



**HAL**  
open science

# Bayesian Off-Policy Evaluation and Learning for Large Action Spaces

Imad Aouali, Victor-Emmanuel Brunel, David Rohde, Anna Korba

► **To cite this version:**

Imad Aouali, Victor-Emmanuel Brunel, David Rohde, Anna Korba. Bayesian Off-Policy Evaluation and Learning for Large Action Spaces. 2024. hal-04606070

**HAL Id: hal-04606070**

**<https://hal.science/hal-04606070v1>**

Preprint submitted on 9 Jun 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

---

# Bayesian Off-Policy Evaluation and Learning for Large Action Spaces

---

**Imad Aouali**

CREST, ENSAE, IP Paris, France  
Criteo AI Lab, Paris, France

**Victor-Emmanuel Brunel**

CREST, ENSAE, IP Paris, France

**David Rohde**

Criteo AI Lab, Paris, France

**Anna Korba**

CREST, ENSAE, IP Paris, France

## Abstract

In interactive systems, actions are often correlated, presenting an opportunity for more sample-efficient off-policy evaluation (OPE) and learning (OPL) in large action spaces. We introduce a unified Bayesian framework to capture these correlations through structured and informative priors. In this framework, we propose  $\text{sDM}$ , a generic Bayesian approach designed for OPE and OPL, grounded in both algorithmic and theoretical foundations. Notably,  $\text{sDM}$  leverages action correlations without compromising computational efficiency. Moreover, inspired by online Bayesian bandits, we introduce Bayesian metrics that assess the average performance of algorithms across multiple problem instances, deviating from the conventional worst-case assessments. We analyze  $\text{sDM}$  in OPE and OPL, highlighting the benefits of leveraging action correlations. Empirical evidence showcases the strong performance of  $\text{sDM}$ .

## 1 Introduction

An off-policy contextual bandit [Dudík et al., 2011] is a practical framework for improving decision-making using logged data that records an agent’s interactions with an online environment [Bottou et al., 2013]. This data consists of context-action-reward tuples generated during online interactions where the agent observes a context, takes an action, and receives a reward depending on that context and action. Such data can be leveraged to enhance the agent’s performance through two main tasks: off-policy evaluation (OPE) [Dudík et al., 2011], which predicts the expected reward of a new policy using logged data, and off-policy learning (OPL) [Swaminathan and Joachims, 2015a], which learns an improved policy using the OPE estimator. Many interactive systems can be modeled as contextual bandits. For example, in online advertising, the context is the user’s features, the action is a product choice, and the reward is the click-through rate (CTR).

When the number of actions is small, existing methods, often based on the inverse propensity scoring (IPS) estimator [Horvitz and Thompson, 1952], are reliable and come with theoretical guarantees [Dudík et al., 2011, Dudík et al., 2012, Dudík et al., 2014, Wang et al., 2017, Farajtabar et al., 2018, Su et al., 2019, 2020, Metelli et al., 2021, Swaminathan and Joachims, 2015a, London and Sandler, 2019, Aouali et al., 2023a]. However, IPS can suffer from high bias and variance. The bias arises because the logging policy might have deficient support [Sachdeva et al., 2020], and the variance is due to potentially high values of the importance weights [Swaminathan et al., 2017]. Both issues are more pronounced in large action spaces. To address this, Saito and Joachims [2022] recently proposed MIPS, an IPS variant for OPE in large action spaces. However, MIPS requires access to informative embeddings that capture entirely the causal effects of actions on rewards. Hence, recent variants of MIPS propose to learn such embeddings directly from logged data Peng et al. [2023],

Sachdeva et al. [2023], Cief et al. [2024], or to relax the embedding causal effect assumption Taufiq et al. [2024], Saito et al. [2023].

An alternative to IPS is the direct method (DM) [Jeunen and Goethals, 2021], which does not rely on importance weighting. Instead, DMs learn a reward model and use it to evaluate policies. DMs tend to introduce less variance than IPS, but they can suffer from higher bias if the model is not well-suited for the data. The theoretical properties of DMs are underexplored, with a few exceptions [Jin et al., 2021, Hong et al., 2023]. Yet, they are widely used in practice, especially in recommender systems [Sakhi et al., 2020, Jeunen and Goethals, 2021, Aouali et al., 2022b,a]. This work aims to develop the algorithmic and theoretical foundations for DMs, improving our understanding of their performance in OPE and OPL from a Bayesian perspective. The goal is to enhance the efficiency of DMs through informative priors, particularly in scenarios with large action spaces. For intuition, consider how MIPS [Saito and Joachims, 2022] extended IPS to large action spaces by incorporating auxiliary information in the form of embeddings. We aim to achieve a similar extension for DMs by integrating auxiliary information into existing DMs and demonstrating how their performance can be improved. However, unlike MIPS’s extension of IPS, we provide theoretical and empirical evidence for both OPE and OPL, rather than focusing only on OPE. Also, our assumptions differ from those made by MIPS and its variants.

**Contributions.** We introduce a Bayesian direct method, called structured Direct Method (sDM), which uses informative priors to share data (reward) information among actions. Essentially, when one action is observed, it enhances sDM’s knowledge about similar actions, significantly boosting statistical efficiency while maintaining computational efficiency. This makes sDM scalable to large action spaces, addressing a key limitation of standard DMs. Moreover, we provide a comprehensive step-by-step framework to assist practitioners in using sDM for OPE and OPL. To analyze sDM, we introduce Bayesian metrics for OPE and OPL that assess average performance across various problem instances, diverging from frequentist worst-case evaluations. Unlike frequentist metrics, these Bayesian metrics capture the benefits of incorporating informative priors, enhancing our understanding of DMs, especially given the relatively underdeveloped statistical analysis of DMs in OPE and OPL. Finally, we evaluate sDM using both synthetic and MovieLens data.

**Related work.** Bayesian models for OPL have been explored by Lazaric and Ghavamzadeh [2010], Hong et al. [2023]. However, these studies incorporate Bayesian techniques in multi-task learning, distinguishing them from our more focused single-task approach. Additionally, the models used in these prior works differ from the ones we employ. While certain similarities can be drawn with Hong et al. [2023] in terms of analysis, our approach offers a broader scope. Note that Hong et al. [2023] also introduced a Bayesian metric for assessing suboptimality, yet their metric evaluates performance in an environment drawn from the prior given a fixed sample set, whereas ours measures average performance across all environments sampled from the prior and all sample sets. The latter is more common in Bayesian analyses (e.g., Bayesian regret in Russo and Van Roy [2014]). Also, in this work, we cover OPE which goes beyond the scope of Lazaric and Ghavamzadeh [2010], Hong et al. [2023] that focused on OPL. For a more extensive discussion of related works, we refer the reader to Appendix B.

The paper is organized as follows. Section 2 provides some background on OPE and OPL. Section 3 describes our method, sDM. Section 4 focuses on the linear case of sDM. In Section 5 we analyze sDM through finite sample bounds. In Section 6, we show that sDM enjoys favorable performance.

## 2 Background

Let  $[n] = \{1, \dots, n\}$  for any positive integer  $n$ . Random variables are denoted with capital letters, and their realizations with the respective lowercase letters, except for Greek letters. Let  $X, Y$  be two random variables, with a slight abuse of notation, the distribution (or density) of  $X \mid Y = y$  evaluated at  $x$  is denoted by  $p(x \mid y)$ . Let  $a_1, \dots, a_n$  be  $n$  vectors of  $\mathbb{R}^d$ ,  $a = (a_i)_{i \in [n]} \in \mathbb{R}^{nd}$  is their  $nd$ -dimensional concatenation,  $\otimes$  denotes the Kronecker product, and  $\mathcal{O}$  the big-O notation. Finally, for any vector  $a \in \mathbb{R}^d$  and any positive-definite matrix  $\Sigma \in \mathbb{R}^{d \times d}$ , we define  $\|a\|_\Sigma = \sqrt{a^\top \Sigma a}$ .

Agents are represented by stochastic policies  $\pi \in \Pi$ , where  $\Pi$  is the set of policies. Precisely, given a context  $x \in \mathbb{R}^d$ ,  $\pi(\cdot \mid x)$  is a probability distribution over a finite action set  $\mathcal{A} = [K]$ , where  $K$  is the number of actions. Agents interact with a contextual bandit environment over  $n$  rounds as follows. In round  $i \in [n]$ , the agent observes a context  $X_i \sim \nu$ , where  $\nu$  is a distribution whose support  $\mathcal{X}$  is a

compact subset of  $\mathbb{R}^d$ . Then the agent takes an *action*  $A_i \sim \pi_0(\cdot | X_i)$  from the action set  $\mathcal{A}$ , where  $\pi_0$  is the *logging policy*. Finally, assuming a parametric model, the agent receives a *stochastic reward*  $R_i \sim p(\cdot | X_i; \theta_{*,A_i})$ . Here  $p(\cdot | x; \theta_{*,a})$  is the *reward distribution* for action  $a$  in context  $x$ , where  $\theta_{*,a} \in \mathbb{R}^d$  is an *unknown  $d$ -dimensional parameter* of action  $a$ . Let  $\theta_* = (\theta_{*,a})_{a \in \mathcal{A}} \in \mathbb{R}^{dK}$  be the concatenation of action parameters. Then,  $r(x, a; \theta_*) = \mathbb{E}_{R \sim p(\cdot | x; \theta_{*,a})} [R]$  is the *reward function* that outputs the expected reward of action  $a$  in context  $x$ . Finally, the goal is to find a policy  $\pi \in \Pi$  that maximizes the value function  $V(\pi; \theta_*) = \mathbb{E}_{X \sim \nu} \mathbb{E}_{A \sim \pi(\cdot | X)} [r(X, A; \theta_*)]$ .

Let  $S = (X_i, A_i, R_i)_{i \in [n]}$  be a set of random variables drawn i.i.d. as  $X_i \sim \nu$ ,  $A_i \sim \pi_0(\cdot | X_i)$  and  $R_i \sim p(\cdot | X_i; \theta_{*,A_i})$ ; that we will refer to as the sample set. The goal of off-policy evaluation (OPE) and learning (OPL) is to build an estimator  $\hat{V}(\pi, S)$  of  $V(\pi; \theta_*)$  and then use  $\hat{V}(\pi, S)$  to find a policy  $\hat{\pi} \in \Pi$  that maximizes  $V(\cdot; \theta_*)$ . In OPE, it is common to use *inverse propensity scoring (IPS)* [Horvitz and Thompson, 1952, Dudík et al., 2012], which leverages importance sampling to estimate the value  $V(\pi; \theta_*)$  as  $\hat{V}_{\text{IPS}}(\pi, S) = \frac{1}{n} \sum_{i \in [n]} \frac{\pi(A_i | X_i)}{\pi_0(A_i | X_i)} R_i$  where for any  $(x, a) \in \mathcal{X} \times \mathcal{A}$ ,  $\frac{\pi(a|x)}{\pi_0(a|x)}$  are the *importance weights*. In practice, IPS can suffer high variance, especially when the action space is large [Swaminathan et al., 2017, Saito and Joachims, 2022]. Moreover, while IPS is unbiased under the assumption that the logging policy has full support, it can induce a high bias when such an assumption is violated [Sachdeva et al., 2020], which is again likely when the action space is large. IPS, also, assumes access to the logging policy  $\pi_0$ . An alternative approach to IPS is to use a *direct method (DM)* [Jeunen and Goethals, 2021], that relies on a reward model  $\hat{r}$  to estimate the value  $V(\pi; \theta_*)$  as

$$\hat{V}_{\text{DM}}(\pi, S) = \frac{1}{n} \sum_{i \in [n]} \sum_{a \in \mathcal{A}} \pi(a | X_i) \hat{r}(X_i, a), \quad (1)$$

where  $\hat{r}(x, a)$  is an estimation of  $r(x, a; \theta_*)$ . DM estimators may exhibit modeling bias, but they generally have lower variance than IPS [Saito and Joachims, 2022]. Another advantage of DM is its practical utility without assuming access to the logging policy  $\pi_0$  [Jeunen and Goethals, 2021, Aouali et al., 2022b, Hong et al., 2023]. Also, DMs can be incorporated into a Bayesian framework, where informative priors can be used to enhance statistical efficiency. This allows for the development of scalable methods suitable for large action spaces, as demonstrated in our work.

### 3 Structured direct method

#### 3.1 Structured priors

**Pitfalls of non-structured priors.** Before presenting sDM, we first describe the pitfalls of using the following widely used standard prior,

$$\begin{aligned} \theta_a &\sim \mathcal{N}(\mu_a, \Sigma_a), & \forall a \in \mathcal{A}, \\ R | \theta, X, A &\sim \mathcal{N}(\phi(X)^\top \theta_A, \sigma^2), \end{aligned} \quad (2)$$

where  $\phi(X)$  outputs a  $d$ -dimensional representation of the context  $X \in \mathcal{X}$ ,  $\mathcal{N}(\mu_a, \Sigma_a)$  is the prior density of  $\theta_a$ , and  $\sigma^2$  is the reward noise. In this prior, each action  $a$  corresponds to a parameter  $\theta_a$ . Then, under the prior in (2), the posterior on an action parameter is also a multivariate Gaussian and writes  $\theta_a | S \sim \mathcal{N}(\hat{\mu}_a, \hat{\Sigma}_a)$ , where  $\hat{\Sigma}_a^{-1} = \Sigma_a^{-1} + G_a$  and  $\hat{\Sigma}_a^{-1} \hat{\mu}_a = \Sigma_a^{-1} \mu_a + B_a$ , with  $G_a = \sigma^{-2} \sum_{i \in [n]} \mathbb{I}_{\{A_i=a\}} \phi(X_i) \phi(X_i)^\top$  and  $B_a = \sigma^{-2} \sum_{i \in [n]} \mathbb{I}_{\{A_i=a\}} R_i \phi(X_i)$ . Note that  $G_a$  and  $B_a$  only use the portion of the sample set  $S$  where action  $a$  was observed. Thus, observations relative to other actions  $b \neq a$  are not involved in the posterior inference relative to action  $a$ . This leads to statistical inefficiency since it is not even guaranteed that the sample set  $S$  covers all the actions. In particular, the posterior  $\theta_a | S$  of an unseen action  $a$  would be exactly the prior  $\mathcal{N}(\mu_a, \Sigma_a)$  since  $G_a = 0_{d \times d}$  and  $B_a = 0_d$  in that case. To mitigate this, we make the assumption that actions can be correlated, which is realistic in practice such as in modern recommender systems. Now one can think of capturing correlations by simply modeling the joint posterior density of  $(\theta_a)_{a \in \mathcal{A}} | S$ . Unfortunately, this is computationally expensive when  $K$  is large.

**Structured priors.** To model action correlations, we assume that there exists an *unknown  $d'$ -dimensional latent parameter*,  $\psi \in \mathbb{R}^{d'}$ , which is sampled from a *latent prior*  $q(\cdot)$ , such as  $\psi \sim q(\cdot)$ . Then the correlations between actions arise since they are derived from the same latent parameter  $\psi$ .

Precisely, action parameters  $\theta_a$  are sampled conditionally independently from a *conditional prior*  $p_a$  as  $\theta_a | \psi \sim p_a(\cdot; f_a(\psi))$  for any  $a \in \mathcal{A}$ , where  $p_a$  is parametrized by  $f_a(\psi)$  and  $f_a : \mathbb{R}^{d'} \rightarrow \mathbb{R}^d$  is a *known* function that encodes the *hierarchical* relation between action parameters  $\theta_a$  and the latent parameter  $\psi$ . In particular, the function  $f_a$  can capture sparsity. That is, when  $\theta_a$  only depends on a few coordinates of  $\psi$ . Moreover,  $p_a$  incorporates uncertainty due to model misspecification when  $\theta_a$  is not a deterministic function of  $\psi$ ; that is, when  $\theta_a \neq f_a(\psi)$ . Finally, the reward distribution for action  $a$  in context  $x$  is  $p(\cdot | x; \theta_a)$  and it only depends on  $x$  and  $\theta_a$ . To summarize, the prior is

$$\begin{aligned} \psi &\sim q(\cdot), \\ \theta_a | \psi &\sim p_a(\cdot; f_a(\psi)), \\ R | \psi, \theta, X, A &\sim p(\cdot | X; \theta_A). \end{aligned} \quad \forall a \in \mathcal{A}, \quad (3)$$

This prior is general as assuming the existence of such latent parameter  $\psi$  is mild (Appendix E.6). To be able to derive the posterior of our prior in (3), we assume that (i)  $(X, A)$  is independent of  $\psi$  and given  $\psi$ ,  $(X, A)$  is independent of  $\theta$ . (ii) Given  $\psi$ , the  $\theta_a$ , for all  $a \in \mathcal{A}$  are independent.

Now, we discuss how to perform OPE and OPL under the structured prior in (3).

