

Learning-Augmented Ski Rental with Discrete Distributions: A Bayesian Approach

Bosun Kang^{1,2} Hyejun Park^{1,3} Chenglin Fan^{1,3*}

¹Algorithmic Foundations of Data Science Lab

²Department of Liberal Studies,

Seoul National University, South Korea

³Department of Computer Science and Engineering,

Seoul National University, South Korea

{bosun1043, 1253spring, fanchenglin}@snu.ac.kr

Abstract

We revisit the classic ski rental problem through the lens of Bayesian decision-making and machine-learned predictions. While traditional algorithms minimize worst-case cost without assumptions, and recent learning-augmented approaches leverage noisy forecasts with robustness guarantees, our work unifies these perspectives. We propose a discrete Bayesian framework that maintains exact posterior distributions over the time horizon, enabling principled uncertainty quantification and seamless incorporation of expert priors. Our algorithm achieves prior-dependent competitive guarantees and gracefully interpolates between worst-case and fully-informed settings. Our extensive experimental evaluation demonstrates superior empirical performance across diverse scenarios, achieving near-optimal results under accurate priors while maintaining robust worst-case guarantees. This framework naturally extends to incorporate multiple predictions, non-uniform priors, and contextual information, highlighting the practical advantages of Bayesian reasoning in online decision problems with imperfect predictions.

Introduction

The ski rental problem is a foundational challenge in online algorithms, capturing the fundamental trade-off between short-term flexibility and long-term commitment under uncertainty. A decision-maker must repeatedly choose between renting an item at a per-use cost or purchasing it upfront at a fixed cost, without knowing how long the item will be needed. Originally introduced by Karlin et al. (1994), this problem models a wide range of real-world scenarios, including cloud resource provisioning, equipment leasing, and inventory management.

In the classical setting, no deterministic algorithm can achieve a competitive ratio better than 2, and randomization improves this to $\frac{e}{e-1} \approx 1.582$ (Chrobak et al. 1991). However, these guarantees are worst-case and fail to leverage any structure or prior information about the underlying demand. Recent work on *learning-augmented algorithms* has proposed integrating machine-learned predictions to improve online performance while retaining robustness (Lyk-

ouris and Vassilvitskii 2018; Mitzenmacher 2020; Purohit, Svitkina, and Kumar 2018). Yet existing approaches often rely on point estimates, regime heuristics, or hand-crafted thresholds, which are brittle in the face of prediction noise and difficult to analyze rigorously.

This paper introduces a Bayesian framework for the ski rental problem that is both learning-augmented and fully probabilistic. Our approach maintains a posterior distribution over the unknown time horizon and updates it exactly each day based on observed survival. This enables decisions based on expected utility rather than point forecasts, and provides a principled mechanism for incorporating prior knowledge—be it empirical, structured, or even adversarially incorrect. Throughout the paper we assume the unknown horizon is bounded by a known M .

Contributions

This work makes several key contributions. First, we introduce the first learning-augmented ski rental algorithm based on exact discrete Bayesian inference, removing the need for point estimates or regime switching. Second, we establish theoretical guarantees on competitive ratio and regret under classical priors, showing robustness to prior misspecification. Third, extensive experiments on synthetic and real-world priors confirm consistent outperformance over deterministic, randomized, and prediction-based baselines—even under noisy or adversarial conditions. Our Bayesian formulation is modular and extensible, supporting multiple priors, structured uncertainty, and continuous-time variants, while adapting to multi-modal priors via decisions informed by mode weights and positions.

Technical Overview

Notation. Consider an unknown season length $T \in \{1, \dots, M\}$ with prior mass function $\pi_k \triangleq \Pr[T = k]$ for $k = 1, \dots, M$. Using k to denote a realized trip length (horizon). Let $b > 1$ be the one-time *buy cost* and, unless stated otherwise, the per-day *rental cost* is normalized to 1.

It maintains the posterior over T conditioned on survival ($T \geq t$) and compares the expected remaining rental cost with b .

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

1. **Posterior update:** $\Pr(T = k \mid T \geq t) = \frac{\pi_k}{\sum_{j=t}^M \pi_j}$ for $k \geq t$.
2. **Expected rental cost:** $E_{\text{rent}}(t) = \sum_{k=t}^M \Pr(T = k \mid T \geq t) \cdot (k - t + 1)$.
3. **Decision:** If $b \leq E_{\text{rent}}(t)$, buy; else, rent and proceed to day $t + 1$.

Units. With per-day rental cost normalized to 1, b and $k - t + 1$ are directly comparable:

$$E_{\text{rent}}(t) = \sum_{k=t}^M \Pr[T = k \mid T \geq t] \cdot (k - t + 1). \quad (1)$$

This fully Bayesian strategy naturally quantifies uncertainty, makes decisions by expected-utility maximization, and seamlessly integrates any prior information. In the sections that follow, we formalize the algorithm, prove its competitive guarantees, and explore its empirical performance and extensions.

