



# Optimal Stopping with Multi-dimensional Comparative Loss Aversion

Linda Cai<sup>(✉)</sup>, Joshua Gardner, and S. Matthew Weinberg

Princeton University, Princeton, NJ 08544, USA  
tcai@princeton.edu

**Abstract.** Motivated by behavioral biases in human decision makers, recent work by [11] explores the effects of loss aversion and reference dependence on the prophet inequality problem, where an online decision maker sees candidates one by one in sequence and must decide immediately whether to select the current candidate or forego it and lose it forever. In their model, the online decision-maker forms a reference point equal to the best candidate previously rejected, and the decision-maker suffers from loss aversion based on the quality of their reference point, and a parameter  $\lambda$  that quantifies their loss aversion. We consider the same prophet inequality setup, but with candidates that have *multiple features*. The decision maker still forms a reference point, and still suffers loss aversion in comparison to their reference point as a function of  $\lambda$ , but now their reference point is a (hypothetical) combination of the best candidate seen so far *in each feature*.

Despite having the same basic prophet inequality setup and model of loss aversion, conclusions in our multi-dimensional model differs considerably from the one-dimensional model of [11]. For example, [11] gives a tight closed-form on the competitive ratio that an online decision-maker can achieve as a function of  $\lambda$ , for any  $\lambda \geq 0$ . In our multi-dimensional model, there is a sharp phase transition: if  $k$  denotes the number of dimensions, then when  $\lambda \cdot (k - 1) \geq 1$ , *no non-trivial competitive ratio is possible*. On the other hand, when  $\lambda \cdot (k - 1) < 1$ , we give a tight bound on the achievable competitive ratio (similar to [11]). As another example, [11] uncovers an exponential improvement in their competitive ratio for the random-order vs. worst-case prophet inequality problem. In our model with  $k \geq 2$  dimensions, the gap is at most a constant-factor. We uncover several additional key differences in the multi- and single-dimensional models.

**Keywords:** Optimal Stopping · Reference Dependence · Loss Aversion

## 1 Introduction

Reference dependence describes the tendency for human decision-makers to compare options against a previously seen or anticipated reference point, such that an option inferior to the reference point is perceived as being of a lower quality

than it would be had the person simply not been aware of the reference point [22]. Similarly, loss aversion describes the tendency of human decision-makers to be more sensitive to losses than to gains of equivalent magnitude. Both phenomena have robust empirical support [10], and have been observed in the context of online decision-making [19].

Recent work of [11], has considered these behavioral biases in the context of optimal stopping, constructing biased agents who suffer from loss aversion in comparing their chosen candidate against previously seen and foregone candidates. Prophet inequalities are a canonical optimal stopping problem, where an agent sees  $n$  candidates, each with a scalar value drawn independently from known distributions. The candidates are presented online, such that whenever the agent sees a candidate they must either accept it, halting their search, or forego it and thus lose the candidate forever. In the event they select a candidate, the utility they receive is exactly the value of the candidate. Motivated by the role of reference points and loss aversion, [11] pose a modified prophet inequality where the utility received by the agent is discounted based on their reference point. Specifically, the biased agent also remembers the best candidate they have previously seen, and perceives new candidates as having decreased utility if the best candidate previously seen had a higher value.

The elegant model of [11] captures scenarios in which candidates are well-represented by a single scalar (such as monetary value). Indeed, their model is partially motivated by the empirical findings of [19] concerning the unwillingness of investors to sell stocks below the price at which they were purchased. Similarly, their model captures the phenomenon of biased decision-makers stopping much sooner than optimal, once again mirroring the empirical results of [19]. Intuitively, due to the anticipation of experiencing regret over foregone candidates, biased decision-makers are likely to settle earlier than they should. In this respect, the theoretical model of [11] reflects human behavior relatively well, at least in the case of decisions involving scalar-valued candidates.

However, their model does not capture decisions in which candidates have multiple features of interest. Notably, many key decisions in life including career, romantic partner, or house typically have many dimensions. Moreover, even simple consumer decisions such as selecting which movie to watch or magazine subscription to purchase often have multiple dimensions. Experimental research by [2] suggests that human decision makers suffer from “comparative loss aversion” in which they compare options along multiple dimensions, and perceive disadvantages as being more significant than advantages due to loss aversion, thus decreasing received utility from the option chosen. Consider the case of selecting a house: suppose that house A is the most aesthetically pleasing, house B is in the best location, and house C has the most space. Regardless of which house a person selects, due to comparative loss aversion, the mere awareness of the other options decreases the satisfaction they experience. Empirical research [3,8,14] on the phenomenon of “choice overload” (where people become less satisfied with their choice when they are presented with a larger set of options) further supports the existence of comparative loss aversion.

One interpretation of this phenomenon that has been popularized in psychological literature like *The Paradox of Choice* [20] is that people combine the best features of options they have previously seen into an imagined alternative of superior quality. Later empirical work by [17] found that “regret is related to the comparison between the alternative chosen and the union of the positive attributes of the alternatives rejected”. Intuitively, compared to an imagined alternative combining the best attributes seen, the option chosen is much less satisfying.

In our work, we model this phenomenon by associating each candidate  $i$  with a  $k$ -dimensional vector  $\sigma^{(i)} = (\sigma_1^{(i)}, \dots, \sigma_k^{(i)})$  rather than a single scalar. The agent’s value for a single candidate is just the sum over the agent’s value for the candidate on each dimension, namely,  $\|\sigma^{(i)}\|_1$ . So for an unbiased decision-maker, this is still just an instance of the classic prophet inequality. As in [11], our agent is biased, and maintains a reference point  $s^{(i)}$  at time  $i$  based on options seen so far. The agent suffers from loss aversion parameterized by  $\lambda$ , and experiences (and anticipates) a loss equal to  $\lambda \cdot (\|\sigma^{(i)}\|_1 - \|s^{(i)}\|_1)$  upon selecting candidate  $i$ . The key distinction between our model and that of [11] is in how the reference point is built. In [11], the reference point is simply the vector of a candidate  $j \leq i$  maximizing the overall value  $\|\sigma^{(j)}\|_1$ . In our model, the reference point is instead a hypothetical super-candidate that combines the best candidate seen so far for each feature. Formally, for each dimension  $\ell$ , the value of the reference point is  $s_\ell^{(i)} = \max_{j \leq i} \{\sigma_\ell^{(j)}\}$ . As previously discussed, this enables our model to better capture the multi-dimensional problem of purchasing a house, versus the single-dimensional problem of buying/selling stocks that are well-captured by the model of [11]. To reiterate: our model is identical to that of [11], *except* for the choice of reference point. Surprisingly, this single change from a single-dimensional reference point to a multi-dimensional reference point drastically changes conclusions in the model.

## 1.1 Summary of Results

Our main findings are that an agent who suffers from comparative loss aversion (i.e. our model, the multi-dimensional case) suffers in a fundamentally different manner than an agent who only suffers from reference-dependent loss aversion (i.e. [11], the single-dimensional case). We now overview several results supporting this. In describing our results, “gambler” refers to the online decision-maker in the prophet inequality setting, and “prophet” refers to the offline decision-maker (who knows all options and simply picks the best).