### 3.2 Off-policy evaluation and learning

**Off-policy evaluation.** OPE aims at estimating the value function  $V(\pi; \theta_*)$  using the sample set  $S$ . In DMs, the estimator  $\hat{V}_{\text{DM}}$  in (1) requires access to the learned reward  $\hat{r}(x, a) \approx r(x, a; \theta_*)$ . In our Bayesian setting, this requires access to the action posterior  $\theta_a | S$  under model (3). The expected reward is then estimated as  $\hat{r}(x, a) = \mathbb{E}[r(x, a; \theta) | S]$  for any  $(x, a) \in \mathcal{X} \times \mathcal{A}$ , and this estimate is plugged into  $\hat{V}_{\text{DM}}$  in (1) to estimate  $V(\pi; \theta_*)$ . Therefore, we only need to derive the unnormalized posterior density of the action parameter  $\theta_a$ ,  $p(\theta_a | S)$ , under the structured prior in (3), which reads

$$p(\theta_a | S) = \int_{\psi} p(\theta_a | \psi, S) p(\psi | S) d\psi, \quad (4)$$

where  $\psi | S$  is the latent posterior and  $\theta_a | \psi, S$  is the *conditional* action posterior. To compute  $p(\theta_a | S)$ , we first compute  $p(\theta_a | \psi, S)$  and  $p(\psi | S)$  and then integrate out  $\psi$  following (4). First,

$$p(\theta_a | \psi, S) \propto \mathcal{L}_a(\theta_a) p_a(\theta_a; f_a(\psi)), \quad (5)$$

with  $\mathcal{L}_a(\theta_a) = \prod_{(X, A, R) \in S_a} p(R | X; \theta_a)$  is the likelihood of observations of action  $a$  ( $S_a = (X_i, A_i, R_i)_{i \in [n], A_i = a}$  is the subset of  $S$  where  $A_i = a$ ). Similarly,

$$p(\psi | S) \propto \prod_{b \in \mathcal{A}} \int_{\theta_b} \mathcal{L}_b(\theta_b) p_b(\theta_b; f_b(\psi)) d\theta_b q(\psi), \quad (6)$$

This allows us to further develop (4) as

$$p(\theta_a | S) \propto \int_{\psi} \mathcal{L}_a(\theta_a) p_a(\theta_a; f_a(\psi)) \prod_{b \in \mathcal{A}} \int_{\theta_b} \mathcal{L}_b(\theta_b) p_b(\theta_b; f_b(\psi)) d\theta_b q(\psi) d\psi. \quad (7)$$

All the quantities inside the integrals in (7) are given (the parameters of  $p_a$  and  $q$ ) or tractable (the terms in  $\mathcal{L}_a$ ). Thus, if these integrals can be computed, then the posterior can be fully characterized in closed-form which we will do in Section 4 in the fully linear case. Otherwise, the posterior should be approximated, e.g. through approximate sampling methods, see Appendix E.1 for more details.

**Greedy off-policy learning.** OPL aims at finding a policy  $\pi$  that maximizes the value  $V(\pi; \theta_*)$ . We act greedy with respect to our estimator  $\hat{V}_{\text{DM}}(\cdot, S)$  and define the learned policy as the one maximizing it:  $\hat{\pi}_{\text{G}} = \arg\max_{\pi \in \Pi} \hat{V}_{\text{DM}}(\pi, S)$ . If the set of policies  $\Pi$  contains deterministic policies, then

$$\hat{\pi}_{\text{G}}(a | x) = \mathbb{1}\{a = \arg\max_{b \in \mathcal{A}} \hat{r}(x, b)\}. \quad (8)$$

Thus we do not adopt the common pessimism approach [Jin et al., 2021]. In pessimism, one constructs confidence intervals of the reward estimate  $\hat{r}(x, a)$  of the form  $|r(x, a; \theta) - \hat{r}(x, a)| \leq u(x, a)$ , and then defines the learned policy as  $\hat{\pi}_{\text{P}}(a | x) = \mathbb{1}\{a = \arg\max_{b \in \mathcal{A}} \hat{r}(x, b) - u(x, b)\}$ . When  $u(x, a)$  does not depend on  $x$  and  $a$ , pessimistic and greedy policies are the same. The advantage of one over another depends on the evaluation metric. In this work, we use a novel metric, Bayesian suboptimality (BSO), defined in Section 5. It assesses the average performance of algorithms across multiple problem instances rather than focusing on the worst-case. We justify that the Greedy policy is more suitable for BSO optimization than pessimism theoretically in Section 5 and empirically in Appendix F.5.

## 4 Application: linear hierarchy

In this section, we motivate the use of linear functions  $f_a$  for the structured prior (3). Let  $W_a \in \mathbb{R}^{d \times d'}$  be the *mixing matrix* for action  $a$ , we define  $f_a(v) = W_a v$  for any  $v \in \mathbb{R}^{d'}$ . Then,

$$\theta_a \mid \psi \sim p_a(\cdot; W_a \psi), \quad \forall a \in \mathcal{A}. \quad (9)$$

This is an important model, in both practice and theory since linear models often lead to closed-form solutions that are both computationally tractable and allow theoretical analysis.

### 4.1 Applications

To highlight the generality of (9), we provide some problems where it can be used.

**Mixed-effect modeling.** (9) allows modeling that action parameters depend on a linear mixture of effect parameters. Precisely, let  $J$  be the number of effects and assume that  $d' = dJ$  so that the latent parameter  $\psi$  is the concatenation of  $J$ ,  $d$ -dimensional effect parameters,  $\psi_j \in \mathbb{R}^d$ , such as  $\psi = (\psi_j)_{j \in [J]} \in \mathbb{R}^{dJ}$ . Moreover, assume that for any  $a \in \mathcal{A}$ ,  $W_a = w_a^\top \otimes I_d \in \mathbb{R}^{d \times dJ}$  where  $w_a = (w_{a,j})_{j \in [J]} \in \mathbb{R}^J$  are the *mixing weights* of action  $a$ . Then,  $W_a \psi = \sum_{j \in [J]} w_{a,j} \psi_j$  for any  $a \in \mathcal{A}$ . Sparsity, i.e., when an action  $a$  only depends on a subset of effects, is captured through the mixing weights  $w_a$ :  $w_{a,j} = 0$  when action  $a$  is independent of the  $j$ -th effect parameter  $\psi_j$  and  $w_{a,j} \neq 0$  otherwise. Also, the level of dependence between action  $a$  and effect  $j$  is quantified by the absolute value of  $w_{a,j}$ . This Mixed-effect model [Aouali et al., 2023b] can be used in numerous applications. In movie recommendation,  $\theta_a$  is the parameter of movie  $a$ ,  $\psi_j$  is the parameter of theme  $j$  (adventure, romance, etc.), and  $w_{a,j}$  quantifies the relevance of theme  $j$  to movie  $a$ . Similarly, in clinical trials, a drug is a combination of multiple ingredients, each with a specific dosage. Then  $\theta_a$  is the parameter of drug  $a$ ,  $\psi_j$  is the parameter of ingredient  $j$ , and  $w_{a,j}$  is the dosage of ingredient  $j$  in drug  $a$ .

**Low-rank modeling.** (9) can also model the case where the dimension of the latent parameter  $\psi$  is much smaller than that of the action parameters  $\theta_a$ , i.e., when  $d' \ll d$ . Again, this is captured through the mixing matrices  $W_a$ . Precisely, when  $W_a$  is low-rank. This improves both statistical and computational efficiency as we show in Section 5.

### 4.2 Closed-form solutions for sDM

We further assume that the latent prior  $q(\cdot) = \mathcal{N}(\cdot; \mu, \Sigma)$  is Gaussian with mean  $\mu \in \mathbb{R}^{d'}$  and covariance  $\Sigma \in \mathbb{R}^{d' \times d'}$ , and the conditional prior  $p_a(\cdot; f_a(\psi)) = \mathcal{N}(\cdot; W_a \psi, \Sigma_a)$  is Gaussian with mean  $f_a(\psi) = W_a \psi \in \mathbb{R}^d$  and covariance  $\Sigma_a \in \mathbb{R}^{d \times d}$ . The reward distribution  $p(\cdot \mid x; \theta_a)$  is also linear-Gaussian as  $\mathcal{N}(\cdot; \phi(x)^\top \theta_a, \sigma^2)$ , where  $\phi(\cdot)$  outputs a  $d$ -dimensional representation of  $x$  and  $\sigma > 0$  is the observation noise. The whole prior induces a Gaussian graphical model

$$\begin{aligned} \psi &\sim \mathcal{N}(\mu, \Sigma), \\ \theta_a \mid \psi &\sim \mathcal{N}(W_a \psi, \Sigma_a), \quad \forall a \in \mathcal{A}, \\ R \mid \psi, \theta, X, A &\sim \mathcal{N}(\phi(X)^\top \theta_A, \sigma^2). \end{aligned} \quad (10)$$

The conditional action posterior is known in closed-form as  $\theta_a \mid \psi, S \sim \mathcal{N}(\tilde{\mu}_a, \tilde{\Sigma}_a)$ , with

$$\tilde{\Sigma}_a^{-1} = \Sigma_a^{-1} + G_a, \quad \tilde{\Sigma}_a^{-1} \tilde{\mu}_a = \Sigma_a^{-1} W_a \psi + B_a, \quad (11)$$

where  $G_a = \sigma^{-2} \sum_{i \in [n]} \mathbb{1}\{A_i = a\} \phi(X_i) \phi(X_i)^\top$  and  $B_a = \sigma^{-2} \sum_{i \in [n]} \mathbb{1}\{A_i = a\} R_i \phi(X_i)$ . This posterior has the standard form except that the prior mean  $W_a \psi$  now depends on the latent parameter  $\psi$ . Similarly, the effect posterior is known in closed-form as  $\psi \mid S \sim \mathcal{N}(\bar{\mu}, \bar{\Sigma})$ , where

$$\bar{\Sigma}^{-1} = \Sigma^{-1} + \sum_{a \in \mathcal{A}} W_a^\top (\Sigma_a^{-1} - \Sigma_a^{-1} \tilde{\Sigma}_a \Sigma_a^{-1}) W_a, \quad \bar{\Sigma}^{-1} \bar{\mu} = \Sigma^{-1} \mu + \sum_{a \in \mathcal{A}} W_a^\top \Sigma_a^{-1} \tilde{\Sigma}_a B_a. \quad (12)$$

The latent posterior precision  $\bar{\Sigma}^{-1}$  is the sum of the latent prior precision  $\Sigma^{-1}$  and the sum of all learned action precisions  $\Sigma_a^{-1} - \Sigma_a^{-1} \tilde{\Sigma}_a \Sigma_a^{-1}$  weighted by  $W_a^\top \cdot W_a$ . The contribution of a learned



action precision to the global latent precision is proportional to  $W_a^\top W_a$ . A similar intuition can be used to interpret the latent posterior  $\bar{\mu}$ . Finally, from (7), the action posterior is  $\theta_a \mid S \sim \mathcal{N}(\hat{\mu}_a, \hat{\Sigma}_a)$ ,

$$\hat{\Sigma}_a = \tilde{\Sigma}_a + \tilde{\Sigma}_a \Sigma_a^{-1} W_a \bar{\Sigma} W_a^\top \Sigma_a^{-1} \tilde{\Sigma}_a, \quad \hat{\mu}_a = \tilde{\Sigma}_a (\Sigma_a^{-1} W_a \bar{\mu} + B_a). \quad (13)$$

To see why this is more beneficial than the standard prior in (2), notice that the mean and covariance of the posterior of action  $a$ ,  $\hat{\mu}_a$  and  $\hat{\Sigma}_a$ , are now computed using the mean and covariance of the latent posterior,  $\bar{\mu}$  and  $\bar{\Sigma}$ . But  $\bar{\mu}$  and  $\bar{\Sigma}$  are learned using the interactions with all the actions in  $S$ . Thus  $\hat{\mu}_a$  and  $\hat{\Sigma}_a$  are also learned using the interactions with all the actions in  $S$ , in contrast with the non-structured case (2) where they were learned using only the interaction with action  $a$ . The additional computational cost of considering the structured prior in (10) is small. The computational and space complexities are  $\mathcal{O}(K((d^2 + d'^2)(d + d')))$  and  $\mathcal{O}(Kd^2)$ , respectively. For example, when  $d' = \mathcal{O}(d)$ , the above complexities become  $\mathcal{O}(Kd^3)$  and  $\mathcal{O}(Kd^2)$ , respectively. This is exactly the cost of the standard prior in (2). In contrast, this strictly improves the computational efficiency of jointly modeling the action parameters, where the complexities are  $\mathcal{O}(K^3d^3)$  and  $\mathcal{O}(K^2d^2)$  since the joint posterior of  $(\theta_a)_{a \in \mathcal{A}} \mid S$  would require converting and storing a  $dK \times dK$  covariance matrix.

Now that we have a closed-form update for the posterior, we can compute reward estimates. By (10), we have  $r(x, a; \theta) = \phi(x)^\top \theta$ . Thus  $\hat{r}(x, a) = \mathbb{E}[r(x, a; \theta) \mid S] = \phi(x)^\top \hat{\mu}_a$  for any  $(x, a) \in \mathcal{X} \times \mathcal{A}$ , that we can plug in (1). Then, the learned policy is  $\hat{\pi}_G(a \mid x) = \mathbb{1}\{a = \operatorname{argmax}_{b \in \mathcal{A}} \phi(x)^\top \hat{\mu}_b\}$ .

*Remark 4.1.* sDM with this linear Gaussian hierarchy can be used even with data generated from non-linear rewards, and we empirically investigate its robustness to misspecification. We found that this model performs well even if the true rewards are not generated from a linear-Gaussian distribution. Additionally, an extension of sDM to logistic-Bernoulli rewards is provided in [Appendix E.1](#).

## 5 Analysis

We now analyze our proposed method sDM. All proofs are in [Appendix E](#). We first introduce new Bayesian metrics for OPE and OPL, which are of independent interest.

### 5.1 Bayesian metrics

While it is standard to discuss and present OPE before OPL, we first introduce our Bayesian metric for OPL, called Bayesian suboptimality, since it is relatively easier to explain, allowing us to build intuition before presenting the Bayesian mean squared error metric for OPE.

**Bayesian suboptimality (BSO).** In OPL, we compare the learned policy  $\hat{\pi}$  to the optimal  $\pi_* = \operatorname{argmax}_{\pi \in \Pi} V(\pi; \theta_*)$ , using the suboptimality (SO) defined as  $\operatorname{so}(\hat{\pi}; \theta_*) = V(\pi_*; \theta_*) - V(\hat{\pi}; \theta_*)$ . This metric is standard in both offline contextual bandits and reinforcement learning [[Jin et al., 2021](#)]. It is appropriate when the environment is generated by some ground truth, unique and fixed  $\theta_*$ . It can be used to evaluate any policy  $\hat{\pi}$ , e.g. learnt in a frequentist manner (e.g., MLE) or a Bayesian one (e.g., ours). However, it is not the most suitable when the environment is generated by a random variable  $\theta_*$ . Motivated by the recent development of Bayesian analyses for online bandits using the Bayes regret [[Russo and Van Roy, 2014](#)], we introduce in our offline setting a novel metric, namely *Bayes suboptimality* defined as

$$\operatorname{BSO}(\hat{\pi}) = \mathbb{E}[V(\pi_*; \theta_*) - V(\hat{\pi}; \theta_*)], \quad (14)$$

where the expectation is taken over all random variables: the sample set  $S$  and  $\theta_*$ , which is treated as a random variable sampled from the prior. The BSO can be computed in two ways. One method involves taking the expectation under the prior  $\theta_*$ , followed by taking an expectation under data generated from a fixed environment  $\theta_*$  as  $S \mid \theta_*$ . The other method involves taking an expectation under the data  $S$ , followed by taking an expectation under the posterior  $\theta_* \mid S$ . The BSO is a reasonable metric for assessing the average performance of algorithms across multiple environments, due to the expectation over  $\theta_*$ . It is also known that Bayes regret captures the benefits of using informative priors [[Aouali et al., 2023b](#)], and this is similarly achieved by the BSO ([Appendix E.2](#)).

Note that [Hong et al. \[2023\]](#) also examined a concept of Bayesian suboptimality defined as  $V(\pi_*; \theta_*) - V(\hat{\pi}; \theta_*) \mid S$ , where the sample set  $S$  is fixed and randomness only comes from  $\theta_*$  sampled from the posterior. In contrast, our definition aligns with traditional online bandit settings (Bayesian regret [[Russo and Van Roy, 2014](#)]) by averaging the suboptimality across both the sample

set  $S$  and parameter  $\theta_*$  sampled from the prior. This leads to a different metric that captures the average performance under various data realizations and parameter draws. Notably, one major difference is that our metric favors the greedy policy (Section 5.2), while they use pessimism [Hong et al., 2023, Sections 4 and 5].

**Bayesian mean squared error (BMSE).** In OPE, we assess the quality of estimator  $\hat{r}$  using the mean squared error (MSE). Then, similarly to the BSO, we can define the BMSE of the estimator  $\hat{r}$  for a fixed action  $a$  and context  $x$  as

$$\text{BMSE}(\hat{r}(x, a)) = \mathbb{E}[(\hat{r}(x, a) - r(x, a; \theta_*))^2]. \quad (15)$$

The BMSE diverges from the MSE in that the expectation in BMSE is calculated over the sample set  $S$  and across the true parameter  $\theta_*$  sampled from the prior.

## 5.2 Theoretical results

Our theory relies on a well-specified prior assumption, where action parameters  $\theta_{*,a}$  and true rewards are assumed to be drawn from the structured prior (10). We start by presenting our OPE results where we bound the BMSE of sDM. Then, we present our main result, a bound on the BSO of sDM in OPL.

**Theorem 5.1** (OPE Result). *Let  $x \in \mathcal{X}$  and  $a \in \mathcal{A}$ , the BMSE of sDM under prior (10) satisfies  $\text{BMSE}(\hat{r}(x, a)) \leq \mathbb{E}[\|\phi(x)\|_{\hat{\Sigma}_a}^2]$ .*

For any context  $x$  and action  $a$ , the BMSE of sDM is bounded by  $\mathbb{E}[\|\phi(x)\|_{\hat{\Sigma}_a}^2]$ . This is expected, given that we assume access to well-specified prior and likelihood, eliminating bias and directly linking BMSE to the variance. Essentially, estimation accuracy increases as the posterior covariance of action  $a$ ,  $\hat{\Sigma}_a$ , diminishes in the direction of  $\phi(x)$ . For the standard non-structured prior in (2), this would only happen if the context  $x$  and action  $a$  appear frequently in the sample set  $S$ . However, sDM’s use of structured priors ensures lower covariance in the direction of  $\phi(x)$  even without observing context  $x$  and action  $a$  together. This is because sDM calculates posterior covariance for action  $a$  using all observed contexts and actions, which reduces its variance. We will now present our main OPL result.

