
Off-policy evaluation for slate recommendation

Adith Swaminathan

Microsoft Research, Redmond
adswamin@microsoft.com

Akshay Krishnamurthy

University of Massachusetts, Amherst
akshay@cs.umass.edu

Alekh Agarwal

Microsoft Research, New York
alekha@microsoft.com

Miroslav Dudík

Microsoft Research, New York
mdudik@microsoft.com

John Langford

Microsoft Research, New York
jcl@microsoft.com

Damien Jose

Microsoft, Redmond
dajose@microsoft.com

Imed Zitouni

Microsoft, Redmond
izitouni@microsoft.com

Abstract

This paper studies the evaluation of policies that recommend an ordered set of items (e.g., a ranking) based on some context—a common scenario in web search, ads, and recommendation. We build on techniques from combinatorial bandits to introduce a new practical estimator that uses logged data to estimate a policy’s performance. A thorough empirical evaluation on real-world data reveals that our estimator is accurate in a variety of settings, including as a subroutine in a learning-to-rank task, where it achieves competitive performance. We derive conditions under which our estimator is unbiased—these conditions are weaker than prior heuristics for slate evaluation—and experimentally demonstrate a smaller bias than parametric approaches, even when these conditions are violated. Finally, our theory and experiments also show exponential savings in the amount of required data compared with general unbiased estimators.

1 Introduction

In recommendation systems for e-commerce, search, or news, we would like to use the data collected during operation to test new content-serving algorithms (called *policies*) along metrics such as revenue and number of clicks [4, 25]. This task is called *off-policy evaluation*. General approaches, namely *inverse propensity scores* (IPS) [13, 18], require unrealistically large amounts of logged data to evaluate whole-page metrics that depend on multiple recommended items, which happens when showing ranked lists. The key challenge is that the number of possible lists (called *slates*) is combinatorially large. As a result, the policy being evaluated is likely to choose different slates from those recorded in the logs most of the time, unless it is very similar to the data-collection policy. This challenge is fundamental [34], so any off-policy evaluation method that works with large slates needs to make some structural assumptions about the whole-page metric or the user behavior.

Previous work on off-policy evaluation and whole-page optimization improves the probability of match between logging and evaluation by restricting attention to small slate spaces [35, 26], introducing assumptions that allow for partial matches between the proposed and observed slates [27], or assuming that the policies used for logging and evaluation are similar [4, 32]. Another line of work constructs parametric models of slate quality [8, 16, 14] (see also Sec. 4.3 of [17]). While these approaches require less data, they can have large bias, and their use in practice requires an expensive trial-and-error cycle involving weeks-long A/B tests to develop new policies [20]. In this paper we

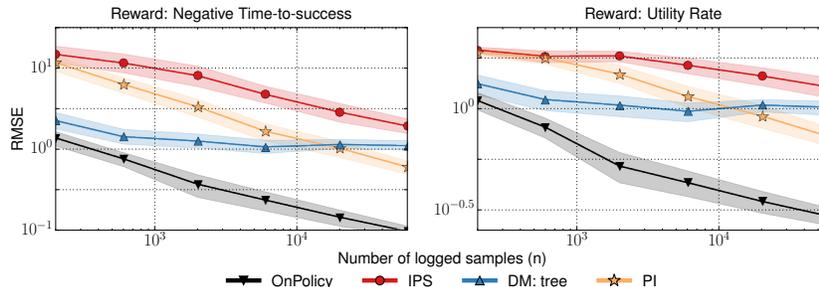


Figure 1: Off-policy evaluation of two whole-page user-satisfaction metrics on proprietary search engine data. Average RMSE of different estimators over 50 runs on a log-log scale. Our method (PI) achieves the best performance with moderate data sizes. The unbiased IPS method suffers high variance, and direct modeling (DM) of the metrics suffers high bias. ONPOLICY is the expensive choice of deploying the policy, for instance, in an A/B test.

design a method more robust to problems with bias and with only modest data requirements, with the goal of substantially shortening this cycle and accelerating the policy development process.

We frame the slate recommendation problem as a combinatorial generalization of *contextual bandits* [3, 23, 13]. In combinatorial contextual bandits, for each *context*, a policy selects a *slate* consisting of component *actions*, after which a *reward* for the entire slate is observed. In web search, the context is the search query augmented with a user profile, the slate is the search results page consisting of a list of retrieved documents (actions), and example reward metrics are page-level measures such as time-to-success, NDCG (position-weighted relevance), or other measures of user satisfaction. As input we receive contextual bandit data obtained by some *logging policy*, and our goal is to estimate the reward of a new *target policy*. This off-policy setup differs from online learning in contextual bandits, where the goal is to adaptively maximize the reward in the presence of an explore-exploit trade-off [5].

Inspired by work in *combinatorial* and *linear bandits* [7, 31, 11], we propose an estimator that makes only a weak assumption about the evaluated metric, while exponentially reducing the data requirements in comparison with IPS. Specifically, we posit a *linearity assumption*, stating that the slate-level reward (e.g., time to success in web search) decomposes additively across actions, but the action-level rewards are not observed. Crucially, the action-level rewards are allowed to depend on the context, and we do not require that they be easily modeled from the features describing the context. In fact, our method is completely agnostic to the representation of contexts.

We make the following contributions:

1. The *pseudoinverse estimator* (PI) for off-policy evaluation: a general-purpose estimator from the combinatorial bandit literature, adapted for off-policy evaluation. When ranking ℓ out of m items under the linearity assumption, PI typically requires $\mathcal{O}(\ell m / \varepsilon^2)$ samples to achieve error at most ε —an exponential gain over the $m^{\Omega(\ell)}$ sample complexity of IPS. We provide distribution-dependent bounds based on the overlap between logging and target policies.
2. Experiments on real-world search ranking datasets: The strong performance of the PI estimator provides, to our knowledge, the first demonstration of high-quality off-policy evaluation of whole-page metrics, comprehensively outperforming prior baselines (see Fig. 1).
3. Off-policy optimization: We provide a simple procedure for learning to rank (L2R) using the PI estimator to impute action-level rewards for each context. This allows direct optimization of whole-page metrics via pointwise L2R approaches, *without requiring pointwise feedback*.

Related work Large state spaces have typically been studied in the online, or on-policy, setting. Some works assume specific parametric (e.g., linear) models relating the metrics to the features describing a slate [2, 31, 15, 10, 29]; this can lead to bias if the model is inaccurate (e.g., we might not have access to sufficiently predictive features). Others posit the same linearity assumption as we do, but further assume a *semi-bandit* feedback model where the rewards of all actions on the slate

are revealed [19, 22, 21]. While much of the research focuses on on-policy setting, the off-policy paradigm studied in this paper is often preferred in practice since it might not be possible to implement low-latency updates needed for online learning, or we might be interested in many different metrics and require a manual review of their trade-offs before deploying new policies.

At a technical level, the PI estimator has been used in online learning [7, 31, 11], but the analysis there is tailored to the specific data collection policies used by the learner. In contrast, we provide distribution-dependent bounds without any assumptions on the logging or target policy.

2 Setting and notation

In combinatorial contextual bandits, a decision maker repeatedly interacts as follows:

1. the decision maker observes a *context* x drawn from a distribution $D(x)$ over some space X ;
2. based on the context, the decision maker chooses a *slate* $\mathbf{s} = (s_1, \dots, s_\ell)$ consisting of *actions* s_j , where a position j is called a *slot*, the number of slots is ℓ , actions at position j come from some space $A_j(x)$, and the slate \mathbf{s} is chosen from a set of allowed slates $S(x) \subseteq A_1(x) \times \dots \times A_\ell(x)$;
3. given the context and slate, a reward $r \in [-1, 1]$ is drawn from a distribution $D(r | x, \mathbf{s})$; rewards in different rounds are independent, conditioned on contexts and slates.

The context space X can be infinite, but the set of actions is finite. We assume $|A_j(x)| = m_j$ for all contexts $x \in X$ and define $m := \max_j m_j$ as the maximum number of actions per slot. The goal of the decision maker is to *maximize the reward*. The decision maker is modeled as a *stochastic policy* π that specifies a conditional distribution $\pi(\mathbf{s} | x)$ (a deterministic policy is a special case). The *value* of a policy π , denoted $V(\pi)$, is defined as the expected reward when following π :

$$V(\pi) := \mathbb{E}_{x \sim D} \mathbb{E}_{\mathbf{s} \sim \pi(\cdot | x)} \mathbb{E}_{r \sim D(\cdot | x, \mathbf{s})} [r] . \quad (1)$$

To simplify derivations, we extend the conditional distribution π into a distribution over triples (x, \mathbf{s}, r) as $\pi(x, \mathbf{s}, r) := D(r | x, \mathbf{s}) \pi(\mathbf{s} | x) D(x)$. With this shorthand, we have $V(\pi) = \mathbb{E}_\pi[r]$.

To finish this section, we introduce notation for the expected reward for a given context and slate, which we call the *slate value*, and denote as:

$$V(x, \mathbf{s}) := \mathbb{E}_{r \sim D(\cdot | x, \mathbf{s})} [r] . \quad (2)$$

Example 1 (Cartesian product). Consider the optimization of a news portal where the reward is the whole-page advertising revenue. Context x is the user profile, slate is the news-portal page with slots corresponding to news sections,¹ and actions are the articles. The set of valid slates is the Cartesian product $S(x) = \prod_{j \leq \ell} A_j(x)$. The number of valid slates is exponential in ℓ : $|S(x)| = \prod_{j \leq \ell} m_j$.

Example 2 (Ranking). Consider web search and ranking. Context x is the query along with user profile. Actions correspond to search items (such as webpages). The policy chooses ℓ of m items, where the set $A(x)$ of m items for a context x is chosen from a corpus by a filtering step (e.g., a database query). We have $A_j(x) = A(x)$ for all $j \leq \ell$, but the allowed slates $S(x)$ have no repetitions. The number of valid slates is exponential in ℓ : $|S(x)| = m! / (m - \ell)! = m^{\Omega(\ell)}$. A reward could be the *negative time-to-success*, i.e., negative of the time taken by the user to find a relevant item.

2.1 Off-policy evaluation and optimization

In the *off-policy* setting, we have access to the *logged data* $(x_1, \mathbf{s}_1, r_1), \dots, (x_n, \mathbf{s}_n, r_n)$ collected using a past policy μ , called the *logging policy*. *Off-policy evaluation* is the task of estimating the value of a new policy π , called the *target policy*, using the logged data. *Off-policy optimization* is the harder task of finding a policy $\hat{\pi}$ that achieves maximal reward.

There are two standard approaches for off-policy evaluation. The *direct method* (DM) uses the logged data to train a (parametric) model $\hat{r}(x, \mathbf{s})$ for predicting the expected reward for a given context and slate. $V(\pi)$ is then estimated as

$$\hat{V}_{\text{DM}}(\pi) = \frac{1}{n} \sum_{i=1}^n \sum_{\mathbf{s} \in S(x_i)} \hat{r}(x_i, \mathbf{s}) \pi(\mathbf{s} | x_i) . \quad (3)$$

¹For simplicity, we do not discuss the more general setting of showing multiple articles in each news section.

The direct method is often biased due to mismatch between model assumptions and ground truth.

The second approach, which is provably unbiased (under modest assumptions), is the *inverse propensity score* (IPS) estimator [18]. The IPS estimator re-weights the logged data according to ratios of slate probabilities under the target and logging policy. It has two common variants:

$$\hat{V}_{\text{IPS}}(\pi) = \frac{1}{n} \sum_{i=1}^n r_i \cdot \frac{\pi(\mathbf{s}_i | x_i)}{\mu(\mathbf{s}_i | x_i)}, \quad \hat{V}_{\text{wIPS}}(\pi) = \sum_{i=1}^n r_i \cdot \frac{\pi(\mathbf{s}_i | x_i)}{\mu(\mathbf{s}_i | x_i)} / \left(\sum_{i=1}^n \frac{\pi(\mathbf{s}_i | x_i)}{\mu(\mathbf{s}_i | x_i)} \right). \quad (4)$$

wIPS generally has better variance with an asymptotically zero bias. The variance of both estimators grows linearly with $\frac{\pi(\mathbf{s} | x)}{\mu(\mathbf{s} | x)}$, which can be $\Omega(|S(x)|)$. This is prohibitive when $|S(x)| = m^{\Omega(\ell)}$.

3 Our approach

The IPS estimator is minimax optimal [34], so its exponential variance is unavoidable in the worst case. We circumvent this hardness by positing an assumption on the structure of rewards. Specifically, we assume that the slate-level reward is a sum of unobserved action-level rewards that depend on the context, the action, and the position on the slate, but not on the other actions on the slate.

Formally, we consider *slate indicator vectors* in $\mathbb{R}^{\ell m}$ whose components are indexed by pairs (j, a) of slots and possible actions in them. A slate is described by an indicator vector $\mathbf{1}_{\mathbf{s}} \in \mathbb{R}^{\ell m}$ whose entry at position (j, a) is equal to 1 if the slate \mathbf{s} has action a in slot j , i.e., if $s_j = a$. The above assumption is formalized as follows:

Assumption 1 (Linearity Assumption). For each context $x \in X$ there exists an (unknown) *intrinsic reward* vector $\phi_x \in \mathbb{R}^{\ell m}$ such that the slate value satisfies $V(x, \mathbf{s}) = \mathbf{1}_{\mathbf{s}}^T \phi_x = \sum_{j=1}^{\ell} \phi_x(j, s_j)$.

The slate indicator vector can be viewed as a feature vector, representing the slate, and ϕ_x can be viewed as a *context-specific* weight vector. The assumption refers to the fact that the value of a slate is a linear function of its feature representation. However, note that this linear dependence is allowed to be completely different across contexts, because we make no assumptions on how ϕ_x depends on x , and in fact our method does not even attempt to accurately estimate ϕ_x . Being agnostic to the form of ϕ_x is the key departure from the direct method and parametric bandits.

While Assumption 1 rules out interactions among different actions on a slate,² its ability to vary intrinsic rewards arbitrarily across contexts captures many common metrics in information retrieval, such as the *normalized discounted cumulative gain* (NDCG) [6], a common metric in web ranking:

Example 3 (NDCG). For a slate \mathbf{s} , we first define $\text{DCG}(x, \mathbf{s}) := \sum_{j=1}^{\ell} \frac{2^{\text{rel}(x, s_j)} - 1}{\log_2(j+1)}$ where $\text{rel}(x, a)$ is the relevance of document a on query x . Then $\text{NDCG}(x, \mathbf{s}) := \text{DCG}(x, \mathbf{s}) / \text{DCG}^*(x)$ where $\text{DCG}^*(x) = \max_{\mathbf{s} \in S(x)} \text{DCG}(x, \mathbf{s})$, so NDCG takes values in $[0, 1]$. Thus, NDCG satisfies Assumption 1 with $\phi_x(j, a) = (2^{\text{rel}(x, a)} - 1) / \log_2(j + 1) \text{DCG}^*(x)$.

In addition to Assumption 1, we also make the standard assumption that the logging policy puts non-zero probability on all slates that can be potentially chosen by the target policy. This assumption is also required for IPS, otherwise unbiased off-policy evaluation is impossible [24].

Assumption 2 (Absolute Continuity). The off-policy evaluation problem satisfies the *absolute continuity* assumption if $\mu(\mathbf{s} | x) > 0$ whenever $\pi(\mathbf{s} | x) > 0$ with probability one over $x \sim D$.

3.1 The pseudoinverse estimator

Using Assumption 1, we can now apply the techniques from the combinatorial bandit literature to our problem. In particular, our estimator closely follows the recipe of Cesa-Bianchi and Lugosi [7], albeit with some differences to account for the off-policy and contextual nature of our setup. Under Assumption 1, we can view the recovery of ϕ_x for a given context x as a linear regression problem. The covariates $\mathbf{1}_{\mathbf{s}}$ are drawn according to $\mu(\cdot | x)$, and the reward follows a linear model, conditional on \mathbf{s} and x , with ϕ_x as the “weight vector”. Thus, we can write the MSE of an estimate \mathbf{w} as $\mathbb{E}_{\mathbf{s} \sim \mu(\cdot | x)} \mathbb{E}_{r \sim D(\cdot | \mathbf{s}, x)} [(\mathbf{1}_{\mathbf{s}}^T \mathbf{w} - r)^2]$, or more compactly as $\mathbb{E}_{\mu} [(\mathbf{1}_{\mathbf{s}}^T \mathbf{w} - r)^2 | x]$, using our definition of μ as a distribution over triples (x, \mathbf{s}, r) . We estimate ϕ_x by the MSE minimizer with the smallest

²We discuss limitations of Assumption 1 and directions to overcome them in Sec. 5.

norm, which can be written in closed form as

$$\bar{\phi}_x = \left(\mathbb{E}_\mu[\mathbf{1}_s \mathbf{1}_s^T | x] \right)^\dagger \mathbb{E}_\mu[r \mathbf{1}_s | x] , \quad (5)$$

where \mathbf{M}^\dagger is the Moore-Penrose pseudoinverse of a matrix \mathbf{M} . Note that this idealized ‘‘estimator’’ $\bar{\phi}_x$ uses conditional expectations over $\mathbf{s} \sim \mu(\cdot | x)$ and $r \sim D(\cdot | \mathbf{s}, x)$. To simplify the notation, we write $\mathbf{\Gamma}_{\mu,x} := \mathbb{E}_\mu[\mathbf{1}_s \mathbf{1}_s^T | x] \in \mathbb{R}^{\ell m \times \ell m}$ to denote the (uncentered) covariance matrix for our regression problem, appearing on the right-hand side of Eq. (5). We also introduce notation for the second term in Eq. (5) and its empirical estimate: $\boldsymbol{\theta}_{\mu,x} := \mathbb{E}_\mu[r \mathbf{1}_s | x]$, and $\hat{\boldsymbol{\theta}}_i := r_i \mathbf{1}_{s_i}$.

