanges*, 16(1): 24–47.

O'Donoghue, T.; and Rabin, M. 2008. Procrastination on Long-Term Projects. *Journal of Economic Behavior & Organization*, 66(2): 161–175.

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

Tang, P.; Teng, Y.; Wang, Z.; Xiao, S.; and Xu, Y. 2017. Computational Issues in Time-Inconsistent Planning. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, 3665–3671. Association for the Advancement of Artificial Intelligence.

Thaler, R. 1980. Toward a Positive Theory of Consumer Choice. *Journal of Economic Behavior and Organization*, 1(1): 39–60.