
Intrinsic Robustness of Prophet Inequality to Strategic Reward Signaling

Wei Tang*

Haifeng Xu†

Ruimin Zhang‡

Derek Zhu§

Abstract

Prophet inequality concerns a basic optimal stopping problem and states that simple threshold stopping policies — i.e., accepting the first reward larger than a certain threshold — can achieve tight $1/2$ -approximation to the optimal prophet value. Motivated by its economic applications, this paper studies the robustness of this approximation to natural *strategic manipulations* in which each random reward is associated with a self-interested player who may selectively reveal his realized reward to the searcher in order to maximize his probability of being selected.

We say a threshold policy is α (-strategically)-robust if it (a) achieves the α -approximation to the prophet value for strategic players; and (b) meanwhile remains a $1/2$ -approximation in the standard non-strategic setting. Starting with a characterization of each player’s optimal information revealing strategy, we demonstrate the intrinsic robustness of prophet inequalities to strategic reward signaling through the following results: (1) for arbitrary reward distributions, there is a threshold policy that is $\frac{1-1/e}{2}$ -robust, and this ratio is tight; (2) for i.i.d. reward distributions, there is a threshold policy that is $1/2$ -robust, which is tight for the setting; and (3) for log-concave (but non-identical) reward distributions, the $1/2$ -robustness can also be achieved under certain regularity assumptions.⁵

1 Introduction

The prophet inequality of Krengel and Sucheston [37] is a foundational framework for the theory of optimal stopping problems and sequential decision-making. In the classic prophet inequality, a searcher faces a finite sequence of non-negative, real-valued and independent random variables X_1, \dots, X_N with known distributions H_i from which a reward of value X_i is drawn sequentially for $i = 1, \dots, N$. Once a random reward is realized, the searcher decides whether to accept the realized reward and stop searching, or reject the reward and proceed to the next reward. The searcher’s objective is to maximize the value of the accepted reward. The performance of the searcher’s stopping policy is evaluated against a *prophet value* which equals to the ex-post maximum realized reward. The classic and elegant result of Samuel-Cahn [42] showed that a simple static threshold policy achieves at least half of the prophet value, and surprisingly, this bound is the best possible even among dynamic policies. Samuel-Cahn’s policy uses the threshold that is the median of the distribution of the highest prizes, and then accepts the first realized reward that exceeds this threshold. The existence of this $1/2$ -approximation is now known as the *prophet inequality*.

Recently, there is a regained interest of the prophet inequality due to its beautiful connection to online mechanism design (see, e.g., [31, 16]), and many different settings of the prophet inequality

*Chinese University of Hong Kong. Email: weitang@cuhk.edu.hk.

†University of Chicago. Email: haifengxu@uchicago.edu.

‡University of Chicago. Email: ruimin@uchicago.edu.

§University of Chicago. Email: dzhuh8@uchicago.edu.

⁵The full version of this paper can be found at <https://arxiv.org/pdf/2409.18269>

have been studied (see the survey by [40, 19]). For example, Kleinberg and Weinberg [36] extend the prophet inequality to all matroid constraints and similar to [42], they also show that a threshold stopping policy with the threshold equal to half of the expected maximum reward can also lead to the optimal $1/2$ -approximation. Indeed, it is now well-known that there exists a range of thresholds for the classic prophet inequality that can achieve the optimal $1/2$ -approximation (see Definition 2.2).

An important assumption in the classic prophet inequality is that the distribution of each random variable is an inanimate object and, once the searcher reaches it, it will fully disclose its realized reward to the searcher. Yet this may not be the case in many real-world applications where each distribution may often be associated with a *strategic player*⁶ who may have incentives to selectively disclose information to maximize his own probability of being chosen by the searcher. This is usually the case when information is not controlled by nature but by humans or algorithms. Such examples are ample in economic activities. For instance, when a recruiter searches for the best candidate for a position by sequentially interviewing a set of job applicants, each job applicant is naturally a strategic player and would want to control how much information they disclose about their characteristics (including strengths, weaknesses, personality, and experience) to the recruiter (the searcher) so as to maximize the probability of being hired. Similarly, when a venture capitalist searches for the most promising startup to invest by sequentially visiting each startup, the startups are strategic players. They have more accurate information about their own potential and can control how much of this information they disclose to attract the investment from the venture capitalist (the searcher). Finally, in the real estate market, when a buyer searches for the most attractive property to purchase, the property sellers are strategic players who can control the amount of information they disclose about the property’s condition and potential returns to the buyer (the searcher).

Motivated by real-world applications like above, we introduce and study a natural variant of the prophet inequality where each reward distribution is associated with a strategic player who can decide what information about the realized reward will be disclosed to the searcher. To capture such partial information revelation, we adopt the standard framework of information design (also known as Bayesian persuasion [35, 10]) and assume that each player can selectively disclose reward information by implementing an information strategy – often referred to as an *experiment* or *signaling scheme* [35, 10] – which stochastically maps the realized reward (unobservable to the searcher) to a random signal (observable to the searcher). For the motivating applications above, such revelation strategies encode answers to standard interview questions (e.g., experiences, behavioral questions) prepared in advance by job candidates, the presentations prepared by startups for the VC, or the property’s brochures prepared by property sellers. Each player aims to maximize his probability of being chosen by the searcher, leading to a (constant-sum) competition among them.

Our Results. In this work, we stand in the searcher’s shoes and look to understand how robust the classic prophet inequalities are to the above strategic player behaviors. Like most previous work in this space, we restrict our attention to threshold stopping policies. We start by characterizing players’ optimal information revealing strategies. Given any threshold stopping policy with threshold T , the optimal information strategy of player i with prior reward distribution H_i has the following clean threshold structure (albeit using a threshold different from T): there exists a reward cutoff t_i such that player i simply reveals whether $X_i \geq t_i$ or $X_i < t_i$. Moreover, t_i satisfies $\mathbb{E}[X_i | X_i \geq t_i] = T$. In words, player i simply pools all the “good rewards” together to make their expectation barely pass the threshold T .⁷ This characterization allows us to reduce players’ strategic behaviors to a related prophet inequality problem with binary reward supports. Our later analysis hence only needs to further investigate how well the thresholds for classic prophet inequalities perform in this (related) additional problem.