Thus, our regression estimator (5) is simply $\bar{\phi}_x = \mathbf{\Gamma}_{\mu,x}^\dagger \boldsymbol{\theta}_{\mu,x}$. Under Assumptions 1 and 2, it is easy to show that $V(x, \mathbf{s}) = \mathbf{1}_s^T \bar{\phi}_x = \mathbf{1}_s^T \mathbf{\Gamma}_{\mu,x}^\dagger \boldsymbol{\theta}_{\mu,x}$. Replacing $\boldsymbol{\theta}_{\mu,x}$ with $\hat{\boldsymbol{\theta}}_i$ motivates the following estimator for $V(\pi)$, which we call the *pseudoinverse estimator* or PI:

$$\hat{V}_{\text{PI}}(\pi) = \frac{1}{n} \sum_{i=1}^n \sum_{\mathbf{s} \in S} \pi(\mathbf{s} | x_i) \mathbf{1}_s^T \mathbf{\Gamma}_{\mu,x_i}^\dagger \hat{\boldsymbol{\theta}}_i = \frac{1}{n} \sum_{i=1}^n r_i \cdot \mathbf{q}_{\pi,x_i}^T \mathbf{\Gamma}_{\mu,x_i}^\dagger \mathbf{1}_{s_i} . \quad (6)$$

In Eq. (6) we have expanded the definition of $\hat{\boldsymbol{\theta}}_i$ and introduced the notation $\mathbf{q}_{\pi,x}$ for the expected slate indicator under π conditional on x , $\mathbf{q}_{\pi,x} := \mathbb{E}_\pi[\mathbf{1}_s | x]$. The summation over \mathbf{s} required to obtain \mathbf{q}_{π,x_i} in Eq. (6) can be replaced by a small sample. We can also derive a weighted variant of PI:

$$\hat{V}_{\text{wPI}}(\pi) = \frac{\sum_{i=1}^n r_i \cdot \mathbf{q}_{\pi,x_i}^T \mathbf{\Gamma}_{\mu,x_i}^\dagger \mathbf{1}_{s_i}}{\sum_{i=1}^n \mathbf{q}_{\pi,x_i}^T \mathbf{\Gamma}_{\mu,x_i}^\dagger \mathbf{1}_{s_i}} . \quad (7)$$

We prove the following unbiasedness property in Appendix A.

Proposition 1. *If Assumptions 1 and 2 hold, then the estimator \hat{V}_{PI} is unbiased, i.e., $\mathbb{E}_{\mu^n}[\hat{V}_{\text{PI}}] = V(\pi)$, where the expectation is over the n logged examples sampled i.i.d. from μ .*

As special cases, PI reduces to IPS when $\ell = 1$, and simplifies to $\sum_{i=1}^n r_i/n$ when $\pi = \mu$ (see Appendix C). To build further intuition, we consider the settings of Examples 1 and 2, and simplify the PI estimator to highlight the improvement over IPS.

Example 4 (PI for a Cartesian product when μ is a product distribution). The PI estimator for the Cartesian product slate space, when μ factorizes across slots as $\mu(\mathbf{s} | x) = \prod_j \mu(s_j | x)$, simplifies to

$$\hat{V}_{\text{PI}}(\pi) = \frac{1}{n} \sum_{i=1}^n r_i \cdot \left(\sum_{j=1}^{\ell} \frac{\pi(s_{ij} | x_i)}{\mu(s_{ij} | x_i)} - \ell + 1 \right) ,$$

by Prop. 2 in Appendix D. Note that unlike IPS, which divides by probabilities of whole slates, the PI estimator only divides by probabilities of actions appearing in individual slots. Thus, the magnitude of each term of the outer summation is only $\mathcal{O}(\ell m)$, whereas the IPS terms are $m^{\Omega(\ell)}$.

Example 5 (PI for rankings with $\ell = m$ and uniform logging). In this case,

$$\hat{V}_{\text{PI}}(\pi) = \frac{1}{n} \sum_{i=1}^n r_i \cdot \left(\sum_{j=1}^{\ell} \frac{\pi(s_{ij} | x_i)}{1/(m-1)} - m + 2 \right) ,$$

by Prop. 4 in Appendix E.1. The summands are again $\mathcal{O}(\ell m) = \mathcal{O}(m^2)$.

3.2 Deviation analysis

So far, we have shown that PI is unbiased under our assumptions and overcomes the deficiencies of IPS in specific examples. We now derive a finite-sample error bound, based on the overlap between π and μ . We use Bernstein’s inequality, for which we define the variance and range terms:

$$\sigma^2 := \mathbb{E}_{x \sim D} \left[\mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{q}_{\pi,x} \right] , \quad \rho := \sup_x \sup_{\mathbf{s}: \mu(\mathbf{s}|x) > 0} \left| \mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{1}_s \right| . \quad (8)$$

The quantity σ^2 bounds the variance whereas ρ bounds the range. They capture the ‘‘average’’ and ‘‘worst-case’’ mismatch between μ and π . They equal one when $\pi = \mu$ (see Appendix C), and yield the following deviation bound:

Theorem 1. Under Assumptions 1 and 2, let σ^2 and ρ be defined as in Eq. (8). Then, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$,

$$\left| \hat{V}_{\text{PI}}(\pi) - V(\pi) \right| \leq \sqrt{\frac{2\sigma^2 \ln(2/\delta)}{n}} + \frac{2(\rho + 1) \ln(2/\delta)}{3n} .$$

We observe that this finite sample bound is structurally different from the regret bounds studied in the prior works on combinatorial bandits. The bound incorporates the extent of overlap between π and μ so that we have higher confidence in our estimates when the logging and evaluation policies are similar—an important consideration in off-policy evaluation.

While the bound might look complicated, it simplifies if we consider the class of ε -uniform logging policies. Formally, for any policy μ , define $\mu_\varepsilon(\mathbf{s} \mid x) = (1 - \varepsilon)\mu(\mathbf{s} \mid x) + \varepsilon\nu(\mathbf{s} \mid x)$, with ν being the uniform distribution over the set $S(x)$. For suitably small ε , such logging policies are widely used in practice. We have the following corollary for these policies, proved in Appendix E:

Corollary 1. In the settings of Example 1 or Example 2, if the logging is done with μ_ε for some $\varepsilon > 0$, we have $|\hat{V}_{\text{PI}}(\pi) - V(\pi)| \leq \mathcal{O}(\sqrt{\varepsilon^{-1} \ell m/n})$.

Again, this turns the $\Omega(m^\ell)$ data dependence of IPS into $O(m\ell)$. The key step in the proof is the bound on a certain norm of Γ_ν^\dagger , similar to the bounds of Cesa-Bianchi and Lugosi [7], but our results are a bit sharper.

4 Experiments

We empirically evaluate the performance of the pseudoinverse estimator for ranking problems. We first show that PI outperforms prior works in a comprehensive semi-synthetic study using a public dataset. We then use our estimator for *off-policy optimization*, i.e., to learn ranking policies, competitively with supervised learning that uses more information. Finally, we demonstrate substantial improvements on proprietary data from search engine logs for two user-satisfaction metrics used in practice: *time-to-success* and *utility rate*, which do not satisfy the linearity assumption. More detailed results are deferred to Appendices F and G. All of our code is available online.³

4.1 Semi-synthetic evaluation

Our semi-synthetic evaluation uses labeled data from the Microsoft Learning to Rank Challenge dataset [30] (MSLR-WEB30K) to create a contextual bandit instance. Queries form the contexts x and actions a are the available documents. The dataset contains over 31K queries, each with up to 1251 judged documents, where the query-document pairs are judged on a 5-point scale, $rel(x, a) \in \{0, \dots, 4\}$. Each pair (x, a) has a feature vector $\mathbf{f}(x, a)$, which can be partitioned into title and body features ($\mathbf{f}_{\text{title}}$ and \mathbf{f}_{body}). We consider two slate rewards: NDCG from Example 3, and the *expected reciprocal rank*, ERR [9], which *does not* satisfy linearity, and is defined as

$$\text{ERR}(x, \mathbf{s}) := \sum_{r=1}^{\ell} \frac{1}{r} \prod_{i=1}^{r-1} (1 - R(s_i)) R(s_r) , \quad \text{where } R(a) = \frac{2^{rel(x,a)} - 1}{2^{maxrel} - 1} \text{ with } maxrel = 4.$$

To derive several distinct logging and target policies, we first train two lasso regression models, called $lasso_{\text{title}}$ and $lasso_{\text{body}}$, and two regression tree models, called $tree_{\text{title}}$ and $tree_{\text{body}}$, to predict relevances from $\mathbf{f}_{\text{title}}$ and \mathbf{f}_{body} , respectively. To create the logs, queries x are sampled uniformly, and the set $A(x)$ consists of the top m documents according to $tree_{\text{title}}$. The logging policy is parametrized by a model, either $tree_{\text{title}}$ or $lasso_{\text{title}}$, and a scalar $\alpha \geq 0$. It samples from a multinomial distribution over documents $p_\alpha(a|x) \propto 2^{-\alpha \lfloor \log_2 rank(x,a) \rfloor}$ where $rank(x, a)$ is the rank of document a for query x according to the corresponding model. Slates are constructed slot-by-slot, sampling *without replacement* according to p_α . Varying α interpolates between uniformly random and deterministic logging. Thus, all logging policies are based on the models derived from $\mathbf{f}_{\text{title}}$. We consider two deterministic target policies based on the two models derived from \mathbf{f}_{body} , i.e., $tree_{\text{body}}$ and $lasso_{\text{body}}$, which select the top ℓ documents according to the corresponding model. The four base models are fairly distinct: on average fewer than 2.75 documents overlap among top 10 (see Appendix H).

³https://github.com/adith387/slates_semisynt_expts

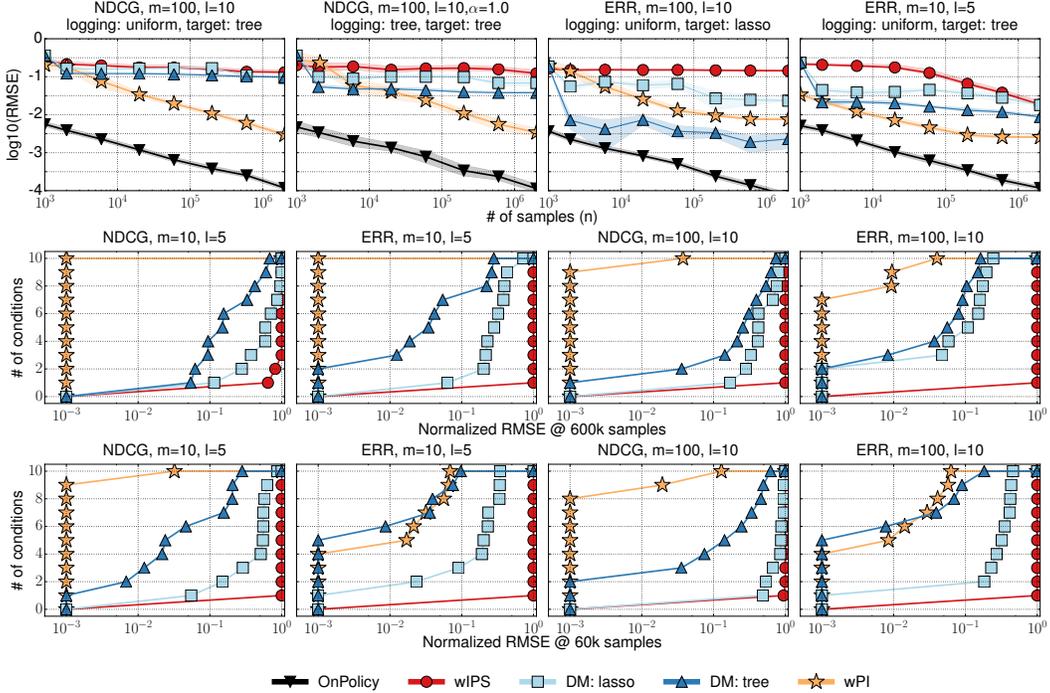


Figure 2: *Top*: RMSE of various estimators under four experimental conditions (see Appendix F for all 40 conditions). *Middle*: CDF of normalized RMSE at 600k samples; each plot aggregates over 10 logging-target combinations; closer to top-left is better. *Bottom*: Same as middle but at 60k samples.

We compare the weighted estimator wPI with the direct method (DM) and weighted IPS (wIPS). (Weighted variants outperformed the unweighted ones.) We implement two variants of DM: regression trees and lasso, each trained on the first $n/2$ examples and using the remaining $n/2$ examples for evaluation according to Eq. (3). We also include an aspirational baseline, ONPOLICY, which corresponds to deploying the target policy as in an A/B test and returning the average of observed rewards. This is the expensive alternative we wish to avoid.

We evaluate the estimators by recording the root mean square error (RMSE) as a function of the number of samples, averaged over at least 25 independent runs. We do this for 40 different experimental conditions, considering two reward metrics, two slate-space sizes, and 10 combinations of target and logging policies (including the choice of α). The top row of Fig. 2 shows results for four representative conditions (see Appendix F for all results), while the middle and bottom rows aggregate across conditions. To produce the aggregates, we shift and rescale the RMSE of all methods, at 600k (middle row) or 60k (bottom row) samples, so the best performance is at 0.001 and the worst is at 1.0 (excluding ONPOLICY). (We use 0.001 instead of 0.0 to allow plotting on a log scale.) The aggregate plots display the cumulative distribution function of these normalized RMSE values across 10 target-logging combinations, keeping the metric and the slate-space size fixed.

The pseudoinverse estimator wPI easily dominates wIPS across all experimental conditions, as can be seen in Fig. 2 (top) and in Appendix F. While wIPS and IPS are (asymptotically) unbiased even without linearity assumption, they both suffer from a large variance caused by the slate size. The variance and hence the mean square error of wIPS and IPS grows exponentially with the slate size, so they perform poorly beyond the smallest slate sizes. DM performs well in some cases, especially with few samples, but often plateaus or degrades eventually as it overfits on the logging distribution, which is different from the target. While wPI does not always outperform DM methods (e.g., Fig. 2, top row, second from right), it is the only method that works robustly across all conditions, as can be seen in the aggregate plots. In general, choosing between DM and wPI is largely a matter of bias-variance tradeoff. DM can be particularly good with very small data sizes, because of its low variance, and in those settings it is often the best choice. However, PI performs comprehensively better given enough data (see Fig. 2, middle row).

In the top row of Fig. 2, we see that, as expected, wPI is biased for the ERR metric since ERR does not satisfy linearity. The right two panels also demonstrate the effect of varying m and ℓ . While wPI deteriorates somewhat for the larger slate space, it still gives a meaningful estimate. In contrast, wPS fails to produce any meaningful estimate in the larger slate space and its RMSE barely improves with more data. Finally, the left two plots in the top row show that wPI is fairly insensitive to the amount of stochasticity in logging, whereas DM improves with more overlap between logging and target.

4.2 Semi-synthetic policy optimization

We now show how to use the pseudoinverse estimator for off-policy optimization. We leverage pointwise learning to rank (L2R) algorithms, which learn a scoring function for query-document pairs by fitting to relevance labels. We call this the *supervised* approach, as it requires relevance labels.

Instead of requiring relevance labels, we use the pseudoinverse estimator to convert page-level reward into per-slot reward components—the estimates of $\phi_x(j, a)$ —and these become targets for regression. Thus, the pseudoinverse estimator enables pointwise L2R to optimize whole-page metrics even without relevance labels. Given a contextual bandit dataset $\{(x_i, \mathbf{s}_i, r_i)\}_{i \leq n}$ collected by the logging policy μ , we begin by creating the estimates of ϕ_{x_i} : $\hat{\phi}_i = \Gamma_{\mu, x_i}^\dagger \hat{\theta}_i$, turning the i -th contextual bandit example into ℓm regression examples. The trained regression model is used to create a slate, starting with the highest scoring slot-action pair, and continuing greedily (excluding the pairs with the already chosen slots or actions). This procedure is detailed in Appendix G. Note that without the linearity assumptions, our imputed regression targets might not lead to the best possible learned policy, but we still expect to adapt somewhat to the slate-level metric.

We use the MSLR-WEB10K dataset [30] to compare our approach with benchmarked results [33] for NDCG@3 (i.e., $\ell = 3$).⁴ This dataset contains 10k queries, over 1.2M relevance judgments, and up to 908 judged documents per query. The state-of-the-art *listwise* L2R method on this dataset is a highly tuned variant of LambdaMART [1] (with an ensemble of 1000 trees, each with up to 70 leaves).

We use the provided 5-fold split and always train on bandit data collected by uniform logging from four folds, while evaluating with supervised data on the fifth. We compare our approach, titled PI-OPT, against the supervised approach (SUP), trained to predict the *gains*, equal to $2^{rel(x,a)} - 1$, computed using annotated relevance judgements in the training folds (predicting raw relevances was inferior). Both PI-OPT and SUP train gradient boosted regression trees (with 1000 trees, each with up to 70 leaves). Additionally, we also experimented with the ERR metric.

The average test-set performance (computed using ground-truth relevance judgments for each test set) across the 5-folds is reported in Table 1. Our method, PI-OPT is competitive with the supervised baseline SUP for NDCG, and is substantially superior for ERR. A different transformation instead of gains might yield a stronger supervised baseline for ERR, but this only illustrates the key benefit of PI-OPT: *the right pointwise targets are automatically inferred for any whole-page metric*. Both PI-OPT and SUP are slightly worse than LambdaMART for NDCG@3, but they are arguably not as highly tuned, and PI-OPT only uses the slate-level metric.

Table 1: Comparison of L2R approaches optimizing NDCG@3 and ERR@3. LambdaMART is a tuned list-wise approach. SUP and PI-OPT use the same pointwise L2R learner; SUP uses 8×10^5 relevance judgments, PI-OPT uses 10^7 samples (under uniform logging) with page-level rewards.

