
Dueling Posterior Sampling for Preference-Based Reinforcement Learning

Ellen R. Novoseller¹, Yibing Wei¹, Yanan Sui², Yisong Yue¹, Joel W. Burdick¹

¹Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125

²School of Aerospace Engineering, Tsinghua University, Beijing, China 100084

{enovoseller, ywwei, yyue}@caltech.edu, ysui@tsinghua.edu.cn, jwb@robotics.caltech.edu

Abstract

In preference-based reinforcement learning (RL), an agent interacts with the environment while receiving preferences instead of absolute feedback. While there is increasing research activity in preference-based RL, the design of formal frameworks that admit tractable theoretical analysis remains an open challenge. Building upon ideas from preference-based bandit learning and posterior sampling in RL, we present DUELING POSTERIOR SAMPLING (DPS), which employs preference-based posterior sampling to learn both the system dynamics and the underlying utility function that governs the preference feedback. As preference feedback is provided on trajectories rather than individual state-action pairs, we develop a Bayesian approach for the credit assignment problem, translating preferences to a posterior distribution over state-action reward models. We prove an asymptotic Bayesian no-regret rate for DPS with a Bayesian linear regression credit assignment model. This is the first regret guarantee for preference-based RL to our knowledge. We also discuss possible avenues for extending the proof methodology to other credit assignment models. Finally, we evaluate the approach empirically, showing competitive performance against existing baselines.

human preferences as feedback. While many RL algorithms assume the existence of a numerical reward signal, in settings involving humans, it is often unclear how to define a reward signal that accurately reflects optimal system-human interaction. For instance, in autonomous driving (Basu et al., 2017) and robotics (Argall et al., 2009; Akrouer et al., 2012), users can have difficulty with both specifying numerical reward functions and providing demonstrations of desired behavior. Moreover, a misspecified reward function can result in “reward hacking” (Amodei et al., 2016), in which undesirable actions achieve high rewards. In such situations, the user’s preferences could more reliably measure her intentions.

This work studies the problem of preference-based reinforcement learning (PBRL), in which the RL agent executes a pair of trajectories of interaction with the environment, and the user provides (noisy) pairwise preference feedback, revealing which of the two trajectories is preferred. Though the study of PBRL has seen increased interest in recent years (Christiano et al., 2017; Wirth et al., 2017), it remains an open challenge to design formal frameworks that admit tractable theoretical analysis. While the preference-based bandit setting—in which the agent observes preferences between selected actions—has seen significant theoretical progress (e.g., Yue et al. (2012); Zoghi et al. (2014); Ailon et al. (2014); Szörényi et al. (2015); Dudík et al. (2015); Zoghi et al. (2015); Ramamohan et al. (2016); Wu and Liu (2016); Sui et al. (2017, 2018b)), the PBRL setting is more challenging, as the environment’s dynamics can stochastically translate the agent’s policies for interaction (analogous to actions in the bandit setting) to the observed trajectories.

In this paper, we present the DUELING POSTERIOR SAMPLING (DPS) algorithm, which uses preference-based posterior sampling to tackle PBRL in the Bayesian regime. Posterior sampling (Thompson, 1933), also called Thompson sampling, is a Bayesian model-based approach to balancing exploration and exploitation, which enables the algorithm to efficiently learn models

1 INTRODUCTION

Reinforcement learning (RL) agents interact with humans in many domains, from clinical trials (Sui et al., 2018a) to autonomous driving (Sadigh et al., 2017) to human-robot interaction (Kupcsik et al., 2018), and take

of both the environment’s state transition dynamics and reward function. Previous work on posterior sampling in RL (Osband et al., 2013; Gopalan and Mannor, 2015; Agrawal and Jia, 2017; Osband and Van Roy, 2017) is focused on learning from absolute rewards, while we extend posterior sampling to both elicit and learn from trajectory-level preference feedback.

To elicit preference feedback, at every episode of learning, DPS draws two independent samples from the posterior to generate two trajectories. This approach is inspired by the Self-Sparring algorithm proposed for the bandit setting (Sui et al., 2017), but has a quite different theoretical analysis, as we need to incorporate trajectory-level preference learning and state transition dynamics.

Learning from trajectory-level preferences is in general a very challenging problem, as information about the rewards is sparse (often just one bit), is only relative to the pair of trajectories being compared, and does not explicitly include information about actions within trajectories. DPS learns from preference feedback by internally maintaining a Bayesian state-action reward model that explains the preferences; this reward model is a solution to the *temporal credit assignment problem* (Akrouer et al., 2012; Zoghi et al., 2014; Szörényi et al., 2015; Christiano et al., 2017; Wirth et al., 2016, 2017), i.e., determining which of the encountered states and actions are responsible for the trajectory-level preference feedback.

We developed DPS concurrently with an analysis framework for characterizing regret convergence in the episodic setting, based upon information-theoretic techniques for bounding the Bayesian regret of posterior sampling (Russo and Van Roy, 2016). We mathematically integrate Bayesian credit assignment and preference elicitation within the conventional posterior sampling framework, evaluate several credit assignment models, and prove a Bayesian asymptotic no-regret rate for DPS with a Bayesian linear regression credit assignment model. To our knowledge, this is the first PBRL approach with theoretical guarantees. We also demonstrate that DPS delivers competitive performance empirically.

2 RELATED WORK

Posterior sampling. Balancing exploration and exploitation is a key problem in RL. In the episodic learning setting, the agent typically aims to balance exploration and exploitation to minimize its regret, i.e., the gap between the expected total rewards of the agent and the optimal policy. Posterior sampling, first proposed in Thompson (1933), is a Bayesian model-based approach toward achieving this goal, which iterates between (1) updating the posterior of a Bayesian environment model

and (2) sampling from this posterior to select the next policy. In both the bandit and RL settings, posterior sampling has been demonstrated to perform competitively in experiments and enjoy favorable theoretical regret guarantees (Osband and Van Roy, 2017; Osband et al., 2013; Agrawal and Jia, 2017; Chapelle and Li, 2011).

Our approach builds upon two existing posterior sampling algorithms: Self-Sparring (Sui et al., 2017) for preference-based bandit learning (also known as dueling bandits (Yue et al., 2012)) and posterior sampling RL (Osband et al., 2013). Self-Sparring maintains a posterior over each action’s reward, and in each iteration, draws multiple samples from this posterior to “duel” or “spar” via preference elicitation. For each set of sampled rewards, the algorithm executes the action with the highest reward sample, obtaining new preferences to update the model posterior. Sui et al. (2017) prove an asymptotic no-regret guarantee for Self-Sparring with independent Beta-Bernoulli reward models for each action.

Within RL, posterior sampling has been applied to the finite-horizon setting with absolute rewards to learn Bayesian posteriors over both the dynamics and rewards. Each posterior sample yields models of both dynamics and rewards, which are used to compute the optimal policy for the sampled system. This policy is executed to get a roll-out trajectory, used to update the dynamics and reward posteriors. In Osband et al. (2013), the authors show an expected regret of $O(hS\sqrt{AT\log(SAT)})$ after T time-steps, with finite time horizon h and discrete state and action spaces of sizes S and A , respectively.

Our theoretical analysis studies the Bayesian linear regression credit assignment model, which most closely resembles Bayesian reward modeling in the linear bandit setting (Abbasi-Yadkori et al., 2011; Agrawal and Goyal, 2013; Abeille and Lazaric, 2017). While both the PBRL and linear bandit settings apply Bayesian linear regression to recover model parameters, PBRL additionally requires learning the dynamics, determining policies via value iteration, and receiving feedback as preferences between trajectory pairs.

Several regret analyses in the linear bandit domain (Abbasi-Yadkori et al., 2011; Agrawal and Goyal, 2013; Abeille and Lazaric, 2017) rely upon martingale concentration properties introduced in Abbasi-Yadkori et al. (2011), and depend upon a bound that is not applicable in the preference-based setting (see Appendix C). Intuitively, these analyses assume that the agent learns about rewards with respect to every observation’s feature vector. In contrast, the preference-based setting assumes that only the *difference* in the total rewards of two trajectories affects human preferences. Thus, while the algorithm incurs regret with respect to every sampled trajectory, only

differences between compared trajectory feature vectors yield information about rewards.

Our regret analysis takes inspiration from the information-theoretic perspective on Thompson sampling introduced in Russo and Van Roy (2016), a framework for quantifying Bayesian regret in terms of the information gained at each step about the optimal action. This analysis focuses upon upper-bounding the *information ratio*, which quantifies the trade-off between exploration (via the information gain) and exploitation (via the instantaneous regret) at each step. Several studies (Zanette and Sarkar, 2017; Nikolov et al., 2018) consider extensions of this work to the RL setting, but to our knowledge, it has not previously been applied toward preference-based learning.

Preference-based learning. Previous work on PBRL has shown successful performance in a number of applications, including Atari games and the Mujoco environment (Christiano et al., 2017), learning human preferences for autonomous driving (Sadigh et al., 2017), and selecting a robot’s controller parameters (Kupcsik et al., 2018; Akrouf et al., 2014). Yet, to our knowledge, the PBRL literature still lacks theoretical guarantees.

Much of the existing work in PBRL handles a distinct setting from ours. While we seek online regret minimization, several existing algorithms minimize the number of preference queries (Christiano et al., 2017; Wirth et al., 2016). Such algorithms, for instance those which apply deep learning, typically assume that many simulations can be cheaply run between preference queries. In contrast, our setting assumes that experimentation is as expensive as preference elicitation; this could include such domains as adaptive experiment design and human-robot interaction without well-understood human dynamics.

Existing approaches for trajectory-level preference-based RL may be broadly divided into three categories (Wirth, 2017): a) directly optimizing policy parameters (Wilson et al., 2012; Busa-Fekete et al., 2013; Kupcsik et al., 2018); b) modeling action preferences in each state (Förnkrantz et al., 2012); and c) learning a utility function to characterize the rewards, returns, or values of state-action pairs (Wirth and Förnkrantz, 2013a,b; Akrouf et al., 2012; Wirth et al., 2016; Christiano et al., 2017). In c), the utility is often modeled as linear in the trajectory features. If those features are defined in terms of visitations to each state-action pair, then utility directly corresponds to the total (undiscounted) reward.

We adopt the third of these paradigms: PBRL with underlying utility functions. By inferring state-action rewards from preference feedback, one can derive relatively-interpretable reward models and employ such

methods as value iteration. In addition, utility-based approaches may be more sample efficient compared to policy search and preference relation methods (Wirth, 2017), as they extract more information from each observation. Notably, Wilson et al. (2012) learn a Bayesian model over policy parameters, and sample from its posterior to inform actions. From existing PBRL methods, their algorithm perhaps most resembles ours; however, compared to utility-based approaches, policy search methods typically require either more samples or expert knowledge to craft the policy parameters (Wirth et al., 2017; Kupcsik et al., 2018).

Beyond RL, preference-based learning has been the subject of much research. The bandit setting (Yue et al., 2012; Zoghi et al., 2014; Ailon et al., 2014; Szörényi et al., 2015; Dudík et al., 2015; Zoghi et al., 2015; Ramamohan et al., 2016; Wu and Liu, 2016; Sui et al., 2017, 2018b) is closest, as it is essentially a single-state variant of RL. Other settings include: active learning (Sadigh et al., 2017; Houthby et al., 2011; Eric et al., 2008), which is focused exclusively on learning an accurate model rather than maximizing utility of decision-making; learning with more structured preference feedback (Radlinski and Joachims, 2005; Shivaswamy and Joachims, 2012; Raman et al., 2013; Shivaswamy and Joachims, 2015), where the learner receives more than one bit of information per preference elicitation; and batch supervised settings such as learning to rank (Herbrich et al., 1999; Chu and Ghahramani, 2005; Joachims, 2005; Burges et al., 2005; Yue et al., 2007; Burges et al., 2007; Liu, 2009).

3 PROBLEM STATEMENT

Preliminaries. We consider fixed-horizon Markov Decision Processes (MDPs), in which rewards are replaced by preferences over trajectories. This class of MDPs can be represented as a tuple, $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \phi, p, p_0, h)$, where the state space \mathcal{S} and action space \mathcal{A} are finite sets with cardinalities S and A , respectively. The agent episodically interacts with the environment in length- h roll-out trajectories of the form $\tau = \{s_1, a_1, s_2, a_2, \dots, s_h, a_h, s_{h+1}\}$. In the i^{th} iteration, the agent executes two trajectory roll-outs τ_{i1} and τ_{i2} and observes a preference between them; we use the notation $\tau \succ \tau'$ to indicate a preference for trajectory τ over τ' . The initial state is sampled from p_0 , while p defines the transition dynamics: $s_{t+1} \sim p(\cdot | s_t, a_t)$. Finally, the function ϕ captures the preference feedback generation mechanism: $\phi(\tau, \tau') := P(\tau \succ \tau') \in [0, 1]$.

A *policy*, $\pi : \mathcal{S} \times \{1, \dots, h\} \rightarrow \mathcal{A}$, is a (possibly-stochastic) mapping from states and time indices to actions. In each iteration i , the agent selects two policies, π_{i1} and π_{i2} , which are rolled out to obtain trajectories τ_{i1}

and τ_{i2} and preference label y_i . We represent each trajectory as a feature vector, where the features record the number of times each state-action pair is visited. In iteration i , rolled-out trajectories τ_{i1} and τ_{i2} correspond, respectively, to feature vectors $\mathbf{x}_{i1}, \mathbf{x}_{i2} \in \mathbb{R}^d$, where $d := SA$ is the total number of state-action pairs, and the k^{th} element of \mathbf{x}_{ij} , $j \in \{1, 2\}$, is the number of times that τ_{ij} visits state-action pair k . The preference for iteration i is denoted $y_i := \mathbb{I}_{[\tau_{i2} \succ \tau_{i1}]} - \frac{1}{2} \in \{-\frac{1}{2}, \frac{1}{2}\}$, where $\mathbb{I}_{[\cdot]}$ denotes the indicator function, so that $P(y_i = \frac{1}{2}) = 1 - P(y_i = -\frac{1}{2}) = \phi(\tau_{i2}, \tau_{i1}) - \frac{1}{2}$; there are no ties in any comparisons. Lastly, we define $\mathbf{x}_i := \mathbf{x}_{i2} - \mathbf{x}_{i1}$.

Our analysis builds upon two main assumptions. Firstly, we assume the existence of underlying utilities, quantifying the user’s satisfaction with each trajectory:

Assumption 1. *Each trajectory τ has utility $\bar{r}(\tau)$, which decomposes additively: $\bar{r}(\tau) \equiv \sum_{t=1}^h \bar{r}(s_t, a_t)$ for the state-action pairs in τ . Defining $\bar{\mathbf{r}} \in \mathbb{R}^d$ as the vector of all state-action rewards, $\bar{r}(\tau)$ can also be expressed in terms of τ ’s state-action visit counts \mathbf{x} : $\bar{r}(\tau) = \bar{\mathbf{r}}^T \mathbf{x}$.*

Secondly, we assume that the utilities $\bar{r}(\tau)$ are stochastically translated to preferences via the noise model ϕ , such that the probability of observing $\tau_{i2} \succ \tau_{i1}$ is a function of the *difference* in their utilities. Intuitively, the greater the disparity in two trajectories’ utilities, the more accurate the user’s preference between them:

Assumption 2. $P(\tau_{i2} \succ \tau_{i1}) = \phi(\tau_{i2}, \tau_{i1}) = g(\bar{r}(\tau_{i2}) - \bar{r}(\tau_{i1})) + \frac{1}{2} = g(\bar{\mathbf{r}}^T \mathbf{x}_{i2} - \bar{\mathbf{r}}^T \mathbf{x}_{i1}) + \frac{1}{2}$, where $g : \mathbb{R} \rightarrow [-\frac{1}{2}, \frac{1}{2}]$ is a link function such that a) g is non-decreasing, and b) $g(x) = -g(-x)$ to ensure that $P(\tau \succ \tau') = 1 - P(\tau' \succ \tau)$. Note that if $\bar{r}(\tau) = \bar{r}(\tau')$, we have $P(\tau \succ \tau') = \frac{1}{2}$, and that $P(\tau_{i2} \succ \tau_{i1}) > \frac{1}{2} \Leftrightarrow g(\bar{\mathbf{r}}^T \mathbf{x}_i) > 0 \Leftrightarrow \bar{\mathbf{r}}^T \mathbf{x}_{i2} > \bar{\mathbf{r}}^T \mathbf{x}_{i1}$.

For noiseless preferences, $g_{\text{ideal}}(x) := \mathbb{I}_{[x > 0]} - \frac{1}{2}$. Alternatively, the logistic or Bradley-Terry link function is defined as $g_{\log}(x) := [1 + \exp(-x/c)]^{-1} - \frac{1}{2}$ with “temperature” $c \in (0, \infty)$. Our theoretical analysis assumes the linear link function (Ailon et al., 2014): $g_{\text{lin}}(x) := cx$, for $c > 0$ and $x \in [-\frac{1}{2c}, \frac{1}{2c}]$. Then, $\mathbb{E}[y_i] = P(\tau_{i2} \succ \tau_{i1}) - \frac{1}{2} = c\bar{\mathbf{r}}^T(\mathbf{x}_{i2} - \mathbf{x}_{i1})$. Without loss of generality, we set $c = 1$ by subsuming c into $\bar{\mathbf{r}}$. Denote the observation noise associated with g_{lin} on iteration i as η_i , such that $y_i = \bar{\mathbf{r}}^T(\mathbf{x}_{i2} - \mathbf{x}_{i1}) + \eta_i$.

Given a policy π , we can define the standard RL value function as the expected total utility when starting in state s at step j , and following π :

$$V_{\pi,j}(s) = \mathbb{E} \left[\sum_{t=j}^h \bar{r}(s_t, \pi(s_t, t)) \mid s_j = s \right]. \quad (1)$$

The optimal policy π^* is then defined as one

that maximizes the expected value over all input states: $\pi^* = \sup_{\pi} \sum_{s \in \mathcal{S}} p_0(s) V_{\pi,1}(s)$. Note that $\mathbb{E}_{s_1 \sim p_0} [V_{\pi,1}(s_1)] \equiv \mathbb{E}_{\tau \sim (\pi, \mathcal{M})} [\bar{r}(\tau)]$. Given fully specified dynamics and rewards, p and \bar{r} , it is straightforward to apply standard dynamic programming approaches such as value iteration to arrive at the optimal policy under p and \bar{r} . The learning goal, then, is to infer p and \bar{r} to the extent necessary for good decision-making.

Learning problem. We quantify the learning agent’s performance via its cumulative T -step Bayesian regret relative to the optimal policy:

$$\mathbb{E}[\text{REG}(T)] = \mathbb{E} \left\{ \sum_{i=1}^{\lceil T/(2h) \rceil} \sum_{s \in \mathcal{S}} p_0(s) [2V_{\pi^*,1}(s) - V_{\pi_{i1},1}(s) - V_{\pi_{i2},1}(s)] \right\}. \quad (2)$$

To minimize regret, the agent must balance exploration (collecting new data) with exploitation (behaving optimally given current knowledge). Over-exploration of bad trajectories will incur large regret, and under-exploration can prevent convergence to optimality. In contrast to the standard regret formulation in RL, at each iteration we measure regret of both selected policies.

Assumptions. We make two further assumptions. The first imposes a regularity condition upon the noise η_i :

Assumption 3. *The label noise $\eta_i = y_i - \bar{\mathbf{r}}^T \mathbf{x}_i$ is conditionally R -sub-Gaussian, that is, there exists $R \geq 0$ such that $\forall \lambda \in \mathbb{R}$:*

$$\mathbb{E} [e^{\lambda \eta_i} \mid \mathbf{x}_1, \dots, \mathbf{x}_{i-1}, \eta_1, \dots, \eta_{i-1}] \leq \exp \left(\frac{\lambda^2 R^2}{2} \right).$$

Note that bounded, zero-mean noise lying in an interval of length at most $2R$ is R -sub-Gaussian, and that sub-Gaussianity requires $\mathbb{E}[\eta_i \mid \mathbf{x}_1, \dots, \mathbf{x}_i, \eta_1, \dots, \eta_{i-1}] = 0$ (Abbasi-Yadkori et al., 2011). Since $y_i \in \{-\frac{1}{2}, \frac{1}{2}\}$ and $\mathbb{E}[y_i \mid \mathbf{x}_i] = \bar{\mathbf{r}}^T \mathbf{x}_i \in [-\frac{1}{2}, \frac{1}{2}]$, we must have $\eta_i \in [-1, 1]$. Thus, η_i is R -sub-Gaussian with $R \leq 1$, provided that $\mathbb{E}[\eta_i \mid \mathbf{x}_1, \dots, \mathbf{x}_i, \eta_1, \dots, \eta_{i-1}] = 0$. The latter holds by the assumption that $\mathbb{E}[y_i \mid \mathbf{x}_i] = \bar{\mathbf{r}}^T \mathbf{x}_i$.

Assumption 4. *For some known $S_r < \infty$, $\|\bar{\mathbf{r}}\|_2 \leq S_r$.*

Additional notation. For random variables X and X_n , $n \in \mathbb{N}$, $X_n \xrightarrow{D} X$ denotes that X_n converges to X in distribution. For $\mathbf{x} \in \mathbb{R}^d$ and positive definite matrix $B \in \mathbb{R}^{d \times d}$, we define the norm $\|\mathbf{x}\|_B := \sqrt{\mathbf{x}^T B \mathbf{x}}$.

4 ALGORITHM

As outlined in Algorithm 1, DUELING POSTERIOR SAMPLING (DPS) iterates among three steps: (a) sampling two policies π_{i1}, π_{i2} from the Bayesian posteriors

Algorithm 1 DUELING POSTERIOR SAMPLING (DPS)

$\mathcal{H}_0 = \emptyset$ {Initialize history}
 Initialize prior for f_p {Initialize state transition model}
 Initialize prior for f_r {Initialize utility model}
for $i = 1, 2, \dots$ **do**
 $\pi_{i1} \leftarrow \text{ADVANCE}(f_p, f_r)$
 $\pi_{i2} \leftarrow \text{ADVANCE}(f_p, f_r)$
 Sample trajectories τ_{i1} and τ_{i2} from π_{i1} and π_{i2}
 Observe feedback $y_i = \mathbb{I}_{[\tau_{i2} > \tau_{i1}]} - \frac{1}{2}$
 $\mathcal{H}_i = \mathcal{H}_{i-1} \cup (\tau_{i1}, \tau_{i2}, y_i)$
 $f_p, f_r = \text{FEEDBACK}(\mathcal{H}_i, f_p, f_r)$
end for

of the dynamics and utility models (ADVANCE – Algorithm 2); (b) rolling out π_{i1} and π_{i2} to obtain trajectories τ_{i1} and τ_{i2} , and receiving a preference y_i between them; and (c) updating the posterior (FEEDBACK – Algorithm 3). In contrast to conventional posterior sampling with absolute feedback, DPS samples two policies rather than one at each iteration and solves a credit assignment problem to learn from feedback.

ADVANCE (Algorithm 2) samples from the Bayesian posteriors of the dynamics and utility models to select a policy to roll out. The sampled dynamics and utilities form an MDP, for which value iteration derives the optimal policy π under the sample. One can also view π as a random function whose randomness depends on the sampling of the dynamics and utility models. In the Bayesian setting, it can be shown that π is sampled according to its posterior probability of being the optimal policy π^* . Intuitively, peaked (i.e., certain) posteriors lead to less variability when sampling π , which implies less exploration, while diffuse (i.e., uncertain) posteriors lead to greater variability when sampling π , implying more exploration.

FEEDBACK (Algorithm 3) updates the Bayesian posteriors of the dynamics and utility models based on new data. Updating the dynamics posterior is relatively straightforward, as we assume that the dynamics are fully-observed; we model the dynamics prior via a Dirichlet distribution for each state-action pair, with conjugate multinomial observation likelihoods. In contrast, performing Bayesian inference over state-action utilities from trajectory-level feedback is much more challenging. We consider a range of approaches (see Appendix B), and found Bayesian linear regression (Section 4.1) to both perform well and admit tractable analysis within our theoretical framework.

4.1 BAYESIAN LINEAR REGRESSION FOR UTILITY INFERENCE AND CREDIT ASSIGNMENT

Credit assignment is the problem of inferring which state-action pairs are responsible for observed trajectory-

Algorithm 2 ADVANCE: Sample policy from dynamics and utility models

Input: f_p, f_r
 Sample $\tilde{p} \sim f_p(\cdot)$ {Sample MDP transition dynamics parameters from posterior}
 Sample $\tilde{r} \sim f_r(\cdot)$ {Sample utilities from posterior}
 Compute $\pi = \text{argmax}_{\pi} V(\tilde{p}, \tilde{r})$ {Value iteration yields sampled MDP’s optimal policy}
 Return π

Algorithm 3 FEEDBACK: Update dynamics and utility models based on new user feedback

Input: history \mathcal{H}, f_p, f_r
 Apply Bayesian update to f_p , given \mathcal{H} {Update dynamics model given history}
 Apply Bayesian update to f_r , given \mathcal{H} {Update utility model given preferences}
 Return f_p, f_r

level preferences. We detail a Bayesian linear regression approach to addressing this task in our setting.

Let n be the number of iterations, or trajectory pairs, observed so far. Then, the maximum a posteriori (MAP) estimate of the rewards \bar{r} is calculated via ridge regression, similarly to algorithms for the linear bandit setting:

$$\hat{r}_n = M_n^{-1} \sum_{i=1}^{n-1} y_i \mathbf{x}_i, \text{ where} \quad (3)$$

$$M_n = \lambda I + \sum_{i=1}^{n-1} \mathbf{x}_i \mathbf{x}_i^T, \text{ and } \lambda \geq 1. \quad (4)$$

We perform Thompson sampling as in Agrawal and Goyal (2013) and Abeille and Lazaric (2017), such that in iteration n , rewards are sampled from the distribution:

$$\tilde{r}_{n1}, \tilde{r}_{n2} \sim \mathcal{N}(\hat{r}_n, \beta_n(\delta)^2 M_n^{-1}), \text{ where} \quad (5)$$

$$\begin{aligned}
 \beta_n(\delta) &= R \sqrt{2 \log \left(\frac{\det(M_n)^{1/2} \lambda^{-d/2}}{\delta} \right)} + \sqrt{\lambda} S_r \\
 &\leq R \sqrt{d \log \left(\frac{1 + \frac{L^2 n}{d\lambda}}{\delta} \right)} + \sqrt{\lambda} S_r,
 \end{aligned}$$

and where $\delta \in (0, 1)$ is a failure probability and for all n , $\|\mathbf{x}_n\|_2 \leq L$. Note that $L \leq 2h$, since $\|\mathbf{x}_n\|_2 = \|\mathbf{x}_{n2} - \mathbf{x}_{n1}\|_2 \leq \|\mathbf{x}_{n2} - \mathbf{x}_{n1}\|_1 \leq \|\mathbf{x}_{n2}\|_1 + \|\mathbf{x}_{n1}\|_1 = 2h$.

The factor $\beta_n(\delta)$, introduced in Abbasi-Yadkori et al. (2011), is critical to deriving the theoretical guarantees for posterior sampling with linear bandits in Agrawal and Goyal (2013) and Abeille and Lazaric (2017), due to their dependence on Theorems 1 and 2 of Abbasi-Yadkori et al. (2011). Our analysis invokes these results as well. Both of the theorems require any noise in the labels y_n to be sub-Gaussian; in our case, sub-Gaussianity

holds by Assumption 3, as we adopted the linear preference noise model with link function g_{lin} .

Our theoretical analysis is quite different from that for linear bandits in Agrawal and Goyal (2013) and Abeille and Lazaric (2017), because in our setting, observations \mathbf{x}_n are *differences* of trajectory feature vectors, policies are chosen via value iteration, and trajectories are obtained by rolling out RL policies while subject to the environment’s state transition dynamics.

5 THEORETICAL RESULTS

This section sketches our analysis of the asymptotic Bayesian regret of DPS under a Bayesian linear regression credit assignment model. Appendix A details the full proof, while Appendix B.4 discusses possible future extensions to additional credit assignment models.

The analysis follows three main steps: 1) we prove that DPS is asymptotically-consistent, that is, the probability with which DPS selects the optimal policy approaches 1 over time (Appendix A.1); 2) we asymptotically bound the one-sided Bayesian regret for π_{i2} under the setting where, at each iteration i , DPS only selects policy π_{i2} , while policy π_{i1} is sampled from a fixed distribution over policies (Appendix A.2); and lastly, 3) we assume DPS selects policy π_{i2} , while the π_{i1} -distribution is drifting but converging, and then we asymptotically bound the one-sided regret for π_{i2} (Appendix A.3). Due to the asymptotic consistency shown in 1), the policies are indeed sampled from converging distributions, and so the asymptotic regret rate in 3) holds.

This outline is inspired by the analysis for Self-Sparring (Sui et al., 2017); however, because their guarantee is for dueling bandits with independent Beta-Bernoulli reward models for each action, the details of our analysis are completely different from theirs. Below, we give intuition for each of the three portions of the proof.

Asymptotic consistency of DPS. To prove that DPS is asymptotically consistent, we first prove that samples of the dynamics and reward parameters converge in distribution to their true values:

Proposition 1. *The sampled dynamics converge in distribution to their true values as the DPS iteration increases.*

Proof sketch. Applying standard concentration inequalities to the Dirichlet dynamics posterior, one can show that the sampled dynamics converge in distribution to their true values if every state-action pair is visited infinitely-often. The latter condition can be proven via contradiction: assuming that certain state-action pairs are visited finitely-often, DPS does not receive new in-

formation about their rewards. Examining their reward posteriors, we show that DPS is guaranteed to eventually sample high enough rewards in the unvisited state-actions that its policies will attempt to reach them.

We also show that with high probability, the sampled rewards exhibit asymptotic consistency:

Proposition 2. *With probability $1 - \delta$, where δ is a parameter of the Bayesian linear regression model, the sampled rewards converge in distribution to the true reward parameters, $\bar{\mathbf{r}}$, as the DPS iteration increases.*

Proof sketch. We leverage Theorem 2 from Abbasi-Yadkori et al. (2011) (Lemma 4 in Appendix A.4): under stated conditions and for any $\delta > 0$, with probability $1 - \delta$ and for all $i > 0$, $\|\hat{\mathbf{r}}_i - \bar{\mathbf{r}}\|_{M_i} \leq \beta_i(\delta)$. This result defines a high-confidence ellipsoid, which can be linked to the posterior sampling distribution. We demonstrate that it suffices to show that all eigenvalues of the posterior covariance matrix, $\beta_i(\delta)^2 M_i^{-1}$, converge in distribution to zero. This statement is proven via contradiction: we analyze the behavior of posterior sampling if this does not hold. The probability of failure δ comes entirely from Theorem 2 in Abbasi-Yadkori et al. (2011).

From the asymptotic consistency of the dynamics and reward samples, it is straightforward to show that the sampled policies converge to the optimal policy:

Theorem 1. *With probability $1 - \delta$, the sampled policies π_{i1}, π_{i2} converge in distribution to the optimal policy, π^* , as $i \rightarrow \infty$. That is, $P(\pi_{i1} = \pi^*) \rightarrow 1$ and $P(\pi_{i2} = \pi^*) \rightarrow 1$ as $i \rightarrow \infty$.*

Bounding the one-sided regret under a fixed π_{i1} -distribution. To analyze the Bayesian regret of DPS, we adapt the information-theoretic posterior sampling analysis in Russo and Van Roy (2016) to the PBRL setting. In comparison to Russo and Van Roy’s work, this requires accounting for preference feedback and incorporating state transition dynamics. Their analysis hinges upon defining a quantity called the *information ratio*, which captures the trade-off between exploration and exploitation. In our setting, we define the information ratio corresponding to the one-sided regret of π_{i2} as:

$$\Gamma_i := \frac{\mathbb{E}_i[y_i^* - y_i]^2}{I_i(\pi^*; (\pi_{i2}, \tau_{i1}, \tau_{i2}, \mathbf{x}_{i2} - \mathbf{x}_{i1}, y_i))},$$

where y_i is the label in iteration i , y_i^* is the label in iteration i given $\pi_{i2} = \pi^*$, $I(\cdot; \cdot)$ denotes mutual information, and the subscripts i in $\mathbb{E}_i[\cdot]$ and $I_i(\cdot; \cdot)$ indicate conditioning upon the history, as formalized in Appendix A.2. The ratio Γ_i is between the squared instantaneous one-sided regret of π_{i2} (exploitation) and the information gained about the optimal policy (exploration).