We begin by reviewing the classical ski rental problem, which forms the foundation for our Bayesian extension.

Related Work

Classical Results. The ski rental problem was formalized by Karlin et al. (1994), who showed a competitive ratio lower bound of 2 for deterministic algorithms, achievable via a simple threshold strategy. Chrobak et al. (1991) proved that randomization improves this to $e/(e - 1) \approx 1.582$. Extensions include multi-item (Fleischer 2001) and dynamic-cost settings (Young 2000).

Learning-Augmented Algorithms. Recent work integrates predictions into online algorithms, with frameworks balancing consistency and robustness (Lykouris and Vassilvitskii 2018; Mitzenmacher 2020). For ski rental, Kumar, Purohit, and Svitkina (2018) proposed prediction-dependent thresholds, while Anand et al. (2022) studied multiple predictions. More recent advances extend to offline problems such as matching (Dinitz et al. 2021), clustering (Ergun et al. 2021), and sorting (Bai and Coester 2023). Inspired by learned index structures (Kraska et al. 2018), robust techniques have emerged for combining multiple predictions (Anand et al. 2022; Antoniadis et al. 2023), alongside new applications in binary search (Dinitz et al. 2024), warm-started decision-making (Blum and Srinivas 2025), and learning-informed dynamic graphs (Brand et al. 2024).

Bayesian Approaches. Bayesian methods are well-established in offline problems (e.g., Thompson sampling (Agrawal and Goyal 2013)), but competitive analyses for online settings remain rare.

Our Work. To the best of our knowledge, our work is the first to introduce a fully Bayesian framework for the ski rental problem with distributional predictions, establishing provable guarantees under a discrete prior.

Preliminaries

We briefly review the classical ski rental problem and then present our discrete Bayesian extension.

Classical Ski Rental

At each day $t = 1, 2, \dots$, a skier chooses to:

- Rent: pay \$1 for the day,
- Buy: pay a one-time cost $b > 1$ for unlimited use.

The offline optimum for an unknown season length T is $\text{OPT}(T) = \min\{T, b\}$. An online algorithm \mathcal{A} incurs cost $\text{ALG}(\mathcal{A}, T)$, and its competitive ratio is

$$\text{CR}(\mathcal{A}) = \sup_{T \geq 1} \frac{\text{ALG}(\mathcal{A}, T)}{\text{OPT}(T)}.$$

No deterministic algorithm achieves $\text{CR} < 2$, but randomization attains the optimal $e/(e - 1) \approx 1.582$.

Discrete Bayesian Extension

We assume a finite horizon M and prior $\Pr[T = k] = \pi_k$ for $k = 1, \dots, M$, with $\sum_{k=1}^M \pi_k = 1$. Conditioned on survival ($T \geq t$), the posterior becomes

$$\Pr[T = k \mid T \geq t] = \frac{\pi_k}{\sum_{j=t}^M \pi_j}, \quad k = t, \dots, M.$$

Define the expected remaining rental cost:

$$E_{\text{rent}}(t) = \sum_{k=t}^M \Pr[T = k \mid T \geq t] \cdot (k - t + 1).$$

Our policy purchases on the first day t such that $b \leq E_{\text{rent}}(t)$, and continues renting otherwise, reassessing daily.

Computational Complexity

This section provides a detailed analysis of the computational complexity of the proposed Bayesian decision procedure, including its time, space, and sparse-support implementation aspects.

- 1) **Time:** Computing the normalization factor and the expected-cost sum each day takes $O(M)$ time, resulting in a worst-case total time complexity of $O(M^2)$.
- 2) **Space:** Storing the prior and posterior distributions requires $O(M)$ memory. All arithmetic operations involve nonnegative sums and divisions, ensuring numerical stability.
- 3) **Sparse-Support Implementation:** When the prior has n (with $n \leq M$) nonzero entries (k_i, π_{k_i}) , we can optimize the algorithm by indexing the support points using a balanced BST or hash map, maintaining prefix sums over survival probabilities for $O(\log n)$ query time, and computing $E_{\text{rent}}(t)$ by iterating only over the n support points. These optimizations reduce the overall runtime to $O(n \log n)$ without affecting the competitive guarantees.

With these preliminaries established, we now present the detailed Bayesian decision procedure that leverages this probabilistic framework to make optimal online decisions.