Armed with the above characterization, next we turn to analyze the intrinsic robustness of the classic prophet inequality. We say a threshold stopping policy with threshold T is α -strategically-robust if it retains the $1/2$ -approximation in the classic setting and meanwhile also can achieve α -approximation in the strategic scenario when all the players optimally reveal information. Our first main result is that, for arbitrary reward distributions, the threshold policy with T equaling the *half expected max* threshold of Kleinberg and Weinberg [36] is $\frac{1-1/e}{2}$ -robust (Theorem 4.1). Moreover,

⁶We refer to the one that controls the reward information as the “player”, and refer to the one that decides when to stop as the “searcher”.

⁷More formally, these response strategies form a subgame perfect Nash equilibrium among the sequential-move game of players (which turns out to be unique as far as players’ utilities are concerned) (see Remark 3.2).

this competitive ratio of $\frac{1-1/e}{2} \approx 0.316$ is the best among all known threshold policies that can secure the $1/2$ -approximation in the classic non-strategic setting (Proposition 4.2). This suggests that this well-studied threshold policy can also perform robustly well even when players are all strategic, illustrating the intrinsic robustness of the classic prophet inequality.

When the reward distributions are identical – referred to as *IID distributions* which have been extensively studied in literature [41, 18, 6, 30, 1, 33] – we show that there exists a threshold in the range of known thresholds (Definition 2.2) that is $1/2$ -robust (Theorem 5.1). Moreover, this $1/2$ approximation ratio is optimal among all possible threshold stopping policies for any IID distributions (Proposition 5.2). Finally, when the reward distributions are not identical but log-concave, under certain regularity conditions we show that any threshold between the expected max and the median of the highest reward [42] is $1/2$ -robust (Theorem 5.4). Note that log-concave reward distributions have been considered in previous works on prophet inequality (see, e.g., [6, 18]); moreover, it is satisfied by a wide range of distributions (e.g., normal, uniform, Gamma, Beta, Laplace, etc., [8]) and is a widely adopted assumption in algorithmic game theory (see, e.g., [15]).

Additional Related Work. Prophet inequality is a fundamental problem in optimal stopping theory which was introduced in the 70s [38, 37]. Recently there has been a growing interest in prophet inequalities, generalizing the problem to different settings [2, 7, 18, 21, 22, 36, 17, 4, 6, 20]. Our discussion here cannot do justice to its rich literature; hence, we refer interested readers to the recent survey by Lucier [40], Correa et al. [19] for comprehensive overviews of recent developments and its connections to economic problems. This work takes an informational perspective and associate each distribution in the prophet inequality with a strategic player that strategically signals reward information to the searcher. We follow the information design literature [35, 10] and allow players to design signaling schemes to influence the searcher’s beliefs about the realized rewards. In our considered game, players are competing with each other for the selection of the searcher. The players’ game thus also shares similarity with competitive information design [26, 28, 5, 23, 34].

Conceptually, our work also relates to the rapidly growing literature on strategic machine learning (see, e.g., [45, 32, 29, 24, 43, 46, 9]), which studies learning from strategic data providers. We also study similar strategic reward providers, albeit in a different algorithmic problem, i.e., optimal stopping. More generally, our work subscribes to the literature on information design in sequential decision-making. In particular, our paper relates to the recently increased interests on using online learning approaches to study the regret minimization when an information-advantaged player repeatedly interacts with an information-disadvantaged player [11, 14, 25, 44, 47] without knowing their preferences. Instead of focusing on regret minimization, we use the competitive ratio (as conventionally used in prophet inequalities) to measure the searcher’s policy. It is worth mentioning that [30] also studies strategic information revealing in prophet inequality problems. In their setting, a centralized player strategically discloses reward information to the searcher, while players in our setting form a decentralized game and each player acts on his own to maximize his payoff.

2 Preliminary

In this section, we first revisit the formulation of the classic prophet inequality problem, and then formally introduce our setting as its natural variant where each distribution is associated with a strategic player who will be strategically signaling their reward information.

Classic prophet inequality. In standard settings, a searcher faces a finite sequence of known distributions $H_{1:N} \triangleq (H_i)_{i \in [N]}$ of N non-negative independent random variables. The outcomes (i.e., the rewards) $X_i \sim H_i$ ⁸ for $i \in [N]$ are revealed sequentially to the searcher for $i = 1, \dots, N$. Let $\lambda_i \triangleq \mathbb{E}_{H_i}[X_i]$ denote the mean reward for distribution H_i . Upon seeing a reward, the searcher decides whether to accept the observed reward and stop the process, or irrevocably move on to the next reward. The searcher’s goal is to maximize the expected accepted reward. The searcher’s expected payoff cannot be more than that of a prophet who knows in advance the realizations of the rewards, namely, X_1, \dots, X_N . We denote by $\text{OPT} \triangleq \mathbb{E}_{H_1, \dots, H_N} [\max_{i \in [N]} X_i]$ the prophet value. Given a stopping policy q , let $X^{(q)}$ be the accepted reward. The stopping policy q is said to be α -approximate (of the prophet value) if the following holds for the searcher’s expected payoff for

⁸Here H_i denotes the cumulative distribution function (CDF) of the reward X_i

any input distributions $H_{1:N}$:

$$\mathbb{E}\left[X^{(q)}\right] \geq \alpha \cdot \text{OPT}.$$

The above statement is often referred to as the prophet inequality. For any stopping policy q , the largest constant $\alpha \in (0, 1]$ that satisfies the above inequality is named the competitive ratio. Classic results of [38, 37] elegantly show that there exist simple threshold stopping policies that achieve the competitive ratio $1/2$, and moreover the ratio of $1/2$ is tight.⁹

Definition 2.1 (Threshold Stopping Policy). *A threshold stopping policy is an online algorithm which pre-computes a threshold value T as a function of the distributions $H_{1:N}$, and then accepts the first reward X_i whose value is no smaller than the threshold, i.e., $X_i \geq T$.*

It is well-known that multiple thresholds can achieve the optimal $1/2$ -approximation ratio.