**Theorem 5.2** (OPL Result). *Let  $\delta \in (0, 1)$  and  $\pi_*(x)$  be the optimal action for context  $x$ . Then the BSO of sDM under prior (10) satisfies*

$$\text{BSO}(\hat{\pi}_G) \leq \beta(d, \delta) \mathbb{E}[\|\phi(X)\|_{\hat{\Sigma}_{\pi_*(X)}}], \quad (16)$$

where  $\beta(d, \delta) = \sqrt{2/\pi} e^{-\alpha(d, \delta)^2/2} + \alpha(d, \delta)$  for  $\alpha(d, \delta) = \sqrt{d + 2\sqrt{d \log 1/\delta} + 2 \log 1/\delta}$ .

Scaling of the bound in Theorem 5.2 aligns with existing frequentist results [Jin et al., 2021, Theorem 4.4]. The main differences lie in the constants, which in our case reflect the benefits of using informative priors (Appendix E.2), and the fact that this rate is achieved using greedy policies (8). This contrasts with the frequentist setting where pessimism is used [Jin et al., 2021] and known to be optimal [Jin et al., 2021, Theorem 4.7]. In fact, greedy policies are optimal when using BSO as a performance metric. Specifically,  $\text{BSO}(\hat{\pi}_G) \leq \text{BSO}(\pi)$  for any policy  $\pi$ , including pessimistic ones [Jin et al., 2021]. Therefore, in the Bayesian setting and when using BSO as a performance metric, greedy policies should always be preferred to pessimistic ones. This fundamental difference is proven in Appendix E.5 and it is of independent interest beyond this work. Theorem 5.2 suggests that the BSO primarily depends on the posterior covariance of action  $\pi_*(X)$  in the direction of the context  $\phi(X)$ . That is, when the uncertainty in the posterior distribution of the optimal action  $\pi_*(X)$  is low on average across different contexts  $X$ , problem instances  $\theta_*$ , and sample sets  $S$ , then the BSO bound is correspondingly small. In particular, the tightness of the bound depends on the degree to which the sample set covers the optimal actions on average. Moreover, this bound would roughly scale as  $\mathcal{O}(1/\sqrt{n})$  if we were to make the *well-explored dataset* assumption, commonly used in the literature [Swaminathan et al., 2017, Jin et al., 2021, Hong et al., 2023].

Even without such an assumption, Theorem 5.2 still highlights the advantages of using sDM over the non-structured prior (2). To see this, notice that the parameters of the non-structured prior in (2),  $\mu_a$  and  $\Sigma_a$ , are obtained by marginalizing out  $\psi$  in (10). In this case,  $\mu_a = W_a \mu$  and  $\Sigma_a = \Sigma_a + W_a \Sigma W_a^\top$ . The corresponding posterior covariance is  $\hat{\Sigma}_a^{\text{NS}} = ((\Sigma_a + W_a \Sigma W_a^\top)^{-1} + G_a)^{-1}$ , and is generally larger than the covariance of sDM,  $\hat{\Sigma}_a$  in (13). This is more pronounced when the






Figure 1: The relative reward of the learned policy on synthetic problems with varying  $n$ ,  $K$  and  $d'$ .

number of actions  $K$  is large and when the latent parameters are more uncertain than the action parameters. Thus, the BSO bound of sDM is smaller due to the reduced posterior uncertainty it exhibits. Also, note that even when  $\pi_*(X)$  is unobserved in the sample set  $S$ , sDM’s posterior covariance  $\hat{\Sigma}_{\pi_*(X)}$  can remain small since we use interactions with all actions to compute it. This contrasts with standard non-structured priors (2), where observing  $\pi_*(X)$  is necessary; without such observations, the posterior covariance  $\hat{\Sigma}_{\pi_*(X)}$  would simply be the prior covariance  $\Sigma_{\pi_*(X)}$ .

## 6 Experiments

We evaluate sDM using synthetic and MovieLens datasets. We focus on OPL experiments and use the average reward relative to the optimal policy as our evaluation metric, which is equivalent to BSO. Additional OPE experiments are provided in Appendix F.1 and they lead to similar conclusions.

**Synthetic problems.** We simulate synthetic data using the model in (10) with  $\sigma = 1$ . The contexts  $X$  are sampled uniformly from  $[-1, 1]^d$ , with  $d = 10$ . The matrices  $W_a$  are sampled uniformly from  $[-1, 1]^{d \times d'}$ , where  $d'$  varies as  $d' \in \{5, 10, 20\}$ . We set  $\Sigma = 3I_{d'}$  and  $\Sigma_a = I_d$  (the latent parameters are more uncertain than the action parameters). Finally, the latent mean  $\mu$  is randomly sampled from  $[-1, 1]^{d'}$ . The number of actions varies as  $K \in \{10^2, 10^3\}$ . We use a uniform logging policy to collect data while we defer the results with  $\epsilon$ -greedy logging policies to Appendix F.2.

We evaluate several baselines. First, we use sDM under prior (10). Second, we examine DM (Bayes), which uses the standard non-structured prior (2), where parameters  $\mu_a$  and  $\Sigma_a$  are obtained by marginalizing out the latent parameters  $\psi$  in (10). Thus DM (Bayes) is a standard Bayesian DM that does not capture arm reward correlations. We also include DM (Freq), which estimates  $\theta_{*,a}$  by the MLE. We include IPS [Horvitz and Thompson, 1952], self-normalized IPS (snIPS) [Swaminathan and Joachims, 2015b], and doubly robust (DR) [Dudik et al., 2014], which we optimize to learn the optimal policy. Note that the goal of these experiments is not to compare with IPS and its variants. In particular, our experimental setting is better suited for DMs due to its high number of actions, known reward function, and well-specified priors, which favor DMs over IPS variants. Nevertheless, we include the IPS variants for the sake of completeness only. In Fig. 1, we plot the average relative reward with varying  $K$ ,  $d'$ ,  $n$ . Overall, sDM consistently outperforms the baselines in terms of both sample efficiency and average reward. This performance gap becomes even more significant in scenarios with a high number of actions  $K$  or small sample size  $n$ . These results highlight sDM’s enhanced efficiency in using available logged data, making it particularly beneficial in data-limited situations and scalable to large action spaces.

**Large action spaces.** sDM achieves improved scalability, in the sense that its performance remains robust in large action spaces, by leveraging data more efficiently. While it still learns a  $d$ -dim. parameter for each action  $a$ , it does so by considering interactions with all actions in the sample set  $S$ , instead of only using interactions with the specific action  $a$ . This is crucial, especially given that many actions may not even be observed in  $S$ . To show sDM’s improved scalability, we compared it to the most competitive baseline, DM (Bayes), for varying  $K \in [10, 10^5]$  with  $n = 10^3$ . The results in Fig. 2 reveal that




Figure 2: Comparison of sDM and DM (Bayes) with varying  $K$ .

the performance gap between sDM and DM (Bayes) becomes more significant when the number of actions  $K$  increases. Hence, despite the necessity for sDM to learn distinct parameters for each action, accommodating practical scenarios like recommender systems where unique embeddings are learned for each product, it still enjoys good scalability.

**MovieLens problems.** We also experiment with the MovieLens 1M dataset [Lam and Herlocker, 2016]. This dataset contains 1M ratings, reflecting the interactions of a set of 6,040 users with a collection of 3,952 movies. Our evaluation approach involves the creation of a semi-synthetic environment as follows. Initially, we apply a low-rank factorization technique to the rating matrix, yielding 5-dimensional representations denoted as  $x_u \in \mathbb{R}^5$  for user  $u \in [6040]$  and  $\theta_a \in \mathbb{R}^5$  for movie  $a \in [3952]$ . Subsequently, the movies are treated as actions, and the contexts  $X$  are sampled randomly from the user vectors. The reward associated with movie  $a$  for user  $u$  is modeled as  $\mathcal{N}(x_u^\top \theta_a, 1)$ , effectively representing a proxy for the rating. We use a uniform logging policy to generate logged data and experiments with  $\epsilon$ -greedy logging policies are in Appendix F.2.

A prior is not needed for DM (Freq), IPS, snIPS, and DR. However, for DM (Bayes), a standard prior (2) is inferred from data, where we set  $\mu_a$  to be the mean of movie vectors across all dimensions, and  $\Sigma_a = \text{diag}(v)$ , where  $v$  represents the variance of movie vectors across all dimensions. Unlike the synthetic experiments, the latent structure assumed by sDM is not inherently present in MovieLens. But we learn it by training a Gaussian Mixture Model (GMM) to cluster movies into  $J = 5$  mixture components. This gives rise to the mixed-effect structure described in Section 4, which represents a specific instance of sDM with  $d' = dJ = 25$ . Both DM (Bayes) and sDM use the same subset of data to learn their priors and it is of size  $\{100, 1000\}$  depending on the setting. We conduct two experiments with  $K \in \{10^2, 10^3\}$  randomly selected movies, and the results are in Fig. 3. Even though the latent structure assumed by sDM is absent, sDM still outperforms the baselines.

**Prior misspecification.** We consider the same synthetic problems above, but where the parameters of the prior distribution used to sample the true parameters differ from those employed by sDM. This was achieved by adding uniformly sampled noise from  $[v, v + 0.5]$  to the true prior parameters  $\mu, \Sigma, W_a, \Sigma_a$ , with  $v$  controlling the level of misspecification. We varied  $v \in \{0.5, 1, 1.5\}$  and analyzed its impact on sDM’s performance. For comparison, we included the well-specified sDM and the most competitive baseline, DM (Bayes), while omitting other baselines to reduce clutter. Results are shown in Fig. 4. As expected, sDM’s performance decreases with increasing misspecification. However, note that the entries of the true parameters  $W_a$  and  $\Sigma_a$  and  $\mu$  are in  $[-1, 1]$ , so values  $v \in \{0.5, 1, 1.5\}$  can lead to significant misspecification, yet sDM outperforms the most competitive baseline, especially when  $K$  is large.

**Likelihood misspecification.** We also simulate the case when the true reward distribution differed from the likelihood assumed by sDM. For example, we simulated binary rewards using a Bernoulli-logistic model while sDM used a linear-Gaussian likelihood. Other DMs: DM (Bayes) and DM (Freq) also use a misspecified likelihood model and to emphasize this we add the suffix Lin to all DMs names. Overall, as known in the literature, IPS variants performed better than standard DMs when the likelihood was misspecified. However, this is not the case




Figure 3: Average relative reward of the learned policy on MovieLens with varying  $n$  and  $K$ .




Figure 4: Effect of prior misspecification.




Figure 5: Effect of likelihood misspecification.

for sDM. While misspecification narrowed the performance gap between sDM and baselines, sDM still consistently outperformed them, particularly in low-data regimes or large action spaces. In a few cases, some baselines, e.g., snIPS, showed competitive performance with enough data.

Additional experiments on prior and likelihood misspecification are provided in [Appendix F.3](#). For completeness, we also compare sDM with MIPS [[Saito and Joachims, 2022](#)] in [Appendix F.4](#). However, the goal of this work is not to compare with MIPS or other IPS variants, but to improve DMs. Therefore, we compared our DM (sDM) to standard DMs (DM (Bayes) and DM (Freq)) to highlight the benefits of using structured priors. The inclusion of IPS variants is for completeness only. Finally, [Appendix F.5](#) compares pessimistic and greedy policies in our Bayesian setting, validating our theory.

## 7 Conclusion

We introduced sDM, a novel approach for OPE and OPL with large action spaces that leverages latent structures among actions to enhance statistical and computational efficiency. We provide theoretical justifications for sDM’s improved concentration properties towards the optimal policy and demonstrate the superiority of greedy policies over pessimistic ones in the Bayesian setting. Our experiments validate our theory, showing that the introduced Bayesian metrics highlight the advantages of structured priors, especially in large action spaces with limited data. However, a primary limitation of our theory is the well-specified prior assumption. While we have empirically studied the effect of misspecification, investigating it theoretically is left for future work. Moreover, extending sDM’s applicability to non-linear hierarchies is an important direction for further investigation.

## References

- Yasin Abbasi-Yadkori, David Pal, and Csaba Szepesvari. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems 24*, pages 2312–2320, 2011.
- Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *Proceedings of the 30th International Conference on Machine Learning*, pages 127–135, 2013.
- Imad Aouali. Linear diffusion models meet contextual bandits with large action spaces. In *NeurIPS 2023 Foundation Models for Decision Making Workshop*, 2023.
- Imad Aouali, Amine Benhalloum, Martin Bompaire, Achraf Ait Sidi Hammou, Sergey Ivanov, Benjamin Heymann, David Rohde, Otmane Sakhi, Flavian Vasile, and Maxime Vono. Reward optimizing recommendation using deep learning and fast maximum inner product search. In *proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, pages 4772–4773, 2022a.
- Imad Aouali, Achraf Ait Sidi Hammou, Sergey Ivanov, Otmane Sakhi, David Rohde, and Flavian Vasile. Probabilistic rank and reward: A scalable model for slate recommendation, 2022b.
- Imad Aouali, Victor-Emmanuel Brunel, David Rohde, and Anna Korba. Exponential smoothing for off-policy learning. In *International Conference on Machine Learning*, pages 984–1017. PMLR, 2023a.
- Imad Aouali, Branislav Kveton, and Sumeet Katariya. Mixed-effect thompson sampling. In *International Conference on Artificial Intelligence and Statistics*, pages 2087–2115. PMLR, 2023b.
- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, 2002.
- Heejung Bang and James M Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005.
- Léon Bottou, Jonas Peters, Joaquin Quiñero-Candela, Denis X Charles, D Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Simard, and Ed Snelson. Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research*, 14(11), 2013.

- Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics*, pages 208–214, 2011.
- Matej Cief, Jacek Golebiowski, Philipp Schmidt, Ziawasch Abedjan, and Artur Bekasov. Learning action embeddings for off-policy evaluation. In *European Conference on Information Retrieval*, pages 108–122. Springer, 2024.
- Miroslav Dudík, John Langford, and Lihong Li. Doubly robust policy evaluation and learning. *International Conference on Machine Learning*, 2011.
- Miroslav Dudík, Dumitru Erhan, John Langford, and Lihong Li. Sample-efficient nonstationary policy evaluation for contextual bandits. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, UAI’12, page 247–254, Arlington, Virginia, USA, 2012. AUAI Press.
- Miroslav Dudík, Dumitru Erhan, John Langford, and Lihong Li. Doubly robust policy evaluation and optimization. *Statistical Science*, 29(4):485–511, 2014.
- Mehrdad Farajtabar, Yinlam Chow, and Mohammad Ghavamzadeh. More robust doubly robust off-policy evaluation. In *International Conference on Machine Learning*, pages 1447–1456. PMLR, 2018.
- Alexandre Gilotte, Clément Calauzènes, Thomas Nedelec, Alexandre Abraham, and Simon Dollé. Offline a/b testing for recommender systems. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 198–206, 2018.
- Joey Hong, Branislav Kveton, Manzil Zaheer, Sumeet Katariya, and Mohammad Ghavamzadeh. Multi-task off-policy learning from bandit feedback. In *International Conference on Machine Learning*, pages 13157–13173. PMLR, 2023.
- Daniel G Horvitz and Donovan J Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, 47(260):663–685, 1952.
- Edward L Ionides. Truncated importance sampling. *Journal of Computational and Graphical Statistics*, 17(2):295–311, 2008.
- Olivier Jeunen and Bart Goethals. Pessimistic reward models for off-policy learning in recommendation. In *Fifteenth ACM Conference on Recommender Systems*, pages 63–74, 2021.
- Ying Jin, Zhuoran Yang, and Zhaoran Wang. Is pessimism provably efficient for offline rl? In *International Conference on Machine Learning*, pages 5084–5096. PMLR, 2021.
- Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, Cambridge, MA, 2009.
- Ilja Kuzborskij, Claire Vernade, Andras Gyorgy, and Csaba Szepesvári. Confident off-policy evaluation and selection through self-normalized importance weighting. In *International Conference on Artificial Intelligence and Statistics*, pages 640–648. PMLR, 2021.
- Shyong Lam and Jon Herlocker. MovieLens Dataset. <http://grouplens.org/datasets/movielens/>, 2016.
- Tor Lattimore and Csaba Szepesvari. *Bandit Algorithms*. Cambridge University Press, 2019.
- Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, pages 1302–1338, 2000.
- Alessandro Lazaric and Mohammad Ghavamzadeh. Bayesian multi-task reinforcement learning. In *ICML-27th international conference on machine learning*, pages 599–606. Omnipress, 2010.
- Lihong Li, Wei Chu, John Langford, and Robert Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web*, 2010.