Algorithm 1: Discrete Bayesian Ski Rental

Require: Buy cost $b > 1$, horizon bound $M \geq 1$, prior distribution $\pi = [\pi_1, \pi_2, \dots, \pi_M]$

Ensure: Purchase day t^* (or $M+1$ if no purchase occurs)

- 1: $t \leftarrow 1$
- 2: **Validate/normalize prior:**
- 3: **if** $\sum_{k=1}^M \pi_k = 0$ **then**
- 4: **return** $M+1$ ▷ Invalid prior
- 5: **end if**
- 6: $\pi \leftarrow \pi / \sum_{k=1}^M \pi_k$ ▷ Normalize to a valid pmf
- 7: **while** $t \leq M$ **do**
- 8: Compute $Z_t, p_{t,k}$, and $E_{\text{rent}}(t)$ as in (2)
- 9: **if** $Z_t = 0$ **then**
- 10: **return** $M+1$ ▷ No remaining probability mass
- 11: **end if**
- 12: **if** $b \leq E_{\text{rent}}(t)$ **then**
- 13: **return** t ▷ Buy today
- 14: **else**
- 15: $t \leftarrow t + 1$ ▷ Rent today; re-evaluate tomorrow
- 16: **end if**
- 17: **end while**
- 18: **return** $M+1$ ▷ Never buy within horizon

Bayesian Decision Procedure

Algorithm Description

Our approach makes decisions via exact Bayesian inference over a discrete prior on the unknown horizon. As shown in Algorithm 1, given a buy cost $b > 1$, horizon bound M , and prior $\pi = [\pi_1, \dots, \pi_M]$, the algorithm updates its posterior over $[t, M]$ at each day t , conditioned on survival. It computes the expected rental cost under this posterior and compares it to b . If the expected cost exceeds b , it buys; otherwise, it rents and continues. This repeats until purchase or horizon end. This simple, efficient method adapts to arbitrary discrete priors without relying on heuristics or point estimates, and naturally extends to contextual or time-varying settings.

$$Z_t \triangleq \sum_{k=t}^M \pi_k, \quad (2a)$$

$$p_{t,k} \triangleq \frac{\pi_k}{Z_t} \quad (k = t, \dots, M), \quad (2b)$$

$$E_{\text{rent}}(t) \triangleq \sum_{k=t}^M p_{t,k} \times (k - t + 1). \quad (2c)$$

Algorithm Properties

The algorithm operates by making a daily decision: whether to buy skis on day t or continue renting. On each day t , it compares the known cost of buying, which is b , against the expected cost of continuing to rent from day t onward, denoted by $E_{\text{rent}}(t)$. This expected rental cost is computed with respect to the posterior distribution over ski days, conditional on the fact that skiing has lasted at least until day t . The decision rule is simple: buy as soon as $b \leq E_{\text{rent}}(t)$. This policy provably minimizes the expected total cost under π .

Lemma 1 (Monotonicity of Purchase Incentive). *Under the log-concave prior distribution, the expected remaining rental cost $E_{\text{rent}}(t)$ is non-increasing in t :*

$$E_{\text{rent}}(t+1) \leq E_{\text{rent}}(t).$$

We analyze the performance of this algorithm in terms of its *expected competitive ratio* (ECR), defined as the ratio between the algorithm's expected cost and the expected cost of the offline optimal strategy:

$$\text{ECR} = \frac{\sum_{k=1}^{t^*-1} \pi_k k + \sum_{k=t^*}^M \pi_k (t^* - 1 + b)}{\sum_{k=1}^M \pi_k \min(k, b)}. \quad (3)$$

The offline optimal algorithm, knowing k in advance, incurs cost $\min(k, b)$. Its expected cost is therefore

$$\mathbb{E}_\pi[\text{Cost}_{\text{OPT}}] = \sum_{k=1}^M \pi_k \cdot \min(k, b).$$

Let t^* denote the day on which the Bayes-optimal algorithm chooses to buy. Its expected cost is

$$\mathbb{E}_\pi[\text{Cost}_{\text{ALG}}] = \sum_{k=1}^{t^*-1} \pi_k \cdot k + \sum_{k=t^*}^M \pi_k \cdot (t^* - 1 + b).$$

Thus, the expected competitive ratio becomes:

$$\text{ECR} = \frac{\sum_{k=1}^{t^*-1} \pi_k \cdot k + \sum_{k=t^*}^M \pi_k \cdot (t^* - 1 + b)}{\sum_{k=1}^M \pi_k \cdot \min(k, b)}.$$

This ratio is always at least 1, and equals 1 only when π is concentrated on a single value of k such that the algorithm makes the correct offline choice.

Compared to classical algorithms, this Bayesian approach provides significantly improved performance in expectation. The deterministic strategy of buying on day b has a worst-case competitive ratio of $(2b-1)/b$, which approaches 2 as $b \rightarrow \infty$. The optimal randomized strategy achieves a worst-case ratio of $e/(e-1) \approx 1.58$. In contrast, the Bayes-optimal strategy leverages prior information to minimize expected cost and achieves strictly better performance when the prior is informative. However, this benefit is contingent on the quality of the prior: a poorly specified prior may result in suboptimal outcomes.