**Unbounded Loss.** Perhaps our most surprising finding is a phase transition in  $\lambda$  as a function of  $k$  when comparing the utility of a biased gambler to that of an unbiased gambler or unbiased prophet. Specifically, [11] show that a biased gambler can always guarantee at least a  $\frac{1}{2+\lambda}$ -approximation to the expected utility achieved by an unbiased prophet, and a  $\frac{1}{1+\lambda}$ -approximation to that of an unbiased gambler (and both of these bounds are tight). When  $\lambda \cdot (k-1) < 1$ ,

we find a natural generalization: the biased gambler can guarantee a  $\frac{1-\lambda(k-1)}{2+\lambda}$ -approximation to the unbiased prophet, and a  $\frac{1-\lambda(k-1)}{1+\lambda}$ -approximation to the unbiased gambler, and these are tight for all  $k$  (including  $k = 1$ , where it matches [11]). However, there is a phase transition at  $\lambda \cdot (k-1) = 1$ : *for any  $\lambda \cdot (k-1) \geq 1$ , and any  $C$ , there exist instances where the biased gambler achieves expected utility 1, but the unbiased gambler (and unbiased prophet) achieve expected utility at least  $C$ .* That is, there is no non-trivial guarantee on the competitive ratio as a function only of  $\lambda$  and  $k$ . Interestingly, this phase transition occurs for all  $k > 1$ , but not for  $k = 1$ . This marks a fundamental difference between online decision-making with comparative loss aversion versus only reference-dependent loss aversion.

Given the importance of the phase transition, we will call  $\lambda \cdot (k-1)$  the *feature-amplified bias*. Intuitively, feature-amplified bias measures how fast the gap between the utilities of unbiased and biased agents grows with the number of candidates. As we will see in an illustrative example in Sect. 3, when the feature-amplified bias is greater than one, the ratio between utility of an unbiased prophet and that of a biased gambler grows exponentially with the number of candidates. The formal statements of our main results when feature-amplified bias is less than one can be found in Theorems 2, 3 and 4. For the setting where feature-amplified bias is at least one, the main results can be found in Theorem 1.

**No Improvement for Random Order, Optimal Order, and I.I.D. Distributions.** One key finding of [11] is that the competitive ratio of a biased gambler compared to an unbiased gambler/prophet drastically improves (the loss is logarithmic in  $\lambda$  instead of linear in  $\lambda$ ) when the candidates are revealed in random order. This of course holds in our model when  $k = 1$  (because it is identical to [11]). When  $k > 1$ , however, this phenomenon no longer occurs. Specifically, for all  $k \geq 2$ , *even if the instance is i.i.d.:* (a) if feature-amplified bias  $\lambda \cdot (k-1) \geq 1$ , then the expected reward of the biased gambler and unbiased prophet/gambler can be unbounded, and (b) when feature-amplified bias  $\lambda \cdot (k-1) < 1$ , the worst-case ratio between the biased gambler and unbiased prophet/gambler is within a constant-factor of the worst-case ratio without the i.i.d. assumption.<sup>1</sup> The formal statements of our results for the i.i.d setting can be found in Theorems 5 and 6.

## 1.2 Related Work

Since its introduction, optimal stopping has been pertinent to the study of economic behaviors with immediate applications in a wide range of economic domains such as search theory, financial trading, auction design, etc. Various models have been proposed based on the optimal stopping paradigm including

<sup>1</sup> Of course, we do not nail down the tight competitive ratio for the gambler in the i.i.d. setting, because even the tight ratio between the gambler and prophet without loss aversion is quite involved. The tight competitive ratio between the unbiased gambler and unbiased prophet for the random order (but not i.i.d.) setting is still unknown.

prophet inequalities (value distribution known), secretary problems (value distribution unknown) and Pandora’s box problem (where a cost is associated with inspection of options).

The model that is most relevant to our present work is that of the prophet inequality, where the agents are given the distribution of their value for the candidates *a priori*. In the seminal work of [13], it was shown that for prophet inequalities where the agent is limited to picking a single candidate, the ratio between the expected largest value (namely, the prophet’s pick) and the expected value of the candidate an agent playing optimally selects online is at most 2. Furthermore, this ratio is tight. [18] further show that by using a threshold based algorithm, the online agent can also guarantee that the expected value of the candidate they select is a  $1/2$ -approximation of the expected optimal value over all candidates. Prophet inequalities have since been further studied under a variety of feasibility constraints on the set of selected candidates [5, 12, 15, 16], and a variety of arrival order constraints [1, 4, 6, 7, 9, 23].

Although optimal stopping may have direct applications in human decision making, very little research has devoted to behavioral models of optimal stopping. Recently, [11] constructed a model which factors in reference dependence and loss aversion, however to our knowledge there is no other prior work exploring behavioral phenomena in the context of optimal stopping.

### 1.3 Organization

The rest of the paper will be organized as follows. In Sect. 2, we discuss our model that incorporates comparative loss aversion and introduce notations that will be used throughout the paper. In Sect. 3, we present an illustrative example that highlights the difference between our model and the single dimensional model presented in [11]. In Sect. 4, we analyze our model under various ordering constraints, with a focus on the gap between the utility of the biased *gambler* and that of rational agents. In Sect. 5 we discuss the limitations of our model and possible directions for future work.

## 2 Model and Preliminary

### 2.1 Modeling Comparative Loss Aversion

**Multidimensional Optimal Stopping.** We now construct a multi-dimensional model of the optimal stopping problem such that candidates have multiple dimensions of value. First, we must formally define the multi-dimensional optimal stopping problem, using slight modifications of the classic definition of the problem in one dimension. Specifically, let there be  $n$  candidates each of which has  $k$  dimensions of value that are of interest to a decision making agent, i.e. there are  $n$  vector valued candidates, such that their value vectors are  $k$ -dimensional. Note that the use of vector values differs from traditional constructions of the problem in which scalar values are used, i.e. in the traditional problem  $k$  is 1.

We will use  $\mathcal{F}_t$  with support on  $\mathbb{R}_+^k$  to denote the prior distribution of the gambler's value vector for the  $t^{th}$  candidate. Let  $\sigma$  be an ordered sequence of  $n$  candidates, where the  $t^{th}$  candidate has a  $k$ -dimensional vector value  $\sigma^{(t)} \in \mathbb{R}^k$  drawn from  $\mathcal{F}_t$ . Formally, we will use  $\mathcal{F} = \times_{t \in [n]} \mathcal{F}_t$  to denote the joint distribution of the value for all candidates, such that  $\sigma$  is drawn from  $\mathcal{F}$ . Moreover, let it be the case that the agent sees the realized value of  $\sigma^{(t)}$  only when they reach step  $t$  in the sequence. Upon reaching the  $t^{th}$  candidate,  $\sigma^{(t)}$ , the agent must either accept them - thus halting their search - or decline them, continuing their search and losing the candidate forever. The agent attempts to select  $\sigma^{(t)}$  to maximize their utility. We now need to reason about what an agent's utility function should look like.

