

Optimal Stopping with Multi-Dimensional Comparative Loss Aversion

Linda Cai* Joshua Gardner† S. Matthew Weinberg‡

Abstract

Motivated by behavioral biases in human decision makers, recent work by Kleinberg et al. [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 λ 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 λ , 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 Kleinberg et al. [11]. For example, Kleinberg et al. [11] gives a tight closed-form on the competitive ratio that an online decision-maker can achieve as a function of λ , 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.

*Princeton University, tcai@princeton.edu.

†Princeton University, jgardner2051@gmail.com.

‡Princeton University, smweinberg@princeton.edu. Supported by NSF CAREER CCF-1942497 and a Google Research Award.

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 Kleinberg et al. [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 Kleinberg et al. [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 Schunk and Winter [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 Kleinberg et al. [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 Brenner et al. [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 [8, 3, 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 Sagi and Friedland [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 $\vec{\sigma}^{(i)} = (\sigma_1^{(i)}, \dots, \sigma_k^{(i)})$ 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, $\|\vec{\sigma}^{(i)}\|_1$. 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 $\vec{s}^{(i)}$ at time i based on options seen so far. The agent suffers from loss aversion parameterized by λ , and experiences (and anticipates) a loss equal to $\lambda \cdot (\|\vec{\sigma}^{(i)}\|_1 - \|\vec{s}^{(i)}\|_1)$ upon selecting candidate i . The key distinction between our model and that of [11] is in how the reference point is built. In Kleinberg et al. [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 ℓ , 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 λ as a function of k when comparing the utility of a biased gambler to that of an unbiased gambler or unbiased prophet. Specifically, Kleinberg et al. [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 λ 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

section 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 4.4, 4.5 and 4.6. For the setting where feature-amplified bias is at least one, the main results can be found in theorem 4.3.

No Improvement for Random Order, Optimal Order, and I.I.D. Distributions. One key finding of Kleinberg et al. [11] is that the competitive ratio of a biased gambler compared to an unbiased gambler/prophet drastically improves (the loss is logarithmic in λ instead of linear in λ) 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.¹ The formal statements of our results for the i.i.d setting can be found in theorems 4.17 and 4.19.

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 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 Krengel and Sucheston [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. Samuel-Cahn [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 (Kleinberg and Weinberg [12], Rubinstein [15], Rubinstein and Singla [16], Duetting et al. [5]), and a variety of arrival order constraints (Esfandiari et al. [7], Ehsani et al. [6], Yan [23], Beyhaghi et al. [1], Hill and Kertz [9], Correa et al. [4]).

Although optimal stopping may have direct applications in human decision making, very little research has devoted to behavioral models of optimal stopping. Recently, Kleinberg et al. [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.

¹Of 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.

1.3 Organization.

The rest of the paper will be organized as follows. In section 2, we discuss our model that incorporates comparative loss aversion and introduce notations that will be used throughout the paper. In section 3, we present an illustrative example that highlights the difference between our model and the single dimensional model presented in [11]. In section 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 Section 5 we discuss the limitations of our model and possible directions for future work. In Appendix A and Appendix B, we provide proofs that are omitted from Section 4. Lastly, in Appendix C we present monotonicity results which were all analogous to those of the one dimensional case presented in Kleinberg et al. [11].

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 σ 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 σ 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 Kleinberg et al. [11], let the parameter $\lambda \geq 0$ serve as a multiplier representing our agent’s loss aversion. In particular, let λ 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 $\vec{s}^{(t)}$ and the j^{th} entry of the super candidate at step t as $\vec{s}_j^{(t)}$. We define the super candidate in step t element-wise as follows:

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

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

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

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 :

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

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 $\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]$ and $\mathbb{E}_{\mathcal{F}}[U_{g_r}^*]$ 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 $\mathbb{E}_{\mathcal{F}}[U_{p_b}]$ and $\mathbb{E}_{\mathcal{F}}[U_{p_r}]$ 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 2.1 (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 2.2 (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 2.3 (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.

Claim 2.4. 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 [13, 12] that the expected value from a value-optimal online algorithm is a $\frac{1}{2}$ approximation of the expected value from an offline prophet. \square

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 σ 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 σ 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 σ a priori for this example. The rational gambler is essentially the same as the rational prophet in this case since they know σ 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 ≤ 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 $\mathbb{E}_{\mathcal{F}}[U_{g_b}^*] \leq 1$. As a result, the biased gambler's prophet utility ratio, $\frac{\mathbb{E}_{\mathcal{F}}[U_{p_r}]}{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]} = \frac{U_{p_r}(\sigma)}{U_{g_b}(\sigma)}$, 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 ∞ 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 ∞ .

As we will soon show in [Proposition 4.1](#) and [Proposition 4.2](#), our present adversarial example with candidates that are excellent on alternating dimensions can be generalized to any k and λ 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 4.15](#) and [Proposition 4.16](#), 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 σ 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 λ 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 λ .

Missing proofs and the complete version of the proof sketches in [Section 4.1](#) and [Section 4.2](#) can be found in [Appendix A](#) and [Appendix B](#) respectively.

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 $\lambda(k-1) = 1$, we obtain a more conservative result that prophet and online utility ratio can be linear in terms of n .

Proposition 4.1. *When the feature-amplified bias $\lambda \cdot (k-1) > 1$, for any $n \in \mathbb{N}^+$, there exists a sequence σ with n candidates such that when the distribution $\mathcal{F}_0(n)$ takes the values of σ 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 $\Omega\left((\lambda(k-1))^{\frac{n}{k}-1}\right)$.*

Proof Sketch. The proposition can be proven by considering the following deterministic distribution, where $\beta = \lambda \cdot (k-1)$:

$$\begin{array}{cccccc} (1, 0, \dots, 0), & (0, 1, 0, \dots, 0), & \dots, & \dots, & (0, \dots, 0, 1), \\ (\beta, 0, \dots, 0), & (0, \beta, 0, \dots, 0), & \dots, & \dots, & (0, \dots, 0, \beta), \\ (\beta^2, 0, \dots, 0), & (0, \beta^2, 0, \dots, 0), & \dots, & \dots, & (0, \dots, 0, \beta^2), \\ & & \dots & & \\ (\beta^{\lceil n/k \rceil - 1}, 0, \dots, 0), & (0, \beta^{\lceil n/k \rceil - 1}, 0, \dots, 0), & \dots, & (0, \dots, 0, \beta^{\lceil n/k \rceil - 1}, 0, \dots, 0). \end{array}$$

While the rational prophet and gambler are able to get utility $\beta^{\lceil n/k \rceil - 1}$ by selecting an option in the last row, the biased gambler could only get utility 1, in particular by selecting the first option. \square

Proposition 4.2. *When the feature-amplified bias $\lambda \cdot (k-1) = 1$, for any $n \in \mathbb{N}^+$, there exists a sequence σ with n candidates such that when the distribution $\mathcal{F}_0(n)$ takes the values of σ 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 $\Omega\left(\frac{n}{k}\right)$.*

Proof Sketch. The proposition can be proven by considering the following deterministic distribution:

$$\begin{array}{cccccc} (1, 0, \dots, 0), & (0, 1, 0, \dots, 0), & \dots, & \dots, & (0, \dots, 0, 1), \\ (2, 0, \dots, 0), & (0, 2, 0, \dots, 0), & \dots, & \dots, & (0, \dots, 0, 2), \\ (3, 0, \dots, 0), & (0, 3, 0, \dots, 0), & \dots, & \dots, & (0, \dots, 0, 3), \\ & & \dots & & \\ (\lceil \frac{n}{k} \rceil, 0, \dots, 0), & (0, \lceil \frac{n}{k} \rceil, 0, \dots, 0), & \dots, & (0, \dots, 0, \lceil \frac{n}{k} \rceil, 0, \dots, 0). \end{array}$$

Here, the rational prophet and gambler are able to get utility $\lceil n/k \rceil$, while the biased gambler could get utility no better than 1. \square

Theorem 4.3. *When the feature-amplified bias $\lambda \cdot (k-1) \geq 1$, the prophet and online competitive ratio are both equal to ∞ .*

Proof. By [Proposition 4.1](#) and [Proposition 4.2](#) we know that when $\lambda \cdot (k-1) \geq 1$, for any $n \in \mathbb{N}^+$, there exists a distribution $\mathcal{F}_0(n)$ such that the prophet and online utility ratio are $\Omega\left(\frac{n}{k}\right)$ (note that $\Omega\left((\lambda(k-1))^{\frac{n}{k}-1}\right) = \omega\left(\frac{n}{k}\right)$). Thus for any $T \in \mathbb{N}$, there exists a $n \in \mathbb{N}^+$ such that the prophet

and online utility of $\mathcal{F}_0(n)$ both exceed T . Since the prophet competitive ratio $\max_{\mathcal{F}} \frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}$ is at least $\max_{n \in \mathbb{N}^+} \frac{\mathbb{E}_{\mathcal{F}_0(n)}[U_{pr}]}{\mathbb{E}_{\mathcal{F}_0(n)}[U_{gb}^*]}$, any $T \in \mathbb{N}$ serves as a lower bound to the prophet utility gap, and thus the prophet competitive ratio is equal to ∞ . Similarly, the online competitive ratio is equal to ∞ . \square

On the other hand, when the feature-amplified bias $\lambda(k-1) < 1$, 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 4.4. *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 4.5. *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 4.6. *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 4.4](#) and [Theorem 4.5](#) follows a similar outline to the single dimensional case presented in Kleinberg et al. [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 4.4](#), we require an additional term when bounding prophet utility ratio (which will later be used in [Theorem 4.5](#)) compared to a similar claim in [11].

We will start by defining threshold based algorithms.

Definition 4.7 (Thresholding algorithms.). *We define a thresholding algorithm parameterized by $\alpha \in (0, 1)$ and denoted A^α as follows. First, A^α 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^α as $\mathbb{E}_{\mathcal{F}}[U_{gb}^\alpha]$.*

We will now reason about the biased gambler's utility from using a thresholding algorithm A^α . 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 σ . This utility bound can then be used to bound the expected utility of the biased gambler given the prior distribution of values.

Claim 4.8. When the feature-amplified bias $\lambda \cdot (k - 1) < 1$, for any sequence σ , 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 $U_{g_b}(\sigma, t) \geq \|\sigma^{(t)}\|_1 - \lambda(k - 1) \cdot T$.

Claim 4.9. When the feature-amplified bias $\lambda \cdot (k - 1) < 1$, for any prior distribution \mathcal{F} and any $\alpha \in (0, 1)$, $\mathbb{E}_{\mathcal{F}}[U_{g_b}^\alpha] \geq ((1 + \lambda)\alpha - k\lambda) \cdot T + (1 - \alpha) \cdot \sum_{t \in [n]} \mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t}[(\|\sigma^{(t)}\|_1 - T)^+]$.

Proof Sketch. Let T be the threshold where $\alpha = \Pr[V^* \geq T]$, then T is the threshold used by the algorithm A^α . Observe that when the gambler uses the algorithm A^α , with probability α the algorithm will select some item, and with probability $1 - \alpha$ the algorithm will select nothing.

Conditioned on A^α selecting no item, we know that all items have value $< T$. Thus even at the end of the algorithm, the gambler's value for the super candidate must not exceed $k \cdot T$. This means that the gambler's utility is at least $-\lambda \cdot k \cdot T$. On the other hand, we can prove by using [Claim 4.8](#) that conditioned on A^α selecting an item at step t , the gambler's utility is at least $(1 - \lambda(k - 1)) \cdot T + \frac{\mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t}[(\|\sigma^{(t)}\|_1 - T)^+]}{\Pr[\|\sigma^{(t)}\|_1 \geq T]}$. The proof of the claim then follows a standard argument for prophet inequality where we take the expectation of agent utility over all events (for all t : selecting an item at step t ; selecting no item). \square

We can now bound the utility ratio between a biased gambler using algorithm A^α and the rational prophet, which can then be used to prove [Theorem 4.4](#) and [Theorem 4.5](#). It turns out that [Theorem 4.4](#) and [Theorem 4.5](#) requires two different ways of bounding the the utility ratio.