Metric	LambdaMART	uniformly random	SUP	PI-OPT
NDCG@3	0.457	0.152	0.438	0.421
ERR@3	—	0.096	0.311	0.321

4.3 Real-world experiments

We finally evaluate all methods using logs collected from a popular search engine. The dataset consists of search queries, for which the logging policy randomly (non-uniformly) chooses a slate of

⁴Our dataset here differs from the dataset MSLR-WEB30K used in Sec. 4.1. There our goal was to study realistic problem dimensions, e.g., constructing length-10 rankings out of 100 candidates. Here, we use MSLR-WEB10K, because it is the largest dataset with public benchmark numbers by state-of-the-art approaches (specifically LambdaMART).

size $\ell = 5$ from a small pre-filtered set of documents of size $m \leq 8$. After preprocessing, there are 77 unique queries and 22K total examples, meaning that for each query, we have logged impressions for many of the available slates. As before, we create the logs by sampling queries uniformly at random, and using a logging policy that samples uniformly from the slates shown for this query.

We consider two page-level metrics: time-to-success (TTS) and UTILITYRATE. TTS measures the number of seconds between presenting the results and the first satisfied click from the user, defined as any click for which the user stays on the linked page for sufficiently long. TTS value is capped and scaled to $[0, 1]$. UTILITYRATE is a more complex page-level metric of user satisfaction. It captures the interaction of a user with the page as a timeline of events (such as clicks) and their durations. The events are classified as revealing a positive or negative utility to the user and their contribution is proportional to their duration. UTILITYRATE takes values in $[-1, 1]$.

We evaluate a target policy based on a logistic regression classifier trained to predict clicks and using the predicted probabilities to score slates. We restrict the target policy to pick among the slates in our logs, so we know the ground truth slate-level reward. Since we know the query distribution, we can calculate the target policy’s value exactly, and measure RMSE relative to this true value.

We compare our estimator (PI) with three baselines similar to those from Sec. 4.1: DM, IPS and ONPOLICY. DM uses regression trees over roughly 20,000 slate-level features.

Fig. 1 from the introduction shows that PI provides a consistent multiplicative improvement in RMSE over IPS, which suffers due to high variance. Starting at moderate sample sizes, PI also outperforms DM, which suffers due to substantial bias.

5 Discussion

In this paper we have introduced a new estimator (PI) for off-policy evaluation in combinatorial contextual bandits under a linearity assumption on the slate-level rewards. Our theoretical and empirical analysis demonstrates the merits of the approach. The empirical results show a favorable bias-variance tradeoff. Even in datasets and metrics where our assumptions are violated, the PI estimator typically outperforms all baselines. Its performance, especially at smaller sample sizes, could be further improved by designing doubly-robust variants [12] and possibly also incorporating weight clipping [34].

One promising approach to relax Assumption 1 is to posit a decomposition over pairs (or tuples) of slots to capture higher-order interactions such as diversity. More generally, one could replace slate spaces by arbitrary compact convex sets, as done in linear bandits. In these settings, the pseudoinverse estimator could still be applied, but tight sample-complexity analysis is open for future research.

References

- [1] Nima Asadi and Jimmy Lin. Training efficient tree-based models for document ranking. In *European Conference on Advances in Information Retrieval*, 2013.
- [2] Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 2002.
- [3] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 2002.
- [4] Léon Bottou, Jonas Peters, Joaquin Quiñero-Candela, Denis Charles, 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*, 2013.
- [5] Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 2012.
- [6] Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. Learning to rank using gradient descent. In *International Conference on Machine Learning*, 2005.
- [7] Nicolo Cesa-Bianchi and Gábor Lugosi. Combinatorial bandits. *Journal of Computer and System Sciences*, 2012.

- [8] Olivier Chapelle and Ya Zhang. A dynamic Bayesian network click model for web search ranking. In *International Conference on World Wide Web*, 2009.
- [9] Olivier Chapelle, Donald Metzler, Ya Zhang, and Pierre Grinspan. Expected reciprocal rank for graded relevance. In *Conference on Information and Knowledge Management*, 2009.
- [10] Wei Chu, Lihong Li, Lev Reyzin, and Robert E Schapire. Contextual bandits with linear payoff functions. In *Artificial Intelligence and Statistics*, 2011.
- [11] Varsha Dani, Thomas P. Hayes, and Sham M. Kakade. The price of bandit information for online optimization. In *Advances in Neural Information Processing Systems*, 2008.
- [12] Miroslav Dudík, John Langford, and Lihong Li. Doubly robust policy evaluation and learning. In *International Conference on Machine Learning*, 2011.
- [13] Miroslav Dudík, Dumitru Erhan, John Langford, and Lihong Li. Doubly robust policy evaluation and optimization. *Statistical Science*, 2014.
- [14] Georges E. Dupret and Benjamin Piwowarski. A user browsing model to predict search engine click data from past observations. In *SIGIR Conference on Research and Development in Information Retrieval*, 2008.
- [15] Sarah Filippi, Olivier Cappé, Aurélien Garivier, and Csaba Szepesvári. Parametric bandits: The generalized linear case. In *Advances in Neural Information Processing Systems*, 2010.
- [16] Fan Guo, Chao Liu, Anitha Kannan, Tom Minka, Michael Taylor, Yi-Min Wang, and Christos Faloutsos. Click chain model in web search. In *International Conference on World Wide Web*, 2009.
- [17] Katja Hofmann, Lihong Li, Filip Radlinski, et al. Online evaluation for information retrieval. *Foundations and Trends in Information Retrieval*, 2016.
- [18] Daniel G Horvitz and Donovan J Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 1952.
- [19] Satyen Kale, Lev Reyzin, and Robert E Schapire. Non-stochastic bandit slate problems. In *Advances in Neural Information Processing Systems*, 2010.
- [20] Ron Kohavi, Roger Longbotham, Dan Sommerfield, and Randal M Henne. Controlled experiments on the web: survey and practical guide. *Knowledge Discovery and Data Mining*, 2009.
- [21] Akshay Krishnamurthy, Alekh Agarwal, and Miroslav Dudík. Efficient contextual semi-bandit learning. *Advances in Neural Information Processing Systems*, 2016.
- [22] Branislav Kveton, Zheng Wen, Azin Ashkan, and Csaba Szepesvári. Tight regret bounds for stochastic combinatorial semi-bandits. In *Artificial Intelligence and Statistics*, 2015.
- [23] John Langford and Tong Zhang. The epoch-greedy algorithm for multi-armed bandits with side information. In *Advances in Neural Information Processing Systems*, 2008.
- [24] John Langford, Alexander Strehl, and Jennifer Wortman. Exploration scavenging. In *International Conference on Machine Learning*, 2008.
- [25] Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *International Conference on World Wide Web*, 2010.
- [26] Lihong Li, Wei Chu, John Langford, and Xuanhui Wang. Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms. In *International Conference on Web Search and Data Mining*, 2011.
- [27] Lihong Li, Imed Zitouni, and Jin Young Kim. Toward predicting the outcome of an a/b experiment for search relevance. In *International Conference on Web Search and Data Mining*, 2015.
- [28] Kaare Brandt Petersen, Michael Syskind Pedersen, et al. The matrix cookbook. *Technical University of Denmark*, 2008.
- [29] Lijing Qin, Shouyuan Chen, and Xiaoyan Zhu. Contextual combinatorial bandit and its application on diversified online recommendation. In *International Conference on Data Mining*, 2014.
- [30] Tao Qin and Tie-Yan Liu. Introducing LETOR 4.0 datasets. *arXiv:1306.2597*, 2013.

- [31] Paat Rusmevichientong and John N Tsitsiklis. Linearly parameterized bandits. *Mathematics of Operations Research*, 2010.
- [32] Adith Swaminathan and Thorsten Joachims. Counterfactual risk minimization: Learning from logged bandit feedback. In *International Conference on Machine Learning*, 2015.
- [33] Niek Tax, Sander Bockting, and Djoerd Hiemstra. A cross-benchmark comparison of 87 learning to rank methods. *Information Processing and Management*, 2015.
- [34] Yu-Xiang Wang, Alekh Agarwal, and Miroslav Dudik. Optimal and adaptive off-policy evaluation in contextual bandits. In *International Conference on Machine Learning*, 2017.
- [35] Yue Wang, Dawei Yin, Luo Jie, Pengyuan Wang, Makoto Yamada, Yi Chang, and Qiaozhu Mei. Beyond ranking: Optimizing whole-page presentation. In *International Conference on Web Search and Data Mining*, pages 103–112, 2016.

A Proof of Proposition 1

Lemma 2. *If Assumption 1 holds and $\mu(\mathbf{s} | x) > 0$, then $V(x, \mathbf{s}) = \mathbf{1}_s^T \mathbf{\Gamma}_{\mu, x}^\dagger \boldsymbol{\theta}_{\mu, x}$.*

Proof. Fix one x for the entirety of the proof. Recall from Sec. 3.1 that

$$V(x, \mathbf{s}) = \mathbf{1}_s^T \boldsymbol{\phi}_x .$$

Let $N = |\text{supp } \mu(\cdot | x)|$ be the size of the support of $\mu(\cdot | x)$ and let $\mathbf{M} \in \{0, 1\}^{N \times m\ell}$ denote the binary matrix with rows $\mathbf{1}_s^T$ for each $\mathbf{s} \in \text{supp } \mu(\cdot | x)$. Thus $\mathbf{M}\boldsymbol{\phi}_x$ is the vector enumerating $V(x, \mathbf{s})$ over \mathbf{s} for which $\mu(\mathbf{s} | x) > 0$. Let $\text{Null}(\mathbf{M})$ denote the null space of \mathbf{M} and $\mathbf{\Pi}$ be the projection on $\text{Null}(\mathbf{M})$. Let $\boldsymbol{\phi}_x^* = (\mathbf{I} - \mathbf{\Pi})\boldsymbol{\phi}_x$. Then clearly, $\mathbf{M}\boldsymbol{\phi}_x = \mathbf{M}\boldsymbol{\phi}_x^*$, and hence, for any $\mathbf{s} \in \text{supp } \mu(\cdot | x)$,

$$V(x, \mathbf{s}) = \mathbf{1}_s^T \boldsymbol{\phi}_x^* . \quad (9)$$

We will now show that $\boldsymbol{\phi}_x^* = \mathbf{\Gamma}_{\mu, x}^\dagger \boldsymbol{\theta}_{\mu, x}$, which will complete the proof.

Recall from Sec. 3.1 that

$$\boldsymbol{\theta}_{\mu, x} = \mathbf{\Gamma}_{\mu, x} \boldsymbol{\phi}_x . \quad (10)$$

Next note that $\mathbf{\Gamma}_{\mu, x}$ is symmetric positive semidefinite by definition, so

$$\text{Null}(\mathbf{\Gamma}_{\mu, x}) = \{\mathbf{v} : \mathbf{v}^T \mathbf{\Gamma}_{\mu, x} \mathbf{v} = 0\} = \{\mathbf{v} : \mathbf{1}_s^T \mathbf{v} = 0 \text{ for all } \mathbf{s} \in \text{supp } \mu(\cdot | x)\} = \text{Null}(\mathbf{M})$$

where the first step follows by positive semi definiteness of $\mathbf{\Gamma}_{\mu, x}$, the second step is from the definition of $\mathbf{\Gamma}_{\mu, x}$, and the final step from the definition of \mathbf{M} . Since $\text{Null}(\mathbf{\Gamma}_{\mu, x}) = \text{Null}(\mathbf{M})$, we have from Eq. (10) that $\boldsymbol{\theta}_x = \mathbf{\Gamma}_{\mu, x} \boldsymbol{\phi}_x^*$, but, importantly, this also implies $\boldsymbol{\phi}_x^* \perp \text{Null}(\mathbf{\Gamma}_{\mu, x})$, so by the definition of the pseudoinverse,

$$\mathbf{\Gamma}_{\mu, x}^\dagger \boldsymbol{\theta}_x = \boldsymbol{\phi}_x^* .$$

This proves Lemma 2, since for any \mathbf{s} with $\mu(\mathbf{s} | x) > 0$, we argued that $V(x, \mathbf{s}) = \mathbf{1}_s^T \boldsymbol{\phi}_x^* = \mathbf{1}_s^T \mathbf{\Gamma}_{\mu, x}^\dagger \boldsymbol{\theta}_x$. \square

Proof of Prop. 1. Note that it suffices to analyze the expectation of a single term in the estimator, that is

$$\sum_{\mathbf{s} \in S} \pi(\mathbf{s} | x_i) \mathbf{1}_s^T \mathbf{\Gamma}_{\mu, x_i}^\dagger \hat{\boldsymbol{\theta}}_i .$$

First note that $\mathbb{E}_{(\mathbf{s}_i, r_i) \sim \mu(\cdot, \cdot | x_i)} [\hat{\boldsymbol{\theta}}_i] = \boldsymbol{\theta}_{x_i}$, because

$$\mathbb{E}_{(\mathbf{s}_i, r_i) \sim \mu(\cdot, \cdot | x_i)} [\hat{\boldsymbol{\theta}}_i(j, a)] = \mathbb{E}_{(\mathbf{s}_i, r_i) \sim \mu(\cdot, \cdot | x_i)} [r_i \mathbf{1}\{s_j = a\}] = \boldsymbol{\theta}_{x_i}(j, a) .$$

The remainder follows by Lemma 2:

$$\begin{aligned} \mathbb{E} \left[\sum_{\mathbf{s} \in S} \pi(\mathbf{s} | x_i) \mathbf{1}_s^T \mathbf{\Gamma}_{\mu, x_i}^\dagger \hat{\boldsymbol{\theta}}_i \right] &= \mathbb{E}_{x_i \sim D} \left[\sum_{\mathbf{s} \in S} \pi(\mathbf{s} | x_i) \mathbf{1}_s^T \mathbf{\Gamma}_{\mu, x_i}^\dagger \mathbb{E}_{(\mathbf{s}_i, r_i) \sim \mu(\cdot, \cdot | x_i)} [\hat{\boldsymbol{\theta}}_i] \right] \\ &= \mathbb{E}_{x_i \sim D} \left[\sum_{\mathbf{s} \in S} \pi(\mathbf{s} | x_i) \mathbf{1}_s^T \mathbf{\Gamma}_{\mu, x_i}^\dagger \boldsymbol{\theta}_{x_i} \right] \\ &= \mathbb{E}_{x_i \sim D} \left[\sum_{\mathbf{s} \in S} \pi(\mathbf{s} | x_i) V(x_i, \mathbf{s}) \right] = V(\pi) . \quad \square \end{aligned}$$

B Proof of Theorem 1

Proof. The proof is based on an application of Bernstein's inequality to the centered sum

$$\sum_{i=1}^n \left[\mathbf{q}_{\pi, x_i}^T \mathbf{\Gamma}_{\mu, x_i}^\dagger \hat{\boldsymbol{\theta}}_i - V(\pi) \right] .$$

The fact that this quantity is centered is directly from Prop. 1. We must compute both the second moment and the range to apply Bernstein's inequality. By independence, we can focus on just one term, so we will drop the subscript i . First, bound the variance:

$$\begin{aligned}
\text{Var} \left[\mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \hat{\boldsymbol{\theta}} \right] &\leq \mathbb{E}_\mu \left[\left(\mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \hat{\boldsymbol{\theta}} \right)^2 \right] \\
&= \mathbb{E}_\mu \left[\left(\mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger r \mathbf{1}_s \right)^2 \right] \\
&\leq \mathbb{E}_\mu \left[\left(\mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{1}_s \right)^2 \right] \\
&= \mathbb{E}_{x \sim D} \left[\mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbb{E}_{\mathbf{s} \sim \mu(\cdot|x)} \left[\mathbf{1}_s \mathbf{1}_s^T \right] \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{q}_{\pi,x} \right] \\
&= \mathbb{E}_{x \sim D} \left[\mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{\Gamma}_{\mu,x} \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{q}_{\pi,x} \right] \\
&= \mathbb{E}_{x \sim D} \left[\mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{q}_{\pi,x} \right] \\
&= \sigma^2 .
\end{aligned}$$

Thus the per-term variance is at most σ^2 . We now bound the range, again focusing on one term,

$$\begin{aligned}
\left| \mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \hat{\boldsymbol{\theta}} - V(\pi) \right| &\leq \left| \mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \hat{\boldsymbol{\theta}} \right| + 1 \\
&= \left| \mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger r \mathbf{1}_s \right| + 1 \\
&\leq \left| \mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{1}_s \right| + 1 \\
&\leq \rho + 1
\end{aligned}$$

The first line here is the triangle inequality, coupled with the fact that since rewards are bounded in $[-1, 1]$, so is $V(\pi)$. The second line is from the definition of $\hat{\boldsymbol{\theta}}$, while the third follows because $r \in [-1, 1]$. The final line follows from the definition of ρ .

Now, we may apply Bernstein's inequality, which says that for any $\delta \in (0, 1)$, with probability at least $1 - \delta$,

$$\left| \sum_{i=1}^n \left[\mathbf{q}_{\pi,x_i}^T \mathbf{\Gamma}_{\mu,x_i}^\dagger \hat{\boldsymbol{\theta}}_i - V(\pi) \right] \right| \leq \sqrt{2n\sigma^2 \ln(2/\delta)} + \frac{2(\rho+1) \ln(2/\delta)}{3} .$$

The theorem follows by dividing by n . □

C Pseudo-inverse estimator when $\pi = \mu$

In this section we show that when the target policy coincides with logging (i.e., $\pi = \mu$), we have $\sigma^2 = \rho = 1$, i.e., the bound of Theorem 1 is independent of the number of actions and slots. Indeed, in Claim 2 we will see that the estimator actually simplifies to taking an empirical average of rewards which are bounded in $[-1, 1]$. Before proving Claim 2 we prove one supporting claim:

Claim 1. *For any policy μ and context x , we have $\mathbf{q}_{\mu,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{1}_s = 1$ for all $s \in \text{supp } \mu(\cdot | x)$.*

Proof. To simplify the exposition, write \mathbf{q} and $\mathbf{\Gamma}$ instead of a more verbose $\mathbf{q}_{\mu,x}$ and $\mathbf{\Gamma}_{\mu,x}$.