Definition 2.2 ($1/2$ -approximation Threshold Spectrum). *Let $T_{\text{KW}} \triangleq \mathbb{E}[\max_i X_i]/2$ (Kleinberg and Weinberg [36]), let T^* satisfy $T^* = \sum_{i \in [N]} \mathbb{E}[(X_i - T^*)^+]$, and let $T_{\text{SC}} \triangleq \sup\{t : \Pr[\max_i X_i \geq t] \geq 1/2\}$ (Samuel-Cahn [42]). Then any threshold stopping policy with threshold $T \in [T_{\text{KW}}, \max\{T_{\text{SC}}, T^*\}]$ guarantees $1/2$ -approximation to the prophet value.*¹⁰

Prophet inequality with strategic reward signaling. In the strategic setting, each distribution H_i may be associated with a strategic player that governs how much information about the realized reward he would like to reveal to the searcher once the searcher reaches his distribution. Formally, upon arriving at the reward distribution H_i , the searcher does not directly observe the realized reward $X_i \sim H_i$. Instead, she observes an information signal, designed by the player i with distribution H_i , that is correlated with the reward X_i . We follow the literature in information design [35] to model this strategic reward signaling: Each player i chooses a signaling scheme $\phi_i(\cdot | x) \in \Delta(\Sigma_i)$, where Σ_i is a measurable signal space and $\phi_i(\sigma | x) \in [0, 1]$ specifies the conditional probability of a signal $\sigma \in \Sigma_i$ that will be sent to the searcher when the reward $X_i = x \sim H_i$ is realized. Notice that, upon seeing a signal $\sigma \sim \phi_i(\cdot | x)$, together with the prior information H_i from which the reward is realized, the searcher can update her Bayesian belief about the underlying realized value X_i , and then decides whether to stop and choose player i , or reject i to continue her search.

Each player i is competing with each other for the final selection from the searcher. Namely, a player obtains payoff 1 if his reward is accepted and payoff 0 if his reward is not accepted.¹¹ Each player's goal is to design an information revealing strategy that maximizes his probability of being chosen by the searcher. We below give two simple examples of information revealing strategies.

Example 2.3 (Examples of Information Strategies). (1) *No information strategy: the signal is completely uninformative (e.g., the distribution $s\phi_i(\cdot | x)$ is a Dirac delta function on a single signal, i.e., $|\Sigma_i| = 1$), hence the searcher infers an expected reward of $\lambda_i = \mathbb{E}[X_i]$ as her perceived reward from player i ;* (2) *Full information revealing strategy: the signal perfectly reveals player i 's value to the searcher (i.e., $\phi_i(\sigma \equiv x | x) = 1$ for every realized $X_i = x$, and $\Sigma_i = \text{supp}(H_i)$)*

In this work, standing in the searcher's shoes, we are interested in how threshold stopping policies perform under players' optimal strategic reward signaling.

Game Timeline. The timeline of our prophet inequality with strategic players problem, where the searcher employs a threshold stopping policy, can be detailed as follows: (1) Knowing $H_{1:N}$, the searcher first announces a threshold stopping policy with threshold T that is a function of $H_{1:N}$; (2) Knowing threshold T , each player then picks a signaling scheme (also known as an *experiment* in economics literature [35, 10]) to reveal partial information about the underlying reward; and (3) The searcher learns all players' information strategies, and then conducts a search based on her threshold

⁹See [40, 19] for the hardness example for the $1/2$ -approximation.

¹⁰There also exists other thresholds that could lead to $1/2$ -approximation, e.g., any $T \in [\min\{\underline{T}_{\text{SC}}, T_{\text{KW}}\}, T_{\text{KW}}]$ where $\underline{T}_{\text{SC}} \triangleq \inf\{t : \Pr[\max_i X_i \geq t] \geq 1/2\}$. However, using these thresholds requires to modify the policy defined in Definition 2.1 to be a *strict* stopping policy to obtain $1/2$ -approximation (i.e., searcher accepts a first reward that is strictly larger than threshold). Moreover, under strict stopping policy, there may exist no Nash equilibrium among players' game that we formulate shortly.

¹¹All of our results hold if each player i has payoff v_i if his reward is accepted and payoff u_i if his reward is not accepted as long as $v_i > u_i$.

stopping policy with threshold T (i.e., accepting the first player whose posterior mean of his reward distribution, given his revealed information signal, exceeds the threshold T). The assumption that the searcher knows information strategies as well as realized signals is commonly adopted in the information design literature [35, 10], and is also well motivated for the domains of our interest. For instance, when startups persuade VCs or property sellers persuade buyers, the signaling scheme could correspond to startups' product exhibitions or sellers' promotion brochures which determine what information the searcher could see. Misreporting realized signals corresponds to revealing untrue information, which not only violates regulation policies and but also causes the players to lose credibility in the long term.

Slightly abusing the notation, we also use $X^{(q)}$ to denote the accepted reward given the searcher's stopping policy q under the strategic reward signaling, and let $u^s(q) \triangleq \mathbb{E}[X^{(q)}]$ be the searcher's expected payoff. Notice that here the expectation is not only over the randomness of the distributions $H_{1:N}$, but also the information strategies $\{\phi_i(\cdot | x)\}$. Anticipating the players' strategic behavior, the searcher wants a stopping policy that can still guarantee a good performance competing against the prophet value OPT .

As mentioned earlier, we are particularly interested in how previously studied threshold stopping policies described in Definition 2.2 perform under the players' strategic reward signaling. To formalize our goal, we introduce the following notion of strategic robustness.