**The Utility of Agents with Comparative Loss Aversion.** In keeping with the notation of [11], let the parameter  $\lambda \geq 0$  serve as a multiplier representing our agent's loss aversion. In particular, let  $\lambda$  be used to amplify the difference between the chosen candidate and the "reference candidate" along any dimensions for which the reference candidate is of higher value. Observe that in the case of our model, reasoning about the reference candidate is somewhat complex since an agent with comparative loss aversion compares each option against all previously seen along different dimensions independently. Regardless, it is without loss of generality to collapse the highest values previously seen along each dimension into single vector and consider this to be the reference candidate.

Accordingly, our agent remembers the best value previously seen along each dimension and combines these to form their reference candidate, which we refer to as the "super candidate", since is an upperbound on the quality of all previously seen candidates and combines all of their best features. Observe that this super candidate changes over time as we see new candidates that are better than all previously seen along some dimension. Due to the agent's reference dependence, when the super candidate is better than the agent's chosen candidate on any dimension, this detracts from the utility they receive. Let us denote the super candidate at step  $t$  in the sequence as  $\mathbf{s}^{(t)}$  and the  $j^{th}$  entry of the super candidate at step  $t$  as  $\mathbf{s}_j^{(t)}$ . We define the super candidate in step  $t$  element-wise as follows:

$$\mathbf{s}_j^{(t)} = \max_{w:w \in [t]} (\sigma_j^{(w)})$$

We define the utility received by a biased gambler, who suffers from loss aversion  $\lambda$  and who selects the  $t^{th}$  candidate in sequence  $\sigma$ , as follows:

$$U_{gb}(\sigma, t) = \sum_j \left( \sigma_j^{(t)} - \lambda(\mathbf{s}_j^{(t)} - \sigma_j^{(t)}) \right) = \|\sigma^{(t)}\|_1 - \lambda \left( \|\mathbf{s}^{(t)}\|_1 - \|\sigma^{(t)}\|_1 \right),$$

and we define the difference between the value selected by the biased gambler and their utility as the biased gambler's *regret*.

Recall that, by construction, we have already restricted all values to be non-negative, allowing us to make use of L1 norms. Moreover, let's compare this to

the utility received by biased prophet who selects option t:

$$U_{pb}(\sigma, t) = \sum_j \left( \sigma_j^{(t)} - \lambda(\mathbf{s}_j^{(n)} - \sigma_j^{(t)}) \right) = \|\sigma^{(t)}\|_1 - \lambda \left( \|\mathbf{s}^{(n)}\|_1 - \|\sigma^{(t)}\|_1 \right)$$

Note that the only difference between the utility of the biased gambler and biased prophet is that the super option of the biased prophet is constructed from all n candidates, whereas that of the biased gambler is constructed only from the t candidates it has seen so far.

Next let us reason about the utilities received by a rational gambler and rational prophet, respectively. In particular, observe that since they do not suffer from any reference dependence or loss aversion, their payoff is just equal to the value of the option they select:

$$U_{pr}(\sigma, t) = U_{gr}(\sigma, t) = \sum_j \sigma_j^{(t)} = \|\sigma^{(t)}\|_1$$

Furthermore, assuming that our agents seek to maximize their utilities, the payoff of a rational prophet does not need parameter t since it selects t deterministically to maximize realized utility. Namely, it suffices to denote the utility of a rational prophet as follows:

$$U_{pr}(\sigma) = \max_{t:t \in [n]} (U_{pr}(\sigma, t))$$

Similarly, since the biased prophet sees all options and selects t deterministically to maximize realized utility, we can remove the parameter t here as well:

$$U_{pb}(\sigma) = \max_{t:t \in [n]} (U_{pb}(\sigma, t))$$

Now that the model has been developed, we can proceed to begin analyzing it. Next we will develop notations to describe the expected utility of biased and rational agents, and their relative differences. These notations will be used throughout the paper.

## 2.2 Preliminaries

In keeping with standard conventions, we use  $\mathbb{E}_{\mathcal{F}}$  to denote the expectation with respect to the distribution  $\mathcal{F}$ . To abbreviate notation, we will use  $\mathbb{E}_{\mathcal{F}}[U_{gb}^*]$  and  $\mathbb{E}_{\mathcal{F}}[U_{gr}^*]$  to denote the expected utility of the biased gambler and a rational gambler, respectively, when they use a utility optimal online algorithm. Similarly, we will use  $\mathbb{E}_{\mathcal{F}}[U_{pb}]$  and  $\mathbb{E}_{\mathcal{F}}[U_{pr}]$  to denote the expected utility of the biased prophet and rational prophet, respectively. In addition, we will use  $V^*$  to denote the maximum value of all candidates, namely,  $V^* = \max_{t \in [n]} \|\sigma^{(t)}\|_1$ . We will use  $S_j^*$

to denote the maximum value on the  $j^{th}$  dimension, namely,  $S_j^* = \max_{t \in [n]} \sigma_j^{(t)}$ .

Next, we will create several definitions to describe the relationship between utility of the biased gambler and the utility of rational agents, which will be frequently referenced in subsequent sections. Informally, the prophet/online utility

ratio describes how well the biased gambler can do compared to the rational prophet/gambler *given a specific prior distribution  $\mathcal{F}$* . On the other hand, the prophet/online competitive ratio describes how well the biased gambler can do compared to the rational prophet/gambler *given the most adversarial prior distribution  $\mathcal{F}$* .

**Definition 1 (Prophet Utility Ratio).** *We will use the term prophet utility ratio to denote the ratio between the expected utility of a rational prophet and the expected optimal utility from a biased gambler under the prior distribution  $\mathcal{F}$ . More formally, we define the prophet utility ratio as  $\frac{\mathbb{E}_{\mathcal{F}}[U_{p_r}]}{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]}$ .*

Note that the Prophet Utility Ratio captures the loss that the biased gambler experiences *both due to bias and due to making decisions online*.

**Definition 2 (Online Utility Ratio).** *We will use the term online utility ratio to denote the ratio between the expected optimal utility of a rational gambler and the expected optimal utility from a biased gambler under the prior distribution is  $\mathcal{F}$ . Formally, we define the online utility ratio as  $\frac{\mathbb{E}_{\mathcal{F}}[U_{g_r}^*]}{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]}$ .*

Note that the Online Utility Ratio captures the loss that the biased gambler experiences *just due to bias*.

**Definition 3 (Prophet/Online Competitive Ratio).** *Given a setting (such as adversarial arrival order or i.i.d prior distribution), we define the terms prophet competitive ratio and online competitive ratio as the maximum prophet and online utility ratio possible in the setting. Specifically, in the adversarial arrival order setting, the prophet and online competitive ratio are formally defined as  $\max_{\mathcal{F}} \frac{\mathbb{E}_{\mathcal{F}}[U_{p_r}]}{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]}$  and  $\max_{\mathcal{F}} \frac{\mathbb{E}_{\mathcal{F}}[U_{g_r}^*]}{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]}$ , respectively. In the i.i.d prior distribution setting, we will denote  $\mathcal{U}$  as the set of all i.i.d distributions. The prophet and online competitive ratio in the i.i.d prior distribution setting are then formally defined as  $\max_{\mathcal{F} \in \mathcal{U}} \frac{\mathbb{E}_{\mathcal{F}}[U_{p_r}]}{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]}$  and  $\max_{\mathcal{F} \in \mathcal{U}} \frac{\mathbb{E}_{\mathcal{F}}[U_{g_r}^*]}{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]}$ .*

As stated in the introduction, we define  $\lambda \cdot (k - 1)$  to be the *feature-amplified bias*. We will soon see that whether the feature-amplified bias  $\lambda \cdot (k - 1)$  is greater or less than one is critical in determining whether the biased gambler has a non trivial prophet/online competitive ratio.