When Bayesian Outperforms Classical Algorithms

Having established the core algorithm and its properties, we now analyze specific scenarios where the Bayesian approach significantly outperforms classical algorithms, demonstrating its practical advantages under various prior distributions.

Case 1: Uniform Prior (Fixed Range)

We analyze the scenario where the trip length k is uniformly distributed over a fixed, finite range, $k \in \{1, 2, \dots, N\}$:

$$\pi_k = \frac{1}{N}, \quad \text{for } k \in \{1, 2, \dots, N\}.$$

Bayesian Algorithm's Behavior. The Bayesian algorithm makes its decision entirely on *Day 1*. Since the expected remaining rental cost $E_{\text{rent}}(t)$ is non-increasing in t (as shown by the Monotonicity Lemma), if the algorithm does not buy on Day 1, it will *never* buy.

The expected rental cost on Day 1 is given by the expectation of the uniform prior:

$$E_{\text{rent}}(1) = \sum_{k=1}^N \pi_k \cdot k = \frac{1}{N} \sum_{k=1}^N k = \frac{N+1}{2}.$$

Therefore, the algorithm chooses to buy on Day 1 ($t^* = 1$) if and only if

$$b \leq E_{\text{rent}}(1),$$

and otherwise never buys ($t^* = N + 1$) when

$$b > E_{\text{rent}}(1).$$

The following theorem holds.

Theorem 1 (Uniform Prior). *If $\pi_k = 1/N$ for $k \in \{1, \dots, N\}$, the Bayesian algorithm achieves an ECR of*

$$ECR = \begin{cases} 1, & \text{if } N \leq b, \\ \frac{N(N+1)}{b(2N-b+1)}, & \text{if } b < N < 2b-1, \\ \frac{2N}{2N-b+1}, & \text{if } N \geq 2b-1. \end{cases}$$

Case 2: Geometric Prior with Fixed Support

This setting assumes the trip length k is drawn from a geometric distribution truncated to the range $k \in \{1, 2, \dots, N\}$.

Prior Distribution and Expected Costs The prior distribution over trip lengths is defined as:

$$\pi_k = \frac{p(1-p)^{k-1}}{1-(1-p)^N}, \quad \text{for } k = 1, 2, \dots, N,$$

where p denotes the success probability of the geometric distribution.

$$ECR = \frac{\sum_{\tau=1}^N \pi_{\tau} [\min(t^*(\tau), \tau) - 1 + b \cdot \mathbf{1}_{\{t^*(\tau) \leq \tau\}}]}{\sum_{\tau=1}^N \pi_{\tau} \min(\tau, b)}.$$

where $t^*(\tau)$ is the stopping time when the true horizon is τ , determined by:

$$t^*(\tau) = \min \left\{ t \leq \tau : b \leq \sum_{k=t}^N \frac{\pi_k}{Z_t} (k-t+1) \right\}.$$

Let E_1 denote the expected trip length (the mean of the truncated geometric distribution):

$$E_1 \triangleq E_{\text{rent}}(1) = \sum_{k=1}^N k \pi_k.$$

We have

$$E_{\text{rent}}(k+1) = \frac{1}{p} - \frac{(N-k)(1-p)^{N-k}}{1-(1-p)^{N-k}}$$

is decreasing in k when N is larger than some constant.

Theorem 2 (Truncated Geometric Prior). *If the prior is given by $\pi_k \propto p(1-p)^{k-1}$ truncated to $k \leq N > 2$, the Bayesian algorithm achieves an Expected Competitive Ratio*

$$ECR = \frac{\min(E_1, b)}{\sum_{k=1}^N \pi_k \min(k, b)}.$$

Case 3: Truncated Gaussian Prior

We assume T is drawn from a truncated Gaussian prior on $\{1, \dots, N\}$:

Prior Distribution. Let μ and σ be the mean and standard deviation of the untruncated Gaussian. The truncated Gaussian prior is defined over $\{1, 2, \dots, N\}$ as:

$$\pi_k = \frac{e^{-\frac{(k-\mu)^2}{2\sigma^2}}}{\sum_{j=1}^N e^{-\frac{(j-\mu)^2}{2\sigma^2}}}, \quad \text{for } k \in \{1, 2, \dots, N\}.$$

This prior is highly informative and places most of its probability mass near the mean μ . A discretized truncated Gaussian prior $\pi_k \propto \exp\left(-\frac{(k-\mu)^2}{2\sigma^2}\right)$ is log-concave; a log-concave discrete PMF has an increasing discrete hazard rate $h(t) = \frac{\pi_t}{\sum_{k=t}^N \pi_k}$, and an increasing hazard rate (IHR) implies that the mean residual life is nonincreasing. Hence the Bayesian algorithm makes its decision entirely on *Day 1*.