Lemma 4.10. When the feature-amplified bias $\lambda \cdot (k - 1) < 1$, for any prior distribution \mathcal{F} , there exists an α such that $\frac{\mathbb{E}_{\mathcal{F}}[U_{g_b}^\alpha]}{\mathbb{E}_{\mathcal{F}}[U_{pr}]} \geq (1 - \lambda(k - 1)) \cdot \max\left\{\frac{\gamma}{1 + \lambda + k}, \frac{1}{2 + \lambda}\right\}$, 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 $V^* = \max_t \|\sigma^{(t)}\|_1$).

Proof Sketch. The key step in this lemma is to bound the term $\sum_{t \in [n]} \mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t}[(\|\sigma^{(t)}\|_1 - T)^+]$ from

[Claim 4.9](#).

Let us call $(\|\sigma^{(t)}\|_1 - T)^+$ the surplus from the t^{th} candidate. Note that $\sum_{t \in [n]} \mathbb{E}_{\|\sigma^{(t)}\|_1 \sim \mathcal{F}_t}[(\|\sigma^{(t)}\|_1 - T)^+]$ is the sum over the expected surplus from each candidate. Thus, it must be at least the expected surplus from the candidate with the largest overall value. Moreover, it must be at least the sum over the expected surplus from candidates that have the largest value on each individual dimension. Formally, the following two inequalities always hold:

$$\sum_{t \in [n]} \mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t}[(\|\sigma^{(t)}\|_1 - T)^+] \geq \mathbb{E}_{\mathcal{F}}[(V^* - T)^+] \geq \mathbb{E}_{\mathcal{F}}[V^* - T], \quad (1)$$

and

$$\sum_{t \in [n]} \mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t}[(\|\sigma^{(t)}\|_1 - T)^+] \geq \sum_{j=1}^k \mathbb{E}_{\mathcal{F}}[(S_j^* - T)^+] \geq \mathbb{E}_{\mathcal{F}}\left[\sum_{j=1}^k S_j^*\right] - k \cdot T \geq \gamma \cdot \mathbb{E}_{\mathcal{F}}[V^*] - k \cdot T. \quad (2)$$

[Equation \(1\)](#) and [Equation \(2\)](#) can then be used to bound the utility of the biased gambler, and in particular its relation to the utility of the rational prophet. \square

We are now ready to prove [Theorem 4.4](#) and [Theorem 4.5](#).

Proof. (of [Theorem 4.4](#)) The first line of the theorem follows directly by [Lemma 4.10](#) 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 $\frac{2 + \lambda}{1 - \lambda(k-1)}$. \square

Proof. (of [Theorem 4.5](#)) 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 4.4](#) 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

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

By the definition of Δ , we know that $\mathbb{E}_{\mathcal{F}}[U_{gr}^*] = \frac{1}{\Delta} \cdot \mathbb{E}_{\mathcal{F}}[U_{pr}]$. Similarly, by the definition of γ , $\mathbb{E}_{\mathcal{F}}\left[\sum_{j=1}^k S_j^*\right] = \gamma \cdot \mathbb{E}_{\mathcal{F}}[V^*] = \gamma \cdot \mathbb{E}_{\mathcal{F}}[U_{pr}]$. 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

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

Notice that inside the max bracket, the first term monotonically increases with Δ , while the second term monotonically decreases with Δ , 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 Δ . \square

We remain to discuss [Theorem 4.6](#), 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 Sketch. (of [Theorem 4.6](#)) Let the first $n - 1$ candidates have the following deterministic value:

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

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 ϵ 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 ϵ approaches 0 and n approaches ∞ . Similarly, the rational gambler can simply always select the randomized option and get utility that approaches and $\frac{1 + \lambda}{1 - \lambda(k - 1)}$. \square

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 ([Claim 4.13](#)), 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 4.14](#)) that for each deterministic distribution (i.e. each distribution such that values drawn from it are equal to some sequence σ 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 [Section 3](#) and [Section 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 4.14](#) 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 4.11. A sequence σ is called *succinct* if there are no two identical candidates in the sequence and no candidate's value is a zero vector.

Definition 4.12. The representation $r(\sigma)$ of a sequence σ is the sequence of unique elements in σ ordered in the ascending order of their first occurrence in σ . (e.g. $r(1, 2, 1, 2, 3) = 1, 2, 3$.)

Claim 4.13. $U_{pr}(\sigma) = U_{pr}(r(\sigma))$ and $U_{gb}(\sigma) = U_{gb}(r(\sigma))$.

Proof. It is clear that $U_{pr}(\sigma) = U_{pr}(r(\sigma))$ since $r(\sigma)$ contains all the unique elements in σ . Let a be an arbitrary candidate in σ , let I_a be the set of all indices t where $\sigma^{(t)} = a$ and let t_a be the minimum index in I_a . Let σ' be the same sequence as σ except for $\|\sigma^{(t)}\|_1 = 0$ for all $t \in I_a \setminus \{t_a\}$. We will show that $U_{gb}(\sigma') = U_{gb}(\sigma)$.

Firstly, for any index t , the super candidate $s^{(t)}$ is only affected by the set of unique candidates at indices $< t$. Therefore $s^{(t)}$ takes the same value in σ and σ' for any $t \notin I_a \setminus \{t_a\}$. As a result, $U_{gb}(\sigma, t) = U_{gb}(\sigma', t)$ for any $t \notin I_a \setminus \{t_a\}$.

Also, when seeing the same candidate value, having seen more candidates in the past can only decrease the biased gambler's utility. Thus for any $t \in I_a$, $U_{gb}(\sigma, t_a) \geq U_{gb}(\sigma, t)$. We conclude that $U_{gb}(\sigma) = \max_{t \in [n]} U_{gb}(\sigma, t) = \max_{t \in [n]} U_{gb}(\sigma', t) = U_{gb}(\sigma')$.

By applying the above construction repeatedly, we can arrive at a sequence σ^* where the only non zero entries are the first occurrence of each candidate in σ . $r(\sigma)$ is simply the result of removing candidates that have value 0 from σ^* . Clearly, this operation does not affect the biased gambler's maximum utility. Thus $U_{gb}(r(\sigma)) = U_{gb}(\sigma^*) = U_{gb}(\sigma)$. □

Lemma 4.14. For any $\epsilon \in (0, 1)$, any succinct sequence σ 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 $\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq (1 - \epsilon) \cdot \frac{U_{pr}(\sigma)}{U_{gb}(\sigma)}$.

Proof. Given an $x \in (0, 1)$ that we will set later, consider a distribution \mathcal{D} that takes value σ_i (i^{th} candidate in the sequence σ) with probability $\frac{x^{i-1}(1-x)}{1-x^m}$ for all $i \in [m]$. First we will verify that \mathcal{D} is a valid distribution (namely, the probability sums up to be 1):

$$\sum_{i=1}^m \frac{x^{i-1}(1-x)}{1-x^m} = \frac{1-x}{1-x^m} \sum_{i=1}^m x^{i-1} = \frac{1-x^m}{1-x} \cdot \frac{1-x}{1-x^m} = 1.$$

Now, let $\mathcal{F} = \times_{i \in [n]} \mathcal{D}$, where $n = m^{\alpha(m-1)} \cdot \log^\alpha(m)$ for some $\alpha > 1$ that we will set later. Let $\tilde{\sigma}$ be a sequence drawn from \mathcal{F} , from [Claim 4.13](#) we know that as long as $r(\tilde{\sigma}) = \sigma$, then $U_{pr}(\tilde{\sigma}) = U_{pr}(\sigma)$ and $U_{gb}(\tilde{\sigma}) = U_{gb}(\sigma)$. Furthermore, by construction the maximum value in a sequence $\tilde{\sigma}$ that is drawn from \mathcal{F} is at most the maximum value in σ , consequently, $U_{gb}(\tilde{\sigma}) \leq U_{pr}(\tilde{\sigma}) \leq U_{pr}(\sigma)$.

Therefore

$$\begin{aligned}\frac{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]}{\mathbb{E}_{\mathcal{F}}[U_{p_r}]} &\leq \frac{\Pr_{\tilde{\sigma} \sim \mathcal{F}}[r(\tilde{\sigma}) = \sigma] \cdot U_{g_b}(\sigma) + \Pr_{\tilde{\sigma} \sim \mathcal{F}}[r(\tilde{\sigma}) \neq \sigma] \cdot U_{p_r}(\sigma)}{\Pr_{\tilde{\sigma} \sim \mathcal{F}}[r(\tilde{\sigma}) = \sigma] \cdot U_{p_r}(\sigma)} \\ &\leq \frac{U_{g_b}(\sigma)}{U_{p_r}(\sigma)} + \frac{\Pr_{\tilde{\sigma} \sim \mathcal{F}}[r(\tilde{\sigma}) \neq \sigma]}{\Pr_{\tilde{\sigma} \sim \mathcal{F}}[r(\tilde{\sigma}) = \sigma]}.\end{aligned}$$

Now we want to bound the probability that $r(\tilde{\sigma}) \neq \sigma$. There are two ways for this to happen, either $\tilde{\sigma}$ does not contain all candidates in σ , or $r(\tilde{\sigma})$ is ordered differently than σ . Formally, by union bound,

$$\Pr_{\tilde{\sigma} \sim \mathcal{F}}[r(\tilde{\sigma}) \neq \sigma] \leq \sum_{i \in [m]} \Pr_{\tilde{\sigma} \sim \mathcal{F}}[\sigma_i \notin r(\tilde{\sigma})] + \sum_{i \in [m-1]} \Pr_{\tilde{\sigma} \sim \mathcal{F}}[\sigma_{i+1} \text{ occur before } \sigma_i \text{ in } r(\tilde{\sigma}) \mid \sigma_i, \sigma_{i+1} \in r(\tilde{\sigma})].$$

Firstly, we bound the probability that $\tilde{\sigma}$ does not contain σ_i as follows

$$\Pr_{\tilde{\sigma} \sim \mathcal{F}}[\sigma_i \notin r(\tilde{\sigma})] = \left(1 - \frac{x^{i-1}(1-x)}{1-x^m}\right)^n \leq \left(1 - \frac{x^{m-1}(1-x)}{1-x^m}\right)^n \leq (1-x^{m-1})^n.$$

Next, let $r_{i,i+1}(\tilde{\sigma})$ denote the smallest index j such that $\tilde{\sigma}_j = \sigma_i$ or σ_{i+1} .

$$\begin{aligned}\Pr_{\tilde{\sigma} \sim \mathcal{F}}[\sigma_{i+1} \text{ occur before } \sigma_i \text{ in } r(\tilde{\sigma}) \mid \sigma_i, \sigma_{i+1} \in r(\tilde{\sigma})] &= \Pr_{\tilde{\sigma} \sim \mathcal{F}}[\tilde{\sigma}_{r_{i,i+1}(\tilde{\sigma})} = \sigma_{i+1} \mid \sigma_i, \sigma_{i+1} \in r(\tilde{\sigma})] \\ &= \Pr_{v \sim \mathcal{D}}[v = \sigma_{i+1} \mid v \in \{\sigma_{i+1}, \sigma_i\}] \\ &= \frac{\frac{x^{i-1}(1-x)}{1-x^m}}{\frac{x^{i-1}(1-x)}{1-x^m} + \frac{x^i(1-x)}{1-x^m}} = \frac{x}{x+1}.\end{aligned}$$