Finally, the following claim that directly follows from classical prophet inequality results will be referenced and used throughout proofs in the rest of the paper.

**Proposition 1.** *For any prior distribution  $\mathcal{F}$  and any assumption about candidate arrival order,  $\mathbb{E}_{\mathcal{F}}[U_{g_r}^*] \geq \frac{1}{2} \cdot \mathbb{E}_{\mathcal{F}}[U_{p_r}]$ .*

*Proof.* In our model, rational agents' utilities are simply their value for the option they select. Therefore our claim follows directly from the classical prophet inequality result [12, 13] that the expected value from a value-optimal online algorithm is a  $\frac{1}{2}$  approximation of the expected value from an offline prophet.

### 3 Unbounded Loss: A Motivating Example

Prior to proving more general claims, we will first show an adversarial example that illustrates a fundamental difference between our multi-dimensional optimal stopping problem with biased agents and the single dimensional model. Specifically, in the single dimensional model, the biased gambler attain worse utility than rational agents only when there is *uncertainty*. (Imagine if the biased agent already knows all values of the candidates, they can simply pick the one with the largest value and suffers no regret). However, in our multi-dimensional model, we will show an example where the biased gambler attains arbitrarily worse utility than rational agents, even when the prior distribution is *deterministic*.

*Example 1.* Consider  $n$  candidates arriving in adversarial order, where the biased gambler considers  $k = 2$  features of candidates and has loss aversion parameter  $\lambda = 2$ . Furthermore, let  $\sigma$  be the following sequence of candidate values:

$$(1, 0), (0, 1), (2, 0), (0, 2), (2^2, 0), (0, 2^2), \dots, (2^{n/2-1}, 0), (0, 2^{n/2-1}),$$

and let the prior distribution  $\mathcal{F}$  be deterministic – where the realized sequence from  $\mathcal{F}$  is exactly  $\sigma$  with probability one.

In our example, the overall value of the candidates increases exponentially over time. Consequently, the rational prophet will receive high utility by simply selecting one of the last two candidates. However, the fact that the candidates are excellent on alternating dimensions causes the biased gambler to suffer from increasing regret as they see more candidates. In fact, should the biased gambler choose a later candidate rather than an earlier candidate, the increment in their regret roughly cancels out the increment in their value for the candidate.

With this intuition in mind, let's formally reason about the behaviors of our agents. Recall that the gambler is aware of the prior distribution  $\mathcal{F}$  and thus is aware of  $\sigma$  a priori for this example. The rational gambler is essentially the same as the rational prophet in this case since they know  $\sigma$  a priori and experience no regret. Thus, the expected utility of both the rational prophet and the rational gambler is simply  $2^{n/2-1}$ .

Under our stylized model, the biased gambler only experiences regret for realized options, completely disregarding any future options they may be aware of a priori. Consequently, the biased gambler has utility 1 if they just take the first candidate and are not presented other candidates, and they have utility  $\leq 0$  if they take any other candidate. (e.g. If the gambler decides to select let's say the candidate  $(4, 0)$ , then the super candidate has value  $(4, 2)$ , which means that the biased gambler's utility is  $\|(4, 0)\|_1 - 2 \cdot (\|(4, 2)\|_1 - \|(4, 0)\|_1) = 4 - 2 \cdot (6 - 4) = 0$ ). Therefore we know  $\mathbb{E}_{\mathcal{F}}[U_{gb}^*] \leq 1$ . As a result, the biased gambler's prophet utility ratio,  $\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} = \frac{U_{pr}(\sigma)}{U_{gb}(\sigma)}$ , and online utility ratio,  $\frac{\mathbb{E}_{\mathcal{F}}[U_{gr}^*]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} = \frac{U_{gr}(\sigma)}{U_{gb}(\sigma)}$ , are both equal to  $2^{n/2-1}$ , and thus tend toward  $\infty$  as  $n$  increases. Recall that prophet and online utility ratios lower bound the prophet and online competitive ratio, respectively, implying that both competitive ratios must also be  $\infty$ .

As we will soon show in Proposition 2 and Proposition 3, our present adversarial example with candidates that are excellent on alternating dimensions can be generalized to any  $k$  and  $\lambda$  such that the feature-amplified bias  $\lambda \cdot (k - 1)$  is at least one in the adversarial setting. Furthermore, as we will show in Proposition 5 and Proposition 6, we can extend our example to i.i.d setting (and thus to random and optimal candidate ordering setting) by carefully constructing the prior distribution so that with high probability, the sequence of realized candidate values presented to the agent is adversarial and (after deleting subsequent duplicates) similar to the sequence  $\sigma$  in this example. Our capability to generalize the adversarial example under *any* arrival order constraint drives our model's departure from the single dimensional model. Further, our adversarial example for the i.i.d setting is highly randomized, which shows that the adversarial examples are not an artifact of deterministic distributions.

Lastly, this example illustrates an additional property unique to our model: the biased gambler may outperform the biased prophet. Notice that in our example, the biased prophet is forced to see both  $(2^{n/2-1}, 0)$  and  $(0, 2^{n/2-1})$  as realized values, resulting in a reference point  $(2^{n/2-1}, 2^{n/2-1})$  – much larger than any value they could have chosen. Consequently, the biased prophet has negative utility and is worse off compared to the biased gambler. In contrast, in [11]'s model (where  $k = 1$ ), the biased prophet is guaranteed to attain at least as much utility as that of online agents.

## 4 Effects of Comparative Loss Aversion

In this section we formally analyze the effect of comparative loss aversion on the biased gambler's utility. Specifically, we compare the biased gambler's utility with those of both the rational prophet and rational gambler. We will focus on the two extremes of arrival order constraints: adversarial order and optimal order (even more specifically, we consider sequences drawn from an i.i.d. prior distribution, where any arrival order is optimal). Under both ordering constraints, we show a phase transition when the feature-amplified bias  $\lambda \cdot (k - 1)$  is equal to one. When the feature-amplified bias  $\lambda(k - 1) \geq 1$ , the biased gambler's prophet and online competitive ratio are both equal to infinity. On the other hand, when the feature-amplified bias  $\lambda(k - 1) < 1$ , the biased gambler can achieve a competitive ratio that is only dependent on  $\lambda$  and  $k$ , and of similar magnitude under both ordering constraints.

It is surprising that our model exhibits similar behavior in terms of competitive ratios regardless of constraints on arrival order. In contrast, [11] shows that for the single dimensional model, assuming the arrival order to be random rather than adversarial results in an exponential improvement in utility guarantees in terms of  $\lambda$ .

Missing proofs in Sect. 4.1 and Sect. 4.2 can be found in the full version of the paper<sup>2</sup>.

---

<sup>2</sup> <https://arxiv.org/abs/2309.14555>.