- Ben London and Ted Sandler. Bayesian counterfactual risk minimization. In *International Conference on Machine Learning*, pages 4125–4133. PMLR, 2019.
- Alberto Maria Metelli, Alessio Russo, and Marcello Restelli. Subgaussian and differentiable importance sampling for off-policy evaluation and learning. *Advances in Neural Information Processing Systems*, 34:8119–8132, 2021.
- Jie Peng, Hao Zou, Jiashuo Liu, Shaoming Li, Yibao Jiang, Jian Pei, and Peng Cui. Offline policy evaluation in large action spaces via outcome-oriented action grouping. In *Proceedings of the ACM Web Conference 2023*, pages 1220–1230, 2023.
- James M Robins and Andrea Rotnitzky. Semiparametric efficiency in multivariate regression models with missing data. *Journal of the American Statistical Association*, 90(429):122–129, 1995.
- Daniel Russo and Benjamin Van Roy. Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 39(4):1221–1243, 2014.
- Daniel Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, and Zheng Wen. A tutorial on Thompson sampling. *Foundations and Trends in Machine Learning*, 11(1):1–96, 2018.
- Noveen Sachdeva, Yi Su, and Thorsten Joachims. Off-policy bandits with deficient support. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 965–975, 2020.
- Noveen Sachdeva, Lequn Wang, Dawen Liang, Nathan Kallus, and Julian McAuley. Off-policy evaluation for large action spaces via policy convolution. *arXiv preprint arXiv:2310.15433*, 2023.
- Yuta Saito and Thorsten Joachims. Off-policy evaluation for large action spaces via embeddings. *arXiv preprint arXiv:2202.06317*, 2022.
- Yuta Saito, Qingyang Ren, and Thorsten Joachims. Off-policy evaluation for large action spaces via conjunct effect modeling. In *international conference on Machine learning*, pages 29734–29759. PMLR, 2023.
- Otmame Sakhi, Stephen Bonner, David Rohde, and Flavian Vasile. Blob: A probabilistic model for recommendation that combines organic and bandit signals. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 783–793, 2020.
- Otmame Sakhi, Nicolas Chopin, and Pierre Alquier. Pac-bayesian offline contextual bandits with guarantees. *arXiv preprint arXiv:2210.13132*, 2022.
- Claudia Schillings, Björn Sprungk, and Philipp Wacker. On the convergence of the laplace approximation and noise-level-robustness of laplace-based monte carlo methods for bayesian inverse problems. *Numerische Mathematik*, 145:915–971, 2020.
- David J Spiegelhalter and Steffen L Lauritzen. Sequential updating of conditional probabilities on directed graphical structures. *Networks*, 20(5):579–605, 1990.
- Yi Su, Lequn Wang, Michele Santacatterina, and Thorsten Joachims. Cab: Continuous adaptive blending for policy evaluation and learning. In *International Conference on Machine Learning*, pages 6005–6014. PMLR, 2019.
- Yi Su, Maria Dimakopoulou, Akshay Krishnamurthy, and Miroslav Dudík. Doubly robust off-policy evaluation with shrinkage. In *International Conference on Machine Learning*, pages 9167–9176. PMLR, 2020.
- Adith Swaminathan and Thorsten Joachims. Batch learning from logged bandit feedback through counterfactual risk minimization. *The Journal of Machine Learning Research*, 16(1):1731–1755, 2015a.
- Adith Swaminathan and Thorsten Joachims. The self-normalized estimator for counterfactual learning. *advances in neural information processing systems*, 28, 2015b.

- Adith Swaminathan, Akshay Krishnamurthy, Alekh Agarwal, Miro Dudik, John Langford, Damien Jose, and Imed Zitouni. Off-policy evaluation for slate recommendation. *Advances in Neural Information Processing Systems*, 30, 2017.
- Muhammad Faaiz Taufiq, Arnaud Doucet, Rob Cornish, and Jean-Francois Ton. Marginal density ratio for off-policy evaluation in contextual bandits. *Advances in Neural Information Processing Systems*, 36, 2024.
- William R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3-4):285–294, 1933.
- Lequn Wang, Akshay Krishnamurthy, and Aleksandrs Slivkins. Oracle-efficient pessimism: Offline policy optimization in contextual bandits. *arXiv preprint arXiv:2306.07923*, 2023.
- Yu-Xiang Wang, Alekh Agarwal, and Miroslav Dudik. Optimal and adaptive off-policy evaluation in contextual bandits. In *International Conference on Machine Learning*, pages 3589–3597. PMLR, 2017.
- Neil Weiss. *A Course in Probability*. Addison-Wesley, 2005.
- Yinglun Zhu, Dylan J Foster, John Langford, and Paul Mineiro. Contextual bandits with large action spaces: Made practical. In *International Conference on Machine Learning*, pages 27428–27453. PMLR, 2022.
- Shi Zong, Hao Ni, Kenny Sung, Nan Rosemary Ke, Zheng Wen, and Branislav Kveton. Cascading bandits for large-scale recommendation problems. *arXiv preprint arXiv:1603.05359*, 2016.



## Organization of the supplementary material

The supplementary material is organized as follows.

- In [Appendix A](#), we provide a detailed notation.
- In [Appendix B](#), we provide an extended related work discussion.
- In [Appendix C](#), we outline the posterior derivations for the standard prior in (2).
- In [Appendix D](#), we outline the posterior derivations under the structured prior in (10).
- In [Appendix E](#), we prove the claims made in [Section 5](#).
- In [Appendix F](#), we present supplementary experiments.

## A Detailed notation

For any positive integer  $n$ , we define  $[n]$  as the set  $1, 2, \dots, n$ . We represent the identity matrix of dimension  $d \times d$  as  $I_d$ . In our notation, unless explicitly stated otherwise, the  $i$ -th coordinate of a vector  $v$  is denoted as  $v_i$ . However, if the vector is already indexed, such as  $x_i$ , we use the notation  $x_{i,j}$  to represent the  $j$ -th entry of the vector  $x_i$ . When dealing with matrices, we refer to the  $(i, j)$ -th entry of a matrix  $A$  as  $A_{i,j}$ . Also,  $\lambda_1(A)$  and  $\lambda_d(A)$  refer to the maximum and minimum eigenvalues of matrix  $A$ , respectively. Moreover, for any positive-definite matrix  $A$  and vector  $x$ , we let  $\|x\|_A = \sqrt{x^\top A x}$ .

Now, consider a collection of  $n$  vectors, denoted as  $a_1 \in \mathbb{R}^d, a_2 \in \mathbb{R}^d, \dots, a_n \in \mathbb{R}^d$ . Then we use  $a = (a_i)_{i \in [n]} \in \mathbb{R}^{nd}$  to represent a  $nd$ -dimensional vector formed by concatenating the vectors  $a_1, a_2, \dots, a_n$ . The operator  $\text{Vec}(\cdot)$  is used to vectorize a matrix or a set of vectors. Let  $A_1, A_2, \dots, A_n$  be a collection of  $n$  matrices, each of dimension  $d \times d$ . The notation  $\text{diag}((A_i)_{i \in [n]}) \in \mathbb{R}^{nd \times nd}$  represents a block diagonal matrix where  $A_1, A_2, \dots, A_n$  are the main-diagonal blocks. Similarly,  $(A_i)_{i \in [n]} \in \mathbb{R}^{nd \times d}$  denotes the  $nd \times d$  matrix formed by concatenating matrices  $A_1, A_2, \dots, A_n$ . To provide a more visual representation, consider a collection of vectors  $a_1 \in \mathbb{R}^d, a_2 \in \mathbb{R}^d, \dots, a_n \in \mathbb{R}^d$  and matrices  $A_1 \in \mathbb{R}^{d \times d}, A_2 \in \mathbb{R}^{d \times d}, \dots, A_n \in \mathbb{R}^{d \times d}$ . We can represent them as follows

$$(a_i)_{i \in [n]} = \begin{pmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{pmatrix} \in \mathbb{R}^{nd}, \quad \text{diag}((A_i)_{i \in [n]}) = \begin{pmatrix} A_1 & 0 & \cdots & 0 \\ 0 & A_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_n \end{pmatrix} \in \mathbb{R}^{nd \times nd},$$

$$(A_i)_{i \in [n]} = \begin{pmatrix} A_1 \\ A_2 \\ \vdots \\ A_n \end{pmatrix} \in \mathbb{R}^{nd \times d}.$$

## B Extended related work

Online learning under uncertainty is often modeled within the framework of contextual bandits [[Lattimore and Szepesvari, 2019](#), [Li et al., 2010](#), [Chu et al., 2011](#)]. Naturally, learning within this framework follows an online paradigm. However, practical applications often entail a large action space and a significant emphasis on short-term gains. This presents a challenge where the agent needs to exhibit a *risk-averse* behavior, contradicting the foundational principle of online algorithms designed to explore actions for long-term benefit [[Auer et al., 2002](#), [Thompson, 1933](#), [Russo et al., 2018](#), [Agrawal and Goyal, 2013](#), [Abbasi-Yadkori et al., 2011](#)]. While several practical online algorithms have emerged to efficiently explore the action space in contextual bandits [[Zong et al., 2016](#), [Zhu et al., 2022](#), [Aouali et al., 2023b](#), [Aouali, 2023](#)], a notable gap exists in the quest for an offline procedure capable of optimizing decision-making using historical data. Fortunately, we are often equipped with a sample set collected from past interactions. Leveraging this data, agents can enhance their policies offline [[Swaminathan and Joachims, 2015a](#), [London and Sandler, 2019](#), [Sakhi et al., 2022](#)], consequently improving the overall system performance. This study is primarily

concerned with this offline or *off-policy* formulation of contextual bandits [Dudík et al., 2011, Dudík et al., 2012, Dudík et al., 2014, Wang et al., 2017, Farajtabar et al., 2018]. There are two main tasks in off-policy contextual bandits, first, off-policy evaluation (OPE) [Dudík et al., 2011] which involves estimating policy performance using historical data, simulating as if these evaluations were performed while the policy is interacting with the environment in real-time. Subsequently, the derived estimator is refined to closely approximate the optimal policy, a process known as off-policy learning (OPL) [Swaminathan and Joachims, 2015a]. Next, we review both OPE and OPL.

## B.1 Off-policy evaluation

Recent years have witnessed a surge in interest in OPE, with several key works contributing to this field [Dudík et al., 2011, Dudík et al., 2012, Dudík et al., 2014, Wang et al., 2017, Farajtabar et al., 2018, Su et al., 2019, 2020, Metelli et al., 2021, Kuzborskij et al., 2021, Saito and Joachims, 2022, Sakhi et al., 2020, Jeunen and Goethals, 2021]. The literature on OPE can be broadly categorized into three main approaches. The first approach referred to as the direct method (DM) [Jeunen and Goethals, 2021, Hong et al., 2023], involves learning a model that approximates the expected reward. This model is then used to estimate the performance of evaluated policies. The second approach, inverse propensity scoring (IPS) [Horvitz and Thompson, 1952, Dudík et al., 2012], aims to estimate the reward of evaluated policies by correcting for the preference bias of the logging policy in the sample set. IPS is unbiased when there’s an assumption that the evaluation policy is absolutely continuous with respect to the logging policy. However, it can exhibit high variance and substantial bias when this assumption is violated [Sachdeva et al., 2020]. Various techniques have been introduced to address the variance issue, such as clipping importance weights [Ionides, 2008, Swaminathan and Joachims, 2015a], smoothing them [Aouali et al., 2023a], self-normalization of weights [Swaminathan and Joachims, 2015b], etc. [Gilotte et al., 2018]. The third approach, known as doubly robust (DR) [Robins and Rotnitzky, 1995, Bang and Robins, 2005, Dudík et al., 2011, Dudík et al., 2014, Farajtabar et al., 2018], combines elements of DM and IPS, helping to reduce variance. Assessing the accuracy of an OPE estimator, denoted as  $\hat{R}_n(\pi)$ , is typically done using the mean squared error (MSE). It may be relevant to note that Metelli et al. [2021] argued that high-probability concentration rates should be preferred over the MSE to evaluate OPE estimators as they provide non-asymptotic guarantees. In this work, we presented a direct method for OPE, for which we could derive both an MSE and high-probability concentration bound under our assumptions. However, we focused more on OPL.

## B.2 Off-policy learning

**OPL under IPS.** Prior research on OPL predominantly revolved around developing learning principles influenced by generalization bounds for IPS. Swaminathan and Joachims [2015a] proposed to penalize the IPS estimator with a variance term leading to a learning principle that promotes policies that achieve high estimated reward and exhibit minimal empirical variance. In a recent development, London et al. [London and Sandler, 2019] made a connection between PAC-Bayes theory and OPL that inspired several works in the same vein [Sakhi et al., 2022, Aouali et al., 2023a]. London and Sandler [2019] proposed a learning principle that favors policies with high estimated reward while keeping the parameter close to that of the logging policy in terms of  $L_2$  distance. In particular, this learning principle is scalable. Then Sakhi et al. [2022], Aouali et al. [2023a] took a different direction of deriving tractable generalization bounds and optimizing them directly as they are. Finally, Wang et al. [2023] recently introduced an efficient learning principle with guarantees for a specific choice of their hyperparameter. All of these learning principles are related to the concept of pessimism which we discuss next.

**OPL under DM.** The majority of OPL approaches in contextual bandits rely on the principle of pessimism [Jeunen and Goethals, 2021, Jin et al., 2021, Hong et al., 2023]. In essence, these methods construct confidence intervals for reward estimates  $\hat{r}(x, a)$ , that hold simultaneously for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$  and satisfy the following condition

$$|r(x, a; \theta) - \hat{r}(x, a)| \leq u(x, a), \quad \forall (x, a) \in \mathcal{X} \times \mathcal{A}$$

Subsequently, the policy learned with this pessimistic approach is defined as

$$\hat{\pi}_p(a | x) = \mathbb{1}\{a = \operatorname{argmax}_{b \in \mathcal{A}} \hat{r}(x, b) - u(x, b)\}. \quad (17)$$

Note that when the function  $u(x, a)$  is independent of both  $x$  and  $a$ , the pessimistic and greedy policies become equivalent. However, in this work, we highlight the critical role of the selected metric and the underlying assumptions about  $\hat{r}$  in choosing between pessimistic and greedy policies. In particular, we introduce the concept of Bayesian suboptimality (BSO), as defined in [Section 5.1](#). BSO evaluates the average performance of algorithms across multiple problem instances, and in this context, the Greedy policy emerges as a more suitable choice for minimizing the BSO. A detailed comparison of pessimistic and greedy policies is presented in [Appendix E.5](#).

In addition to not using pessimism, our work also considers structured contextual bandit problems, upon which we build informative priors. Notably, while structured problems have been previously investigated in the context of multi-task learning by [Lazaric and Ghavamzadeh \[2010\]](#), [Hong et al. \[2023\]](#), these works primarily focused on cases where there are multiple contextual bandit instances that bear similarity to one another. Interaction with one bandit instance contributes to the agent's understanding of other instances. In contrast, our work tackles a single contextual bandit instance with a large action space and considers an inherent structure among actions within this bandit instance. Furthermore, the structure we address in the single-instance problem is different and more general compared to the ones explored in the aforementioned multi-task learning works.

## C Posterior derivations under standard priors

Here we derive the posterior under the standard prior in (2). These are standard derivations and we present them here for the sake of completeness. But first, we state the following standard assumption that allows posterior derivations.

**Assumption C.1** (Independence).  $(X, A)$  is independent of  $\theta$ , and the  $\theta_a$ , for  $a \in \mathcal{A}$  are independent.

*Derivation of  $p(\theta_a | S)$  for the standard prior in (2).* We start by recalling the standard prior in (2)

$$\begin{aligned} \theta_a &\sim \mathcal{N}(\mu_a, \Sigma_a), & \forall a \in \mathcal{A}, \\ R | \theta, X, A &\sim \mathcal{N}(\phi(X)^\top \theta_A, \sigma^2), \end{aligned} \quad (18)$$

where  $\mathcal{N}(\mu_a, \Sigma_a)$  is the prior on the action parameter  $\theta_a$ . Let  $\theta = (\theta_a)_{a \in \mathcal{A}} \in \mathbb{R}^{dK}$ ,  $\Sigma_A = \text{diag}(\Sigma_a)_{a \in \mathcal{A}} \in \mathbb{R}^{dK \times dK}$  and  $\mu_A = (\mu_a)_{a \in \mathcal{A}} \in \mathbb{R}^{dK}$ . Also, let  $u_a \in \{0, 1\}^K$  be the binary vector representing the action  $a$ . That is,  $u_{a,a} = 1$  and  $u_{a,a'} = 0$  for all  $a' \neq a$ . Then we can rewrite the model in (18) as

$$\begin{aligned} \theta &\sim \mathcal{N}(\mu_A, \Sigma_A), \\ R | \theta, X, A &\sim \mathcal{N}((u_A \otimes \phi(X))^\top \theta, \sigma^2). \end{aligned} \quad (19)$$

Then the *joint* action posterior  $p(\theta | S)$  decomposes as

$$\begin{aligned} p(\theta | S) &= p(\theta | (X_i, A_i, R_i)_{i \in [n]}) \stackrel{(i)}{\propto} p((R_i)_{i \in [n]} | \theta, (X_i, A_i)_{i \in [n]}) p(\theta | (X_i, A_i)_{i \in [n]}), \\ &\stackrel{(ii)}{=} p((R_i)_{i \in [n]} | \theta, (X_i, A_i)_{i \in [n]}) p(\theta) \stackrel{(iii)}{=} \prod_{i \in [n]} p(R_i | \theta, X_i, A_i) p(\theta), \\ &\stackrel{(iv)}{=} \prod_{i \in [n]} \mathcal{N}(R_i; (u_{A_i} \otimes \phi(X_i))^\top \theta, \sigma^2) \mathcal{N}(\theta; \mu_A, \Sigma_A), \\ &= \exp \left[ -\frac{1}{2} \left( v \sum_{i=1}^n (R_i^2 - 2R_i (u_{A_i} \otimes \phi(X_i))^\top \theta + ((u_{A_i} \otimes \phi(X_i))^\top \theta)^2) + \theta^\top \Lambda_A \theta - 2\theta^\top \Lambda_A \mu_A + \mu_A^\top \Lambda_A \mu_A \right) \right], \\ &\propto \exp \left[ -\frac{1}{2} \left( \theta^\top \left( v \sum_{i=1}^n (u_{A_i} \otimes \phi(X_i))(u_{A_i} \otimes \phi(X_i))^\top + \Lambda_A \right) \theta - 2\theta^\top \left( v \sum_{i=1}^n (u_{A_i} \otimes \phi(X_i))^\top R_i + \Lambda_A \mu_A \right) \right) \right], \\ &= \exp \left[ -\frac{1}{2} \left( \theta^\top \left( v \sum_{i=1}^n (u_{A_i} u_{A_i}^\top \otimes \phi(X_i) \phi(X_i)^\top) + \Lambda_A \right) \theta - 2\theta^\top \left( v \sum_{i=1}^n (u_{A_i} \otimes \phi(X_i))^\top R_i + \Lambda_A \mu_A \right) \right) \right], \\ &\stackrel{(v)}{\propto} \mathcal{N} \left( \theta; \hat{\mu}_A, \left( \hat{\Lambda}_A \right)^{-1} \right). \end{aligned}$$

In (i), we apply Bayes rule. In (ii), we use that  $\theta$  is independent of  $(X, A)$ , and (iii) follows from the assumption that  $R_i \mid \theta, X_i, A_i$  are i.i.d. Finally, in (iv), we replace the distribution by their Gaussian form, and in (v), we set  $\hat{\Lambda}_{\mathcal{A}} = v \sum_{i=1}^n (u_{A_i} u_{A_i}^\top \otimes \phi(X_i) \phi(X_i)^\top) + \Lambda_{\mathcal{A}}$ , and  $\hat{\mu}_a = \hat{\Lambda}_{\mathcal{A}}^{-1} (v \sum_{i=1}^n (u_{A_i} \otimes \phi(X_i))^\top R_i + \Lambda_{\mathcal{A}} \mu_{\mathcal{A}})$ . Now notice that  $\hat{\Lambda}_{\mathcal{A}} = \text{diag}(\Sigma_a^{-1} + G_a)_{a \in \mathcal{A}}$ . Thus,  $p(\theta \mid S) = \mathcal{N}(\theta, \hat{\mu}_{\mathcal{A}}, \hat{\Lambda}_{\mathcal{A}}^{-1})$  where  $\hat{\mu}_{\mathcal{A}} = (\hat{\mu}_a)_{a \in \mathcal{A}}$  and  $\hat{\Lambda}_{\mathcal{A}} = \text{diag}(\hat{\Lambda}_a)_{a \in \mathcal{A}}$ , with

$$\hat{\Lambda}_a = \Sigma_a^{-1} + G_a, \quad \hat{\Lambda}_a \hat{\mu}_a = \Sigma_a^{-1} \mu_a + B_a.$$

Since the covariance matrix of  $p(\theta \mid S)$  is diagonal by block, we know that the marginals  $\theta_a \mid S$  also have a Gaussian density  $p(\theta_a \mid S) = \mathcal{N}(\theta_a; \hat{\mu}_a, \hat{\Sigma}_a)$  where  $\hat{\Sigma}_a = \hat{\Lambda}_a^{-1}$ .