When π_{i1} is drawn from a fixed distribution, we show that analogously to Russo and Van Roy (2016), the Bayesian one-sided regret $\mathbb{E}[\text{REG}_2(T)]$ for π_{i2} can be bounded in terms of an upper bound on Γ_i :

Lemma 12. *If $\Gamma_i \leq \bar{\Gamma}$ almost surely for each $i \in \{1, \dots, N\}$, where N is the number of DPS iterations (over which the policies π_{i2} take $T = Nh$ actions), then:*

$$\mathbb{E}[\text{REG}_2(T)] = \mathbb{E}[\text{REG}_2(Nh)] \leq \sqrt{\bar{\Gamma}H(\pi^*)N},$$

where $H(\pi^*)$ is the entropy of the optimal policy π^* . Because there are at most A^{Sh} deterministic policies, $H(\pi^*) \leq \log |A^{Sh}| = Sh \log A$. Substituting this,

$$\mathbb{E}[\text{REG}_2(T)] \leq \sqrt{\bar{\Gamma}ShN \log A} = \sqrt{\bar{\Gamma}ST \log A}.$$

We show that Γ_i can be asymptotically upper-bounded such that $\lim_{i \rightarrow \infty} \Gamma_i \leq \frac{SA}{2}$, and consequently:

Theorem 2. *If the policy π_{i1} is drawn from a fixed distribution for all i , then for the competing policy π_{i2} , DPS achieves a one-sided asymptotic Bayesian regret rate of $S\sqrt{\frac{AT \log A}{2}}$.*

The bounds in Lemma 12 and Theorem 2 are asymptotic rather than finite-time, due to the convergence in distribution of the dynamics. If the dynamics are known a priori, then these would be finite-time guarantees; in fact, to prove Lemma 12, we first show that under known dynamics, $\Gamma_i \leq \frac{SA}{2}$ for all i , and then extend the analysis to prove that under converging dynamics, the result still holds asymptotically. Note that in the PBRL setting, it is significantly more difficult to learn the rewards via credit assignment than to learn the dynamics, which are assumed to be fully-observed. Thus, in practice, we expect that DPS would learn the dynamics much faster than the rewards, and so it is reasonable to consider convergence of the dynamics model only asymptotically.

Bounding the one-sided regret under a converging π_{i1} -distribution. Finally, we assume that the distribution of π_{i1} is no longer fixed, but rather converges to some fixed distribution over deterministic policies. To asymptotically bound the one-sided regret incurred by π_{i2} , we leverage that when two discrete random variables converge in distribution, such that $X_n \xrightarrow{D} X$ and $Y_n \xrightarrow{D} Y$, their mutual information also converges: $\lim_{n \rightarrow \infty} I(X_n, Y_n) = I(X, Y)$. This fact allows us to bound the one-sided regret for π_{i2} as follows:

Lemma 17. *Assume that the sampling distribution of π_{i1} converges to a fixed probability distribution. Then, the information ratio Γ_i corresponding to π_{i2} 's one-sided regret $\mathbb{E}[\text{REG}_2(T)]$ satisfies $\lim_{i \rightarrow \infty} \Gamma_i \leq \frac{SA}{2}$.*

Combining Lemma 17 with the asymptotic consistency of sampled policies as shown in Theorem 1, $P(\pi_{i1} = \pi^*) \rightarrow 1$, yields our main theoretical result:

Theorem 3. *With probability $1 - \delta$, where δ is a parameter of the Bayesian linear regression model, the expected Bayesian regret $\mathbb{E}[\text{REG}(T)]$ of DPS achieves an asymptotic rate of $S\sqrt{2AT \log A}$.*

Discussion. The specific theoretical results presented yield a high-probability asymptotic Bayesian no-regret rate for DPS under Bayesian linear regression credit assignment. The proof consists of first demonstrating that the algorithm is asymptotically consistent, and then analyzing its information ratio to characterize the Bayesian regret. We adopted this information-theoretic perspective because we found it more amenable to preference-based feedback than other prevalent methods from the linear bandits literature.

In particular, while several existing regret analyses for posterior sampling with linear bandits (Agrawal and Goyal, 2013; Abeille and Lazaric, 2017) are based upon martingale concentration properties derived in Abbasi-Yadkori et al. (2011), we found that these techniques cannot readily extend to the preference-feedback setting (Appendix C). These linear bandit analyses assume that each observation \mathbf{x}_i that incurs regret contributes fully toward learning the rewards. In contrast, we assume that while regret is incurred with respect to the observations $\mathbf{x}_{i1}, \mathbf{x}_{i2}$, learning occurs only with respect to observation differences, $\mathbf{x}_i = \mathbf{x}_{i2} - \mathbf{x}_{i1}$. In preference-based learning settings, it is common to make such assumptions as $P(\tau_{i2} \succ \tau_{i1}) = g(\mathbf{x}_{i2} - \mathbf{x}_{i1})$, for some function g . In comparison to the martingale-based techniques, the information ratio provides a more direct method for quantifying the trade-off between exploration and exploitation.

Theoretically analyzing other credit assignment models, in addition to Bayesian linear regression, is an important direction for future work. We conjecture that our proof methodology could extend toward other asymptotically-consistent credit assignment models. Indeed, recent work (Dong and Van Roy, 2018) has analyzed the information ratio for more general link functions, including for logistic bandits. It would be interesting to study the information ratio's behavior under general link functions, as well as to characterize its relationship to the dynamics model's convergence. It would also be interesting to develop methodology for extending the analysis to achieve finite-time convergence guarantees.

6 EXPERIMENTS

We validate the empirical performance of DPS in three simulated domains with varying degrees of preference




Figure 1: Empirical performance of DPS; each simulated environment is shown under the two least-noisy user preference models evaluated. The plots show DPS with three credit assignment models: Gaussian process regression (GPR), Bayesian linear regression, and a Gaussian process preference model. PSRL is an upper bound that receives numerical rewards, while EPMC is a baseline. Plots display the mean \pm one standard deviation over 100 runs of each algorithm tested. The remaining user noise models are plotted in Appendix D. For RiverSwim and Random MDPs, normalization is with respect to the total reward achieved by the optimal policy. Overall, we see that DPS performs well and is robust to the choice of credit assignment model.

noise and using three alternative credit assignment models. We find that DPS generally performs well and compares favorably against standard PBRL baselines.

Experimental setup. We evaluate on three simulated environments: RiverSwim and random MDPs (described in Osband et al. (2013)) and the Mountain Car problem as detailed in Wirth (2017). The RiverSwim environment has six states and two actions (actions 0 and 1); the optimal policy always chooses action 1, which maximizes the probability of reaching a goal state-action pair. Meanwhile, a suboptimal policy—yielding a small reward compared to the goal—is quickly and easily discovered and incentivizes the agent to always select action 0. The algorithm must demonstrate sufficient exploration to have hope of discovering the optimal policy quickly.

In the second environment, we generate random MDPs with 10 states and 5 actions. The transition dynamics and rewards are respectively generated from Dirichlet (all parameters set to 0.1) and exponential (rate parameter = 5) distributions. These distribution parameters were cho-

sen to generate MDPs with sparse dynamics and rewards. For each random MDP, the sampled reward values were shifted and normalized so that the minimum reward is zero and their mean is one.

Thirdly, in the Mountain Car problem, an under-powered car in a valley must reach the top of a hill by accelerating in both directions to build its momentum. The state space is two-dimensional (position and velocity), while there are three actions (left, right, and neutral). Our implementation begins each episode in a uniformly-random state and has a maximum episode length of 500. We discretize the state space into 10 states in each dimension. Each episode terminates either when the car reaches the goal or after 500 steps, and rewards are -1 in every step.

In each environment, preferences between trajectory pairs were generated by (noisily) comparing their total accrued rewards; this reward information was hidden from the learning algorithm, which observed only the trajectory preferences and state transitions. For trajectories τ_i and τ_j with total rewards $\bar{r}(\tau_i)$ and $\bar{r}(\tau_j)$, we con-

sider two models for generating preferences: a) a logistic model, $P(\tau_i \succ \tau_j) = \{1 + \exp[-(\bar{r}(\tau_i) - \bar{r}(\tau_j))/c]\}^{-1}$, and b) a linear model, $P(\tau_i \succ \tau_j) = (\bar{r}(\tau_i) - \bar{r}(\tau_j))/c$, where in both cases, the temperature c controls the degree of noisiness. In the linear case, c is assumed to be large enough that $P(\tau_i \succ \tau_j) \in [0, 1]$. Note that in ties where $\bar{r}(\tau_i) = \bar{r}(\tau_j)$, preferences are uniformly-random.

Methods compared. We evaluate DPS under three credit assignment models (Appendix B): 1) Bayesian linear regression, 2) Gaussian process regression, and 3) a Gaussian process preference model. User noise generated via the logistic model has noise levels: $c \in \{10, 2, 1, 0.001\}$ for RiverSwim and random MDPs and $c \in \{100, 20, 10, 0.001\}$ for the Mountain Car. We selected higher values of c for the Mountain Car because $|\bar{r}(\tau_i) - \bar{r}(\tau_j)|$ has a wider range. Additionally, we evaluate the linear preference noise model with $c = 2h\Delta\bar{r}$, where $\Delta\bar{r}$ is the difference between the maximum and minimum element of \bar{r} for each MDP; this choice of c guarantees that $P(\tau_i \succ \tau_j) \in [0, 1]$, but yields noisier preferences than the logistic noise models considered.

As discussed in Section 2, many existing PBRL algorithms handle a somewhat distinct setting from ours, as they assume access to a simulator between preference queries and/or prioritize minimizing preference queries rather than online regret. As a baseline, we evaluate the Every-Visit Preference Monte Carlo (EPMC) algorithm with probabilistic credit assignment (Wirth and Furnkranz, 2013b; Wirth, 2017). While EPMC does not require simulations between preference queries, it has several limitations, including: 1) the exploration approach always takes uniformly-random actions with some probability, and thus, the authors’ plots do not depict online reward accumulation, and 2) EPMC assumes that compared trajectories start in the same state. Lastly, we compare against the posterior sampling RL algorithm (PSRL) from Osband et al. (2013), which receives the true numerical rewards at each step, and thus upper-bounds the achievable performance of a preference-based algorithm.

Results. Figure 1 depicts performance curves for the three environments, each with two noise models (Appendix D contains additional results and details). DPS performs well in all simulations, and significantly outperforms the EPMC baseline. In RiverSwim, most credit assignment models perform best in the second-to-least-noisy case (logistic noise, $c = 1$), since it is harder to escape the local minimum under the least-noisy preferences. We also see that DPS is competitive with PSRL, which has access to the full cardinal rewards at each state-action. Additionally, while our theoretical guarantees for DPS assume fixed-horizon episodes, the Moun-

tain Car results demonstrate that it also succeeds with variable episode lengths. Finally, the performance of DPS is robust to the choice of credit assignment model, and in fact using Gaussian processes (for which we do not have an end-to-end regret analysis) often leads to the best empirical performance. These results suggest that DPS is a practically-promising approach that can robustly incorporate many models as subroutines.

7 CONCLUSION

This work investigates the preference-based reinforcement learning problem, in which an RL agent receives comparative preferences instead of absolute real-valued rewards as feedback. We develop the DUELING POSTERIOR SAMPLING (DPS) algorithm, which optimizes policies in a highly efficient and flexible way. To our knowledge, DPS is the first preference-based RL algorithm with a regret guarantee. DPS also performs well in our simulations, making it both a theoretically-justified and practically-promising algorithm.

There are many directions for future work. Assumptions governing the user’s preferences, such as requiring an underlying utility model, could be relaxed. It would also be interesting to extend our theoretical analysis to additional credit assignment approaches and to pursue finite-time guarantees. We expect that DPS would perform well with any asymptotically-consistent reward model that sufficiently captures users’ preference behavior, and hope to develop models that are tractable with larger state and action spaces. For instance, incorporating kernelized input spaces could further improve sample efficiency.

Acknowledgments

This work was supported by NIH grant EB007615 and an Amazon graduate fellowship.

References

- Y. Abbasi-Yadkori, D. Pal, and C. Szepesvari. Improved algorithms for linear stochastic bandits. In *NeurIPS*, 2011.
- M. Abeille and A. Lazaric. Linear Thompson sampling revisited. *Electronic Journal of Statistics*, 11(2), 2017.
- S. Agrawal and N. Goyal. Thompson sampling for contextual bandits with linear payoffs. In *ICML*, 2013.
- S. Agrawal and R. Jia. Optimistic posterior sampling for reinforcement learning: Worst-case regret bounds. In *NeurIPS*, 2017.
- N. Ailon, Z. Karnin, and T. Joachims. Reducing dueling bandits to cardinal bandits. In *ICML*, 2014.
- R. Akrou, M. Schoenauer, and M. Sebag. APRIL: Active preference learning-based reinforcement learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2012.

- R. Akrou, M. Schoenauer, M. Sebag, and J.-C. Souplet. Programming by feedback. In *ICML*, volume 32, 2014.
- D. Amodei, C. Olah, J. Steinhardt, P. Christiano, et al. Concrete problems in AI safety. *arXiv preprint*, 2016.
- B. D. Argall, S. Chernova, M. Veloso, and B. Browning. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5), 2009.
- C. Basu, Q. Yang, D. Hungerman, M. Sinahal, and A. D. Dragan. Do you want your autonomous car to drive like you? In *Int. Conf. on Human-Robot Interaction*. IEEE, 2017.
- P. Billingsley. *Convergence of probability measures*. John Wiley & Sons, 1968.
- C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, et al. Learning to rank using gradient descent. In *ICML*, 2005.
- C. J. Burges, R. Ragno, and Q. V. Le. Learning to rank with nonsmooth cost functions. In *NeurIPS*, 2007.
- R. Busa-Fekete, B. Szörényi, P. Weng, W. Cheng, and E. Hüllermeier. Preference-based evolutionary direct policy search. In *ICRA Workshop on Autonomous Learning*, 2013.
- O. Chapelle and L. Li. An empirical evaluation of Thompson sampling. In *NeurIPS*, 2011.
- P. F. Christiano, J. Leike, T. Brown, et al. Deep reinforcement learning from human preferences. In *NeurIPS*, 2017.
- W. Chu and Z. Ghahramani. Preference learning with Gaussian processes. In *ICML*, 2005.
- T. M. Cover and J. A. Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- S. Dong and B. Van Roy. An information-theoretic analysis for Thompson sampling with many actions. In *NeurIPS*, 2018.
- S. Dong, T. Ma, and B. Van Roy. On the performance of Thompson sampling on logistic bandits. In *COLT*, 2019.
- M. Dudík, K. Hofmann, R. E. Schapire, A. Slivkins, and M. Zoghi. Contextual dueling bandits. In *COLT*, 2015.
- B. Eric, N. D. Freitas, and A. Ghosh. Active preference learning with discrete choice data. In *NeurIPS*, 2008.
- J. Fürnkranz, E. Hüllermeier, W. Cheng, et al. Preference-based reinforcement learning: A formal framework and a policy iteration algorithm. *Machine learning*, 89(1-2), 2012.
- A. Gopalan and S. Mannor. Thompson sampling for learning parameterized Markov decision processes. In *COLT*, 2015.
- R. Herbrich, T. Graepel, and K. Obermayer. Support vector learning for ordinal regression. In *International Conference on Artificial Neural Networks*, volume 1. IET, 1999.
- N. Houlsby, F. Huszár, et al. Bayesian active learning for classification and preference learning. *arXiv*, 2011.
- T. Jaksch, R. Ortner, and P. Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11, 2010.
- T. Joachims. A support vector method for multivariate performance measures. In *ICML*. ACM, 2005.
- A. Kupcsik, D. Hsu, and W. S. Lee. Learning dynamic robot-to-human object handover from human feedback. In *Robotics research*. Springer, 2018.
- T.-Y. Liu. Learning to rank for information retrieval. *Foundations and Trends® in Information Retrieval*, 3(3), 2009.
- N. Nikolov, J. Kirschner, et al. Information-directed exploration for deep reinforcement learning. *arXiv*, 2018.
- I. Osband and B. Van Roy. Why is posterior sampling better than optimism for reinforcement learning? In *ICML*, 2017.
- I. Osband, D. Russo, et al. (More) efficient reinforcement learning via posterior sampling. In *NeurIPS*, 2013.
- F. Radlinski and T. Joachims. Query chains: Learning to rank from implicit feedback. In *SIGKDD*. ACM, 2005.
- S. Y. Ramamohan, A. Rajkumar, and S. Agarwal. Dueling bandits: Beyond Condorcet winners to general tournament solutions. In *NeurIPS*, 2016.
- K. Raman, T. Joachims, P. Shivaswamy, and T. Schnabel. Stable coactive learning via perturbation. In *ICML*, 2013.
- C. E. Rasmussen and C. K. Williams. Gaussian processes for machine learning. *The MIT Press*, 2(3):4, 2006.
- D. Russo and B. Van Roy. An information-theoretic analysis of Thompson sampling. *The Journal of Machine Learning Research*, 17(1), 2016.
- D. Sadigh, A. D. Dragan, S. Sastry, and S. A. Seshia. Active preference-based learning of reward functions. In *Robotics: Science and Systems*, 2017.
- P. Shivaswamy and T. Joachims. Online structured prediction via coactive learning. In *ICML*, 2012.
- P. Shivaswamy and T. Joachims. Coactive learning. *Journal of Artificial Intelligence Research*, 53, 2015.
- Y. Sui, V. Zhuang, J. W. Burdick, and Y. Yue. Multi-dueling bandits with dependent arms. In *UAI*, 2017.
- Y. Sui, V. Zhuang, J. Burdick, et al. Stagewise safe Bayesian optimization with Gaussian processes. In *ICML*, 2018a.
- Y. Sui, M. Zoghi, K. Hofmann, and Y. Yue. Advancements in dueling bandits. In *IJCAI*, 2018b.
- B. Szörényi, R. Busa-Fekete, A. Paul, and E. Hüllermeier. Online rank elicitation for Plackett-Luce: A dueling bandits approach. In *NeurIPS*, 2015.
- W. R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4), 1933.
- A. Wilson, A. Fern, et al. A Bayesian approach for policy learning from trajectory preference queries. In *NeurIPS*, 2012.
- C. Wirth. *Efficient Preference-based Reinforcement Learning*. PhD thesis, Technische Universität, 2017.
- C. Wirth and J. Fürnkranz. EPMC: Every visit preference Monte Carlo for reinforcement learning. In *Asian Conference on Machine Learning*, 2013a.
- C. Wirth and J. Fürnkranz. A policy iteration algorithm for learning from preference-based feedback. In *International Symposium on Intelligent Data Analysis*. Springer, 2013b.
- C. Wirth, J. Fürnkranz, and G. Neumann. Model-free preference-based reinforcement learning. In *AAAI*, 2016.
- C. Wirth, R. Akrou, G. Neumann, and J. Fürnkranz. A survey of preference-based reinforcement learning methods. *The Journal of Machine Learning Research*, 18(1), 2017.
- H. Wu and X. Liu. Double Thompson sampling for dueling bandits. In *NeurIPS*, 2016.
- Y. Yue, T. Finley, et al. A support vector method for optimizing average precision. In *Int. SIGIR conf. on research and development in information retrieval*. ACM, 2007.

- Y. Yue, J. Broder, et al. The k-armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5), 2012.
- A. Zanette and R. Sarkar. Information directed reinforcement learning. Technical report, Technical report, 2017.
- M. Zoghi, S. Whiteson, et al. Relative upper confidence bound for the k-armed dueling bandit problem. In *ICML*, 2014.
- M. Zoghi, Z. S. Karnin, S. Whiteson, and M. De Rijke. Copeland dueling bandits. In *NeurIPS*, 2015.

APPENDICES

A DERIVATION OF THE ASYMPTOTIC REGRET RATE

As outlined in Section 5, the analysis follows three main steps:

1. Prove that DPS is asymptotically-consistent, that is, over time, the probability that DPS selects the optimal policy approaches 1 (Appendix A.1).
2. Assume that in each iteration i , policy π_{i1} is drawn from a fixed distribution while policy π_{i2} is selected by DPS. Then, asymptotically bound the one-sided regret rate for π_{i2} (Appendix A.2).
3. Assume that policy π_{i1} is drawn from a drifting but converging distribution while policy π_{i2} is selected by DPS. Then, asymptotically bound the one-sided regret rate for π_{i2} (Appendix A.3).

Finally, Appendix A.4 combines these results to asymptotically bound the expected regret rate.

A.1 ASYMPTOTIC CONSISTENCY OF DPS

We show asymptotic consistency of DPS in three parts: 1) samples from the model posterior over transition dynamics parameters converge in distribution to the true transition probabilities; 2) samples from the reward posterior converge in distribution to the true utilities; and 3) consequently, the sampled policies converge in distribution to the optimal policy.

State transition dynamics are modeled independently for each state-action pair. For a given state-action pair, a Dirichlet model estimates the probability of transitioning to each possible subsequent state. The prior and posterior distributions are both Dirichlet; because the Dirichlet and multinomial distributions are conjugate, each state-action pair’s posterior can be updated easily using the observed transitions from that state-action. Each time that DPS draws a sample from the dynamics distribution, values are sampled for all S^2A transition parameters, $\{P(s_{t+1} = s' | s_t = s, a_t = a) | s, s' \in \mathcal{S}, a \in \mathcal{A}\}$.

Notation. This section uses the following notation. Let $\bar{\mathbf{p}} \in \mathbb{R}^{S^2A}$ be the vector containing all true state transition dynamics parameters, $\{P(s_{t+1} = s' | s_t = s, a_t = a) | s, s' \in \mathcal{S}, a \in \mathcal{A}\}$. Let $\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{p}}_{i2} \in \mathbb{R}^{S^2A}$ be the two posterior samples of the transition dynamics $\bar{\mathbf{p}}$ in iteration i . Similarly, $\bar{\mathbf{r}} \in \mathbb{R}^{SA}$ is the vector of true reward parameters, while $\tilde{\mathbf{r}}_{i1}, \tilde{\mathbf{r}}_{i2} \in \mathbb{R}^{SA}$ are posterior samples of $\bar{\mathbf{r}}$ in iteration i . For a random variable X and a sequence of random variables (X_n) , $n \in \mathbb{N}$, $X_n \xrightarrow{D} X$ denotes that X_n converges to X in distribution, while $X_n \xrightarrow{P} X$ denotes that X_n converges to X in probability. Notation for the value function and for policies given by value iteration follows.

Definition 1 (Value function given transition dynamics, rewards, and a policy). *Define $V(\mathbf{p}, \mathbf{r}, \pi)$ as the value function over a length- h episode—i.e., the expected total reward in the episode—under transition dynamics $\mathbf{p} \in \mathbb{R}^{S^2A}$, rewards $\mathbf{r} \in \mathbb{R}^{SA}$, and policy π :*

$$V(\mathbf{p}, \mathbf{r}, \pi) = \sum_{s \in \mathcal{S}} p_0(s) \mathbb{E} \left[\sum_{t=1}^h \bar{r}(s_t, \pi(s_t, t)) \mid s_1 = s, \bar{\mathbf{p}} = \mathbf{p}, \bar{\mathbf{r}} = \mathbf{r} \right].$$

Definition 2 (Optimal deterministic policy given transition dynamics and rewards). *Define $\pi_{vi}(\mathbf{p}, \mathbf{r}) := \operatorname{argmax}_{\pi} V(\mathbf{p}, \mathbf{r}, \pi)$ as the optimal deterministic policy given transition dynamics $\mathbf{p} \in \mathbb{R}^{S^2A}$ and rewards $\mathbf{r} \in \mathbb{R}^{SA}$ (breaking ties randomly if multiple deterministic policies achieve the maximum). Note that $\pi_{vi}(\mathbf{p}, \mathbf{r})$ can be found via finite-horizon value iteration: defining $V_{\pi,t}(s)$ as in (1), set $V_{\pi,h+1}(s) := 0$ for each $s \in \mathcal{S}$ and use the Bellman equation to calculate $V_{\pi,t}(s)$ successively for $t \in \{h, h-1, \dots, 1\}$ given \mathbf{p} and \mathbf{r} :*

$$\pi(s, t) = \operatorname{argmax}_{a \in \mathcal{A}} \left[\bar{r}(s, a) + \sum_{s' \in \mathcal{S}} P(s_{t+1} = s' | s_t = s, a_t = a) V_{\pi,t+1}(s') \right],$$

$$V_{\pi,t}(s) = \sum_{a \in \mathcal{A}} \mathbb{I}_{[\pi(s,t)=a]} \left[\bar{r}(s, a) + \sum_{s' \in \mathcal{S}} P(s_{t+1} = s' | s_t = s, a_t = a) V_{\pi,t+1}(s') \right].$$

As value iteration results in only deterministic policies, of which there are finitely-many (more precisely, there are A^{S^h}), the maximum argument $\pi_{vi}(\mathbf{p}, \mathbf{r}) := \operatorname{argmax}_{\pi} V(\mathbf{p}, \mathbf{r}, \pi)$ is taken over a finite policy class.

Finally, we define notation for the eigenvectors and eigenvalues of the matrix $M_i := \lambda I + \sum_{k=1}^{i-1} \mathbf{x}_k \mathbf{x}_k^T$ (see Equation (4)).

Definition 3 (Eigenvalue notation). Let $\lambda_j^{(i)}$ refer to the j^{th} -largest eigenvalue of M_i , and $\mathbf{v}_j^{(i)}$ denote its corresponding eigenvector. Note that M_i^{-1} also has eigenvectors $\mathbf{v}_j^{(i)}$, with corresponding eigenvalues $\frac{1}{\lambda_j^{(i)}}$. Because M_i is positive definite, the eigenvectors $\{\mathbf{v}_j^{(i)}\}$ form an orthonormal basis, and $\lambda_j^{(i)} > 0$ for all i, j .

We demonstrate convergence in distribution of the sampled transition dynamics parameters. First, Lemma 1 shows that if every state-action pair is visited infinitely-often, the desired result holds. Then, Lemma 3 completes the argument by showing that DPS indeed visits each state-action pair infinitely-often.

Lemma 1. *If every state-action pair is visited infinitely-often, then the sampled transition dynamics parameters converge in distribution to their true values: $\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{p}}_{i2} \xrightarrow{D} \bar{\mathbf{p}}$ as $i \rightarrow \infty$, where \xrightarrow{D} denotes convergence in distribution.*

Proof. Denote the $d = SA$ state-action pairs as $\tilde{s}_1, \dots, \tilde{s}_d$. At a particular DPS episode, let n_j be the number of visits to \tilde{s}_j and n_{jk} be the number of observed transitions from \tilde{s}_j to the k^{th} subsequent state. For the j^{th} state-action pair at iteration i , let $\bar{\mathbf{p}}^{(j)}, \tilde{\mathbf{p}}^{(j)}, \hat{\mathbf{p}}^{(j)}, \hat{\mathbf{p}}^{\prime(j)} \in \mathbb{R}^S$ be the true, sampled, MAP, and maximum likelihood dynamics parameters, respectively (hiding the dependency on the DPS episode $i1$ or $i2$ for the latter three quantities); thus, $[\bar{\mathbf{p}}^{(j)}]_k$ denotes the true probability of transitioning from state-action pair \tilde{s}_j to the k^{th} state, and analogously for the k^{th} elements of $\tilde{\mathbf{p}}^{(j)}, \hat{\mathbf{p}}^{(j)}$, and $\hat{\mathbf{p}}^{\prime(j)}$. Then, from the Dirichlet model,

$$[\hat{\mathbf{p}}^{(j)}]_k = \frac{n_{jk} + \alpha_{jk,0}}{n_j + \sum_{m=1}^S \alpha_{jm,0}},$$

where the prior for $\bar{\mathbf{p}}^{(j)}$ is $\frac{1}{\sum_{m=1}^S \alpha_{jm,0}} [\alpha_{j1,0}, \dots, \alpha_{jS,0}]^T$ for user-defined hyperparameters $\alpha_{jk,0} > 0$. Meanwhile, the maximum likelihood is given by $[\hat{\mathbf{p}}^{\prime(j)}]_k = \frac{n_{jk}}{\max(n_j, 1)}$ (this is equivalent to $[\hat{\mathbf{p}}^{(j)}]_k$, except with the prior parameters set to zero). Consider the sampled dynamics at state-action pair \tilde{s}_j . For any $\varepsilon > 0$,

$$\begin{aligned} P(\|\tilde{\mathbf{p}}^{(j)} - \bar{\mathbf{p}}^{(j)}\|_1 \geq \varepsilon) &= P(\|\tilde{\mathbf{p}}^{(j)} - \hat{\mathbf{p}}^{(j)} + \hat{\mathbf{p}}^{(j)} - \hat{\mathbf{p}}^{\prime(j)} + \hat{\mathbf{p}}^{\prime(j)} - \bar{\mathbf{p}}^{(j)}\|_1 \geq \varepsilon) \\ &\stackrel{(a)}{\leq} P(\|\tilde{\mathbf{p}}^{(j)} - \hat{\mathbf{p}}^{(j)}\|_1 + \|\hat{\mathbf{p}}^{(j)} - \hat{\mathbf{p}}^{\prime(j)}\|_1 + \|\hat{\mathbf{p}}^{\prime(j)} - \bar{\mathbf{p}}^{(j)}\|_1 \geq \varepsilon) \\ &\leq P\left(\|\tilde{\mathbf{p}}^{(j)} - \hat{\mathbf{p}}^{(j)}\|_1 \geq \frac{\varepsilon}{3} \cup \|\hat{\mathbf{p}}^{(j)} - \hat{\mathbf{p}}^{\prime(j)}\|_1 \geq \frac{\varepsilon}{3} \cup \|\hat{\mathbf{p}}^{\prime(j)} - \bar{\mathbf{p}}^{(j)}\|_1 \geq \frac{\varepsilon}{3}\right) \\ &\stackrel{(b)}{\leq} P\left(\|\tilde{\mathbf{p}}^{(j)} - \hat{\mathbf{p}}^{(j)}\|_1 \geq \frac{\varepsilon}{3}\right) + P\left(\|\hat{\mathbf{p}}^{(j)} - \hat{\mathbf{p}}^{\prime(j)}\|_1 \geq \frac{\varepsilon}{3}\right) + P\left(\|\hat{\mathbf{p}}^{\prime(j)} - \bar{\mathbf{p}}^{(j)}\|_1 \geq \frac{\varepsilon}{3}\right), \quad (6) \end{aligned}$$

where (a) holds due to the triangle inequality and (b) follows from the union bound. For each term in (6), we will upper-bound the quantity in terms of n_j and show that it decays as $n_j \rightarrow \infty$, that is, as \tilde{s}_j is visited infinitely-often. For the first term, we will achieve this bound via Chebyshev's inequality:

$$\begin{aligned} P\left(\|\tilde{\mathbf{p}}^{(j)} - \hat{\mathbf{p}}^{(j)}\|_1 \geq \frac{\varepsilon}{3}\right) &\leq P\left(\bigcup_{k=1}^S \left\{ \left| [\tilde{\mathbf{p}}^{(j)}]_k - [\hat{\mathbf{p}}^{(j)}]_k \right| \geq \frac{\varepsilon}{3S} \right\}\right) \stackrel{(a)}{\leq} \sum_{k=1}^S P\left(\left| [\tilde{\mathbf{p}}^{(j)}]_k - [\hat{\mathbf{p}}^{(j)}]_k \right| \geq \frac{\varepsilon}{3S}\right) \\ &\stackrel{(b)}{\leq} \sum_{k=1}^S \frac{9S^2}{\varepsilon^2} \operatorname{Var}\left[[\tilde{\mathbf{p}}^{(j)}]_k\right], \end{aligned}$$

where (a) follows from the union bound and (b) is an application of Chebyshev's inequality. For a Dirichlet variable X with parameters $(\alpha_1, \dots, \alpha_S)$, $\alpha_k > 0$ for each k , the variance of the k^{th} component X_k is given by:

$$\operatorname{Var}[X_k] = \frac{\tilde{\alpha}_k(1 - \tilde{\alpha}_k)}{1 + \sum_{m=1}^S \alpha_m} \leq \frac{1}{2} * \frac{1}{1 + \sum_{m=1}^S \alpha_m},$$

where $\tilde{\alpha}_k := \frac{\alpha_k}{\sum_{m=1}^S \alpha_m}$. In the DPS algorithm, $\tilde{\mathbf{p}}^{(j)}$ is drawn from a Dirichlet distribution with parameters $(\alpha_{j1}, \dots, \alpha_{jS}) = (\alpha_{j1,0} + n_{j1}, \dots, \alpha_{jS,0} + n_{jS})$, and so,

$$\text{Var} \left[[\tilde{\mathbf{p}}^{(j)}]_k \right] \leq \frac{1}{2} * \frac{1}{1 + \sum_{m=1}^S \alpha_{jm}} = \frac{1}{2} * \frac{1}{1 + \sum_{m=1}^S (\alpha_{jm,0} + n_{jm})} \leq \frac{1}{2} * \frac{1}{1 + \sum_{m=1}^S n_{jm}} = \frac{1}{2(1 + n_j)},$$

and so,

$$P \left(\|\tilde{\mathbf{p}}^{(j)} - \hat{\mathbf{p}}^{(j)}\|_1 \geq \frac{\varepsilon}{3} \right) \leq \sum_{k=1}^S \frac{9S^2}{\varepsilon^2} \frac{1}{2(1 + n_j)} = \frac{9S^3}{2\varepsilon^2(1 + n_j)}.$$