We now extend the basic framework to handle more complex scenarios, including multiple predictions, adaptive learning, and contextual information.

Algorithmic Extensions and Variants

Multiple Predictions Algorithm

Given multiple independent predictions $\hat{T}_1, \dots, \hat{T}_n$ about the unknown horizon T , each with an associated uncertainty level σ_i , we seek to compute a refined posterior distribution over T that integrates all available information. We assume each prediction is a noisy observation of the true horizon, modeled as a Gaussian centered at T with variance σ_i^2 . Starting from a uniform prior, we sequentially update the posterior using Bayes' rule. Each prediction contributes a likelihood term, and the updates ensure that the final posterior distribution reflects all predictions in a statistically optimal way. The resulting posterior can then be used to make decisions in the ski rental problem with improved accuracy.

Our approach above not only leverages each predictor's uncertainty but also achieves a near-optimal decision quality in expectation.

Adaptive Prior Learning

In many real-world scenarios, the ski rental problem is encountered repeatedly with varying horizons drawn from an unknown distribution. We propose an adaptive algorithm that learns this prior online, starting from a uniform initialization and updating it each round using the observed buying time T_r via an exponential moving average with learning rate α . This refinement enables adaptation to the underlying horizon distribution and improved decision-making, aligning with sequential or learning-augmented online optimization frameworks (Gummadi and Pratap 2021; Lattanzi et al. 2020).

Algorithm 2: Bayesian Ski Rental with Multiple Predictions

Require: Predictions $\{\hat{T}_1, \dots, \hat{T}_n\}$ and their accuracies $\{\sigma_1, \dots, \sigma_n\}$
Ensure: Combined posterior distribution

- 1: **Initialize:** Uniform prior $\pi_k^{(0)} = 1/M$
- 2: **for** $i = 1$ to n **do**
- 3: **Compute Likelihood:** $L_i(k) = \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{(k-\hat{T}_i)^2}{2\sigma_i^2}\right)$
- 4: **Update Posterior:** $\pi_k^{(i)} \propto \pi_k^{(i-1)} \cdot L_i(k)$
- 5: **Normalize:** $\pi_k^{(i)} = \pi_k^{(i)} / \sum_{j=1}^M \pi_j^{(i)}$
- 6: **end for**
- 7: **return** $\pi^{(n)}$

Algorithm 3: Adaptive Bayesian Ski Rental

Require: Number of rounds R , learning rate α
Ensure: Sequence of purchasing decisions

- 1: **Initialize:** $\pi^{(1)} = \text{Uniform}(1, M)$
- 2: **for** round $r = 1$ to R **do**
- 3: Run Bayesian algorithm using prior $\pi^{(r)}$
- 4: Observe realized horizon T_r
- 5: **Update:** $\pi_k^{(r+1)} = (1 - \alpha_r)\pi_k^{(r)} + \alpha_r \cdot \mathbf{1}[k = T_r]$
- 6: **end for**

Theorem 3 (Adaptive Regret Bound). *With learning rate*

$$\alpha_r = \min \left\{ 1, \sqrt{\frac{\log M}{r}} \right\},$$

the adaptive algorithm incurs regret:

$$\text{Regret}_R = \sum_{r=1}^R [\text{Cost}_r - \text{OPT}_r] = O(\sqrt{R \log M}).$$

The theorem shows that the adaptive algorithm achieves cumulative regret $O(\sqrt{R \log M})$ over R rounds, where Cost_r is the incurred cost and OPT_r is the offline optimal cost for horizon T_r .

When contextual information $x \in \mathcal{X}$ is available, the prior can be conditioned on x through a softmax parameterization:

Definition 1 (Contextual Prior). Given context x , the prior is defined as

$$\pi_k(x) = \frac{\exp(\theta_k^\top \phi(x))}{\sum_{j=1}^M \exp(\theta_j^\top \phi(x))},$$

where $\phi(x)$ is a feature map and θ_k is the parameter associated with horizon k .

This parametrization lets the prior depend smoothly on x , effectively modeling $\pi(\cdot | x)$ as a categorical exponential-family distribution. The expected remaining rental cost then becomes

$$\mathbb{E}[T - t + 1 | x, T \geq t],$$

so contextual information directly shifts the decision boundary through the induced posterior.

Experiments

We structure our experimental evaluation around the following questions:

Q1. Robustness to prior misspecification. How stable is the Bayesian algorithm when the assumed prior differs from the true horizon distribution?

Q2. Performance under perfect prior knowledge. How close does the algorithm get to the offline optimal when the prior is accurate?