#### 4.1 Utility Loss Under Adversarial Arrival Order

Firstly, we generalize Example 1 to show that when the feature-amplified bias  $\lambda \cdot (k-1) \geq 1$ , as  $n$  grows arbitrarily large, the prophet and online utility ratios (namely, the ratio between expected utility from the biased gambler and the expected utility from the rational prophet or gambler) can be arbitrarily large. In particular, we construct a prior distribution  $\mathcal{F}_0(n)$ , parameterized by the number of candidates  $n$  such that, when the feature-amplified bias  $\lambda \cdot (k-1) > 1$ , both the prophet and online utility ratio are exponential in terms of  $n$ . When the feature-amplified bias  $\lambda \cdot (k-1) = 1$ , we obtain a more conservative result that prophet and online utility ratio can be linear in terms of  $n$ .

**Proposition 2.** *When the feature-amplified bias  $\lambda \cdot (k-1) > 1$ , for any  $n \in \mathbb{N}^+$ , there exists a sequence  $\sigma$  with  $n$  candidates such that when the distribution  $\mathcal{F}_0(n)$  takes the values of  $\sigma$  deterministically, i.e. with probability 1, and when our prior distribution  $\mathcal{F}$  equals  $\mathcal{F}_0(n)$ , the prophet utility ratio  $\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}$  and the online utility ratio  $\frac{\mathbb{E}_{\mathcal{F}}[U_{gr}^*]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}$  are both equal to  $\Omega\left((\lambda(k-1))^{\frac{n}{k}-1}\right)$ .*

**Proposition 3.** *When the feature-amplified bias  $\lambda \cdot (k-1) = 1$ , for any  $n \in \mathbb{N}^+$ , there exists a sequence  $\sigma$  with  $n$  candidates such that when the distribution  $\mathcal{F}_0(n)$  takes the values of  $\sigma$  deterministically, i.e. with probability 1, and when our prior distribution  $\mathcal{F}$  equals  $\mathcal{F}_0(n)$ , the prophet utility ratio  $\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}$  and the online utility ratio  $\frac{\mathbb{E}_{\mathcal{F}}[U_{gr}^*]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}$  are both equal to  $\Omega\left(\frac{n}{k}\right)$ .*

**Theorem 1.** *When the feature-amplified bias  $\lambda \cdot (k-1) \geq 1$ , the prophet and online competitive ratio are both equal to  $\infty$ .*

On the other hand, when the feature-amplified bias  $\lambda \cdot (k-1) < 1$ , we prove that the biased gambler, by using a simple threshold based algorithm, can get at least a  $\frac{1-\lambda(k-1)}{2+\lambda}$  fraction of the prophet's utility and a  $\frac{1-\lambda(k-1)}{1+\lambda}$  fraction of the optimal rational gambler's utility regardless of the value of  $n$ . We also show that both of these ratios are tight. More formally, we prove the following theorems.

**Theorem 2.** *When the feature-amplified bias  $\lambda \cdot (k-1) < 1$ , for any prior distribution  $\mathcal{F}$ , the inverse of the prophet utility ratio  $\frac{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}{\mathbb{E}_{\mathcal{F}}[U_{pr}]} \geq (1-\lambda(k-1)) \cdot \max\left\{\frac{\gamma}{1+\lambda+k}, \frac{1}{2+\lambda}\right\}$ , where  $\gamma = \frac{\mathbb{E}_{\mathcal{F}}[\sum_{j=1}^k S_j^*]}{\mathbb{E}_{\mathcal{F}}[V^*]}$ . Consequently, the prophet competitive ratio is at most  $\frac{2+\lambda}{1-\lambda(k-1)}$ .*

**Theorem 3.** *When the feature-amplified bias  $\lambda \cdot (k-1) < 1$ , for any prior distribution  $\mathcal{F}$ , the online utility ratio  $\frac{\mathbb{E}_{\mathcal{F}}[U_{gr}^*]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \leq \frac{1+\lambda}{(1-\lambda(k-1))}$ . Consequently, the online competitive ratio is at most  $\frac{1+\lambda}{1-\lambda(k-1)}$ .*

**Theorem 4.** *When the feature-amplified bias  $\lambda \cdot (k-1) < 1$ , for any  $\delta \in (0, 1)$ , there exists a prior distribution  $\mathcal{F}$  such that the prophet utility ratio  $\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq$*

$\frac{2+\lambda}{1-\lambda(k-1)} \cdot (1-\delta)$  and that the online utility ratio  $\frac{\mathbb{E}_{\mathcal{F}}[U_{gr}^*]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq \frac{1+\lambda}{1-\lambda(k-1)} \cdot (1-\delta)$ . Consequently, the prophet competitive ratio is at least  $\frac{2+\lambda}{1-\lambda(k-1)}$  and the online competitive ratio is at least  $\frac{1+\lambda}{1-\lambda(k-1)}$ .

Our proof for Theorem 2 and Theorem 3 follows a similar outline to the single dimensional case presented in [11]. However, our analysis deals with the additional complexity from having multiple dimensions, and often requires a tighter analysis for intermediate terms in order to find the exact competitive ratios. For instance, in Theorem 2, we require an additional term when bounding prophet utility ratio (which will later be used in Theorem 3) compared to a similar claim in [11].

We will start by defining threshold based algorithms.

**Definition 4 (Thresholding algorithms.).** We define a thresholding algorithm parameterized by  $\alpha \in (0, 1)$  and denoted  $A^\alpha$  as follows. First,  $A^\alpha$  sets threshold  $T$  such that  $\alpha = \Pr[V^* \geq T]$  (recall that  $V^*$  is defined as the maximum value of all candidates). Next, during the selection process, the algorithm will select the first candidate whose value  $\|\sigma^{(t)}\|_1$  is at least the threshold  $T$ . We will denote the expected utility from  $A^\alpha$  as  $\mathbb{E}_{\mathcal{F}}[U_{gb}^\alpha]$ .

We will now reason about the biased gambler's utility from using a thresholding algorithm  $A^\alpha$ . In the following two claims, we first bound the utility (and regret) of a biased gambler that uses the thresholding algorithm when they encounter a particular sequence of values  $\sigma$ . This utility bound can then be used to bound the expected utility of the biased gambler given the prior distribution of values.

**Claim 1.** When the feature-amplified bias  $\lambda \cdot (k-1) < 1$ , for any sequence  $\sigma$ , for any threshold  $T$ , and for timestep  $t$  such that  $t$  is the smallest timestep where  $\|\sigma^{(t)}\|_1 \geq T$ , it must be the case that  $U_{gb}(\sigma, t) \geq \|\sigma^{(t)}\|_1 - \lambda(k-1) \cdot T$ .

**Claim 2.** When the feature-amplified bias  $\lambda \cdot (k-1) < 1$ , for any prior distribution  $\mathcal{F}$  and any  $\alpha \in (0, 1)$ ,  $\mathbb{E}_{\mathcal{F}}[U_{gb}^\alpha] \geq ((1+\lambda)\alpha - k\lambda) \cdot T + (1-\alpha) \cdot \sum_{t \in [n]} \mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t} [(\|\sigma^{(t)}\|_1 - T)^+]$ .