Considering the second term in (6),

$$\begin{aligned} P \left(\|\hat{\mathbf{p}}^{(j)} - \hat{\mathbf{p}}'^{(j)}\|_1 \geq \frac{\varepsilon}{3} \right) &\leq P \left(\bigcup_{k=1}^S \left\{ \left| [\hat{\mathbf{p}}^{(j)}]_k - [\hat{\mathbf{p}}'^{(j)}]_k \right| \geq \frac{\varepsilon}{3S} \right\} \right) \stackrel{(a)}{\leq} \sum_{k=1}^S P \left(\left| [\hat{\mathbf{p}}^{(j)}]_k - [\hat{\mathbf{p}}'^{(j)}]_k \right| \geq \frac{\varepsilon}{3S} \right) \\ &\stackrel{(b)}{\leq} \sum_{k=1}^S P \left(\frac{\alpha_{jk,0} + \sum_{m=1}^S \alpha_{jm,0}}{n_j + \sum_{m=1}^S \alpha_{jm,0}} \geq \frac{\varepsilon}{3S} \right), \end{aligned}$$

where (a) holds via the union bound and (b) follows for $n_j \geq 1$ because when $n_j \geq 1$:

$$\begin{aligned} \left| [\hat{\mathbf{p}}^{(j)}]_k - [\hat{\mathbf{p}}'^{(j)}]_k \right| &= \left| \frac{n_{jk} + \alpha_{jk,0}}{n_j + \sum_{m=1}^S \alpha_{jm,0}} - \frac{n_{jk}}{n_j} \right| = \left| \frac{\alpha_{jk,0}}{n_j + \sum_{m=1}^S \alpha_{jm,0}} - \frac{n_{jk} \sum_{m=1}^S \alpha_{jm,0}}{n_j (n_j + \sum_{m=1}^S \alpha_{jm,0})} \right| \\ &\leq \frac{\alpha_{jk,0}}{n_j + \sum_{m=1}^S \alpha_{jm,0}} + \frac{n_{jk}}{n_j} \frac{\sum_{m=1}^S \alpha_{jm,0}}{n_j + \sum_{m=1}^S \alpha_{jm,0}} \leq \frac{\alpha_{jk,0} + \sum_{m=1}^S \alpha_{jm,0}}{n_j + \sum_{m=1}^S \alpha_{jm,0}}. \end{aligned}$$

For the third term in (6), we apply the following concentration inequality for Dirichlet variables (see Appendix C.1 in Jaksch et al. (2010)):

$$P(\|\hat{\mathbf{p}}'^{(j)} - \bar{\mathbf{p}}^{(j)}\|_1 \geq \varepsilon) \leq (2^S - 2) \exp \left(\frac{-n_j \varepsilon^2}{2} \right).$$

Therefore:

$$P \left(\|\hat{\mathbf{p}}^{(j)} - \bar{\mathbf{p}}^{(j)}\|_1 \geq \frac{\varepsilon}{3} \right) \leq (2^S - 2) \exp \left(\frac{-n_j \varepsilon^2}{18} \right).$$

Thus, to upper-bound (6), for any $\varepsilon > 0$:

$$P(\|\tilde{\mathbf{p}}^{(j)} - \bar{\mathbf{p}}^{(j)}\|_1 \geq \varepsilon) \leq \frac{9S^3}{2\varepsilon^2(n_j + 1)} + \sum_{k=1}^S P \left(\frac{\alpha_{jk,0} + \sum_{m=1}^S \alpha_{jm,0}}{n_j + \sum_{m=1}^S \alpha_{jm,0}} \geq \frac{\varepsilon}{3S} \right) + (2^S - 2) \exp \left(\frac{-n_j \varepsilon^2}{18} \right).$$

On the right hand side, the first and third terms clearly decay as $n_j \rightarrow \infty$. The middle term is identically zero for n_j large enough, since the $\alpha_{jk,0}$ values are user-defined constants. Given this inequality, it is clear that for any $\varepsilon > 0$, as $n_j \rightarrow \infty$, $P(\|\tilde{\mathbf{p}}^{(j)} - \bar{\mathbf{p}}^{(j)}\|_1 \geq \varepsilon) \rightarrow 0$. If every state-action pair is visited infinitely-often, then $n_j \rightarrow \infty$ for each j , and so $\tilde{\mathbf{p}}^{(j)}$ converges in probability to $\bar{\mathbf{p}}^{(j)}$: $\tilde{\mathbf{p}}^{(j)} \xrightarrow{P} \bar{\mathbf{p}}^{(j)}$. Convergence in probability implies convergence in distribution, the desired result. \square

To continue proving that DPS's model of the transition dynamics converges, we next prove the intermediate result that the magnitude of the reward MAP estimate, $\|\hat{\mathbf{r}}_n\|_2$, is uniformly upper-bounded:

Lemma 2. *Across all $n \geq 1$, there exists some $b < \infty$ such that estimated reward at DPS trial n is bounded by b : $\|\hat{\mathbf{r}}_n\|_2 \leq b$.*

Proof. Recall that the MAP reward estimate $\hat{\mathbf{r}}_n$ is the solution to a ridge regression problem:

$$\hat{\mathbf{r}}_n = \arg \inf_{\mathbf{r}} \left\{ \sum_{i=1}^{n-1} (\mathbf{x}_i^T \mathbf{r} - y_i)^2 + \lambda \|\mathbf{r}\|_2^2 \right\} = \arg \inf_{\mathbf{r}} \left\{ \sum_{i=1}^{n-1} \left[(\mathbf{x}_i^T \mathbf{r} - y_i)^2 + \frac{1}{n-1} \lambda \|\mathbf{r}\|_2^2 \right] \right\}. \quad (7)$$

We will prove the desired result by contradiction. Assuming that there exists no upper bound b , we will identify a subsequence $(\hat{\mathbf{r}}_{n_i})$ of MAP estimates whose lengths increase unboundedly, but whose directions converge. Then, we will show that such vectors fail to minimize the objective in (7), achieving a contradiction.

Because there exist finitely-many state-action pairs, there are finitely-many possible length- h trajectories. Hence, the vector \mathbf{x}_i in Equation (7) can take finitely-many possible values. The binary labels y_i take values in $\{-\frac{1}{2}, \frac{1}{2}\}$. Note that for $\mathbf{r} = 0$, $(\mathbf{x}_i^T \mathbf{r} - y_i)^2 + \frac{1}{n-1} \lambda \|\mathbf{r}\|_2^2 = \frac{1}{4}$. We prove the desired statement by contradiction: assume that there is no $b < \infty$ such that $\|\hat{\mathbf{r}}_n\|_2 \leq b$ for all n . Then, the sequence $\hat{\mathbf{r}}_1, \hat{\mathbf{r}}_2, \dots$ must have a subsequence indexed by (n_i) such that $\lim_{i \rightarrow \infty} \|\hat{\mathbf{r}}_{n_i}\|_2 = \infty$. Consider the sequence of unit vectors $\frac{\hat{\mathbf{r}}_{n_i}}{\|\hat{\mathbf{r}}_{n_i}\|_2}$. This sequence lies within the compact set of unit vectors in \mathbb{R}^d , and so it must have a convergent subsequence; we index this subsequence of the sequence (n_i) by (n_{i_j}) . Then, the sequence $(\hat{\mathbf{r}}_{n_{i_j}})$ is such that $\lim_{j \rightarrow \infty} \|\hat{\mathbf{r}}_{n_{i_j}}\|_2 = \infty$ and $\lim_{j \rightarrow \infty} \frac{\hat{\mathbf{r}}_{n_{i_j}}}{\|\hat{\mathbf{r}}_{n_{i_j}}\|_2} = \hat{\mathbf{r}}_{unit}$, where $\hat{\mathbf{r}}_{unit} \in \mathbb{R}^d$ is a fixed unit vector.

For any \mathbf{x}_i such that $|\mathbf{x}_i^T \hat{\mathbf{r}}_{unit}| \neq 0$, $\lim_{n_{i_j} \rightarrow \infty} (\mathbf{x}_i^T \hat{\mathbf{r}}_{n_{i_j}} - y_i)^2 = \infty$, and so the corresponding terms in (7) approach infinity. However, a lower value of the optimization objective in (7) can be realized by replacing $\hat{\mathbf{r}}_{n_{i_j}}$ with the assignment $\mathbf{r} = 0$. Meanwhile, for any \mathbf{x}_i such that $|\mathbf{x}_i^T \hat{\mathbf{r}}| = 0$, replacing $\hat{\mathbf{r}}_{n_{i_j}}$ with $\mathbf{r} = 0$ would also decrease the value of the optimization objective in (7). Therefore, for large j , $\mathbf{r} = 0$ results in a smaller objective function value than $\hat{\mathbf{r}}_{n_{i_j}}$. This is a contradiction, and so the elements of the sequence $\hat{\mathbf{r}}_{n_{i_j}}$ cannot have arbitrarily-large magnitudes. Thus, the elements of the original sequence $\hat{\mathbf{r}}_i$ also cannot become arbitrarily large, and $\|\hat{\mathbf{r}}_i\| \leq b$ for some $b < \infty$. \square

To finish proving convergence of the transition dynamics Bayesian model, we show that every state-action pair is visited infinitely-often.

Lemma 3. *Under DPS, every state-action pair is visited infinitely-often.*

Proof. The proof proceeds by assuming that there exists a state-action pair that is visited only finitely-many times. This assumption will lead to a contradiction¹: once this state-action pair is no longer visited, the reward model posterior is no longer updated with respect to it. Then, DPS is guaranteed to eventually sample a high enough reward for this state-action that the resultant policy will prioritize visiting it.

First we note that DPS is guaranteed to reach at least one state-action pair infinitely often: given our problem's finite state and action spaces, at least one state-action pair must be visited infinitely-often during DPS execution. If all state-actions are *not* visited infinitely-often, there must exist a state-action pair (s, a) such that s is visited infinitely-often, while (s, a) is not. Otherwise, if all actions are selected infinitely-often in all infinitely-visited states, the finitely-visited states are unreachable (in which case these states are irrelevant to the learning process and regret minimization, and can be ignored). Without loss of generality, we label this state-action pair (s, a) as \tilde{s}_1 . To reach a contradiction, it suffices to show that \tilde{s}_1 is visited infinitely-often.

Let \mathbf{r}_1 be the reward vector with a reward of 1 in state-action pair \tilde{s}_1 and rewards of zero elsewhere. From Definition 2, $\pi_{vi}(\tilde{\mathbf{p}}, \mathbf{r}_1)$ is the policy that maximizes the expected number of visits to \tilde{s}_1 under dynamics $\tilde{\mathbf{p}}$ and reward vector \mathbf{r}_1 :

$$\pi_{vi}(\tilde{\mathbf{p}}, \mathbf{r}_1) = \operatorname{argmax}_{\pi} V(\tilde{\mathbf{p}}, \mathbf{r}_1, \pi),$$

where $V(\tilde{\mathbf{p}}, \mathbf{r}_1, \pi)$ is the expected total reward of a length- h trajectory under $\tilde{\mathbf{p}}, \mathbf{r}_1$, and π , or equivalently (by definition of \mathbf{r}_1), the expected number of visits to state-action \tilde{s}_1 .

¹Note that in finite-horizon MDPs, the concept of visiting a state finitely-many times is not the same as that of a transient state in an infinite Markov chain, because: 1) due to a finite horizon, the state is resampled from the initial state distribution $p_0(s)$ every h timesteps, and 2) the policy—which determines which state-action pairs can be reached in an episode—is also resampled every h timesteps.

We next show that there exists a $\rho > 0$ such that $P(\pi = \pi_{vi}(\tilde{\boldsymbol{p}}, \boldsymbol{r}_1)) > \rho$ for all possible values of $\tilde{\boldsymbol{p}}$. That is, for any sampled parameters $\tilde{\boldsymbol{p}}$, the probability of selecting policy $\pi_{vi}(\tilde{\boldsymbol{p}}, \boldsymbol{r}_1)$ is uniformly lower-bounded, implying that DPS must eventually select $\pi_{vi}(\tilde{\boldsymbol{p}}, \boldsymbol{r}_1)$.

Let \tilde{r}_j be the sampled reward associated with state-action pair \tilde{s}_j in a particular DPS episode, for each state-action $j \in \{1, \dots, d\}$, with $d = SA$. We show that conditioned on $\tilde{\boldsymbol{p}}$, there exists $v > 0$ such that if \tilde{r}_1 exceeds $\max\{v\tilde{r}_2, v\tilde{r}_3, \dots, v\tilde{r}_d\}$, then value iteration returns the policy $\pi_{vi}(\tilde{\boldsymbol{p}}, \boldsymbol{r}_1)$, which is the policy maximizing the expected amount of time spent in \tilde{s}_1 . This can be seen by setting $v := \frac{h}{\rho_1}$, where h is the time horizon and ρ_1 is the expected number of visits to \tilde{s}_1 under $\pi_{vi}(\tilde{\boldsymbol{p}}, \boldsymbol{r}_1)$. Under this definition of v , the event $\{\tilde{r}_1 \geq \max\{v\tilde{r}_2, v\tilde{r}_3, \dots, v\tilde{r}_d\}\}$ is equivalent to $\{\tilde{r}_1 \rho_1 \geq h \max\{\tilde{r}_2, \tilde{r}_3, \dots, \tilde{r}_d\}\}$; the latter inequality implies that given $\tilde{\boldsymbol{p}}$ and $\tilde{\boldsymbol{r}}$, the expected reward accumulated solely in state-action \tilde{s}_1 exceeds the reward gained by repeatedly (during all h time-steps) visiting the state-action pair in the set $\{\tilde{s}_2, \dots, \tilde{s}_d\}$ having the highest sampled reward. Clearly, in this situation, value iteration results in the policy $\pi_{vi}(\tilde{\boldsymbol{p}}, \boldsymbol{r}_1)$.

Next we show that $v = \frac{h}{\rho_1}$ is continuous in the sampled dynamics $\tilde{\boldsymbol{p}}$ by showing that ρ_1 is continuous in $\tilde{\boldsymbol{p}}$. Recall that ρ_1 is defined as expected number of visits to \tilde{s}_1 under $\pi_{vi}(\tilde{\boldsymbol{p}}, \boldsymbol{r}_1)$. This is equivalent to the expected reward for following $\pi_{vi}(\tilde{\boldsymbol{p}}, \boldsymbol{r}_1)$ under dynamics $\tilde{\boldsymbol{p}}$ and rewards \boldsymbol{r}_1 :

$$\rho_1 = V(\tilde{\boldsymbol{p}}, \boldsymbol{r}_1, \pi_{vi}(\tilde{\boldsymbol{p}}, \boldsymbol{r}_1)) = \max_{\pi} V(\tilde{\boldsymbol{p}}, \boldsymbol{r}_1, \pi). \quad (8)$$

The value of any policy π is continuous in the transition dynamics parameters, and so $V(\tilde{\boldsymbol{p}}, \boldsymbol{r}_1, \pi)$ is continuous in $\tilde{\boldsymbol{p}}$. The maximum in (8) is taken over the finite set of deterministic policies; because a maximum over a finite number of continuous functions is also continuous, ρ_1 is continuous in $\tilde{\boldsymbol{p}}$.

Next, recall that a continuous function on a compact set achieves its maximum and minimum values on that set. The set of all possible dynamics parameters $\tilde{\boldsymbol{p}}$ is such that for each state-action pair j , $\sum_{k=1}^S p_{jk} = 1$ and $p_{jk} \geq 0 \forall k$; the set of all possible vectors $\tilde{\boldsymbol{p}}$ is clearly closed and bounded, and hence compact. Therefore, v achieves its maximum and minimum values on this set, and so for any $\tilde{\boldsymbol{p}}$, $v \in [v_{\min}, v_{\max}]$, where $v_{\min} > 0$ (v is nonnegative by definition, and $v = 0$ is impossible, as it would imply that \tilde{s}_1 is unreachable).

Then, $P(\pi = \pi_{vi}(\tilde{\boldsymbol{p}}, \boldsymbol{r}_1))$ can then be expressed in terms of v and the parameters of the reward posterior. Firstly,

$$P(\pi = \pi_{vi}(\tilde{\boldsymbol{p}}, \boldsymbol{r}_1)) \geq P(\tilde{r}_1 > \max\{v\tilde{r}_2, v\tilde{r}_3, \dots, v\tilde{r}_d\}) \geq \prod_{j=2}^d P(\tilde{r}_1 > v\tilde{r}_j) = \prod_{j=2}^d [1 - P(\tilde{r}_1 - v\tilde{r}_j \leq 0)].$$

In the n^{th} DPS iteration, the sampled rewards are drawn from a jointly Gaussian posterior: $\tilde{\boldsymbol{r}} \sim \mathcal{N}(\boldsymbol{\mu}^{(n)}, \Sigma^{(n)})$ for some $\boldsymbol{\mu}^{(n)}$ and $\Sigma^{(n)}$, where $[\boldsymbol{\mu}^{(n)}]_j = \mu_j^{(n)}$ and $[\Sigma^{(n)}]_{jk} = \Sigma_{jk}^{(n)}$. Then, $(\tilde{r}_1 - v\tilde{r}_j) \sim \mathcal{N}(\mu_1^{(n)} - v\mu_j^{(n)}, \Sigma_{11}^{(n)} + v^2\Sigma_{jj}^{(n)} - 2v\Sigma_{1j}^{(n)})$, so that:

$$P(\pi_{n1} = \pi_{vi}(\tilde{\boldsymbol{p}}, \boldsymbol{r}_1)) \geq \prod_{j=2}^d \left[1 - \Phi \left(\frac{-\mu_1^{(n)} + v\mu_j^{(n)}}{\sqrt{\Sigma_{11}^{(n)} + v^2\Sigma_{jj}^{(n)} - 2v\Sigma_{1j}^{(n)}}} \right) \right] = \prod_{j=2}^d \Phi \left(\frac{\mu_1^{(n)} - v\mu_j^{(n)}}{\sqrt{\Sigma_{11}^{(n)} + v^2\Sigma_{jj}^{(n)} - 2v\Sigma_{1j}^{(n)}}} \right), \quad (9)$$

where Φ is the standard Gaussian cumulative distribution function. For the right-hand expression in (9) to have a lower bound greater than zero, the argument of $\Phi(\cdot)$ must be lower-bounded. It suffices to upper-bound the numerator's magnitude and to lower-bound the denominator above zero for each product factor j and over all iterations n .

The numerator can be upper-bounded using Lemma 2. Since $\boldsymbol{\mu}^{(n)}$ is equal to the MAP reward estimate at iteration n , $\|\boldsymbol{\mu}^{(n)}\|_2 \leq b$, and so $|\mu_1^{(n)}|, |\mu_j^{(n)}| \leq b$. Because $0 < v \leq v_{\max}$, $|\mu_1 - v\mu_j| \leq |\mu_1^{(n)}| + v|\mu_j^{(n)}| \leq (1 + v_{\max})b$.

To lower-bound the denominator, first note that the reward model's posterior covariance is equal to $\beta_n(\delta)^2 M_n^{-1}$, with M_n and $\beta_n(\delta)$ as defined in Equations (4) and (5), respectively; however, because $\beta_n(\delta)^2$ is non-decreasing in n , it suffices to prove the statement while ignoring the $\beta_n(\delta)^2$ factor. Thus, to prove this lemma, we can set $\Sigma^{(n)} := M_n^{-1}$.

Let $\boldsymbol{w}_j \in \mathbb{R}^d$ be a vector with 1 in the first position, $-v$ in the j^{th} position for some $j \in \{2, \dots, d\}$, and zero elsewhere:

$$\boldsymbol{w}_j = [1, 0, \dots, 0, -v, 0, \dots, 0]^T. \quad (10)$$

The denominator in (9) can be expressed in terms of \mathbf{w}_j : $\Sigma_{11}^{(n)} + v^2 \Sigma_{jj}^{(n)} - 2v \Sigma_{1j}^{(n)} = \mathbf{w}_j^T \Sigma^{(n)} \mathbf{w}_j$. Recall from Definition 3 that the eigenvectors of $\Sigma^{(n)}$ are $\mathbf{v}_1^{(n)}, \dots, \mathbf{v}_d^{(n)}$, with corresponding eigenvalues $(\lambda_1^{(n)})^{-1}, \dots, (\lambda_d^{(n)})^{-1}$. We can write \mathbf{w}_j in terms of the orthonormal basis formed by the eigenvectors $\{\mathbf{v}_k^{(n)}\}$:

$$\mathbf{w}_j = \sum_{k=1}^d \alpha_k^{(n)} \mathbf{v}_k^{(n)}, \quad (11)$$

for some coefficients $\alpha_k^{(n)} \in \mathbb{R}$. Using (11), the square of the denominator in (9) can now be written as:

$$\begin{aligned} \Sigma_{11}^{(n)} + v^2 \Sigma_{jj}^{(n)} - 2v \Sigma_{1j}^{(n)} &= \mathbf{w}_j^T \Sigma^{(n)} \mathbf{w}_j = \left(\sum_{k=1}^d \alpha_k^{(n)} \mathbf{v}_k^{(n)T} \right) \left(\sum_{l=1}^d \frac{1}{\lambda_l^{(n)}} \mathbf{v}_l^{(n)} \mathbf{v}_l^{(n)T} \right) \left(\sum_{m=1}^d \alpha_m^{(n)} \mathbf{v}_m^{(n)} \right) \\ &\stackrel{(a)}{=} \sum_{k=1}^d \left(\alpha_k^{(n)} \right)^2 \frac{1}{\lambda_k^{(n)}} \stackrel{(b)}{\geq} \left(\alpha_{k_0}^{(n)} \right)^2 \frac{1}{\lambda_{k_0}^{(n)}}, \end{aligned} \quad (12)$$

where equality (a) follows by orthonormality of the eigenvector basis, and (b) holds for any $k_0 \in \{1, \dots, d\}$ due to positivity of the eigenvalues $(\lambda_k)^{-1}$. Therefore, to show that the denominator is bounded away from zero, it suffices to show that for every n , there exists some k_0 such that $(\alpha_{k_0}^{(n)})^2 (\lambda_{k_0}^{(n)})^{-1}$ is bounded away from zero.

To prove the previous statement, note that by definition of M_n , the eigenvalues $(\lambda_k^{(n)})^{-1}$ are non-increasing in n . Below, we will show that for any eigenvalue $(\lambda_k^{(n)})^{-1}$ such that $\lim_{n \rightarrow \infty} (\lambda_k^{(n)})^{-1} = 0$, the first element of its corresponding eigenvector, $[\mathbf{v}_k^{(n)}]_1$, also converges to zero. Since the first element of \mathbf{w}_j equals 1, (10) implies that there must exist some k_0 such that $[\mathbf{v}_{k_0}^{(n)}]_1 \not\rightarrow 0$ and $\alpha_{k_0}^{(n)}$ is bounded away from 0. If these implications did not hold, then \mathbf{w}_j would not have a value of 1 in its first element, contradicting its definition. These observations imply that for every n , there must be some k_0 such that as $n \rightarrow \infty$, $(\lambda_{k_0}^{(n)})^{-1} \not\rightarrow 0$ and $\alpha_{k_0}^{(n)}$ is bounded away from zero.

Let X_n denote the observation matrix after $n - 1$ observations: $X_n := [\mathbf{x}_1 \ \dots \ \mathbf{x}_{n-1}]^T$. Then, $\Sigma^{(n)} = M_n^{-1} = (X_n^T X_n + \lambda I)^{-1}$. The matrices M_n^{-1} and $X_n^T X_n$ have the same eigenvectors. Meanwhile, for each eigenvalue $(\lambda_i^{(n)})^{-1}$ of M_n^{-1} , $X_n^T X_n$ has an eigenvalue $\nu_i^{(n)} := \lambda_i^{(n)} - \lambda \geq 0$ corresponding to the same eigenvector. We aim to characterize the eigenvectors of M_n^{-1} whose eigenvalues approach zero. Since these eigenvectors are identical to those of $X_n^T X_n$ whose eigenvalues approach infinity, we consider the latter instead.

We assume that all finitely-visited state-action pairs (including \tilde{s}_1) occur in the first $m < n - 1$ iterations. Without loss of generality, we index these finitely-visited state-action pairs from 1 to $r \geq 1$, so that the finitely-visited state-actions are: $\{\tilde{s}_1, \tilde{s}_2, \dots, \tilde{s}_r\}$. Let $X_{1:m} \in \mathbb{R}^{m \times d}$ denote the matrix containing the first m rows of X_n , while $X_{m+1:n} \in \mathbb{R}^{n-m \times d}$ denotes the remaining rows of X_n . With this notation,

$$X_n^T X_n = \sum_{i=1}^{n-1} \mathbf{x}_i \mathbf{x}_i^T = X_{1:m}^T X_{1:m} + X_{m+1:n}^T X_{m+1:n}.$$

Because the first r state-action pairs, $\{\tilde{s}_1, \tilde{s}_2, \dots, \tilde{s}_r\}$, are unvisited after iteration m , the first r elements of \mathbf{x}_i are zero for all $i > m$. Therefore, $X_{m+1:n}^T X_{m+1:n}$ can be written in the following block matrix form:

$$X_{m+1:n}^T X_{m+1:n} = \begin{bmatrix} O_{r \times r} & O_{r \times (d-r)} \\ O_{(d-r) \times r} & A_n \end{bmatrix},$$

where $O_{a \times b}$ denotes the all-zero matrix with dimensions $a \times b$. The matrix A_n includes elements that are unbounded as $n \rightarrow \infty$. In particular, the diagonal elements of A_n approach infinity as $n \rightarrow \infty$. We can write $X_n^T X_n$ in the following block matrix form:

$$X_n^T X_n = X_{1:m}^T X_{1:m} + X_{m+1:n}^T X_{m+1:n} = \begin{bmatrix} [X_{1:m}^T X_{1:m}]_{(1:r,1:r)} & [X_{1:m}^T X_{1:m}]_{(1:r,r+1:d)} \\ [X_{1:m}^T X_{1:m}]_{(r+1:d,1:r)} & [X_{1:m}^T X_{1:m}]_{(r+1:d,r+1:d)} + A_n \end{bmatrix} := \begin{bmatrix} B & C \\ C^T & D_n \end{bmatrix},$$

where $M_{(a,b,c:d)}$ denotes the submatrix of M obtained by extracting rows a through b and columns c through d . Matrices B and C only depend upon $X_{1:m}$, and so are fixed as n increases, while matrix D_n contains values that grow towards infinity with increasing n . In particular, all elements along D_n 's diagonal are unbounded. Intuitively, in the limit, B and C are close to zero compared to D_n , and $X_n^T X_n$ (when normalized) increasingly resembles a matrix in which only the bottom-right block is nonzero. This intuitive notion is formalized next.

Consider an eigenpair $(\mathbf{v}_i^{(n)}, \nu_i^{(n)})$ of $X_n^T X_n$ such that $\lim_{n \rightarrow \infty} \nu_i^{(n)} = \infty$. We show that the first element of $\mathbf{v}_i^{(n)}$ must approach 0. Let $\mathbf{v}_i^{(n)} = \begin{bmatrix} \mathbf{z}_i^{(n)T} & \mathbf{q}_i^{(n)T} \end{bmatrix}^T$, where $\mathbf{z}_i^{(n)} \in \mathbb{R}^m$ and $\mathbf{q}_i^{(n)} \in \mathbb{R}^{n-1-m}$. We see that:

$$(X_n^T X_n) \mathbf{v}_i^{(n)} = X_n^T X_n \begin{bmatrix} \mathbf{z}_i^{(n)} \\ \mathbf{q}_i^{(n)} \end{bmatrix} = \begin{bmatrix} B & C \\ C^T & D_n \end{bmatrix} \begin{bmatrix} \mathbf{z}_i^{(n)} \\ \mathbf{q}_i^{(n)} \end{bmatrix} = \begin{bmatrix} B\mathbf{z}_i^{(n)} + C\mathbf{q}_i^{(n)} \\ C^T\mathbf{z}_i^{(n)} + D_n\mathbf{q}_i^{(n)} \end{bmatrix} = \lambda_i^{(n)} \begin{bmatrix} \mathbf{z}_i^{(n)} \\ \mathbf{q}_i^{(n)} \end{bmatrix}.$$

Dividing both sides by $\nu_i^{(n)}$,

$$\frac{1}{\nu_i^{(n)}} X_n^T X_n \begin{bmatrix} \mathbf{z}_i^{(n)} \\ \mathbf{q}_i^{(n)} \end{bmatrix} = \begin{bmatrix} \frac{1}{\nu_i^{(n)}} (B\mathbf{z}_i^{(n)} + C\mathbf{q}_i^{(n)}) \\ \frac{1}{\nu_i^{(n)}} (C^T\mathbf{z}_i^{(n)} + D_n\mathbf{q}_i^{(n)}) \end{bmatrix} = \begin{bmatrix} \mathbf{z}_i^{(n)} \\ \mathbf{q}_i^{(n)} \end{bmatrix}.$$

In the upper matrix block: $\lim_{n \rightarrow \infty} \nu_i^{(n)} = \infty$, B and C are fixed as n increases, and $\mathbf{z}_i^{(n)}$ and $\mathbf{q}_i^{(n)}$ have upper-bounded elements because $\mathbf{v}_i^{(n)}$ is a unit vector. Thus, $\lim_{n \rightarrow \infty} \frac{1}{\nu_i^{(n)}} (B\mathbf{z}_i^{(n)} + C\mathbf{q}_i^{(n)}) = \mathbf{0}$. In particular, the first element of $\mathbf{z}_i^{(n)}$ converges to zero, and so the same is true of $\mathbf{v}_i^{(n)}$.

As justified above, this result implies that for each iteration n , there exists an index $k_0 \in \{1, \dots, d\}$ such that the right-hand side of (12) has a lower bound above zero. This completes the proof that the denominator of the fraction in (9) does not decay to zero. As a result, there exists some $\rho > 0$ such that $P(\pi = \pi_{vi}(\tilde{\mathbf{p}}, \mathbf{r}_1)) \geq \rho > 0$.

