The Prophet Paradox: Optimal Stopping with Multi-Dimensional Comparative Loss Aversion

JOSH GARDNER

In the optimal stopping problem, an online decision maker sees options one by one in sequence and must decide immediately whether to select the current option or forego it and lose it forever. Recent work has explored the effects of loss aversion and reference dependence on the optimal stopping problem. We extend on this work by considering a modified version of the problem, in which options have multiple features and agents are behaviorally biased. Under this scheme, a biased agent compares the current option against the best option previously seen for each feature. In this sense, its reference point is the combination of the best features of all previously seen options.

Our modified version of the optimal stopping problem is mathematically simple and grounded in behavioral research. Significantly, it has bizarre properties. In particular, using the language of prophet inequalities, if both are biased, then there exist scenarios in which the performance of the prophet(i.e. an offline decision maker) is a lower bound on that of the gambler(the online decision maker) for all strategies. Furthermore, there are sequences such that adding options to the tail of the sequence strictly decreases the performance of a biased prophet, as well as sequences in which improving the value of all options individually leads to a decrease in the performance of the prophet. Lastly, throughout we discuss behavioral analogs of our results, most notably the paradox of choice, in which adding more options decreases a person's happiness with their chosen option.

Manuscript submitted for review to the 23rd ACM Conference on Economics & Computation (EC'22).

1 INTRODUCTION

The present work builds upon Kleinberg et al. [4], extending their results on behaviorally biased optimal stopping to a more general model in which candidates have value along multiple dimensions, i.e. are vector valued. For sake of clarity, let us first revisit their model. The main novelty of Kleinberg et al. [4] was the incorporation of reference dependence and loss aversion into their model of optimal stopping. 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 [8]. 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[3], and have been observed in the context of online decision making [5].

In the model of Kleinberg et al. [4], an agent suffers from both of these biases and must select a candidate online from a sequence in which each candidate has a single scalar value. Intuitively, this model captures scenarios in which candidates are assessed by a single scalar such as monetary value. Indeed, their results directly mirror the empirical findings of Schunk and Winter [5] 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 is rationally optimal, once again mirroring the empirical results in [5]. Informally, due to the anticipation of experiencing regret over foregone options, people are likely to settle earlier than they should. In this respect, the theoretical model of Kleinberg et al. [4] 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, such as choice of career, romantic partner, or home typically have many dimensions. Moreover, even simple consumer decisions such as selecting which movie to watch or magazine to subscribe to often have multiple dimensions. Experimental research by Brenner et al. [1] 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. Such findings have been used to explain a phenomenon referred to as the paradox of choice [6] in which the addition of extra options decreases the happiness a person experiences with the option they select. Adding additional options cannot decrease the utility received by a rational agent, so this phenomenon represents a significant divergence between human behavior and rational behavior. 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 is the cheapest. Regardless of which house a person selects, due to comparative loss aversion, the mere awareness of the other options will decrease the satisfaction they experience.

Modeling comparative loss aversion in optimal stopping. We now construct a model of the optimal stopping problem in which an agent suffers from comparative loss aversion. First, we must formally define the optimal stopping problem, using classic definitions with slight modifications. In particular, let σ be an ordered sequence of n candidates, where the t-th candidate has a k-dimensional vector value $\vec{v}^{(t)} \in R^k$ such that the value at index j, denoted $\vec{v}^{(t)}_j$, is drawn from some distribution $\mathcal{F}_j$. (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.) Moreover, let it be the case that the agent sees the realized value of $\vec{v}^{(t)}$ only when they reach t-th candidate in the sequence σ. Upon reaching the t-th candidate, the agent must either accept them - thus halting their search - or decline them, continuing their search and losing the candidate forever.