□

## D Posterior derivations under structured priors

Here we derive the posteriors under the structured prior in (10). Precisely, we derive the latent posterior density of  $\psi \mid S$ , the conditional posterior density of  $\theta \mid S, \psi$ . Then, we derive the marginal posterior  $\theta \mid S$ . Posterior derivations rely on the following assumption.

**Assumption D.1** (Structured Independence). (i)  $(X, A)$  is independent of  $\psi$  and given  $\psi$ ,  $(X, A)$  is independent of  $\theta$ . (ii) Given  $\psi$ , the  $\theta_a$ , for all  $a \in \mathcal{A}$  are independent.

### D.1 Latent posterior

*Derivation of  $p(\psi \mid S)$ .* First, recall that our model in (10) reads

$$\begin{aligned} \psi &\sim \mathcal{N}(\mu, \Sigma), \\ \theta_a \mid \psi &\sim \mathcal{N}(W_a \psi, \Sigma_a), \quad \forall a \in \mathcal{A}, \\ R \mid \psi, \theta, X, A &\sim \mathcal{N}(\phi(X)^\top \theta_{\mathcal{A}}, \sigma^2). \end{aligned} \quad (20)$$

Then we first rewrite it as

$$\begin{aligned} \psi &\sim \mathcal{N}(\mu, \Sigma), \\ \theta \mid \psi &\sim \mathcal{N}(W_{\mathcal{A}} \psi, \Sigma_{\mathcal{A}}), \\ R \mid \psi, \theta, X, A &\sim \mathcal{N}((u_{\mathcal{A}} \otimes \phi(X))^\top \theta, \sigma^2). \end{aligned} \quad (21)$$

Then the latent posterior is

$$\begin{aligned} p(\psi \mid (X_i, A_i, R_i)_{i \in [n]}) &\propto p((R_i)_{i \in [n]} \mid \psi, (X_i, A_i)_{i \in [n]}) p(\psi \mid (X_i, A_i)_{i \in [n]}), \\ &\stackrel{(i)}{=} p((R_i)_{i \in [n]} \mid \psi, (X_i, A_i)_{i \in [n]}) q(\psi), \\ &= \int_{\theta} p((R_i)_{i \in [n]}, \theta \mid \psi, (X_i, A_i)_{i \in [n]}) d\theta q(\psi), \\ &= \int_{\theta} p((R_i)_{i \in [n]} \mid \psi, \theta, (X_i, A_i)_{i \in [n]}) p(\theta \mid \psi, (X_i, A_i)_{i \in [n]}) d\theta q(\psi), \\ &\stackrel{(ii)}{=} \int_{\theta} p((R_i)_{i \in [n]} \mid \psi, \theta, (X_i, A_i)_{i \in [n]}) p(\theta \mid \psi) d\theta q(\psi), \end{aligned}$$

In (i), we use that  $(X, A)$  is independent of  $\psi$ , which follows from [Assumption D.1](#). Similarly, in (ii), we use that  $\theta$  is conditionally independent of  $(X, A)$  given  $\psi$ . Now we know that given  $\theta$ ,  $R_i \mid X_i, A_i$  are i.i.d. and hence  $p((R_i)_{i \in [n]} \mid \psi, \theta, (X_i, A_i)_{i \in [n]}) = \prod_{a \in \mathcal{A}} \mathcal{L}_a(\theta_a)$ . Moreover,  $\theta_a$  for  $a \in \mathcal{A}$  are conditionally independent given  $\psi$ . Thus  $p(\theta \mid \psi) = \prod_{a \in \mathcal{A}} p_a(\theta_a; f_a(\psi))$ , where we

also used that  $\theta_a \mid \psi \sim p_a(\cdot; f_a(\psi))$ . This leads to

$$\begin{aligned} p(\psi \mid (X_i, A_i, R_i)_{i \in [n]}) &\propto \int_{\theta} \prod_{a \in \mathcal{A}} \mathcal{L}_a(\theta_a) p_a(\theta_a; f_a(\psi)) d\theta q(\psi), \\ &\stackrel{(i)}{=} \prod_{a \in \mathcal{A}} \int_{\theta_a} \mathcal{L}_a(\theta_a) \mathcal{N}(\theta_a; W_a \psi, \Sigma_a) d\theta_a \mathcal{N}(\psi; \mu, \Sigma), \\ &\stackrel{(ii)}{=} \prod_{a \in \mathcal{A}} \int_{\theta_a} \left( \prod_{i \in I_a} \mathcal{N}(R_i; \phi(X_i)^\top \theta_a, \sigma^2) \right) \mathcal{N}(\theta_a; W_a \psi, \Sigma_a) d\theta_a \mathcal{N}(\psi; \mu, \Sigma). \end{aligned}$$

In (i), we notice that  $\theta = (\theta_a)_{a \in \mathcal{A}}$  and apply Fubini's Theorem. In (ii), we let  $I_a = |S_a| = \sum_{i \in [n]} \mathbb{1}\{A_i = a\}$  as the number of times action  $a$  appears in the sample set  $S$ . Now let  $h_a(\psi) = \int_{\theta_a} \left( \prod_{i \in I_a} \mathcal{N}(R_i; \phi(X_i)^\top \theta_a, \sigma^2) \right) \mathcal{N}(\theta_a; W_a \psi, \Sigma_a) d\theta$ . Then we have that

$$p(\psi \mid S) \propto \prod_{a \in \mathcal{A}} h_a(\psi) \mathcal{N}(\psi; \mu, \Sigma). \quad (22)$$

We start by computing  $h_a$ . To reduce clutter, let  $v = \sigma^{-2}$  and  $\Lambda_a = \Sigma_a^{-1}$ . Then we compute  $h_a$  as

$$\begin{aligned} h_a(\psi) &= \int_{\theta_a} \left( \prod_{i \in I_a} \mathcal{N}(R_i; \phi(X_i)^\top \theta_a, \sigma^2) \right) \mathcal{N}(\theta_a; W_a \psi, \Sigma_a) d\theta_a, \\ &\propto \int_{\theta_a} \exp \left[ -\frac{1}{2} v \sum_{i \in I_a} (R_i - \phi(X_i)^\top \theta_a)^2 - \frac{1}{2} (\theta_a - W_a \psi)^\top \Lambda_a (\theta_a - W_a \psi) \right] d\theta_a, \\ &= \int_{\theta_a} \exp \left[ -\frac{1}{2} \left( v \sum_{i \in I_a} (R_i^2 - 2R_i \theta_a^\top \phi(X_i) + (\theta_a^\top \phi(X_i))^2) + \theta_a^\top \Lambda_a \theta_a - 2\theta_a^\top \Lambda_a W_a \psi \right. \right. \\ &\quad \left. \left. + (W_a \psi)^\top \Lambda_a (W_a \psi) \right) \right] d\theta_a, \\ &\propto \int_{\theta_a} \exp \left[ -\frac{1}{2} \left( \theta_a^\top \left( v \sum_{i \in I_a} \phi(X_i) \phi(X_i)^\top + \Lambda_a \right) \theta_a - 2\theta_a^\top \left( v \sum_{i \in I_a} R_i \phi(X_i) + \Lambda_a W_a \psi \right) \right. \right. \\ &\quad \left. \left. + (W_a \psi)^\top \Lambda_a (W_a \psi) \right) \right] d\theta_a. \end{aligned}$$

Now recall that  $G_a = v \sum_{i \in I_a} \phi(X_i) \phi(X_i)^\top$  and  $B_a = v \sum_{i \in I_a} R_i \phi(X_i)$  and let  $V_a = (G_a + \Lambda_a)^{-1}$ ,  $U_a = V_a^{-1}$ , and  $\beta_a = V_a (B_a + \Lambda_a W_a \psi)$ . Then have that  $U_a V_a = V_a U_a = I_d$ , and thus

$$\begin{aligned} f(\psi) &\propto \int_{\theta_a} \exp \left[ -\frac{1}{2} \left( \theta_a^\top U_a \theta_a - 2\theta_a^\top U_a V_a (B_a + \Lambda_a W_a \psi) + (W_a \psi)^\top \Lambda_a (W_a \psi) \right) \right] d\theta_a, \\ &= \int_{\theta_a} \exp \left[ -\frac{1}{2} \left( \theta_a^\top U_a \theta_a - 2\theta_a^\top U_a \beta_a + (W_a \psi)^\top \Lambda_a (W_a \psi) \right) \right] d\theta_a, \\ &= \int_{\theta_a} \exp \left[ -\frac{1}{2} \left( (\theta_a - \beta_a)^\top U_a (\theta_a - \beta_a) - \beta_a^\top U_a \beta_a + (W_a \psi)^\top \Lambda_a (W_a \psi) \right) \right] d\theta_a, \\ &\propto \exp \left[ -\frac{1}{2} \left( -\beta_a^\top U_a \beta_a + (W_a \psi)^\top \Lambda_a (W_a \psi) \right) \right], \\ &= \exp \left[ -\frac{1}{2} \left( -(B_a + \Lambda_a W_a \psi)^\top V_a (B_a + \Lambda_a W_a \psi) + (W_a \psi)^\top \Lambda_a (W_a \psi) \right) \right], \\ &\propto \exp \left[ -\frac{1}{2} \left( \psi^\top W_a^\top (\Lambda_a - \Lambda_a V_a \Lambda_a) W_a \psi - 2\psi^\top (W_a^\top \Lambda_a V_a B_a) \right) \right], \\ &\propto \mathcal{N}(\psi; \bar{\mu}_a, \bar{\Sigma}_a), \end{aligned}$$

where

$$\begin{aligned} \bar{\Sigma}_a^{-1} &= W_a^\top (\Lambda_a - \Lambda_a V_a \Lambda_a) W_a = W_a^\top (\Sigma_a^{-1} - \Sigma_a^{-1} (G_a + \Sigma_a^{-1})^{-1} \Sigma_a^{-1}) W_a, \\ \bar{\Sigma}_a^{-1} \bar{\mu}_a &= W_a^\top \Lambda_a V_a B_a = W_a^\top \Sigma_a^{-1} (G_a + \Sigma_a^{-1})^{-1} B_a. \end{aligned} \quad (23)$$

However, we know from (22) that  $p(\psi | S) \propto \prod_{a \in \mathcal{A}} h_a(\psi) \mathcal{N}(\psi; \mu, \Sigma)$ . But  $h_a(\psi)$  is proportional to  $\mathcal{N}(\psi; \bar{\mu}_a, \bar{\Sigma}_a)$  for any  $a$ . Thus  $p(\psi | S)$  can be seen as the product of  $K + 1$  multivariate Gaussian distributions  $\mathcal{N}(\mu, \Sigma)$  and  $\mathcal{N}(\bar{\mu}_a, \bar{\Sigma}_a)$  for  $a \in \mathcal{A}$  (the terms  $h_a$ ). Thus,  $p(\psi | S)$  is also a multivariate Gaussian distribution  $\mathcal{N}(\bar{\mu}, \bar{\Sigma}^{-1})$ , with

$$\bar{\Sigma}^{-1} = \Sigma^{-1} + \sum_{a \in \mathcal{A}} \bar{\Sigma}_a^{-1} = \Sigma^{-1} + \sum_{a \in \mathcal{A}} W_a^\top (\Sigma_a^{-1} - \Sigma_a^{-1}(G_a + \Sigma_a^{-1})^{-1} \Sigma_a^{-1}) W_a, \quad (24)$$

$$\bar{\Sigma}^{-1} \bar{\mu} = \Sigma^{-1} \mu + \sum_{a \in \mathcal{A}} \bar{\Sigma}_a^{-1} \bar{\mu}_a = \Sigma^{-1} \mu + \sum_{a \in \mathcal{A}} W_a^\top \Sigma_a^{-1} (G_a + \Sigma_a^{-1})^{-1} B_a. \quad (25)$$

□

## D.2 Conditional posterior

*Derivation of  $p(\theta_a | \psi, S)$ .* Let  $v = \sigma^{-2}$ ,  $\Lambda_a = \Sigma_a^{-1}$ . We consider the model rewritten in (21), then the *joint* action posterior  $p(\theta | S)$  decomposes as

$$\begin{aligned} p(\theta | \psi, S) &= p(\theta | \psi, (X_i, A_i, R_i)_{i \in [n]}) \stackrel{(i)}{\propto} p((R_i)_{i \in [n]} | \theta, \psi, (X_i, A_i)_{i \in [n]}) p(\theta | \psi, (X_i, A_i)_{i \in [n]}), \\ &\stackrel{(ii)}{=} p((R_i)_{i \in [n]} | \theta, (X_i, A_i)_{i \in [n]}) p(\theta | \psi) \stackrel{(iii)}{=} \prod_{i \in [n]} p(R_i | \theta, X_i, A_i) p(\theta | \psi), \\ &\stackrel{(iv)}{=} \prod_{i \in [n]} \mathcal{N}(R_i; (u_{A_i} \otimes \phi(X_i))^\top \theta, \sigma^2) \mathcal{N}(\theta; W_{\mathcal{A}} \psi, \Sigma_{\mathcal{A}}), \\ &= \exp \left[ -\frac{1}{2} \left( v \sum_{i=1}^n (R_i^2 - 2R_i (u_{A_i} \otimes \phi(X_i))^\top \theta + ((u_{A_i} \otimes \phi(X_i))^\top \theta)^2) + \theta^\top \Lambda_{\mathcal{A}} \theta - 2\theta^\top \Lambda_{\mathcal{A}} W_{\mathcal{A}} \psi \right. \right. \\ &\quad \left. \left. + \psi^\top W_{\mathcal{A}}^\top \Lambda_{\mathcal{A}} W_{\mathcal{A}} \psi \right) \right], \\ &\propto \exp \left[ -\frac{1}{2} \left( \theta^\top \left( v \sum_{i=1}^n (u_{A_i} \otimes \phi(X_i)) (u_{A_i} \otimes \phi(X_i))^\top + \Lambda_{\mathcal{A}} \right) \theta - 2\theta^\top \left( v \sum_{i=1}^n (u_{A_i} \otimes \phi(X_i))^\top R_i + \Lambda_{\mathcal{A}} W_{\mathcal{A}} \psi \right) \right) \right], \\ &= \exp \left[ -\frac{1}{2} \left( \theta^\top \left( v \sum_{i=1}^n (u_{A_i} u_{A_i}^\top \otimes \phi(X_i) \phi(X_i)^\top) + \Lambda_{\mathcal{A}} \right) \theta - 2\theta^\top \left( v \sum_{i=1}^n (u_{A_i} \otimes \phi(X_i))^\top R_i + \Lambda_{\mathcal{A}} W_{\mathcal{A}} \psi \right) \right) \right], \\ &\stackrel{(v)}{\propto} \mathcal{N} \left( \theta; \tilde{\mu}_{\mathcal{A}}, \left( \tilde{\Lambda}_{\mathcal{A}} \right)^{-1} \right), \end{aligned}$$

where we use Bayes rule in (i), (ii) uses two assumptions. First, Given  $\theta, X, A, R$  is independent of  $\psi$ . Second, given  $\psi, \theta$  is independent of  $(X, A)$ . Moreover, (iii) follows from the assumption that  $R_i | \theta, X_i, A_i$  are i.i.d. Finally, in iv, we replace the distribution by their Gaussian form, and in (v), we set  $\tilde{\Lambda}_{\mathcal{A}} = v \sum_{i=1}^n (u_{A_i} u_{A_i}^\top \otimes \phi(X_i) \phi(X_i)^\top) + \Lambda_{\mathcal{A}}$ , and  $\tilde{\mu}_{\mathcal{A}} = \tilde{\Lambda}_{\mathcal{A}}^{-1} \left( v \sum_{i=1}^n (u_{A_i} \otimes \phi(X_i))^\top R_i + \Lambda_{\mathcal{A}} W_{\mathcal{A}} \psi \right)$ . Now notice that  $\tilde{\Lambda}_{\mathcal{A}} = \text{diag}(\Sigma_a^{-1} + G_a)_{a \in \mathcal{A}}$ . Thus,  $p(\theta | S) = \mathcal{N}(\theta, \tilde{\mu}_{\mathcal{A}}, \Lambda_{\mathcal{A}}^{-1})$  where  $\tilde{\mu}_{\mathcal{A}} = (\tilde{\mu}_a)_{a \in \mathcal{A}}$  and  $\Lambda_{\mathcal{A}} = \text{diag}(\tilde{\Lambda}_a)_{a \in \mathcal{A}}$ , with

$$\begin{aligned} \tilde{\Lambda}_a &= \Sigma_a^{-1} + G_a, \\ \tilde{\Lambda}_a \tilde{\mu}_a &= \Sigma_a^{-1} W_a \psi + B_a. \end{aligned}$$