In consequence, DPS is guaranteed to infinitely-often sample pairs $(\tilde{\mathbf{p}}, \pi)$ such that $\pi = \pi_{vi}(\tilde{\mathbf{p}}, \mathbf{r}_1)$. As a result, DPS infinitely-often samples policies that prioritize reaching \tilde{s}_1 as quickly as possible. Such a policy always takes action a in state s . Furthermore, because s is visited infinitely-often, either a) $p_0(s) > 0$ or b) the infinitely-visited state-action pairs include a path with a nonzero probability of reaching s . In case a), since the initial state distribution is fixed, the MDP will infinitely-often begin in state s under the policy $\pi = \pi_{vi}(\tilde{\mathbf{p}}, \mathbf{r}_1)$, and so \tilde{s}_1 will be visited infinitely-often. In case b), due to Lemma 1, the transition dynamics parameters for state-actions along the path to s converge to their true values (intuitively, the algorithm knows how to reach s). In episodes with the policy $\pi = \pi_{vi}(\tilde{\mathbf{p}}, \mathbf{r}_1)$, DPS is thus guaranteed to reach \tilde{s}_1 infinitely-often. Since DPS selects $\pi_{vi}(\tilde{\mathbf{p}}, \mathbf{r}_1)$ infinitely-often, it must reach \tilde{s}_1 infinitely-often. This presents a contradiction, and so every state-action pair must be visited infinitely-often. \square

Thus, by the direct combination of Lemmas 1 and 3, we arrive at the following result.

Proposition 1. *Under DPS, the sampled dynamics $\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{p}}_{i2}$ converge in distribution to the true dynamics: $\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{p}}_{i2} \xrightarrow{D} \bar{\mathbf{p}}$, where \xrightarrow{D} denotes convergence in distribution.*

Next, we show that the sampled rewards converge in distribution to their true values. Our analysis will use Theorem 2 from Abbasi-Yadkori et al. (2011), which is repeated below. Recall that Eq.s (3)-(4) define the MAP reward estimate.

Lemma 4 (Theorem 2 from Abbasi-Yadkori et al. (2011)). *Let $\{F_i\}_{i=0}^\infty$ be a filtration. Let $\{\eta_i\}_{i=1}^\infty$ be a real-valued stochastic process such that η_i is F_i -measurable and η_i is conditionally R -sub-Gaussian for some $R \geq 0$. Let $\{\mathbf{x}_i\}$ be an \mathbb{R}^d -valued stochastic process such that \mathbf{x}_i is F_{i-1} -measurable. Define $y_i := \mathbf{x}_i^T \bar{\mathbf{r}} + \eta_i$, and assume that $\|\bar{\mathbf{r}}\|_2 \leq S_r$ and $\|\mathbf{x}_i\|_2 \leq L$. Then, for any $\delta > 0$, with probability at least $1 - \delta$, for all $i > 0$, $\|\hat{\mathbf{r}}_i - \bar{\mathbf{r}}\|_{M_i} \leq \beta_i(\delta)$, where:*

$$\beta_i(\delta) = R \sqrt{2 \log \left(\frac{\det(M_i)^{1/2} \lambda^{-d/2}}{\delta} \right)} + \sqrt{\lambda} S_r \leq R \sqrt{d \log \left(\frac{1 + \frac{L^2 i}{d\lambda}}{\delta} \right)} + \sqrt{\lambda} S_r.$$

Note that in the present case, $L \leq 2h$, since:

$$\|\mathbf{x}_i\|_2 = \|\mathbf{x}_{i2} - \mathbf{x}_{i1}\|_2 \stackrel{(a)}{\leq} \|\mathbf{x}_{i2} - \mathbf{x}_{i1}\|_1 \leq \|\mathbf{x}_{i2}\|_1 + \|\mathbf{x}_{i1}\|_1 \stackrel{(b)}{=} 2h,$$

where (a) holds because $\|\mathbf{x}\|_2 \leq \|\mathbf{x}\|_1$ for any $\mathbf{x} \in \mathbb{R}^d$, and (b) holds because \mathbf{x}_{i1} and \mathbf{x}_{i2} each count an episode's visits to every state-action pair, and so their elements are non-negative integers summing to h .

Proposition 2. *With probability $1 - \delta$, where δ is a parameter of the Bayesian linear regression model, the sampled rewards $\tilde{\mathbf{r}}_{i1}, \tilde{\mathbf{r}}_{i2}$ converge in distribution to the true reward parameters, $\tilde{\mathbf{r}}_{i1}, \tilde{\mathbf{r}}_{i2} \xrightarrow{D} \bar{\mathbf{r}}$, as $i \rightarrow \infty$.*

Proof. This is a direct implication of Lemmas 5 and 11, both proven below. \square

Lemma 5. *If $\frac{\beta_i(\delta)^2}{\lambda_d^{(i)}} \xrightarrow{D} 0$ as $i \rightarrow \infty$, where $\lambda_d^{(i)}$ is the minimum eigenvalue of M_i and \xrightarrow{D} denotes convergence in distribution, then $\tilde{\mathbf{r}}_{i1}, \tilde{\mathbf{r}}_{i2} \xrightarrow{D} \bar{\mathbf{r}}$ with probability $1 - \delta$.*

Proof. From Lemma 4, with probability at least $1 - \delta$, $\hat{\mathbf{r}}_i$ belongs to a confidence ellipsoid centered at $\bar{\mathbf{r}}$: $\|\hat{\mathbf{r}}_i - \bar{\mathbf{r}}\|_{M_i} \leq \beta_i(\delta)$. We show that under this high-probability event, $\tilde{\mathbf{r}}_{i1}, \tilde{\mathbf{r}}_{i2} \xrightarrow{D} \bar{\mathbf{r}}$. Similarly to the high-probability confidence ellipsoid from Lemma 4, the Thompson sampling covariance matrix is also defined by $\beta_i(\delta)$ and M_i :

$$\tilde{\mathbf{r}}_{i1}, \tilde{\mathbf{r}}_{i2} \sim \mathcal{N}(\hat{\mathbf{r}}_i, \beta_i(\delta)^2 M_i^{-1}). \quad (13)$$

Letting $\mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, I)$ be independent for each i , we can equivalently express $\tilde{\mathbf{r}}_{i1}$ (and similarly, $\tilde{\mathbf{r}}_{i2}$) as:

$$\tilde{\mathbf{r}}_{i1} = \hat{\mathbf{r}}_i + \beta_i(\delta) M_i^{-\frac{1}{2}} \mathbf{z}_i, \quad (14)$$

since the random variable in (14) has the same distribution as (13). The quantity $\|\tilde{\mathbf{r}}_{i1} - \hat{\mathbf{r}}_i\|_{M_i}$ can be rewritten as:

$$\|\tilde{\mathbf{r}}_{i1} - \hat{\mathbf{r}}_i\|_{M_i} = \left\| \beta_i(\delta) M_i^{-\frac{1}{2}} \mathbf{z}_i \right\|_{M_i} = \beta_i(\delta) \sqrt{\mathbf{z}_i^T M_i^{-\frac{1}{2}} M_i M_i^{-\frac{1}{2}} \mathbf{z}_i} = \beta_i(\delta) \|\mathbf{z}_i\|_2.$$

Because the probability distribution of $\|\mathbf{z}_i\|_2$ is fixed, there exists some fixed $a > 0$ such that with probability at least $1 - \delta$, $\|\mathbf{z}_i\|_2 \leq a$. So, for each i , with probability at least $1 - \delta$,

$$\|\tilde{\mathbf{r}}_{i1} - \hat{\mathbf{r}}_i\|_{M_i} = \beta_i(\delta) \|\mathbf{z}_i\|_2 \leq \beta_i(\delta) a.$$

Assuming that the high-probability event in Lemma 4 occurs, we combine the previous inequality with $\|\hat{\mathbf{r}}_i - \bar{\mathbf{r}}\|_{M_i} \leq \beta_i(\delta)$ to obtain that for each i , with probability at least $1 - \delta$,

$$\|\tilde{\mathbf{r}}_{i1} - \bar{\mathbf{r}}\|_{M_i} \leq \|\tilde{\mathbf{r}}_{i1} - \hat{\mathbf{r}}_i\|_{M_i} + \|\hat{\mathbf{r}}_i - \bar{\mathbf{r}}\|_{M_i} \leq (a + 1)\beta_i(\delta).$$

Squaring both sides and dividing by $\beta_i(\delta)$ yields that for each i , with probability at least $1 - \delta$,

$$\frac{1}{\beta_i(\delta)^2} (\tilde{\mathbf{r}}_{i1} - \bar{\mathbf{r}})^T M_i (\tilde{\mathbf{r}}_{i1} - \bar{\mathbf{r}}) \leq (a + 1)^2.$$

By assumption, $\frac{\lambda_d^{(i)}}{\beta_i(\delta)^2} \xrightarrow{D} \infty$ as $i \rightarrow \infty$. Recall from Definition 3 that $\mathbf{v}_j^{(i)}, j \in \{1, \dots, d\}$, represent the eigenvectors of M_i corresponding to the eigenvalues $\lambda_j^{(i)}$. Then, with probability at least $1 - \delta$ for each i :

$$\begin{aligned} \frac{1}{\beta_i(\delta)^2} (\tilde{\mathbf{r}}_{i1} - \bar{\mathbf{r}})^T M_i (\tilde{\mathbf{r}}_{i1} - \bar{\mathbf{r}}) &= \frac{1}{\beta_i(\delta)^2} (\tilde{\mathbf{r}}_{i1} - \bar{\mathbf{r}})^T \left(\sum_{j=1}^d \lambda_j^{(i)} \mathbf{v}_j^{(i)} \mathbf{v}_j^{(i)T} \right) (\tilde{\mathbf{r}}_{i1} - \bar{\mathbf{r}}) \\ &= \frac{1}{\beta_i(\delta)^2} \sum_{j=1}^d \lambda_j^{(i)} \left((\tilde{\mathbf{r}}_{i1} - \bar{\mathbf{r}})^T \mathbf{v}_j^{(i)} \right)^2 \leq (a + 1)^2. \end{aligned} \quad (15)$$

Since for each j , $\frac{\lambda_j^{(i)}}{\beta_i(\delta)^2} \rightarrow \infty$ as $i \rightarrow \infty$, and $\mathbf{v}_j^{(i)}$ is an orthonormal basis, the constant bound of $(a + 1)^2$ in (15) is violated if we do not have $\tilde{\mathbf{r}}_{i1} - \bar{\mathbf{r}} \xrightarrow{D} \mathbf{0}$. Equation (15) must hold with probability at least $1 - \delta$ independently for

each iteration i , with the $(1 - \delta)$ -probability due entirely to randomness in the Thompson sampling distribution, (13). Therefore, it follows that $\tilde{\mathbf{r}}_{i1} \xrightarrow{D} \bar{\mathbf{r}}$. The proof that $\tilde{\mathbf{r}}_{i2} \xrightarrow{D} \bar{\mathbf{r}}$ with high probability is identical. \square

The next result enables us to leverage convergence in distribution of the dynamics samples, $\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{p}}_{i2} \xrightarrow{D} \bar{\mathbf{p}}$ (as guaranteed by Proposition 1), in characterizing the impact of sampled policies upon convergence of the reward model.

Lemma 6. *Let $f : \mathbb{R}^{S^2 A} \times \mathbb{R}^{SA} \rightarrow \mathbb{R}$ be a function of transition dynamics $\mathbf{p} \in \mathbb{R}^{S^2 A}$ and reward vector $\mathbf{r} \in \mathbb{R}^{SA}$, $f(\mathbf{p}, \mathbf{r})$, where f is continuous in \mathbf{p} and uniformly-continuous in \mathbf{r} . Assume that Proposition 1 holds: $\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{p}}_{i2} \xrightarrow{D} \bar{\mathbf{p}}$. Then, for any $\delta, \varepsilon > 0$, there exists i' such that for $i > i'$, $|f(\bar{\mathbf{p}}, \mathbf{r}) - f(\tilde{\mathbf{p}}_{ij}, \mathbf{r})| < \varepsilon$ for any unit vector \mathbf{r} and $j \in \{1, 2\}$ with probability at least $1 - \delta$.*

Proof. The proof is identical for $j = 1$ and $j = 2$, and so without loss of generality, we set $j = 1$. Applying Proposition 1, $\tilde{\mathbf{p}}_{i1} \xrightarrow{D} \bar{\mathbf{p}}$. By continuity of f , we can apply Fact 1 from Appendix A.5 to obtain that $f(\tilde{\mathbf{p}}_{i1}, \mathbf{r}) \xrightarrow{D} f(\bar{\mathbf{p}}, \mathbf{r})$ for any \mathbf{r} . Further applying Fact 2 from Appendix A.5, $f(\tilde{\mathbf{p}}_{i1}, \mathbf{r}) \xrightarrow{P} f(\bar{\mathbf{p}}, \mathbf{r})$ for any \mathbf{r} . By definition of convergence in probability, given δ , there exists i_r such that for $i \geq i_r$:

$$|f(\tilde{\mathbf{p}}_{i1}, \mathbf{r}) - f(\bar{\mathbf{p}}, \mathbf{r})| < \frac{1}{3}\varepsilon \text{ with probability at least } 1 - \delta. \quad (16)$$

To obtain a high-probability bound that applies over all unit vectors \mathbf{r} , we use compactness of the set of unit vectors. Any infinite cover of a compact set has a finite subcover; in particular, for any $\delta' > 0$, the set of unit vectors in \mathbb{R}^d has a finite cover of the form $\{\mathcal{B}(\mathbf{r}_1, \delta'), \dots, \mathcal{B}(\mathbf{r}_K, \delta')\}$, where $\{\mathbf{r}_1, \dots, \mathbf{r}_K\}$ are unit vectors, and $\mathcal{B}(\mathbf{r}, \delta') := \{\mathbf{r}' \in \mathbb{R}^d \mid \|\mathbf{r}' - \mathbf{r}\|_2 < \delta'\}$ is the d -dimensional sphere of radius δ' centered at \mathbf{r} . Thus, there exists a finite set of unit vectors $\mathcal{U} = \{\mathbf{r}_1, \dots, \mathbf{r}_K\}$ such that for any unit vector \mathbf{r}' , $\|\mathbf{r}_i - \mathbf{r}'\|_2 < \delta'$ for some $i \in \{1, \dots, K\}$. Because f is uniformly-continuous in \mathbf{r} , for any transition dynamics \mathbf{p} , there exists $\delta_p > 0$ such that for any two unit vectors \mathbf{r}, \mathbf{r}' such that $\|\mathbf{r} - \mathbf{r}'\|_2 < \delta_p$:

$$|f(\mathbf{p}, \mathbf{r}) - f(\mathbf{p}, \mathbf{r}')| < \frac{1}{3}\varepsilon. \quad (17)$$

Without loss of generality, for each \mathbf{p} , define $\delta_p := \sup x$ such that $\|\mathbf{r} - \mathbf{r}'\|_2 < x$ implies $|f(\mathbf{p}, \mathbf{r}) - f(\mathbf{p}, \mathbf{r}')| \leq \frac{1}{6}\varepsilon$. Then, because f is continuous in \mathbf{p} , δ_p is also continuous in \mathbf{p} . Because the set of all possible transition probability vectors \mathbf{p} is compact, and a continuous function over a compact set achieves its minimum value, there exists $\delta_{\min} > 0$ such that $\delta_p \geq \delta_{\min} > 0$ over all \mathbf{p} . We can define \mathcal{U} such that $\delta' \leq \delta_{\min}$; then, for any unit vector \mathbf{r}' , there exists $\mathbf{r} \in \mathcal{U}$ such that $\|\mathbf{r} - \mathbf{r}'\|_2 < \delta_{\min}$, and thus (17) holds for any \mathbf{p} .

By (16), for each $\mathbf{r}_j \in \mathcal{U}$, there exists there exists i_{r_j} such that for $i \geq i_{r_j}$: $\|f(\tilde{\mathbf{p}}_{i1}, \mathbf{r}) - f(\bar{\mathbf{p}}, \mathbf{r})\|_2 < \frac{1}{3}\varepsilon$ with probability at least $1 - \delta$. Because \mathcal{U} is a finite set, there exists $i' > \max\{i_{r_1}, \dots, i_{r_K}\}$ such that for $\mathbf{r} \in \mathcal{U}$ and $i > i'$:

$$|f(\tilde{\mathbf{p}}_{i1}, \mathbf{r}) - f(\bar{\mathbf{p}}, \mathbf{r})| < \frac{1}{3}\varepsilon \text{ for each } \mathbf{r} \in \mathcal{U} \text{ with probability at least } 1 - \delta. \quad (18)$$

Therefore, for any unit vector \mathbf{r}' , there exists $\mathbf{r} \in \mathcal{U}$ such that $\|\mathbf{r} - \mathbf{r}'\|_2 < \delta' \leq \delta_{\min}$, and with probability at least $1 - \delta$ for $i > i'$:

$$|f(\bar{\mathbf{p}}, \mathbf{r}') - f(\tilde{\mathbf{p}}_{i1}, \mathbf{r}')| = |f(\bar{\mathbf{p}}, \mathbf{r}') - f(\bar{\mathbf{p}}, \mathbf{r}) + f(\bar{\mathbf{p}}, \mathbf{r}) - f(\tilde{\mathbf{p}}_{i1}, \mathbf{r}) + f(\tilde{\mathbf{p}}_{i1}, \mathbf{r}) - f(\tilde{\mathbf{p}}_{i1}, \mathbf{r}')| \quad (19)$$

$$\stackrel{(a)}{\leq} |f(\bar{\mathbf{p}}, \mathbf{r}') - f(\bar{\mathbf{p}}, \mathbf{r})| + |f(\bar{\mathbf{p}}, \mathbf{r}) - f(\tilde{\mathbf{p}}_{i1}, \mathbf{r})| + |f(\tilde{\mathbf{p}}_{i1}, \mathbf{r}) - f(\tilde{\mathbf{p}}_{i1}, \mathbf{r}')| \quad (20)$$

$$\stackrel{(b)}{\leq} \frac{1}{3}\varepsilon + \frac{1}{3}\varepsilon + \frac{1}{3}\varepsilon = \varepsilon, \quad (21)$$

where (a) holds due to the triangle inequality, and (b) holds via (17) and (18), where we showed that there exists δ_{\min} such that $0 < \delta_{\min} \leq \delta_p$ for all possible transition dynamics parameters \mathbf{p} . \square

Applying Lemma 6 to the two functions $V(\mathbf{p}, \mathbf{r}, \pi_{v_i}(\mathbf{p}, \mathbf{r})) = \max_{\pi'} V(\mathbf{p}, \mathbf{r}, \pi')$ and $V(\mathbf{p}, \mathbf{r}, \pi)$, for any fixed policy π , yields the following result.

Lemma 7. For any $\varepsilon, \delta > 0$, any policy π , and any unit reward vector \mathbf{r} , both of the following hold with probability at least $1 - \delta$ for sufficiently-large i and $j \in \{1, 2\}$:

$$\begin{aligned} |V(\bar{\mathbf{p}}, \mathbf{r}, \pi) - V(\tilde{\mathbf{p}}_{ij}, \mathbf{r}, \pi)| &< \varepsilon \\ |V(\bar{\mathbf{p}}, \mathbf{r}, \pi_{vi}(\bar{\mathbf{p}}, \mathbf{r})) - V(\tilde{\mathbf{p}}_{ij}, \mathbf{r}, \pi_{vi}(\tilde{\mathbf{p}}_{ij}, \mathbf{r}))| &< \varepsilon. \end{aligned}$$

Proof. Both statements follow by applying Lemma 6. First, consider the function $f_1(\mathbf{p}, \mathbf{r}) := V(\mathbf{p}, \mathbf{r}, \pi)$ for a fixed policy π . The value function $V(\mathbf{p}, \mathbf{r}, \pi)$ is continuous in both \mathbf{p} and \mathbf{r} . Furthermore, it is linear in \mathbf{r} and therefore uniformly-continuous in \mathbf{r} : for a linear function $g(\mathbf{z}) = \mathbf{a}^T \mathbf{z}$ and for any $\varepsilon' > 0$, if $\delta' := \frac{\varepsilon'}{\|\mathbf{a}\|}$, then for any $\mathbf{z}_1, \mathbf{z}_2$ such that $\|\mathbf{z}_1 - \mathbf{z}_2\| < \delta'$:

$$|g(\mathbf{z}_1) - g(\mathbf{z}_2)| = |\mathbf{a}^T(\mathbf{z}_1 - \mathbf{z}_2)| \leq \|\mathbf{a}\|_2 \|\mathbf{z}_1 - \mathbf{z}_2\|_2 < \|\mathbf{a}\|_2 \delta' = \varepsilon'.$$

Thus, f_1 satisfies the conditions of Lemma 6 for any fixed π , and so for $i > i_\pi$, $|V(\bar{\mathbf{p}}, \mathbf{r}, \pi) - V(\tilde{\mathbf{p}}_{ij}, \mathbf{r}, \pi)| < \varepsilon$ with probability at least $1 - \delta$. Because there are finitely-many deterministic policies π , we can set $i > \max_\pi i_\pi$, so that the statement holds jointly over all π .

Next, let $f_2(\mathbf{p}, \mathbf{r}) = \max_\pi V(\mathbf{p}, \mathbf{r}, \pi) = V(\mathbf{p}, \mathbf{r}, \pi_{vi}(\mathbf{p}, \mathbf{r}))$. A maximum over finitely-many continuous functions is continuous, and a maximum over finitely-many uniformly-continuous functions is uniformly-continuous. Therefore, f_2 also satisfies the conditions of Lemma 6. \square

We will show convergence in distribution of the reward samples, $\tilde{\mathbf{r}}_{i1}, \tilde{\mathbf{r}}_{i2} \xrightarrow{D} \bar{\mathbf{r}}$, by applying Lemma 5 and demonstrating that $\frac{1}{\beta_i(\delta)^2} \lambda_d^{(i)} \rightarrow \infty$ as $i \rightarrow \infty$. This result is proven by contradiction: intuitively, if $\frac{1}{\beta_i(\delta)^2} \lambda_d^{(i)}$ is upper-bounded, then DPS has a lower-bounded probability of selecting policies that increase $\lambda_d^{(i)}$. Importantly, abbreviating $\lambda_d^{(i)}$'s eigenvector $\mathbf{v}_d^{(i)}$ as \mathbf{v} , this proof is contingent upon there existing a pair of policies π_1, π_2 such that:

$$|\mathbb{E}[\mathbf{x}_i^T \mathbf{v} \mid \pi_{i1} = \pi_1, \pi_{i2} = \pi_2]| = |\mathbb{E}[(\mathbf{x}_{i2} - \mathbf{x}_{i1})^T \mathbf{v} \mid \pi_{i1} = \pi_1, \pi_{i2} = \pi_2]| \stackrel{(a)}{=} |V(\bar{\mathbf{p}}, \mathbf{v}, \pi_1) - V(\bar{\mathbf{p}}, \mathbf{v}, \pi_2)| > 0,$$

where (a) holds because the value function $V(\bar{\mathbf{p}}, \mathbf{v}, \pi)$ gives the expected total reward of π under the reward vector \mathbf{v} . In other words, the proof will require,

$$\max_{\pi_1, \pi_2} |\mathbb{E}[\mathbf{x}_i^T \mathbf{v} \mid \pi_{i1} = \pi_1, \pi_{i2} = \pi_2]| = \max_{\pi_1, \pi_2} |V(\bar{\mathbf{p}}, \mathbf{v}, \pi_1) - V(\bar{\mathbf{p}}, \mathbf{v}, \pi_2)| > 0. \quad (22)$$

If this does not hold, then it is impossible to select a pair of policies under which the observation \mathbf{x}_i is not expected to be orthogonal to the eigenvector \mathbf{v} .

We argue that without loss of generality, (22) can be assumed to hold for all eigenvectors of M_i . Note that if (22) does not hold, then $\mathbb{E}[\mathbf{x}_{i1}^T \mathbf{v} \mid \pi_{i1} = \pi] = V(\bar{\mathbf{p}}, \mathbf{v}, \pi)$ is fixed for all π . Given $\bar{\mathbf{p}}$, by linearity of the value function V in the rewards, any \mathbf{v} -directed component of \mathbf{r} does not affect policy selection: $V(\bar{\mathbf{p}}, \mathbf{r}, \pi) = V(\bar{\mathbf{p}}, \mathbf{r}^{\mathbf{v}} + \mathbf{r}^{\mathbf{v}\perp}, \pi) = V(\bar{\mathbf{p}}, \mathbf{r}^{\mathbf{v}}, \pi) + V(\bar{\mathbf{p}}, \mathbf{r}^{\mathbf{v}\perp}, \pi)$, where $\mathbf{r}^{\mathbf{v}}$ is the projection of \mathbf{r} onto the \mathbf{v} -direction and $\mathbf{r}^{\mathbf{v}\perp}$ is its orthogonal complement in \mathbb{R}^d . Because $V(\bar{\mathbf{p}}, \mathbf{r}^{\mathbf{v}}, \pi)$ does not depend on π , $\pi_{vi}(\bar{\mathbf{p}}, \mathbf{r}) = \operatorname{argmax}_\pi V(\bar{\mathbf{p}}, \mathbf{r}, \pi) = \operatorname{argmax}_\pi V(\bar{\mathbf{p}}, \mathbf{r}^{\mathbf{v}\perp}, \pi)$.

We call any vector \mathbf{v} which does *not* satisfy (22) an *irrelevant dimension* of the rewards: given $\bar{\mathbf{p}}$, removing the \mathbf{v} -directed component of \mathbf{r} does not influence policy selection. The following lemma demonstrates that such vectors remain irrelevant towards policy selection when $\bar{\mathbf{p}}$ is unknown, but once $\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{p}}_{i2}$ have sufficiently converged to $\bar{\mathbf{p}}$ in distribution.

Lemma 8. For any reward vector $\mathbf{r} \in \mathbb{R}^d$, let \mathbf{r}^{rel} be the projection of \mathbf{r} onto the relevant subspace (for which (22) holds), and \mathbf{r}^\perp be its orthogonal complement in \mathbb{R}^d , such that $\mathbf{r} = \mathbf{r}^{\text{rel}} + \mathbf{r}^\perp$, and \mathbf{r}^\perp belongs to the subspace of irrelevant dimensions (where (22) does not hold). The reward samples on iteration i are $\tilde{\mathbf{r}}_{ij}$, $j \in \{1, 2\}$. Then, for any $\varepsilon, \delta > 0$, there exists i_0 such that for $i > i_0$, with probability at least $1 - \delta$:

$$|V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{ij}, \pi_{ij}) - V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{ij}, \pi_{vi}(\tilde{\mathbf{p}}_{ij}, \tilde{\mathbf{r}}_{ij}^{\text{rel}}))| < \varepsilon.$$

In other words, with respect to the sampled rewards $\tilde{\mathbf{r}}_{ij}$, the expected reward of the selected policy $\pi_{ij} = \pi_{vi}(\tilde{\mathbf{p}}_{ij}, \tilde{\mathbf{r}}_{ij})$ is close to the expected reward of the policy that would have been selected were $\tilde{\mathbf{r}}_{ij}$ replaced by $\tilde{\mathbf{r}}_{ij}^{\text{rel}}$.

Proof. We prove the result for $j = 1$ (the proof is identical for $j = 2$). Because $V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^\perp, \pi)$ is constant for all π , we define $w := V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^\perp, \pi)$ for convenience. First, we show that under the true transition dynamics $\bar{\mathbf{p}}$, the irrelevant dimensions of $\tilde{\mathbf{r}}_{i1}$ do not affect policy selection. For any π ,

$$\begin{aligned} V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1})) &= \max_{\pi} V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}, \pi) \stackrel{(a)}{=} \max_{\pi} [V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}, \pi) + V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^\perp, \pi)] \\ &\stackrel{(b)}{=} V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}, \pi_{vi}(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^{\text{rel}})) + V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^\perp, \pi_{vi}(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^{\text{rel}})) \stackrel{(c)}{=} V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^{\text{rel}})), \end{aligned} \quad (23)$$

where (a) and (c) hold because $\tilde{\mathbf{r}}_{i1} = \tilde{\mathbf{r}}_{i1}^{\text{rel}} + \tilde{\mathbf{r}}_{i1}^\perp$ and the value function is linear in the rewards, and (b) holds because $V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^\perp, \pi) = w$ is constant across all policies π .

To upper-bound $|V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}, \pi_{i1}) - V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}))|$, we write:

$$\begin{aligned} &|V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}, \pi_{i1}) - V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}))| \stackrel{(a)}{=} |V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1})) - (V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}})) + w)| \\ &= |V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1})) - (V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}})) + w) \\ &\quad - V(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1})) + V(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}))| - V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1})) + V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}))| \\ &\quad - (V(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}})) + w) + (V(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}})) + w)| \\ &\stackrel{(b)}{\leq} |V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1})) - V(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}))| + |V(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1})) - V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}))| \\ &\quad + |V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1})) - (V(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}})) + w)| \\ &\quad + |(V(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}})) + w) - (V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}})) + w)| \\ &\stackrel{(c)}{\leq} |V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1})) - V(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}))| + |V(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1})) - V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}))| \\ &\quad + |V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}, \pi_{vi}(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^{\text{rel}})) - V(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}))| \\ &\quad + |V(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}})) - V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}, \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}))|, \end{aligned} \quad (24)$$

where (a) applies $\tilde{\mathbf{r}}_{i1} = \tilde{\mathbf{r}}_{i1}^{\text{rel}} + \tilde{\mathbf{r}}_{i1}^\perp$, linearity of the value function in the rewards, and the definition of w ; (b) rearranges terms and uses the triangle inequality; and (c) applies (23) to line (24), that is, $V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}, \pi_{vi}(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1})) = V(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^{\text{rel}}, \pi_{vi}(\bar{\mathbf{p}}, \tilde{\mathbf{r}}_{i1}^{\text{rel}})) + w$.

Each of the four terms in (25) can be upper-bounded with high probability using Lemma 7. In particular, for large enough i , each term is less than $\frac{1}{4}\varepsilon$ with probability at least $1 - \frac{1}{4}\delta$. Therefore, the desired result holds. \square

Remark 1. *The reward dimensions could be irrelevant due to a number of reasons. For instance, because the elements of $\mathbf{x}_i := \mathbf{x}_{i2} - \mathbf{x}_{i1}$ must sum to zero, the vector $[1, 1, \dots, 1]^T$ must always be orthogonal to every observation \mathbf{x}_i . Alternatively, the MDP's transition dynamics could constrain the expected number of visits to a particular state to be constant regardless of the policy.*

Such constraints result in a subspace of \mathbb{R}^d that is irrelevant to learning the optimal policy once the transition dynamics model has converged sufficiently. Therefore, we only need Lemma 5 to be satisfied for eigenvalues of M_i along relevant dimensions in order to asymptotically select the optimal policy. Thus, we can assume that sampled reward vectors $\tilde{\mathbf{r}}_{i1}, \tilde{\mathbf{r}}_{i2}$ have been projected onto the relevant subspace of \mathbb{R}^d . As a result, in proving that $\frac{\beta_i(\delta)^2}{\lambda_d^{(i)}} \xrightarrow{D} 0$ as $i \rightarrow \infty$, we can assume that all eigenvectors of M_i belong to the relevant subspace without loss of generality. More formally, we assume without loss of generality that all eigenvectors $\{\mathbf{v}_j^{(i)}\}$ of M_i satisfy (22).

In Lemma 11, we will show that as $i \rightarrow \infty$, $\frac{\beta_i(\delta)^2}{\lambda_d^{(i)}} \xrightarrow{D} 0$. Combined with Lemma 5, this proves Proposition 2, that the reward samples are convergent in distribution to $\bar{\mathbf{r}}$. Lemma 11 proves this result via contradiction, by first assuming that there exists an i_0 such that for all $i \geq i_0$, $\frac{\beta_i(\delta)^2}{\lambda_d^{(i)}} \geq \alpha$. The following two lemmas, Lemmas 9-10, are intermediate results leading to Lemma 11, which utilize this contradiction hypothesis as a premise. In particular, Lemma 9 demonstrates that under the contradiction hypothesis, there is a non-decaying probability of sampling rewards $\tilde{\mathbf{r}}_{i1}, \tilde{\mathbf{r}}_{i2}$ that are highly-aligned with the eigenvector $\mathbf{v}_d^{(i)}$ of M_i^{-1} corresponding to its largest eigenvalue, $(\lambda_d^{(i)})^{-1}$.