Definition 2.4 (α -strategically)-robust Stopping Policies). *A stopping policy q is α -robust if it achieves α -approximation to the OPT when players are strategically signaling their rewards, and it remains a $1/2$ -approximation in the standard non-strategic setting.*

3 Warm-up: Characterizing Optimal Information Revealing Strategy

We start our analysis by showing that when the searcher adopts a threshold stopping policy (Definition 2.1), each player's optimal information revealing strategy admits clean characterizations.

Proposition 3.1 (Optimal Information Revealing Strategy). *Given a threshold stopping policy as in Definition 2.1 with threshold T , for each player i :*

- if $T \leq \lambda_i$, then player i 's optimal information revealing strategy is the no information strategy;
- if $T > \lambda_i$, then player i 's optimal information revealing strategy is threshold signaling and determined by a cutoff t_i that satisfies $T = \mathbb{E}[X_i | X_i \geq t_i] = \int_{t_i}^{\infty} x dH_i(x) / (1 - H_i(t_i))$. That is, player i 's optimal signaling scheme sends one of two signals: $X_i \geq t_i$ or $X_i < t_i$.¹²

We highlight the intuition behind Proposition 3.1 below. Under a threshold stopping policy, every player maximizes his utility by maximizing the probability that the signal's posterior expected reward is at least T (hence is selected). When the stopping threshold $T \leq \lambda_i$, this probability is 1 and hence maximized when player i simply reveals no information. When $T > \lambda_i$, this probability is maximized when player i blends the highest rewards together to form a posterior mean just equal to T , which is exactly the scheme described in Proposition 3.1.

Remark 3.1 (Addressing Point Masses). *For ease of presentation, Proposition 3.1 assumes the distribution H_i is continuous. When H_i has point masses, Proposition 3.1 still holds, but with a more subtle description of the pooling cutoff t_i . We provide more detailed discussions in the full version of the paper.*

Remark 3.2. *In game-theoretic terminology, Proposition 3.1 characterizes the subgame Perfect Nash equilibrium (SPNE) for the multi-player sequential game induced by any threshold stopping policy that the searcher commits to. This SPNE happens to enjoy simple structures; indeed, each player's equilibrium strategy is only a function of the threshold T and not on other players' strategies or their order. This clean characterization is a consequence of the simple structure of (static) threshold policies. If the searcher's stopping policy is instead dynamic (i.e., allowing the decision of player i to depend on previously realized rewards), SPNE is also well-defined but will adopt significantly more complex structures. While analyzing the SPNE under these dynamic policies may also be interesting, it is beyond the scope of this work since our focus is to study the power of static*

¹²In the corner case when T is larger than the upper bound of player i 's distribution, player i will never get chosen and thus they have no strategy

threshold policies, which is also a central theme in the study of prophet inequalities. Finally, we note that the optimal strategies characterized in Proposition 3.1 are not unique but they all result in the same utility for every player. This is because player i is actually indifferent on how to disclose the reward when $X_i \leq t_i$, leading to many different but utility-equivalent information strategies.

Equivalent Representation of Optimal Information Revealing Strategy. When $T > \lambda_i$, player i 's optimal information revealing strategy described in Proposition 3.1 pools all the rewards $X_i \sim H_i$ above t_i together to forge a conditional mean value of T , and pools the remaining smaller rewards into another signal. We hence also refer to t_i as the pooling cutoff. This threshold signaling scheme can be equivalently represented as a binary-support distribution G_i supported on two realizations corresponding to the two signals respectively:

$$\Pr_{x \sim G_i}[x = T] = 1 - H_i(t_i), \quad \Pr_{x \sim G_i}[x = a_i] = H_i(t_i) \text{ where } a_i \triangleq \frac{\lambda_i - T(1 - H_i(t_i))}{H_i(t_i)}. \quad (1)$$

In the literature of information design, this distribution G_i is also known as a mean-preserving contraction of prior reward distribution H_i [27, 12, 13].

Viewing the player i 's optimal information strategy as the distribution G_i , one can simplify the interaction between the searcher and player i as the following when the searcher visits player i : a random reward $X_i \sim G_i$ is realized, and the searcher stops if and only if $X_i \geq T$. With this observation, one can also reduce the original interaction to the following simplified protocol: (1) the searcher first decides a stopping threshold T ; (2) each player i chooses the distribution G_i as in Equation (1) according to his prior H_i and the threshold T ,¹³ and (3) the searcher visits each G_i sequentially and stops when the first realized reward $X_i \geq T$ where $X_i \sim G_i$.

We emphasize that in the above reduction, given a threshold stopping policy with threshold T , the searcher's expected payoff under the strategic reward signaling can be computed as $u^s(T) = \mathbb{E}[X^{(q)}]$ where the expectation is over distributions $G_{1:N}$ and $X^{(q)}$ is the first realized $X_i \sim G_i$ such that $X^{(q)} \geq T$.

4 Achieving $\frac{1-1/e}{2}$ -robustness for Arbitrary Distributions

In this section, we show that for any distributions $H_{1:N}$, there exists a $\frac{1-1/e}{2}$ -(-strategically)-robust threshold stopping policy using a threshold within the spectrum in Definition 2.2. The main result in this section is stated below.

Theorem 4.1. *For any distributions $H_{1:N}$, a threshold stopping policy with threshold $T = T_{\text{KW}}$ is $\frac{1-1/e}{2}$ -robust.*

From Definition 2.2, we know that using the threshold stopping policy with threshold T_{KW} can achieve the optimal $1/2$ -approximation in the classic prophet inequality for any distributions $H_{1:N}$. Theorem 4.1 above shows another desired property of the threshold T_{KW} : it can achieve $\frac{1-1/e}{2}$ -approximation even when distributions $H_{1:N}$ are strategically signaling their rewards, thus establishing its $\frac{1-1/e}{2}$ -robustness. Given the optimality of the threshold T_{KW} in the non-strategic setting, it would be intriguing to ask whether this threshold T_{KW} can also achieve $1/2$ -approximation under the strategic setting, or if there exists a threshold within the spectrum in Definition 2.2 that can achieve $1/2$ -approximation under the strategic setting. Below we show that the answer is No. In fact, any threshold stopping policy using a threshold from the spectrum in Definition 2.2 cannot achieve $(\frac{1-1/e}{2} + \varepsilon)$ -approximation for any $\varepsilon > 0$ under strategic reward signaling.