The covariance matrix of  $p(\theta | \psi, S)$  is diagonal by block. Thus  $\theta_a | \psi, S$  for  $a \in \mathcal{A}$  are independent and have a Gaussian density  $p(\theta_a | S) = \mathcal{N}(\theta_a; \tilde{\mu}_a, \tilde{\Sigma}_a)$  where  $\tilde{\Sigma}_a = \tilde{\Lambda}_a^{-1}$ . □

## D.3 Action posterior

*Derivation of  $p(\theta_a | S)$ .* We know that  $\theta_a | S, \psi \sim \mathcal{N}(\tilde{\mu}_a, \tilde{\Sigma}_a)$  and  $\psi | S \sim \mathcal{N}(\bar{\mu}, \bar{\Sigma})$ . Thus the posterior density of  $\theta_a | S$  is also Gaussian since Gaussianity is preserved after marginalization [Koller and Friedman, 2009]. We let  $\theta_a | S \sim \mathcal{N}(\hat{\mu}_a, \hat{\Sigma}_a)$ . Then, we can compute  $\hat{\mu}_a$  and  $\hat{\Sigma}_a$  using



the total expectation and total covariance decompositions. Let  $\Lambda_a = \Sigma_a^{-1}$ . Then we have that

$$\begin{aligned}\tilde{\Sigma}_a &= (G_a + \Lambda_a)^{-1} \\ \mathbb{E}[\theta_a | \psi, S] &= \tilde{\Sigma}_a (B_a + \Lambda_a W_a \psi)\end{aligned}$$

First, given  $S$ ,  $\tilde{\Sigma}_a = (G_a + \Lambda_a)^{-1}$  and  $B_a$  are constant (do not depend on  $\psi$ ). Thus

$$\begin{aligned}\hat{\mu}_a &= \mathbb{E}[\theta_a | S] = \mathbb{E}[\mathbb{E}[\theta_a | \psi, S] | S] = \mathbb{E}_{\psi \sim \mathcal{N}(\bar{\mu}, \bar{\Sigma})} [\tilde{\Sigma}_a (B_a + \Lambda_a W_a \psi)] = \tilde{\Sigma}_a (B_a + \Lambda_a W_a \mathbb{E}_{\psi \sim \mathcal{N}(\bar{\mu}, \bar{\Sigma})}[\psi]), \\ &= \tilde{\Sigma}_a (B_a + \Lambda_a W_a \bar{\mu}).\end{aligned}$$

This concludes the computation of  $\hat{\mu}_a$ . Similarly, given  $S$ ,  $\tilde{\Sigma}_a = (G_a + \Lambda_a)^{-1}$  and  $B_a$  are constant (do not depend on  $\psi$ ), yields two things. First,

$$\mathbb{E}[\text{cov}[\theta_a | \psi, S] | S] = \mathbb{E}[\tilde{\Sigma}_a | S] = \tilde{\Sigma}_a.$$

Second,

$$\begin{aligned}\text{cov}[\mathbb{E}[\theta_a | \psi, S] | S] &= \text{cov}[\tilde{\Sigma}_a \Lambda_a W_a \psi | S] \\ &= \tilde{\Sigma}_a \Lambda_a W_a \text{cov}[\psi | S] W_a^\top \Lambda_a \tilde{\Sigma}_a \\ &= \tilde{\Sigma}_a \Lambda_a W_a \bar{\Sigma} W_a^\top \Lambda_a \tilde{\Sigma}_a.\end{aligned}$$

Finally, the total covariance decomposition [Weiss, 2005] yields that

$$\hat{\Sigma}_a = \text{cov}[\theta_a | S] = \mathbb{E}[\text{cov}[\theta_a | \psi, S] | S] + \text{cov}[\mathbb{E}[\theta_a | \psi, S] | S] = \tilde{\Sigma}_a + \tilde{\Sigma}_a \Lambda_a W_a \bar{\Sigma} W_a^\top \Lambda_a \tilde{\Sigma}_a.$$

This concludes the proof. □

## E Missing discussions and proofs

### E.1 Linear hierarchy with non-linear rewards

In this subsection, the action and latent parameters are generated as in (10). But the reward is no longer Gaussian. Precisely, let  $\text{Ber}(p)$  be a Bernoulli distribution with mean  $p$ , and  $g(u) = 1/(1 + \exp(-u))$  be the logistic function. We define  $p(\cdot | x; \theta_a) = \text{Ber}(g(\phi(x)^\top \theta_a))$  which yields

$$\begin{aligned}\psi &\sim \mathcal{N}(\mu, \Sigma), \\ \theta_a | \psi &\sim \mathcal{N}(W_a \psi, \Sigma_a), \quad \forall a \in \mathcal{A}, \\ R | \psi, \theta, X, A &\sim \text{Ber}(g(\phi(X)^\top \theta_A)).\end{aligned}\tag{26}$$

**Step (A).** In this case, we cannot derive closed-form posteriors. But we can approximate the likelihoods  $\mathcal{L}_a(\cdot)$  by multivariate Gaussian distributions such as

$$\mathcal{L}_a(\cdot) \approx \mathcal{N}(\cdot; \check{\mu}_a, (\check{G}_a)^{-1}).\tag{27}$$

where  $\check{\mu}_a$  and  $\check{G}_a$  are the maximum likelihood estimator (MLE) and the Hessian of the negative log-likelihood  $-\log \mathcal{L}_a(\cdot)$ , respectively:

$$\begin{aligned}\check{\mu}_a &= \text{argmax}_{\theta \in \mathbb{R}^d} \log \mathcal{L}_a(\theta), \\ \check{G}_a &= \sum_{i \in [n]} \mathbb{1}\{A_i = a\} \dot{g}(X_i^\top \check{\mu}_a) \phi(X_i) \phi(X_i)^\top.\end{aligned}$$

The latter approximation is slightly different from the popular Laplace approximation<sup>1</sup> and is sometimes referred to as Bernstein-von-Mises approximation, see [Schillings et al., 2020, Remark 3] for a detailed discussion. It is simpler and more computationally friendly than the Laplace approximation

<sup>1</sup>Laplace approximation's mean is given by the maximum a-posteriori estimate (MAP) of the posterior and its covariance is the inverse Hessian of the negative log posterior density.

in our setting. Indeed, interestingly, (27) allows us to use the Gaussian posteriors (11), (12) and (13) in **Step (A)** in Section 4, except that  $G_a$  and  $B_a$  are replaced by  $\check{G}_a$  and  $\check{G}_a\check{\mu}_a$ , respectively (notice that the inverse  $\check{G}_a^{-1}$  may not exist but it is not used in our formulas). There is a clear intuition behind this. First,  $G_a \leftarrow \check{G}_a$  captures the change of curvature due to the nonlinearity of the mean function  $g$ . Moreover,  $B_a \leftarrow \check{G}_a\check{\mu}_a$  follows from the fact that the MLE of the action parameter  $\theta_a$  with linear rewards in (10) is  $G_a^{-1}B_a$ , and it corresponds to  $\check{\mu}_a$  in this case.

**Step (B).** The reward function under (26) is  $r(x, a; \theta) = g(\phi(x)^\top \theta)$ . Thus our reward estimate is

$$\hat{r}(x, a) = g\left(\frac{\phi(x)^\top \check{\mu}_a}{\sqrt{1+\pi/8\|\phi(x)\|_{\check{\Sigma}_a}}}\right),$$

where we use an approximation of the expectation of the sigmoid function under a Gaussian distribution [Spiegelhalter and Lauritzen, 1990, Appendix A].

**Step (C).** Since  $g$  is increasing, the learned policy is

$$\hat{\pi}_G(a | x) = \mathbb{1}\left\{a = \operatorname{argmax}_{b \in \mathcal{A}} \frac{\phi(x)^\top \check{\mu}_b}{\sqrt{1+\pi/8\|\phi(x)\|_{\check{\Sigma}_b}}}\right\}.$$

## E.2 Frequentist vs. Bayesian suboptimality

Before diving into our main result, a Bayesian suboptimality bound for the structured prior in (3), we motivate its use by analyzing both Bayesian and frequentist suboptimality for the simpler non-structured prior in (2). While this analysis uses the non-structured prior, our main result applies directly to the structured one. The key distinction between Bayesian and frequentist suboptimality lies in their confidence intervals, as discussed in [Hong et al., 2023, Section 4.2]. However, we still provide a self-contained explanation for completeness, considering our slightly different results and proofs. Also, note that Hong et al. [2023] also explored a notion of Bayesian suboptimality that they defined as  $V(\pi_*; \theta_*) - V(\hat{\pi}; \theta_*) | S$ . Precisely, they fix the sample set  $S$  and consider only randomness in  $\theta_*$  sampled from the prior. In contrast, our definition aligns with traditional online bandits (Bayesian regret) by averaging suboptimality across both the sample set and  $\theta_*$  sampled from the prior:  $\mathbb{E}[V(\pi_*; \theta_*) - V(\hat{\pi}; \theta_*)]$ . These metrics are different and have distinct implications. Interestingly, ours favors the greedy policy (Appendix E.5), while theirs incentivizes pessimistic policies.

**Frequentist confidence intervals.** Assume that for any  $a \in \mathcal{A}$ , there exists a parameter  $\theta_{*,a} \in \mathbb{R}^d$  generating the reward as  $R | \theta_*, X, A \sim \mathcal{N}(\phi(X)^\top \theta_{*,A}, \sigma^2)$ . Then the posterior under the non-structured prior in (2) satisfies

$$\mathbb{P}\left(\forall x \in \mathcal{X} : |r(x, a; \theta_*) - \hat{r}(x, a)| \leq \left(\alpha(d, \delta) + \frac{c_a + \|\mu_a\|_2}{\sqrt{\lambda_d(\Sigma_a)}}\right) \|\phi(X)\|_{\check{\Sigma}_a} | \tilde{S}\right) \geq 1 - \delta, \quad (28)$$

where  $\tilde{S} = (X_i, A_i)_{i \in [n]}$ , the randomness in  $\mathbb{P}$  is over the true reward noise  $\mathcal{N}(0, \sigma^2)$ , and  $\alpha(d, \delta) = \sqrt{d + 2\sqrt{d \log \frac{1}{\delta}}} + 2 \log \frac{1}{\delta}$ . Now let us compare this confidence interval to the Bayesian one below.

**Bayesian confidence intervals.** Here we further assume that for any  $a \in \mathcal{A}$ , the action parameter  $\theta_{*,a}$  is generated from a known Gaussian distribution as  $\theta_{*,a} \sim \mathcal{N}(\mu_a, \Sigma_a)$  whose parameters  $\mu_a, \Sigma_a$  match the ones used in the non-structured prior (2). Then the posterior under the non-structured prior in (2) satisfies

$$\mathbb{P}\left(\forall x \in \mathcal{X} : |r(x, a; \theta_*) - \hat{r}(x, a)| \leq \alpha(d, \delta) \|\phi(X)\|_{\check{\Sigma}_a} | S\right) \geq 1 - \delta \quad (29)$$

Here the randomness in  $\mathbb{P}$  is over the action parameter uncertainty  $\mathcal{N}(0, \Sigma_a)$  as we condition on  $S = (X_i, A_i, R_i)_{i \in [n]}$ . The proof is given after Lemma E.1. In Bayesian bounds,  $\theta_{*,a}$  is treated as a random variable with distribution  $\mathcal{N}(\mu_a, \Sigma_a)$  and it is included in the randomness in  $\mathbb{P}$ . The randomization of the data  $(X_i, A_i, R_i)_{i \in [n]}$  is also included in  $\mathbb{P}$ . However, we condition on  $S$ , and hence they are treated as fixed. In contrast with frequentist bounds where  $\theta_{*,a}$  is fixed while the reward is random and included in the randomness in  $\mathbb{P}$  without conditioning on it. The Bayesian confidence intervals capture that the agent has access to the prior on the true action parameters and

hence the additional misspecification term  $\frac{c_a + \|\mu_a\|_2}{\sqrt{\lambda_d(\Sigma_a)}}$  is not present, as opposed to its frequentist counterpart. This is the main reason why BSO captures the benefits of knowing and leveraging informative priors.

**Lemma E.1** (Bayesian bound). *Let  $a \in \mathcal{A}$ ,  $\delta \in (0, 1)$  and  $\alpha(d, \delta) = \sqrt{d + 2\sqrt{d \log \frac{1}{\delta}} + 2 \log \frac{1}{\delta}}$ .*

*Then*

$$\mathbb{P} \left( \forall x \in \mathcal{X} : |r(x, a; \theta_*) - r(x, a; \hat{\mu})| \leq \alpha(d, \delta) \|\phi(X)\|_{\hat{\Sigma}_a} \mid S \right) \geq 1 - \delta \quad (30)$$

*Proof.* First, we have that

$$\begin{aligned} |r(x, a; \theta_*) - r(x, a; \hat{\mu})| &= |\phi(X)^\top \theta_{*,a} - \phi(X)^\top \hat{\mu}_a| = |\phi(X)^\top (\theta_{*,a} - \hat{\mu}_a)|, \\ &= |\phi(X)^\top \hat{\Sigma}_a^{\frac{1}{2}} \hat{\Sigma}_a^{-\frac{1}{2}} (\theta_{*,a} - \hat{\mu}_a)| \leq \|\phi(X)\|_{\hat{\Sigma}_a} \|\theta_{*,a} - \hat{\mu}_a\|_{\hat{\Sigma}_a^{-1}}, \end{aligned}$$

where we use the Cauchy-Schwarz inequality in the last step. Now we assumed that the true action parameters  $\theta_{*,a}$  are random and their prior distribution matches our model in (10). Therefore,  $\theta_{*,a} \mid S$  have the same density as our posterior  $\theta_a \mid S$ , and hence  $\theta_{*,a} \mid S \sim \mathcal{N}(\hat{\mu}_a, \hat{\Sigma}_a)$ . This means that  $\theta_{*,a} - \hat{\mu}_a \mid S \sim \mathcal{N}(0, \hat{\Sigma}_a)$ . Thus,  $\hat{\Sigma}_a^{-\frac{1}{2}}(\theta_{*,a} - \hat{\mu}_a) \sim \mathcal{N}(0, I_d)$ . But notice that  $\|\theta_{*,a} - \hat{\mu}_a\|_{\hat{\Sigma}_a^{-1}} = \|\hat{\Sigma}_a^{-\frac{1}{2}}(\theta_{*,a} - \hat{\mu}_a)\|$ . Thus we apply [Laurent and Massart \[2000, Lemma 1\]](#) and get that

$$\mathbb{P} \left( \|\theta_{*,a} - \hat{\mu}_a\|_{\hat{\Sigma}_a^{-1}} \leq \alpha(d, \delta) \mid S \right) \geq 1 - \delta.$$

Finally using that for any  $x \in \mathcal{A}$ ,  $|r(x, a; \theta_*) - r(x, a; \hat{\mu})| \leq \|\phi(X)\|_{\hat{\Sigma}_a} \|\theta_{*,a} - \hat{\mu}_a\|_{\hat{\Sigma}_a^{-1}}$  concludes the proof.  $\square$

**Lemma E.2** (Frequentist bound). *Let  $a \in \mathcal{A}$ ,  $\delta \in (0, 1)$  and  $\alpha(d, \delta) = \sqrt{d + 2\sqrt{d \log \frac{1}{\delta}} + 2 \log \frac{1}{\delta}}$ .*

*Then*

$$\mathbb{P} \left( \forall x \in \mathcal{X} : |r(x, a; \theta_*) - r(x, a; \hat{\mu})| \leq \left( \alpha(d, \delta) + \frac{c_a + \|\hat{\mu}_a\|_2}{\sqrt{\lambda_d(\Sigma_a)}} \right) \|\phi(X)\|_{\hat{\Sigma}_a} \mid \tilde{S} \right) \geq 1 - \delta, \quad (31)$$

where  $\tilde{S} = (X_i, A_i)_{i \in [n]}$ .