Lemma 9. Assume that for a given iteration i , $\beta_i(\delta)^2 \left(\lambda_d^{(i)}\right)^{-1} \geq \alpha$. Then, the reward samples $\tilde{\mathbf{r}}_{i1}, \tilde{\mathbf{r}}_{i2}$ satisfy:

$$P(\tilde{\mathbf{r}}_{i1}^T \mathbf{v}_d^{(i)} \geq a \max_{j < d} |\tilde{\mathbf{r}}_{i1}^T \mathbf{v}_j^{(i)}|) \geq c(a) > 0, \quad (26)$$

$$P(\tilde{\mathbf{r}}_{i2}^T \mathbf{v}_d^{(i)} \leq -a \max_{j < d} |\tilde{\mathbf{r}}_{i2}^T \mathbf{v}_j^{(i)}|) \geq c(a) > 0, \quad (27)$$

where $c : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a continuous, monotonically-decreasing function.

Proof. Recall that the reward samples $\tilde{\mathbf{r}}_{i1}, \tilde{\mathbf{r}}_{i2}$ are drawn according to (13). We will demonstrate that the reward samples can equivalently be expressed as:

$$\tilde{\mathbf{r}}_{i1} = \hat{\mathbf{r}}_i + \beta_i(\delta) \sum_{j=1}^d \left(\lambda_j^{(i)}\right)^{-\frac{1}{2}} z_{ij} \mathbf{v}_j^{(i)}, z_{ij} \sim \mathcal{N}(0, 1) \text{ i.i.d.}, \quad (28)$$

and similarly for $\tilde{\mathbf{r}}_{i2}$. Similarly to (13), the expression in (28) has a multivariate Gaussian distribution. We take the expectation and covariance of (28) with respect to the variables $\{z_{ij}\}$ to show that they match the expressions in (13):

$$\begin{aligned} \mathbb{E} \left[\hat{\mathbf{r}}_i + \beta_i(\delta) \sum_{j=1}^d \left(\lambda_j^{(i)}\right)^{-\frac{1}{2}} z_{ij} \mathbf{v}_j^{(i)} \right] &= \hat{\mathbf{r}}_i + \beta_i(\delta) \sum_{j=1}^d \left(\lambda_j^{(i)}\right)^{-\frac{1}{2}} \mathbb{E}[z_{ij}] \mathbf{v}_j^{(i)} = \hat{\mathbf{r}}_i, \\ \text{Cov} \left[\hat{\mathbf{r}}_i + \beta_i(\delta) \sum_{j=1}^d \left(\lambda_j^{(i)}\right)^{-\frac{1}{2}} z_{ij} \mathbf{v}_j^{(i)} \right] &\stackrel{(a)}{=} \mathbb{E} \left[\left(\beta_i(\delta) \sum_{j=1}^d \left(\lambda_j^{(i)}\right)^{-\frac{1}{2}} z_{ij} \mathbf{v}_j^{(i)} \right) \left(\beta_i(\delta) \sum_{k=1}^d \left(\lambda_k^{(i)}\right)^{-\frac{1}{2}} z_{ik} \mathbf{v}_k^{(i)} \right)^T \right] \\ &= \beta_i(\delta)^2 \sum_{j=1}^d \sum_{k=1}^d \left(\lambda_j^{(i)}\right)^{-\frac{1}{2}} \left(\lambda_k^{(i)}\right)^{-\frac{1}{2}} \mathbf{v}_j^{(i)} \mathbf{v}_k^{(i)T} \mathbb{E}[z_{ij} z_{ik}] \\ &\stackrel{(b)}{=} \beta_i(\delta)^2 \sum_{j=1}^d \left(\lambda_j^{(i)}\right)^{-1} \mathbf{v}_j^{(i)} \mathbf{v}_j^{(i)T} = \beta_i(\delta)^2 M_i^{-1}, \end{aligned}$$

which match the expectation and covariance in (13). In the above, (a) applies the definition $\text{Cov}[\mathbf{x}] = \mathbb{E}[(\mathbf{x} - \mathbb{E}[\mathbf{x}])(\mathbf{x} - \mathbb{E}[\mathbf{x}])^T]$, and (b) holds because $\mathbb{E}[z_{ij} z_{ik}] = \text{Cov}[z_{ij} z_{ik}] = \delta_{jk}$, where δ_{jk} is the Kronecker delta function.

Next, we show that the probability that $\tilde{\mathbf{r}}_{i1}$ is arbitrarily-aligned with $\mathbf{v}_d^{(i)}$ is lower-bounded above zero: that is, there exists $c : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that for any $a > 0$, $P(\tilde{\mathbf{r}}_{i1}^T \mathbf{v}_d^{(i)} \geq a \max_{j < d} |\tilde{\mathbf{r}}_{i1}^T \mathbf{v}_j^{(i)}|) \geq c(a) > 0$. This can be shown by bounding the terms $|\tilde{\mathbf{r}}_{i1}^T \mathbf{v}_j^{(i)}|$, $j < d$, and $\tilde{\mathbf{r}}_{i1}^T \mathbf{v}_d^{(i)}$. Firstly, the term $|\tilde{\mathbf{r}}_{i1}^T \mathbf{v}_j^{(i)}|$, $j < d$, can be upper-bounded:

$$\begin{aligned} |\tilde{\mathbf{r}}_{i1}^T \mathbf{v}_j^{(i)}| &\stackrel{(a)}{=} \left| \hat{\mathbf{r}}_i^T \mathbf{v}_j^{(i)} + \beta_i(\delta) \sum_{k=1}^d \left(\lambda_k^{(i)}\right)^{-\frac{1}{2}} z_{ik} \mathbf{v}_k^{(i)T} \mathbf{v}_j^{(i)} \right| \stackrel{(b)}{=} \left| \hat{\mathbf{r}}_i^T \mathbf{v}_j^{(i)} + \beta_i(\delta) \left(\lambda_j^{(i)}\right)^{-\frac{1}{2}} z_{ij} \right| \\ &\leq |\hat{\mathbf{r}}_i^T \mathbf{v}_j^{(i)}| + \beta_i(\delta) \left(\lambda_j^{(i)}\right)^{-\frac{1}{2}} |z_{ij}| \stackrel{(c)}{\leq} \|\hat{\mathbf{r}}_i\|_2 \|\mathbf{v}_j^{(i)}\|_2 + \beta_i(\delta) \left(\lambda_j^{(i)}\right)^{-\frac{1}{2}} |z_{ij}| \\ &\stackrel{(d)}{\leq} b + \beta_i(\delta) \left(\lambda_j^{(i)}\right)^{-\frac{1}{2}} |z_{ij}|, \end{aligned}$$

where (a) applies (28), (b) follows from orthonormality of the eigenbasis, (c) follows from the Cauchy-Schwarz inequality, and (d) uses that $\|\hat{\mathbf{r}}_i\|_2 \leq b$ (Lemma 2). Similarly, $\tilde{\mathbf{r}}_{i1}^T \mathbf{v}_d^{(i)}$ can be lower-bounded:

$$\begin{aligned} \tilde{\mathbf{r}}_{i1}^T \mathbf{v}_d^{(i)} &\stackrel{(a)}{=} \hat{\mathbf{r}}_i^T \mathbf{v}_d^{(i)} + \beta_i(\delta) \sum_{j=1}^d \left(\lambda_j^{(i)}\right)^{-\frac{1}{2}} z_{ij} \mathbf{v}_j^{(i)T} \mathbf{v}_d^{(i)} \stackrel{(b)}{=} \hat{\mathbf{r}}_i^T \mathbf{v}_d^{(i)} + \beta_i(\delta) \left(\lambda_d^{(i)}\right)^{-\frac{1}{2}} z_{i1} \\ &\geq -|\hat{\mathbf{r}}_i^T \mathbf{v}_d^{(i)}| + \beta_i(\delta) \left(\lambda_d^{(i)}\right)^{-\frac{1}{2}} z_{i1} \stackrel{(c)}{\geq} -\|\hat{\mathbf{r}}_i\|_2 \|\mathbf{v}_d^{(i)}\|_2 + \beta_i(\delta) \left(\lambda_d^{(i)}\right)^{-\frac{1}{2}} z_{i1} \\ &\stackrel{(d)}{\geq} -b + \beta_i(\delta) \left(\lambda_d^{(i)}\right)^{-\frac{1}{2}} z_{i1}, \end{aligned}$$

where as before, (a) applies (28), (b) follows from orthonormality of the eigenbasis, (c) follows from the Cauchy-Schwarz inequality, and (d) holds via Lemma 2. Given these upper and lower bounds, the probability $P\left(\tilde{\mathbf{r}}_{i1}^T \mathbf{v}_d^{(i)} \geq a \max_{j < d} |\tilde{\mathbf{r}}_{i1}^T \mathbf{v}_j^{(i)}|\right)$ can be lower-bounded:

$$\begin{aligned} P\left(\tilde{\mathbf{r}}_{i1}^T \mathbf{v}_d^{(i)} \geq a \max_{j < d} |\tilde{\mathbf{r}}_{i1}^T \mathbf{v}_j^{(i)}|\right) &\stackrel{(a)}{\geq} P\left(-b + \beta_i(\delta) \left(\lambda_d^{(i)}\right)^{-\frac{1}{2}} z_{i1} \geq a \max_{j < d} \left[b + \beta_i(\delta) \left(\lambda_j^{(i)}\right)^{-\frac{1}{2}} |z_{ij}| \right]\right) \\ &= P\left(z_{i1} \geq \frac{b\sqrt{\lambda_d^{(i)}}}{\beta_i(\delta)} + a \max_{j < d} \left[\frac{b\sqrt{\lambda_d^{(i)}}}{\beta_i(\delta)} + \sqrt{\frac{\lambda_d^{(i)}}{\lambda_j^{(i)}}} |z_{ij}| \right]\right) \\ &\stackrel{(b)}{\geq} P\left(z_{i1} \geq \frac{b}{\sqrt{\alpha}} + a \max_{j < d} \left[\frac{b}{\sqrt{\alpha}} + |z_{ij}| \right]\right) \\ &= P\left(z_{i1} \geq \frac{b(1+a)}{\sqrt{\alpha}} + a \max_{j < d} |z_{ij}| \right) := c(a) > 0, \end{aligned}$$

where (a) results from the upper and lower bounds derived above, and (b) follows because $\frac{\lambda_d^{(i)}}{\lambda_j^{(i)}} \leq 1$ and $\beta_i(\delta) \left(\lambda_d^{(i)}\right)^{-\frac{1}{2}} \geq \sqrt{\alpha}$ by assumption. The function $c(a) > 0$ is continuous and decreasing in a .

By identical arguments, $P(\tilde{\mathbf{r}}_{i2}^T \mathbf{v}_d^{(i)} \leq -a \max_{j < d} |\tilde{\mathbf{r}}_{i2}^T \mathbf{v}_j^{(i)}|) \geq c(a)$. Thus, for any $a > 0$ and set of eigenvectors $\mathbf{v}_j^{(i)}$:

$$\begin{aligned} P(\tilde{\mathbf{r}}_{i1}^T \mathbf{v}_d^{(i)} \geq a \max_{j < d} |\tilde{\mathbf{r}}_{i1}^T \mathbf{v}_j^{(i)}|) &\geq c(a) > 0, \\ P(\tilde{\mathbf{r}}_{i2}^T \mathbf{v}_d^{(i)} \leq -a \max_{j < d} |\tilde{\mathbf{r}}_{i2}^T \mathbf{v}_j^{(i)}|) &\geq c(a) > 0. \end{aligned}$$

□

Next, we show that given sampled rewards $\tilde{\mathbf{r}}_{i1}, \tilde{\mathbf{r}}_{i2}$ that are highly-aligned with the eigenvector $\mathbf{v}_d^{(i)}$ of M_i as in Lemma 9, there is a lower-bounded probability of sampling trajectories that have non-zero projection onto this eigenvector.

Lemma 10. *Assume that there exists i_0 such that for $i > i_0$, $\beta_i(\delta)^2 \left(\lambda_d^{(i)}\right)^{-1} \geq \alpha$. Then, there exists $i' \geq i_0$ and a constant $c' > 0$ such that for $i > i'$:*

$$\mathbb{E} \left[\left| \mathbf{x}_i^T \mathbf{v}_d^{(i)} \right| \right] \geq c' > 0, \quad (29)$$

where $c' > 0$ depends only on the MDP parameters $\bar{\mathbf{p}}$ and $\bar{\mathbf{r}}$, so that in particular, (29) holds for any eigenvector $\mathbf{v}_d^{(i)}$.

Proof. By Lemma 9, (26) and (27) both hold. We will refer to the events in (26) and (27), $\{\tilde{\mathbf{r}}_{i1}^T \mathbf{v}_d^{(i)} \geq a \max_{j < d} |\tilde{\mathbf{r}}_{i1}^T \mathbf{v}_j^{(i)}|\}$ and $\{\tilde{\mathbf{r}}_{i2}^T \mathbf{v}_d^{(i)} \leq -a \max_{j < d} |\tilde{\mathbf{r}}_{i2}^T \mathbf{v}_j^{(i)}|\}$, as events $A(a)$ and $B(a)$, respectively. From Lemma 9, $A(a)$ and $B(a)$ have positive probability for any a .

We will show that by setting a to a large-enough value, under events $A(a)$ and $B(a)$, value iteration samples policies π_{i1}, π_{i2} such that $\mathbb{E} \left[\left| \mathbf{x}_i^T \mathbf{v}_d^{(i)} \right| \right] \geq c' > 0$ for some $c' > 0$, for sufficiently-high i and for any unit vector $\mathbf{v}_d^{(i)}$.

First, note that under events $A(a)$ and $B(a)$, as $a \rightarrow \infty$, $\frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2} \rightarrow \mathbf{v}_d^{(i)}$ and $\frac{\tilde{\mathbf{r}}_{i2}}{\|\tilde{\mathbf{r}}_{i2}\|_2} \rightarrow -\mathbf{v}_d^{(i)}$. Let $\varepsilon > 0$. Under event $A(a)$ for sufficiently-large a , $\left\| \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2} - \mathbf{v}_d^{(i)} \right\|_2 < \varepsilon$. Define $a_{\min,1}(\varepsilon, \mathbf{v}_1^{(i)}, \dots, \mathbf{v}_d^{(i)})$ as the minimum value of a such that $A(a)$ implies $\left\| \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2} - \mathbf{v}_d^{(i)} \right\|_2 \leq \frac{\varepsilon}{2}$, given the eigenbasis $\{\mathbf{v}_1^{(i)}, \dots, \mathbf{v}_d^{(i)}\}$. Because the inequality defining $A(a)$ is continuous in a , $\tilde{\mathbf{r}}_{i1}$, and the eigenbasis $\{\mathbf{v}_1^{(i)}, \dots, \mathbf{v}_d^{(i)}\}$, the function $a_{\min,1}(\varepsilon, \mathbf{v}_1^{(i)}, \dots, \mathbf{v}_d^{(i)})$ is also continuous in the eigenbasis $\{\mathbf{v}_1^{(i)}, \dots, \mathbf{v}_d^{(i)}\}$. Because $a_{\min,1}(\varepsilon, \mathbf{v}_1^{(i)}, \dots, \mathbf{v}_d^{(i)})$ is positive for all $\{\mathbf{v}_1^{(i)}, \dots, \mathbf{v}_d^{(i)}\}$, and the set of all eigenbases $\{\mathbf{v}_1^{(i)}, \dots, \mathbf{v}_d^{(i)}\}$ is compact, there exists $a_{\min,1}(\varepsilon)$ such that for any eigenbasis, if $A(a)$ holds for $a \geq a_{\min,1}(\varepsilon)$, then $\left\| \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2} - \mathbf{v}_d^{(i)} \right\|_2 < \varepsilon$.

By the same arguments, there exists $a_{\min,2}(\varepsilon)$ such that for any eigenbasis, if $B(a)$ holds for $a \geq a_{\min,2}(\varepsilon)$, then $\left\| \frac{\tilde{\mathbf{r}}_{i2}}{\|\tilde{\mathbf{r}}_{i2}\|_2} - (-\mathbf{v}_d^{(i)}) \right\|_2 < \varepsilon$. Taking $a_{\min}(\varepsilon) := \max\{a_{\min,1}(\varepsilon), a_{\min,2}(\varepsilon)\}$, then for any $a \geq a_{\min}(\varepsilon)$, under events $A(a)$ and $B(a)$, both $\left\| \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2} - \mathbf{v}_d^{(i)} \right\|_2 < \varepsilon$ and $\left\| \frac{\tilde{\mathbf{r}}_{i2}}{\|\tilde{\mathbf{r}}_{i2}\|_2} - (-\mathbf{v}_d^{(i)}) \right\|_2 < \varepsilon$ hold.

Next, we will show that by setting ε small enough, the inequality $\left\| \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2} - \mathbf{v}_d^{(i)} \right\|_2 < \varepsilon$ implies that the expected reward accrued by π_{i1} with respect to $\mathbf{v}_d^{(i)}$, that is, $V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{i1})$, is close to the maximum possible expected reward with respect to $\mathbf{v}_d^{(i)}$, $\max_{\pi} V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi) = V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{vi}(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}))$. (The same approach yields an equivalent result for $\tilde{\mathbf{r}}_{i2}$.)

Assume that $\left\| \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2} - \mathbf{v}_d^{(i)} \right\|_2 < \varepsilon$, and let $\varepsilon' > 0$. We will show that $|V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{vi}(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)})) - V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{i1})| < \varepsilon'$ for small-enough ε and when $\tilde{\mathbf{p}}_{i1}$ has sufficiently converged in distribution to $\bar{\mathbf{p}}$ (as is guaranteed to occur by Proposition 1):

$$\begin{aligned}
& \left| V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{vi}(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)})) - V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{i1}) \right| \stackrel{(a)}{=} \left| V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{vi}(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)})) - V\left(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{vi}\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)\right) \right| \\
&= \left| V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{vi}(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)})) - V\left(\bar{\mathbf{p}}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}, \pi_{vi}\left(\bar{\mathbf{p}}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)\right) \right. \\
&\quad + V\left(\bar{\mathbf{p}}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}, \pi_{vi}\left(\bar{\mathbf{p}}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)\right) - V\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}, \pi_{vi}\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)\right) \\
&\quad + V\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}, \pi_{vi}\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)\right) - V\left(\bar{\mathbf{p}}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}, \pi_{vi}\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)\right) \\
&\quad \left. + V\left(\bar{\mathbf{p}}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}, \pi_{vi}\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)\right) - V\left(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{vi}\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)\right) \right| \\
&\stackrel{(b)}{\leq} \left| V\left(\bar{\mathbf{p}}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}, \pi_{vi}\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)\right) - V\left(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{vi}\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)\right) \right| \tag{30} \\
&\quad + \left| V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{vi}(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)})) - V\left(\bar{\mathbf{p}}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}, \pi_{vi}\left(\bar{\mathbf{p}}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)\right) \right| \tag{31} \\
&\quad + \left| V\left(\bar{\mathbf{p}}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}, \pi_{vi}\left(\bar{\mathbf{p}}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)\right) - V\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}, \pi_{vi}\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)\right) \right| \tag{32} \\
&\quad + \left| V\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}, \pi_{vi}\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)\right) - V\left(\bar{\mathbf{p}}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}, \pi_{vi}\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)\right) \right|, \tag{33}
\end{aligned}$$

where (a) uses that $\pi_{i1} = \pi_{vi}(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1})$ by definition, and also that positive scaling of the reward argument of $\pi_{vi}(\mathbf{p}, \mathbf{r})$ does not affect its output; and (b) applies the triangle inequality. Next, we will show that each of (30)-(33) can be upper-bounded by $\frac{1}{4}\varepsilon'$ (for (32) and (33) with high probability) by appropriate choice of ε and by utilizing that $\tilde{\mathbf{p}}_{i1} \xrightarrow{D} \bar{\mathbf{p}}$ (Proposition 1).

Beginning with (30), because $V(\mathbf{p}, \mathbf{r}, \pi)$ is linear in \mathbf{r} , it is uniformly continuous in \mathbf{r} for fixed transition dynamics \mathbf{p} and policy π . So, for fixed dynamics $\bar{\mathbf{p}}$ and policy π and for any reward vector \mathbf{r} , there exists ε_{π} such that if $\|\mathbf{r} - \mathbf{r}'\| < \varepsilon_{\pi}$, then $|V(\bar{\mathbf{p}}, \mathbf{r}, \pi) - V(\bar{\mathbf{p}}, \mathbf{r}', \pi)| < \frac{1}{4}\varepsilon'$. Because there are finitely-many deterministic policies, there exists $\varepsilon_1 > 0$ such that $\varepsilon_1 \leq \varepsilon_{\pi}$ for all π . Therefore, for any policy π , if $\|\mathbf{r} - \mathbf{r}'\|_2 < \varepsilon_1$, then $|V(\bar{\mathbf{p}}, \mathbf{r}, \pi) - V(\bar{\mathbf{p}}, \mathbf{r}', \pi)| < \frac{1}{4}\varepsilon'$. The expression in (30) is thus upper-bounded by $\frac{1}{4}\varepsilon'$ if $\varepsilon < \varepsilon_1$.

To upper-bound (31), observe that $V(\mathbf{p}, \mathbf{r}, \pi_{vi}(\mathbf{p}, \mathbf{r})) = \max_{\pi} V(\mathbf{p}, \mathbf{r}, \pi)$ is also uniformly continuous in \mathbf{r} for fixed transition dynamics \mathbf{p} : the maximum over finitely-many uniformly continuous functions is also uniformly continuous. Thus, there exists $\varepsilon_2 > 0$ such that if $\|\mathbf{r} - \mathbf{r}'\|_2 < \varepsilon_2$, then $|V(\bar{\mathbf{p}}, \mathbf{r}, \pi_{vi}(\bar{\mathbf{p}}, \mathbf{r})) - V(\bar{\mathbf{p}}, \mathbf{r}', \pi_{vi}(\bar{\mathbf{p}}, \mathbf{r}'))| < \frac{1}{4}\varepsilon'$. The expression in (31) is thus upper-bounded by $\frac{1}{4}\varepsilon'$ if $\varepsilon < \varepsilon_2$.

To obtain high-probability upper bounds for (32) and (33), we apply Lemma 7. For sufficiently-high i , each of the

following holds with probability at least $1 - \frac{1}{2}\delta'$:

$$\begin{aligned} & \left| V\left(\bar{\mathbf{p}}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}, \pi_{vi}\left(\bar{\mathbf{p}}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)\right) - V\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}, \pi_{vi}\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)\right) \right| < \frac{1}{4}\varepsilon', \\ & \left| V\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}, \pi\right) - V\left(\bar{\mathbf{p}}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}, \pi\right) \right| < \frac{1}{4}\varepsilon', \end{aligned}$$

where the second statement holds for any policy π , and in particular for $\pi = \pi_{vi}\left(\tilde{\mathbf{p}}_{i1}, \frac{\tilde{\mathbf{r}}_{i1}}{\|\tilde{\mathbf{r}}_{i1}\|_2}\right)$.

We combine the bounds for (30)-(33), setting $\varepsilon < \min\{\varepsilon_1, \varepsilon_2\}$ and $i > i'$. Thus, for any $\varepsilon', \delta' > 0$, we have shown that by setting ε small enough and taking $i > i'$:

$$\left| V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{vi}(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)})) - V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{i1}) \right| < \varepsilon' \text{ with probability at least } 1 - \delta'.$$

Combining with the analogous result for π_{i2} and $-\mathbf{v}_d^{(i)}$ yields that for any $\varepsilon', \delta' > 0$, there exists sufficiently-small ε and large-enough i' such that for $i > i'$:

$$\begin{aligned} & \left| V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{vi}(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)})) - V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{i1}) \right| < \varepsilon' \text{ with probability at least } 1 - \delta', \text{ and} \\ & \left| V(\bar{\mathbf{p}}, -\mathbf{v}_d^{(i)}, \pi_{vi}(\bar{\mathbf{p}}, -\mathbf{v}_d^{(i)})) - V(\bar{\mathbf{p}}, -\mathbf{v}_d^{(i)}, \pi_{i2}) \right| < \varepsilon' \text{ with probability at least } 1 - \delta'. \end{aligned} \quad (34)$$

Next, we will set ε' to a small enough number to achieve $\left| \mathbb{E}[\mathbf{x}_i^T \mathbf{v}_d^{(i)}] \right| > \varepsilon' > 0$. Firstly, note that that $\left| \mathbb{E}[\mathbf{x}_i^T \mathbf{v}_d^{(i)}] \right|$ is maximized when setting $\pi_{i1} = \pi_{vi}(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)})$ and $\pi_{i2} = \pi_{vi}(\bar{\mathbf{p}}, -\mathbf{v}_d^{(i)})$:

$$\begin{aligned} \max_{\pi_1, \pi_2} \left| \mathbb{E}[\mathbf{x}_i^T \mathbf{v}_d^{(i)}] \mid \pi_{i1} = \pi_1, \pi_{i2} = \pi_2 \right| &= \max_{\pi_1, \pi_2} \left| \mathbb{E}[\mathbf{x}_{i1}^T \mathbf{v}_d^{(i)} - \mathbf{x}_{i2}^T \mathbf{v}_d^{(i)}] \mid \pi_{i1} = \pi_1, \pi_{i2} = \pi_2 \right| \\ &= \left| \mathbb{E}[\mathbf{x}_{i1}^T \mathbf{v}_d^{(i)} - \mathbf{x}_{i2}^T \mathbf{v}_d^{(i)} \mid \pi_{i1} = \operatorname{argmax}_{\pi} \mathbb{E}[\mathbf{x}_{i1}^T \mathbf{v}_d^{(i)}], \pi_{i2} = \operatorname{argmin}_{\pi} \mathbb{E}[\mathbf{x}_{i2}^T \mathbf{v}_d^{(i)}]] \right| \\ &= \left| \mathbb{E}[\mathbf{x}_{i1}^T \mathbf{v}_d^{(i)} - \mathbf{x}_{i2}^T \mathbf{v}_d^{(i)} \mid \pi_{i1} = \operatorname{argmax}_{\pi} \mathbb{E}[\mathbf{x}_{i1}^T \mathbf{v}_d^{(i)}], \pi_{i2} = \operatorname{argmax}_{\pi} \mathbb{E}[\mathbf{x}_{i2}^T (-\mathbf{v}_d^{(i)})]] \right| \\ &= \left| \mathbb{E}[\mathbf{x}_{i1}^T \mathbf{v}_d^{(i)} - \mathbf{x}_{i2}^T \mathbf{v}_d^{(i)} \mid \pi_{i1} = \operatorname{argmax}_{\pi} V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi), \pi_{i2} = \operatorname{argmax}_{\pi} V(\bar{\mathbf{p}}, -\mathbf{v}_d^{(i)}, \pi)] \right| \\ &= \left| \mathbb{E}[\mathbf{x}_{i1}^T \mathbf{v}_d^{(i)} - \mathbf{x}_{i2}^T \mathbf{v}_d^{(i)} \mid \pi_{i1} = \pi_{vi}(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}), \pi_{i2} = \pi_{vi}(\bar{\mathbf{p}}, -\mathbf{v}_d^{(i)})] \right|. \end{aligned}$$

From Lemma 8 and Remark 1, we can assume without loss of generality that for all $\mathbf{v}_d^{(i)}$, $\max_{\pi_1, \pi_2} \left| \mathbb{E}[\mathbf{x}_i^T \mathbf{v}_d^{(i)}] \mid \pi_{i1} = \pi_1, \pi_{i2} = \pi_2 \right| > 0$. Because $\left| \mathbb{E}[\mathbf{x}_i^T \mathbf{v}_d^{(i)}] \mid \pi_{i1} = \pi_1, \pi_{i2} = \pi_2 \right|$ is continuous in $\mathbf{v}_d^{(i)}$ for fixed π_1, π_2 , and a maximum over finitely-many continuous functions is continuous, $\max_{\pi_1, \pi_2} \left| \mathbb{E}[\mathbf{x}_i^T \mathbf{v}_d^{(i)}] \mid \pi_{i1} = \pi_1, \pi_{i2} = \pi_2 \right|$ is also continuous in $\mathbf{v}_d^{(i)}$. Because $\mathbf{v}_d^{(i)}$ belongs to the compact set of unit vectors, the expression achieves a minimum positive value on the set of possible $\mathbf{v}_d^{(i)}$, and thus, there exists $\eta > 0$ such that $\max_{\pi_1, \pi_2} \left| \mathbb{E}[\mathbf{x}_i^T \mathbf{v}_d^{(i)}] \mid \pi_{i1} = \pi_1, \pi_{i2} = \pi_2 \right| \geq \eta > 0$. Setting $\varepsilon' := \frac{\eta}{3}$:

$$\begin{aligned} 0 < 3\varepsilon' = \eta &\leq \max_{\pi_1, \pi_2} \left| \mathbb{E}[\mathbf{x}_i^T \mathbf{v}_d^{(i)}] \mid \pi_{i1} = \pi_1, \pi_{i2} = \pi_2 \right| = \left| \max_{\pi_1} \mathbb{E}[\mathbf{x}_{i1}^T \mathbf{v}_d^{(i)} \mid \pi_{i1} = \pi_1] - \min_{\pi_2} \mathbb{E}[\mathbf{x}_{i2}^T \mathbf{v}_d^{(i)} \mid \pi_{i2} = \pi_2] \right| \\ &= \left| \max_{\pi_1} V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_1) - \min_{\pi_2} V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_2) \right| = \left| \max_{\pi_1} V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_1) + \max_{\pi_2} [-V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_2)] \right| \\ &= \left| \max_{\pi_1} V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_1) + \max_{\pi_2} V(\bar{\mathbf{p}}, -\mathbf{v}_d^{(i)}, \pi_2) \right| = \left| V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{vi}(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)})) + V(\bar{\mathbf{p}}, -\mathbf{v}_d^{(i)}, \pi_{vi}(\bar{\mathbf{p}}, -\mathbf{v}_d^{(i)})) \right| \\ &\stackrel{(a)}{\leq} \left| V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{i1}) + V(\bar{\mathbf{p}}, -\mathbf{v}_d^{(i)}, \pi_{i2}) \right| + 2\varepsilon', \end{aligned}$$

where (a) holds with probability at least $1 - 2\delta'$ by (34). Rearranging terms, with probability at least $1 - 2\delta'$,

$$\varepsilon' < \left| V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{i1}) + V(\bar{\mathbf{p}}, -\mathbf{v}_d^{(i)}, \pi_{i2}) \right| = \left| V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{i1}) - V(\bar{\mathbf{p}}, \mathbf{v}_d^{(i)}, \pi_{i2}) \right| = \left| \mathbb{E}[\mathbf{x}_i^T \mathbf{v}_d^{(i)}] \right|.$$

This implies that:

$$\mathbb{E} \left[\left| \mathbf{x}_i^T \mathbf{v}_d^{(i)} \right| \right] \stackrel{(a)}{\geq} \left| \mathbb{E} \left[\mathbf{x}_i^T \mathbf{v}_d^{(i)} \right] \right| \geq c' > 0 \text{ for some positive } c' \text{ and for all } i > i' \text{ and } \mathbf{v}_d^{(i)},$$

where (a) holds via Jensen's inequality and $c' := \varepsilon'$. \square

We are now equipped to complete the proof of asymptotic consistency of the reward model.