The bulk of the proof is in deriving an explicit expression for $\mathbf{\Gamma}^\dagger$. We begin by expressing $\mathbf{\Gamma}$ in a suitable basis. Since $\mathbf{\Gamma}$ is the matrix of second moments and \mathbf{q} is the vector of first moments of $\mathbf{1}_s$, the matrix $\mathbf{\Gamma}$ can be written as

$$\mathbf{\Gamma} = \mathbf{V} + \mathbf{q}\mathbf{q}^T$$

where \mathbf{V} is the covariance matrix of $\mathbf{1}_s$, i.e., $\mathbf{V} := \mathbb{E}_{\mathbf{s} \sim \mu(\cdot|x)} \left[(\mathbf{1}_s - \mathbf{q})(\mathbf{1}_s - \mathbf{q})^T \right]$. Assume that the rank of \mathbf{V} is r and consider the eigenvalue decomposition of \mathbf{V}

$$\mathbf{V} = \sum_{i=1}^r \lambda_i \mathbf{u}_i \mathbf{u}_i^T = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T ,$$

where $\lambda_i > 0$ and vectors \mathbf{u}_i are orthonormal; we have grouped the eigenvalues into the diagonal matrix $\mathbf{\Lambda} := \text{diag}(\lambda_1, \dots, \lambda_r)$ and eigenvectors into the matrix $\mathbf{U} := (\mathbf{u}_1 \ \mathbf{u}_2 \ \dots \ \mathbf{u}_r)$.

We next argue that $\mathbf{q} \notin \text{Range}(\mathbf{V})$. To see this, note that the all-ones-vector $\mathbf{1}$ is in the null space of \mathbf{V} because, for any valid slate \mathbf{s} , we have $\mathbf{1}_s^T \mathbf{1} = \ell$ and thus also for the convex combination \mathbf{q} we have $\mathbf{q}^T \mathbf{1} = \ell$, which means that

$$\mathbf{1}^T \mathbf{V} \mathbf{1} = \mathbb{E}_{\mathbf{s} \sim \mu(\cdot|x)} [\mathbf{1}^T (\mathbf{1}_s - \mathbf{q})(\mathbf{1}_s - \mathbf{q})^T \mathbf{1}] = 0 .$$

Now, since $\mathbf{1} \perp \text{Range}(\mathbf{V})$ and $\mathbf{q}^T \mathbf{1} = \ell$, we have that $\mathbf{q} \notin \text{Range}(\mathbf{V})$. In particular, we can write \mathbf{q} in the form

$$\mathbf{q} = \sum_{i=1}^r \beta_i \mathbf{u}_i + \alpha \mathbf{n} = (\mathbf{U} \ \mathbf{n}) \begin{pmatrix} \boldsymbol{\beta} \\ \alpha \end{pmatrix} \quad (11)$$

where $\alpha \neq 0$ and $\mathbf{n} \in \text{Null}(\mathbf{V})$ is a unit vector. Note that $\mathbf{n} \perp \mathbf{u}_i$ since $\mathbf{u}_i \perp \text{Null}(\mathbf{V})$. Thus, the second moment matrix $\boldsymbol{\Gamma}$ can be written as

$$\boldsymbol{\Gamma} = \mathbf{V} + \mathbf{q}\mathbf{q}^T = (\mathbf{U} \ \mathbf{n}) \begin{pmatrix} \boldsymbol{\Lambda} + \boldsymbol{\beta}\boldsymbol{\beta}^T & \alpha\boldsymbol{\beta} \\ \alpha\boldsymbol{\beta}^T & \alpha^2 \end{pmatrix} (\mathbf{U} \ \mathbf{n})^T . \quad (12)$$

Let $\mathbf{Q} \in \mathbb{R}^{(r+1) \times (r+1)}$ denote the middle matrix in the factorization of Eq. (12):

$$\mathbf{Q} := \begin{pmatrix} \boldsymbol{\Lambda} + \boldsymbol{\beta}\boldsymbol{\beta}^T & \alpha\boldsymbol{\beta} \\ \alpha\boldsymbol{\beta}^T & \alpha^2 \end{pmatrix} . \quad (13)$$

This matrix is a representation of $\boldsymbol{\Gamma}$ with respect to the basis $\{\mathbf{u}_1, \dots, \mathbf{u}_r, \mathbf{n}\}$. Since $\mathbf{q} \notin \text{Range}(\mathbf{V})$, the rank of $\boldsymbol{\Gamma}$ and that of \mathbf{Q} is $r + 1$. Thus, \mathbf{Q} is invertible and

$$\boldsymbol{\Gamma}^\dagger = (\mathbf{U} \ \mathbf{n}) \mathbf{Q}^{-1} (\mathbf{U} \ \mathbf{n})^T . \quad (14)$$

To obtain \mathbf{Q}^{-1} , we use the following identity (see [28]):

$$\begin{pmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{pmatrix}^{-1} = \begin{pmatrix} \mathbf{M}^{-1} & -\mathbf{M}^{-1}\mathbf{A}_{12}\mathbf{A}_{22}^{-1} \\ -\mathbf{A}_{22}^{-1}\mathbf{A}_{21}\mathbf{M}^{-1} & \mathbf{A}_{22}^{-1}\mathbf{A}_{21}\mathbf{M}^{-1}\mathbf{A}_{12}\mathbf{A}_{22}^{-1} + \mathbf{A}_{22}^{-1} \end{pmatrix} , \quad (15)$$

where $\mathbf{M} := \mathbf{A}_{11} - \mathbf{A}_{12}\mathbf{A}_{22}^{-1}\mathbf{A}_{21}$ is the Schur complement of \mathbf{A}_{22} . The identity of Eq. (15) holds whenever \mathbf{A}_{22} and its Schur complement \mathbf{M} are both invertible. In the block representation of Eq. (13), we have $\mathbf{A}_{22} = \alpha^2 \neq 0$ and

$$\mathbf{M} = (\boldsymbol{\Lambda} + \boldsymbol{\beta}\boldsymbol{\beta}^T) - (\alpha\boldsymbol{\beta})\alpha^{-2}(\alpha\boldsymbol{\beta}^T) = \boldsymbol{\Lambda} ,$$

so Eq. (15) can be applied to obtain \mathbf{Q}^{-1} :

$$\begin{aligned} \mathbf{Q}^{-1} &= \begin{pmatrix} \boldsymbol{\Lambda} + \boldsymbol{\beta}\boldsymbol{\beta}^T & \alpha\boldsymbol{\beta} \\ \alpha\boldsymbol{\beta}^T & \alpha^2 \end{pmatrix}^{-1} = \begin{pmatrix} \boldsymbol{\Lambda}^{-1} & -\boldsymbol{\Lambda}^{-1}(\alpha\boldsymbol{\beta})\alpha^{-2} \\ -\alpha^{-2}(\alpha\boldsymbol{\beta}^T)\boldsymbol{\Lambda}^{-1} & \alpha^{-2}(\alpha\boldsymbol{\beta}^T)\boldsymbol{\Lambda}^{-1}(\alpha\boldsymbol{\beta})\alpha^{-2} + \alpha^{-2} \end{pmatrix} \\ &= \begin{pmatrix} \boldsymbol{\Lambda}^{-1} & -\alpha^{-1}\boldsymbol{\Lambda}^{-1}\boldsymbol{\beta} \\ -\alpha^{-1}\boldsymbol{\beta}^T\boldsymbol{\Lambda}^{-1} & \alpha^{-2}(1 + \boldsymbol{\beta}^T\boldsymbol{\Lambda}^{-1}\boldsymbol{\beta}) \end{pmatrix} . \end{aligned} \quad (16)$$

Next, we will evaluate $\boldsymbol{\Gamma}^\dagger \mathbf{q}$, using the factorizations in Eqs. (14) and (11), and substituting Eq. (16) for \mathbf{Q}^{-1} :

$$\begin{aligned} \boldsymbol{\Gamma}^\dagger \mathbf{q} &= (\mathbf{U} \ \mathbf{n}) \mathbf{Q}^{-1} (\mathbf{U} \ \mathbf{n})^T (\mathbf{U} \ \mathbf{n}) \begin{pmatrix} \boldsymbol{\beta} \\ \alpha \end{pmatrix} \\ &= (\mathbf{U} \ \mathbf{n}) \mathbf{Q}^{-1} \begin{pmatrix} \boldsymbol{\beta} \\ \alpha \end{pmatrix} \\ &= (\mathbf{U} \ \mathbf{n}) \begin{pmatrix} \boldsymbol{\Lambda}^{-1}\boldsymbol{\beta} - \boldsymbol{\Lambda}^{-1}\boldsymbol{\beta} \\ -\alpha^{-1}\boldsymbol{\beta}^T\boldsymbol{\Lambda}^{-1}\boldsymbol{\beta} + \alpha^{-1}(1 + \boldsymbol{\beta}^T\boldsymbol{\Lambda}^{-1}\boldsymbol{\beta}) \end{pmatrix} \\ &= (\mathbf{U} \ \mathbf{n}) \begin{pmatrix} \mathbf{0} \\ \alpha^{-1} \end{pmatrix} \\ &= \alpha^{-1} \mathbf{n} . \end{aligned}$$

To finish the proof, we consider any $\mathbf{s} \in \text{supp } \mu(\cdot|x)$ and consider the decomposition of $\mathbf{1}_s$ in the basis $\{\mathbf{u}_1, \dots, \mathbf{u}_r, \mathbf{n}\}$. First, note that $(\mathbf{1}_s - \mathbf{q}) \perp \text{Null}(\mathbf{V})$ since

$$\text{Null}(\mathbf{V}) = \{ \mathbf{v} : \mathbb{E}_{\mathbf{s} \sim \mu(\cdot|x)} [((\mathbf{1}_s - \mathbf{q})^T \mathbf{v})^2] = 0 \} = \{ \mathbf{v} : (\mathbf{1}_s - \mathbf{q})^T \mathbf{v} = 0 \text{ for all } \mathbf{s} \in \text{supp } \mu(\cdot|x) \} .$$

Thus, $(\mathbf{1}_s - \mathbf{q}) \in \text{Range}(\mathbf{V})$. Therefore, we obtain

$$\mathbf{q}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{1}_s = \alpha^{-1} \mathbf{n}^T \mathbf{1}_s = \alpha^{-1} \mathbf{n}^T (\mathbf{1}_s - \mathbf{q}) + \alpha^{-1} \mathbf{n}^T \mathbf{q} = 0 + \alpha^{-1} \alpha = 1 ,$$

where the third equality follows because $(\mathbf{1}_s - \mathbf{q}) \perp \mathbf{n}$ and the decomposition in Eq. (11) shows that $\mathbf{n}^T \mathbf{q} = \alpha$. \square

Claim 2. If $\pi = \mu$ then $\sigma^2 = \rho = 1$ and $\hat{V}_{\text{PI}}(\pi) = \hat{V}_{\text{PI}}(\mu) = \frac{1}{n} \sum_{i=1}^n r_i$.

Proof. From Claim 1

$$\mathbf{q}_{\mu,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{q}_{\mu,x} = \mathbb{E}_{\mathbf{s} \sim \mu(\cdot|x)} [\mathbf{q}_{\mu,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{1}_s] = 1 .$$

Taking expectation over x then yields $\sigma^2 = 1$. Equality $\rho = 1$ follows immediately from plugging Claim 1 into the definition of ρ . The final statement of Claim 2 follows by applying Claim 1 to a single term of $\hat{V}_{\text{PI}}(\mu)$:

$$\mathbf{q}_{\mu,x_i}^T \mathbf{\Gamma}_{\mu,x_i}^\dagger r_i \mathbf{1}_{s_i} = r_i . \quad \square$$

D A product slate space under a product logging distribution

Proposition 2. Consider the product slate space where $S(x) = A_1(x) \times \dots \times A_\ell(x)$ and assume that the logging policy picks any $\mathbf{s} \in S(x)$ with non-zero probability and factorizes across the slots as $\mu(\mathbf{s} | x) = \prod_j \mu(s_j | x)$. For any policy π , any $\mathbf{s} \in S(x)$, and any $r \in [-1, 1]$ we then have

$$\mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger r \mathbf{1}_s = r \cdot \left[\sum_{j=1}^{\ell} \frac{\pi(s_j | x)}{\mu(s_j | x)} - \ell + 1 \right] . \quad (17)$$

Proof. The proof uses Claim 1 and the identities introduced in its proof. As in the proof of Claim 1, write \mathbf{q} and $\mathbf{\Gamma}$ instead of $\mathbf{q}_{\mu,x}$ and $\mathbf{\Gamma}_{\mu,x}$, and let $\mathbf{V} := \mathbb{E}_{\mathbf{s} \sim \mu(\cdot|x)} [(\mathbf{1}_s - \mathbf{q})(\mathbf{1}_s - \mathbf{q})^T]$. Thus, $\mathbf{\Gamma} = \mathbf{V} + \mathbf{q}\mathbf{q}^T$. It suffices to show that for any $\mathbf{s}, \mathbf{s}' \in S(x)$,

$$\mathbf{1}_{s'}^T \mathbf{\Gamma}^\dagger \mathbf{1}_s = \sum_{j=1}^{\ell} \frac{\mathbf{1}\{s'_j = s_j\}}{\mu(s_j | x)} - \ell + 1 . \quad (18)$$

Pick $\mathbf{s}, \mathbf{s}' \in S(x) = \text{supp } \mu(\cdot | x)$. By Claim 1, we have $\mathbf{q}^T \mathbf{\Gamma}^\dagger \mathbf{1}_s = \mathbf{1}_{s'}^T \mathbf{\Gamma}^\dagger \mathbf{q} = \mathbf{q}^T \mathbf{\Gamma}^\dagger \mathbf{q} = 1$, so

$$\begin{aligned} \mathbf{1}_{s'}^T \mathbf{\Gamma}^\dagger \mathbf{1}_s &= (\mathbf{1}_{s'} - \mathbf{q})^T \mathbf{\Gamma}^\dagger (\mathbf{1}_s - \mathbf{q}) + \underbrace{\mathbf{q}^T \mathbf{\Gamma}^\dagger \mathbf{1}_s}_{=1} + \underbrace{\mathbf{1}_{s'}^T \mathbf{\Gamma}^\dagger \mathbf{q}}_{=1} - \underbrace{\mathbf{q}^T \mathbf{\Gamma}^\dagger \mathbf{q}}_{=1} \\ &= (\mathbf{1}_{s'} - \mathbf{q})^T \mathbf{\Gamma}^\dagger (\mathbf{1}_s - \mathbf{q}) + 1 . \end{aligned} \quad (19)$$

Similar to the reasoning at the end of the proof of Claim 1, we know that $(\mathbf{1}_s - \mathbf{q}) \in \text{Range}(\mathbf{V})$ and $(\mathbf{1}_{s'} - \mathbf{q}) \in \text{Range}(\mathbf{V})$. The factorization of $\mathbf{\Gamma}^\dagger$ in Eqs. (14) and (16) therefore yields

$$\begin{aligned} (\mathbf{1}_{s'} - \mathbf{q})^T \mathbf{\Gamma}^\dagger (\mathbf{1}_s - \mathbf{q}) &= (\mathbf{1}_{s'} - \mathbf{q})^T (\mathbf{U}\mathbf{\Lambda}^{-1}\mathbf{U}^T) (\mathbf{1}_s - \mathbf{q}) \\ &= (\mathbf{1}_{s'} - \mathbf{q})^T \mathbf{V}^\dagger (\mathbf{1}_s - \mathbf{q}) , \end{aligned} \quad (20)$$

where the last step follows from the fact that $\mathbf{V} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$, and so $\mathbf{V}^\dagger = \mathbf{U}\mathbf{\Lambda}^{-1}\mathbf{U}^T$.

To finish the proof, we study the structure of \mathbf{V} and \mathbf{V}^\dagger . First, let \mathbf{q}_j denote the block of \mathbf{q} corresponding to the j th slot. Its a th entry corresponds to the probability $\mu(s_j = a | x)$. Since the values s_j are conditionally independent, conditioned on x , the covariance matrix \mathbf{V} takes form $\mathbf{V} = \text{diag}_{j=1,\dots,\ell} \mathbf{V}_j$, where $\mathbf{V}_j = (\text{diag}_{a \in A_j(x)} q_{j,a}) - \mathbf{q}_j \mathbf{q}_j^T$ is the covariance matrix of the multinomial distribution described by \mathbf{q}_j . Thus,

$$\mathbf{V}^\dagger = \text{diag}_{j=1,\dots,\ell} \mathbf{V}_j^\dagger . \quad (21)$$

It can be directly verified that the pseudoinverse of \mathbf{V}_j takes form

$$\mathbf{V}_j^\dagger = \mathbf{P}_j (\text{diag}_{a \in A_j(x)} q_{j,a}^{-1}) \mathbf{P}_j , \quad (22)$$

where $\mathbf{P}_j := \mathbf{I}_j - \mathbf{1}_j \mathbf{1}_j^T / m_j$, and \mathbf{I}_j is the $m_j \times m_j$ identity matrix, and $\mathbf{1}_j$ the m_j -dimensional all-ones vector. To verify that Eq. (22) holds, first note that \mathbf{P}_j is the projection matrix on $\text{Range}(\mathbf{V}_j)$. Then set $\mathbf{V}'_j := \mathbf{P}_j (\text{diag}_{a \in A_j(x)} q_{j,a}^{-1}) \mathbf{P}_j$, and directly verify that $\mathbf{V}'_j \mathbf{V}_j = \mathbf{P}_j$ and $\mathbf{V}_j \mathbf{V}'_j = \mathbf{P}_j$. The first identity can be verified as follows:

$$\mathbf{V}'_j \mathbf{V}_j = \mathbf{P}_j (\text{diag}_{a \in A_j(x)} q_{j,a}^{-1}) \mathbf{P}_j \mathbf{V}_j = \mathbf{P}_j (\text{diag}_{a \in A_j(x)} q_{j,a}^{-1}) \mathbf{V}_j = \mathbf{P}_j (\mathbf{I}_j - \mathbf{1}_j \mathbf{q}_j^T) = \mathbf{P}_j .$$

The second identity follows similarly.