*Proof.* Similarly to [Lemma E.1](#), we have that  $|r(x, a; \theta_*) - r(x, a; \hat{\mu})| \leq \|\phi(X)\|_{\hat{\Sigma}_a} \|\theta_{*,a} - \hat{\mu}_a\|_{\hat{\Sigma}_a^{-1}}$ . Note that  $(X_i, A_i)$  and  $\theta_{*,a}$  are fixed; the randomness only comes from  $R_i$  for  $i \in [n]$ . Thus  $\hat{\Sigma}_a$  is fixed while  $\hat{\mu}_a$  is random. Keeping this in mind, we have that  $R_i \sim \mathcal{N}(\phi(X_i)^\top \theta_{*,a}, \sigma^2)$  and thus

$$\sum_{i \in [n]} \mathbb{I}_{\{A_i=a\}} \phi(X_i) R_i \sim \mathcal{N} \left( \sum_{i \in [n]} \mathbb{I}_{\{A_i=a\}} \phi(X_i) \phi(X_i)^\top \theta_{*,a}, \sigma^2 \sum_{i \in [n]} \mathbb{I}_{\{A_i=a\}} \phi(X_i) \phi(X_i)^\top \right). \quad (32)$$

But  $\sum_{i \in [n]} \mathbb{I}_{\{A_i=a\}} \phi(X_i) R_i = \sigma^2 B_a$  and  $\sum_{i \in [n]} \mathbb{I}_{\{A_i=a\}} \phi(X_i) \phi(X_i)^\top = \sigma^2 G_a$ . Thus, we get that

$$\sigma^2 B_a \sim \mathcal{N}(\sigma^2 G_a \theta_{*,a}, \sigma^4 G_a), \quad (33)$$

and hence

$$B_a \sim \mathcal{N}(G_a \theta_{*,a}, G_a). \quad (34)$$

Now since  $\Sigma_a, \mu_a$  and  $\hat{\Sigma}_a$  are fixed, we have that

$$\hat{\mu}_a = \hat{\Sigma}_a (\Sigma_a^{-1} \mu_a + B_a) \sim \mathcal{N}(\hat{\Sigma}_a (\Sigma_a^{-1} \mu_a + G_a \theta_{*,a}), \hat{\Sigma}_a G_a \hat{\Sigma}_a). \quad (35)$$

But notice that  $G_a \theta_{*,a} = (G_a + \Sigma_a^{-1}) \theta_{*,a} - \Sigma_a^{-1} \theta_{*,a} = \hat{\Sigma}_a^{-1} \theta_{*,a} - \Sigma_a^{-1} \theta_{*,a}$  and hence (35) becomes

$$\hat{\mu}_a \sim \mathcal{N}(\hat{\Sigma}_a \Sigma_a^{-1} (\mu_a - \theta_{*,a}) + \theta_{*,a}, \hat{\Sigma}_a G_a \hat{\Sigma}_a). \quad (36)$$

Thus

$$\hat{\mu}_a - \theta_{*,a} \sim \mathcal{N}(\hat{\Sigma}_a \Sigma_a^{-1} (\mu_a - \theta_{*,a}), \hat{\Sigma}_a G_a \hat{\Sigma}_a). \quad (37)$$

Multiplying by  $\hat{\Sigma}_a^{-\frac{1}{2}}$  leads to

$$\hat{\Sigma}_a^{-\frac{1}{2}} (\hat{\mu}_a - \theta_{*,a}) \sim \mathcal{N}(\hat{\Sigma}_a^{\frac{1}{2}} \Sigma_a^{-1} (\mu_a - \theta_{*,a}), \hat{\Sigma}_a^{\frac{1}{2}} G_a \hat{\Sigma}_a^{\frac{1}{2}}). \quad (38)$$

Now we rewrite (38) as

$$\hat{\Sigma}_a^{-\frac{1}{2}} (\hat{\mu}_a - \theta_{*,a}) = \hat{\Sigma}_a^{\frac{1}{2}} \Sigma_a^{-1} (\mu_a - \theta_{*,a}) + \hat{\Sigma}_a^{\frac{1}{2}} G_a^{\frac{1}{2}} Z, \quad Z \sim \mathcal{N}(0, I_d). \quad (39)$$

Thus we have that

$$\begin{aligned} \|\hat{\Sigma}_a^{-\frac{1}{2}} (\hat{\mu}_a - \theta_{*,a})\|_2 &= \|\hat{\mu}_a - \theta_{*,a}\|_{\hat{\Sigma}_a^{-1}} = \|\hat{\Sigma}_a^{\frac{1}{2}} \Sigma_a^{-1} (\mu_a - \theta_{*,a}) + \hat{\Sigma}_a^{\frac{1}{2}} G_a^{\frac{1}{2}} Z\|_2, \\ &\leq \|\hat{\Sigma}_a^{\frac{1}{2}} \Sigma_a^{-1} (\mu_a - \theta_{*,a})\|_2 + \|\hat{\Sigma}_a^{\frac{1}{2}} G_a^{\frac{1}{2}} Z\|_2, \\ &\leq \frac{\|\mu_a - \theta_{*,a}\|_2}{\sqrt{\lambda_d(\Sigma_a)}} + \|Z\|_2. \end{aligned} \quad (40)$$

Finally, since  $Z \sim \mathcal{N}(0, I_d)$ , we apply [Laurent and Massart \[2000, Lemma 1\]](#) and get that

$$\mathbb{P}(\|Z\|_2 \leq \alpha(d, \delta)) \geq 1 - \delta.$$

Combining this with (40) leads to

$$\mathbb{P}\left(\|\theta_{*,a} - \hat{\mu}_a\|_{\hat{\Sigma}_a^{-1}} \leq \frac{\|\mu_a - \theta_{*,a}\|_2}{\sqrt{\lambda_d(\Sigma_a)}} + \alpha(d, \delta) \mid (X_i, A_i)_{i \in [n]}\right) \geq 1 - \delta.$$

Finally using that for any  $x \in \mathcal{A}$ ,  $|r(x, a; \theta_*) - r(x, a; \hat{\mu})| \leq \|\phi(X)\|_{\hat{\Sigma}_a} \|\theta_a - \hat{\mu}_a\|_{\hat{\Sigma}_a^{-1}}$  and using that  $\|\theta_{*,a} - \hat{\mu}_a\|_2 \leq c_a + \|\hat{\mu}_a\|_2$  concludes the proof.  $\square$

### E.3 Main result

In this section, we prove [Theorem 5.2](#). Recall that we make the following well-specified prior assumption.

**Assumption E.3** (Well-specified priors). Action parameters  $\theta_{*,a}$  and rewards are drawn from (10).

First, given  $x \in \mathcal{X}$ , by definition of the optimal policy, we know that it is deterministic. That is, there exists  $a_{x, \theta_*} \in [K]$  such that  $\pi_*(a_{x, \theta_*} \mid x) = 1$ . To simplify the notation and since  $\pi_*$  is deterministic, we let  $\pi_*(x) = a_{x, \theta_*}$ . Also, we know that the greedy policy is deterministic in  $\hat{a}_x = \operatorname{argmax}_{b \in \mathcal{A}} \hat{r}(x, b)$ . That is  $\hat{\pi}_G(\hat{a}_x \mid x) = 1$ . Similarly, we let  $\hat{\pi}_G(x) = \hat{a}_x$ . Moreover, we let  $\Phi(x, a) = e_a \otimes \phi(X) \in \mathbb{R}^{dK}$  where  $e_a \in \mathbb{R}^K$  is the indicator vector of action  $a$ , such that  $e_{a,b} = 0$  for any  $b \in \mathcal{A} \setminus \{a\}$  and  $e_{a,a} = 1$ . Also, recall that  $\hat{\mu} = (\hat{\mu}_a)_{a \in \mathcal{A}}$  is the concatenation of the posterior means.

$$\begin{aligned} \text{BSO}(\hat{\pi}_G) &= \mathbb{E}[V(\pi_*; \theta_*) - V(\hat{\pi}; \theta_*)], \\ &= \mathbb{E}[r(X, \pi_*(X); \theta_*) - r(X, \hat{\pi}_G(X); \theta_*)], \\ &= \mathbb{E}[r(X, \pi_*(X); \theta_*) - r(X, \hat{\pi}_G(X); \hat{\mu}) + r(X, \hat{\pi}_G(X); \hat{\mu}) - r(X, \hat{\pi}_G(X); \theta_*)], \\ &\leq \mathbb{E}[r(X, \pi_*(X); \theta_*) - r(X, \pi_*(X); \hat{\mu}) + r(X, \hat{\pi}_G(X); \hat{\mu}) - r(X, \hat{\pi}_G(X); \theta_*)], \\ &\leq \mathbb{E}[r(X, \pi_*(X); \theta_*) - r(X, \pi_*(X); \hat{\mu})] + \mathbb{E}[r(X, \hat{\pi}_G(X); \hat{\mu}) - r(X, \hat{\pi}_G(X); \theta_*)]. \end{aligned}$$

Now we start by proving that  $\mathbb{E} [r(X, \hat{\pi}_G(X); \hat{\mu}) - r(X, \hat{\pi}_G(X); \theta_*)] = 0$ . This is achieved as follows

$$\begin{aligned}
\mathbb{E} [r(X, \hat{\pi}_G(X); \hat{\mu}) - r(X, \hat{\pi}_G(X); \theta_*)] &= \mathbb{E} [\mathbb{E} [r(X, \hat{\pi}_G(X); \hat{\mu}) - r(X, \hat{\pi}_G(X); \theta_*) \mid X, S]] , \\
&= \mathbb{E} [\mathbb{E} [\phi(X)^\top \hat{\mu}_{\hat{\pi}_G(X)} - \phi(X)^\top \theta_{*, \hat{\pi}_G(X)} \mid X, S]] , \\
&\stackrel{(i)}{=} \mathbb{E} [\mathbb{E} [\Phi(X, \hat{\pi}_G(X))^\top \hat{\mu} - \Phi(X, \hat{\pi}_G(X))^\top \theta_* \mid X, S]] , \\
&\stackrel{(ii)}{=} \mathbb{E} [\Phi(X, \hat{\pi}_G(X))^\top \mathbb{E} [\hat{\mu} - \theta_* \mid X, S]] , \\
&\stackrel{(iii)}{=} \mathbb{E} [\Phi(X, \hat{\pi}_G(X))^\top (\hat{\mu} - \mathbb{E} [\theta_* \mid X, S])] , \\
&\stackrel{(iv)}{=} 0 .
\end{aligned}$$

In (i), we used that by definition of  $\Phi(x, a) = e_a \otimes \phi(X) \in \mathbb{R}^{dK}$ , we have  $\phi(x)^\top \hat{\mu}_a = \Phi(x, a)^\top \hat{\mu}$  for any  $(x, a)$ , and the same holds for  $\theta_*$ . In (ii), we used that  $\Phi(X, \hat{\pi}_G(X))$  is deterministic given  $X$  and  $S$ . In (iii), we used that  $\hat{\mu}$  is deterministic given  $X$  and  $S$ . Finally, in (iv), we used that  $\mathbb{E} [\theta_* \mid X, S] = \mathbb{E} [\theta_* \mid S] = \hat{\mu}$ , which follows from the assumption that  $\theta_*$  does not depend on  $X$  and the assumption that  $\theta_*$  is drawn from the prior, and hence when conditioned on  $S$ , it is drawn from the posterior whose mean is  $\hat{\mu}$ . Therefore,  $\mathbb{E} [r(X, \hat{\pi}_G(X); \hat{\mu}) - r(X, \hat{\pi}_G(X); \theta_*)] = 0$  which leads to

$$\text{BSO}(\hat{\pi}_G) \leq \mathbb{E} [r(X, \pi_*(X); \theta_*) - r(X, \pi_*(X); \hat{\mu})] .$$

Now let  $\delta \in (0, 1)$ , we define the following high-probability events

$$E_a = \left\{ \forall x \in \mathcal{X} : |r(x, a; \theta_*) - r(x, a; \hat{\mu})| \leq \alpha(d, \delta) \|\phi(x)\|_{\hat{\Sigma}_a} \right\}, \quad \forall a \in \mathcal{A},$$

where  $\alpha(d, \delta) = \sqrt{d + 2\sqrt{d \log \frac{1}{\delta}} + 2 \log \frac{1}{\delta}}$ . Then we decompose  $\text{BSO}(\hat{\pi}_G)$  as

$$\begin{aligned}
\text{BSO}(\hat{\pi}_G) &\leq \mathbb{E} [r(X, \pi_*(X); \theta_*) - r(X, \pi_*(X); \hat{\mu})] , \\
&\leq \mathbb{E} [|r(X, \pi_*(X); \theta_*) - r(X, \pi_*(X); \hat{\mu})|] , \\
&\leq \mathbb{E} [|r(X, \pi_*(X); \theta_*) - r(X, \pi_*(X); \hat{\mu})| \mathbb{1}\{E_{\pi_*(X)}\}] + \mathbb{E} [|r(X, \pi_*(X); \theta_*) - r(X, \pi_*(X); \hat{\mu})| \mathbb{1}\{\bar{E}_{\pi_*(X)}\}] , \\
&\leq \alpha(d, \delta) \mathbb{E} [\|\phi(X)\|_{\hat{\Sigma}_{\pi_*(X)}}] + \mathbb{E} [|r(X, \pi_*(X); \theta_*) - r(X, \pi_*(X); \hat{\mu})| \mathbb{1}\{\bar{E}_{\pi_*(X)}\}] .
\end{aligned}$$

Now we deal with the term  $\mathbb{E} [|r(X, \pi_*(X); \theta_*) - r(X, \pi_*(X); \hat{\mu})| \mathbb{1}\{\bar{E}_{\pi_*(X)}\}]$ . Let  $Z_a = r(X, a; \theta_*) - r(X, a; \hat{\mu})$ , so that  $Z_{\pi_*(X)} = r(X, \pi_*(X); \theta_*) - r(X, \pi_*(X); \hat{\mu})$ . Then we have that

$$\begin{aligned}
\mathbb{E} [|r(X, \pi_*(X); \theta_*) - r(X, \pi_*(X); \hat{\mu})| \mathbb{1}\{\bar{E}_{\pi_*(X)}\}] &= \mathbb{E} [|Z_{\pi_*(X)}| \mathbb{1}\{|Z_{\pi_*(X)}| > \alpha(d, \delta) \|\phi(X)\|_{\hat{\Sigma}_{\pi_*(X)}}\}] , \\
&= \mathbb{E} \left[ \mathbb{E} [|Z_{\pi_*(X)}| \mathbb{1}\{|Z_{\pi_*(X)}| > \alpha(d, \delta) \|\phi(X)\|_{\hat{\Sigma}_{\pi_*(X)}}\} \mid X, S] \right] , \\
&\stackrel{(i)}{\leq} \mathbb{E} \left[ \frac{2}{\|\phi(X)\|_{\hat{\Sigma}_{\pi_*(X)}} \sqrt{2\pi}} \int_{u=\alpha(d, \delta) \|\phi(X)\|_{\hat{\Sigma}_{\pi_*(X)}}}^{\infty} u \exp \left[ -\frac{u^2}{2\|\phi(X)\|_{\hat{\Sigma}_{\pi_*(X)}}^2} \right] du \right] , \\
&\stackrel{(ii)}{\leq} \mathbb{E} \left[ \|\phi(X)\|_{\hat{\Sigma}_{\pi_*(X)}} \frac{2}{\sqrt{2\pi}} \int_{u=\alpha(d, \delta)}^{\infty} u \exp \left[ -\frac{u^2}{2} \right] du \right] , \\
&\stackrel{(iii)}{\leq} \sqrt{\frac{2}{\pi}} \exp \left( -\frac{\alpha(d, \delta)^2}{2} \right) \mathbb{E} [\|\phi(X)\|_{\hat{\Sigma}_{\pi_*(X)}}] ,
\end{aligned}$$

where (i) follows from the facts that  $Z_{\pi_*(X)} \mid X, S \sim \mathcal{N}(0, \|\phi(X)\|_{\hat{\Sigma}_{\pi_*(X)}})$ . In (ii), we use the change of variables  $u \leftarrow u / \|\phi(X)\|_{\hat{\Sigma}_{\pi_*(X)}}$ . Finally, in (iii), we compute the integral. This leads to the following result

$$\text{BSO}(\hat{\pi}_G) \leq \left( \sqrt{\frac{2}{\pi}} \exp \left( -\frac{\alpha(d, \delta)^2}{2} \right) + \alpha(d, \delta) \right) \mathbb{E} [\|\phi(X)\|_{\hat{\Sigma}_{\pi_*(X)}}] . \quad (41)$$

(42)

This concludes the proof.

#### E.4 Proof of Theorem 5.1

*Proof.* The proof is similar to that of Theorem 5.2. Precisely, we have that

$$\begin{aligned} \text{BMSE}(\hat{r}(x, a)) &= \mathbb{E}[(\hat{r}(x, a) - r(x, a; \theta_*))^2], \\ &= \mathbb{E}[\mathbb{E}[(\hat{r}(x, a) - r(x, a; \theta_*))^2 | S]], \\ &= \mathbb{E}[\mathbb{E}[(\phi(x)^\top \hat{\mu}_a - \phi(x)^\top \theta_{*,a})^2 | S]], \\ &= \mathbb{E}[\mathbb{E}[(\phi(x)^\top \theta_{*,a} - \phi(x)^\top \hat{\mu}_a)^2 | S]]. \end{aligned}$$

But we assumed that the true action parameters  $\theta_{*,a}$  are random and their prior distribution matches our model in (10). Therefore,  $\theta_{*,a} | S$  have the same density as our posterior  $\theta_a | S$ , and hence  $\theta_{*,a} | S \sim \mathcal{N}(\hat{\mu}_a, \hat{\Sigma}_a)$ . This means that  $\phi(x)^\top \theta_{*,a} | S \sim \mathcal{N}(\phi(x)^\top \hat{\mu}_a, \|\phi(x)\|_{\hat{\Sigma}_a}^2)$ . But  $\mathbb{E}[(\phi(x)^\top \theta_{*,a} - \phi(x)^\top \hat{\mu}_a)^2 | S]$  is exactly the variance of  $\phi(x)^\top \theta_{*,a} | S$  and thus it is equal to  $\|\phi(x)\|_{\hat{\Sigma}_a}^2$ . Taking the expectation concludes the proof.  $\square$

#### E.5 Optimality of greedy policies

Here, we show that Greedy policy  $\hat{\pi}_G$  should be preferred to any other choice of policies when considering the BSO as our performance metric. This is because  $\hat{\pi}_G$  minimizes the BSO. To see this, note that by definition the Greedy policy  $\hat{\pi}_G$  is deterministic, that is for any context  $x \in \mathcal{X}$ , there exists  $\hat{a}_G$ , such that  $\hat{\pi}_G(\hat{a}_G | x) = 1$ . Thus, for any context  $x \in \mathcal{X}$ , we simplify the notation by letting  $\hat{\pi}_G(x)$  denote the action that has a mass equal to 1. Then, we have that

$$\mathbb{E}_{A \sim \hat{\pi}_G(\cdot | x)} [\mathbb{E}_{\theta_*} [r(x, A; \theta_*) | S]] = \mathbb{E}_{\theta_*} [r(x, \hat{\pi}_G(x); \theta_*) | S] \geq \mathbb{E} [r(x, a; \theta_*) | S] \quad \forall x, a \in \mathcal{X} \times \mathcal{A}. \quad (43)$$

where this follows from the definition of  $\hat{r}(x, a) = \mathbb{E} [r(x, a; \theta) | S]$ , the definition of  $\hat{\pi}_G$  and the fact that  $\theta_*$  is sampled from the prior, which leads to  $\mathbb{E} [r(x, a; \theta) | S] = \mathbb{E} [r(x, a; \theta_*) | S]$ . Now (43) holds for any  $x \in \mathcal{X}$  and  $a \in \mathcal{A}$ , and hence it holds in expectation under  $X \sim \nu$  and  $A \sim \pi(\cdot | X)$  for any policy  $\pi$ . That is,

$$\mathbb{E}_{X \sim \nu, A \sim \hat{\pi}_G(\cdot | X)} [\mathbb{E}_{\theta_*} [r(x, A; \theta_*) | S]] \geq \mathbb{E}_{X \sim \nu, A \sim \pi(\cdot | X)} [\mathbb{E}_{\theta_*} [r(x, A; \theta_*) | S]]. \quad (44)$$

Taking another expectation w.r.t. the sample set  $S$  and using Fubini's theorem and the tower rule leads to  $\mathbb{E} [V(\hat{\pi}_G; \theta_*)] \geq \mathbb{E} [V(\pi; \theta_*)]$  for any stationary policy  $\pi$ . Then, subtracting  $\mathbb{E} [V(\pi_*; \theta_*)]$  from both sides of the previous inequality yields that the BSO is minimized by  $\hat{\pi}_G$  compared to any stationary policy  $\pi$ , in particular, compared to the policy  $\pi_p$  induced by pessimism.