Proposition 4.2 (Tightness of Theorem 4.1). *There exist distributions $H_{1:N}$ such that no threshold from the spectrum in Definition 2.2 can achieve α -robustness where $\alpha > \frac{1-(1-1/(N-1))^{N-1}}{2}$.*

Notice that $\lim_{N \rightarrow \infty} \frac{1-(1-1/(N-1))^{N-1}}{2} = \frac{1-1/e}{2}$. Thus, the above Proposition 4.2 establishes the tightness of the results in Theorem 4.1.

Remark 4.1. *We point out a subtlety in the above lower bound, which leads to an intriguing open problem. Proposition 4.2 shows that any threshold within the spectrum in Definition 2.2 cannot*

¹³When $T \leq \lambda_i$, distribution G_i is a point mass at mean value λ_i .

achieve $(\frac{1-1/e}{2} + \varepsilon)$ -approximation under strategic reward signaling. However, this does not rule out the possibility of having a threshold outside that spectrum that achieves $1/2$ -robustness. This is an interesting open question to resolve, though necessarily challenging since it is even already quite non-trivial to prove $1/2$ -approximation in the non-strategic setting for thresholds outside the spectrum of Definition 2.2, let alone achieving $1/2$ -approximation simultaneously in both worlds.

Theorem 4.1 holds for all distributions, regardless of being discrete or continuous. In the remainder of this subsection, we present the proof of Theorem 4.1 only for the continuous distribution case. The proof of this case carries our core ideas but is cleaner to present.

4.1 (Partial) Proof of Theorem 4.1: The Case of Continuous Distributions

Our proof starts by upper bounding the prophet value OPT .¹⁴

Lemma 4.3 (Upper Bounding OPT via Pooling Cutoffs). *Given a threshold stopping policy with threshold T , let t_i be the pooling cutoff for each player i defined as in Proposition 3.1. Let $I \triangleq \arg \max_{i \in [N]} t_i$, then we have $\text{OPT} \leq H_I(t_I)t_I + \sum_i T(1 - H_i(t_i))$.*

Proof of Lemma 4.3. Let us fix a threshold T , and let t_i be the pooling cutoff for player i defined in Proposition 3.1. Define $b_i \triangleq (X_i - t_i)^+$. By definition, for each i , we have $X_i \leq t_i + b_i$. Thus,

$$\begin{aligned} \text{OPT} &= \mathbb{E}_{H_{1:N}} \left[\max_i X_i \right] \leq \max_i t_i + \sum_i \mathbb{E}_{H_i} [b_i] \\ &\stackrel{(a)}{\leq} \max_i t_i + \sum_i (T - t_i) \cdot (1 - H_i(t_i)) \\ &\stackrel{(b)}{\leq} t_I H_I(t_I) + \sum_i T(1 - H_i(t_i)), \end{aligned}$$

where inequality (a) follows from the definition of pooling cutoff t_i in Proposition 3.1 and inequality (b) follows from the definition of I , $t_i \geq 0$, and rearranging the terms. \square

With the above upper bound of prophet value, we have the following results:

Lemma 4.4. *Let T^\dagger satisfy $\prod_{i=1}^N H_i(t_i^\dagger) = (\frac{N-1}{N})^N$ where t_i^\dagger is defined in Proposition 3.1 with threshold T^\dagger , then searcher's expected payoff $u^s(T^\dagger) \geq T^\dagger \left(1 - \prod_{i=1}^N H_i(t_i^\dagger) \right) \geq \frac{1 - (\frac{N-1}{N})^N}{2} \cdot \text{OPT}$.¹⁵*

Proof of Lemma 4.4. Given a stopping threshold T , by Proposition 3.1, we can lower bound the searcher's expected payoff as follows: $u^s(T) \geq T \cdot \left(1 - \prod_{i=1}^N H_i(t_i) \right)$. Thus, together with

¹⁴All our analysis in the main text implicitly consider the case where $\max_i \lambda_i < \text{OPT}/2$. If we have $\max_i \lambda_i \geq \text{OPT}/2$, then the searcher can simply choose the threshold $T = \text{OPT}/2$ which lies in the range of the thresholds defined in Definition 2.2 to obtain payoff $u^s(T) \geq \text{OPT}/2$. To see this, by Proposition 3.1, if $\max_i \lambda_i \geq \text{OPT}/2 = T$, then there exists at least one player j whose $\lambda_j \geq \text{OPT}/2$ will choose the no information revealing strategy, and the searcher surely obtains a payoff no smaller than $\text{OPT}/2$.

¹⁵Note that for non-continuous distributions T^\dagger always exists but is defined a bit differently, please see the full version of the paper for more details.

Lemma 4.3, we have

$$\begin{aligned}
\frac{u^s(T)}{\text{OPT}} &\geq \frac{T \cdot \left(1 - \prod_{i=1}^N H_i(t_i)\right)}{\text{OPT}} \\
&\stackrel{(a)}{\geq} \frac{T \cdot \left(1 - \prod_{i=1}^N H_i(t_i)\right)}{t_I H_I(t_I) + \sum_i T(1 - H_i(t_i))} \\
&\stackrel{(b)}{\geq} \frac{1 - \prod_{i=1}^N H_i(t_i)}{H_I(t_I) + \sum_i (1 - H_i(t_i))} \\
&\stackrel{(c)}{\geq} \frac{1 - \prod_{i=1}^N H_i(t_i)}{N + 1 - \sum_i H_i(t_i)} \\
&\stackrel{(d)}{\geq} \frac{1 - \prod_{i=1}^N H_i(t_i)}{N + 1 - (\prod_{i=1}^N H_i(t_i))^{\frac{1}{N}}},
\end{aligned}$$

where inequality (a) is by Lemma 4.3; inequality (b) is by the fact that $t_i \leq T$ for all $i \in [N]$; inequality (c) is due to $H_I(t_I) \leq 1$; and inequality (d) is by the AM-GM inequality. Now consider the function $f(x) \triangleq \frac{1-x}{N+1-Nx^{\frac{1}{N}}}$ over $x \in [0, 1]$. By choosing $x^\dagger = (\frac{N-1}{N})^N$, we have $f(x^\dagger) = \frac{1-(\frac{N-1}{N})^N}{2}$. This implies that we have $T^\dagger \left(1 - \prod_i H_i(t_i^\dagger)\right) \geq \frac{1-(\frac{N-1}{N})^N}{2} \cdot \text{OPT}$. \square