Lemma 11. *As $i \rightarrow \infty$, $\frac{\beta_i(\delta)^2}{\lambda_d^{(i)}} \xrightarrow{D} 0$, where $\lambda_d^{(i)}$ is the minimum eigenvalue of M_i .*

Proof. We will first show that $\liminf_{i \rightarrow \infty} \beta_i(\delta)^2 \left(\lambda_d^{(i)} \right)^{-1} = 0$ via a proof by contradiction. Assume that:

$$\liminf_{i \rightarrow \infty} \beta_i(\delta)^2 \left(\lambda_d^{(i)} \right)^{-1} = 2\alpha > 0. \quad (35)$$

If (35) is true, then there exists i_0 such that for all $i \geq i_0$, $\beta_i(\delta)^2 \left(\lambda_d^{(i)} \right)^{-1} \geq \alpha$. In the following, we assume that $i \geq i_0$. Since $\beta_i(\delta)$ increases at most logarithmically in i , it suffices to show that $\lambda_d^{(i)}$ increases at least linearly on average to achieve a contradiction with (35).

Under the contradiction hypothesis, Lemmas 9 and 10 both hold. Due to Lemma 10, DPS will infinitely-often, and at a non-decaying rate, sample trajectory pairs such that $|\mathbf{x}_i^T \mathbf{v}_d^{(i)}|$ is lower-bounded away from zero. At iteration n , we analyze the effect of this guarantee upon $\lambda_d^{(n)}$. Note that $\lambda_d^{(n)}$ corresponds to the eigenvector $\mathbf{v}_d^{(n)}$ of M_n , and so:

$$\lambda_d^{(n)} = \mathbf{v}_d^{(n)T} M_n \mathbf{v}_d^{(n)} \stackrel{(a)}{=} \mathbf{v}_d^{(n)T} \left(\lambda I + \sum_{i=1}^{n-1} \mathbf{x}_i \mathbf{x}_i^T \right) \mathbf{v}_d^{(n)} = \lambda + \sum_{i=1}^{n-1} \left(\mathbf{x}_i^T \mathbf{v}_d^{(n)} \right)^2, \quad (36)$$

where (a) follows from the definition of M_n . Note that while the right-hand side expression of (36) depends upon $\mathbf{x}_i^T \mathbf{v}_d^{(n)}$ for $i < n$, (29) depends upon $\mathbf{x}_i^T \mathbf{v}_d^{(i)}$. Clearly, if $\mathbf{v}_d^{(i)}$ remained constant in i , then the combination of (29) and (36) would suffice to prove that $\lambda_d^{(n)} \rightarrow \infty$ with at least a linear on-average rate; however, $\mathbf{v}_d^{(i)}$ can vary with i over the entire space of unit vectors in \mathbb{R}^d .

We leverage that $\mathbf{v}_d^{(i)}$ is a unit vector, that the set of unit vectors in \mathbb{R}^d is compact, and that any infinite cover of a compact set has a finite subcover. In particular, for any $\varepsilon > 0$, there exist sets $S_1, \dots, S_K \subset \mathbb{R}^d$, $K < \infty$, such that:

1. For $\mathbf{v} \in \mathbb{R}^d$ such that $\|\mathbf{v}\|_2 = 1$, $\mathbf{v} \in S_k$ for some $k \in \{1, \dots, K\}$, and
2. If $\mathbf{v}_1, \mathbf{v}_2 \in S_k$, then $\|\mathbf{v}_1 - \mathbf{v}_2\| < \varepsilon$.

We will show that there exists a sequence $(n_i) \in \mathbb{N}$ such that $\mathbf{v}_d^{(n_i)} \in S_k$ for fixed $k \in \{1, \dots, K\}$, with the events $\mathbf{v}_d^{(n_i)} \in S_k$ corresponding to the indices (n_i) occurring at some non-decaying frequency. Then, by appropriately choosing ε , we will use (36) and the mutual proximity of the vectors $\mathbf{v}_d^{(n_i)}$ to show that $\lambda_d^{(n)}$ increases with an at-least linear rate.

Observe that for any number of total iterations N , there exists an integer $k \in \{1, \dots, K\}$ such that $\mathbf{v}_d^{(i)} \in S_k$ during at least $\frac{N}{K}$ iterations. Because K is a constant and $\frac{N}{K}$ is linear in N , the number of iterations in which $\mathbf{v}_d^{(i)} \in S_k$ is at least linear in N for some k . The right-hand sum in (36) can then be divided according to the indices (n_i) and the remaining indices:

$$\lambda_d^{(n_j)} = \lambda + \sum_{i=1}^{n_j-1} \left(\mathbf{x}_i^T \mathbf{v}_d^{(n_i)} \right)^2 = \lambda + \sum_{i=1}^{j-1} \left(\mathbf{x}_{n_i}^T \mathbf{v}_d^{(n_i)} \right)^2 + \sum_{j=1; j \notin \{n_1, n_2, \dots, n_{j-1}\}}^{n_j-1} \left(\mathbf{x}_j^T \mathbf{v}_d^{(n_i)} \right)^2. \quad (37)$$

The latter sum in (37) is non-decreasing in n_j , as all of its terms are non-negative. In the former sum, $\|\mathbf{v}_d^{(n_j)} - \mathbf{v}_d^{(n_i)}\| < \varepsilon$ for each $i \in \{1, \dots, j-1\}$. Defining $\boldsymbol{\delta} := \mathbf{v}_d^{(n_j)} - \mathbf{v}_d^{(n_i)}$, so that $\|\boldsymbol{\delta}\|_2 \leq \varepsilon$:

$$\left(\mathbf{x}_{n_i}^T \mathbf{v}_d^{(n_i)}\right)^2 = \left(\mathbf{x}_{n_i}^T \left(\mathbf{v}_d^{(n_i)} + \boldsymbol{\delta}\right)\right)^2 = \left(\mathbf{x}_{n_i}^T \mathbf{v}_d^{(n_i)} + \mathbf{x}_{n_i}^T \boldsymbol{\delta}\right)^2 \geq \left(\left|\mathbf{x}_{n_i}^T \mathbf{v}_d^{(n_i)}\right| - \left|\mathbf{x}_{n_i}^T \boldsymbol{\delta}\right|\right)^2. \quad (38)$$

By the Cauchy-Schwarz inequality, $|\mathbf{x}_{n_i}^T \boldsymbol{\delta}| \leq \|\boldsymbol{\delta}\|_2 * \|\mathbf{x}_{n_i}\|_2 \leq 2\varepsilon h$, where h is the trajectory horizon. Because (29) requires that $\mathbb{E} \left[\left| \mathbf{x}_{n_i}^T \mathbf{v}_d^{(n_i)} \right| \right] \geq c' > 0$, one can choose ε small enough that $\mathbb{E} \left[\left| \mathbf{x}_{n_i}^T \mathbf{v}_d^{(n_i)} \right| - \left| \mathbf{x}_{n_i}^T \boldsymbol{\delta} \right| \right] \geq c' - 2\varepsilon h \geq c'' > 0$, and:

$$\mathbb{E} \left[\left(\mathbf{x}_{n_i}^T \mathbf{v}_d^{(n_i)} \right)^2 \right] \stackrel{(a)}{\geq} \mathbb{E} \left[\left(\left| \mathbf{x}_{n_i}^T \mathbf{v}_d^{(n_i)} \right| - \left| \mathbf{x}_{n_i}^T \boldsymbol{\delta} \right| \right)^2 \right] \stackrel{(b)}{\geq} \mathbb{E} \left[\left| \mathbf{x}_{n_i}^T \mathbf{v}_d^{(n_i)} \right| - \left| \mathbf{x}_{n_i}^T \boldsymbol{\delta} \right| \right]^2 \geq (c'')^2 > 0,$$

where (a) takes expectations of both sides of (38), and (b) follows from Jensen's inequality. Merging this result with (37) implies that $\lambda_d^{(n_j)}$ is expected to increase at least linearly on average, according to the positive constant c'' , over the indices (n_i) . Recall that there always exists an S_k such that the number of times when $\mathbf{v}_d^{(i)} \in S_k$ is at least linear in the total number of iterations N . Thus, the rate at which indices (n_i) occur is always (at least) linear in N on average, and $\lambda_d^{(n_j)}$ increases at least linearly in N in expectation.

We demonstrate that $\liminf_{i \rightarrow \infty} \frac{\beta_i(\delta)^2}{\lambda_d^{(i)}} = 0$ holds: the numerator of $\frac{\beta_i(\delta)^2}{\lambda_d^{(i)}}$ is the square of a quantity that increases at most logarithmically in i , while the denominator increases at least linearly in i on average. This contradicts the assumption in (35), and so $\liminf_{i \rightarrow \infty} \beta_i(\delta)^2 \left(\lambda_d^{(i)} \right)^{-1} = 0$ must hold.

Finally, we leverage that $\liminf_{i \rightarrow \infty} \beta_i(\delta)^2 \left(\lambda_d^{(i)} \right)^{-1} = 0$ to show that $\lim_{i \rightarrow \infty} \beta_i(\delta)^2 \left(\lambda_d^{(i)} \right)^{-1} = 0$. Consider the following two possible cases: 1) $\beta_i(\delta)^2 \left(\lambda_d^{(i)} \right)^{-1}$ converges to zero in probability, and 2) $\beta_i(\delta)^2 \left(\lambda_d^{(i)} \right)^{-1}$ does not converge to zero in probability. In case 1), because convergence in probability implies convergence in distribution, the desired result holds.

In case 2), there exists some $\varepsilon > 0$ such that $P \left(\beta_i(\delta)^2 \left(\lambda_d^{(i)} \right)^{-1} \geq \varepsilon \right) \not\rightarrow 0$. In this case, one can apply the same arguments used to show that $\liminf_{i \rightarrow \infty} \beta_i(\delta)^2 \left(\lambda_d^{(i)} \right)^{-1} = 0$, but specifically over time indices where $\beta_i(\delta)^2 \left(\lambda_d^{(i)} \right)^{-1} \geq \varepsilon$. Due to the non-convergence in probability, these time indices must occur at some non-decaying rate, and so the same analysis applies. Thus, $\lambda_d^{(i)}$ increases in i with at least a minimum linear average rate, while $\beta_i(\delta)$ increases at most logarithmically in i . This violates the non-convergence assumption of case 2), resulting in a contradiction. Therefore, only case 1) can hold. □

From the asymptotic consistency of the dynamics and reward samples, one can show that the sampled policies converge in distribution to the optimal policy:

Theorem 1. *With probability $1 - \delta$, the sampled policies π_{i1}, π_{i2} converge in distribution to the optimal policy, π^* , as $i \rightarrow \infty$. That is, $P(\pi_{i1} = \pi^*) \rightarrow 1$ and $P(\pi_{i2} = \pi^*) \rightarrow 1$ as $i \rightarrow \infty$.*

Proof. It suffices to show that $P(\pi_{i1} = \pi^*) \rightarrow 1$ as $i \rightarrow \infty$, as the proof is identical for π_{i2} . From Propositions 1 and 2, respectively, we have that $\tilde{\mathbf{p}}_{i1} \xrightarrow{D} \bar{\mathbf{p}}$ and that $\tilde{\mathbf{r}}_{i1} \xrightarrow{D} \bar{\mathbf{r}}$ with probability $1 - \delta$. We proceed under the assumption that $\tilde{\mathbf{r}}_{i1} \xrightarrow{D} \bar{\mathbf{r}}$, i.e., that the probability- $(1 - \delta)$ event occurs.

By Fact 1 in Appendix A.5, for each fixed π , $V(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}, \pi) \xrightarrow{D} V(\bar{\mathbf{p}}, \bar{\mathbf{r}}, \pi)$, as value functions are continuous in the dynamics and reward parameters. Applying Fact 2 in Appendix A.5, for each fixed π and $\varepsilon > 0$:

$$P(|V(\tilde{\mathbf{p}}_{i1}, \tilde{\mathbf{r}}_{i1}, \pi) - V(\bar{\mathbf{p}}, \bar{\mathbf{r}}, \pi)| > \varepsilon) \rightarrow 0 \text{ as } i \rightarrow \infty. \quad (39)$$

Next, we set the value of ε to be less than half of the smallest gap between the value of the optimal policy and the value of any suboptimal policy:

$$\varepsilon < \frac{1}{2} \left[\max_{\pi} V(\bar{\boldsymbol{p}}, \bar{\boldsymbol{r}}, \pi) - \max_{\pi \text{ s.t. } V(\bar{\boldsymbol{p}}, \bar{\boldsymbol{r}}, \pi) < \max_{\pi'} V(\bar{\boldsymbol{p}}, \bar{\boldsymbol{r}}, \pi')} V(\bar{\boldsymbol{p}}, \bar{\boldsymbol{r}}, \pi) \right].$$

Then, the probability of selecting a non-optimal policy can be upper-bounded by a quantity that decays with i :

$$\begin{aligned} P(\pi_{i1} \neq \pi^*) &\stackrel{(a)}{\leq} P \left(\bigcup_{\pi} \{ |V(\tilde{\boldsymbol{p}}_{i1}, \tilde{\boldsymbol{r}}_{i1}, \pi) - V(\bar{\boldsymbol{p}}, \bar{\boldsymbol{r}}, \pi)| > \varepsilon \} \right) \\ &\stackrel{(b)}{\leq} \sum_{\pi} P(|V(\tilde{\boldsymbol{p}}_{i1}, \tilde{\boldsymbol{r}}_{i1}, \pi) - V(\bar{\boldsymbol{p}}, \bar{\boldsymbol{r}}, \pi)| > \varepsilon) \stackrel{(c)}{\rightarrow} 0 \text{ as } i \rightarrow \infty, \end{aligned}$$

where (a) follows from the definition of ε , (b) follows from the union bound, and (c) holds due to (39). □

A.2 BOUNDING THE ONE-SIDED REGRET WHEN ONE POLICY IS DRAWN FROM A FIXED DISTRIBUTION

In this portion of the analysis, we assume that in each iteration i , policy π_{i1} is drawn from a fixed distribution over Π , the set of deterministic policies. In this setting, we only consider the one-sided regret incurred by the policies $\{\pi_{i2}\}$, $i \geq 1$. Let N be the total number of iterations of DPS, so that the total number of actions taken by $\{\pi_{i2}\}$, $i \geq 1$, is $T = Nh$. Let \boldsymbol{x}_i^* be the trajectory obtained in iteration i if $\pi_{i2} = \pi^*$. Then, the expected one-sided regret of the policies $\{\pi_{i2}\}$ is the expected loss in total utility due to selecting suboptimal policies:

$$\mathbb{E}[\text{REG}_2(T)] = \mathbb{E}[\text{REG}_2(Nh)] := \mathbb{E} \left[\sum_{i=1}^N \bar{\boldsymbol{r}}^T(\boldsymbol{x}_i^* - \boldsymbol{x}_{i2}) \right].$$

Recall that the preference outcome on iteration i , $y_i \in \{-\frac{1}{2}, \frac{1}{2}\}$, can be written as:

$$y_i = \bar{\boldsymbol{r}}^T(\boldsymbol{x}_{i2} - \boldsymbol{x}_{i1}) + \eta_i,$$

for zero-mean noise η_i . Define the outcome of selecting policy $\pi_{i2} = \pi^*$ in iteration i as $y_i^* \in \{-\frac{1}{2}, \frac{1}{2}\}$:

$$y_i^* = \bar{\boldsymbol{r}}^T(\boldsymbol{x}_i^* - \boldsymbol{x}_{i1}) + \eta_i^*,$$

for zero-mean noise η_i^* . Importantly, the difference $y_i^* - y_i$ is equal to the instantaneous one-sided regret in expectation:

$$\begin{aligned} \mathbb{E}[\text{REG}_2(T)] &:= \mathbb{E} \left[\sum_{i=1}^N \bar{\boldsymbol{r}}^T(\boldsymbol{x}_i^* - \boldsymbol{x}_{i2}) \right] = \mathbb{E} \left[\sum_{i=1}^N (\bar{\boldsymbol{r}}^T(\boldsymbol{x}_i^* - \boldsymbol{x}_{i1}) - \bar{\boldsymbol{r}}^T(\boldsymbol{x}_{i2} - \boldsymbol{x}_{i1})) \right] \\ &\stackrel{(a)}{=} \mathbb{E} \left[\sum_{i=1}^N (y_i^* - y_i) \right], \end{aligned} \tag{40}$$

where equality (a) holds because the noise terms η_i and η_i^* are zero-mean.

Define the one-sided algorithm's history at iteration i as $\mathcal{H}_i^{(2)} = \{Z_1, Z_2, \dots, Z_i\}$, where $Z_i = (\pi_{i2}, \tau_{i1}, \tau_{i2}, \boldsymbol{x}_{i2} - \boldsymbol{x}_{i1}, y_i)$. Analogously to Russo and Van Roy (2016), we establish notation for probabilities and information-theoretic quantities while conditioning on the history $\mathcal{H}_{i-1}^{(2)}$. In particular, $P_i(\cdot) := P(\cdot | \mathcal{H}_{i-1}^{(2)})$ and $\mathbb{E}_i[\cdot] := \mathbb{E}[\cdot | \mathcal{H}_{i-1}^{(2)}]$. With respect to the history, the entropy of a random variable X is $H_i(X) := -\sum_x P_i(X=x) \log P_i(X=x)$, while two random variables X and Y have mutual information $I_i(X; Y) := H_i(X) - H_i(X|Y)$. Lastly, $D(P||Q)$ is the Kullback-Leibler divergence between two discrete probability distributions P and Q .

The *information ratio* is then defined as:

$$\Gamma_i := \frac{\mathbb{E}_i[y_i^* - y_i]^2}{I_i(\pi^*; (\pi_{i2}, \tau_{i1}, \tau_{i2}, \mathbf{x}_{i2} - \mathbf{x}_{i1}, y_i))}. \quad (41)$$

This definition is analogous to the information ratio defined in Russo and Van Roy (2016), but while Russo and Van Roy (2016) study the bandit setting with absolute feedback, our definition is adapted to the preference-based RL setting. Note that the numerator is the square of the expected instantaneous one-sided regret, while the denominator is the information gained about the optimal policy in iteration i .

In Russo and Van Roy (2016), the authors express the expected Bayesian regret of Thompson sampling in the bandit setting in terms of the information ratio (Proposition 1, Russo and Van Roy (2016)). The following lemma adapts this result to the PBRL setting:

Lemma 12. *If there exists an upper bound $\bar{\Gamma}$ such that $\Gamma_i \leq \bar{\Gamma}$ almost surely for each $i \in \{1, \dots, N\}$, where N is the number of DPS iterations (i.e., pairs of trajectories), then the policies $\{\pi_{i2}\}$ take $T = Nh$ total actions, and:*

$$\mathbb{E}[\text{REG}_2(T)] = \mathbb{E}[\text{REG}_2(Nh)] \leq \sqrt{\bar{\Gamma} H(\pi^*) N},$$

where $H(\pi^*)$ is the entropy of the optimal policy π^* . Note that since there are at most A^{Sh} possible deterministic policies, $H(\pi^*) \leq \log |A^{Sh}| = Sh \log A$. Substituting this,

$$\mathbb{E}[\text{REG}_2(T)] \leq \sqrt{\bar{\Gamma} Sh N \log A} = \sqrt{\bar{\Gamma} ST \log A}.$$

Proof.

$$\begin{aligned} \mathbb{E}[\text{REG}_2(T)] &= \sum_{i=1}^N \mathbb{E}[y_i^* - y_i] \stackrel{(a)}{=} \sum_{i=1}^N \mathbb{E}_{\mathcal{H}_{i-1}^{(2)}} \mathbb{E}_i[y_i^* - y_i] = \sum_{i=1}^N \mathbb{E}_{\mathcal{H}_{i-1}^{(2)}} \sqrt{\Gamma_i I_i(\pi^*; (\pi_{i2}, \tau_{i1}, \tau_{i2}, \mathbf{x}_{i2} - \mathbf{x}_{i1}, y_i))} \\ &\leq \sqrt{\bar{\Gamma}} \sum_{i=1}^N \mathbb{E}_{\mathcal{H}_{i-1}^{(2)}} \sqrt{I_i(\pi^*; (\pi_{i2}, \tau_{i1}, \tau_{i2}, \mathbf{x}_{i2} - \mathbf{x}_{i1}, y_i))} \\ &\stackrel{(b)}{\leq} \sqrt{\bar{\Gamma} N \sum_{i=1}^N \left[\mathbb{E}_{\mathcal{H}_{i-1}^{(2)}} \sqrt{I_i(\pi^*; (\pi_{i2}, \tau_{i1}, \tau_{i2}, \mathbf{x}_{i2} - \mathbf{x}_{i1}, y_i))} \right]^2} \\ &\stackrel{(c)}{\leq} \sqrt{\bar{\Gamma} N \sum_{i=1}^N \mathbb{E}_{\mathcal{H}_{i-1}^{(2)}} I_i(\pi^*; (\pi_{i2}, \tau_{i1}, \tau_{i2}, \mathbf{x}_{i2} - \mathbf{x}_{i1}, y_i))}, \end{aligned} \quad (42)$$

where (a) results from the tower property of conditional expectation, (b) follows from the Cauchy-Schwarz inequality, and (c) follows from Jensen's inequality. It remains to upper-bound the summation in (42) by $H(\pi^*)$. Defining $Z_i := (\pi_{i2}, \tau_{i1}, \tau_{i2}, \mathbf{x}_{i2} - \mathbf{x}_{i1}, y_i)$, the summation terms are equal to $\mathbb{E}_{\mathcal{H}_{i-1}^{(2)}} I_i(\pi^*; Z_i) = I(\pi^*; Z_i | Z_1, \dots, Z_{i-1})$, where the last equality comes from applying the definitions of I_i and of conditional mutual information. Therefore,

$$\begin{aligned} \sum_{i=1}^N \mathbb{E}_{\mathcal{H}_{i-1}^{(2)}} I_i(\pi^*; (\pi_{i2}, \tau_{i1}, \tau_{i2}, \mathbf{x}_{i2} - \mathbf{x}_{i1}, y_i)) &= \sum_{i=1}^N I(\pi^*; Z_i | Z_1, \dots, Z_{i-1}) \stackrel{(a)}{=} I(\pi^*; Z_1, \dots, Z_i) \\ &\stackrel{(b)}{=} H(\pi^*) - H(\pi^* | Z_1, \dots, Z_i) \stackrel{(c)}{\leq} H(\pi^*), \end{aligned}$$

where (a) results from the chain rule for mutual information, (b) is a standard identity resulting from the definitions of entropy and mutual information (Cover and Thomas, 2012), and (c) follows from the non-negativity of entropy. \square

Corollary 1. *If $\Gamma_i \leq \bar{\Gamma}$ almost surely for each $i \in \{i_0, \dots, N\}$ for some index $i_0 \geq 1$, then DPS achieves an asymptotic one-sided regret rate of $\sqrt{\bar{\Gamma} Sh N \log A} = \sqrt{\bar{\Gamma} ST \log A}$.*

Proof. The one-sided regret can be decomposed into two sums:

$$\mathbb{E}[\text{REG}_2(T)] = \sum_{i=1}^N \mathbb{E}[y_i^* - y_i] = \sum_{i=1}^{i_0-1} \mathbb{E}[y_i^* - y_i] + \sum_{i=i_0}^N \mathbb{E}[y_i^* - y_i].$$

The second term can be upper-bounded via the same arguments used to prove Lemma 12. The first term is a constant for iteration $i \geq i_0$, as it does not depend on the performance of the algorithm for iterations $i \geq i_0$. \square

Next, we turn to upper-bounding the information ratio. First, Lemma 13 demonstrates in Thompson sampling, the probability of selecting a policy is equal to that policy's posterior probability of being optimal.

Lemma 13. *When π_{i_2} is selected via Thompson sampling—i.e., by drawing respective samples $\tilde{\mathbf{p}}$ and $\tilde{\mathbf{r}}$ from the dynamics and reward posteriors, and performing value iteration to obtain $\pi_{i_2} = \pi_{v_i}(\tilde{\mathbf{p}}, \tilde{\mathbf{r}})$ —then for each deterministic policy π , $P_i(\pi_{i_2} = \pi) = P_i(\pi^* = \pi)$.*

Proof. In Thompson sampling, the MDP parameters are sampled from the model posterior, that is, according to their posterior probability of being the true MDP parameters. Let $\bar{\mathbf{m}} = [\bar{\mathbf{r}}^T, \bar{\mathbf{p}}^T]^T \in \mathbb{R}^{S^2 A + SA}$ be the vector of true MDP parameters, where $\bar{\mathbf{r}} \in \mathbb{R}^{SA}$ is the vector of reward parameters and $\bar{\mathbf{p}} \in \mathbb{R}^{S^2 A}$ is the vector of dynamics parameters. Let $p_i(\mathbf{m})$ be the posterior probability density of $\bar{\mathbf{m}}$, given the history; then, Thompson sampling samples parameter vectors $\mathbf{m} \in \mathbb{R}^{S^2 A + SA}$ according to $p_i(\mathbf{m})$. Finally, for any deterministic policy π , let $m(\pi) \subset \mathbb{R}^{S^2 A + SA}$ give the set of all MDP parameters for which value iteration yields the policy π . Then:

$$P_i(\pi_{i_2} = \pi) \stackrel{(a)}{=} \int_{m(\pi)} p_i(\mathbf{m}) d\mathbf{m} \stackrel{(b)}{=} P_i(\pi^* = \pi),$$

where (a) holds because Thompson sampling selects policies by sampling $\mathbf{m} \sim p_i(\cdot)$ and then applying value iteration to \mathbf{m} , and (b) follows from integrating over the posterior probability of all MDP parameter vectors resulting in π . \square

To upper-bound the information ratio, we next express its numerator and denominator using the optimal policy's distribution via the following lemma (analogous to Proposition 2 in Russo and Van Roy (2016)):

Lemma 14. *Recall that Π is the set of deterministic policies. The following two statements hold:*

$$\mathbb{E}_i[y_i^* - y_i] = \sum_{\pi \in \Pi} P_i(\pi^* = \pi) \{ \mathbb{E}_i[y_i | \pi^* = \pi_{i_2} = \pi] - \mathbb{E}_i[y_i | \pi_{i_2} = \pi] \}$$

and

$$I_i(\pi^*; (\pi_{i_2}, \tau_{i_1}, \tau_{i_2}, \mathbf{x}_{i_2} - \mathbf{x}_{i_1}, y_i)) \geq \sum_{\pi, \pi' \in \Pi} P_i(\pi^* = \pi) P_i(\pi^* = \pi') D(P_i(y_i | \pi_{i_2} = \pi, \pi^* = \pi') || P_i(y_i | \pi_{i_2} = \pi)),$$

where $D(p||q)$ is the Kullback-Leibler divergence between discrete probability distributions p and q .

Proof. The numerator of the information ratio can be written as:

$$\begin{aligned} \mathbb{E}_i[y_i^* - y_i] &= \sum_{\pi \in \Pi} P_i(\pi^* = \pi) \mathbb{E}_i[y_i^* | \pi^* = \pi] - \sum_{\pi \in \Pi} P_i(\pi_{i_2} = \pi) \mathbb{E}_i[y_i | \pi_{i_2} = \pi] \\ &\stackrel{(a)}{=} \sum_{\pi \in \Pi} P_i(\pi^* = \pi) (\mathbb{E}_i[y_i^* | \pi^* = \pi] - \mathbb{E}_i[y_i | \pi_{i_2} = \pi]) \\ &= \sum_{\pi \in \Pi} P_i(\pi^* = \pi) \{ \mathbb{E}_i[y_i | \pi^* = \pi_{i_2} = \pi] - \mathbb{E}_i[y_i | \pi_{i_2} = \pi] \}, \end{aligned}$$

where (a) follows from Lemma 13. The denominator of the information ratio can meanwhile be lower-bounded:

$$\begin{aligned}
I_i(\pi^*; (\pi_{i2}, \tau_{i1}, \tau_{i2}, \mathbf{x}_{i2} - \mathbf{x}_{i1}, y_i)) &\stackrel{(a)}{=} I_i(\pi^*; \pi_{i2}) + I_i(\pi^*; y_i | \pi_{i2}) + I_i(\pi^*; (\mathbf{x}_{i2} - \mathbf{x}_{i1}, \tau_{i1}, \tau_{i2}) | \pi_{i2}, y_i) \\
&\stackrel{(b)}{\geq} I_i(\pi^*; y_i | \pi_{i2}) = \sum_{\pi \in \Pi} P_i(\pi_{i2} = \pi) I_i(\pi^*; y_i | \pi_{i2} = \pi) \\
&\stackrel{(c)}{=} \sum_{\pi \in \Pi} P_i(\pi^* = \pi) I_i(\pi^*; y_i | \pi_{i2} = \pi),
\end{aligned}$$

where (a) follows from the chain rule for mutual information, (b) results from the non-negativity of mutual information, and (c) is a consequence of Lemma 13. We next apply the following information-theoretic identity (Fact 6 in Russo and Van Roy (2016)), which holds for discrete random variables X and Y :

$$I(X; Y) = \sum_x P(X = x) D(P(Y | X = x) || P(Y)).$$

From this, one obtains:

$$\begin{aligned}
I_i(\pi^*; y_i | \pi_{i2} = \pi) &= \sum_{\pi' \in \Pi} P_i(\pi^* = \pi' | \pi_{i2} = \pi) D(P_i(y_i | \pi^* = \pi', \pi_{i2} = \pi) || P_i(y_i | \pi_{i2} = \pi)) \\
&\stackrel{(a)}{=} \sum_{\pi' \in \Pi} P_i(\pi^* = \pi') D(P_i(y_i | \pi^* = \pi', \pi_{i2} = \pi) || P_i(y_i | \pi_{i2} = \pi)),
\end{aligned}$$

where (a) holds because by definition of Thompson sampling, π^* and π_{i2} are independent given the history. Therefore:

$$\begin{aligned}
I_i(\pi^*; (\pi_{i2}, \tau_{i1}, \tau_{i2}, \mathbf{x}_{i2} - \mathbf{x}_{i1}, y_i)) &\geq \sum_{\pi \in \Pi} P_i(\pi^* = \pi) \sum_{\pi' \in \Pi} P_i(\pi^* = \pi') D(P_i(y_i | \pi^* = \pi', \pi_{i2} = \pi) || P_i(y_i | \pi_{i2} = \pi)) \\
&= \sum_{\pi, \pi' \in \Pi} P_i(\pi^* = \pi) P_i(\pi^* = \pi') D(P_i(y_i | \pi_{i2} = \pi, \pi^* = \pi') || P_i(y_i | \pi_{i2} = \pi)),
\end{aligned}$$

which is the desired result. □

The next lemma asymptotically upper-bounds the information ratio Γ_i for the one-sided regret. It is inspired by the analysis in Russo and Van Roy (2016) for the linear bandit setting (Proposition 5 in Russo and Van Roy (2016)); however, extending this result to the PBRL setting requires accounting for the dynamics, which complicates the analysis.

Lemma 15. *Consider the one-sided regret when π_{i1} is drawn from a fixed distribution. The information ratio Γ_i satisfies:*

$$\lim_{i \rightarrow \infty} \Gamma_i \leq \frac{d}{2} = \frac{SA}{2}.$$

Thus, for any $\varepsilon > 0$ and a sufficiently-large iteration index i , $\Gamma_i \leq \frac{SA}{2} + \varepsilon$.