#### E.6 Discussing the main assumption

In this section, we discuss our main assumption of the existence of the latent parameter  $\psi$ . Let us take the Gaussian case, where we assume that

$$\begin{aligned} \psi &\sim \mathcal{N}(\mu, \Sigma), \\ \theta_a | \psi &\sim \mathcal{N}(W_a \psi, \Sigma_a), \quad \forall a \in \mathcal{A}. \end{aligned} \quad (45)$$

Now we discuss if (45) is as mild as assuming that the action parameters are jointly sampled as  $\theta \sim \mathcal{N}(\mu_{\text{ACT}}, \Sigma_{\text{ACT}})$ , where  $\mu_{\text{ACT}} \in \mathbb{R}^{dK}$  and  $\Sigma_{\text{ACT}} \in \mathbb{R}^{dK \times dK}$ . To do so, assume that  $\theta \sim \mathcal{N}(0, \Sigma_{\text{ACT}})$ , where we set  $\mu_{\text{ACT}} = 0$  to simplify notation. Then we want to prove that there exists  $\psi \sim \mathcal{N}(0, I_{d'})$ , with  $\mu \in \mathbb{R}^{d'}$  for some  $d' \geq 1$  and  $W \in \mathbb{R}^{dK \times d'}$  such that  $\theta | \psi \sim \mathcal{N}(W\psi, \hat{\Sigma})$ , where  $\hat{\Sigma}$  is diagonal by  $d \times d$  blocks. First, note that  $(\theta, \psi)$  is also Gaussian with a mean equal to 0 and covariance  $\Omega = \begin{pmatrix} \Sigma_{\text{ACT}} & \Omega_2 \\ \Omega_2^\top & I_{d'} \end{pmatrix} \in \mathbb{R}^{(dK+d') \times (dK+d')}$ . Then it suffices to find  $W$  such that  $\theta - W\psi \perp\!\!\!\perp \psi$ . That is,  $W$  such that  $\text{cov}(\theta - W\psi, \psi) = \text{cov}(\theta, \psi) - W \text{cov}(\psi, \psi) = \Omega_2 - W = 0$ , which leads to  $W = \Omega_2$ . Using this conditional independence and that  $W = \Omega_2$ , we get that the distribution of  $\theta - W\psi | \psi$  is the same as the distribution of  $\theta - W\psi$  and it writes  $\mathcal{N}(0, \Sigma_{\text{ACT}} - WW^\top)$ . Thus,



$\theta \mid \psi \sim \mathcal{N}(W\psi, \Sigma_{\text{ACT}} - WW^\top)$ . But  $\Sigma_{\text{ACT}}$  is positive-definite and hence there exists  $\lambda > 0$  such that  $\Sigma_{\text{ACT}} - \lambda I_{dK}$  is symmetric and positive semi-definite. Therefore,  $\Sigma_{\text{ACT}} - \lambda I_{dK}$  can be written as  $\Sigma_{\text{ACT}} - \lambda I_{dK} = BB^\top$  where  $B \in \mathbb{R}^{dK \times d'}$  is such that  $d' \geq \text{rank}(\Sigma_{\text{ACT}} - \lambda I_{dK})$  and setting  $W = B$  concludes the proof.

## F Additional experiments

As mentioned in [Appendix F](#), our experiments were conducted on internal machines with 30 CPUs and thus they required a moderate amount of computation. These experiments are also reproducible with minimal computational resources.

### F.1 OPE experiments

We consider the same synthetic setting as in [Section 6](#), but now we fix a target policy  $\pi$  and assess algorithm performance in evaluating the value function of that target policy using the mean squared error (MSE) averaged over 50 problem instances sampled from the prior. This can be seen as a proxy for our theoretical metric, Bayesian MSE (BMSE). The target policy is defined as  $\epsilon$ -greedy with  $\epsilon = 0.5$ ; it chooses the best action with probability 0.5 and a random action with probability 0.5.

Results are shown in [Fig. 6](#). While sDM maintains its advantage over baselines, particularly in low-data settings, an interesting observation emerged. IPS-based methods (snIPS and DR) outperformed standard direct methods (DM (Bayes) and DM (Freq)) in the OPE experiments, which was not the case in OPL experiments. The only direct method that outperformed them in OPE is sDM.

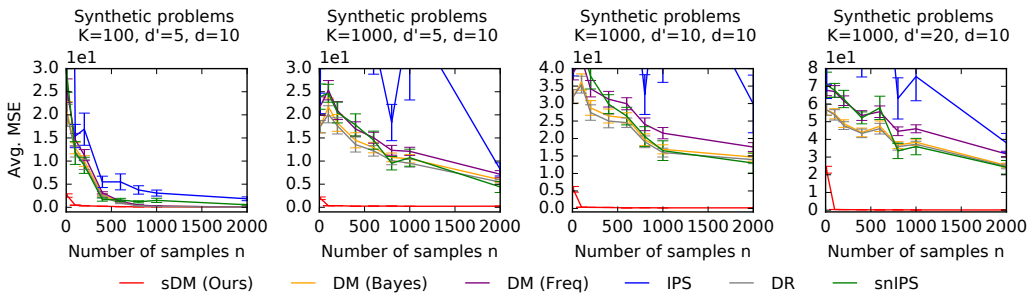


Figure 6: The average MSE of an  $\epsilon$ -greedy target policy on synthetic problems with varying  $n$ ,  $K$  and  $d'$ .

### F.2 Different logging policy: $\epsilon$ -greedy

Here we consider the setting in [Section 6](#) and provide additional results with the  $\epsilon$ -greedy logging policy, where  $\epsilon = 0.5$ . We consider both synthetic and MovieLens datasets. The results are shown in [Fig. 7](#). The conclusions are similar to those in [Section 6](#), except that IPS outperforms the other variants DR and snIPS when the logging policy is performing well. Also, the results for the MovieLens problems are given in [Fig. 8](#) and the conclusions are similar to those made in [Section 6](#).




Figure 7: The relative reward of the learned policy on **synthetic problems** with varying  $n$ ,  $K$  and  $d'$  and varying logging policies.

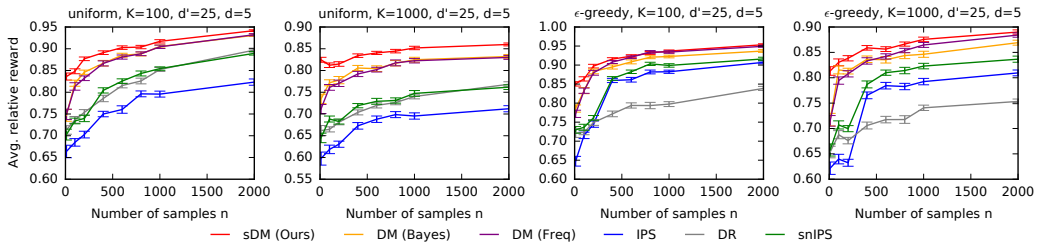


Figure 8: The relative reward of the learned policy on **MovieLens problems** with varying  $n$  and  $K$  and varying logging policies.

### F.3 Robustness to misspecification

We strengthened our evaluation by assessing sDM’s robustness to prior/model misspecification. In these experiments, the true data-generating process (same as the synthetic experiments in Section 6) differed from sDM’s assumptions in four different ways:

**Misspecified prior means (Fig. 9).** This was achieved by adding uniformly sampled noise from  $[v, v+0.5]$  to the true prior mean parameters  $\mu$  and  $W_a$ , with  $v$  controlling the level of misspecification. We varied  $v \in \{0.5, 1, 1.5\}$  and analyzed its impact on sDM’s performance.

**Misspecified prior covariances (Fig. 10).** This is achieved by adding uniformly sampled noise from  $[v, v+0.5]$  to the true prior covariance parameters  $\Sigma$  and  $\Sigma_a$ , with  $v$  controlling the level of misspecification. We varied  $v \in \{0.5, 1, 1.5\}$  and analyzed its impact on sDM’s performance.

**Misspecified prior means and covariances (Fig. 11).** This is achieved by adding uniformly sampled noise from  $[v, v+0.5]$  to both the true prior mean and covariance parameters  $\mu, \Sigma, W_a, \Sigma_a$ , with  $v$  controlling the level of misspecification. We varied  $v \in \{0.5, 1, 1.5\}$  and analyzed its impact on sDM’s performance. For comparison, we included the well-specified sDM and the most competitive baseline, DM (Bayes), while omitting other baselines to reduce clutter. In all figures (Figs. 9 to 11) sDM’s performance decreases with increasing misspecification, yet sDM with misspecification still outperforms the most competitive baseline, especially when  $K$  is large. We also observe that the impact of prior covariance misspecification is less significant compared to prior mean misspecification.

**Misspecified likelihood (Fig. 12).** We also simulate the case when the true reward distribution differed from the likelihood assumed by sDM. For example, we simulated binary rewards using a

Bernoulli-logistic model while sDM used a linear-Gaussian likelihood. Other DMs: DM (Bayes) and DM (Freq) also use a misspecified likelihood model and to emphasize this we add the suffix Lin to all DMs names. Overall, as known in the literature, IPS variants performed better than standard DMs when the likelihood was misspecified. However, this is not the case for sDM. While misspecification narrowed the performance gap between sDM and baselines, sDM still consistently outperformed them, particularly in low-data regimes or large action spaces. In a few cases, some baselines, e.g., snIPS, showed competitive performance with enough data.




Figure 9: Effect of prior mean misspecification: The relative reward of the learned policy on synthetic problems using both misspecified prior means and covariances with varying  $n$  and  $K$ .




Figure 10: Effect of prior covariance misspecification: The relative reward of the learned policy on synthetic problems using both misspecified prior means and covariances with varying  $n$  and  $K$ .




Figure 11: Effect of prior mean and covariance misspecification: The relative reward of the learned policy on synthetic problems using both misspecified prior means and covariances with varying  $n$  and  $K$ .




Figure 12: Effect of likelihood misspecification: The relative reward of the learned policy on synthetic problems using misspecified likelihood with varying  $n$  and  $K$ .

#### F.4 Comparison to MIPS

MIPS was designed for OPE only and thus we consider the OPE setting in [Appendix F.1](#). MIPS assumes access to embeddings  $E_i$  in the sample set  $S$  and defines its estimator as follows

$$\hat{V}^{\text{MIPS}}(\pi, S) = \frac{1}{n} \sum_{i=1}^n \frac{\pi(E_i | X_i)}{\pi_0(E_i | X_i)} R_i = \frac{1}{n} \sum_{i=1}^n w(X_i, E_i) R_i,$$

where the sample set  $S = \{(X_i, A_i, E_i, R_i)\}_{i=1}^n$  now includes action embeddings  $E_i$  for each data point  $i \in [n]$ , and the marginal importance weight is given by  $w(x, e) = \frac{\pi(e|x)}{\pi_0(e|x)} = \frac{\sum_a p(e|x,a)\pi(a|x)}{\sum_a p(e|x,a)\pi_0(a|x)}$ , where  $p(\cdot | x, a)$  is the embedding distribution such that  $E_i \sim p(\cdot | X_i, A_i)$ . They assume that  $R$  is independent of  $A$  given  $X$  and  $E$  and hence we should define embeddings that have a direct effect on the reward. From our prior (10), these correspond to the true action parameters  $\theta_*$ . In fact, creating MIPS’s structure from our prior in (10) is not straightforward, but we approach it as follows. We define the embedding distribution as  $p(\cdot | x, a) = \mathcal{N}(\cdot; \mu_a, \Sigma_a)$ , where  $\mu_a$  and  $\Sigma_a$  are obtained by marginalizing out the latent parameters  $\psi$  in (10). Additionally, we set  $A_i = \theta_{*, A_i}$ , having access to  $\theta_*$  since these are synthetic problems. This is the best version of MIPS that we could create from our prior (10). However, we stress again that this comparison is only provided for completeness as the setting may be favorable to sDM. In this setting, sDM outperforms MIPS.

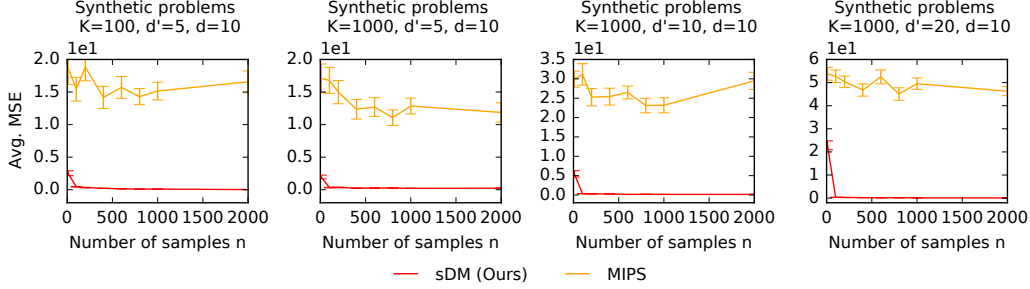


Figure 13: Comparison of sDM and MIPS in OPE: The average MSE of an  $\epsilon$ -greedy policy on synthetic problems with varying  $n$ ,  $K$  and  $d'$ .

### E.5 Comparison of greedy and pessimistic policies

To validate our theory that a greedy policy should be preferred over the commonly adopted pessimistic policy in our Bayesian setting, we used a performance metric averaged over multiple bandit problems sampled from the prior. To verify this, we considered the same OPL synthetic setting as in Section 6 and compared sDM with a greedy policy to sDM with a pessimistic policy. Recall that a greedy policy with respect to our reward estimate writes

$$\hat{\pi}_G(a | x) = \mathbb{1}\{a = \operatorname{argmax}_{b \in \mathcal{A}} \hat{r}(x, b)\}, \quad (46)$$

while a pessimistic one writes

$$\hat{\pi}_P(a | x) = \mathbb{1}\{a = \operatorname{argmax}_{b \in \mathcal{A}} \hat{r}(x, b) - u(x, a)\}, \quad (47)$$

where  $u(x, a) = \alpha(d, \delta) \|\phi(X)\|_{\hat{\Sigma}_a}$  with  $\alpha(d, \delta) = \sqrt{d + 2\sqrt{d \log \frac{1}{\delta}} + 2 \log \frac{1}{\delta}}$  and this is derived in (29) in Appendix E.2. As predicted by our theory, the results show that the greedy policy has better average performance over multiple bandit instances sampled from the prior.




Figure 14: Comparison of sDM with greedy policy and sDM with pessimistic policy in OPL: The average MSE of an  $\epsilon$ -greedy policy on synthetic problems with varying  $n$ ,  $K$ , and  $d'$ .

## G Broader impact

This work contributes to the development and analysis of practical algorithms for offline learning to act under uncertainty. While our generic setting and algorithms have broad potential applications, the specific downstream social impacts are inherently dependent on the chosen application.

## 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: All claims are supported by the theory in [Section 5](#) (with proofs provided in the appendix) and experiments in [Section 6](#). The limitations of direct methods were also discussed in the introduction.

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: Limitations were discussed in the conclusion in [Section 7](#).

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: Assumptions are mentioned in the main text and restated in the appendix ([Assumption D.1](#), [Assumption E.3](#)), where proofs are also provided.

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: [Yes]

Justification: Information needed to reproduce the main experimental results of the paper is described in [Section 6](#) and [Appendix F](#).

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: [Yes]

Justification: The code for the main experiments is shared in the supplementary material.

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: [Yes]

Justification: All experimental details are described in [Section 6](#) and [Appendix F](#).

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: [Yes]

Justification: Standard error bars are included in the figures.

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: [Yes]

Justification: As mentioned in [Appendix F](#), our experiments were conducted on internal machines with 30 CPUs and thus they required a moderate amount of computation. These experiments are also reproducible with minimal computational resources.

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: This work contributes to the development and theoretical analysis of offline learning to act under uncertainty and it adheres to the 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: [Yes]

Justification: Broader Impacts are discussed in [Appendix G](#).

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: Our paper is mainly theoretical and the used data is either simulated or openly accessible. Thus, we believe that our work poses no such risks.

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: [Yes]

Justification: To the best of our knowledge, all relevant and used papers were cited.

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](https://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: [Yes]

Justification: We include our code as supplementary material, with all details needed for reproducibility given in [Section 6](#) and [Appendix F](#).

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 work does not involve crowdsourcing nor research with 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: This work does not involve crowdsourcing nor research with human subjects.

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.