Q3. Noisy single predictions. Can the algorithm maintain performance when only a biased or noisy point prediction is available?

Q4. Multi-modal prior distributions. Does the algorithm adapt to complex, multi-peaked distributions without explicit mode detection?

Experimental Setup

We evaluate all methods with buy cost $b = 100$, horizon bound $M = 500$, and 10,000 Monte Carlo trials per configuration. We compare the Bayesian algorithm against deterministic thresholding, the optimal randomized strategy, point-prediction purchase, and the learning-augmented strategy of Kumar et al. (2024). Performance is measured via competitive ratio (CR) and success rate.

Q1. Robustness to Prior Misspecification

We test three Gaussian uncertainty regimes ($\sigma/\mu \in \{0.42, 0.33, 0.31\}$) and introduce misspecification by perturbing the mean, variance, and distributional shape. As shown in Figure 1, performance degrades only mildly: the average cost increase is 5.3%, and the worst-case degradation reaches 18.7% under extreme mean errors. Variance and model-form errors have negligible impact.

Q2. Performance Under Perfect Prior Knowledge

Table 1 summarizes performance when the assumed prior matches the true distribution. The Bayesian method achieves near-optimal CR ≈ 1.02 with a 98.7% success rate, significantly outperforming classical baselines.

Q3. Noisy Single Predictions

We next evaluate robustness when only a single noisy prediction \hat{T} is provided. Predictions follow $\hat{T} \sim \mathcal{N}(\alpha T, (\beta T)^2)$ with $\beta = 0.3$ and bias $\alpha \in [0.5, 2.0]$. Table 2 shows that the Bayesian method degrades smoothly from CR 1.05 to 1.43 as bias increases, consistently outperforming the point-prediction baseline.

Q4. Multi-modal Prior Distributions

Finally, we evaluate adaptability under complex multi-modal priors. As shown in Table 3, even when the density exhibits multiple peaks, the algorithm aligns its purchase threshold with the effective tail mass captured by the survival function, requiring no explicit mode identification.




Figure 1: Robustness under prior misspecification. (a) Performance degradation across uncertainty regimes remains small even at high total variation (TV) distances. (b) Mean errors have the largest impact, while variance and model errors are negligible. (c) CR distribution remains stable across TV bins.

Experiment 1: Robustness to Prior Misspecification

We analyze how performance changes when the assumed prior is distorted. It shows that the competitive ratio remains highly stable across a wide range of perturbations: even at large total-variation distances, the average cost increases by only 5.3% and the worst-case degradation under severe mean shifts reaches just 18.7%, while variance and shape errors have negligible effect.

This robustness arises in the regime where $E[T] > b$, so the buy decision depends mainly on the expected horizon rather than finer details of the prior. Once this boundary is crossed, moderate distortions leave the optimal action essentially unchanged. For comparison, Table 1 shows that with a perfectly specified prior the algorithm achieves near-optimal performance ($CR \approx 1.02$), confirming that robustness does not come at the cost of optimality.

Experiment 2: Perfect Prior Knowledge We evaluate performance when the algorithm’s assumed prior matches the true horizon distribution.

Setup. The true horizon is drawn from a known prior,

$$T \sim \pi, \\ \pi \in \{\text{Unif}[1, 500], \mathcal{N}(100, 30^2), \text{Exp}(0.01)\}.$$

With perfect information, the Bayesian threshold rule uses the exact posterior and therefore coincides with the optimal stopping rule for each distribution.

Across all prior families, the Bayesian algorithm achieves near-optimal performance, obtaining a mean competitive ratio of 1.02 and a 98.7% success rate. In contrast, deterministic and randomized strategies produce significantly higher costs ($CR = 1.85$ and 1.58), while point-prediction methods remain noticeably suboptimal ($CR = 1.16$). These results demonstrate that the Bayesian policy not only retains robustness under misspecification (Experiment 1) but also fully realizes its advantage when accurate prior information is available. Taken together, the evidence highlights that principled uncertainty modeling yields consistent gains across both noisy and well-specified environments.

| Algorithm | Mean CR | 95% CI | Success Rate |
|------------------|---------|----------------|--------------|
| Bayesian | 1.023 | [1.021, 1.025] | 98.7% |
| Randomized | 1.582 | [1.577, 1.587] | 67.3% |
| Deterministic | 1.847 | [1.839, 1.855] | 52.1% |
| Prediction-based | 1.156 | [1.150, 1.162] | 81.4% |

Table 1: Performance under perfect prior knowledge (10,000 trials).




Figure 2: Competitive ratio under perfect prior knowledge. The Bayesian method achieves $CR \approx 1.02$ across all priors, significantly outperforming classical algorithms.