Furthermore, in keeping with the notation of Kleinberg et al. [4], 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. Furthermore, this aligns with behavioral intuitions about how comparative loss aversion may work:

The existence of multiple alternatives makes it easy for us to imagine alternatives that don't exist—alternatives that combine the attractive features of the ones that do exist. And to the extent that we engage our imaginations in this way, we will be even less satisfied with the alternative we end up choosing.
-Barry Schwartz, The Paradox of Choice[6]

At all times, the agent imagines precisely such an alternative and uses it as its reference point when considering newly presented candidates. The agent remembers the best value previously seen for each feature 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. Moreover, 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 $\vec{s}_j^{(t)} = \max_{w: w \in [t]} (\vec{v}_j^{(w)})$. For the purpose of simplifying our analysis, we restrict all values to be non-negative. Moreover, we define the utility received by a biased agent, who suffers from loss aversion λ and who selects the t-th candidate in sequence σ, as follows:

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

For sake of brevity and semantic clarity let us refer to $U(\sigma, t)$ simply as utility, since it is the experienced payoff of the agent, even if it is biased by loss aversion and reference dependence. Compare this to the payoff received by a fully rational agent who does not suffer from reference dependence or loss aversion and selects candidate i:

$$V(t) = \sum_j \vec{v}_j^{(t)} = \|v^{(t)}\|_1$$

Similarly, for the sake of brevity, let us refer to this payoff simply as the "value" since it is simply the L1 norm of the value vector of the item and does not depend on previously seen items. Semantically, it is the objective value of the item, in contrast with its subjective utility.

Lastly, let us frame the problem in the traditional language of prophet inequalities, i.e. in terms of "the prophet" and "the gambler." Specifically, let the prophet have full knowledge of the candidates ahead of time and thus decide as if they are an offline decision maker. Note that for the purpose of analysis, the prophet is fully equivalent to an offline decision maker and will be interchangeably referred to as such. In contrast, let the gambler be an online decision maker, who only knows candidates as they are revealed one by one. Note that since the prophet is an offline decision maker, they just choose whichever candidate maximizes their utility. As a consequence in most analyses of optimal stopping, the utility of the prophet is used as an upper bound on the attainable utility of the gambler. As we will soon see, when multidimensional reference dependence is factored in, this is no longer necessarily the case.

High Level Summary of Novel Findings. We have successfully constructed a behaviorally based model of the optimal stopping problem in which agents suffer from comparative loss aversion. The purpose of the present work is to explore the properties of this new model and to determine bounds on the performance of the agent for various sequences. Our quantitative analysis of the model yields results which are of interest in the qualitative analysis of behavior and decision making. To the extent to which our model is an accurate representation of human decision making, it makes interesting and at times startling predictions.

We now provide a summary of our main quantitative results. First, comparative loss aversion causes the prophet to behave irrationally when candidates have more than one dimension. In Kleinberg et al. [4], there was no need to distinguish between a behaviorally biased prophet and a rational prophet since they behaved in exactly the same way. However, in the multidimensional case, a biased prophet can receive dramatically decreased utility compared to a rational one. Stranger still, there exist sequences for which, regardless of strategy used, the gambler always performs at least as well as the prophet and indeed outperforms the prophet for the vast majority of strategies. We dub this strange phenomenon in which the gambler can be better off than the prophet "the prophet paradox." Additionally, we find that adding options can make both the gambler and the prophet worse off, namely we find that our model exhibits the paradox of choice. Lastly, we find that it exhibits an even stranger paradox in which by strictly improving all options in a sequence, we harm the utility of both the gambler and the prophet. We name this paradox "the paradox of quality."

We perform extensive qualitative analysis, constructing sequences which we believe are likely to mirror many seen in real world scenarios. We show that the strange behaviors of our model also occur on these somewhat more mundane sequences. Behaviorally our model yields two key insights. First, options that are sufficiently excellent on differing dimensions will cause the paradox of choice to occur, and as the magnitude of their excellence increases, our agent will paradoxically be less satisfied with them(i.e. the paradox of quality occurs). Second, we observe that the mere awareness of more choices should result in all the startling behaviors of the model, and thus observe that the proliferation of choice and information in modern society may have unexpected adverse effects. Lastly, we find direct analogs of these insights in behavioral literature and also discuss shortcomings of our model.

2 THE PROPHET PARADOX AND RELATED RESULTS

We will start our analysis by considering some prophet inequalities. Prior to exploring these however, we need to establish a distinction between a "biased prophet" and an "unbiased prophet." Namely, we need to distinguish between a prophet who suffers from behavioral biases and one who does not. Such a distinction has been completely unnecessary in prior work, however it is critical for understanding the present model and its properties.