Proof. Let $K := |\Pi|$. Index all the deterministic policies as π_1, \dots, π_K , and define $B^{(i)} \in \mathbb{R}^{K \times K}$ with j^{th} element:

$$B_{jk}^{(i)} = \sqrt{P_i(\pi^* = \pi_j) P_i(\pi^* = \pi_k)} \left(\mathbb{E}_i[y_i | \pi_{i2} = \pi_j, \pi^* = \pi_k] - \mathbb{E}_i[y_i | \pi_{i2} = \pi_j] \right). \quad (43)$$

The numerator and denominator of Γ_i can both be expressed in terms of $B^{(i)}$. Applying Lemma 14 to the numerator, the instantaneous regret can be written as follows:

$$\begin{aligned}
\mathbb{E}_i[y_i^* - y_i] &= \sum_{\pi \in \Pi} P_i(\pi^* = \pi) \{ \mathbb{E}_i[y_i | \pi^* = \pi_{i2} = \pi] - \mathbb{E}_i[y_i | \pi_{i2} = \pi] \} \\
&= \sum_{j=1}^K P_i(\pi^* = \pi_j) \{ \mathbb{E}_i[y_i | \pi^* = \pi_{i2} = \pi_j] - \mathbb{E}_i[y_i | \pi_{i2} = \pi_j] \} \\
&= \sum_{j=1}^K B_{jj}^{(i)} = \text{Tr}(B^{(i)}).
\end{aligned}$$

Applying Lemma 14 to the denominator of Γ_i in Equation (41) yields:

$$\begin{aligned} I_i(\pi^*; (\pi_{i2}, \tau_{i1}, \tau_{i2}, \mathbf{x}_{i2} - \mathbf{x}_{i1}, y_i)) &\geq \sum_{\pi, \pi' \in \Pi} P_i(\pi^* = \pi) P_i(\pi^* = \pi') D(P_i(y_i | \pi_{i2} = \pi, \pi^* = \pi') || P_i(y_i | \pi_{i2} = \pi)) \\ &= \sum_{j, k=1}^K P_i(\pi^* = \pi_j) P_i(\pi^* = \pi_k) D(P_i(y_i | \pi_{i2} = \pi_j, \pi^* = \pi_k) || P_i(y_i | \pi_{i2} = \pi_j)). \end{aligned}$$

We convert the Kullback-Leibler divergence to a difference of expectations by applying Fact 9 from Russo and Van Roy (2016), restated here: for any probability distributions P and Q such that P is absolutely continuous with respect to Q , any random variable X taking values on the set \mathcal{X} , and any $g : \mathcal{X} \rightarrow \mathbb{R}$ such that $\sup g - \inf g \leq 1$,

$$D(P || Q) \geq 2(\mathbb{E}_P[g(X)] - \mathbb{E}_Q[g(X)])^2. \quad (44)$$

Thus, we have that:

$$D(P_i(y_i | \pi_{i2} = \pi_j, \pi^* = \pi_k) || P_i(y_i | \pi_{i2} = \pi_j)) \geq 2\left(\mathbb{E}_i[y_i | \pi_{i2} = \pi_j, \pi^* = \pi_k] - \mathbb{E}_i[y_i | \pi_{i2} = \pi_j]\right)^2,$$

where we applied (44) with $g(x) = x$; this definition of g satisfies the requirement that $\sup g - \inf g \leq 1$, since its argument is $y_i \in \{-\frac{1}{2}, \frac{1}{2}\}$. As a result:

$$\begin{aligned} I_i(\pi^*; (\pi_{i2}, \tau_{i1}, \tau_{i2}, \mathbf{x}_{i2} - \mathbf{x}_{i1}, y_i)) &\geq 2 \sum_{j, k=1}^K P_i(\pi^* = \pi_j) P_i(\pi^* = \pi_k) \left(\mathbb{E}_i[y_i | \pi_{i2} = \pi_j, \pi^* = \pi_k] - \mathbb{E}_i[y_i | \pi_{i2} = \pi_j]\right)^2 \\ &= 2 \sum_{j, k=1}^K \left(B_{jk}^{(i)}\right)^2 = 2\|B^{(i)}\|_F^2. \end{aligned}$$

Combining these results gives that $\Gamma_i \leq \frac{\text{Tr}(B^{(i)})^2}{2\|B^{(i)}\|_F^2}$. As shown in Russo and Van Roy (2016) (Fact 10), for any square matrix $B \in \mathbb{R}^{m \times m}$, $\text{Tr}(B) \leq \sqrt{\text{Rank}(B)}\|B\|_F$. Thus:

$$\Gamma_i \leq \frac{\text{Tr}(B^{(i)})^2}{2\|B^{(i)}\|_F^2} \leq \frac{1}{2} \text{Rank}(B^{(i)}).$$

The problem is therefore reduced to upper-bounding $\text{Rank}(B^{(i)})$. First, we will show that under known transition dynamics, $\text{Rank}(B^{(i)}) \leq d = SA$. Subsequently, we will demonstrate that as the sampled dynamics parameters converge in distribution to their true values (which occurs by Proposition 1), $\lim_{i \rightarrow \infty} \text{Rank}(B^{(i)}) \leq d$.

Recall the definition of $B^{(i)}$ in (43). We will first show that under known transition dynamics, we can define a set of vectors $\mathbf{u}_1, \dots, \mathbf{u}_K, \mathbf{v}_1, \dots, \mathbf{v}_K \in \mathbb{R}^d$, such that $B_{jk}^{(i)} = \mathbf{u}_k^T \mathbf{v}_j$. Recalling that $y_i = \bar{\mathbf{r}}^T(\mathbf{x}_{i2} - \mathbf{x}_{i1}) + \eta_i$, and assuming that the dynamics $\bar{\mathbf{p}}$ are known,

$$\begin{aligned} \mathbb{E}_i[y_i | \pi_{i2} = \pi_j, \pi^* = \pi_k, \bar{\mathbf{p}}] &= \mathbb{E}_i[\bar{\mathbf{r}}^T(\mathbf{x}_{i2} - \mathbf{x}_{i1}) + \eta_i | \pi_{i2} = \pi_j, \pi^* = \pi_k, \bar{\mathbf{p}}] \\ &\stackrel{(a)}{=} \mathbb{E}_i[\bar{\mathbf{r}}^T(\mathbf{x}_{i2} - \mathbf{x}_{i1}) | \pi_{i2} = \pi_j, \pi^* = \pi_k, \bar{\mathbf{p}}] \\ &\stackrel{(b)}{=} \mathbb{E}_i[\bar{\mathbf{r}} | \pi_{i2} = \pi_j, \pi^* = \pi_k, \bar{\mathbf{p}}]^T (\mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] - \mathbb{E}_i[\mathbf{x}_{i1}]) \\ &\stackrel{(c)}{=} \mathbb{E}_i[\bar{\mathbf{r}} | \pi^* = \pi_k, \bar{\mathbf{p}}]^T (\mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] - \mathbb{E}_i[\mathbf{x}_{i1}]), \end{aligned} \quad (45)$$

where equality (a) holds because the noise η_i is zero-mean. Equality (b) holds because, by assumption, \mathbf{x}_{i1} is drawn from a fixed distribution, independently of π^* , π_{i2} , or $\bar{\mathbf{r}}$, while given π_{i2} and known dynamics, the distribution of \mathbf{x}_{i2}

is fully-determined and independent of π^* and $\bar{\mathbf{r}}$. Equality (c) holds because, conditioned upon $\mathcal{H}_{i-1}^{(2)}$, $\bar{\mathbf{r}}$ and π_{i2} are independent (as the history fully determines the distribution of π_{i2}). By similar arguments,

$$\begin{aligned}\mathbb{E}_i[y_i | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] &= \mathbb{E}_i[\bar{\mathbf{r}}^T(\mathbf{x}_{i2} - \mathbf{x}_{i1}) + \eta_i | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] = \mathbb{E}_i[\bar{\mathbf{r}}^T(\mathbf{x}_{i2} - \mathbf{x}_{i1}) | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] \\ &= \mathbb{E}_i[\bar{\mathbf{r}} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}]^T (\mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] - \mathbb{E}_i[\mathbf{x}_{i1}]) \\ &= \mathbb{E}_i[\bar{\mathbf{r}} | \bar{\mathbf{p}}]^T (\mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] - \mathbb{E}_i[\mathbf{x}_{i1}]).\end{aligned}\quad (46)$$

Subtracting the latter two quantities yields:

$$\begin{aligned}\mathbb{E}_i[y_i | \pi_{i2} = \pi_j, \pi^* = \pi_k, \bar{\mathbf{p}}] - \mathbb{E}_i[y_i | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] \\ = (\mathbb{E}_i[\bar{\mathbf{r}} | \pi^* = \pi_k, \bar{\mathbf{p}}] - \mathbb{E}_i[\bar{\mathbf{r}} | \bar{\mathbf{p}}])^T (\mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] - \mathbb{E}_i[\mathbf{x}_{i1}]).\end{aligned}\quad (47)$$

Applying this result, under the case of known dynamics, $B_{jk}^{(i)}$ can be written as:

$$\begin{aligned}B_{jk}^{(i)} &= \sqrt{P_i(\pi^* = \pi_j | \bar{\mathbf{p}})P_i(\pi^* = \pi_k | \bar{\mathbf{p}})} \left(\mathbb{E}_i[y_i | \pi_{i2} = \pi_j, \pi^* = \pi_k, \bar{\mathbf{p}}] - \mathbb{E}_i[y_i | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] \right) \\ &= \sqrt{P_i(\pi^* = \pi_k | \bar{\mathbf{p}})} (\mathbb{E}_i[\bar{\mathbf{r}} | \pi^* = \pi_k, \bar{\mathbf{p}}] - \mathbb{E}_i[\bar{\mathbf{r}} | \bar{\mathbf{p}}])^T \sqrt{P_i(\pi^* = \pi_j | \bar{\mathbf{p}})} (\mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] - \mathbb{E}_i[\mathbf{x}_{i1}]) \\ &:= \mathbf{u}_k^T \mathbf{v}_j,\end{aligned}$$

where in the last equality, we define $\mathbf{u}_k := \sqrt{P_i(\pi^* = \pi_k | \bar{\mathbf{p}})} (\mathbb{E}_i[\bar{\mathbf{r}} | \pi^* = \pi_k, \bar{\mathbf{p}}] - \mathbb{E}_i[\bar{\mathbf{r}} | \bar{\mathbf{p}}])$ and $\mathbf{v}_j := \sqrt{P_i(\pi^* = \pi_j | \bar{\mathbf{p}})} (\mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] - \mathbb{E}_i[\mathbf{x}_{i1}])$. Therefore:

$$B^{(i)} = \begin{bmatrix} \mathbf{u}_1^T \mathbf{v}_1 & \dots & \mathbf{u}_K^T \mathbf{v}_1 \\ \vdots & \ddots & \vdots \\ \mathbf{u}_1^T \mathbf{v}_K & \dots & \mathbf{u}_K^T \mathbf{v}_K \end{bmatrix} = \begin{bmatrix} \mathbf{v}_1^T \\ \vdots \\ \mathbf{v}_K^T \end{bmatrix} [\mathbf{u}_1 \quad \dots \quad \mathbf{u}_K]. \quad (48)$$

Because $B^{(i)}$ can be written as the product of a $K \times d$ matrix and a $d \times K$ matrix, $B^{(i)}$ can have rank at most d . To reach this result, however, we assumed that the MDP transition dynamics are known. To complete the proof, we now show that the result still holds asymptotically as the sampled dynamics converge in distribution to the true dynamics (Proposition 1 guarantees that this convergence occurs). Note that we only used our assumed knowledge of the dynamics to arrive at the equalities in lines (45) and (46). In both cases, knowledge of $\bar{\mathbf{p}}$ is used to treat \mathbf{x}_{i2} and $\bar{\mathbf{r}}$ as conditionally independent given π_{i2} . In particular, (45) uses that:

$$\mathbb{E}_i[\bar{\mathbf{r}}^T \mathbf{x}_{i2} | \pi_{i2} = \pi_j, \pi^* = \pi_k, \bar{\mathbf{p}}] = \mathbb{E}_i[\bar{\mathbf{r}} | \pi_{i2} = \pi_j, \pi^* = \pi_k, \bar{\mathbf{p}}]^T \mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}].$$

We show that in general, when the transition dynamics $\bar{\mathbf{p}}$ are not known but learned, that:

$$\mathbb{E}_i[\bar{\mathbf{r}}^T \mathbf{x}_{i2} | \pi_{i2} = \pi_j, \pi^* = \pi_k] \xrightarrow{i \rightarrow \infty} \mathbb{E}_i[\bar{\mathbf{r}} | \pi_{i2} = \pi_j, \pi^* = \pi_k]^T \mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}]. \quad (49)$$

Let $p_i(\cdot)$ be the posterior probability density of the transition dynamics parameters; then, $p_i(\cdot)$ is also the density with which the dynamics $\tilde{\mathbf{p}}$ are sampled at iteration i . Because each dynamics parameter converges in distribution to its true value—and there are finitely-many dynamics parameters—the distribution $p_i(\cdot)$ converges uniformly to $\delta(\tilde{\mathbf{p}} = \bar{\mathbf{p}})$, where $\delta(\cdot)$ denotes the Dirac-delta distribution. Therefore,

$$\begin{aligned}P_i(\mathbf{x}_{i2} = \cdot | \pi_{i2} = \pi_j, \pi^* = \pi_k, \bar{\mathbf{r}}) &\stackrel{(a)}{=} \int_{\tilde{\mathbf{p}}} P_i(\mathbf{x}_{i2} = \cdot | \pi_{i2} = \pi_j, \pi^* = \pi_k, \bar{\mathbf{r}}, \bar{\mathbf{p}} = \tilde{\mathbf{p}}) p_i(\tilde{\mathbf{p}}) d\tilde{\mathbf{p}} \\ &\xrightarrow{i \rightarrow \infty} \int_{\tilde{\mathbf{p}}} P_i(\mathbf{x}_{i2} = \cdot | \pi_{i2} = \pi_j, \pi^* = \pi_k, \bar{\mathbf{r}}, \bar{\mathbf{p}} = \tilde{\mathbf{p}}) \delta(\tilde{\mathbf{p}} = \bar{\mathbf{p}}) d\tilde{\mathbf{p}} \\ &\stackrel{(b)}{=} P_i(\mathbf{x}_{i2} = \cdot | \pi_{i2} = \pi_j, \pi^* = \pi_k, \bar{\mathbf{r}}, \bar{\mathbf{p}}) \stackrel{(c)}{=} P_i(\mathbf{x}_{i2} = \cdot | \pi_{i2} = \pi_j, \bar{\mathbf{p}}),\end{aligned}\quad (50)$$

where (a) integrates over the posterior probability density of each possible dynamics parameter vector $\tilde{\mathbf{p}}$, (b) utilizes the sifting property of the Dirac-delta function, and (c) follows because the distribution of \mathbf{x}_{i2} is fully determined

given the dynamics and policy π_{i2} . For any discrete random variables X and $X_n, n \geq 1$, defined over a finite set \mathcal{X} , convergence in distribution $X_n \xrightarrow{D} X$ is equivalent to $\lim_{n \rightarrow \infty} P(X_n = x) = P(X = x)$ for each $x \in \mathcal{X}$. Since \mathcal{X} is a finite set, one also has convergence in expectation:

$$\mathbb{E}[X_n] = \sum_{x \in \mathcal{X}} xP(X_n = x) \xrightarrow{n \rightarrow \infty} \sum_{x \in \mathcal{X}} xP(X = x) = \mathbb{E}[X]. \quad (51)$$

Combining (50) and (51) yields the result,

$$\mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \pi^* = \pi_k, \bar{\mathbf{r}}] \xrightarrow{i \rightarrow \infty} \mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}]. \quad (52)$$

In consequence,

$$\begin{aligned} \mathbb{E}_i[\bar{\mathbf{r}}^T \mathbf{x}_{i2} | \pi_{i2} = \pi_j, \pi^* = \pi_k] &\stackrel{(a)}{=} \mathbb{E}_i[\bar{\mathbf{r}}^T \mathbb{E}_i[\bar{\mathbf{r}}^T \mathbf{x}_{i2} | \pi_{i2} = \pi_j, \pi^* = \pi_k, \bar{\mathbf{r}}] | \pi_{i2} = \pi_j, \pi^* = \pi_k] \\ &= \mathbb{E}_i[\bar{\mathbf{r}}^T \mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \pi^* = \pi_k, \bar{\mathbf{r}}] | \pi_{i2} = \pi_j, \pi^* = \pi_k] \\ &\xrightarrow{(b)} \mathbb{E}_i[\bar{\mathbf{r}} | \pi_{i2} = \pi_j, \pi^* = \pi_k]^T \mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] \text{ as } i \rightarrow \infty, \end{aligned} \quad (53)$$

where (a) follows from the tower property of expectation, and (b) is an application of (52). This proves the desired statement, (49). Repeating the same analysis as in (49)-(53), but removing the conditioning on $\pi^* = \pi_k$ yields:

$$\mathbb{E}_i[\bar{\mathbf{r}}^T \mathbf{x}_{i2} | \pi_{i2} = \pi_j] \rightarrow \mathbb{E}_i[\bar{\mathbf{r}} | \pi_{i2} = \pi_j]^T \mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] \text{ as } i \rightarrow \infty. \quad (54)$$

The analysis in (45)-(47) can be repeated, replacing the knowledge of $\bar{\mathbf{p}}$ with the asymptotic relations in (53) and (54):

$$\begin{aligned} \mathbb{E}_i[y_i | \pi_{i2} = \pi_j, \pi^* = \pi_k] &= \mathbb{E}_i[\bar{\mathbf{r}}^T (\mathbf{x}_{i2} - \mathbf{x}_{i1}) + \eta_i | \pi_{i2} = \pi_j, \pi^* = \pi_k] \\ &= \mathbb{E}_i[\bar{\mathbf{r}}^T (\mathbf{x}_{i2} - \mathbf{x}_{i1}) | \pi_{i2} = \pi_j, \pi^* = \pi_k] \\ &\xrightarrow{i \rightarrow \infty} \mathbb{E}_i[\bar{\mathbf{r}} | \pi_{i2} = \pi_j, \pi^* = \pi_k]^T (\mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] - \mathbb{E}_i[\mathbf{x}_{i1}]) \\ &= \mathbb{E}_i[\bar{\mathbf{r}} | \pi^* = \pi_k]^T (\mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] - \mathbb{E}_i[\mathbf{x}_{i1}]), \text{ and similarly,} \\ \mathbb{E}_i[y_i | \pi_{i2} = \pi_j] &= \mathbb{E}_i[\bar{\mathbf{r}}^T (\mathbf{x}_{i2} - \mathbf{x}_{i1}) + \eta_i | \pi_{i2} = \pi_j] = \mathbb{E}_i[\bar{\mathbf{r}}^T (\mathbf{x}_{i2} - \mathbf{x}_{i1}) | \pi_{i2} = \pi_j] \\ &\xrightarrow{i \rightarrow \infty} \mathbb{E}_i[\bar{\mathbf{r}} | \pi_{i2} = \pi_j]^T (\mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] - \mathbb{E}_i[\mathbf{x}_{i1}]) \\ &= \mathbb{E}_i[\bar{\mathbf{r}}]^T (\mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] - \mathbb{E}_i[\mathbf{x}_{i1}]). \end{aligned}$$

Then, the difference between these two quantities asymptotically becomes:

$$\mathbb{E}_i[y_i | \pi_{i2} = \pi_j, \pi^* = \pi_k] - \mathbb{E}_i[y_i | \pi_{i2} = \pi_j] \xrightarrow{i \rightarrow \infty} (\mathbb{E}_i[\bar{\mathbf{r}} | \pi^* = \pi_k] - \mathbb{E}_i[\bar{\mathbf{r}}])^T (\mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] - \mathbb{E}_i[\mathbf{x}_{i1}]).$$

Asymptotically, $B_{jk}^{(i)}$ can then be expressed as:

$$\begin{aligned} B_{jk}^{(i)} &= \sqrt{P_i(\pi^* = \pi_j)P_i(\pi^* = \pi_k)} \left(\mathbb{E}_i[y_i | \pi_{i2} = \pi_j, \pi^* = \pi_k] - \mathbb{E}_i[y_i | \pi_{i2} = \pi_j] \right) \\ &\xrightarrow{i \rightarrow \infty} \sqrt{P_i(\pi^* = \pi_k)} (\mathbb{E}_i[\bar{\mathbf{r}} | \pi^* = \pi_k] - \mathbb{E}_i[\bar{\mathbf{r}}])^T \sqrt{P_i(\pi^* = \pi_j)} (\mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] - \mathbb{E}_i[\mathbf{x}_{i1}]) \\ &:= \mathbf{u}_k'^T \mathbf{v}_j', \end{aligned}$$

where in the last equality, we define $\mathbf{u}_k' := \sqrt{P_i(\pi^* = \pi_k)} (\mathbb{E}_i[\bar{\mathbf{r}} | \pi^* = \pi_k] - \mathbb{E}_i[\bar{\mathbf{r}}])$ and $\mathbf{v}_j' := \sqrt{P_i(\pi^* = \pi_j)} (\mathbb{E}_i[\mathbf{x}_{i2} | \pi_{i2} = \pi_j, \bar{\mathbf{p}}] - \mathbb{E}_i[\mathbf{x}_{i1}])$.

Thus, one can write $B^{(i)} = B_a^{(i)} + B_b^{(i)}$, where $B_a^{(i)}$ has jk^{th} element $\mathbf{u}_k'^T \mathbf{v}_j'$. As in (48), $B_a^{(i)}$ can be written as a product of a $K \times d$ matrix with a $d \times K$ matrix, and so its rank is at most d . Meanwhile, the elements of $B_b^{(i)}$ decay to zero as $i \rightarrow \infty$. For any $\varepsilon > 0$ and sufficiently-high i , the information ratio becomes upper-bounded by:

$$\Gamma_i \leq \frac{\text{Tr}(B^{(i)})^2}{2\|B^{(i)}\|_F^2} = \frac{\text{Tr}(B_a^{(i)} + B_b^{(i)})^2}{2\|B_a^{(i)} + B_b^{(i)}\|_F^2} \stackrel{(a)}{\leq} \frac{\text{Tr}(B_a^{(i)})^2}{2\|B_a^{(i)}\|_F^2} + \varepsilon \stackrel{(b)}{\leq} \frac{1}{2} \text{Rank}(B_a^{(i)}) + \varepsilon \stackrel{(c)}{\leq} \frac{d}{2} + \varepsilon = \frac{SA}{2} + \varepsilon,$$

where (a) follows because $\text{Tr}(B)$ and $\|B\|_F$ are both continuous in the elements of B , and the elements of $B_b^{(i)}$ approach zero; (b) follows from Fact 10 in Russo and Van Roy (2016), as described earlier for the case with known dynamics; and (c) holds by definition of $B_a^{(i)}$. □

Combining Corollary 1 and Lemma 15 yields the asymptotic regret rate for a fixed π_{i1} -distribution:

Theorem 2. *If the policy π_{i1} is drawn from a fixed distribution for each iteration i , then for the competing policy π_{i2} , DPS achieves a one-sided asymptotic Bayesian regret rate of:*

$$S\sqrt{\frac{AT \log A}{2}}.$$

Proof. This result is a direct consequence of Corollary 1 and Lemma 15; the asymptotic bound $\bar{\Gamma} \leq \frac{SA}{2} + \varepsilon$ for any ε is substituted into the expression in Corollary 1. □

A.3 BOUNDING THE ONE-SIDED REGRET WHEN ONE POLICY IS DRAWN FROM A DRIFTING AND CONVERGING DISTRIBUTION

We now assume that the distribution of π_{i1} is no longer fixed, but rather, that the sampled policies π_{i1} converge in distribution toward some fixed probability distribution over Π , the set of deterministic policies. We will asymptotically bound the one-sided regret incurred by π_{i2} in this case. To do so, we will leverage that when two discrete random variables converge in distribution, their mutual information also converges:

Lemma 16. *Let X_n and Y_n , $n \in \mathbb{N}$, be two sequences of discrete random variables that take values on the finite sets \mathcal{X} and \mathcal{Y} , respectively. If $X_n \xrightarrow{D} X$ and $Y_n \xrightarrow{D} Y$, then $\lim_{n \rightarrow \infty} I(X_n; Y_n) = I(X; Y)$.*

Proof. Firstly, note that $X_n \xrightarrow{D} X$ and $Y_n \xrightarrow{D} Y$ imply that jointly, $(X_n, Y_n) \xrightarrow{D} (X, Y)$. Let $P_n(x)$ and $P(x)$ be the probability distributions of X_n and X , respectively. Because the variables are discrete, $X_n \xrightarrow{D} X$ implies that $P_n(x) \rightarrow P(x)$ for each $x \in \mathcal{X}$, and similarly for Y and (X, Y) .

We express the mutual information as a sum of entropies: $I(X; Y) = H(X) + H(Y) - H(X, Y)$ (Cover and Thomas, 2012). It suffices to show that if $X_n \xrightarrow{D} X$, then $\lim_{n \rightarrow \infty} H(X_n) = H(X)$: because Y and (X, Y) are also discrete random variables, this would imply that similarly, $\lim_{n \rightarrow \infty} H(Y_n) = H(Y)$ and $\lim_{n \rightarrow \infty} H(X_n, Y_n) = H(X, Y)$.

By definition, the entropies of X and X_n are:

$$H(X) = - \sum_{x \in \mathcal{X}} P(x) \log P(x) \text{ and}$$

$$H(X_n) = - \sum_{x \in \mathcal{X}} P_n(x) \log P_n(x).$$

Because $P_n(x) \rightarrow P(x)$ for each $x \in \mathcal{X}$, for any $x \in \mathcal{X}$ and $\delta > 0$, there exists $N_x \in \mathbb{N}$ such that for all $n \geq N_x$, $|P_n(x) - P(x)| < \delta$. Let $N = \max_{x \in \mathcal{X}} N_x$. Then, for all $n > N$ and all $x \in \mathcal{X}$, $|P_n(x) - P(x)| < \delta$.

Choose any $\varepsilon' > 0$. Since $f(z) = z \log z$ is continuous for $z \geq 0$ (with $0 \log 0 := 0$), there exists $\delta > 0$ such that if $|P_n(x) - P(x)| < \delta$, then $|P_n(x) \log P_n(x) - P(x) \log P(x)| < \varepsilon'$. We choose a δ that satisfies this condition.

Then, there exists N such that for all $n > N$ and for all $x \in \mathcal{X}$, $|P_n(x) \log P_n(x) - P(x) \log P(x)| < \varepsilon'$. Finally, choose $\varepsilon > 0$ and set $\varepsilon' \leq \frac{\varepsilon}{|\mathcal{X}|}$, so that for all $n > N$:

$$\begin{aligned} \left| H(X) - H(X_n) \right| &= \left| \sum_{x \in \mathcal{X}} P(x) \log P(x) - \sum_{x \in \mathcal{X}} P_n(x) \log P_n(x) \right| \leq \sum_{x \in \mathcal{X}} \left| P(x) \log P(x) - P_n(x) \log P_n(x) \right| \\ &\leq \sum_{x \in \mathcal{X}} \varepsilon' = \varepsilon' |\mathcal{X}| \leq \frac{\varepsilon}{|\mathcal{X}|} |\mathcal{X}| = \varepsilon. \end{aligned}$$

So for any $\varepsilon > 0$, there exists N such that for all $n > N$, $|H(X) - H(X_n)| \leq \varepsilon$. This proves that $\lim_{n \rightarrow \infty} H(X_n) = H(X)$, and therefore that $\lim_{n \rightarrow \infty} I(X_n; Y_n) = I(X; Y)$. \square

Armed with this convergence of mutual information, the one-sided regret for π_{i2} can be bounded as follows:

Lemma 17. *Assume that π_{i1} is drawn from a distribution that is drifting and converging, that is, π_{i1} converges in distribution to some fixed probability distribution. From Lemma 10, if π_{i1} is drawn from a fixed distribution, then asymptotically, its information ratio is bounded by $\bar{\Gamma}_{\pi_{i1} \text{ fixed}} \leq \frac{SA}{2}$. In the case of a drifting/converging π_{i1} distribution, the information ratio Γ_i for π_{i2} 's one-sided regret satisfies $\lim_{i \rightarrow \infty} \Gamma_i \leq \bar{\Gamma}_{\pi_{i1} \text{ fixed}} \leq \frac{SA}{2}$.*

Proof. By Lemma 15, under a fixed π_{i1} distribution, the information ratio corresponding to the one-sided regret for π_{i2} is asymptotically upper-bounded: $\lim_{i \rightarrow \infty} \Gamma_{i, \pi_{i1} \text{ fixed}} \leq \bar{\Gamma}_{\pi_{i1} \text{ fixed}}$, where $\Gamma_{i, \pi_{i1} \text{ fixed}}$ is the information ratio at iteration i when the distribution of π_{i1} is fixed.

The denominator of the information ratio, $I_i(\pi^*; (\pi_{i2}, \tau_{i1}, \tau_{i2}, \mathbf{x}_{i2} - \mathbf{x}_{i1}, y_i))$, is a mutual information between discrete random variables, as there are finitely-many possible policies, trajectories, and preference outcomes; therefore, by Lemma 16, it converges to the values that it would have under the fixed distribution to which π_{i1} converges.

The numerator of the information ratio is the square of the expected instantaneous one-sided regret, $\mathbb{E} \left[\bar{\mathbf{r}}^T (\mathbf{x}_i^* - \mathbf{x}_{i2}) \mid \mathcal{H}_{i-1}^{(2)} \right]^2$. Conditioned upon the history $\mathcal{H}_{i-1}^{(2)}$, this does not depend upon the action \mathbf{x}_{i1} , and thus is unaffected by the distribution of \mathbf{x}_{i1} (recall that the \mathbf{x}_{i1} -dependency cancels in the regret formulation in (40)).

Thus, the information ratio can be asymptotically upper-bounded:

$$\lim_{i \rightarrow \infty} \Gamma_i \leq \lim_{i \rightarrow \infty} \Gamma_{i, \pi_{i1} \text{ fixed}} \leq \bar{\Gamma}_{\pi_{i1} \text{ fixed}}.$$

This means that for all $\varepsilon > 0$, there exists i_0 such that for all $i > i_0$, $\Gamma_i \leq \bar{\Gamma}_{\pi_{i1} \text{ fixed}} + \varepsilon \leq \frac{SA}{2} + \varepsilon$. \square

A.4 OBTAINING THE ASYMPTOTIC REGRET RATE

By combining Lemma 17 with previous results, we obtain the final asymptotic Bayesian regret rate.

Theorem 3. *With probability $1 - \delta$, where δ is a parameter of the Bayesian linear regression model, the Bayesian regret $\mathbb{E}[\text{REG}(T)]$ of DPS achieves an asymptotic rate of $S\sqrt{2AT \log A}$.*

Proof. Combining Theorem 1, Theorem 2, and Lemma 17, the Bayesian one-sided regrets $\mathbb{E}[\text{REG}_1(T)]$ and $\mathbb{E}[\text{REG}_2(T)]$ of policies π_{i1} and π_{i2} respectively each achieve asymptotic rates of $S\sqrt{\frac{AT \log A}{2}}$. The total regret is the sum of the regret contributions from policies π_{i1} and π_{i2} , and so asymptotically, $\mathbb{E}[\text{REG}(T)]$ increases at a rate of at most $S\sqrt{2AT \log A}$. \square

A.5 FACTS ABOUT CONVERGENCE IN DISTRIBUTION

Recall that for a random variable X and a sequence of random variables (X_n) , $n \in \mathbb{N}$, $X_n \xrightarrow{D} X$ denotes that X_n converges to X in distribution, while $X_n \xrightarrow{P} X$ denotes that X_n converges to X in probability. We apply the following two facts about convergence in distribution:

Fact 1 (Billingsley (1968)). *For random variables $\mathbf{x}, \mathbf{x}_n \in \mathbb{R}^d$, where $n \in \mathbb{N}$, and any continuous function $g : \mathbb{R}^d \rightarrow \mathbb{R}$, if $\mathbf{x}_n \xrightarrow{D} \mathbf{x}$, then $g(\mathbf{x}_n) \xrightarrow{D} g(\mathbf{x})$.*

Fact 2 (Billingsley (1968)). *For random variables $\mathbf{x}_n \in \mathbb{R}^d$, $n \in \mathbb{N}$, and constant vector $\mathbf{c} \in \mathbb{R}^d$, $\mathbf{x}_n \xrightarrow{D} \mathbf{c}$ is equivalent to $\mathbf{x}_n \xrightarrow{P} \mathbf{c}$. Convergence in probability means that for any $\varepsilon > 0$, $P(\|\mathbf{x}_n - \mathbf{c}\|_2 \geq \varepsilon) \rightarrow 0$ as $n \rightarrow \infty$.*

B CREDIT ASSIGNMENT MODELS

This appendix contains the mathematical details of the credit assignment models evaluated in our experiments. Afterward, we also discuss possible avenues for extending our regret analysis techniques to additional credit assignment models besides Bayesian linear regression.