We can now bound the utility ratio between a biased gambler using algorithm  $A^\alpha$  and the rational prophet, which can then be used to prove Theorem 2 and Theorem 3. It turns out that Theorem 2 and Theorem 3 requires two different ways of bounding the the utility ratio.

**Lemma 1.** When the feature-amplified bias  $\lambda \cdot (k-1) < 1$ , for any prior distribution  $\mathcal{F}$ , there exists an  $\alpha$  such that  $\frac{\mathbb{E}_{\mathcal{F}}[U_{gb}^\alpha]}{\mathbb{E}_{\mathcal{F}}[U_{pr}]} \geq (1-\lambda(k-1)) \cdot \max \left\{ \frac{\gamma}{1+\lambda+k}, \frac{1}{2+\lambda} \right\}$ , where  $\gamma = \frac{\mathbb{E}_{\mathcal{F}}[\sum_{j=1}^k S_j^*]}{\mathbb{E}_{\mathcal{F}}[V^*]}$  (recall that  $S_j^*$  denotes the maximum value in the sequence along dimension  $j$  and  $V^* = \max_t \|\sigma^{(t)}\|_1$ ).

We are now ready to prove Theorem 2 and Theorem 3.

*Proof* (of Theorem 2). The first line of the theorem follows directly by Lemma 1 and the fact that  $\mathbb{E}_{\mathcal{F}}[U_{gb}^*] \geq \mathbb{E}_{\mathcal{F}}[U_{gb}^\alpha]$ . (Because  $\mathbb{E}_{\mathcal{F}}[U_{gb}^*]$  is defined as the expected utility from a biased gambler that uses the utility optimal online algorithm,  $\mathbb{E}_{\mathcal{F}}[U_{gb}^*]$  is at least the expected utility from a biased gambler using some specific online algorithm).

Now we will derive the second line of the theorem. From the first line of the theorem we know that for any prior distribution  $\mathcal{F}$ ,

$$\frac{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}{\mathbb{E}_{\mathcal{F}}[U_{pr}]} \geq (1 - \lambda(k-1)) \cdot \max \left\{ \frac{\gamma}{1 + \lambda + k}, \frac{1}{2 + \lambda} \right\} \geq \frac{1 - \lambda(k-1))}{2 + \lambda}.$$

Hence, for any prior distribution  $\mathcal{F}$ , the prophet utilit ratio  $\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}$  is at most  $\frac{2+\lambda}{1-\lambda(k-1)}$ . We conclude that the prophet competitive ratio  $\max_{\mathcal{F}} \frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}$  is at most  $\frac{2+\lambda}{1-\lambda(k-1)}$ .

*Proof* (of Theorem 3). Let  $\beta = \frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}$  and let  $\Delta = \frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gr}^*]}$ . Since  $\frac{\mathbb{E}_{\mathcal{F}}[U_{gr}^*]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} = \frac{\beta}{\Delta}$ , proving our theorem is equivalent to proving that  $\frac{\Delta}{\beta} \geq \frac{(1-\lambda(k-1))}{1+\lambda}$ , and we will do exactly this for the rest of the proof.

Firstly, let  $\gamma = \frac{\mathbb{E}_{\mathcal{F}}[\sum_{j=1}^k S_j^*]}{\mathbb{E}_{\mathcal{F}}[V^*]}$ , then by Theorem 2 we know that  $\frac{1}{\beta} = \frac{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}{\mathbb{E}_{\mathcal{F}}[U_{pr}]} \geq (1 - \lambda(k-1)) \cdot \frac{\gamma}{1 + \lambda + k}$ .

Next, because  $\mathbb{E}_{\mathcal{F}}[U_{gb}^*]$  represents the highest expected utility a biased gambler can gain from using an online algorithm,  $\mathbb{E}_{\mathcal{F}}[U_{gb}^*]$  must be at least the expected utility a biased gambler can gain from using the value-optimal online algorithm, which is equal to the optimal utility of a rational gambler minus the loss aversion the biased gambler experience from using the value-optimal online algorithm. Recall that  $S_j^*$  is defined as the maximum value on the  $j^{th}$  dimension among all candidates, thus the value of super candidate the biased gambler experiences loss aversion against must be at most  $\sum_{j=1}^k S_j^*$ . Thus

$$\begin{aligned} \mathbb{E}_{\mathcal{F}}[U_{gb}^*] &\geq \mathbb{E}_{\mathcal{F}} \left[ U_{gr}^* - \lambda \cdot \left( \sum_{j=1}^k S_j^* - U_{gr}^* \right) \right] \\ &= (1 + \lambda) \cdot \mathbb{E}_{\mathcal{F}}[U_{gr}^*] - \lambda \cdot \mathbb{E}_{\mathcal{F}} \left[ \sum_{j=1}^k S_j^* \right]. \end{aligned}$$

By the definition of  $\Delta$ , we know that  $\mathbb{E}_{\mathcal{F}}[U_{gr}^*] = \frac{1}{\Delta} \cdot \mathbb{E}_{\mathcal{F}}[U_{pr}]$ . Similarly, by the definition of  $\gamma$ ,  $\mathbb{E}_{\mathcal{F}}[\sum_{j=1}^k S_j^*] = \gamma \cdot \mathbb{E}_{\mathcal{F}}[V^*] = \gamma \cdot \mathbb{E}_{\mathcal{F}}[U_{pr}]$ . Therefore combining these two statements with the previous inequality we know the following:

$$\mathbb{E}_{\mathcal{F}}[U_{gb}^*] \geq \left( \frac{1 + \lambda}{\Delta} - \lambda \cdot \gamma \right) \cdot \mathbb{E}_{\mathcal{F}}[U_{pr}].$$

Re-expressing algebraically, we have  $\frac{1}{\beta} = \frac{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}{\mathbb{E}_{\mathcal{F}}[U_{pr}]} \geq \left( \frac{1 + \lambda}{\Delta} - \lambda \cdot \gamma \right) = \frac{1 + \lambda - \lambda \cdot \gamma \cdot \Delta}{\Delta}$ .

Now we have two bounds on  $\frac{1}{\beta}$ : we know that  $\frac{1}{\beta} \geq (1 - \lambda(k-1)) \cdot \frac{\gamma}{1+\lambda+k}$  and that  $\frac{1}{\beta} \geq \frac{1+\lambda-\lambda\cdot\gamma\cdot\Delta}{\Delta}$ . By these two bounds, we conclude that

$$\frac{\Delta}{\beta} = \Delta \cdot \frac{1}{\beta} \geq \max \left\{ (1 - \lambda(k-1)) \cdot \frac{\Delta \cdot \gamma}{1 + \lambda + k}, 1 + \lambda - \lambda \cdot \gamma \cdot \Delta \right\}.$$

Notice that inside the max bracket, the first term monotonically increases with  $\Delta$ , while the second term monotonically decreases with  $\Delta$ , therefore the minimum of the max of two terms occurs when the two terms are equal, namely, when  $(1 - \lambda(k-1)) \cdot \frac{\Delta \cdot \gamma}{1 + \lambda + k} = 1 + \lambda - \lambda \cdot \gamma \cdot \Delta$ . This is achieved when  $\Delta = \frac{1+\lambda+k}{\gamma \cdot (1+\lambda)}$ , at which point

$$(1 - \lambda(k-1)) \cdot \frac{\Delta \cdot \gamma}{1 + \lambda + k} = 1 + \lambda - \lambda \cdot \gamma \cdot \Delta = \frac{1 - \lambda(k-1)}{1 + \lambda}.$$

We thus conclude that it is always the case that  $\frac{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}{\mathbb{E}_{\mathcal{F}}[U_{gr}^*]} = \frac{\Delta}{\beta} \geq \frac{1 - \lambda(k-1)}{1 + \lambda}$  regardless of the value of  $\Delta$ .