In the one dimensional case(i.e. that explored in Kleinberg et al. [4]), there is no difference between a biased prophet and a unbiased one. The payoff maximizing candidate for both the biased and rational prophet is just the maximum valued candidate. In particular since it is the maximum valued candidate, there is no loss aversion factor to subtract for the biased prophet and thus their utility is the same as the payoff received by the unbiased prophet. Furthermore, in the one dimensional case, the utility of biased prophet is an upper bound on the utility of the biased gambler: since selecting the maximum valued candidate is the best the gambler can possibly do, it follows that the gambler's utility is upper bounded by that of the prophet.

However, the jump from one dimension to multiple dimensions causes a fundamental shift in the problem, such that a biased prophet can receive significantly lower utility than an unbiased prophet. To explore this phenomenon, there are two prophet inequalities which are of particular interest: (1) the ratio of the value of the candidate the biased gambler selects to the value of the candidate that the biased prophet selects, (2) the ratio of the utility of the biased gambler, denoted $U_g(\sigma, t)$, to the utility of the biased prophet, denoted $U_p(\sigma, q)$, where we let t denote the candidate that the biased gambler selects and let r denote the candidate that the biased prophet selects.

Theorem 2.1 (The Prophet Paradox). *There exist adversarial sequences such that the utility received by a biased prophet is a lower bound on the utility received by a biased gambler for all possible strategies, in spite of the values of their chosen candidates being equivalent. Formally, there exists σ such that $V(t) = V(r)$ and $U_g(\sigma, t) \geq U_p(\sigma, r)$.*

Proof. Consider an adversarial sequence σ of length k, such that there are k candidates each of which has a k-dimensional value vector. Moreover, let q be some fixed constant such that $q > 1$, and let the value vector of the i-th candidate be q at index i and 0 everywhere else. Semantically, q describes the "quality" of our set of options, i.e. since all options have the same value, q directly determines the maximum attainable value. Our sequence appears as follows:

$$\begin{aligned} v^{(1)} &= [q, 0, ..., 0, 0] \\ v^{(2)} &= [0, q, ..., 0, 0] \\ &\cdots \\ v^{(k-1)} &= [0, 0, ..., q, 0] \\ v^{(k)} &= [0, 0, ..., 0, q] \end{aligned}$$

Trivially, since the prophet sees all candidates, it follows that the "super candidate" imagined by the biased prophet is $s^{(k)} = [q, q, \ldots, q] \in R^k$. Since all the value vectors for the candidate set are q at only one index, it follows that the biased prophet receives the same utility regardless of which candidate, r, they select. Specifically, we have:

$$U_p(\sigma, r) = \|v^{(r)}\|_1 - \lambda\left(\|s^{(k)}\|_1 - \|v^{(r)}\|_1\right) = q\left(1 - \lambda(k-1)\right)$$

Observe that, in contrast, the gambler who selects candidate t sees only t candidates total and thus their super candidate has a value vector with q's in the first t indices and zeros thereafter. Thus, he gambler receives the following utility:

$$U_g(\sigma, t) = \|v^{(t)}\|_1 - \lambda\left(\|s^{(t)}\|_1 - \|v^{(t)}\|_1\right) = q\left(1 - \lambda(t-1)\right)$$

Since $t \in [k]$, for $k \geq 2$, $q > 1$ and $\lambda > 0$, we have:

$$q\left(1 - \lambda(k-1)\right) \leq q\left(1 - \lambda(t-1)\right) \implies U_p(\sigma, r) \leq U_g(\sigma, t)$$

□

This result is somewhat bizarre because the prophet is worse off then the gambler. Regardless of which candidate the prophet selects, its value is exactly the same as that of the candidate which the gambler selects, however the prophet's is always the same or worse than that of the gambler. Moreover, since the gambler's utility is a decreasing function of their chosen option's position in the sequence, the sooner they stop, the better off they are. In particular, if they take the first candidate, they will receive positive utility, otherwise, they will receive negative utility. Their utility will become increasingly negative the longer that they wait before settling on a candidate, such that choosing the last possible candidate is the worst possible strategy. In contrast, a biased prophet always does exactly as bad as this worst case since they see all candidates. In this case, there is no benefit to being able to see the future. If anything, for this sequence, the prophet's clairvoyance is a curse.