Combining Eqs. (20), (21) and (22) yields

$$\begin{aligned} (\mathbf{1}_{s'} - \mathbf{q})^T \mathbf{V}^\dagger (\mathbf{1}_s - \mathbf{q}) &= \sum_{j=1}^{\ell} (\mathbf{1}_{s'_j} - \mathbf{q}_j)^T \mathbf{V}_j^\dagger (\mathbf{1}_{s_j} - \mathbf{q}_j) \\ &= \sum_{j=1}^{\ell} (\mathbf{1}_{s'_j} - \mathbf{q}_j)^T \mathbf{P}_j (\text{diag}_{a \in A_j(x)} q_{j,a}^{-1}) \mathbf{P}_j (\mathbf{1}_{s_j} - \mathbf{q}_j) \\ &= \sum_{j=1}^{\ell} (\mathbf{1}_{s'_j} - \mathbf{q}_j)^T (\text{diag}_{a \in A_j(x)} q_{j,a}^{-1}) (\mathbf{1}_{s_j} - \mathbf{q}_j) \\ &= \sum_{j=1}^{\ell} (\mathbf{1}\{s'_j = s_j\} q_{j,s_j}^{-1} - 1) \\ &= \sum_{j=1}^{\ell} \left(\frac{\mathbf{1}\{s'_j = s_j\}}{\mu(s_j | x)} - 1 \right) . \end{aligned}$$

Plugging this back into Eq. (19) then proves Eq. (18). \square

E Proof of Corollary 1

For a given logging policy μ and context x , let

$$\bar{\rho}_{\mu,x} := \sup_{\mathbf{s} \in \text{supp } \mu(\cdot | x)} \mathbf{1}_s^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{1}_s .$$

This quantity can be viewed as a norm of $\mathbf{\Gamma}_{\mu,x}^\dagger$ with respect to the set of slates chosen by μ with non-zero probability. It can be used to bound σ^2 and ρ , and thus to bound an error of \hat{V}_{PI} :

Proposition 3. *For any logging policy μ and target policy π that is absolutely continuous with respect to μ , we have*

$$\sigma^2 \leq \rho \leq \sup_x \bar{\rho}_{\mu,x} .$$

Proof. Recall that

$$\sigma^2 = \mathbb{E}_{x \sim D} [\mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{q}_{\pi,x}] , \quad \rho = \sup_x \sup_{\mathbf{s} \in \text{supp } \mu(\cdot | x)} |\mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{1}_s| .$$

To see that $\sigma^2 \leq \rho$ note that

$$\mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{q}_{\pi,x} = \mathbb{E}_{\mathbf{s} \sim \pi(\cdot | x)} [\mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{1}_s] \leq \rho$$

where the last inequality follows by the absolute continuity of π with respect to μ . It remains to show that $\rho \leq \sup_x \bar{\rho}_{\mu,x}$.

First, by positive semi-definiteness of $\mathbf{\Gamma}_{\mu,x}^\dagger$ and from the definition of $\bar{\rho}_{\mu,x}$, we have that for any slates $\mathbf{s}, \mathbf{s}' \in \text{supp } \mu(\cdot | x)$ and any $z \in \{-1, 1\}$

$$z \mathbf{1}_{\mathbf{s}'}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{1}_s \leq \frac{\mathbf{1}_s^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{1}_s + \mathbf{1}_{\mathbf{s}'}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{1}_{\mathbf{s}'}}{2} \leq \max\{\mathbf{1}_s^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{1}_s, \mathbf{1}_{\mathbf{s}'}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{1}_{\mathbf{s}'}\} \leq \bar{\rho}_{\mu,x} .$$

Therefore, for any π absolutely continuous with respect to μ and any $\mathbf{s} \in \text{supp } \mu(\cdot | x)$, we have

$$|\mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{1}_s| = \max_{z \in \{-1, 1\}} \mathbb{E}_{\mathbf{s}' \sim \pi(\cdot | x)} [z \mathbf{1}_{\mathbf{s}'}^T \mathbf{\Gamma}_{\mu,x}^\dagger \mathbf{1}_s] \leq \bar{\rho}_{\mu,x} .$$

Taking a supremum over x and $\mathbf{s} \in \text{supp } \mu(\cdot | x)$, we obtain $\rho \leq \sup_x \bar{\rho}_{\mu,x}$. \square

We next derive bounds on $\bar{\rho}_{\mu,x}$ for uniformly-random policies in the ranking example. Then we prove a translation theorem, which allows translating the bound for uniform distributions into a bound for the ε -uniform distributions. Finally, we put these results together to prove Corollary 1.

E.1 Uniform logging distribution over rankings

Let $\mathbf{1}_j \in \mathbb{R}^{\ell m}$ be the vector that is all-ones on the actions in the j -th position and zeros elsewhere. Similarly, let $\mathbf{1}_a \in \mathbb{R}^{\ell m}$ be the vector that is all-ones on the action a in all positions and zeros elsewhere. Finally, let $\mathbf{1} \in \mathbb{R}^{\ell m}$ be the all-ones vector. We also use $\mathbf{I}_j = \text{diag}(\mathbf{1}_j)$ to denote the diagonal matrix with all-ones on the actions in the j -th position and zeros elsewhere.

Proposition 4. *Consider the ranking setting where for each x there is a set $A(x)$ such that $A_j(x) = A(x)$ and where all slates $\mathbf{s} \in A(x)^\ell$ without repetitions are legal. Let ν denote the uniform logging policy over these slates. If $\ell < m$, then $\bar{\rho}_{\nu,x} = m\ell - \ell + 1$ and*

$$\mathbf{\Gamma}_{\nu,x}^\dagger = \left(\frac{1}{\ell^2} - \frac{m-1}{m(m-\ell)} \right) \cdot \mathbf{1}\mathbf{1}^T + (m-1)\mathbf{I} - \frac{m-1}{m} \sum_j \mathbf{1}_j \mathbf{1}_j^T + \frac{m-1}{m-\ell} \sum_a \mathbf{1}_a \mathbf{1}_a^T,$$

and for $\ell = m$, we have $\bar{\rho}_{\nu,x} = m^2 - 2m + 2$ and

$$\mathbf{\Gamma}_{\nu,x}^\dagger = \frac{1}{m} \cdot \mathbf{1}\mathbf{1}^T + (m-1)\mathbf{I} - \frac{m-1}{m} \sum_j \mathbf{1}_j \mathbf{1}_j^T - \frac{m-1}{m} \sum_a \mathbf{1}_a \mathbf{1}_a^T.$$

For $\ell = m$, we have for any policy π , any $\mathbf{s} \in S(x)$, and any $r \in [-1, 1]$ that

$$\mathbf{q}_{\pi,x}^T \mathbf{\Gamma}_{\nu,x}^\dagger r \mathbf{1}_s = r \cdot \left[\sum_{j=1}^{\ell} \frac{\pi(s_j | x)}{1/(m-1)} - m + 2 \right]. \quad (23)$$

Proof. Throughout the proof we will write $\mathbf{\Gamma}$ instead of the more verbose $\mathbf{\Gamma}_{\nu,x}$. Note that for ranking and the uniform distribution we have

$$\Gamma(j, a; k, a') = \begin{cases} \frac{1}{m} & \text{if } j = k \text{ and } a = a' \\ \frac{1}{m(m-1)} & \text{if } j \neq k \text{ and } a \neq a' \\ 0 & \text{otherwise.} \end{cases}$$

Thus, for any \mathbf{z}

$$\begin{aligned} \mathbf{z}^T \mathbf{\Gamma} \mathbf{z} &= \sum_{j,a} \frac{z_{j,a}^2}{m} + \frac{1}{m(m-1)} \sum_{j \neq k, a \neq a'} z_{j,a} z_{k,a'} \\ &= \frac{1}{m} \|\mathbf{z}\|_2^2 + \frac{1}{m(m-1)} \left((\mathbf{z}^T \mathbf{1})^2 - \sum_j (\mathbf{z}^T \mathbf{1}_j)^2 - \sum_a (\mathbf{z}^T \mathbf{1}_a)^2 + \|\mathbf{z}\|_2^2 \right) \\ &= \frac{1}{m(m-1)} \left((\mathbf{z}^T \mathbf{1})^2 - \sum_j (\mathbf{z}^T \mathbf{1}_j)^2 - \sum_a (\mathbf{z}^T \mathbf{1}_a)^2 + m \|\mathbf{z}\|_2^2 \right). \end{aligned} \quad (24)$$

Let $\mathbf{1}_{\mathcal{J}} \in \mathbb{R}^\ell$ and $\mathbf{1}_{\mathcal{A}} \in \mathbb{R}^m$ be all-ones vectors in the respective spaces and $\mathbf{I}_{\mathcal{J}} \in \mathbb{R}^{\ell \times \ell}$ and $\mathbf{I}_{\mathcal{A}} \in \mathbb{R}^{m \times m}$ be identity matrices in the respective spaces. We can rewrite the quadratic form

described by Γ as

$$\begin{aligned}
m(m-1)\Gamma &= \mathbf{1}\mathbf{1}^T - \sum_j \mathbf{1}_j \mathbf{1}_j^T - \sum_a \mathbf{1}_a \mathbf{1}_a^T + m\mathbf{I} \\
&= (\mathbf{1}_{\mathcal{J}} \mathbf{1}_{\mathcal{J}}^T) \otimes (\mathbf{1}_{\mathcal{A}} \mathbf{1}_{\mathcal{A}}^T) - \mathbf{I}_{\mathcal{J}} \otimes (\mathbf{1}_{\mathcal{A}} \mathbf{1}_{\mathcal{A}}^T) - (\mathbf{1}_{\mathcal{J}} \mathbf{1}_{\mathcal{J}}^T) \otimes \mathbf{I}_{\mathcal{A}} + m \cdot \mathbf{I}_{\mathcal{J}} \otimes \mathbf{I}_{\mathcal{A}} \\
&= \ell m \cdot \frac{\mathbf{1}_{\mathcal{J}} \mathbf{1}_{\mathcal{J}}^T}{\ell} \otimes \frac{\mathbf{1}_{\mathcal{A}} \mathbf{1}_{\mathcal{A}}^T}{m} - m \cdot \mathbf{I}_{\mathcal{J}} \otimes \frac{\mathbf{1}_{\mathcal{A}} \mathbf{1}_{\mathcal{A}}^T}{m} - \ell \cdot \frac{\mathbf{1}_{\mathcal{J}} \mathbf{1}_{\mathcal{J}}^T}{\ell} \otimes \mathbf{I}_{\mathcal{A}} + m \cdot \mathbf{I}_{\mathcal{J}} \otimes \mathbf{I}_{\mathcal{A}} \\
&= \ell(m-1) \cdot \frac{\mathbf{1}_{\mathcal{J}} \mathbf{1}_{\mathcal{J}}^T}{\ell} \otimes \frac{\mathbf{1}_{\mathcal{A}} \mathbf{1}_{\mathcal{A}}^T}{m} - m \cdot \mathbf{I}_{\mathcal{J}} \otimes \left(\frac{\mathbf{1}_{\mathcal{A}} \mathbf{1}_{\mathcal{A}}^T}{m} - \mathbf{I}_{\mathcal{A}} \right) - \ell \cdot \frac{\mathbf{1}_{\mathcal{J}} \mathbf{1}_{\mathcal{J}}^T}{\ell} \otimes \left(\mathbf{I}_{\mathcal{A}} - \frac{\mathbf{1}_{\mathcal{A}} \mathbf{1}_{\mathcal{A}}^T}{m} \right) \\
&= \ell(m-1) \cdot \frac{\mathbf{1}_{\mathcal{J}} \mathbf{1}_{\mathcal{J}}^T}{\ell} \otimes \frac{\mathbf{1}_{\mathcal{A}} \mathbf{1}_{\mathcal{A}}^T}{m} \\
&\quad + m \cdot \left(\mathbf{I}_{\mathcal{J}} - \frac{\mathbf{1}_{\mathcal{J}} \mathbf{1}_{\mathcal{J}}^T}{\ell} \right) \otimes \left(\mathbf{I}_{\mathcal{A}} - \frac{\mathbf{1}_{\mathcal{A}} \mathbf{1}_{\mathcal{A}}^T}{m} \right) + (m-\ell) \cdot \frac{\mathbf{1}_{\mathcal{J}} \mathbf{1}_{\mathcal{J}}^T}{\ell} \otimes \left(\mathbf{I}_{\mathcal{A}} - \frac{\mathbf{1}_{\mathcal{A}} \mathbf{1}_{\mathcal{A}}^T}{m} \right). \quad (25)
\end{aligned}$$

Next, we would like to argue that Eq. (25) is an eigendecomposition. For this, we just need to show that each of the three Kronecker products in Eq. (25) equals a projection matrix in $\mathbb{R}^{\ell m}$, and that the ranges of the projection matrices are orthogonal. The first property follows, because if \mathbf{P}_1 and \mathbf{P}_2 are projection matrices then so is $\mathbf{P}_1 \otimes \mathbf{P}_2$. The second property follows, because for $\mathbf{P}_1, \mathbf{P}'_1$ (square of the same dimension) and $\mathbf{P}_2, \mathbf{P}'_2$ (square of the same dimension) such that either ranges of \mathbf{P}_1 and \mathbf{P}'_1 are orthogonal or ranges of \mathbf{P}_2 and \mathbf{P}'_2 are orthogonal, we obtain that the ranges of $\mathbf{P}_1 \otimes \mathbf{P}_2$ and $\mathbf{P}'_1 \otimes \mathbf{P}'_2$ are orthogonal.

Now we are ready to derive the pseudo-inverse. We distinguish two cases.

Case $\ell < m$: We directly invert the eigenvalues in Eq. (25) to obtain

$$\begin{aligned}
\Gamma^\dagger &= \frac{m}{\ell} \cdot \frac{\mathbf{1}_{\mathcal{J}} \mathbf{1}_{\mathcal{J}}^T}{\ell} \otimes \frac{\mathbf{1}_{\mathcal{A}} \mathbf{1}_{\mathcal{A}}^T}{m} + (m-1) \cdot \left(\mathbf{I}_{\mathcal{J}} - \frac{\mathbf{1}_{\mathcal{J}} \mathbf{1}_{\mathcal{J}}^T}{\ell} \right) \otimes \left(\mathbf{I}_{\mathcal{A}} - \frac{\mathbf{1}_{\mathcal{A}} \mathbf{1}_{\mathcal{A}}^T}{m} \right) \\
&\quad + \frac{m-1}{1-\ell/m} \cdot \frac{\mathbf{1}_{\mathcal{J}} \mathbf{1}_{\mathcal{J}}^T}{\ell} \otimes \left(\mathbf{I}_{\mathcal{A}} - \frac{\mathbf{1}_{\mathcal{A}} \mathbf{1}_{\mathcal{A}}^T}{m} \right) \\
&= \frac{1}{\ell^2} \cdot \mathbf{1}\mathbf{1}^T + (m-1) \cdot \left(\mathbf{I}_{\mathcal{J}} + \frac{\mathbf{1}_{\mathcal{J}} \mathbf{1}_{\mathcal{J}}^T}{m-\ell} \right) \otimes \left(\mathbf{I}_{\mathcal{A}} - \frac{\mathbf{1}_{\mathcal{A}} \mathbf{1}_{\mathcal{A}}^T}{m} \right) \\
&= \left(\frac{1}{\ell^2} - \frac{m-1}{m(m-\ell)} \right) \cdot \mathbf{1}\mathbf{1}^T + (m-1)\mathbf{I} - \frac{m-1}{m} \sum_j \mathbf{1}_j \mathbf{1}_j^T + \frac{m-1}{m-\ell} \sum_a \mathbf{1}_a \mathbf{1}_a^T.
\end{aligned}$$

Recall that Eq. (25) involves $m(m-1)\Gamma$. To obtain $\bar{\rho}$, we again evaluate $\mathbf{1}_{\mathbf{s}'}^T \Gamma^\dagger \mathbf{1}_{\mathbf{s}}$ for any $\mathbf{s} \in S(x)$. We write $A_{\mathbf{s}}$ for the set of actions appearing on the slate \mathbf{s} :

$$\begin{aligned}
\mathbf{1}_{\mathbf{s}'}^T \Gamma^\dagger \mathbf{1}_{\mathbf{s}} &= \left(\frac{1}{\ell^2} - \frac{m-1}{m(m-\ell)} \right) \cdot (\mathbf{1}_{\mathbf{s}'}^T \mathbf{1})(\mathbf{1}^T \mathbf{1}_{\mathbf{s}}) + (m-1) \mathbf{1}_{\mathbf{s}'}^T \mathbf{1}_{\mathbf{s}} - \frac{m-1}{m} \sum_j (\mathbf{1}_{\mathbf{s}'}^T \mathbf{1}_j)(\mathbf{1}_j^T \mathbf{1}_{\mathbf{s}}) \\
&\quad + \frac{m-1}{m-\ell} \sum_a (\mathbf{1}_{\mathbf{s}'}^T \mathbf{1}_a)(\mathbf{1}_a^T \mathbf{1}_{\mathbf{s}}) \\
&= \left(\frac{1}{\ell^2} - \frac{m-1}{m(m-\ell)} \right) \cdot \ell^2 + \sum_j \frac{\mathbf{1}\{s'_j = s_j\}}{1/(m-1)} - \frac{m-1}{m} \cdot \ell \\
&\quad + \frac{m-1}{m-\ell} \sum_a \mathbf{1}\{a \in A_{\mathbf{s}'}\} \mathbf{1}\{a \in A_{\mathbf{s}}\} \quad (26) \\
&= 1 - \frac{(m-1)(\ell^2 + m\ell - \ell^2)}{m(m-\ell)} + \sum_j \frac{\mathbf{1}\{s'_j = s_j\}}{1/(m-1)} + \frac{m-1}{m-\ell} \cdot |A_{\mathbf{s}'} \cap A_{\mathbf{s}}| \\
&= 1 - \frac{m-1}{m-\ell} \cdot \ell + \sum_j \frac{\mathbf{1}\{s'_j = s_j\}}{1/(m-1)} + \frac{m-1}{m-\ell} \cdot |A_{\mathbf{s}} \cap A_{\mathbf{s}'}|,
\end{aligned}$$

where Eq. (26) follows because $\mathbf{1}^T \mathbf{1}_s = \ell$ and $\mathbf{1}_j^T \mathbf{1}_s = 1$ for any valid slate \mathbf{s} . By setting $\mathbf{s}' = \mathbf{s}$, we obtain $\bar{\rho} = 1 + \ell(m-1) = m\ell - \ell + 1$.