We remain to discuss Theorem 4, which we prove by constructing a worst case prior distribution. Our construction combines a generalization of Example 1 with the canonical two item example showing a 1/2-competitive ratio for classical prophet inequality. In our construction, all but the last candidate has deterministic values constructed similarly to those in Example 1 (but the value of the candidates grow much slower due to the fact that the feature-amplified bias is less than one). The last candidate has a high value with a very small probability, which is similar to constructions for classical prophet inequality.

*Proof (Proof Sketch)* (of Theorem 4). Let the first  $n-1$  candidates have the following deterministic value:

$$\begin{aligned} (1, 0, \dots, 0), & \quad (0, 1, 0, \dots, 0), & \dots, & \quad (0, \dots, 0, 1), \\ (1 + \beta, 0, \dots, 0), & \quad (0, 1 + \beta, 0, \dots, 0), & \dots, & \quad (0, \dots, 0, 1 + \beta), \\ (1 + \beta + \beta^2, 0, \dots, 0), & \quad (0, 1 + \beta + \beta^2, 0, \dots, 0), & \dots, & \quad (0, \dots, 0, 1 + \beta + \beta^2), \\ & \quad \dots \\ (\sum_{j=0}^{\lceil n/k \rceil - 1} \beta^j, 0, \dots, 0), & \quad (0, \sum_{j=0}^{\lceil n/k \rceil - 1} \beta^j, 0, \dots, 0), & \dots, & \quad (0, \dots, 0, \sum_{j=0}^{\lceil n/k \rceil - 1} \beta^j). \end{aligned}$$

Then, the last randomized candidate is equal to

$$\left( (1 - \epsilon)(1 + \lambda) \cdot \sum_{j=0}^{\lceil n/k \rceil - 1} \beta^j / \epsilon, 0, \dots, 0 \right)$$

with probability  $\epsilon$  and is equal to  $(0, 0, \dots, 0)$  with probability  $1 - \epsilon$ . One can then verify that the biased gambler can only get utility at most 1, since by the time the biased gambler reach the last randomized candidate, their imaginary super candidate is too strong for the biased gambler to gain positive expected utility

from the randomized candidate. However, the rational prophet can simply select the maximum between the best deterministic option and the randomized option, getting utility that approaches  $\frac{2+\lambda}{1-\lambda(k-1)}$  as  $\epsilon$  approaches 0 and  $n$  approaches  $\infty$ . Similarly, the rational gambler can simply always select the randomized option and get utility that approaches and  $\frac{1+\lambda}{1-\lambda(k-1)}$ .

## 4.2 No Improvement Under Relaxed Arrival Order

We now show that most of our results for adversarial arrival order can be used to derive analogous results for the case where the prior distribution is i.i.d.

Specifically, in the next claim (Proposition 4), we will show that removing repeated candidate values from a sequence does not affect the biased gambler's optimal utility. We will then use the claim to establish in a subsequent lemma (Lemma 2) that for each deterministic distribution (i.e. each distribution such that values drawn from it are equal to some sequence  $\sigma$  with probability 1), we can find an i.i.d prior distribution  $\mathcal{F}$  with a similar realized sequence (after deleting duplicates), and hence similar prophet and online utility ratio. As we have discussed in Sect. 3 and Sect. 4.1, the construction of the worst case instances for the adversarial arrival order setting is either a fully deterministic sequence, or with few randomized options. Consequently, Lemma 2 can be leveraged to translate all lower bounds on the utility ratios in adversarial setting into lower bounds in the i.i.d setting.

**Definition 5.** A sequence  $\sigma$  is called *succinct* if there are no two identical candidates in the sequence and no candidate's value is a zero vector.

**Definition 6.** The representation  $r(\sigma)$  of a sequence  $\sigma$  is the sequence of unique elements in  $\sigma$  ordered in the ascending order of their first occurrence in  $\sigma$ . (e.g.  $r(1, 2, 1, 2, 3) = 1, 2, 3$ ).

**Proposition 4.**  $U_{p_r}(\sigma) = U_{p_r}(r(\sigma))$  and  $U_{g_b}(\sigma) = U_{g_b}(r(\sigma))$ .

**Lemma 2.** For any  $\epsilon \in (0, 1)$ , any succinct sequence  $\sigma$  such that its length  $m = |\sigma|$  is greater than 1, there exists a distribution  $\mathcal{F}$  where all candidates are i.i.d with support on the set  $\{\sigma^{(1)}, \dots, \sigma^{(m)}\}$  such that the prophet utility ratio  $\frac{\mathbb{E}_{\mathcal{F}}[U_{p_r}]}{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]} \geq (1 - \epsilon) \cdot \frac{U_{p_r}(\sigma)}{U_{g_b}(\sigma)}$ .

We can now use Lemma 2 to prove that when the candidate values are i.i.d, it is also the case that when the feature-amplified bias  $\lambda(k-1) \geq 1$ , prophet and online utility ratios can get progressively worse as the number of candidates  $n$  increases, and as a result, both the prophet and online competitive ratio are equal to  $\infty$ .

**Proposition 5.** When the feature-amplified bias  $\lambda \cdot (k-1) > 1$ , for all  $n \in \mathbb{N}^+$ , there exists a prior distribution  $\mathcal{F}_0(n)$  where all candidates are i.i.d such that when  $\mathcal{F} = \mathcal{F}_0(n)$ , the prophet utility ratio  $\frac{\mathbb{E}_{\mathcal{F}}[U_{p_r}]}{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]}$  and the online utility ratio  $\frac{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]}{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]}$  are both equal to  $\Omega\left((\lambda(k-1))^{f(\lambda, k) \cdot \log^{\frac{1}{2}} n}\right)$ , where  $f(\lambda, k)$  is a function that only depends on  $\lambda$  and  $k$ .

**Proposition 6.** *When the feature-amplified bias  $\lambda \cdot (k-1) = 1$ , for all  $n \in \mathbb{N}^+$ , there exists a prior distribution  $\mathcal{F}_0(n)$  where all candidates are i.i.d such that when  $\mathcal{F} = \mathcal{F}_0(n)$ , the prophet utility ratio  $\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}$  and the online utility ratio  $\frac{\mathbb{E}_{\mathcal{F}}[U_{gr}^*]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}$  are both equal to  $\Omega\left(\frac{\log n}{k \cdot \log \log n}\right)$ .*

**Theorem 5.** *In the i.i.d prior distribution setting, when the feature-amplified bias  $\lambda \cdot (k-1) \geq 1$ , the prophet and online competitive ratio are both equal to  $\infty$ .*

On the other hand, when the feature-amplified bias  $\lambda(k-1) < 1$ , we can utilize Lemma 2 to show a lower bound for the prophet and online competitive ratios that matches those in the adversarial setting within constant factors.