Now we set $x = m^{-\alpha}$, and conclude that

$$\begin{aligned}\Pr_{\tilde{\sigma} \sim \mathcal{F}}[r(\tilde{\sigma}) \neq \sigma] &\leq m \cdot (1-x^{m-1})^n + (m-1) \cdot \frac{x}{x+1} \\ &\leq m \cdot \left(1 - m^{-\alpha(m-1)}\right)^{m^{\alpha(m-1)} \cdot \log^{\alpha}(m)} + (m-1) \cdot m^{-\alpha} \\ &\leq m \cdot m^{-\alpha} + (m-1) \cdot m^{-\alpha} \\ &\leq m^{-(\alpha-1)} + m^{-(\alpha-1)}.\end{aligned}$$

By choosing $\alpha = \log_m\left(\frac{U_{p_r}(\sigma)}{\epsilon \cdot U_{g_b}(\sigma)}\right) + 2$, the above entity is $< \epsilon \cdot \frac{U_{g_b}(\sigma)}{U_{p_r}(\sigma)}$, thus

$$\frac{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]}{\mathbb{E}_{\mathcal{F}}[U_{p_r}]} \leq \frac{U_{g_b}(\sigma)}{U_{p_r}(\sigma)} + \frac{\Pr_{\tilde{\sigma} \sim \mathcal{F}}[r(\tilde{\sigma}) \neq \sigma]}{\Pr_{\tilde{\sigma} \sim \mathcal{F}}[r(\tilde{\sigma}) = \sigma]} \leq \frac{U_{g_b}(\sigma)}{U_{p_r}(\sigma)} + \frac{\epsilon \cdot \frac{U_{g_b}(\sigma)}{U_{p_r}(\sigma)}}{1 - \epsilon \cdot \frac{U_{g_b}(\sigma)}{U_{p_r}(\sigma)}}.$$

Since $\frac{U_{g_b}(\sigma)}{U_{p_r}(\sigma)}$ is at most 1, we conclude that $\frac{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]}{\mathbb{E}_{\mathcal{F}}[U_{p_r}]} \leq \left(1 + \frac{\epsilon}{1-\epsilon}\right) \cdot \frac{U_{g_b}(\sigma)}{U_{p_r}(\sigma)}$. Consequently,

$$\frac{\mathbb{E}_{\mathcal{F}}[U_{p_r}]}{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]} \geq (1-\epsilon) \cdot \frac{U_{p_r}(\sigma)}{U_{g_b}(\sigma)}.$$

□

We can now use [Lemma 4.14](#) 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 ∞ .

Proposition 4.15. *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 $\Omega\left((\lambda(k-1))^{f(\lambda,k) \cdot \log^{\frac{1}{2}} n}\right)$, where $f(\lambda, k)$ is a function that only depends on λ and k .*

Proof Sketch. Let $\sigma(m)$ be the adversarial sequence considered in [Proposition 4.1](#) of length m . Then $\frac{U_{pr}(\sigma(m))}{U_{gb}(\sigma(m))} = \Omega\left((\lambda(k-1))^{\frac{m}{k}-1}\right)$. By construction $\sigma(m)$ contains distinct entries, and thus is succinct. Thus by [Lemma 4.14](#), we can construct a randomized and i.i.d distribution \mathcal{F} with $n = m^{\alpha(m-1)} \cdot \log^{\alpha}(m)$ candidates where $\alpha = \log_m\left(\frac{(\lambda(k-1))^{\frac{m}{k}-1}}{\epsilon}\right) + 2$ that has a similar prophet and online utility ratio (within ϵ multiplicative difference). A careful relation between n and m then proves the claim. \square

Proposition 4.16. *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 $\Omega\left(\frac{\log n}{k \cdot \log \log n}\right)$.*

Proof Sketch. The proof follows similarly to [Proposition 4.16](#), except that we consider the adversarial sequence $\sigma(m)$ from [Proposition 4.2](#), which has a prophet and online utility ratio of $\Omega\left(\frac{m}{k}\right)$. \square

Theorem 4.17. *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 ∞ .*

On the other hand, when the feature-amplified bias $\lambda(k-1) < 1$, we can utilize [Lemma 4.14](#) to show a lower bound for the prophet and online competitive ratios that matches those in the adversarial setting within constant factors.

Claim 4.18. *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 $\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq (1 - \delta) \cdot \frac{1}{(1 - \lambda(k-1))}$.*

Proof Sketch. We will consider a sequence of length m that is similar to the mostly deterministic distribution considered in [theorem 4.6](#), except that the last candidate, which has randomized value, is excluded. Specifically, let σ be the following sequence:

$$\begin{array}{lll} (1, 0, \dots, 0), & (0, 1, 0, \dots, 0), & \dots, (0, \dots, 0, 1), \\ (1 + \beta, 0, \dots, 0), & (0, 1 + \beta, 0, \dots, (1 - \lambda(k-1))(1 + \delta)s, 0), & \dots, (0, \dots, 0, 1 + \beta), \\ (1 + \beta + \beta^2, 0, \dots, 0), & (0, 1 + \beta + \beta^2, 0, \dots, 0), & \dots, (0, \dots, 0, 1 + \beta + \beta^2), \\ \dots & \dots & \dots \\ (\sum_{i=0}^{\lceil \frac{m}{k} \rceil - 1} \beta^i, 0, \dots, 0), & (0, \sum_{i=0}^{\lceil \frac{m}{k} \rceil - 1} \beta^i, 0, \dots, 0), & \dots, (0, \dots, 0, \sum_{i=0}^{\lceil \frac{m}{k} \rceil - 1} \beta^i). \end{array}$$

For a deterministic distribution equal to σ , the prophet utility ratio is $\frac{1}{1-\lambda(k-1)}$. We will then apply [Lemma 4.14](#) to obtain an i.i.d distribution with a larger number of candidates that has a similar prophet utility ratio. \square

Theorem 4.19. *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 $\frac{1}{4} \cdot \frac{1+\lambda}{(1-\lambda(k-1))}$.*

Proof. By [Claim 4.18](#), there exists an i.i.d prior distribution \mathcal{F} such that $\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq (1-\delta) \cdot \frac{1}{(1-\lambda(k-1))}$. Combining this with [Claim 2.4](#), we have $\frac{\mathbb{E}_{\mathcal{F}}[U_{gr}^*]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq \frac{(1-\delta)}{2} \cdot (1-\lambda(k-1))$.

Since $\lambda(k-1) < 1$ and $k > 1$, it must be the case that $2+\lambda < 3$ and that $1+\lambda < 2$. Hence $\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq \frac{(1-\delta)}{3} \cdot \frac{2+\lambda}{1-\lambda(k-1)}$ and $\frac{\mathbb{E}_{\mathcal{F}}[U_{gr}^*]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq \frac{(1-\delta)}{4} \cdot \frac{1+\lambda}{1-\lambda(k-1)}$. This implies that the upper bound matches the lower bounds proved for adversarial setting within constant factors. \square

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] Hedyeh Beyhaghi, Negin Golrezaei, Renato Paes Leme, Martin Pal, and Balasubramanian Sivan. Improved approximations for free-order prophets and second-price auctions. *CoRR*, abs/1807.03435, 2018. URL <http://arxiv.org/abs/1807.03435>. 4
- [2] Lyle Brenner, Yuval Rottenstreich, and Sanjay Sood. Comparison, grouping, and preference. *Psychological Science*, 10(3):225–229, 1999. 2
- [3] Alexander Chernev, Ulf Böckenholt, and Joseph Goodman. Choice overload: A conceptual review and meta-analysis. *Journal of Consumer Psychology*, 25(2):333–358, 2015. 2

- [4] José R. Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Posted price mechanisms for a random stream of customers. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, pages 169–186, 2017. doi: 10.1145/3033274.3085137. URL <https://doi.org/10.1145/3033274.3085137>. 4
- [5] Paul Duetting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. 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, October 15-17, 2017*, pages 540–551, 2017. doi: 10.1109/FOCS.2017.56. URL <https://doi.org/10.1109/FOCS.2017.56>. 4
- [6] Soheil Ehsani, MohammadTaghi Hajiaghayi, Thomas Kesselheim, and Sahil Singla. 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, January 7-10, 2018*, pages 700–714, 2018. doi: 10.1137/1.9781611975031.46. URL <https://doi.org/10.1137/1.9781611975031.46>. 4
- [7] Hossein Esfandiari, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Morteza Monemizadeh. Prophet secretary. In *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*, pages 496–508, 2015. doi: 10.1007/978-3-662-48350-3_42. URL https://doi.org/10.1007/978-3-662-48350-3_42. 4
- [8] Graeme A Haynes. Testing the boundaries of the choice overload phenomenon: The effect of number of options and time pressure on decision difficulty and satisfaction. *Psychology & Marketing*, 26(3):204–212, 2009. 2
- [9] Theodore P. Hill and Robert P. Kertz. Comparisons of stop rule and supremum expectations of i.i.d. random variables. *The Annals of Probability*, 10:336–345, 1982. 4
- [10] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. In *Handbook of the fundamentals of financial decision making: Part I*, pages 99–127. World Scientific, 2013. 2
- [11] Jon Kleinberg, Robert Kleinberg, and Sigal Oren. Optimal stopping with behaviorally biased agents: The role of loss aversion and changing reference points. *arXiv preprint arXiv:2106.00604*, 2021. 1, 2, 3, 4, 5, 9, 11, 18, 30, 32
- [12] Robert Kleinberg and S. Matthew Weinberg. Matroid prophet inequalities. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 123–136, 2012. doi: 10.1145/2213977.2213991. URL <https://doi.org/10.1145/2213977.2213991>. 4, 8
- [13] Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. *Advances in Probability and Related Topics*, 4:197–266, 1978. 4, 8
- [14] Jeong-Yeol Park and SooCheong Shawn Jang. Confused by too many choices? choice overload in tourism. *Tourism Management*, 35:1–12, 2013. 2

- [15] Aviad Rubinstein. 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, June 18-21, 2016*, pages 324–332, 2016. doi: 10.1145/2897518.2897540. URL <https://doi.org/10.1145/2897518.2897540>. 4
- [16] Aviad Rubinstein and Sahil Singla. Combinatorial prophet inequalities. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 1671–1687, 2017. doi: 10.1137/1.9781611974782.110. URL <https://doi.org/10.1137/1.9781611974782.110>. 4
- [17] Adi Sagi and Nehemia Friedland. The cost of richness: The effect of the size and diversity of decision sets on post-decision regret. *Journal of personality and social psychology*, 93(4):515, 2007. 2
- [18] Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent non-negative random variables. *Annals of Probability*, 12(4):1213–1216, 1984. 4
- [19] Daniel Schunk and Joachim Winter. The relationship between risk attitudes and heuristics in search tasks: A laboratory experiment. *Journal of Economic Behavior & Organization*, 71(2): 347–360, 2009. 2
- [20] Barry Schwartz. *The paradox of choice: Why more is less*. Ecco New York, 2004. 2
- [21] Itamar Simonson and Amos Tversky. Choice in context: Tradeoff contrast and extremeness aversion. *Journal of marketing research*, 29(3):281–295, 1992. 18
- [22] Amos Tversky and Daniel Kahneman. Loss aversion in riskless choice: A reference-dependent model. *The quarterly journal of economics*, 106(4):1039–1061, 1991. 2
- [23] Qiqi Yan. Mechanism design via correlation gap. In Dana Randall, editor, *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 710–719. SIAM, 2011. doi: 10.1137/1.9781611973082.56. URL <https://doi.org/10.1137/1.9781611973082.56>. 4

A Omitted Proofs of Section 4.1

Proposition 4.1. *When the feature-amplified bias $\lambda \cdot (k - 1) > 1$, for any $n \in N^+$, there exists a sequence σ with n candidates such that when the distribution $\mathcal{F}_0(n)$ takes the values of σ 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 $\Omega\left((\lambda(k - 1))^{\frac{n}{k} - 1}\right)$.*

Proof. Let $\beta = \lambda \cdot (k - 1)$, consider a sequence σ with n candidates where the t^{th} candidate takes value $\sigma^{(t)}$ such that for dimension $j = (t \bmod k)$, $\sigma_j^{(t)} = \beta^{\lceil \frac{t}{k} \rceil - 1}$ and for any dimension $j \neq (t \bmod k)$, $\sigma_j^{(t)} = 0$. Below is a more intuitive representation of the sequence σ :

$$\begin{array}{cccccc}
(1, 0, \dots, 0), & (0, 1, 0, \dots, 0), & \dots, & \dots, & (0, \dots, 0, 1), \\
(\beta, 0, \dots, 0), & (0, \beta, 0, \dots, 0), & \dots, & \dots, & (0, \dots, 0, \beta), \\
(\beta^2, 0, \dots, 0), & (0, \beta^2, 0, \dots, 0), & \dots, & \dots, & (0, \dots, 0, \beta^2), \\
& & \dots & & \\
(\beta^{\lceil n/k \rceil - 1}, 0, \dots, 0), & (0, \beta^{\lceil n/k \rceil - 1}, 0, \dots, 0), & \dots, & (0, \dots, 0, \beta^{\lceil n/k \rceil - 1}, 0, \dots, 0).
\end{array}$$

Clearly, if the biased gambler selects the first candidate, they will get utility 1. Further, we will show that the biased gambler gets utility at most 0 whenever it accepts a candidate that is not the first candidate.