Experiment 3: Noisy Single Predictions We evaluate robustness when only a biased and noisy single prediction is provided. Predictions follow

$$\hat{T} \sim \mathcal{N}(\alpha T, (\beta T)^2),$$

with true horizon fixed at $T = 100$, noise level $\beta = 0.3$, and bias $\alpha \in [0.5, 2.0]$.

The competitive ratio rises smoothly from 1.05 to 1.43 as the bias parameter increases from 0.5 to 2.0, indicating that performance degrades gradually rather than abruptly as pre-

| Bias α | $ \alpha - 1 $ | Bayesian CR | Point Pred. CR |
|---------------|----------------|-------------|----------------|
| 0.5 | 0.5 | 1.05 | 1.29 |
| 0.8 | 0.2 | 1.08 | 1.34 |
| 1.0 | 0.0 | 1.12 | 1.38 |
| 1.2 | 0.2 | 1.21 | 1.47 |
| 1.5 | 0.5 | 1.31 | 1.60 |
| 2.0 | 1.0 | 1.43 | 1.79 |

Table 2: Competitive ratio under noisy single predictions ($T = 100$).

dictions become less accurate. Importantly, no sharp phase transition is observed: even under substantial bias, the algorithm avoids catastrophic over- or under-purchasing by explicitly accounting for predictive uncertainty.

Across all bias levels, the Bayesian method consistently outperforms the point-prediction baseline by 15–30%. This improvement stems from the fact that point-based strategies implicitly treat the prediction as ground truth, whereas the Bayesian approach interprets it as a noisy observation and tempers its influence through the posterior distribution. As a result, extreme prediction errors are down-weighted rather than directly translated into early or delayed purchase decisions.

Overall, this experiment demonstrates that incorporating uncertainty at the distributional level yields robust behavior under noisy single predictions: the algorithm degrades gracefully as bias increases, maintaining reliable performance even when the available forecast is substantially misaligned with the true horizon.

Experiment 4: Multi-modal Prior Distributions We evaluate the algorithm’s behavior under multi-modal horizon distributions, which commonly arise in real-world settings such as weekday–weekend cycles or seasonal demand patterns. For each prior, we describe the distributional structure and the resulting optimal purchase threshold.

- **Bi-modal typical case** ($t^* = 30$). The prior

$$\pi = 0.7\mathcal{N}(10, 3^2) + 0.3\mathcal{N}(25, 5^2)$$

exhibits two dominant modes at 10 and 25. Despite these early peaks, the posterior survival probability $S(t) = \Pr(T \geq t)$ decays slowly in the right tail, keeping $\mathbb{E}[T - t + 1 \mid T \geq t] > b$ until relatively late in the horizon. Consequently, the algorithm delays purchase beyond both modes and selects a coherent threshold at $t^* = 30$.

- **Tri-modal balanced case** ($t^* = 21$). The prior

$$\pi = \frac{1}{3}(\delta_8 + \delta_{20} + \delta_{40})$$

contains three equally weighted peaks. Although the density is highly non-monotonic, the decision depends only on the monotone survival function $S(t)$ rather than on the number or precise locations of modes. As a result, the algorithm produces a single stable threshold at $t^* = 21$ that balances early and late probability mass.

| Case | Dominant modes | Optimal t^* |
|--------------------|----------------|---------------|
| Bi-modal typical | 10, 25 | 30 |
| Tri-modal balanced | 8, 20, 40 | 21 |
| Seasonal peak | 5, 30 | 1 |

Table 3: Summary of multi-modal priors and optimal Bayesian purchase times.

- **Seasonal-peak case** ($t^* = 1$). The prior

$$\pi = 0.3\mathcal{N}(5, 3^2) + 0.7\mathcal{N}(30, 10^2)$$

features occasional short horizons together with a dominant late-season peak. In this setting, the heavy early mass causes a sharp drop in $S(t)$ at small t , so $\mathbb{E}[T - t + 1 \mid T \geq t]$ immediately falls below b . This triggers an immediate purchase decision at $t^* = 1$.

Taken together, these cases highlight a structural property of the Bayesian rule: the decision boundary is governed by the *integrated* survival function

$$S(t) = \Pr(T \geq t),$$

rather than by local density values or mode locations. As a consequence, strong multi-modality does not induce instability, and the algorithm adapts smoothly to the global shape of the horizon distribution.

The stopping rule depends only on the tail mass $S(t) = \Pr(T \geq t)$ through the condition $b \leq \mathbb{E}[T - t + 1 \mid T \geq t]$. Thus local bumps or multiple modes in π do not affect the decision; only the cumulative survival shape determines when the threshold is crossed.