Observe that in the multi-dimensional case, an unbiased prophet would be guaranteed to get non-negative payoff since all candidates have non-negative value. In contrast, a biased prophet can receive arbitrarily negative utility for adversarial sequences. As the number of dimensions considered, k, grows arbitrarily large, the utility received by the biased prophet can grow arbitrarily negative. Using the example provided in theorem 2.1, regardless of which candidate the prophet chooses they just receive payoff of $q(1-\lambda(k-1))$. As k grows arbitrarily large, this payoff becomes arbitrarily negative.

Comparative loss aversion leads to a fundamental breakdown in the traditional semantics of prophet inequalities. As we have seen, it is critical that we distinguish between biased and unbiased prophets. Furthermore, whereas in one dimension, a biased prophet was still guaranteed to always outperform both a biased and unbiased gambler, in multiple dimensions, there is no such guarantee.

2.1 The Paradox of Quality

The example sequence used to show the prophet paradox leads to other strange behaviors as well. In particular, we find that for agents who suffer from comparative loss aversion, it is possible to improve the quality of entire the set of candidates they are exposed to while making their utility decrease.

Definition 2.2 (Higher Quality). If sequences σ and σ' are such that the minimum valued option in σ' is of greater value than the maximum valued option in σ, then we say that σ' is of a higher quality than σ.

Corollary 2.3 (The Paradox of Quality). *There exist σ and σ', where σ' is of higher quality than σ, such that for all r, $U_p(\sigma, r) > U_p(\sigma', r)$, and, for some t, $U_g(\sigma, t) > U_p(\sigma', t)$.*

Proof. Let us simply re-use the example sequence from theorem 2.1 with different settings of the quality q. Recall that in this sequence, all candidates had the same value, namely q. To construct σ simply set q such that $q > 1$, to construct σ' simply set $q' > q$. Moreover, recall the utilities we derived for the biased prophet and gambler, namely $U_p(\sigma, r) = q(1 - \lambda(k - 1))$ and $U_g(\sigma, t) = q(1 - \lambda(t - 1))$, respectively. If $\lambda > \frac{1}{k-1}$, then as the quality q of each candidate improves, the utility of the biased prophet becomes increasingly negative, as does the utility of the biased gambler for all strategies which do not select the first candidate. Thus it follows that $\forall r \in [k]$,$U_p(\sigma, r) > U_p(\sigma', r)$, and for all $t \geq 2$, $U_g(\sigma, t) > U_p(\sigma', t)$. □

This result is somewhat remarkable in that by improving the quality q of each candidate and thus its value, we make the biased prophet increasingly unhappy with whatever candidate they end up choosing. Moreover, the same effect happens for the biased gambler as long as they don't select the first candidate they see. As the quality of their candidates becomes arbitrarily good, their utility becomes arbitrarily bad. In essence, as the objective quality of their options increases, the subjective quality of their options decreases due to comparative loss aversion. This paradox is especially bizarre and, to the extent to which our model reflects human behavior, somewhat concerning.

3 BEHAVIORAL ANALYSIS

Intuitively, information about additional candidates makes our agent less satisfied with their ultimate decision. Moreover, this effects the prophet worse than the gambler; since the gambler will never know the candidates that would come after the one they pick, they can't feel loss over having turned them down. In the case of our biased agent, ignorance truly is bliss. In behavioral literature[6], the phenomenon in which knowledge of additional candidates decreases happiness with one's chosen candidate is referred to as "The Paradox of Choice". We borrow the behavioral terminology here and use it explore additional aspects of the paradox of choice that occur in our model but not in the original model of Kleinberg et al. [4].

Definition 3.1 (Paradox of Choice). If adding a candidate makes the utility of an agent decrease, we refer to this as a paradox of choice.

Note that paradoxes of choice cannot occur for fully rational agents since they will simply disregard additional candidates that do not improve the maximum. Notably, Kleinberg et al. [4] already showed a paradox of choice for the one dimensional case. In particular, in their monotonicity results, they showed that adding candidates to the beginning of a sequence can harm the utility of a biased gambler.