Firstly, observe that for any fixed i , among all the candidates in i^{th} row, the biased gambler's utility for the first candidate in i^{th} row is the highest. This is because the gambler has the same value for all candidates in i^{th} row, however, the super candidate's value increases as the gambler sees more items, and thus the utility of the gambler decreases as they see more items in the same row.

Thus it suffices for us to consider the biased gambler's utility when they accept the first candidate in the i^{th} row of the above matrix for some $i > 0$. In this case, the super candidate $s^{((i-1) \cdot k + 1)} = (\beta^{i-1}, \beta^{i-2}, \dots, \beta^{i-2})$. We conclude that

$$\begin{aligned}
U_{g_b}(\sigma, (i-1) \cdot k + 1) &= \|\sigma^{((i-1) \cdot k + 1)}\|_1 - \lambda \cdot \left(\|s^{((i-1) \cdot k + 1)}\|_1 - \|\sigma^{((i-1) \cdot k + 1)}\|_1 \right) \\
&= \beta^{i-1} - \lambda \cdot (k-1) \cdot \beta^{i-2} = \beta^{i-1} - \beta \cdot \beta^{i-2} = 0.
\end{aligned}$$

Thus, under σ , the optimal strategy for the biased gambler is to accept the very first candidate, yielding $\mathbb{E}_{\mathcal{F}}[U_{g_b}^*] = 1$. Meanwhile, the rational prophet's expected utility $\mathbb{E}_{\mathcal{F}}[U_{p_r}]$ is equal to the expected largest value, which is $\beta^{\lceil \frac{n}{k} \rceil - 1} \geq (\lambda(k-1))^{\frac{n}{k} - 1}$. Thus $\frac{\mathbb{E}_{\mathcal{F}}[U_{p_r}]}{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]} = \Omega\left((\lambda(k-1))^{\frac{n}{k} - 1}\right)$.

By claim 2.4 this immediately implies that $\frac{\mathbb{E}_{\mathcal{F}}[U_{g_r}^*]}{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]} = \Omega\left((\lambda(k-1))^{\frac{n}{k} - 1}\right)$ as well. \square

Proposition 4.2. *When the feature-amplified bias $\lambda \cdot (k-1) = 1$, for any $n \in \mathbb{N}^+$, there exists a sequence σ with n candidates such that when the distribution $\mathcal{F}_0(n)$ takes the values of σ deterministically, i.e. with probability 1, and when our prior distribution \mathcal{F} equals $\mathcal{F}_0(n)$, the prophet utility ratio $\frac{\mathbb{E}_{\mathcal{F}}[U_{p_r}]}{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]}$ and the online utility ratio $\frac{\mathbb{E}_{\mathcal{F}}[U_{g_r}^*]}{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]}$ are both equal to $\Omega\left(\frac{n}{k}\right)$.*

Proof. We will consider a sequence σ with n candidates where the t^{th} candidate takes value $\sigma^{(t)}$ such that for dimension $j = (t \bmod k)$, $\sigma_j^{(t)} = \lceil \frac{t}{k} \rceil$ and for any dimension $j \neq (t \bmod k)$, $\sigma_j^{(t)} = 0$. Below is a more intuitive representation of the sequence σ :

We will consider an example similar to the one in claim 4.1, where σ is equal to the following sequence:

$$\begin{array}{cccccc}
(1, 0, \dots, 0), & (0, 1, 0, \dots, 0), & \dots, & \dots, & (0, \dots, 0, 1), \\
(2, 0, \dots, 0), & (0, 2, 0, \dots, 0), & \dots, & \dots, & (0, \dots, 0, 2), \\
(3, 0, \dots, 0), & (0, 3, 0, \dots, 0), & \dots, & \dots, & (0, \dots, 0, 3), \\
& & \dots & & \\
(\lceil \frac{n}{k} \rceil, 0, \dots, 0), & (0, \lceil \frac{n}{k} \rceil, 0, \dots, 0), & \dots, & (0, \dots, 0, \lceil \frac{n}{k} \rceil, 0, \dots, 0).
\end{array}$$

We will show that the biased gambler gets utility at most 1 no matter what candidate they accept.

For the same reason as in claim 4.1 it suffice for us to consider the biased gambler's utility when they accept the first candidate in the i^{th} row of the above matrix. In this case, the super candidate $s^{((i-1) \cdot k+1)} = (i, i-1, \dots, i-1)$. Thus

$$\begin{aligned} U_{g_b}(\sigma, (i-1) \cdot k+1) &= \|\sigma^{((i-1) \cdot k+1)}\|_1 - \lambda \cdot \left(\|s^{((i-1) \cdot k+1)}\|_1 - \|\sigma^{((i-1) \cdot k+1)}\|_1 \right) \\ &= i - \lambda \cdot (k-1) \cdot (i-1) = i - (i-1) = 1. \end{aligned}$$

Thus, under σ , one of the optimal strategies for the biased gambler is to accept the very first candidate, yielding $\mathbb{E}_{\mathcal{F}}[U_{g_b}^*] = U_{g_b}(\sigma, 1) = 1$. On the other hand, the rational prophet's expected utility $\mathbb{E}_{\mathcal{F}}[U_{p_r}]$ is equal to the expected largest value, which is at least $\frac{n}{k}$. Thus $\frac{\mathbb{E}_{\mathcal{F}}[U_{p_r}]}{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]} = \Omega(\frac{n}{k})$.

By claim 2.4, this immediately implies that $\frac{\mathbb{E}_{\mathcal{F}}[U_{g_r}^*]}{\mathbb{E}_{\mathcal{F}}[U_{g_b}^*]} = \Omega(\frac{n}{k})$ as well. \square

Claim 4.8. *When the feature-amplified bias $\lambda \cdot (k-1) < 1$, for any sequence σ , 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 $U_{g_b}(\sigma, t) \geq \|\sigma^{(t)}\|_1 - \lambda(k-1) \cdot T$.*

Proof. We know by definition that

$$\begin{aligned} U_{g_b}(\sigma, t) &= \|\sigma^{(t)}\|_1 - \lambda \cdot (\|\bar{s}^{(t)}\|_1 - \|\sigma^{(t)}\|_1) \\ &= \|\sigma^{(t)}\|_1 - \lambda \cdot (\|\bar{s}^{(t)}\|_1 - \|\sigma^{(t)}\|_1) = (1 + \lambda) \cdot \|\sigma^{(t)}\|_1 - \lambda \cdot \|\bar{s}^{(t)}\|_1. \end{aligned}$$

Since t is the smallest timestep where $\|\sigma^{(t)}\|_1 \geq T$, we know that $\|\sigma^{(t')}\|_1 < T$ for any $t' < t$. Thus, the value at any particular index of the super candidate at step $t-1$ is at most T . As a result, for any j , $\|\bar{s}_j^{(t)}\|_1 \leq \max(\|\sigma_j^{(t)}\|_1, \|\bar{s}_j^{(t-1)}\|_1) \leq \max(\|\sigma_j^{(t)}\|_1, T)$. This implies that

$$\|\bar{s}^{(t)}\|_1 = \sum_j \bar{s}_j^{(t)} \leq \sum_j \max(\sigma_j^{(t)}, T) = \sum_j T + \sum_j (\sigma_j^{(t)} - T)^+.$$

Observe that since $\|\sigma^{(t)}\|_1 - T \geq 0$, in the case where for all j , $\|\sigma^{(t)}\|_1 - T < 0$, $\sum_j (\sigma_j^{(t)} - T)^+ = 0 \leq \|\sigma^{(t)}\|_1 - T$. On the other hand, when for at least one dimension j^* , $\|\sigma^{(t)}\|_1 - T \geq 0$, we can write $\sum_j (\sigma_j^{(t)} - T)^+ = \sigma_{j^*}^{(t)} - T + \sum_{j \neq j^*} (\sigma_j^{(t)} - T)^+$, which is at most $\sigma_{j^*}^{(t)} - T + \sum_{j \neq j^*} \sigma_j^{(t)} = \|\sigma^{(t)}\|_1 - T$. We thus conclude that

$$\|\bar{s}^{(t)}\|_1 \leq k \cdot T + (\|\sigma^{(t)}\|_1 - T) = \|\sigma^{(t)}\|_1 + (k-1) \cdot T.$$

Combining everything together, we get that

$$\begin{aligned} U_{g_b}(\sigma, t) &= (1 + \lambda) \cdot \|\sigma^{(t)}\|_1 - \lambda \cdot \|\bar{s}^{(t)}\|_1 \geq (1 + \lambda) \cdot \|\sigma^{(t)}\|_1 - \lambda \cdot (\|\sigma^{(t)}\|_1 + (k-1) \cdot T) \\ &\geq \|\sigma^{(t)}\|_1 - \lambda(k-1) \cdot T. \end{aligned}$$

\square

Claim 4.9. *When the feature-amplified bias $\lambda \cdot (k-1) < 1$, for any prior distribution \mathcal{F} and any $\alpha \in (0, 1)$, $\mathbb{E}_{\mathcal{F}}[U_{g_b}^*] \geq ((1 + \lambda)\alpha - k\lambda) \cdot T + (1 - \alpha) \cdot \sum_{t \in [n]} \mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t} [(\|\sigma^{(t)}\|_1 - T)^+]$.*

Proof. Let T be the threshold where $\alpha = \Pr[V^* \geq T]$, then T is the threshold used by the algorithm A^α . Observe that when the gambler uses the algorithm A^α , with probability α the algorithm will select some item, and with probability $1 - \alpha$ the algorithm will select nothing.

Conditioned on A^α selecting no item, we know that all items have value $< T$. Thus even at the end of the algorithm, the gambler's value for the super candidate must not exceed $k \cdot T$. This means that the gambler's utility is at least $-\lambda \cdot k \cdot T$.