Discussion and Future Work

Practical Implications: Our discrete Bayesian framework offers key practical benefits: principled uncertainty quantification via full posterior maintenance, graceful handling of noisy forecasts beyond point estimates, and seamless incorporation of domain knowledge through customizable priors. It naturally supports multi-modal distributions—e.g., weekday/weekend patterns—without explicit regime switching, as confirmed by our experiments. Moreover, it adapts online by updating posteriors as new data arrive, improving responsiveness in dynamic settings.

Limitations and Assumptions: We assume a discrete, finite horizon and stationary costs, which may limit real-world applicability. Extensions to continuous time and dynamic pricing would require new techniques. Episodes are treated as independent, ignoring temporal correlations. Perfect observability of the horizon end is assumed; more realistic settings with partial or delayed feedback would need belief tracking or filtering.

Future Research Directions: Future work includes meta-learning priors from historical data, robust Bayesian optimization for worst-case guarantees, and multi-agent extensions for shared resource settings. Continuous-time models could link to optimal stopping theory. Incorporating high-dimensional contextual information via deep generative priors offers potential for richer, feature-aware decision-making.

Acknowledgements

This work was supported by the New Faculty Startup Fund from Seoul National University.

References

- Agrawal, S.; and Goyal, N. 2013. Thompson Sampling for Contextual Bandits with Linear Payoffs. In *International Conference on Machine Learning*, 127–135. PMLR.
- Anand, K.; Ge, R.; Kumar, A.; and Panigrahi, D. 2022. Online algorithms with multiple predictions. In *International Conference on Machine Learning*, 582–598. PMLR.
- Antoniadis, A.; Coester, C.; Eliás, M.; Polak, A.; and Simon, B. 2023. Mixing predictions for online metric algorithms. In *International Conference on Machine Learning*, 969–983. PMLR.
- Bai, X.; and Coester, C. 2023. Sorting with predictions. *Advances in Neural Information Processing Systems*, 36: 26563–26584.
- Blum, A.; and Srinivas, V. 2025. Competitive strategies to use “warm start” algorithms with predictions. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 3775–3801. SIAM.
- Brand, J. v. d.; Forster, S.; Nazari, Y.; and Polak, A. 2024. On dynamic graph algorithms with predictions. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 3534–3557. SIAM.
- Chrobak, M.; Karloff, H.; Payne, T.; and Vishwanathan, S. 1991. New results on server problems. *SIAM Journal on Discrete Mathematics*, 4(2): 172–181.
- Dinitz, M.; Im, S.; Lavastida, T.; Moseley, B.; Niaparast, A.; and Vassilvitskii, S. 2024. Binary Search with Distributional Predictions. *Advances in Neural Information Processing Systems*, 37: 90456–90472.
- Dinitz, M.; Im, S.; Lavastida, T.; Moseley, B.; and Vassilvitskii, S. 2021. Faster matchings via learned duals. *Advances in Neural information processing systems*, 34: 10393–10406.
- Ergun, J. C.; Feng, Z.; Silwal, S.; Woodruff, D. P.; and Zhou, S. 2021. Learning-augmented k -means clustering. *arXiv preprint arXiv:2110.14094*.
- Fleischer, R. 2001. On the Bahncard problem. *Theoretical Computer Science*, 268(1): 161–174.
- Gummadi, K. T.; and Pratap, R. 2021. Sequential Ski Rental Problem. *arXiv preprint arXiv:2104.06435*.
- Karlin, A. R.; Manasse, M. S.; McGeoch, L. A.; and Owicki, S. 1994. Competitive randomized algorithms for nonuniform problems. *Algorithmica*, 11(6): 542–571.
- Kraska, T.; Beutel, A.; Chi, E. H.; Dean, J.; and Polyzotis, N. 2018. The case for learned index structures. In *Proceedings of the 2018 international conference on management of data*, 489–504.
- Kumar, R.; Purohit, M.; and Svitkina, Z. 2018. Improving online algorithms via ML predictions. In *Advances in Neural Information Processing Systems (NeurIPS)*, 9684–9693.
- Lattanzi, S.; Lavastida, T.; Moseley, B.; and Vassilvitskii, S. 2020. Online scheduling via learned weights. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1859–1877.
- Lykouris, T.; and Vassilvitskii, S. 2018. Competitive Caching with Machine Learned Advice. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, 3302–3311. PMLR.
- Mitzenmacher, M. 2020. Scheduling with predictions and the price of misprediction. In *Proceedings of the 11th Innovations in Theoretical Computer Science Conference (ITCS)*, 14:1–14:18.
- Purohit, M.; Svitkina, Z.; and Kumar, R. 2018. Improving online algorithms via ML predictions. *Advances in Neural Information Processing Systems*, 31.
- Young, N. E. 2000. Online paging against adversarially biased random inputs. *Journal of Algorithms*, 37(1): 218–235.