Case $\ell = m$: Again, we directly invert the eigenvalues in Eq. (25) to obtain

$$\begin{aligned} \Gamma^\dagger &= \frac{1}{\ell^2} \cdot \mathbf{1}\mathbf{1}^T + (m-1) \cdot \left(\mathbf{I}_{\mathcal{J}} - \frac{\mathbf{1}_{\mathcal{J}}\mathbf{1}_{\mathcal{J}}^T}{\ell} \right) \otimes \left(\mathbf{I}_{\mathcal{A}} - \frac{\mathbf{1}_{\mathcal{A}}\mathbf{1}_{\mathcal{A}}^T}{m} \right) \\ &= \frac{1}{m} \cdot \mathbf{1}\mathbf{1}^T + (m-1)\mathbf{I} - \frac{m-1}{m} \sum_j \mathbf{1}_j \mathbf{1}_j^T - \frac{m-1}{m} \sum_a \mathbf{1}_a \mathbf{1}_a^T. \end{aligned}$$

We finish the theorem by evaluating $\mathbf{1}_{s'}^T \Gamma^\dagger \mathbf{1}_s$:

$$\begin{aligned} \mathbf{1}_{s'}^T \Gamma^\dagger \mathbf{1}_s &= \frac{1}{m} \cdot (\mathbf{1}_{s'}^T \mathbf{1})(\mathbf{1}^T \mathbf{1}_s) + (m-1) \mathbf{1}_{s'}^T \mathbf{1}_s - \frac{m-1}{m} \sum_j (\mathbf{1}_{s'}^T \mathbf{1}_j)(\mathbf{1}_j^T \mathbf{1}_s) \\ &\quad - \frac{m-1}{m} \sum_a (\mathbf{1}_{s'}^T \mathbf{1}_a)(\mathbf{1}_a^T \mathbf{1}_s) \\ &= \frac{1}{m} \cdot m^2 + \sum_j \frac{\mathbf{1}\{s'_j = s_j\}}{1/(m-1)} - \frac{m-1}{m} \cdot m - \frac{m-1}{m} \cdot m \\ &= \sum_j \frac{\mathbf{1}\{s'_j = s_j\}}{1/(m-1)} - m + 2. \end{aligned}$$

We obtain $\bar{\rho} = m^2 - 2m + 2$ by setting $\mathbf{s}' = \mathbf{s}$ and Eq. (23) by taking an expectation over $\mathbf{s}' \sim \pi(\cdot | x)$. \square

E.2 Proof of Corollary 1

We need one last technical result in order to establish the proposition.

Claim 3. Let \mathbf{A}, \mathbf{B} be two symmetric positive semi-definite matrices with $\text{Null}(\mathbf{A}) \subseteq \text{Null}(\mathbf{B})$. Then

$$\max_{\mathbf{z} \perp \text{Null}(\mathbf{B}), \mathbf{z} \neq \mathbf{0}} \frac{\mathbf{z}^T \mathbf{B}^\dagger \mathbf{z}}{\mathbf{z}^T \mathbf{A}^\dagger \mathbf{z}} \leq \max_{\mathbf{z} \perp \text{Null}(\mathbf{B}), \mathbf{z} \neq \mathbf{0}} \frac{\mathbf{z}^T \mathbf{A} \mathbf{z}}{\mathbf{z}^T \mathbf{B} \mathbf{z}}.$$

We now provide the proof of Corollary 1, following which we will prove Claim 3.

Proof of Corollary 1. The corollary follows by Prop. 3. The key step is to bound $\bar{\rho}_{\mu_\varepsilon, x}$, for which we invoke Claim 3. Specifically, we apply the claim with $\mathbf{A} = \Gamma_{\nu, x}$ and $\mathbf{B} = \Gamma_{\mu_\varepsilon, x}$. Since

$$\Gamma_{\mu_\varepsilon, x} = (1 - \varepsilon)\Gamma_{\mu, x} + \varepsilon\Gamma_{\nu, x} = (1 - \varepsilon)\mathbb{E}_{\mathbf{s} \sim \mu(\cdot | x)}[\mathbf{1}_s \mathbf{1}_s^T] + \varepsilon\mathbb{E}_{\mathbf{s} \sim \nu(\cdot | x)}[\mathbf{1}_s \mathbf{1}_s^T],$$

we observe that $\text{Null}(\Gamma_{\nu, x}) = \text{Null}(\Gamma_{\mu_\varepsilon, x})$, because the support of $\mu(\cdot | x)$ is always included in the support of $\nu(\cdot | x)$. Now we can invoke Claim 3 with these choices to see that

$$\begin{aligned} \bar{\rho}_{\mu_\varepsilon, x} &= \sup_{\mathbf{s} \in \text{supp } \mu_\varepsilon(\cdot | x)} \mathbf{1}_s^T \Gamma_{\mu_\varepsilon, x}^\dagger \mathbf{1}_s \\ &\leq \sup_{\mathbf{s} \in \text{supp } \nu(\cdot | x)} \mathbf{1}_s^T \Gamma_{\nu, x}^\dagger \mathbf{1}_s \sup_{\mathbf{s} \in \text{supp } \mu_\varepsilon(\cdot | x)} \frac{\mathbf{1}_s^T \Gamma_{\mu_\varepsilon, x}^\dagger \mathbf{1}_s}{\mathbf{1}_s^T \Gamma_{\nu, x}^\dagger \mathbf{1}_s} \\ &\leq \bar{\rho}_{\nu, x} \max_{\mathbf{z} \perp \text{Null}(\Gamma_{\mu_\varepsilon, x}), \mathbf{z} \neq \mathbf{0}} \frac{\mathbf{z}^T \Gamma_{\mu_\varepsilon, x}^\dagger \mathbf{z}}{\mathbf{z}^T \Gamma_{\nu, x}^\dagger \mathbf{z}} \\ &\leq \bar{\rho}_{\nu, x} \max_{\mathbf{z} \perp \text{Null}(\Gamma_{\mu_\varepsilon, x}), \mathbf{z} \neq \mathbf{0}} \frac{\mathbf{z}^T \Gamma_{\nu, x} \mathbf{z}}{\mathbf{z}^T \Gamma_{\mu_\varepsilon, x} \mathbf{z}} \\ &\leq \bar{\rho}_{\nu, x} \frac{\mathbf{z}^T \Gamma_{\nu, x} \mathbf{z}}{\varepsilon \mathbf{z}^T \Gamma_{\nu, x} \mathbf{z}} \\ &= \frac{\bar{\rho}_{\nu, x}}{\varepsilon} \end{aligned}$$

For the product slate space, using Eq. (18), which was proved within the proof of Prop. 2, we have

$$\begin{aligned}\bar{\rho}_{\nu,x} &= \sup_{\mathbf{s} \in \text{supp } \nu(\cdot|x)} \mathbf{1}_s^T \mathbf{\Gamma}_{\nu,x}^\dagger \mathbf{1}_s \\ &= \sup_{\mathbf{s} \in \text{supp } \nu(\cdot|x)} \left[\sum_{j=1}^{\ell} \frac{1}{1/m} - \ell + 1 \right] \\ &= \ell m - \ell + 1 .\end{aligned}$$

For the ranking slate space, using Prop. 4, we also have $\bar{\rho}_{\nu,x} = \mathcal{O}(\ell m)$, so for both the product slate space and ranking slate space, we obtain $\bar{\rho}_{\mu_\varepsilon,x} = \mathcal{O}(\ell m/\varepsilon)$. Finally, plugging this upper bound and Prop. 3 into the statement of Theorem 1 completes the proof. \square

We finally prove Claim 3.

Proof of Claim 3. Let \mathbf{U} be the square root of matrix \mathbf{A} , i.e., \mathbf{U} is a symmetric positive semidefinite matrix with the same eigenvectors as \mathbf{A} , but with eigenvalues that are square root of the corresponding eigenvalues of \mathbf{A} . Similarly, let \mathbf{V} be the square root of matrix \mathbf{B} . Thus, we have $\mathbf{A} = \mathbf{U}\mathbf{U}$ and $\mathbf{A}^\dagger = \mathbf{U}^\dagger\mathbf{U}^\dagger$ and similarly for \mathbf{B} and \mathbf{V} . Let $\mathbf{\Pi}_\mathbf{A} = \mathbf{U}^\dagger\mathbf{U} = \mathbf{U}\mathbf{U}^\dagger$ denote the projection onto the range of \mathbf{A} and $\mathbf{\Pi}_\mathbf{B}$ denote the projection onto the range of \mathbf{B} . Since $\text{Null}(\mathbf{A}) \subseteq \text{Null}(\mathbf{B})$, we have $\text{Range}(\mathbf{A}) \supseteq \text{Range}(\mathbf{B})$. We prove the claim as follows:

$$\max_{\mathbf{z} \perp \text{Null}(\mathbf{B}), \mathbf{z} \neq \mathbf{0}} \frac{\mathbf{z}^T \mathbf{B}^\dagger \mathbf{z}}{\mathbf{z}^T \mathbf{A}^\dagger \mathbf{z}} = \max_{\mathbf{z} \perp \text{Null}(\mathbf{B}), \mathbf{z} \neq \mathbf{0}} \frac{\mathbf{z}^T \mathbf{U}^\dagger \mathbf{U} \mathbf{B}^\dagger \mathbf{U} \mathbf{U}^\dagger \mathbf{z}}{\mathbf{z}^T \mathbf{U}^\dagger \mathbf{U}^\dagger \mathbf{z}} \quad (27)$$

$$\leq \max_{\mathbf{y} \neq \mathbf{0}} \frac{\mathbf{y}^T \mathbf{U} \mathbf{B}^\dagger \mathbf{U} \mathbf{y}}{\mathbf{y}^T \mathbf{y}} \quad (28)$$

$$= \max_{\mathbf{y} \neq \mathbf{0}} \frac{\mathbf{y}^T \mathbf{U} \mathbf{V}^\dagger \mathbf{V}^\dagger \mathbf{U} \mathbf{y}}{\mathbf{y}^T \mathbf{y}} = \max_{\mathbf{y}: \|\mathbf{y}\|_2=1} \|\mathbf{V}^\dagger \mathbf{U} \mathbf{y}\|_2^2 \quad (29)$$

$$= \max_{\mathbf{y}: \|\mathbf{y}\|_2=1} \|\mathbf{U} \mathbf{V}^\dagger \mathbf{y}\|_2^2 \quad (30)$$

$$= \max_{\mathbf{y} \neq \mathbf{0}} \frac{\mathbf{y}^T \mathbf{V}^\dagger \mathbf{U} \mathbf{U} \mathbf{V}^\dagger \mathbf{y}}{\mathbf{y}^T \mathbf{y}} \quad (31)$$

$$= \max_{\mathbf{y} \perp \text{Null}(\mathbf{B}), \mathbf{y} \neq \mathbf{0}} \frac{\mathbf{y}^T \mathbf{V}^\dagger \mathbf{A} \mathbf{V}^\dagger \mathbf{y}}{\mathbf{y}^T \mathbf{y}} \quad (32)$$

$$= \max_{\mathbf{z} \perp \text{Null}(\mathbf{B}), \mathbf{z} \neq \mathbf{0}} \frac{\mathbf{z}^T \mathbf{V} \mathbf{V}^\dagger \mathbf{A} \mathbf{V}^\dagger \mathbf{V} \mathbf{z}}{\mathbf{z}^T \mathbf{V} \mathbf{V} \mathbf{z}} \quad (33)$$

$$= \max_{\mathbf{z} \perp \text{Null}(\mathbf{B}), \mathbf{z} \neq \mathbf{0}} \frac{\mathbf{z}^T \mathbf{A} \mathbf{z}}{\mathbf{z}^T \mathbf{B} \mathbf{z}} .$$

In Eq. (27) we substitute $\mathbf{U}^\dagger\mathbf{U}^\dagger = \mathbf{A}^\dagger$ and also use the fact that $\mathbf{U}\mathbf{U}^\dagger = \mathbf{\Pi}_\mathbf{A}$ and $\mathbf{\Pi}_\mathbf{A}\mathbf{z} = \mathbf{z}$ because $\mathbf{z} \in \text{Range}(\mathbf{B}) \subseteq \text{Range}(\mathbf{A})$. Eq. (28) is obtained by substituting $\mathbf{y} = \mathbf{U}^\dagger\mathbf{z}$ and relaxing the maximization to be over $\mathbf{y} \neq \mathbf{0}$. In Eq. (29) we substitute $\mathbf{V}^\dagger\mathbf{V}^\dagger = \mathbf{B}^\dagger$. In Eq. (30) we use the fact that the operator norm of a matrix and its transpose are equal. In Eq. (31) we substitute $\mathbf{A} = \mathbf{U}\mathbf{U}$ and note that it suffices to consider $\mathbf{y} \perp \text{Null}(\mathbf{B})$ because $\text{Null}(\mathbf{V}^\dagger\mathbf{A}\mathbf{V}^\dagger) = \text{Null}(\mathbf{B})$. In Eq. (32) we use the fact that $\mathbf{z} \mapsto \mathbf{V}\mathbf{z}$ is a bijection on $\text{Range}(\mathbf{B})$, which is an orthogonal complement to $\text{Null}(\mathbf{B})$, so we can substitute $\mathbf{V}\mathbf{z} = \mathbf{y}$. Finally, in Eq. (33) we substitute $\mathbf{B} = \mathbf{V}\mathbf{V}$ and use the fact that $\mathbf{V}^\dagger\mathbf{V} = \mathbf{\Pi}_\mathbf{B}$ and $\mathbf{\Pi}_\mathbf{B}\mathbf{z} = \mathbf{z}$ because $\mathbf{z} \in \text{Range}(\mathbf{B})$. \square

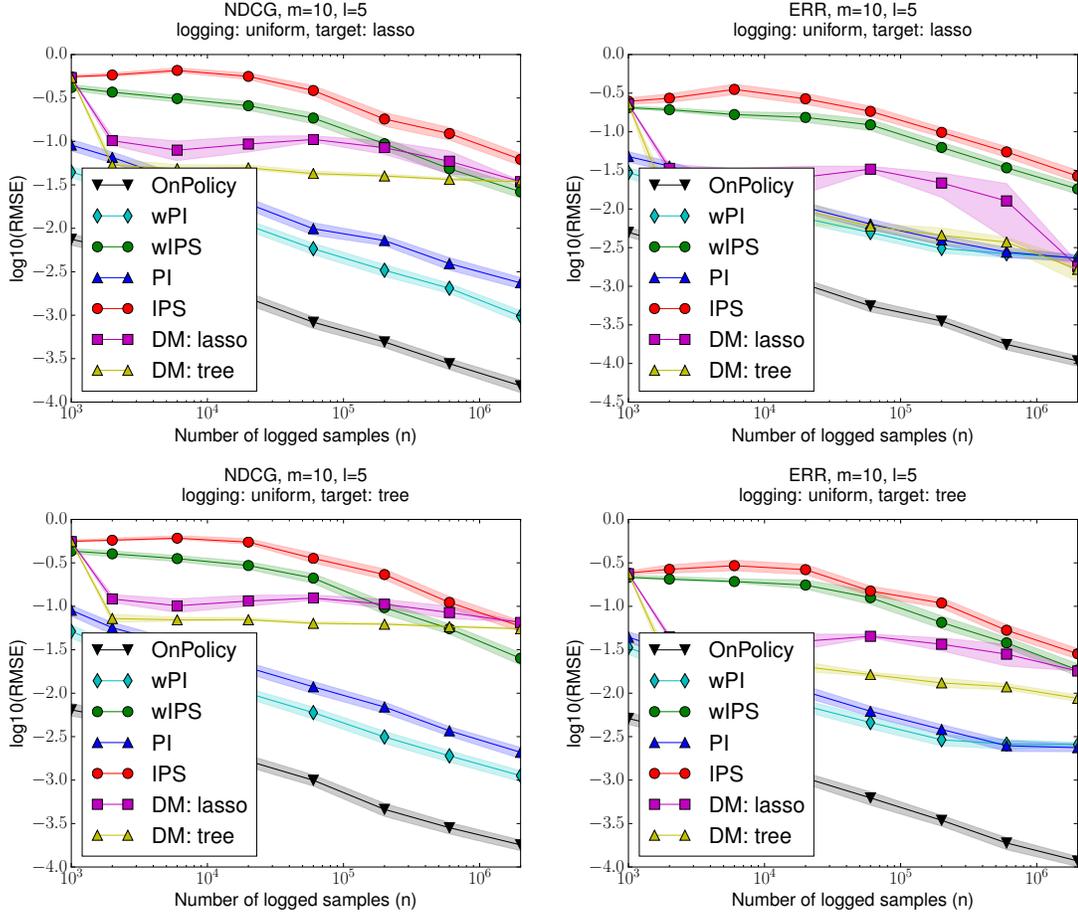


Figure 3: RMSE of value estimators for an increasing logged dataset under a uniform logging policy with slate space $(10, 5)$. Target is $lasso_{body}$ (top panel) and $tree_{body}$ (bottom panel). Metrics are NDCG (left) and ERR (right).