On the other hand, we will prove by using claim 4.8 that conditioned on A^α selecting an item at step t , the gambler's utility is at least $\|\sigma^{(t)}\|_1 - \lambda(k-1) \cdot T$. We will begin by showing a block of formal deduction, which we subsequently unpack in the following paragraphs.

$$\begin{aligned}
\mathbb{E}_{\mathcal{F}}[U_{g_b}^\alpha \mid A^\alpha \text{ selects item } t] &= \mathbb{E}_{\sigma \sim \mathcal{F}} \left[U_{g_b}(\sigma, t) \mid t = \operatorname{argmin}_{i \in [n]} \{\|\sigma^{(i)}\|_1 \geq T\} \right] \\
&\geq \mathbb{E}_{\sigma \sim \mathcal{F}} \left[\|\sigma^{(t)}\|_1 - \lambda(k-1) \cdot T \mid t = \operatorname{argmin}_{i \in [n]} \{\|\sigma^{(i)}\|_1 \geq T\} \right] \\
&\geq (1 - \lambda(k-1)) \cdot T + \mathbb{E}_{\sigma \sim \mathcal{F}} \left[(\|\sigma^{(t)}\|_1 - T)^+ \mid \|\sigma^{(t)}\|_1 \geq T \right] \\
&\geq (1 - \lambda(k-1)) \cdot T + \mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t} \left[(\|\sigma^{(t)}\|_1 - T)^+ \mid \|\sigma^{(t)}\|_1 \geq T \right] \\
&\geq (1 - \lambda(k-1)) \cdot T + \frac{\mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t} [(\|\sigma^{(t)}\|_1 - T)^+]}{\Pr[\|\sigma^{(t)}\|_1 \geq T]}.
\end{aligned}$$

In the above argument, the first line is due to A^α selecting the first candidate with value at least T . Furthermore, the second line follows from the first line due to claim 4.8. The second line to the third line is because the value of the candidates are independent, and thus “ $\sigma^{(t)}$ is the first candidate that has value at least T ” and “ $\sigma^{(t)}$ has value at least T ” have equivalent implications for the value of $\sigma^{(t)}$. The fourth line follows from the third because our formula only involves $\sigma^{(t)}$ and thus the expectation over σ drawn from \mathcal{F} is equivalent to the expectation over $\sigma^{(t)}$ drawn from \mathcal{F}_t . The fourth line to the fifth line is because when $\|\sigma^{(t)}\|_1 < T$, $(\|\sigma^{(t)}\|_1 - T)^+$ takes the value 0, thus

$$\mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t} [(\|\sigma^{(t)}\|_1 - T)^+] = \mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t} [(\|\sigma^{(t)}\|_1 - T)^+ \mid \|\sigma^{(t)}\|_1 \geq T] \cdot \Pr[\|\sigma^{(t)}\|_1 \geq T].$$

We will now bound the expected utility of the biased gambler when they use the algorithm A^α . The argument is presented in the inequalities immediately below, which will be unpacked in the

following paragraphs.

$$\begin{aligned}
\mathbb{E}_{\mathcal{F}}[U_{g_b}^\alpha] &= \Pr[U_{g_b}^\alpha \text{ selects no item}] \cdot (-k\lambda \cdot T) + \sum_{t \in [n]} \Pr[U_{g_b}^\alpha \text{ selects item } j] \cdot \mathbb{E}[U_{g_b}^\alpha \mid A^\alpha \text{ selects item } j] \\
&\geq (1 - \alpha) \cdot (-k\lambda \cdot T) + \sum_{t \in [n]} \Pr[U_{g_b}^\alpha \text{ selects item } t] \cdot \left((1 - \lambda(k - 1)) \cdot T + \frac{\mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t}[(\|\sigma^{(t)}\|_1 - T)^+]}{\Pr[\|\sigma^{(t)}\|_1 \geq T]} \right) \\
&\geq (1 - \alpha) \cdot (-k\lambda \cdot T) + \alpha \cdot (1 - \lambda(k - 1)) \cdot T \\
&\quad + (1 - \alpha) \cdot \sum_{t \in [n]} \Pr[\|\sigma^{(t)}\|_1 \geq T] \cdot \frac{\mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t}[(\|\sigma^{(t)}\|_1 - T)^+]}{\Pr[\|\sigma^{(t)}\|_1 \geq T]} \\
&\geq ((1 + \lambda)\alpha - k\lambda) \cdot T + (1 - \alpha) \cdot \sum_{t \in [n]} \mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t}[(\|\sigma^{(t)}\|_1 - T)^+].
\end{aligned}$$

For the above equation, the first and second line is just summing over all possible conditions and using inequalities that we established previous. The second line to the third line is because $\sum_{t \in [n]} \Pr[U_{g_b}^\alpha \text{ selects item } j] = \alpha$ and the following inequalities,

$$\begin{aligned}
\Pr[U_{g_b}^\alpha \text{ selects item } t] &= \Pr[\forall i < t, \|\sigma^{(i)}\|_1 < T] \cdot \Pr[\|\sigma^{(t)}\|_1 \geq T] \\
&\geq \Pr[\forall i \in [n], \|\sigma^{(i)}\|_1 < T] \cdot \Pr[\|\sigma^{(t)}\|_1 \geq T] = (1 - \alpha) \cdot \Pr[\|\sigma^{(t)}\|_1 \geq T].
\end{aligned}$$

The third line to the fourth line is simply due to a cancellation of terms. \square

Lemma 4.10. *When the feature-amplified bias $\lambda \cdot (k - 1) < 1$, for any prior distribution \mathcal{F} , there exists an α such that $\frac{\mathbb{E}_{\mathcal{F}}[U_{g_b}^\alpha]}{\mathbb{E}_{\mathcal{F}}[U_{p_r}]} \geq (1 - \lambda(k - 1)) \cdot \max\left\{\frac{\gamma}{1 + \lambda + k}, \frac{1}{2 + \lambda}\right\}$, 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 $V^* = \max_t \|\sigma^{(t)}\|_1$).*

Proof. Let $\alpha \in (0, 1)$ be a fixed value that we will set later. By claim 4.9, we know that

$$\mathbb{E}_{\mathcal{F}}[U_{g_b}^\alpha] \geq ((1 + \lambda)\alpha - k\lambda) \cdot T + (1 - \alpha) \cdot \sum_{t \in [n]} \mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t}[(\|\sigma^{(t)}\|_1 - T)^+].$$

In order to further reason about the above quantity, we need to bound $\sum_{t \in [n]} \mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t}[(\|\sigma^{(t)}\|_1 - T)^+]$,

which is what we will do next.

Let us call $(\|\sigma^{(t)}\|_1 - T)^+$ the surplus from the t^{th} candidate. Note that $\sum_{t \in [n]} \mathbb{E}_{\|\sigma^{(t)}\|_1 \sim \mathcal{F}_t}[(\|\sigma^{(t)}\|_1 - T)^+]$

is the sum over the expected surplus from each candidate. Thus, it must be at least the expected surplus from the candidate with the largest overall value. Moreover, it must be at least the sum over the expected surplus from candidates that have the largest value on each individual dimension. Formally, the following two inequalities always hold:

$$\sum_{t \in [n]} \mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t}[(\|\sigma^{(t)}\|_1 - T)^+] \geq \mathbb{E}_{\mathcal{F}}[(V^* - T)^+] \geq \mathbb{E}_{\mathcal{F}}[V^* - T], \tag{3}$$

and

$$\sum_{t \in [n]} \mathbb{E}_{\sigma^{(t)} \sim \mathcal{F}_t} \left[(\|\sigma^{(t)}\|_1 - T)^+ \right] \geq \sum_{j=1}^k \mathbb{E}_{\mathcal{F}} [(S_j^* - T)^+] \geq \mathbb{E}_{\mathcal{F}} \left[\sum_{j=1}^k S_j^* \right] - k \cdot T \geq \gamma \cdot \mathbb{E}_{\mathcal{F}} [V^*] - k \cdot T. \quad (4)$$

Now, we will use these bounds to derive the final bound on the biased gambler's utility from using algorithm A^α . By combining inequality 3 and claim 4.9, we know that

$$\mathbb{E}_{\mathcal{F}} [U_{gb}^\alpha] \geq ((1 + \lambda)\alpha - k\lambda) \cdot T + (1 - \alpha) \cdot \mathbb{E}_{\mathcal{F}} [V^* - T].$$

When we set $\alpha = \frac{\lambda k + 1}{2 + \lambda}$, $((1 + \lambda)\alpha - k\lambda) = 1 - \alpha = \frac{1 - \lambda(k - 1)}{2 + \lambda}$. This means that for this choice of α , the preceding inequality becomes:

$$\mathbb{E}_{\mathcal{F}} [U_{gb}^\alpha] \geq \frac{1 - \lambda(k - 1)}{2 + \lambda} \cdot (T + \mathbb{E}[V^* - T]) \geq \frac{1 - \lambda(k - 1)}{2 + \lambda} \cdot \mathbb{E}_{\mathcal{F}} [V^*] = \frac{1 - \lambda(k - 1)}{2 + \lambda} \cdot \mathbb{E}_{\mathcal{F}} [U_{pr}].$$

On the other hand, by combining inequality 4 and claim 4.9, we know that

$$\begin{aligned} \mathbb{E}_{\mathcal{F}} [U_{gb}^\alpha] &\geq ((1 + \lambda)\alpha - k\lambda) \cdot T + (1 - \alpha) \cdot (\gamma \cdot \mathbb{E}_{\mathcal{F}} [V^*] - k \cdot T) \\ &\geq ((1 + \lambda)\alpha - k \cdot \lambda - (1 - \alpha) \cdot k) \cdot T + (1 - \alpha) \cdot (\gamma \cdot \mathbb{E}_{\mathcal{F}} [V^*]). \end{aligned}$$

When we set $\alpha = \frac{k \cdot (1 + \lambda)}{1 + \lambda + k}$, $((1 + \lambda)\alpha - k \cdot \lambda - (1 - \alpha) \cdot k) = 0$, therefore for this choice of α ,

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

Combining the above two cases we conclude that

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

which implies our lemma statement. \square

Theorem 4.6. *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 $\frac{1 + \lambda}{1 - \lambda(k - 1)}$.*

Proof. Let $\beta = \lambda \cdot (k - 1)$, and fix some positive integer $w \in \mathbb{N}^+$ and some small constant ϵ that we will set later. In our example there will be $n = w \cdot k + 1$ candidates, and all candidates will be deterministic except for the last candidate, i.e. $\sigma^{(1)} \dots \sigma^{(n-1)}$ are drawn from deterministic $\mathcal{F}_1 \dots \mathcal{F}_{n-1}$ but σ^n is drawn from stochastic \mathcal{F}_n . We define the sequence of deterministic values as follows:

$$\begin{aligned}
& (1, 0, \dots, 0), & (0, 1, 0, \dots, 0), & \dots, & (0, \dots, 0, 1), \\
& (1 + \beta, 0, \dots, 0), & (0, 1 + \beta, 0, \dots, 0), & \dots, & (0, \dots, 0, 1 + \beta), \\
& (1 + \beta + \beta^2, 0, \dots, 0), & (0, 1 + \beta + \beta^2, 0, \dots, 0), & \dots, & (0, \dots, 0, 1 + \beta + \beta^2), \\
& \dots & \dots & \dots & \dots \\
& (\sum_{j=0}^{w-1} \beta^j, 0, \dots, 0), & (0, \sum_{j=0}^{w-1} \beta^j, 0, \dots, 0), & \dots, & (0, \dots, 0, \sum_{j=0}^{w-1} \beta^j).
\end{aligned}$$

The last randomized candidate is equal to $\left((1 - \epsilon)(1 + \lambda) \cdot \sum_{j=0}^{w-1} \beta^j / \epsilon, 0, \dots, 0 \right)$ with probability ϵ and is equal to $(0, 0, \dots, 0)$ with probability $1 - \epsilon$. We will prove that given our prior distribution \mathcal{F} , the gambler's expected utility is always at most 1 no matter which candidate they choose.