With the above Lemma 4.4, we are now ready to prove Theorem 4.1:

Proof of Theorem 4.1. From Lemma 4.4, we showed

$$T^\dagger \left(1 - \prod_i H_i(t_i^\dagger)\right) \geq \text{OPT} \cdot \frac{1 - (\frac{N-1}{N})^N}{2},$$

where t_i^\dagger is the pooling cutoff defined in Proposition 3.1 with the threshold T^\dagger . By definition of T^\dagger , we have $\prod_i H_i(t_i^\dagger) = (\frac{N-1}{N})^N \leq 1/e$. Thus we can deduce that $T^\dagger \geq \text{OPT}/2 = T_{\text{KW}}$. Now let t_i^\ddagger be the pooling cutoff when the searcher uses the stopping threshold T_{KW} . Then we have

$$u^s(T_{\text{KW}}) \geq T_{\text{KW}} \cdot \left(1 - \prod_i H_i(t_i^\ddagger)\right) \stackrel{(a)}{\geq} T_{\text{KW}} \cdot \left(1 - \prod_i H_i(t_i^\dagger)\right) \geq \text{OPT} \cdot \frac{1 - 1/e}{2}$$

where inequality (a) is due to $T^\dagger \geq T_{\text{KW}}$, and thus we have $t_i^\dagger \geq t_i^\ddagger$. \square

5 Achieving $\frac{1}{2}$ -robustness for Special Distributions

The preceding section showed the $(1 - 1/e)/2$ -robustness of the T_{KW} -threshold stopping policy for arbitrary reward distributions. In this section, we show that this ratio can be improved to $1/2$ -robustness when the distributions $H_{1:N}$ satisfy certain conditions: (1) IID distributions – all reward distributions are identical, namely, $H \equiv H_i$ for all $i \in [N]$ (see Section 5.1); and (2) Log-concave distributions – reward distribution H_i has log-concave density (see Section 5.2). We also show that $1/2$ -robustness is tight under IID distributions.

5.1 IID Distributions

Our main findings for IID distributions are stated below:

Theorem 5.1. *For any distributions $H_{1:N}$ where $H \equiv H_i, \forall i \in [N]$, a threshold stopping policy with threshold $T = T^*$ is $\frac{1}{2}$ -robust where T^* is defined in Definition 2.2.*

For IID distributions, we show that the searcher is able to achieve a better robustness approximation ratio compared to arbitrary distributions. Below we argue that this $1/2$ -robustness is tight in the sense that there exists no threshold policy that can achieve better robustness approximation ratios.

Proposition 5.2 (Tightness of Theorem 5.1). *There exist IID distributions such that there exists no threshold stopping policy that can achieve α -robustness where $\alpha > \frac{1}{2} + \varepsilon$ for any $\varepsilon > 0$.*

We note that Proposition 5.2 is a slightly stronger lower bound than Proposition 4.2 as it rules out the possibility for *all* possible threshold stopping policies, being within the spectrum in Definition 2.2 or beyond. We prove Proposition 5.2 by constructing a hard instance. In particular, we construct IID distributions with binary support. With this instance, we show that any threshold policy that achieves competitive ratio at least $1/2$ in non-strategic setting will have competitive ratio approaching to 0 in the strategic setting. Please see

A crucial requirement in our robustness study so far is that we insist that the threshold policy should, first of all, remains an $1/2$ -approximation in the non-strategic setting¹⁶, conditioned on which we look for additional guarantee for in the strategic setting. We conclude this section by pointing that if one was willing to give up the $1/2$ -approximation in the non-strategic setting, then it is indeed possible to have a threshold policy that achieves better approximation (specifically, an $(1 - 1/e)$ -approximation) in the strategic setting. While this does not satisfy our robustness requirement, it is useful to note.

Corollary 5.3. *For any IID distributions, there exists a threshold stopping policy that is $(1 - 1/e)$ -approximation under strategic reward signaling. Moreover, there exist IID distributions such that no threshold stopping policy can achieve $(1 - 1/e + \varepsilon)$ -approximation for any $\varepsilon > 0$.*

5.2 Log-concave Heterogeneous Distributions

In this subsection, we show that when the distributions $H_{1:N}$ satisfy certain regularity assumptions, there exist threshold stopping policies with thresholds from Definition 2.2 that can also achieve $1/2$ -robustness. The main result in this section is stated as follows:

Theorem 5.4. *For $\alpha, \beta > 0$, let $F_{\alpha, \beta}$ be the family of distributions with log-concave probability density functions f on support $[0, 1]$ such that $f(1) \geq \alpha$ and $f'(1) \geq -\beta$. If the distributions H_1, \dots, H_N are all from $F_{\alpha, \beta}$ and $N \geq 1 + \frac{\beta}{\alpha^2}$, then we always have $2 \cdot T_{\text{KW}} \leq T_{\text{SC}}$ and any threshold T satisfying $2 \cdot T_{\text{KW}} \leq T \leq T_{\text{SC}}$ is $1/2$ -robust.*

A few remarks on the assumptions in Theorem 5.4 are worth mentioning. First, log-concavity of probability density functions¹⁷ is a commonly used assumption; they include but are not limited to: normal, beta, gamma, and exponential distributions. The restriction to support on $[0, 1]$ is for normalization reasons, hence without loss of generality. The main non-trivial restriction is that this result holds when the number of players N is large enough, formally $N \geq 1 + \frac{\beta}{\alpha^2}$. This condition becomes less restrictive when α (lower bounding $f(1)$) becomes larger and/or β (upper bounding $-f'(1)$) becomes smaller. These together, intuitively, imply that f decreases slowly within $[0, 1]$.