F Supplementary plots for off-policy evaluation on semi-synthetic data

We experimented with several configurations of slate spaces, logging and target policies, and whole-page metrics in the semi-synthetic evaluation setup. This section details the plots for all configurations. The key parameters were:

1. Metric: NDCG or ERR. NDCG satisfies the linearity assumption, while ERR does not.
2. Slate space: $(m, l) = (100, 10)$ or $(10, 5)$.
3. Logging policy: Unif, $lasso_{title}$, $tree_{title}$
4. Target policy: $lasso_{body}$, $tree_{body}$
5. Temperature α : Uniform, Slightly peaked, Very peaked. Uniform corresponds to $\alpha = 0$. For the small slate spaces with $(m, l) = (10, 5)$, $\alpha = 1.0$ creates a slightly peaked logging distribution, while $\alpha = 2.0$ creates a severely peaked logging distribution. For the larger slate spaces with $(m, l) = (100, 10)$, $\alpha = 0.5$ is moderately peaked while $\alpha = 1.0$ is severely peaked.

The plots in Figures 3–12 detail the top row of Figure 2 for all combinations of these parameters.

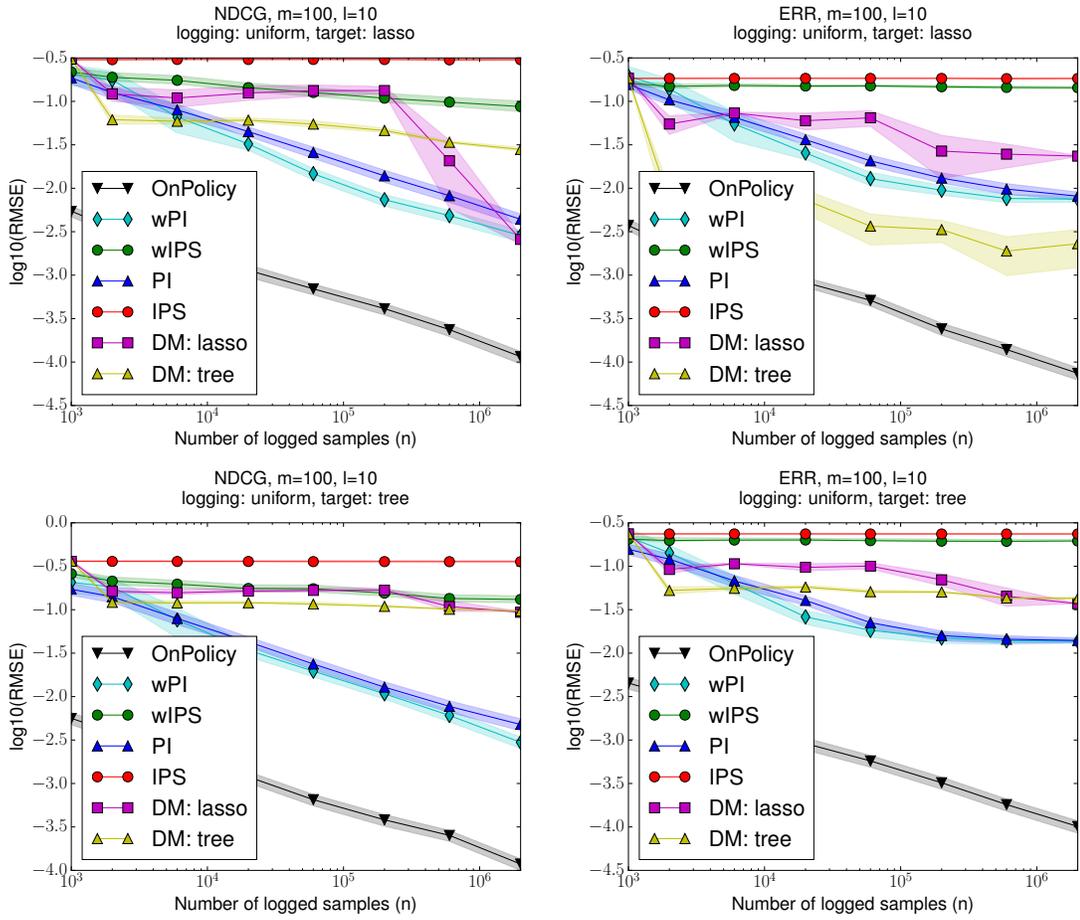


Figure 4: RMSE of value estimators for an increasing logged dataset under a uniform logging policy with slate space (100, 10). Target is $lasso_{body}$ (top panel) and $tree_{body}$ (bottom panel). Metrics are NDCG (left) and ERR (right).

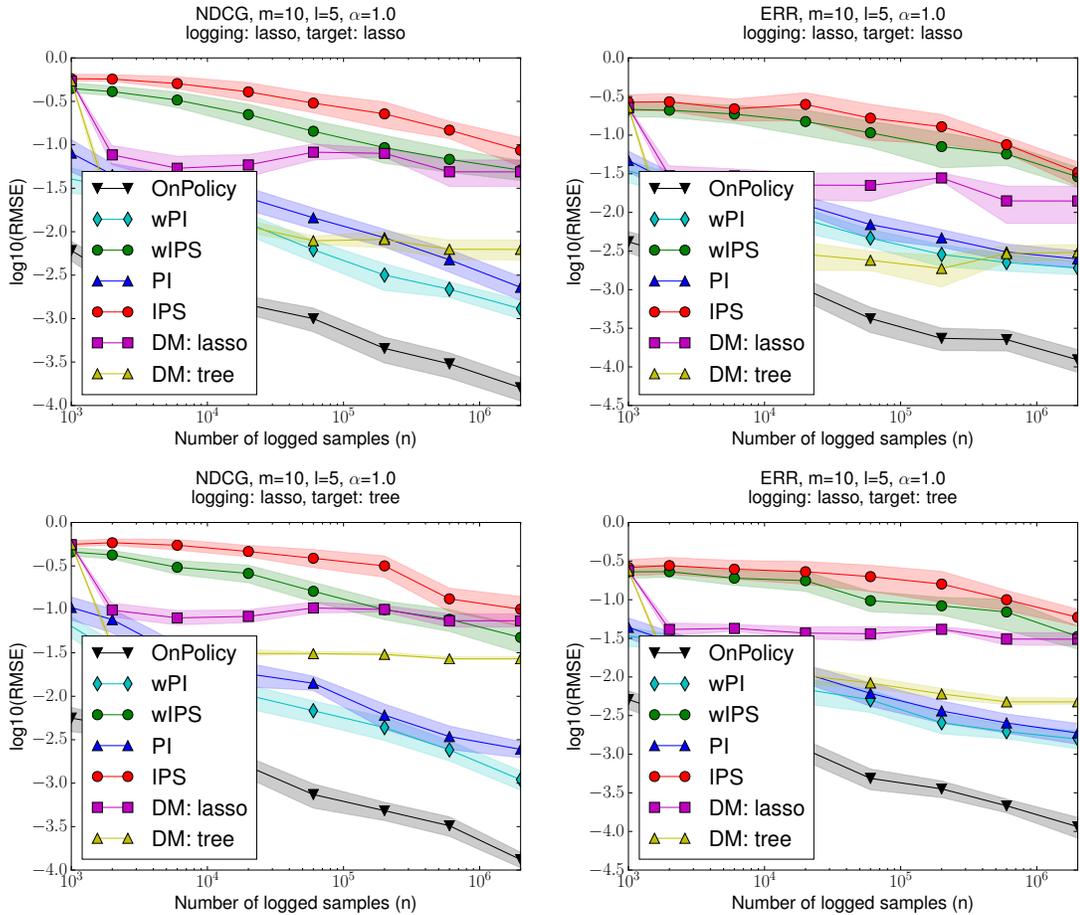


Figure 5: RMSE of value estimators for an increasing logged dataset under a moderately peaked logging policy ($lasso_{title}$, $\alpha = 1.0$) with slate space (10, 5). Target is $lasso_{body}$ (top panel) and $tree_{body}$ (bottom panel). Metrics are NDCG (left) and ERR (right).

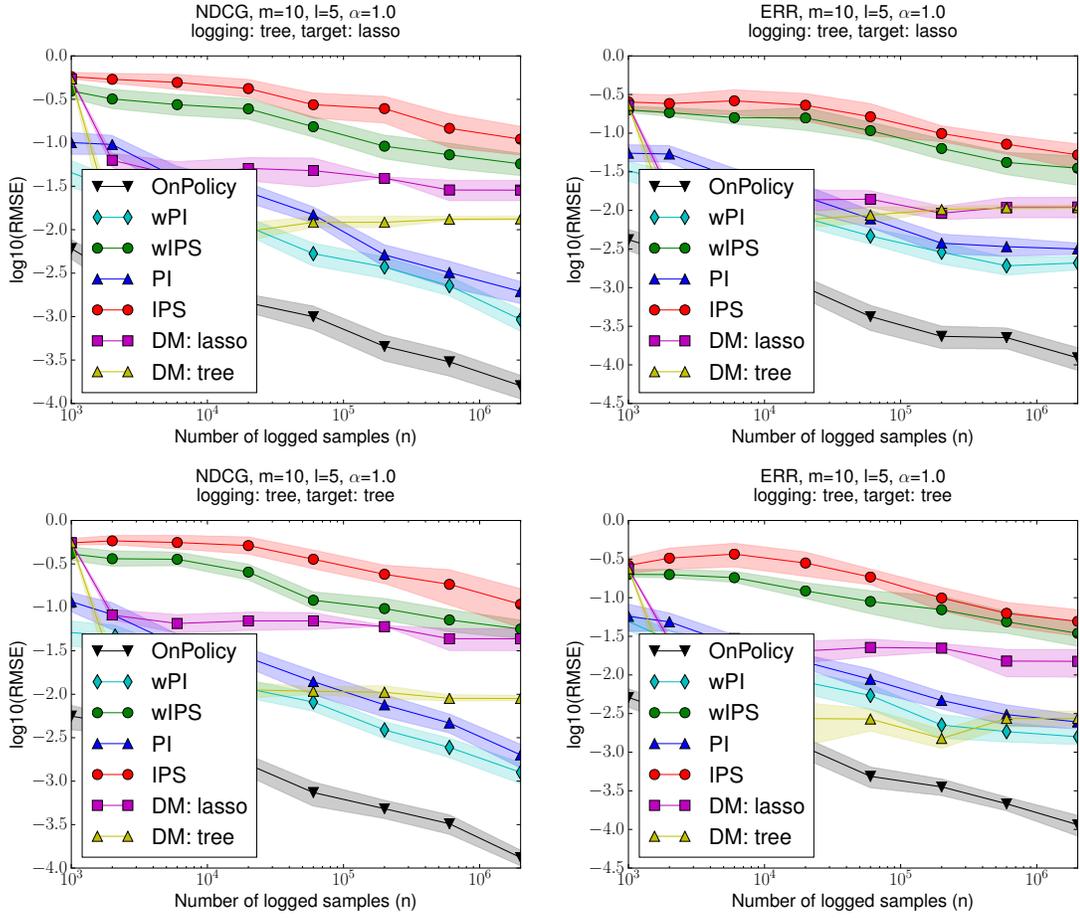


Figure 6: RMSE of value estimators for an increasing logged dataset under a moderately peaked logging policy ($tree_{title}$, $\alpha = 1.0$) with slate space (10, 5). Target is $lasso_{body}$ (top panel) and $tree_{body}$ (bottom panel). Metrics are NDCG (left) and ERR (right).

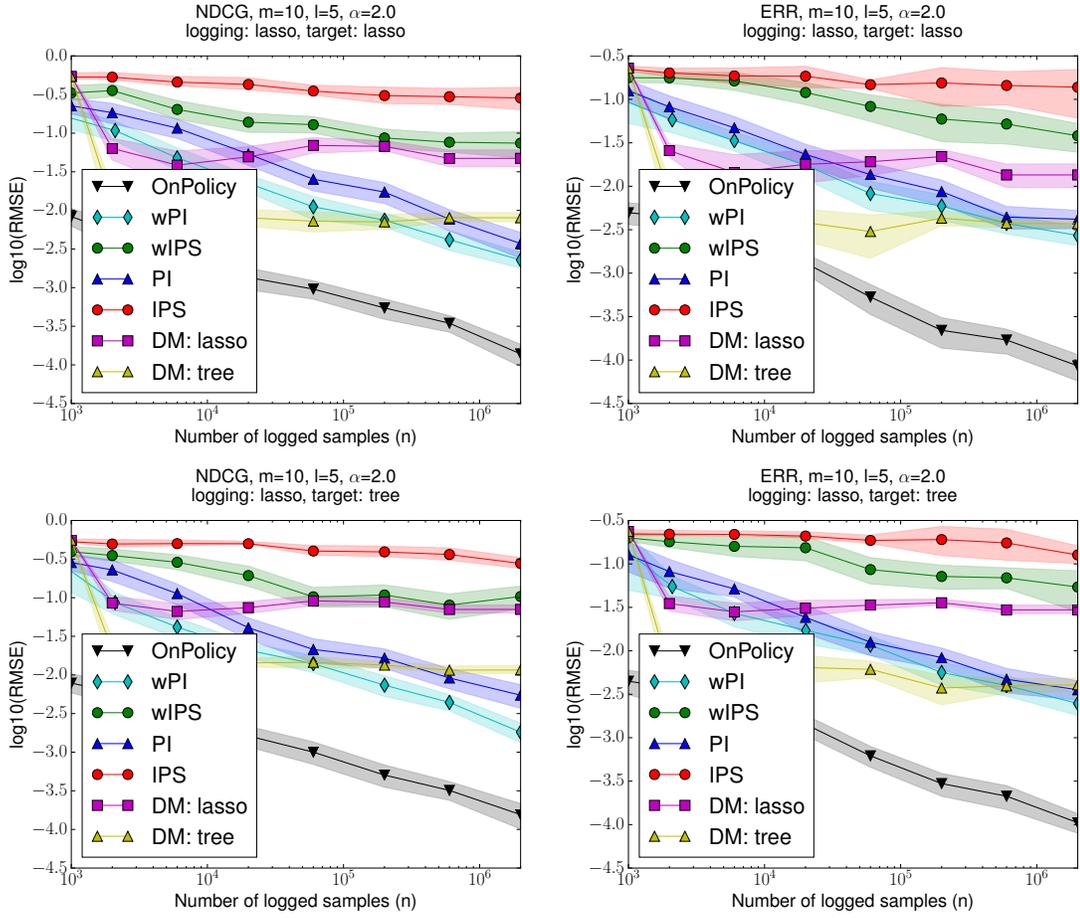


Figure 7: RMSE of value estimators for an increasing logged dataset under a severely peaked logging policy ($lasso_{title}$, $\alpha = 2.0$) with slate space (10, 5). Target is $lasso_{body}$ (top panel) and $tree_{body}$ (bottom panel). Metrics are NDCG (left) and ERR (right).

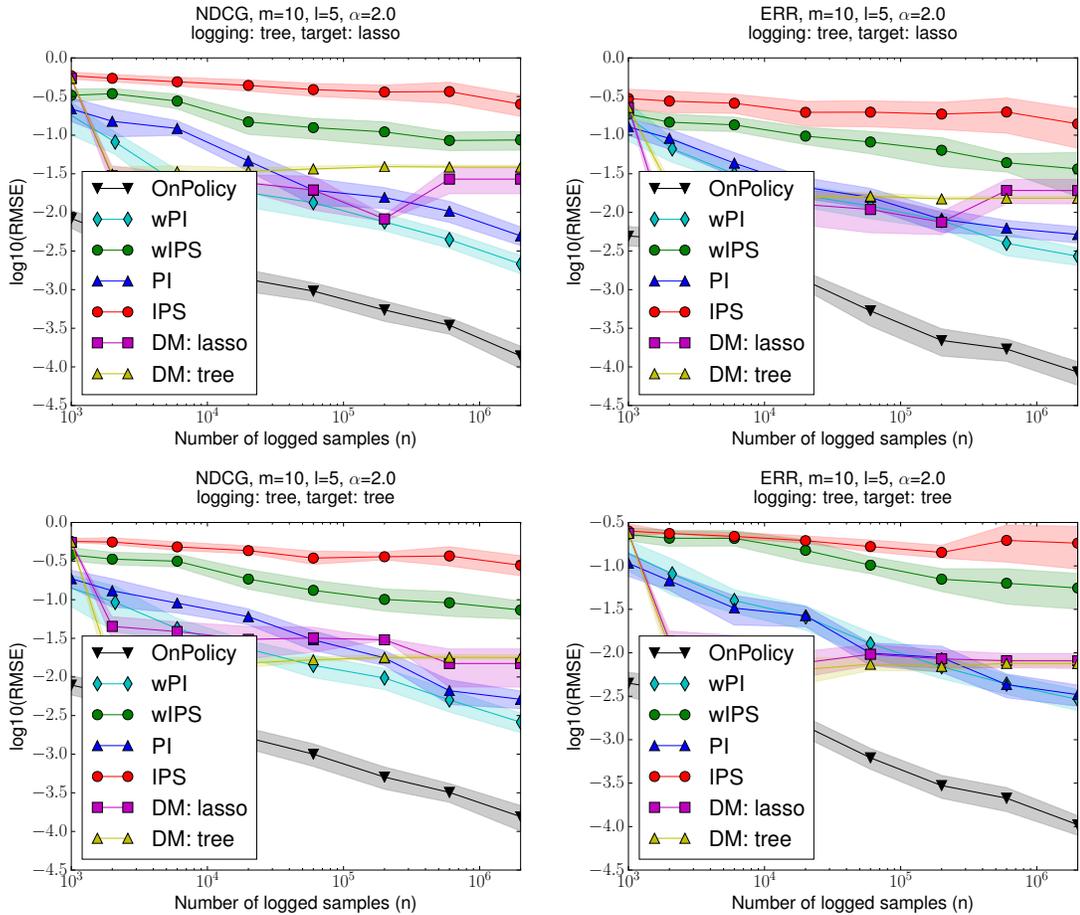


Figure 8: RMSE of value estimators for an increasing logged dataset under a severely peaked logging policy ($tree_{title}$, $\alpha = 2.0$) with slate space (10, 5). Target is $lasso_{body}$ (top panel) and $tree_{body}$ (bottom panel). Metrics are NDCG (left) and ERR (right).