Notice that as with similar constructions in claim 4.1 and claim 4.2, if the gambler picks a deterministic candidate in some row i , then their utility is maximized by picking the first candidate in row i , which results in utility

$$\begin{aligned}
\mathbb{E}_{\sigma \sim \mathcal{F}}[U_{g_b}(\sigma, i \cdot k + 1)] &= \|\sigma^{(i \cdot k + 1)}\|_1 - \lambda \cdot \left(\|\bar{s}^{(i \cdot k + 1)}\|_1 - \|\sigma^{(i \cdot k + 1)}\|_1 \right) = \sum_{j=0}^{i-1} \beta^j - \lambda \cdot (k-1) \cdot \sum_{j=0}^{i-2} \beta^j \\
&= \sum_{j=0}^{i-1} (\lambda(k-1))^j - \lambda \cdot (k-1) \cdot \sum_{j=0}^{i-2} (\lambda(k-1))^j = 1.
\end{aligned}$$

Finally, if the gambler picks the last candidate (i.e. the randomized candidate), then their expected utility is

$$\begin{aligned}
\mathbb{E}_{\sigma \sim \mathcal{F}}[U_{g_b}(\sigma, wk + 1)] &= \epsilon \cdot \left((1 - \epsilon)(1 + \lambda) \cdot \sum_{i=0}^{w-1} \beta^i / \epsilon - \lambda \cdot (k-1) \cdot \sum_{i=0}^{w-1} \beta^i \right) + (1 - \epsilon) \cdot \left(-\lambda \cdot k \cdot \sum_{i=0}^{w-1} \beta^i \right) \\
&\leq (1 - \epsilon) \cdot \left((1 + \lambda) \cdot \sum_{i=0}^{w-1} \beta^i - \lambda \cdot k \cdot \sum_{i=0}^{w-1} \beta^i \right) \\
&= (1 - \epsilon) \cdot (1 - \lambda(k-1)) \cdot \frac{1 - (\lambda(k-1))^w}{1 - \lambda(k-1)} = (1 - \epsilon) \cdot (1 - (\lambda(k-1))^w) < 1.
\end{aligned}$$

Therefore it is not possible for the biased gambler to get utility more than 1, consequently, $\mathbb{E}_{\mathcal{F}}[U_{g_b}^*] \leq 1$. On the other hand, the rational prophet can get expected utility

$$\begin{aligned}
\mathbb{E}_{\mathcal{F}}[U_{p_r}] &= \mathbb{E}[V^*] = \epsilon \cdot \|\sigma^{(wk+1)}\|_1 + (1 - \epsilon) \cdot \|\sigma^{(wk)}\|_1 \\
&= \epsilon \cdot \left((1 - \epsilon)(1 + \lambda) \cdot \sum_{i=0}^{w-1} \beta^i / \epsilon \right) + (1 - \epsilon) \cdot \sum_{i=0}^{w-1} \beta^i \\
&= (1 - \epsilon)(2 + \lambda) \cdot \frac{1 - (\lambda(k-1))^w}{1 - \lambda(k-1)}.
\end{aligned}$$

Similarly, the rational gambler can simply always choose the last candidate, thus

$$\mathbb{E}_{\mathcal{F}}[U_{g_r}^*] \geq \epsilon \cdot \left((1 - \epsilon)(1 + \lambda) \cdot \sum_{i=0}^{w-1} \beta^i / \epsilon \right) = (1 - \epsilon)(1 + \lambda) \cdot \frac{1 - (\lambda(k-1))^w}{1 - \lambda(k-1)}.$$

Let ϵ be such that $(1 - \epsilon)^2 = 1 - \delta$ and let $w = \log_{\lambda(k-1)} \epsilon$. Then $(\lambda(k-1))^w = \epsilon$ and thus

$$\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq \frac{(2 + \lambda) \cdot (1 - \epsilon)^2}{1 - \lambda(k-1)} = \frac{2 + \lambda}{1 - \lambda(k-1)} \cdot (1 - \delta).$$

Similarly,

$$\frac{\mathbb{E}_{\mathcal{F}}[U_{gr}^*]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq \frac{(1 + \lambda) \cdot (1 - \epsilon)^2}{1 - \lambda(k-1)} = \frac{1 + \lambda}{1 - \lambda(k-1)} \cdot (1 - \delta).$$

□

B Omitted Proofs of Section 4.2

Proposition 4.15. *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 $\Omega\left((\lambda(k-1))^{f(\lambda,k) \cdot \log^{\frac{1}{2}} n}\right)$, where $f(\lambda, k)$ is a function that only depends on λ and k .*

Proof. The corollary follows in a relatively straightforward manner from claim 4.1 and lemma 4.14. Let $\sigma(m)$ be the sequence considered in claim 4.1 of length m . Then $\frac{U_{pr}(\sigma(m))}{U_{gb}(\sigma(m))} = \Omega\left((\lambda(k-1))^{\frac{m}{k}-1}\right)$, and in particular $\frac{U_{pr}(\sigma(m))}{U_{gb}(\sigma(m))} \geq (\lambda(k-1))^{\frac{m}{k}-1}$. Observe that $\sigma(m)$ is succinct. Thus by lemma 4.14, for any $\epsilon \in (0, 1)$, there exists a distribution \mathcal{F} with $n = m^{\alpha(m-1)} \cdot \log^{\alpha}(m)$ where $\alpha = \log_m\left(\frac{(\lambda(k-1))^{\frac{m}{k}-1}}{\epsilon}\right) + 2$ such that $\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq (1 - \epsilon) \cdot (\lambda(k-1))^{\frac{m}{k}-1}$. Now in order to get a bound in terms of n we only need to establish the relationship between n and m .

Since $n = m^{\alpha(m-1)} \cdot \log^{\alpha}(m) \leq m^{m \cdot \alpha} = e^{\log m \cdot m \cdot \alpha}$, this means that

$$\begin{aligned} \log n &\leq m \cdot \log m \cdot \left(\log_m \left(\frac{(\lambda(k-1))^{\frac{m}{k}-1}}{\epsilon} \right) + 2 \right) \\ &\leq m \cdot \log \left(\frac{(\lambda(k-1))^{\frac{m}{k}-1}}{\epsilon} \right) + 2 \cdot m \cdot \log m \\ &\leq m \cdot \left(\frac{m}{k} - 1 \right) \cdot \log(\lambda(k-1)) + m \cdot \log \left(\frac{1}{\epsilon} \right) + 2m \cdot \log m. \end{aligned}$$

Let $\epsilon = 1/2$, since $m \geq 2$, we know that $\log m \geq 1$, this means that

$$\begin{aligned} \log n &\leq \frac{m^2}{k} \cdot \log(\lambda(k-1)) + m \cdot \log 2 + 2m \log m \\ &\leq \frac{m^2}{k} \cdot \log(\lambda(k-1)) + 4m \log m \\ &\leq 2 \cdot \max \left(\frac{m^2}{k} \cdot \log(\lambda(k-1)), 4m \log m \right). \end{aligned}$$

Now there are two cases:

1. When $\frac{m^2}{k} \cdot \log(\lambda(k-1)) \geq 4m \log m$, we know that $\log n \leq 2 \cdot \frac{m^2}{k} \cdot \log(\lambda(k-1))$. This implies that $m^2 \geq \frac{k \cdot \log n}{2 \cdot \log(\lambda(k-1))}$ and thus $m \geq \frac{\sqrt{k} \cdot \log^{\frac{1}{2}} n}{\sqrt{2} \cdot \log^{1/2}(\lambda(k-1))}$.
2. When $\frac{m^2}{k} \cdot \log(\lambda(k-1)) < 4m \log m$, we know that $\log n \leq 8m \log m \leq 8m^2$, which means that $m \geq \frac{\log^{\frac{1}{2}} n}{2\sqrt{2}}$.

We conclude that no matter which case we are in, $m \geq \frac{1}{\sqrt{2}} \cdot \log^{\frac{1}{2}} n \cdot \min\left(\frac{\sqrt{k}}{\log^{1/2}(\lambda(k-1))}, \frac{1}{2}\right)$, which means that

$$\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq (1 - \epsilon) \cdot \left(\lambda(k-1)\right)^{\frac{m}{k}-1} = \Omega\left(\left(\lambda(k-1)\right)^{\frac{1}{\sqrt{2k}} \cdot \log^{\frac{1}{2}} n \cdot \min\left(\frac{1}{\log^{1/2}(\lambda(k-1))}, \frac{1}{2\sqrt{k}}\right)} - 1\right).$$

Let $f(\lambda, k) = \frac{1}{\sqrt{2k}} \cdot \min\left(\frac{1}{\log^{1/2}(\lambda(k-1))}, \frac{1}{2\sqrt{k}}\right) - 1$, and we get the desired claim that $\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} = \Omega\left(\left(\lambda(k-1)\right)^{f(\lambda, k) \cdot \log^{\frac{1}{2}} n}\right)$. By claim 2.4, $\frac{\mathbb{E}_{\mathcal{F}}[U_{gr}^*]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} = \Omega\left(\left(\lambda(k-1)\right)^{f(\lambda, k) \cdot \log^{\frac{1}{2}} n}\right)$ as well. \square

Proposition 4.16. *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 $\Omega\left(\frac{\log n}{k \cdot \log \log n}\right)$.*

Proof. The corollary follows in a relatively straightforward manner from claim 4.2 and lemma 4.14. Let $\sigma(m)$ be the sequence considered in claim 4.2 of length m . Then $\frac{U_{pr}(\sigma(m))}{U_{gb}(\sigma(m))} = \Omega\left(\frac{m}{k}\right)$, and in particular $\frac{U_{pr}(\sigma(m))}{U_{gb}(\sigma(m))} \geq \frac{m}{k}$. Observe that $\sigma(m)$ is succinct. Thus by lemma 4.14, for any $\epsilon \in (0, 1)$, there exists a distribution \mathcal{F} with $n = m^{\alpha(m-1)} \cdot \log^{\alpha}(m)$ where $\alpha = \log_m\left(\frac{m}{k \cdot \epsilon}\right) + 2$ such that $\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}^*]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq (1 - \epsilon) \cdot \frac{m}{k}$.

By using a similar analysis as in corollary 4.15, we can compute the relationship between n and m as follows:

$$\log n \leq \log\left(m^{\alpha(m-1)} \cdot \log^{\alpha}(m)\right) \leq \log(m^{\alpha m}) = \alpha \cdot m \log m.$$

Now let us set $\epsilon = 1/2$ and $\beta = \alpha - 3$, then

$$\beta = \left(\log_m\left(\frac{m}{k \cdot \epsilon}\right) + 2\right) - (2 + \log_m m) = \log_m\left(\frac{1}{k \cdot \epsilon}\right) = \log_m\left(\frac{2}{k}\right) = O(1).$$

We conclude that $\alpha = 3 + \beta = 3 + O(1) = \Theta(1)$. We will now show that $m = \Omega\left(\frac{\log n}{\log \log n}\right)$.

1. Assume $m \geq \log n$, then it's clear that $m = \Omega\left(\frac{\log n}{\log \log n}\right)$.
2. Assume $m \leq \log n$, then $\log m \leq \log \log n$, since $\log n \leq \alpha \cdot m \log m$, we conclude that

$$m \geq \frac{\log n}{\alpha \log m} \geq \frac{\log n}{\alpha \log \log n} = \Omega\left(\frac{\log n}{\log \log n}\right).$$

Consequently, $\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq (1 - \epsilon) \cdot \frac{m}{k} = \Omega\left(\frac{\log n}{k \cdot \log \log n}\right)$ and by claim 2.4 $\frac{\mathbb{E}_{\mathcal{F}}[U_{gr}^*]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} = \Omega\left(\frac{\log n}{k \cdot \log \log n}\right)$ as well. \square

Theorem 4.17. *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 ∞ .*

Proof. By corollaries 4.15 and 4.16 we know that when $\lambda \cdot (k-1) \geq 1$, for any $n \in \mathbb{N}^+$, there exists an i.i.d distribution $\mathcal{F}_0(n)$ such that the prophet and online utility ratio of $\mathcal{F}_0(n)$ is an unboundedly increasing function in terms of the number of candidates n . Thus for any $T \in \mathbb{N}$, there exists a $n \in \mathbb{N}^+$ such that the prophet and online utility of $\mathcal{F}_0(n)$ both exceed T .