Observe that in the example used in theorem 2.1, as the number of candidates seen increases, the utility received by the agent decreases. Moreover, observe that, by construction of the example, the order in which the candidates are presented has no effect on the analysis. All items have the same rational utility, yet as the number of items seen increases, the biased utility decreases linearly. This is an extreme example of the paradox of choice at work.

3.1 Applying The Model

We can extend the intuition from theorem 2.1 to create models which appear more behaviorally realistic than the contrived example used in the proof of the theorem, yet which still yield paradoxes of quality and choice. In particular, consider the scenario in which all candidates have the same values for all features except for one salient feature per candidate. There is empirical evidence [2] suggesting that people reduce phenomena to salient features and can only perceive differences between stimuli if their magnitude exceeds certain thresholds. We consider a scenario in which most features across candidates are assigned the same value, and all noticeably better features are assigned another greater value. In particular, we start with the case where each candidate is advantaged along a different dimension.

Differently Advantaged Candidates As in theorem 2.1, consider an adversarial sequence σ in which there are k candidates, each of which has values along k distinct dimensions. Let "a" be the standard value assigned to a non-salient features, i.e. features which are perceived as being normal/unexceptional, and let $a > 0$. Moreover, let q be some fixed constant such that $q > 1$. Let the value vector of the i-th candidate be $a + q$ at index i and a everywhere else. Trivially, since the prophet sees all candidates, it follows that the "super candidate" seen by the biased prophet is $s^{(k)} = (a+q, a+q, \dots, a+q) \in R^k$. Since all the value vectors for the candidate set are $a+q$ at only one index, it follows that the biased prophet(i.e. an online decision maker) receives the same utility $U_p(\sigma, t)$ regardless of selected candidate r. We have:

$$U_p(\sigma, r) = \|v^{(r)}\|_1 - \lambda\bigg(\|s^{(k)}\|_1 - \|v^{(r)}\|_1\bigg) = (ak+q) - \lambda\bigg((a+q)k - (ak+q)\bigg) = ak + q\bigg(1 - \lambda(k-1)\bigg)$$

Similarly for the gambler who stops at candidate r in the sequence, we have:

$$U_g(\sigma, r) = \|v^{(r)}\|_1 - \lambda\left(\|s^{(r)}\|_1 - \|v^{(r)}\|_1\right) = (ar+q) - \lambda\left((a+q)r - (ar+q)\right) = ar + q\left(1 - \lambda(r-1)\right)$$

Observe that the prophet's utility is negative when $\lambda > \frac{1}{k-1}$ and $q \geq \frac{ak}{(\lambda(k-1)-1)}$, and as q grows beyond this point, the prophet's utility grows increasingly negative. Similarly, the gambler's utility is negative when r is greater than 1, $\lambda > \frac{1}{r-1}$, and $q \geq \frac{ar}{(\lambda(r-1)-1)}$. As q grows beyond this point, the prophet's utility grows increasingly negative. Thus, the paradoxes of quality and choice occur for this example as well. At an intuitive level, it appears that the paradoxes occur when candidates are advantaged along different dimensions but nevertheless are otherwise of similar overall quality.

3.2 Behavioral Limitations of the Model

Although the model captures interesting aspects of human behavior such as comparative loss aversion and certain aspects of the paradox of choice, it fails to capture certain key behavioral phenomena. In particular, sometimes adding inferior but similar options can alter people's ultimate decisions. As an example, Simonson and Tversky [7] found when given the choice between a luxury pen and \$6, only 36% of participants chose the pen compared to 46% of participants who were given the choice between \$6, a luxury pen, and a comparatively cheap pen. Observe that in the case of the study, the decision was offline, i.e. analogous to the decision of a prophet in the context of optimal stopping. In contrast with the results of the study, a biased prophet in our model always selects the same candidate as an unbiased one, namely the candidate of maximum value. Adding an additional candidate of lower value should have no effect on which candidate they select(although it may decrease the utility of the candidate they choose).