**Proposition 7.** *When the feature-amplified bias  $\lambda(k-1) < 1$ , for any  $\delta \in (0, 1)$ , there exists an i.i.d prior distribution  $\mathcal{F}$  where the prophet utility ratio  $\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq (1-\delta) \cdot \frac{1}{(1-\lambda(k-1))}$ .*

**Theorem 6.** *When the feature-amplified bias  $\lambda(k-1) < 1$  and  $k > 1$ , for any  $\delta \in (0, 1)$ , there exists an i.i.d prior distribution  $\mathcal{F}$  where the prophet utility ratio  $\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq \frac{(1-\delta)}{3} \cdot \frac{2+\lambda}{1-\lambda(k-1)}$  and the online utility ratio  $\frac{\mathbb{E}_{\mathcal{F}}[U_{gr}^*]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq \frac{(1-\delta)}{4} \cdot \frac{1+\lambda}{1-\lambda(k-1)}$ . Consequently, the prophet competitive ratio is at least  $\frac{1}{3} \cdot \frac{2+\lambda}{(1-\lambda(k-1))}$  and the online competitive ratio is at least  $\frac{1}{4} \cdot \frac{1+\lambda}{(1-\lambda(k-1))}$ .*

## 5 Conclusion and Directions for Future Work

We consider an online decision-maker who suffers from comparative loss aversion. Our model is identical to that of [11], except for how the reference point is formed. This change alone accounts for fundamental differences in conclusions. For example, even the prophet suffers loss in our model, and this enables novel phenomena to occur. Additionally, the loss due to bias in our model can be unboundedly bad, and is not significantly mitigated by assumptions on the prophet inequality arrival order.

Empirically observed biases have only recently been incorporated into research on optimal stopping. Our work takes a step beyond [11] and uncovers new phenomena by changing just one aspect of their model. An important direction for future work is to continue developing models that uncover additional phenomena of behavioral game theory (one such possibility is those studied in [21], where adding inferior options can influence decision-making).

## References

1. Beyhaghi, H., Golrezaei, N., Leme, R.P., Pal, M., Sivan, B.: Improved approximations for free-order prophets and second-price auctions. CoRR abs/1807.03435 (2018). <http://arxiv.org/abs/1807.03435>

2. Brenner, L., Rottenstreich, Y., Sood, S.: Comparison, grouping, and preference. *Psychol. Sci.* **10**(3), 225–229 (1999)
3. Chernev, A., Böckenholt, U., Goodman, J.: Choice overload: a conceptual review and meta-analysis. *J. Consum. Psychol.* **25**(2), 333–358 (2015)
4. Correa, J.R., Foncea, P., Hoeksma, R., Oosterwijk, T., Vredeveld, T.: Posted price mechanisms for a random stream of customers. In: Proceedings of the 2017 ACM Conference on Economics and Computation, EC 2017, Cambridge, MA, USA, 26–30 June 2017, pp. 169–186 (2017). <https://doi.org/10.1145/3033274.3085137>
5. Duetting, P., Feldman, M., Kesselheim, T., Lucier, B.: Prophet inequalities made easy: stochastic optimization by pricing non-stochastic inputs. In: 58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, 15–17 October 2017, pp. 540–551 (2017). <https://doi.org/10.1109/FOCS.2017.56>
6. Ehsani, S., Hajiaghayi, M., Kesselheim, T., Singla, S.: Prophet secretary for combinatorial auctions and matroids. In: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, 7–10 January 2018, pp. 700–714 (2018). <https://doi.org/10.1137/1.9781611975031.46>
7. Esfandiari, H., Hajiaghayi, M., Liaghat, V., Monemizadeh, M.: Prophet secretary. In: Proceedings of the Algorithms - ESA 2015–23rd Annual European Symposium, Patras, Greece, 14–16 September 2015, pp. 496–508 (2015). [https://doi.org/10.1007/978-3-662-48350-3\\_42](https://doi.org/10.1007/978-3-662-48350-3_42)
8. Haynes, G.A.: Testing the boundaries of the choice overload phenomenon: the effect of number of options and time pressure on decision difficulty and satisfaction. *Psychol. Mark.* **26**(3), 204–212 (2009)
9. Hill, T.P., Kertz, R.P.: Comparisons of stop rule and supremum expectations of I.I.D. random variables. *Ann. Probab.* **10**, 336–345 (1982)
10. Kahneman, D., Tversky, A.: Prospect theory: an analysis of decision under risk. In: *Handbook of the Fundamentals of Financial Decision Making: Part I*, pp. 99–127. World Scientific (2013)
11. Kleinberg, J., Kleinberg, R., Oren, S.: Optimal stopping with behaviorally biased agents: the role of loss aversion and changing reference points. arXiv preprint [arXiv:2106.00604](https://arxiv.org/abs/2106.00604) (2021)
12. Kleinberg, R., Weinberg, S.M.: Matroid prophet inequalities. In: Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, 19–22 May 2012, pp. 123–136 (2012). <https://doi.org/10.1145/2213977.2213991>
13. Krengel, U., Sucheston, L.: On semiamarts, amarts, and processes with finite value. *Adv. Probab. Relat. Top.* **4**, 197–266 (1978)
14. Park, J.Y., Jang, S.S.: Confused by too many choices? Choice overload in tourism. *Tour. Manag.* **35**, 1–12 (2013)
15. Rubinstein, A.: Beyond matroids: secretary problem and prophet inequality with general constraints. In: Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, 18–21 June 2016, pp. 324–332 (2016). <https://doi.org/10.1145/2897518.2897540>
16. Rubinstein, A., Singla, S.: Combinatorial prophet inequalities. In: Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, 16–19 January 2017, pp. 1671–1687 (2017). <https://doi.org/10.1137/1.9781611974782.110>
17. Sagi, A., Friedland, N.: The cost of richness: the effect of the size and diversity of decision sets on post-decision regret. *J. Pers. Soc. Psychol.* **93**(4), 515 (2007)

18. Samuel-Cahn, E.: Comparison of threshold stop rules and maximum for independent nonnegative random variables. *Ann. Probab.* **12**(4), 1213–1216 (1984)
19. Schunk, D., Winter, J.: The relationship between risk attitudes and heuristics in search tasks: a laboratory experiment. *J. Econ. Behav. Organ.* **71**(2), 347–360 (2009)
20. Schwartz, B.: *The Paradox of Choice: Why More is Less*. Ecco, New York (2004)
21. Simonson, I., Tversky, A.: Choice in context: tradeoff contrast and extremeness aversion. *J. Mark. Res.* **29**(3), 281–295 (1992)
22. Tversky, A., Kahneman, D.: Loss aversion in riskless choice: a reference-dependent model. *Q. J. Econ.* **106**(4), 1039–1061 (1991)
23. Yan, Q.: Mechanism design via correlation gap. In: Randall, D. (ed.) *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011*, San Francisco, California, USA, 23–25 January 2011, pp. 710–719. SIAM (2011). <https://doi.org/10.1137/1.9781611973082.56>