Recall that the formal definition for competitive ratio is slightly different in the i.i.d prior distribution setting in contrast to the adversarial order setting. Namely, the prophet and online competitive ratios are defined as $\max_{\mathcal{F} \in \mathcal{U}} \frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}$ and $\max_{\mathcal{F} \in \mathcal{U}} \frac{\mathbb{E}_{\mathcal{F}}[U_{gr}^*]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}$, where \mathcal{U} is the set of all i.i.d distribution. Observe that the prophet competitive ratio $\max_{\mathcal{F} \in \mathcal{U}} \frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]}$ is at least $\max_{n \in \mathbb{N}^+} \frac{\mathbb{E}_{\mathcal{F}_0(n)}[U_{pr}]}{\mathbb{E}_{\mathcal{F}_0(n)}[U_{gb}^*]}$, thus any $T \in \mathbb{N}$ serves as a lower bound to the prophet competitive ratio, and thus the prophet competitive ratio is equal to ∞ . Similarly, the online competitive ratio is equal to ∞ . \square

Claim 4.18. *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 $\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq (1 - \delta) \cdot \frac{1}{(1 - \lambda(k-1))}$.*

Proof. We will consider a sequence that is similar to the mostly deterministic distribution considered in theorem 4.6, except that the last candidate, which has randomized value, is excluded. Specifically, let $\beta = \lambda \cdot k$, let w be a parameter that we will set later and let σ be the following sequence:

$$\begin{array}{lll} (1, 0, \dots, 0), & (0, 1, 0, \dots, 0), & \dots, (0, \dots, 0, 1), \\ (1 + \beta, 0, \dots, 0), & (0, 1 + \beta, 0, \dots, (1 - \lambda(k-1))(1 + \delta)s, 0), & \dots, (0, \dots, 0, 1 + \beta), \\ (1 + \beta + \beta^2, 0, \dots, 0), & (0, 1 + \beta + \beta^2, 0, \dots, 0), & \dots, (0, \dots, 0, 1 + \beta + \beta^2), \\ \dots & \dots & \dots \\ (\sum_{i=0}^{w-1} \beta^i, 0, \dots, 0), & (0, \sum_{i=0}^{w-1} \beta^i, 0, \dots, 0), & \dots, (0, \dots, 0, \sum_{i=0}^{w-1} \beta^i). \end{array}$$

As reasoned in theorem 4.6, the biased gambler has utility ≤ 1 when selecting any candidate in this sequence. However, the rational prophet will pick a candidate in the last row, and gains utility $\sum_{i=0}^{w-1} \beta^i = \frac{1 - \beta^w}{1 - \beta}$. Thus,

$$\frac{U_{pr}(\sigma)}{U_{gb}(\sigma)} \geq \frac{1 - \beta^w}{1 - \beta} = \frac{1 - (\lambda(k-1))^w}{1 - \lambda(k-1)}.$$

Let ϵ be such that $(1 - \epsilon)^2 = 1 - \delta$, and let $w = \epsilon \cdot \log_{\lambda(k-1)}$, then $(\lambda(k-1))^w = \epsilon$. Thus

$$\frac{U_{pr}(\sigma)}{U_{gb}(\sigma)} \geq \frac{1 - (\lambda(k-1))^w}{1 - \lambda(k-1)} \geq (1 - \epsilon) \cdot \frac{1}{1 - \lambda(k-1)}.$$

Since all candidates in σ are distinct, σ is succinct, hence by lemma 4.14, there exists a prior distribution \mathcal{F} where all candidates are i.i.d such that

$$\frac{\mathbb{E}_{\mathcal{F}}[U_{pr}]}{\mathbb{E}_{\mathcal{F}}[U_{gb}^*]} \geq (1 - \epsilon) \cdot \frac{U_{pr}(\sigma)}{U_{gb}(\sigma)} = (1 - \epsilon)^2 \cdot \frac{1}{1 - \lambda(k-1)} = (1 - \lambda(k-1)) \cdot (1 - \delta).$$

\square

C Monotonicity Results

Claim C.1. *If $\|s^{(t)}\|_1 < \|s^{(t)}\|_1$ for sequences σ' and σ , where $\sigma^{(t)} = \sigma'^{(t)}$, then $U_{g_b}(\sigma, t) \leq U_{g_b}(\sigma', t)$.*

Proof. From the definition of the biased gambler's utility we have:

$$\|\sigma^{(t)}\|_1 - \lambda \left(\|\bar{s}^{(t)}\|_1 - \|\sigma^{(t)}\|_1 \right) \leq \|\sigma'^{(t)}\|_1 - \lambda \left(\|\bar{s}'^{(t)}\|_1 - \|\sigma'^{(t)}\|_1 \right) \implies U_{g_b}(\sigma, t) \leq U_{g_b}(\sigma', t)$$

□

Claim C.2. *For any super candidate $s^{(t)}$ and candidate $\sigma^{(t)}$, if $\lambda' > \lambda \geq 0$, then, when they select candidate $\sigma^{(t)}$, the biased gambler receives greater utility with loss aversion λ than with loss aversion λ' , namely $U_{g_{b,\lambda}}(\sigma, t) \geq U_{g_{b,\lambda'}}(\sigma, t)$.*

Proof. From the definition of the biased gambler's utility and the fact that $\lambda' > \lambda$, we have:

$$\|\sigma^{(t)}\|_1 - \lambda \left(\|\bar{s}^{(t)}\|_1 - \|\sigma^{(t)}\|_1 \right) \geq \|\sigma^{(t)}\|_1 - \lambda' \left(\|\bar{s}^{(t)}\|_1 - \|\sigma^{(t)}\|_1 \right) \implies U_{g_{b,\lambda}}(\sigma, t) \geq U_{g_{b,\lambda'}}(\sigma, t)$$

□

Next, we conduct analysis regarding Kleinberg et al. [11]'s construct of patience in our multi-dimensional setting. For sake of semantic clarity, we re-iterate their definition here:

Definition C.3 (Patient). *A stopping rule π is more patient than a stopping rule π' if for every realization of the candidates the stopping rule π either selects the same candidate as π' or selects a candidate that is later in the sequence.*

Furthermore, for any stopping rule π , let us define the notation $\pi(\sigma)$ to denote the value of the candidate selected by the π when the realization of the candidates is the sequence σ . Specifically, $\pi_\lambda(\sigma)$ denotes the value of the candidate selected by the optimal λ -biased stopping rule.

Claim C.4. *Given a prior distribution \mathcal{F} , if an optimal multidimensional λ -biased stopping rule π_λ is more patient than another multidimensional stopping rule π , then the expected value of the candidate selected by π_λ is higher than the expected value of the candidate selected by π (i.e. $\mathbb{E}_{\sigma \sim \mathcal{F}}[\pi_\lambda(\sigma)] \geq \mathbb{E}_{\sigma \sim \mathcal{F}}[\pi(\sigma)]$).*

Proof. First, note that we are reasoning about the objective value of the selected candidates not about the utilities that they yield to the agents. Let $t(\sigma)$ and $w(\sigma)$ denote the index of the candidate selected by π and π_λ when the realization of candidates is the sequence σ .

By the definition of patience, for any sequence σ in the support of \mathcal{F} , $t(\sigma) \leq w(\sigma)$. Since π_λ is optimal for a λ -biased agent, for a λ -biased agent playing optimally on σ :

$$\mathbb{E}_{\sigma \sim \mathcal{F}}[U_{g_b}(\sigma, w(\sigma))] \geq \mathbb{E}_{\sigma \sim \mathcal{F}}[U_{g_b}(\sigma, t(\sigma))]$$

Note that this is a statement about utilities not about the objective values of the items selected. Now we can substitute in the definition of U_{g_b} and relate the values of candidates $\sigma^{(t(\sigma))}$ and $\sigma^{(w(\sigma))}$.

$$\mathbb{E}_{\sigma \sim \mathcal{F}} \left[\|\sigma^{(w(\sigma))}\|_1 - \lambda (\|s^{(w(\sigma))}\|_1 - \|\sigma^{(w(\sigma))}\|_1) \right] \geq \mathbb{E}_{\sigma \sim \mathcal{F}} \left[\|\sigma^{(t(\sigma))}\|_1 - \lambda (\|s^{(t(\sigma))}\|_1 - \|\sigma^{(t(\sigma))}\|_1) \right].$$

Re-expressing this algebraically, we have:

$$(1 + \lambda)\mathbb{E}_{\sigma \sim \mathcal{F}}[\|\sigma^{(w(\sigma))}\|_1] - \lambda\mathbb{E}_{\sigma \sim \mathcal{F}}[\|s^{(w(\sigma))}\|_1] \geq (1 + \lambda)\mathbb{E}_{\sigma \sim \mathcal{F}}[\|\sigma^{(t(\sigma))}\|_1] - \lambda\mathbb{E}_{\sigma \sim \mathcal{F}}[\|s^{(t(\sigma))}\|_1].$$

Now assume for contradiction that $\mathbb{E}_{\sigma \sim \mathcal{F}}[\|\sigma^{(w(\sigma))}\|_1] < \mathbb{E}_{\sigma \sim \mathcal{F}}[\|\sigma^{(t(\sigma))}\|_1]$. Observe that since for any sequence σ in the support of \mathcal{F} , $w(\sigma) \geq t(\sigma)$, it follows that $\|s^{(t(\sigma))}\|_1 \leq \|s^{(w(\sigma))}\|_1$ and thus $\mathbb{E}_{\sigma \sim \mathcal{F}}[\|s^{(t(\sigma))}\|_1] \leq \mathbb{E}_{\sigma \sim \mathcal{F}}[\|s^{(w(\sigma))}\|_1]$. Consequently, our assumption violates the above inequality which we know to be true, thus it must be the case that $\mathbb{E}_{\sigma \sim \mathcal{F}}[\|\sigma^{(w(\sigma))}\|_1] > \mathbb{E}_{\sigma \sim \mathcal{F}}[\|\sigma^{(t(\sigma))}\|_1]$. \square

Claim C.5. *The optimal λ -biased stopping rule, π_λ , is more patient than the optimal λ' -biased stopping rule, $\pi_{\lambda'}$, for $\lambda \leq \lambda'$.*

Proof. Let t' denote the random variable representing the index of the candidate $\pi_{\lambda'}$ selects. For contradiction, assume that there exists realization $\sigma^{(1)}, \dots, \sigma^{(t)} = \sigma_1, \dots, \sigma_t$ for which π_λ selects $\sigma^{(t)}$ and t' (the choice of $\pi_{\lambda'}$) is the index of some candidate after t . Let event A be the event where $\sigma^{(1)}, \dots, \sigma^{(t)} = \sigma_1, \dots, \sigma_t$ and let s_t be the super candidate at time step t conditioned on event A . We claim that after seeing $\sigma_1, \dots, \sigma_t$, the λ' -biased agent will get higher expected utility by using a strategy that selects σ_t with probability 1 than by using any other strategy. Thus our assumption must be false. We now prove our claim. Since π_λ is optimal for a λ -biased agent, it must be the case that $\mathbb{E}_{\sigma \sim \mathcal{F}}[U_{g_{b_\lambda}}(\sigma, t) \mid A] \geq \mathbb{E}_{\sigma \sim \mathcal{F}}[U_{g_{b_\lambda}}(\sigma, t') \mid A]$, namely for a λ -biased agent, we have:

$$\|\sigma_t\|_1 - \lambda(\|s_t\|_1 - \|\sigma_t\|_1) \geq \mathbb{E}_{\sigma \sim \mathcal{F}}[\|\sigma^{(t')}\|_1 - \lambda(\|s^{(t')}\|_1 - \|\sigma^{(t')}\|_1) \mid A]. \quad (1)$$

As an intermediate claim, We show that $\|s_t\|_1 - \|\sigma_t\|_1 \leq \mathbb{E}_{\sigma \sim \mathcal{F}}[\|s^{(t')}\|_1 - \|\sigma^{(t')}\|_1 \mid A]$. First, let's consider the case where $\|\sigma_t\|_1 \leq \mathbb{E}_{\sigma \sim \mathcal{F}}[\|\sigma^{(t')}\|_1 \mid A]$. Due to the fact that the value of the super candidate is monotonically increasing, we have $\|s_t\|_1 \leq \mathbb{E}_{\sigma \sim \mathcal{F}}[\|s^{(t')}\|_1 \mid A]$. Combining the two inequalities gives us the desired inequality.