Lemma 3.2. *A biased prophet always selects the same candidate as an unbiased one. More specifically, both always select the value maximizing option.*

Proof. For a sequence σ with n candidates, by definition, the biased prophet selects candidate t such that

$$t = \arg\max_x \left(\|\vec{v}^{(x)}\|_1 - (\|\vec{s}^{(n)}\|_1 - \|\vec{v}^{(x)}\|_1)\right)$$

Observe that $\|\vec{s}^{(n)}\|_1$ is positive and fixed regardless of which candidate is chosen and thus the optimal choice is just whichever candidate t has the highest value $\|\vec{v}^{(t)}\|_1$. Moreover, recall that by construction an unbiased prophet simply selects the candidate with highest value as well. □

The addition of an inferior option will not alter the choice of our prophet. Our model's deviation from observed human behavior suggests that there is still substantial room for improvement. As much as our model produces irrational behaviors, there are still important aspects of human irrationality which it does not capture.

4 DISCUSSION

In this paper, we have created a novel variation of the optimal stopping problem in which a biased agent suffering from comparative loss aversion considers the pros and cons of candidates along multiple dimensions. While it is fundamentally an extension of the model of behaviorally biased optimal stopping developed by Kleinberg et al. [4], our model nonetheless exhibits drastically different behavior, particularly in the phenomena that we have dubbed "the prophet paradox" and "the paradox of quality", respectively.

To the extent to which our model is an accurate reflection of human decision making, it makes some alarming suggestions. Notably, our model suggests that in situations where options are of roughly the same overall quality but have advantages on differing dimensions, it may be best to choose early and avoid considering too many options. Remarkably, as we have seen, in some situations, it may be best to simply take the first option that comes our way. This suggestion violates basic common sense and with good reason. Our model uses an agent who has no knowledge of future options, not even at the distributional level. In real world decisions, such a complete lack of information is uncommon.

As a concrete example, a job seeker who simply selects the first job that comes their way will likely be aware of other potential job opportunities that are out there. In this respect, their reference point is not necessarily jobs that they have been offered and have declined. Rather, it is likely to be some "super job" comprised of the best features of the jobs that they perceive as being attainable. Similar reasoning likely applies to dating, home buying, choosing schools, and any number of other key areas in life.

Earlier in this paper, we used the phrase "clairvoyance is a curse." In reality, it doesn't take a clairvoyant prophet to anticipate future options. In a world with ever growing information and choice available at our fingertips, it is all too easy to imagine the better options lurking around the corner, whether a swipe away on a dating app, a few interviews away at the hot new startup across town, or hiding in home listings online. An abundance of choice and hyper-awareness of options is an inescapable feature of modern life. In line with behavioral research [6], our model suggests the unintuitive notion that the explosive growth in quantity and quality of choices that many of us experience year to year may actually be making us feel worse in spite of making the objective quality of our options and, by extension, our lives, better.

REFERENCES

[1] Lyle Brenner, Yuval Rottenstreich, and Sanjay Sood. 1999. Comparison, grouping, and preference. *Psychological Science* 10, 3 (1999), 225–229.

[2] E Tory Higgins. 1996. Activation: Accessibility, and salience. *Social psychology: Handbook of basic principles* (1996), 133–168.

[3] Daniel Kahneman and Amos Tversky. 2013. Prospect theory: An analysis of decision under risk. In *Handbook of the fundamentals of financial decision making: Part I.* World Scientific, 99–127.

[4] Jon Kleinberg, Robert Kleinberg, and Sigal Oren. 2021. Optimal Stopping with Behaviorally Biased Agents: The Role of Loss Aversion and Changing Reference Points. *arXiv preprint arXiv:2106.00604* (2021).

[5] Daniel Schunk and Joachim Winter. 2009. The relationship between risk attitudes and heuristics in search tasks: A laboratory experiment. *Journal of Economic Behavior & Organization* 71, 2 (2009), 347–360.

[6] Barry Schwartz. 2004. The paradox of choice: Why more is less. Ecco New York.

[7] Itamar Simonson and Amos Tversky. 1992. Choice in context: Tradeoff contrast and extremeness aversion. *Journal of marketing research* 29, 3 (1992), 281–295.

[8] Amos Tversky and Daniel Kahneman. 1991. Loss aversion in riskless choice: A reference-dependent model. *The quarterly journal of economics* 106, 4 (1991), 1039–1061.