Define $\bar{H}(x) \triangleq \prod_{i=1}^N H_i(x)$. Theorem 5.4 follows directly from the following Lemma 5.5 and Lemma 5.6.

Lemma 5.5. *If \bar{H} is convex, then $2 \cdot T_{\text{KW}} \leq T_{\text{SC}}$ and any threshold stopping policy with the threshold T satisfying $2 \cdot T_{\text{KW}} \leq T \leq T_{\text{SC}}$, where $T_{\text{KW}}, T_{\text{SC}}$ are defined as in Definition 2.1, is $1/2$ -robust.*

Lemma 5.6. *If the distributions H_1, \dots, H_N are all from $F_{\alpha, \beta}$ and $N \geq 1 + \frac{\beta}{\alpha^2}$, then \bar{H} is convex.*

6 Conclusion and Future Directions

In this paper, we study a variant of the prophet inequality problem where each random variable is associated with a strategic player who can strategically signal their reward to the searcher. We first fully characterize the optimal information strategy of each player, then we show the threshold stopping policies that can perform robustly well under both the strategic and non-strategic settings.

Our novel consideration of natural strategic manipulations in prophet inequalities open the door for many interesting future directions. First, it is interesting to see if we can improve the $\frac{1-1/e}{2}$ -robustness guarantee for arbitrary reward distributions by using not commonly used thresholds outside the spectrum of Definition 2.2. This may be technically challenging since finding a threshold

¹⁶Note that in non-strategic setting, unlike the case for arbitrary distributions where there exists no policy that can achieve better than $1/2$ -approximation, for IID distributions, $(1 - 1/e)$ -approximation can be achieved by fixed threshold together with some careful probabilistic tie-breaking rule [3] (or 0.7451-approximation with an adaptive threshold policy [39]). However, in our work, we focus on fixed threshold policy with deterministic tie-breaking rule where $1/2$ -approximation is still optimal for IID distributions.

¹⁷A probability density function $f : \mathbb{R} \rightarrow \mathbb{R}^+$ is log-concave if $\log(f)$ is concave.

outside this spectrum with optimal $1/2$ -approximation for the classic non-strategic prophet inequality is already non-trivial. Another interesting direction is to go beyond threshold policies. That leads to a different research theme, not about robustness of the classic prophet inequality, but rather in the search of potentially much more complex best-of-both-world policies. Our preliminary result reveals that a dynamic threshold stopping policy (using different thresholds for different players) has the potential to help the searcher to achieve more than $\frac{1-1/e}{2}$ -approximation under the strategic setting. However, it is still unclear whether $1/2$ -robustness is achievable under threshold stopping policies, no matter whether it is static threshold or dynamic threshold. Finally, even if one only cares about the performance under the strategic reward signaling (and ignore the non-strategic world), it is still unclear which static threshold stopping policy achieves the highest competitive ratio. The authors in [20] study a model with a very similar mathematical structure and proved that this best upper bound is strictly less than $1/2$ in Section 4.1, but an interesting direction is improving this bound further. We note that all our constructed examples in the hardness results (e.g., Proposition 4.2 and Proposition 5.2) do not rule out the existence of such $1/2$ -approximation threshold stopping policies.

Acknowledgment. Haifeng Xu is supported in part by the NSF Award CCF-2303372, Army Research Office Award W911NF-23-1-0030 and ONR Award N00014-23-1-2802. We thank the anonymous reviewers for constructive and helpful comments.

References

- [1] Melika Abolhassani, Soheil Ehsani, Hossein Esfandiari, MohammadTaghi Hajiaghayi, Robert Kleinberg, and Brendan Lucier. Beating 1-1/e for ordered prophets. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 61–71, 2017.
- [2] Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014.
- [3] Nick Arnosti and Will Ma. Tight guarantees for static threshold policies in the prophet secretary problem. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, EC '22, page 242, 2022.
- [4] Makis Arsenis and Robert Kleinberg. Individual fairness in prophet inequalities. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 245–245, 2022.
- [5] Pak Hung Au and Keiichi Kawai. Competitive information disclosure by multiple senders. *Games and Economic Behavior*, 119:56–78, 2020.
- [6] Pablo D Azar, Robert Kleinberg, and S Matthew Weinberg. Prophet inequalities with limited information. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1358–1377. SIAM, 2014.
- [7] Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In *Symposium on Discrete Algorithms (SODA'07)*, pages 434–443, 2007.
- [8] Mark Bagnoli and Ted Bergstrom. Log-concave probability and its applications. In *Rationality and Equilibrium: A Symposium in Honor of Marcel K. Richter*, pages 217–241. Springer, 2006.
- [9] Yahav Bechavod, Katrina Ligett, Steven Wu, and Juba Ziani. Gaming helps! learning from strategic interactions in natural dynamics. In *International Conference on Artificial Intelligence and Statistics*, pages 1234–1242. PMLR, 2021.
- [10] Dirk Bergemann and Stephen Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522, 2016.
- [11] Martino Bernasconi, Matteo Castiglioni, Andrea Celli, Alberto Marchesi, Francesco Trovò, and Nicola Gatti. Optimal rates and efficient algorithms for online bayesian persuasion. In *International Conference on Machine Learning*, pages 2164–2183. PMLR, 2023.
- [12] David Blackwell. Comparison of experiments. In *Proceedings of the second Berkeley symposium on mathematical statistics and probability*, volume 2, pages 93–103. University of California Press, 1951.
- [13] David A Blackwell and Meyer A Girshick. *Theory of games and statistical decisions*. Courier Corporation, 1979.
- [14] Matteo Castiglioni, Andrea Celli, Alberto Marchesi, and Nicola Gatti. Online bayesian persuasion. *Advances in Neural Information Processing Systems*, 33, 2020.
- [15] Shuchi Chawla, Jason D Hartline, and Robert Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, pages 243–251, 2007.
- [16] Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320, 2010.
- [17] Shuchi Chawla, Nikhil Devanur, and Thodoris Lykouris. Static pricing for multi-unit prophet inequalities. *Operations Research*, 2023.
- [18] José 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*, pages 169–186, 2017.