Next, we consider the case where $\|\sigma_t\|_1 > \mathbb{E}_{\sigma \sim \mathcal{F}}[\|\sigma^{(t')}\|_1 \mid A]$. Observe that

$$\mathbb{E}_{\sigma \sim \mathcal{F}}[\|s^{(t')}\|_1 \mid A] - \mathbb{E}_{\sigma \sim \mathcal{F}}[\|\sigma^{(t')}\|_1 \mid A] \geq \|s_t\|_1 - \mathbb{E}_{\sigma \sim \mathcal{F}}[\|\sigma^{(t')}\|_1 \mid A],$$

since $s_t \leq \mathbb{E}_{\sigma \sim \mathcal{F}}[\|s^{(t')}\|_1 \mid A]$. Combining this with the assumption that $\|\sigma_t\|_1 > \mathbb{E}_{\sigma \sim \mathcal{F}}[\|\sigma^{(t')}\|_1 \mid A]$, we get $\mathbb{E}_{\sigma \sim \mathcal{F}}[\|s^{(t')}\|_1 \mid A] - \mathbb{E}_{\sigma \sim \mathcal{F}}[\|\sigma^{(t')}\|_1 \mid A] \geq \|s_t\|_1 - \|\sigma_t\|_1$, as desired.

Having shown the intermediate claim, we can re-express it algebraically as follows:

$$-(\lambda' - \lambda) \cdot (\|s_t\|_1 - \|\sigma_t\|_1) \geq -(\lambda' - \lambda) \cdot \mathbb{E}_{\sigma \sim \mathcal{F}}[(\|s^{(t')}\|_1 - \|\sigma^{(t')}\|_1) \mid A]. \quad (2)$$

Combining statements (1) and (2), we get the following:

$$\|\sigma_t\|_1 - \lambda'(\|s_t\|_1 - \|\sigma_t\|_1) \geq \mathbb{E}_{\sigma \sim \mathcal{F}}[\|\sigma^{(t')}\|_1 - \lambda'(\|s^{(t')}\|_1 - \|\sigma^{(t')}\|_1) \mid A]. \quad (3)$$

By the definition of the biased agents utility, statement (3) is equivalent to the statement that $\mathbb{E}_{\sigma \sim \mathcal{F}}[U_{g_{b_{\lambda'}}}(\sigma, t) \mid A] \geq \mathbb{E}_{\sigma \sim \mathcal{F}}[U_{g_{b_{\lambda'}}}(\sigma, t') \mid A]$, namely that conditioned on event A , a λ' -biased agent is at least as well off selecting t as selecting t' thus contradicting our initial assumption and showing no such $\sigma_1, \dots, \sigma_t$ can exist. \square

Corollary C.6. *If $\lambda \leq \lambda'$, then, when both are playing optimally, the expected value of the candidate selected by a λ -biased agent is higher than the expected value of a candidate selected by a λ' -biased agent (i.e. $\mathbb{E}_{\sigma \sim \mathcal{F}}[\pi_\lambda(\sigma)] \geq \mathbb{E}_{\sigma \sim \mathcal{F}}[\pi_{\lambda'}(\sigma)]$).*

Proof. It follows as a direct consequence of claims C.4 and C.5 when they are taken together. \square

We now move on to consider monotonicity in the number of candidates.

Claim C.7. *Let n, m be two integers such that $m > n$. Let $\mathcal{F}_1, \dots, \mathcal{F}_m$ be m value distributions for each of $n+1$ distinct candidates. Let $\mathcal{F} = \times_{i \in [n]} \mathcal{F}_i$ and $\mathcal{F}' = \times_{i \in [m]} \mathcal{F}_i$. We will use π_λ and π'_λ to denote the optimal λ -biased stopping rule for $\sigma \sim \mathcal{F}$ and $\sigma' \sim \mathcal{F}'$ respectively. We claim that π'_λ is more patient than π_λ .*

Proof. Assume for the sake of contradiction that there exist $\sigma_1, \dots, \sigma_m$ where $\sigma = \sigma_1, \dots, \sigma_n$ and $\sigma' = \sigma_1, \dots, \sigma_m$ such that $\pi_\lambda(\sigma)$ is later in the sequence than $\pi'_\lambda(\sigma')$. Since, π_λ is the optimal λ -biased stopping rule, this assumption implies that

$$U_{g_{b_\lambda}}(\sigma, \pi_\lambda(\sigma)) > U_{g_{b_\lambda}}(\sigma, \pi'_\lambda(\sigma')) = U_{g_{b_\lambda}}(\sigma', \pi'_\lambda(\sigma')).$$

This is a contradiction since a stopping rule for $\sigma' \sim \mathcal{F}'$ can get strictly better utility compared to π'_λ by taking the same action as π_λ and never selecting candidates with index greater than n . \square

Claim C.8. *Let n, m be two integers such that $m > n$. Let $\mathcal{F}_1, \dots, \mathcal{F}_m$ be m value distributions for each of m distinct candidates. Let $\mathcal{F} = \times_{i \in [n]} \mathcal{F}_i$ and $\mathcal{F}' = \times_{i \in [m]} \mathcal{F}_i$. Consider $\sigma \sim \mathcal{F}$ and $\sigma' \sim \mathcal{F}'$ where σ is a prefix of σ' . For any $\lambda \geq 0$ and optimal λ -biased stopping rule π_λ and π'_λ for $\sigma \sim \mathcal{F}$ and $\sigma' \sim \mathcal{F}'$ respectively, $U_{g_{b_\lambda}}(\sigma, \pi_\lambda(\sigma)) \leq U_{g_{b_\lambda}}(\sigma', \pi'_\lambda(\sigma'))$ and $\pi_\lambda(\sigma) \leq \pi'_\lambda(\sigma')$. Semantically, both the utility and expected value of the selected candidate improve monotonically as more candidates are added to the end of the sequence.*

Proof. Since σ is a prefix of σ' , an agent can simply select candidate $\pi_\lambda(\sigma)$ and receive at least as much utility on σ' as on σ , thus $U_{g_{b_\lambda}}(\sigma, \pi_\lambda(\sigma)) \leq U_{g_{b_\lambda}}(\sigma', \pi'_\lambda(\sigma'))$. By claim C.7, the agent is more patient on σ' than on σ ; combining this with claim C.4, it must be the case that the expected value of the candidate selected on σ' is greater than or equal to the expected value of that selected on σ . \square

Following the outline of Kleinberg et al. [11], we now consider adding a candidate at the beginning of the sequence. Such addition can decrease the expected utility of a λ -biased agent with comparative loss aversion, whereas it cannot decrease that of a rational agent.

Claim C.9. *Let $\mathcal{F}_1, \dots, \mathcal{F}_{n+1}$ be $n+1$ value distributions for each of $n+1$ distinct candidates. Let $\mathcal{F} = \times_{i=2}^{n+1} \mathcal{F}_i$ and $\mathcal{F}' = \times_{i \in [n+1]} \mathcal{F}_i$. For any $\lambda \geq 0$ and optimal λ -biased stopping rule π_λ and π'_λ for \mathcal{F} and \mathcal{F}' respectively, consider sequences $\sigma \sim \mathcal{F}$ and $\sigma' \sim \mathcal{F}'$, where $\sigma' = \sigma^{(0)'} + \sigma$. (We can view σ' as the result of appending a candidate, $\sigma^{(0)'}$, to the beginning of σ .) We have $U_{g_b}(\sigma', \pi'_\lambda(\sigma')) \geq \frac{1}{1+\lambda} U_{g_b}(\sigma, \pi_\lambda(\sigma))$.*

Proof. We lower bound $U_{g_b}(\sigma', \pi_\lambda(\sigma'))$ by considering the utilities for two possible stopping rules:

1. Accept the first candidate in σ' , namely $\sigma^{(0)'}$

2. Ignore the first candidate, $\sigma^{(0)'}$ and instead accept the same candidate selected on σ , namely candidate $\pi_\lambda(\sigma)$.

In the first case, since the first candidate seen is selected, there is no regret and the utility of the gambler is $\|\sigma^{(0)'}\|_1$. In the second case, since the same candidate is chosen as under σ and the super candidate can increase by at most the value of the candidate appended to the beginning, the utility is lower bounded by $U_{g_b}(\sigma, \pi_\lambda(\sigma)) - \lambda\|\sigma^{(0)'}\|_1$. Since π_λ is the utility maximizing strategy we thus have:

$$U_{g_b}(\sigma', \pi_\lambda(\sigma')) \geq \max\{U_{g_b}(\sigma, \pi_\lambda(\sigma)) - \lambda\|\sigma^{(0)'}\|_1, \|\sigma^{(0)'}\|\} \quad (1)$$

Consider the case where $U_{g_b}(\sigma, \pi_\lambda(\sigma)) - \lambda\|\sigma^{(0)'}\|_1 \leq \|\sigma^{(0)'}\|_1$, which algebraically re-expressed is $\|\sigma^{(0)'}\|_1 \geq \frac{U_{g_b}(\sigma, \pi_\lambda(\sigma))}{1+\lambda}$. Combining this with inequality (1) yields $U_{g_b}(\sigma', \pi_\lambda(\sigma')) \geq \frac{U_{g_b}(\sigma, \pi_\lambda(\sigma))}{1+\lambda}$. Now consider the case where $U_{g_b}(\sigma, \pi_\lambda(\sigma)) - \lambda\|\sigma^{(0)'}\|_1 > \|\sigma^{(0)'}\|_1$, algebraically re-expressed this is the case where $\|\sigma^{(0)'}\|_1 < \frac{U_{g_b}(\sigma, \pi_\lambda(\sigma))}{1+\lambda}$. Here we get:

$$U_{g_b}(\sigma', \pi_\lambda(\sigma')) \geq U_{g_b}(\sigma, \pi_\lambda(\sigma)) - \lambda\|\sigma^{(0)'}\|_1 \geq U_{g_b}(\sigma, \pi_\lambda(\sigma)) - \lambda \frac{U_{g_b}(\sigma, \pi_\lambda(\sigma))}{1+\lambda} = \frac{U_{g_b}(\sigma, \pi_\lambda(\sigma))}{1+\lambda}$$

Thus, we have proven the first the claim, namely that $U_{g_b}(\sigma', \pi'_\lambda(\sigma')) \geq \frac{U_{g_b}(\sigma, \pi_\lambda(\sigma))}{1+\lambda}$. □