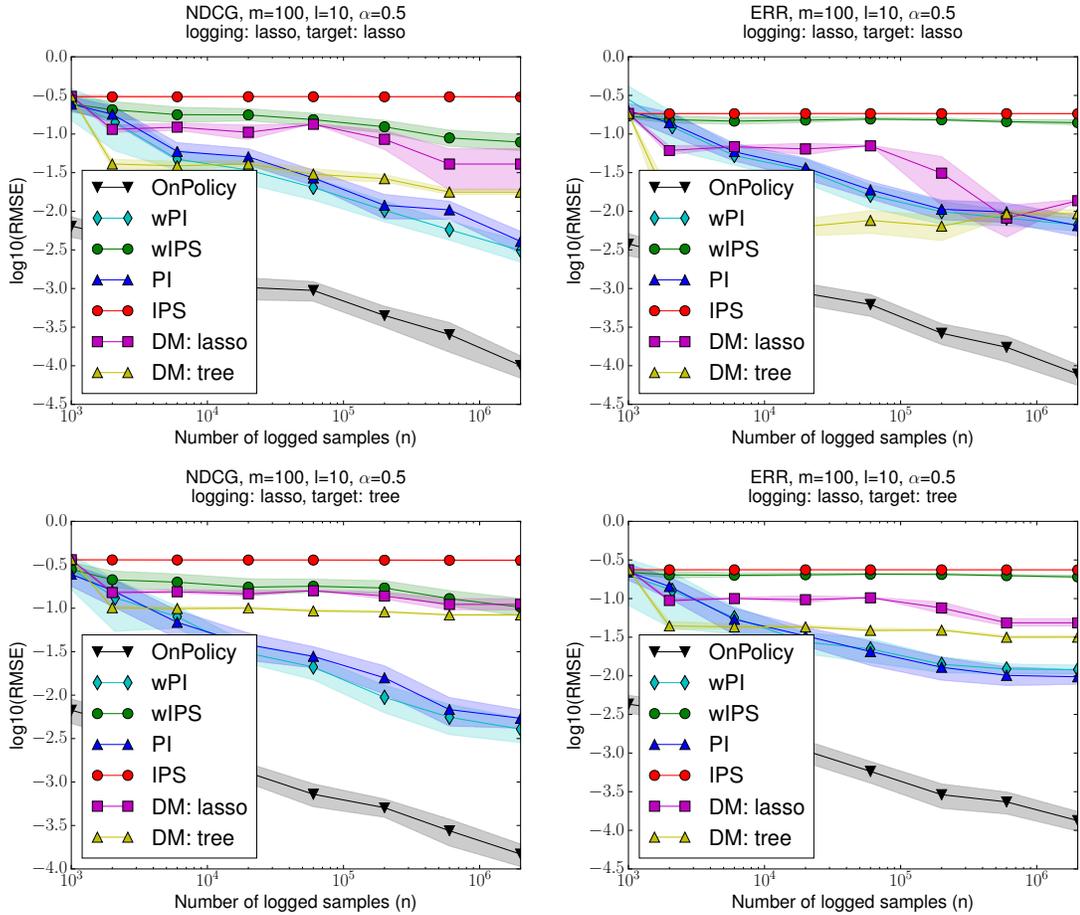


Figure 9: RMSE of value estimators for an increasing logged dataset under a moderately peaked logging policy ($lasso_{title}$, $\alpha = 0.5$) with slate space (100, 10). Target is $lasso_{body}$ (top panel) and $tree_{body}$ (bottom panel). Metrics are NDCG (left) and ERR (right).

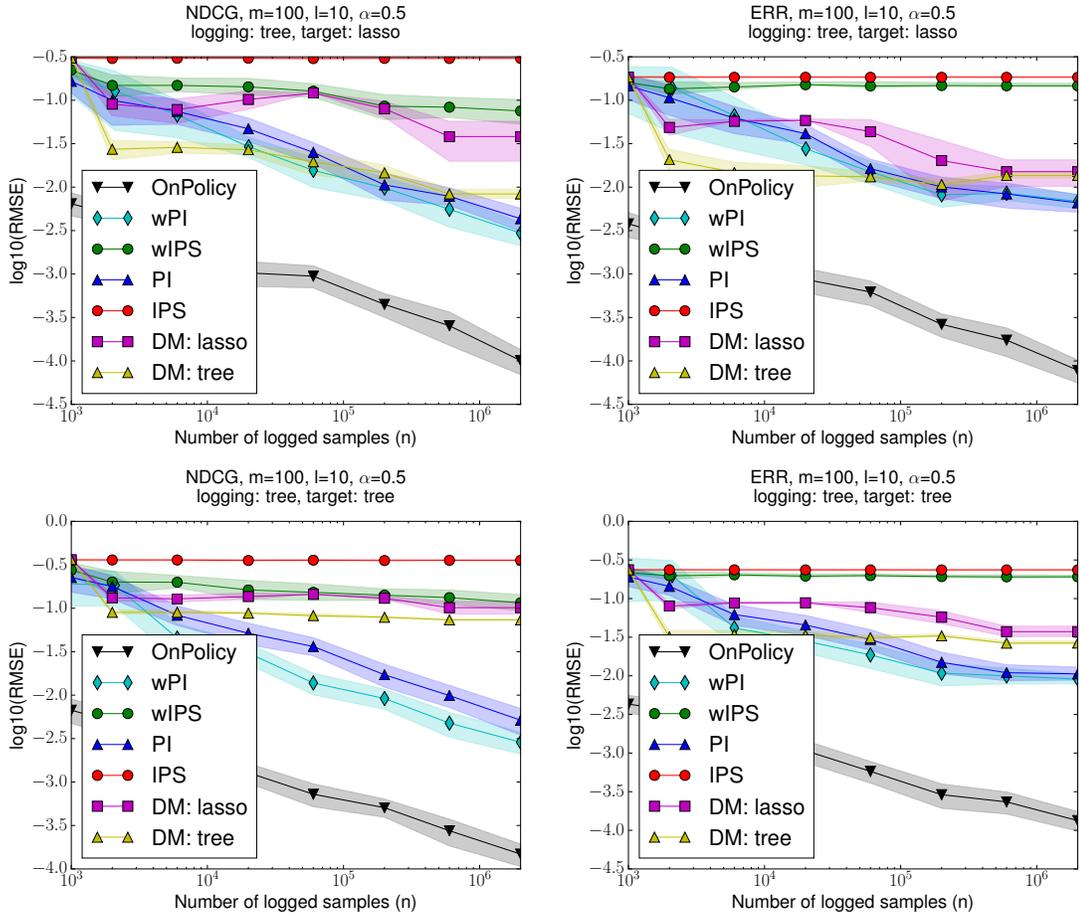


Figure 10: RMSE of value estimators for an increasing logged dataset under a moderately peaked logging policy ($tree_{title}, \alpha = 0.5$) with slate space (100, 10). Target is $lasso_{body}$ (top panel) and $tree_{body}$ (bottom panel). Metrics are NDCG (left) and ERR (right).

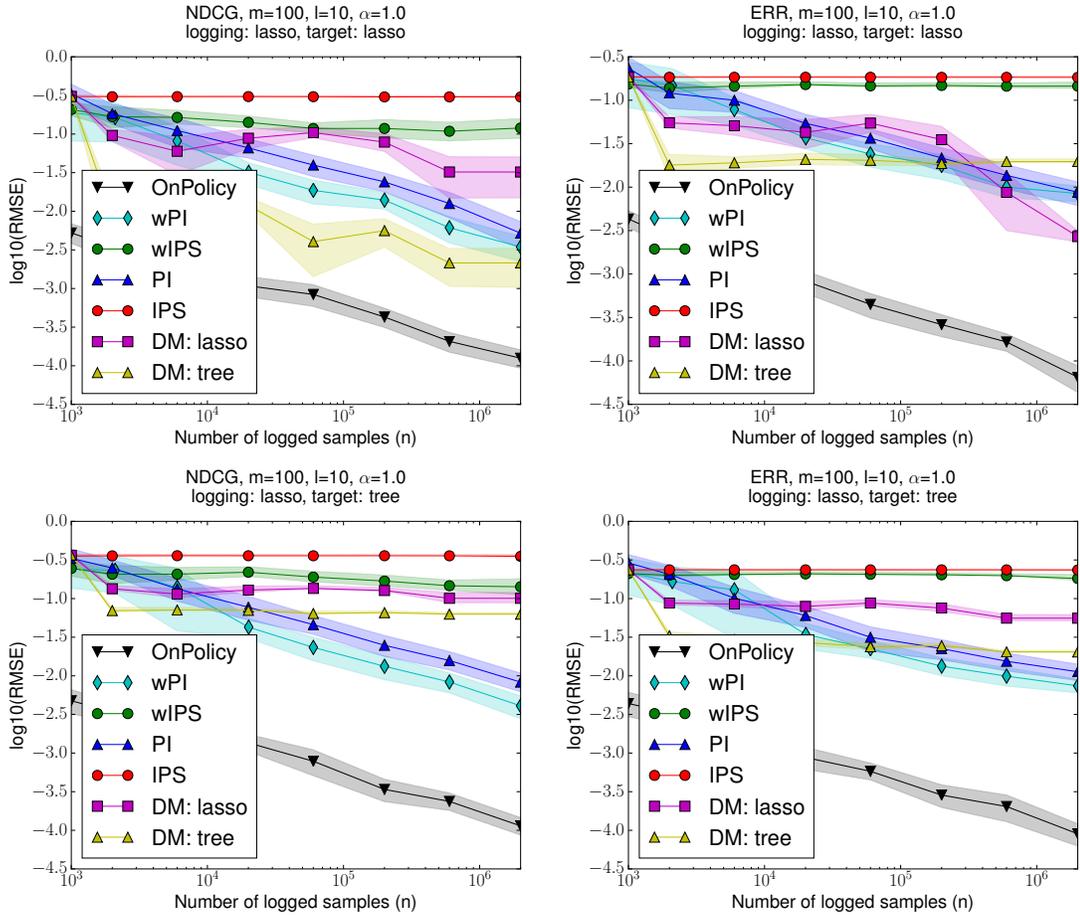


Figure 11: RMSE of value estimators for an increasing logged dataset under a severely peaked logging policy ($lasso_{title}$, $\alpha = 1.0$) with slate space (100, 10). Target is $lasso_{body}$ (top panel) and $tree_{body}$ (bottom panel). Metrics are NDCG (left) and ERR (right).

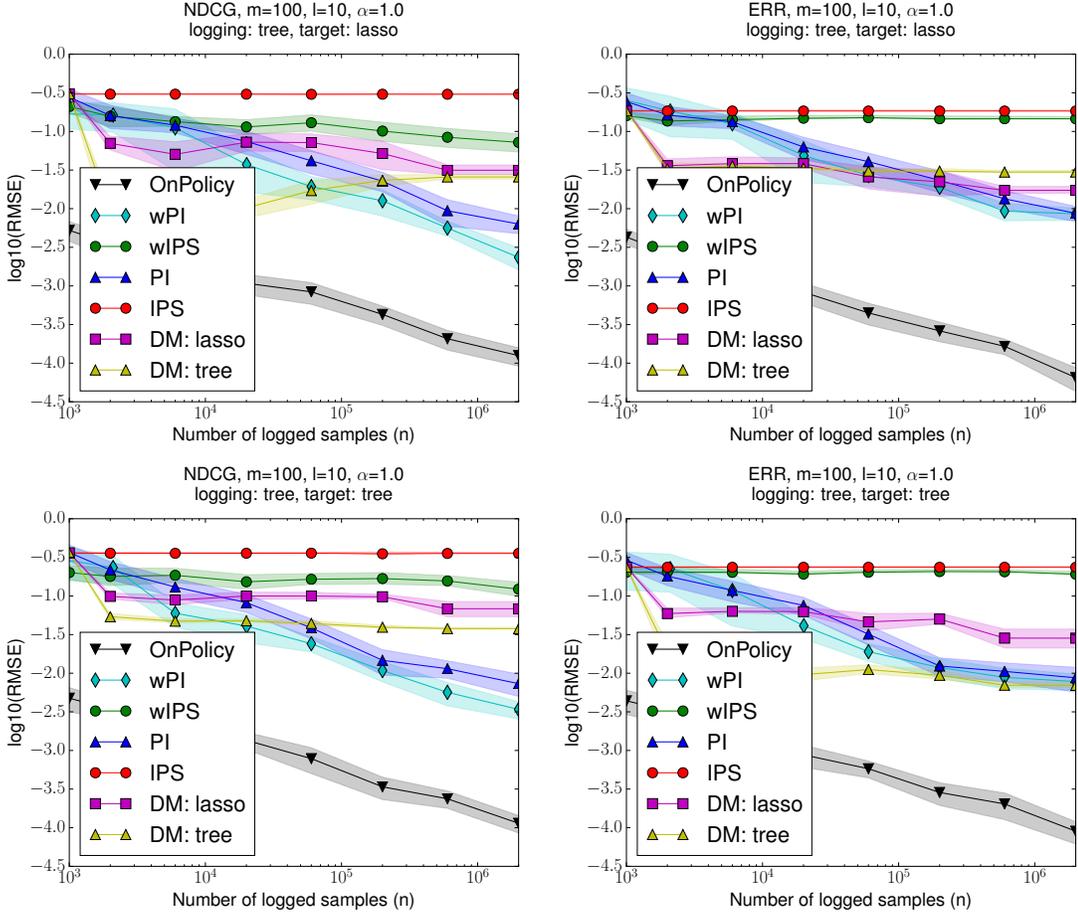


Figure 12: RMSE of value estimators for an increasing logged dataset under a severely peaked logging policy ($tree_{title}, \alpha = 1.0$) with slate space (100, 10). Target is $lasso_{body}$ (top panel) and $tree_{body}$ (bottom panel). Metrics are NDCG (left) and ERR (right).

G Off-policy optimization

For off-policy optimization experiments, we compare two methods – SUP and PI-OPT on the MSLR-WEB10K dataset. Both SUP and PI-OPT uses all the features $\mathbf{f}(x, a)$ for query-document (x, a) pairs in the training fold. SUP uses regression targets as outlined in Section 4.2. We also experimented with regression to raw relevance judgments, this is denoted SUP-*rel*. For PI-OPT, each query-document-*position* triplet produces a regression example (x, a, j) with a concatenated feature vector $\mathbf{f}(x, a, j) := [\mathbf{f}(x, a); \mathbf{1}_j]$ where $\mathbf{1}_j$ is a ℓ -dimensional one-hot encoding of position j . Every logged sample with query x yields an estimate $\hat{\phi}_j(x, a)$ for every candidate document a and position j . These are our natural regression targets. There is one further optimization we can do that is computationally tractable when the set of queries $\{x\}$ is finite. By averaging all estimated $\hat{\phi}_j(x, a)$ for a particular query, we can create a lower-variance target for regression that remains an unbiased estimate of $\phi_j(x, a)$.

Both SUP and PI-OPT employ gradient boosted regression trees with $n = 1000$ tree-ensembles and up to 70 leaves in each tree, to predict their corresponding regression targets. With a trained model, SUP constructs slates in a straightforward way: For any input query x , we score all candidate documents $a \in A(x)$ using the trained model $\mathbf{f}(x, a) \mapsto score(a)$, and sort the scores in descending order. Rankings are constructed using the top- ℓ scoring candidates in order. For PI-OPT, we score every document-position pair $(a, j) \in A(x) \times \{1 \dots \ell\}$, $\mathbf{f}(x, a, j) \mapsto score(a, j)$. Now we greedily pick the highest scored pair (a, j) and insert document a in slot j of the slate. After eliminating all invalid

pairs, $(*, j)$ and $(a, *)$, we repeat this greedy procedure until all positions in the slate are filled. This gives us a computationally efficient, albeit approximate, maximizer of $\operatorname{argmax}_{\mathbf{s}} \sum_{j=1}^{\ell} \operatorname{score}(s_j, j)$.

H Overlap between base-rankers

We use four different base-rankers $\operatorname{lasso}_{\text{title}}$, $\operatorname{lasso}_{\text{body}}$, $\operatorname{tree}_{\text{title}}$, $\operatorname{tree}_{\text{body}}$ in our semi-synthetic experiments to instantiate logging and target policies. In Table 2, we report how similar the top- ℓ rankings ($\ell = 10$) retrieved by these rankers are. We report two metrics for every pair of rankers: the average fraction of documents retrieved in common by both rankers (and its standard deviation), and the Kendall’s tau computed over the union of documents retrieved by either ranker (documents retrieved by one ranker but not the other are assumed to be ranked at $\ell + 1$ in the other ranking).

Table 2: Reporting the difference between the base-rankers $\operatorname{lasso}_{\text{title}}$, $\operatorname{lasso}_{\text{body}}$, $\operatorname{tree}_{\text{title}}$, $\operatorname{tree}_{\text{body}}$ as measured by average overlap of retrieved document sets and Kendall’s tau.

Pair	Overlap		Kendall’s τ	
	Avg.	Std.Dev.	Avg.	Std.Dev.
$(\operatorname{lasso}_{\text{title}}, \operatorname{tree}_{\text{title}})$	0.523	0.216	-0.041	0.307
$(\operatorname{lasso}_{\text{body}}, \operatorname{tree}_{\text{body}})$	0.426	0.236	-0.221	0.322
$(\operatorname{tree}_{\text{title}}, \operatorname{tree}_{\text{body}})$	0.270	0.198	-0.394	0.236
$(\operatorname{tree}_{\text{title}}, \operatorname{lasso}_{\text{body}})$	0.274	0.203	-0.405	0.239
$(\operatorname{lasso}_{\text{title}}, \operatorname{tree}_{\text{body}})$	0.250	0.199	-0.421	0.231
$(\operatorname{lasso}_{\text{title}}, \operatorname{lasso}_{\text{body}})$	0.262	0.202	-0.415	0.233