- [19] Jose Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Recent developments in prophet inequalities. *ACM SIGecom Exchanges*, 17(1):61–70, 2019.
- [20] Yuan Deng, Vahab Mirrokni, and Hanrui Zhang. Posted pricing and dynamic prior-independent mechanisms with value maximizers. *Advances in Neural Information Processing Systems*, 35:24158–24169, 2022.
- [21] Nikhil R Devanur, Kamal Jain, Balasubramanian Sivan, and Christopher A Wilkens. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 29–38, 2011.
- [22] Nikhil R Devanur, Balasubramanian Sivan, and Yossi Azar. Asymptotically optimal algorithm for stochastic adwords. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 388–404, 2012.
- [23] Bolin Ding, Yiding Feng, Chien-Ju Ho, Wei Tang, and Haifeng Xu. Competitive information design for pandora’s box. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 353–381. SIAM, 2023.
- [24] Jinshuo Dong, Aaron Roth, Zachary Schutzman, Bo Waggoner, and Zhiwei Steven Wu. Strategic classification from revealed preferences. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 55–70, 2018.
- [25] Yiding Feng, Wei Tang, and Haifeng Xu. Online bayesian recommendation with no regret. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 818–819, 2022.
- [26] Matthew Gentzkow and Emir Kamenica. Competition in persuasion. *The Review of Economic Studies*, 84(1):300–322, 2016.
- [27] Matthew Gentzkow and Emir Kamenica. A rothschild-stiglitz approach to bayesian persuasion. *American Economic Review*, 106(5):597–601, 2016.
- [28] Matthew Gentzkow and Emir Kamenica. Bayesian persuasion with multiple senders and rich signal spaces. *Games and Economic Behavior*, 104:411–429, 2017.
- [29] Ganesh Ghalme, Vineet Nair, Itay Eilat, Inbal Talgam-Cohen, and Nir Rosenfeld. Strategic classification in the dark. In *International Conference on Machine Learning*, pages 3672–3681. PMLR, 2021.
- [30] Niklas Hahn, Martin Hoefer, and Rann Smorodinsky. Prophet inequalities for bayesian persuasion. In *IJCAI*, pages 175–181, 2020.
- [31] Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *AAAI*, volume 7, pages 58–65, 2007.
- [32] Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. Strategic classification. In *Proceedings of the 2016 ACM conference on innovations in theoretical computer science*, pages 111–122, 2016.
- [33] Theodore P Hill and Robert P Kertz. Comparisons of stop rule and supremum expectations of iid random variables. *The Annals of Probability*, pages 336–345, 1982.
- [34] Safwan Hossain, Tonghan Wang, Tao Lin, Yiling Chen, David C Parkes, and Haifeng Xu. Multi-sender persuasion—a computational perspective. *arXiv preprint arXiv:2402.04971*, 2024.
- [35] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- [36] Robert Kleinberg and Seth Weinberg. Matroid prophet inequalities. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 123–136, 2012.
- [37] Ulrich Krengel and Louis Sucheston. Semiamarts and finite values. 1977.

- [38] Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces*, 4:197–266, 1978.
- [39] Allen Liu, Renato Paes Leme, Martin Pal, Jon Schneider, and Sivan Balasubramanian. Variable decomposition for prophet inequalities and optimal ordering. *EC '21: Proceedings of the 22nd ACM Conference on Economics and Computation*, page 692, 2021.
- [40] Brendan Lucier. An economic view of prophet inequalities. *ACM SIGecom Exchanges*, 16(1):24–47, 2017.
- [41] Sebastian Perez-Salazar, Mohit Singh, and Alejandro Toriello. The iid prophet inequality with limited flexibility. *arXiv preprint arXiv:2210.05634*, 2022.
- [42] Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent non-negative random variables. *the Annals of Probability*, pages 1213–1216, 1984.
- [43] Ravi Sundaram, Anil Vullikanti, Haifeng Xu, and Fan Yao. Pac-learning for strategic classification. *Journal of Machine Learning Research*, 24(192):1–38, 2023.
- [44] Jibang Wu, Zixuan Zhang, Zhe Feng, Zhaoran Wang, Zhuoran Yang, Michael I Jordan, and Haifeng Xu. Sequential information design: Markov persuasion process and its efficient reinforcement learning. In *Proceedings of the 23th ACM Conference on Economics and Computation*, 2022.
- [45] Hanrui Zhang and Vincent Conitzer. Incentive-aware pac learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5797–5804, 2021.
- [46] Tijana Zrnic, Eric Mazumdar, Shankar Sastry, and Michael Jordan. Who leads and who follows in strategic classification? *Advances in Neural Information Processing Systems*, 34:15257–15269, 2021.
- [47] You Zu, Krishnamurthy Iyer, and Haifeng Xu. Learning to persuade on the fly: Robustness against ignorance. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 927–928, 2021.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The abstract and introduction accurately summarize our results. In the main text, we provide all rigorous statements and proofs for these results.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: They are discussed in conclusion section.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: All assumptions are clearly stated, and proofs are rigorously written.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [NA]

Justification: This is a theory work, no experiment is provided.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: No data and code is used.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: No experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: No experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification: No experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: No violation of NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: The paper studies a theoretical problem, to the best of our knowledge, there is no societal impact of the work that we need to address.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Not applicable to our paper.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: No assets used.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: No new assets as this is a theory work.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This project is not a crowd sourcing project and has no human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: No IRB needed.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

The full version of this paper, including detailed proofs, can be found at <https://arxiv.org/pdf/2409.18269>