B.1 BAYESIAN LINEAR REGRESSION

Define $X \in \mathbb{R}^{N \times d}$ as the observation matrix after N preferences, in which the i^{th} row contains observation $\mathbf{x}_i = \mathbf{x}_{i2} - \mathbf{x}_{i1}$, while $\mathbf{y} \in \mathbb{R}^N$ is the vector of corresponding preference labels, with i^{th} element $y_i \in \{-\frac{1}{2}, \frac{1}{2}\}$.

Section 4.1 defines the Bayesian linear regression credit assignment model to which our theoretical guarantees apply. Because the $\beta_i(\delta)$ factor necessary for the theoretical guarantees results in a conservative covariance matrix leading to over-exploration, our simulations implement the following, more practical, variant. We define a Gaussian prior over the reward vector $\mathbf{r} \in \mathbb{R}^d$: $\mathbf{r} \sim \mathcal{N}(0, \lambda^{-1}I)$. The likelihood of the data conditioned upon \mathbf{r} is also Gaussian:

$$p(\mathbf{y}|X, \mathbf{r}; \sigma^2) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp\left(-\frac{1}{2\sigma^2}\|\mathbf{y} - X\mathbf{r}\|^2\right).$$

This conjugate prior and likelihood lead to the following closed-form posterior:

$$\mathbf{r}|X, \mathbf{y}, \sigma^2, \lambda \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma), \text{ where } \boldsymbol{\mu} = (X^T X + \sigma^2 \lambda I)^{-1} X^T \mathbf{y} \text{ and } \Sigma = \sigma^2 (X^T X + \sigma^2 \lambda I)^{-1}.$$

B.2 GAUSSIAN PROCESS REGRESSION

Credit assignment via Gaussian processes (Rasmussen and Williams, 2006) extends the linear credit assignment model in B.1 to larger state and action spaces by generalizing across nearby states and actions. In this and the following section, we consider two Gaussian process-based credit assignment approaches.

To perform credit assignment via Gaussian process regression, we assign binary labels to each trajectory based on whether it is preferred or dominated. We place a Gaussian process prior upon the utilities of the state-action pairs. Using that a trajectory’s total reward is a sum over component state-action utilities, we will show how to perform inference over sums of Gaussian process variables to infer the state-action utilities from each trajectory’s total utility. As the total utility of each trajectory is not observed in practice, the binary preference labels are instead substituted as approximations in their place.

Let $\{\tilde{s}_1, \dots, \tilde{s}_d\}$ denote the $d = SA$ state-action pairs. In this section, the data matrix $Z \in \mathbb{R}^{2N \times d}$ holds all state-action visitation vectors $\mathbf{x}_{k1}, \mathbf{x}_{k2}$, for DPS iterations $k \in \{1, \dots, N\}$. (This contrasts with the other credit assignment methods, which learn from their differences, $\mathbf{x}_{k2} - \mathbf{x}_{k1}$.) Let \mathbf{z}_i^T be the i^{th} row of Z , such that $Z = [\mathbf{z}_1 \dots, \mathbf{z}_{2N}]^T$, and $\mathbf{z}_i = \mathbf{x}_{kj}$ for some DPS iteration k and $j \in \{1, 2\}$, that is, \mathbf{z}_i contains the state-action visit counts for the i^{th} trajectory rollout. In particular, the ij^{th} matrix element $z_{ij} = [Z]_{ij}$ is the number of times that the i^{th} observed trajectory \mathbf{z}_i visits state-action \tilde{s}_j .

The label vector is $\mathbf{y} \in \mathbb{R}^{2N}$, where the i^{th} element y_i is the preference label corresponding to the i^{th} observed trajectory. For instance, if $\tau_{i2} \succ \tau_{i1}$, then \mathbf{x}_{i2} receives a label of $\frac{1}{2}$, while \mathbf{x}_{i1} is labelled $-\frac{1}{2}$. As before, we use $\bar{r}(\tilde{s})$ to denote the true utility of state-action \tilde{s} , with $\bar{r}(\tau)$ being trajectory τ ’s total utility along the state-action pairs it encounters. To infer \bar{r} , we approximate each $\bar{r}(\tau_i)$ with its preference label y_i .

We place a Gaussian process prior upon the rewards \bar{r} : $\bar{r} \sim \mathcal{GP}(\boldsymbol{\mu}_r, K_r)$, where $\boldsymbol{\mu}_r \in \mathbb{R}^d$ is the prior mean and $K_r \in \mathbb{R}^{d \times d}$ is the prior covariance matrix, such that $[K_r]_{ij}$ models the prior covariance between $\bar{r}(\tilde{s}_i)$ and $\bar{r}(\tilde{s}_j)$. We model trajectory τ_i ’s total utility, $\bar{r}(\tau_i)$, as a sum over the latent state-action utilities: $\bar{r}(\tau_i) = \sum_{j=1}^d z_{ij} \bar{r}(\tilde{s}_j)$. Let R_i be a noisy version of $\bar{r}(\tau_i)$: $R_i = \bar{r}(\tau_i) + \varepsilon_i$, where $\varepsilon_i \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ is i.i.d. noise. Then, given rewards \bar{r} , we expect:

$$R_i = \sum_{j=1}^d z_{ij} \bar{r}(\tilde{s}_j) + \varepsilon_i.$$

Because any linear combination of jointly Gaussian variables is Gaussian, R_i is a Gaussian process over the values $\{z_{i1}, \dots, z_{id}\}$. Let $\mathbf{R} \in \mathbb{R}^{2N}$ be the vector with i^{th} element equal to R_i . We will calculate the relevant expectations

and covariances to show that $\bar{\mathbf{r}} \sim \mathcal{GP}(\boldsymbol{\mu}_r, K_r)$ and \mathbf{R} have the following jointly-Gaussian distribution:

$$\begin{bmatrix} \bar{\mathbf{r}} \\ \mathbf{R} \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \boldsymbol{\mu}_r \\ X\boldsymbol{\mu}_r \end{bmatrix}, \begin{bmatrix} K_r & K_r Z^T \\ ZK_r^T & ZK_r Z^T + \sigma_\varepsilon^2 I \end{bmatrix} \right). \quad (55)$$

The standard approach for obtaining a conditional distribution from a joint Gaussian distribution (Rasmussen and Williams, 2006) yields $\bar{\mathbf{r}}|\mathbf{R} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$, where:

$$\boldsymbol{\mu} = \boldsymbol{\mu}_r + K_r Z^T [ZK_r Z^T + \sigma_\varepsilon^2 I]^{-1} (\mathbf{R} - Z\boldsymbol{\mu}_r) \quad (56)$$

$$\Sigma = K_r - K_r Z^T [ZK_r Z^T + \sigma_\varepsilon^2 I]^{-1} ZK_r^T. \quad (57)$$

In practice, we do not observe the variable \mathbf{R} . Instead, \mathbf{R} is approximated with the observed preference labels \mathbf{y} , $\mathbf{R} \approx \mathbf{y}$, to perform credit assignment inference.

Next, we derive the posterior inference equations (56) and (57) used in Gaussian process regression credit assignment. We infer the state-action rewards $\bar{\mathbf{r}}$ given noisy observations \mathbf{R} of the trajectories' total utilities via the following four steps, corresponding to the next four subsections:

- A) Model the state-action utilities $\bar{r}(\tilde{s})$ as a Gaussian process over state-action pairs \tilde{s} .
- B) Model the trajectory utilities \mathbf{R} as a Gaussian process that results from summing the state-action utilities $\bar{r}(\tilde{s})$.
- C) Using the two Gaussian processes defined in A) and B), obtain the covariance matrix between the values of $\{\bar{r}(\tilde{s})|\tilde{s} \in 1, \dots, d\}$ and $\{R_i|i \in 1, \dots, 2N\}$.
- D) Write the joint Gaussian distribution in (55) between the values of $\{\bar{r}(\tilde{s})|\tilde{s} \in 1, \dots, d\}$ and $\{R_i|i \in 1, \dots, 2N\}$, and obtain the posterior distribution of $\bar{\mathbf{r}}$ over all state-actions given \mathbf{R} (Equations (56) and (57)).

B.2.1 The state-action utility Gaussian process

We model the state-action utilities as a Gaussian process over \tilde{s} , with mean $\mathbb{E}[\bar{r}(\tilde{s})] = \mu_r(\tilde{s})$ and covariance kernel $\text{Cov}(\bar{r}(\tilde{s}), \bar{r}(\tilde{s}')) = k_r(\tilde{s}, \tilde{s}')$, for all state-action pairs \tilde{s}, \tilde{s}' . For instance, k_r could be the squared exponential kernel:

$$k_r(\tilde{s}, \tilde{s}') = \sigma_f^2 \exp \left(-\frac{1}{2} \left(\frac{\|\tilde{s} - \tilde{s}'\|}{l} \right)^2 \right) + \sigma_n^2 \delta_{ij}, \quad (58)$$

where σ_f is the signal variance, l is the kernel lengthscale, σ_n is the noise variance, and δ_{ij} is the Kronecker delta function. Thus,

$$\bar{r}(\tilde{s}) \sim \mathcal{GP}(\mu_r(\tilde{s}), k_r(\tilde{s}, \tilde{s}')).$$

Define $\boldsymbol{\mu}_r \in \mathbb{R}^d$ such that the i^{th} element is $[\boldsymbol{\mu}_r]_i = \mu_r(\tilde{s}_i)$, the prior mean of state-action \tilde{s}_i 's utility. Let $K_r \in \mathbb{R}^{d \times d}$ be the covariance matrix over state-action utilities, such that $[K_r]_{ij} = k_r(\tilde{s}_i, \tilde{s}_j)$. Therefore, the reward vector $\bar{\mathbf{r}}$ is also a Gaussian process:

$$\bar{\mathbf{r}} \sim \mathcal{GP}(\boldsymbol{\mu}_r, K_r).$$

B.2.2 The trajectory utility Gaussian process

By assumption, the trajectory utilities $\mathbf{R} \in \mathbb{R}^{2N}$ are sums of the latent state-action utilities via the following relationship between \mathbf{R} and $\bar{\mathbf{r}}$:

$$R(\mathbf{z}_i) := R_i = \sum_{j=1}^d z_{ij} \bar{r}(\tilde{s}_j) + \varepsilon_i,$$

where ε_i are i.i.d. noise variables distributed according to $\mathcal{N}(0, \sigma_\varepsilon^2)$. Note that $R(\mathbf{z}_i)$ is a Gaussian process over $\mathbf{z}_i \in \mathbb{R}^d$ because $\{\bar{r}(\tilde{s}_j), \forall j\}$ are jointly normally distributed by definition of a Gaussian process, and any linear

combination of jointly Gaussian variables has a univariate normal distribution. Next, we calculate the expectation and covariance of \mathbf{R} over the observations. The expectation of the i^{th} element $R_i = R(\mathbf{z}_i)$ can be expressed:

$$\mathbb{E}[R_i] = \mathbb{E} \left[\sum_{j=1}^d z_{ij} \bar{r}(\tilde{s}_j) + \varepsilon_i \right] = \sum_{j=1}^d z_{ij} \mathbb{E}[\bar{r}(\tilde{s}_j)] = \sum_{j=1}^d z_{ij} \mu_r(\tilde{s}_j).$$

The expectation over \mathbf{R} can thus be written as $\mathbb{E}[\mathbf{R}(Z)] = Z\boldsymbol{\mu}_r$. Next, we model the covariance matrix of \mathbf{R} . The ij^{th} element of this matrix is the covariance of $R(\mathbf{z}_i)$ and $R(\mathbf{z}_j)$:

$$\begin{aligned} \text{Cov}(R(\mathbf{z}_i), R(\mathbf{z}_j)) &= \mathbb{E}[R(\mathbf{z}_i)R(\mathbf{z}_j)] - \mathbb{E}[R(\mathbf{z}_i)]\mathbb{E}[R(\mathbf{z}_j)] \\ &= \mathbb{E} \left[\left(\sum_{k=1}^d z_{ik} \bar{r}(\tilde{s}_k) + \varepsilon_i \right) \left(\sum_{m=1}^d z_{jm} \bar{r}(\tilde{s}_m) + \varepsilon_j \right) \right] - \left(\sum_{k=1}^d z_{ik} \mu_r(\tilde{s}_k) \right) \left(\sum_{m=1}^d z_{jm} \mu_r(\tilde{s}_m) \right) \\ &= \sum_{k=1}^d \sum_{m=1}^d z_{ik} z_{jm} \mathbb{E}[\bar{r}(\tilde{s}_k) \bar{r}(\tilde{s}_m)] + \mathbb{E}[\varepsilon_i \varepsilon_j] - \sum_{k=1}^d \sum_{m=1}^d z_{ik} z_{jm} \mu_r(\tilde{s}_k) \mu_r(\tilde{s}_m) \\ &= \sum_{k=1}^d \sum_{m=1}^d \left\{ z_{ik} z_{jm} [\text{Cov}(\bar{r}(\tilde{s}_k), \bar{r}(\tilde{s}_m)) + \mu_r(\tilde{s}_k) \mu_r(\tilde{s}_m)] - z_{ik} z_{jm} \mu_r(\tilde{s}_k) \mu_r(\tilde{s}_m) + \sigma_\varepsilon^2 \mathbb{I}_{[i=j]} \right\} \\ &= \sum_{k=1}^d \sum_{m=1}^d z_{ik} z_{jm} \text{Cov}(\bar{r}(\tilde{s}_k), \bar{r}(\tilde{s}_m)) + \sigma_\varepsilon^2 \mathbb{I}_{[i=j]} \\ &= \sum_{k=1}^d \sum_{m=1}^d z_{ik} z_{jm} k_r(\tilde{s}_k, \tilde{s}_m) + \sigma_\varepsilon^2 \mathbb{I}_{[i=j]} = \mathbf{z}_i^T K_r \mathbf{z}_j + \sigma_\varepsilon^2 \mathbb{I}_{[i=j]}. \end{aligned}$$

We can then write the covariance matrix of \mathbf{R} as K_R , where $[K_R]_{ij} := \text{Cov}(R(\mathbf{z}_i), R(\mathbf{z}_j)) = \mathbf{z}_i^T K_r \mathbf{z}_j + \sigma_\varepsilon^2 \mathbb{I}_{[i=j]}$. From here, it can be seen that $K_R = ZK_r Z^T + \sigma_\varepsilon^2 I$:

$$ZK_r Z^T = \begin{bmatrix} \mathbf{z}_1^T \\ \mathbf{z}_2^T \\ \vdots \\ \mathbf{z}_{2N}^T \end{bmatrix} K_r \begin{bmatrix} \mathbf{z}_1 & \mathbf{z}_2 & \dots & \mathbf{z}_{2N} \end{bmatrix} = \begin{bmatrix} \mathbf{z}_1^T K_r \mathbf{z}_1 & \dots & \mathbf{z}_1^T K_r \mathbf{z}_{2N} \\ \vdots & \ddots & \vdots \\ \mathbf{z}_{2N}^T K_r \mathbf{z}_1 & \dots & \mathbf{z}_{2N}^T K_r \mathbf{z}_{2N} \end{bmatrix} = K_R - \sigma_\varepsilon^2 I.$$

B.2.3 Covariance between state-action and trajectory utilities

We next consider the covariance between $\bar{\mathbf{r}}$ and \mathbf{R} , denoted $K_{r,R}$:

$$[K_{r,R}]_{ij} = \text{Cov}([\bar{\mathbf{r}}]_i, [\mathbf{R}]_j) = \text{Cov}(\bar{r}(\tilde{s}_i), R(\mathbf{z}_j)).$$

This covariance matrix can be expressed in terms of Z , K_r , and $\boldsymbol{\mu}_r$:

$$\begin{aligned}
[K_{r,R}]_{ij} &= \text{Cov}(\bar{r}(\tilde{s}_i), R(\mathbf{z}_j)) = \text{Cov}\left(\bar{r}(\tilde{s}_i), \sum_{k=1}^d z_{jk} \bar{r}(\tilde{s}_k) + \varepsilon_j\right) \\
&= \mathbb{E}\left[\bar{r}(\tilde{s}_i) \sum_{k=1}^d z_{jk} \bar{r}(\tilde{s}_k) + \varepsilon_j \bar{r}(\tilde{s}_i)\right] - \mathbb{E}[\bar{r}(\tilde{s}_i)] \mathbb{E}\left[\sum_{k=1}^d z_{jk} \bar{r}(\tilde{s}_k) + \varepsilon_j\right] \\
&= \sum_{k=1}^d z_{jk} \mathbb{E}[\bar{r}(\tilde{s}_i) \bar{r}(\tilde{s}_k)] - [\boldsymbol{\mu}_r(\tilde{s}_i)] [\mathbf{z}_j^T \boldsymbol{\mu}_r] \\
&= \sum_{k=1}^d z_{jk} \{ \text{Cov}(\bar{r}(\tilde{s}_i), \bar{r}(\tilde{s}_k)) + \mathbb{E}[\bar{r}(\tilde{s}_i)] \mathbb{E}[\bar{r}(\tilde{s}_k)] \} - \boldsymbol{\mu}_r(\tilde{s}_i) \mathbf{z}_j^T \boldsymbol{\mu}_r \\
&= \sum_{k=1}^d z_{jk} [k_r(\tilde{s}_i, \tilde{s}_k) + \boldsymbol{\mu}_r(\tilde{s}_i) \boldsymbol{\mu}_r(\tilde{s}_k)] - \boldsymbol{\mu}_r(\tilde{s}_i) \mathbf{z}_j^T \boldsymbol{\mu}_r \\
&= \sum_{k=1}^d z_{jk} k_r(\tilde{s}_i, \tilde{s}_k) + \boldsymbol{\mu}_r(\tilde{s}_i) \mathbf{z}_j^T \boldsymbol{\mu}_r - \boldsymbol{\mu}_r(\tilde{s}_i) \mathbf{z}_j^T \boldsymbol{\mu}_r = \sum_{k=1}^d z_{jk} k_r(\tilde{s}_i, \tilde{s}_k) = \mathbf{z}_j^T [K_r]_{i,:}^T,
\end{aligned}$$

where $[K_r]_{i,:}^T$ is the column vector obtained by transposing the i^{th} row of K_r . It is evident that $K_{r,R} = K_r Z^T$.

B.2.4 Posterior inference over state-action utilities

Merging the previous three subsections' results, one obtains the following joint probability density between $\bar{\mathbf{r}}$ and \mathbf{R} :

$$\begin{bmatrix} \bar{\mathbf{r}} \\ \mathbf{R} \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \boldsymbol{\mu}_r \\ Z \boldsymbol{\mu}_r \end{bmatrix}, \begin{bmatrix} K_r & K_r Z^T \\ Z K_r^T & Z K_r Z^T + \sigma_\varepsilon^2 I \end{bmatrix}\right).$$

This relationship expresses all components of the joint Gaussian density in terms of Z , K_r , and $\boldsymbol{\mu}_r$, or in other words, in terms of the observed state-action visitation counts (i.e., Z) and the Gaussian process prior on $\bar{\mathbf{r}}$. Via the standard approach for obtaining a conditional distribution from a joint Gaussian distribution, we obtain $\bar{\mathbf{r}} | \mathbf{R} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$, where the expressions for $\boldsymbol{\mu}$ and Σ are given by Equations (56) and (57) above. Substituting \mathbf{y} for \mathbf{R} , we have expressed the conditional posterior density of $\bar{\mathbf{r}}$ in terms of Z , \mathbf{y} , K_r , and $\boldsymbol{\mu}_r$.

B.3 GAUSSIAN PROCESS PREFERENCE MODEL

Finally, we show how to extend the preference-based Gaussian process model defined in Chu and Ghahramani (2005) from the dueling bandit setting to the PBRL setting to perform credit assignment. Similarly to the GP regression model, this approach places a Gaussian prior over possible rewards $\bar{\mathbf{r}}$; in contrast, however, this method explicitly models the likelihood of the observed preferences given the utilities $\bar{\mathbf{r}}$, and thus it is a more theoretically-justified approach for handling preference data.

We elicit a preference feedback dataset $\mathcal{D} = \{\tau_{i2} \succ \tau_{i1} \mid i = 1, \dots, N\}$, where $\tau_{i2} \succ \tau_{i1}$ indicates that trajectory τ_{i2} is preferred to τ_{i1} in preference i . Without loss of generality, we index each trajectory pair i such that $\tau_{i2} \succ \tau_{i1}$. As before, we assume that each state-action pair \tilde{s}_j , $j \in \{1, \dots, d\}$, has a latent, underlying utility $\bar{r}(\tilde{s}_j)$. In vector form, these are written: $\bar{\mathbf{r}} = [\bar{r}(\tilde{s}_1), \bar{r}(\tilde{s}_2), \dots, \bar{r}(\tilde{s}_d)]^T$. We define a Gaussian prior over $\bar{\mathbf{r}}$:

$$p(\bar{\mathbf{r}}) = \frac{1}{(2\pi)^{\frac{d}{2}} |\Sigma|^{\frac{1}{2}}} \exp\left(-\frac{1}{2} \bar{\mathbf{r}}^T \Sigma^{-1} \bar{\mathbf{r}}\right), \tag{59}$$

where $\Sigma \in \mathbb{R}^{d \times d}$ and $[\Sigma]_{ij} = k(\bar{r}(\tilde{s}_i), \bar{r}(\tilde{s}_j))$ for some kernel function k , such as the squared exponential kernel defined in (58). Next, we assume that the likelihood of the i^{th} preference given utilities $\bar{\mathbf{r}}$ takes the following form:

$$P(\tau_{i2} \succ \tau_{i1} \mid \bar{\mathbf{r}}) = g\left(\frac{\bar{r}(\tau_{i2}) - \bar{r}(\tau_{i1})}{c}\right),$$

where $g(\cdot)$ is a monotonically-increasing link function that is bounded between 0 and 1, and $c > 0$ is a model hyperparameter controlling the degree of preference noise. The total return $\bar{r}(\tau_{i1})$ of trajectory τ_{i1} can be written in terms of the corresponding state-action visitation vector \mathbf{x}_{i1} : $\bar{r}(\tau_{i1}) = \bar{\mathbf{r}}^T \mathbf{x}_{i1}$. Thus, the full likelihood expression is:

$$P(\mathcal{D} | \bar{\mathbf{r}}) = \prod_{i=1}^N g(z_i), \quad z_i := \frac{\bar{r}(\tau_{i2}) - \bar{r}(\tau_{i1})}{c} = \frac{\bar{\mathbf{r}}^T (\mathbf{x}_{i2} - \mathbf{x}_{i1})}{c} = \frac{\bar{\mathbf{r}}^T \mathbf{x}_i}{c}. \quad (60)$$

Given the preference dataset \mathcal{D} , we are interested in the posterior probability of $\bar{\mathbf{r}}$:

$$p(\bar{\mathbf{r}} | \mathcal{D}) \propto P(\mathcal{D} | \bar{\mathbf{r}}) p(\bar{\mathbf{r}}),$$

where the expressions for the prior $p(\bar{\mathbf{r}})$ and likelihood $P(\mathcal{D} | \bar{\mathbf{r}})$ are given by Equations (59) and (60), respectively. We sample reward vectors $\tilde{\mathbf{r}}$ from the posterior obtained via the Laplace approximation, $\tilde{\mathbf{r}} \sim \mathcal{N}(\hat{\mathbf{r}}_{\text{MAP}}, \alpha \Sigma_{\text{MAP}})$, where:

$$\hat{\mathbf{r}}_{\text{MAP}} = \operatorname{argmin}_{\mathbf{r}} S(\mathbf{r}), \quad (61)$$

$$\Sigma_{\text{MAP}} = (\nabla_{\mathbf{r}}^2 S(\mathbf{r})|_{\hat{\mathbf{r}}_{\text{MAP}}})^{-1}, \quad (62)$$

and $S(\mathbf{r}) := \frac{1}{2} \mathbf{r}^T \Sigma^{-1} \mathbf{r} - \sum_{i=1}^N \log g(z_i)$ (note that $S(\mathbf{r})$ is equivalent to $-\log p(\mathbf{r}, \mathcal{D})$, neglecting constant terms); lastly, $\alpha > 0$ is a tunable hyperparameter that influences the balance between exploration and exploitation. In order for the Laplace approximation to be valid, $S(\mathbf{r})$ must be a convex function in \mathbf{r} : this guarantees that the optimization problem in (61) is convex and that the covariance matrix defined by (62) is positive semidefinite, and therefore a valid Gaussian covariance matrix. Convexity of $S(\mathbf{r})$ can be established by demonstrating that its Hessian matrix is positive definite. It can be shown that for any \mathbf{r} , $\nabla_{\mathbf{r}}^2 S(\mathbf{r}) = \Sigma^{-1} + \Lambda$, where:

$$\Lambda_{mn} := \frac{1}{c^2} \sum_{i=1}^N [\mathbf{x}_i]_m [\mathbf{x}_i]_n \left[-\frac{g''(z_i)}{g(z_i)} + \left(\frac{g'(z_i)}{g(z_i)} \right)^2 \right], \quad (63)$$

for $\mathbf{x}_i = \mathbf{x}_{i2} - \mathbf{x}_{i1}$. To show that $\nabla_{\mathbf{r}}^2 S(\mathbf{r})$ is positive definite, because the prior covariance Σ is positive definite, it suffices to show that Λ is positive semidefinite. From (63), one can see that:

$$\Lambda = \frac{1}{c^2} \sum_{i=1}^N \left[-\frac{g''(z_i)}{g(z_i)} + \left(\frac{g'(z_i)}{g(z_i)} \right)^2 \right] \mathbf{x}_i \mathbf{x}_i^T.$$

Clearly $\mathbf{x}_i \mathbf{x}_i^T$ is positive semidefinite, and thus we arrive at the following sufficient condition for convexity of $S(\mathbf{r})$:

$$\left[-\frac{g''(z)}{g(z)} + \left(\frac{g'(z)}{g(z)} \right)^2 \right] \geq 0 \text{ for all } z \in \mathbb{R}.$$

This condition is in particular satisfied for the Gaussian link function, $g_{\text{Gaussian}}(\cdot) = \Phi(\cdot)$, where Φ is the standard Gaussian CDF, as well as for the sigmoidal link function, $g_{\text{sig}}(x) := \sigma(x) = \frac{1}{1 + \exp(-x)}$. Our experiments utilize the sigmoidal link function.

B.3.1 Bayesian Logistic Regression

Many of our experiments with the Gaussian process preference model fall under the special case of Bayesian logistic regression, in which $c = 1$, g is the sigmoidal link function, and the prior covariance matrix is diagonal, i.e. $\Sigma = \lambda I$; for instance, the latter condition occurs with the squared exponential kernel defined in (58) when its lengthscale l is set to zero. In this case, the Gaussian prior over possible reward vectors $\mathbf{r} \in \mathbb{R}^d$ is: $\mathbf{r} \sim \mathcal{N}(\mathbf{0}, \lambda I)$, where $\lambda > 0$. Setting the i^{th} preference label y_i equal to 1 if $\tau_{i2} \succ \tau_{i1}$, while $y_i = -1$ if $\tau_{i1} \succ \tau_{i2}$, the logistic regression likelihood is:

$$p(\mathcal{D} | \mathbf{r}) = \prod_{i=1}^N p(\mathbf{x}_i, y_i | \mathbf{r}) = \prod_{i=1}^N \frac{1}{1 + \exp(-y_i \mathbf{x}_i^T \mathbf{r})}.$$

We approximate the posterior, $p(\mathbf{r} | \mathcal{D}) \propto p(\mathcal{D} | \mathbf{r})p(\mathbf{r})$, as Gaussian via the Laplace approximation:

$$\begin{aligned}
 p(\mathbf{r} | \mathcal{D}) &\approx \mathcal{N}(\hat{\mathbf{r}}_{\text{MAP}}, \alpha \Sigma_{\text{MAP}}), \text{ where:} \\
 \hat{\mathbf{r}}_{\text{MAP}} &= \underset{\mathbf{r}}{\operatorname{argmin}} f(\mathbf{r}), \quad f(\mathbf{r}) := -\log p(\mathcal{D}, \mathbf{r}) = -\log p(\mathbf{r}) - \log p(\mathcal{D} | \mathbf{r}), \\
 \Sigma_{\text{MAP}} &= \left(\nabla_{\mathbf{r}}^2 f(\mathbf{r}) \Big|_{\hat{\mathbf{r}}} \right)^{-1}, \text{ where the optimization problem in (64) is convex,}
 \end{aligned} \tag{64}$$

and $\alpha > 0$ is a tunable hyperparameter that influences the balance between exploration and exploitation.

B.4 EXTENDING PROOF TECHNIQUES TO OTHER CREDIT ASSIGNMENT MODELS

Currently, our proof methodology treats only the Bayesian linear regression credit assignment model. Extending it to other credit assignment models, such as the Gaussian process-based and Bayesian logistic regression methods detailed above, is an important direction for future work. Recall that our theoretical analysis follows three main steps:

1. Prove that DPS is asymptotically-consistent, that is, over time, the probability that DPS selects the optimal policy approaches 1 (Appendix A.1).
2. Assume that in each iteration i , policy π_{i1} is drawn from a fixed distribution while policy π_{i2} is selected by DPS. Then, asymptotically bound the one-sided regret for π_{i2} (Appendix A.2).
3. Assume that policy π_{i1} is drawn from a drifting but converging distribution while policy π_{i2} is selected by DPS. Then, asymptotically bound the one-sided regret for π_{i2} (Appendix A.3).

Notably, this proof outline does not depend upon any specific credit assignment model definition, and thus could likely extend to many models. Step 1) requires asymptotic consistency of the credit assignment model; this is not a restrictive requirement, as a non-asymptotically consistent model would not yield sublinear regret. Step 2) fixes one of the two distributions from which policies are drawn; this removes some mathematical difficulties inherent in analyzing preference-based sampling, making the required analysis more similar to the numeric feedback setting. Step 3) depends mainly upon continuity arguments.

The information-theoretic perspective used to prove 2) and 3) likely applies to a wide class of credit assignment models. For instance, recent work has applied the bandit analysis framework in Russo and Van Roy (2016) to bandits with rewards generated via a class of general link functions (Dong and Van Roy, 2018). In particular, the information ratio has been studied for the logistic bandit problem (Dong et al., 2019); we expect that this work could be extended toward analyzing credit assignment via Bayesian logistic regression in the PBRL setting.

Another interesting direction would be to analyze the information ratio for the state transition dynamics model. Bounding this quantity would strengthen the results significantly, since our current analysis only considers dynamics convergence asymptotically. Such a result would also be of independent interest outside of preference-based learning.

Finally, the concept of approximate linearity (Sui et al., 2017) could perhaps help to bridge the gap between the preference and absolute-reward domains, as it has previously done in the bandit setting, and could help to extend existing proof techniques toward a wider class of link functions. In practice, we expect that DPS would perform well with any asymptotically-consistent credit assignment model that sufficiently captures users' preference behavior.

C WHY THE INFORMATION-THEORETIC REGRET ANALYSIS?

Several existing regret analyses in the linear bandit domain (Abbasi-Yadkori et al., 2011; Agrawal and Goyal, 2013; Abeille and Lazaric, 2017) utilize martingale concentration properties introduced by Abbasi-Yadkori et al. (2011). In these analyses, a key step requires sublinearly upper-bounding an expression of the form (e.g. Lemma 11 in Abbasi-Yadkori et al. (2011), Prop. 2 in Abeille and Lazaric (2017)):

$$\sum_{i=1}^n \mathbf{x}_i^T \left(\lambda I + \sum_{s=1}^{i-1} \mathbf{x}_s \mathbf{x}_s^T \right)^{-1} \mathbf{x}_i, \tag{65}$$

where $\lambda \geq 1$ and \mathbf{x}_i is the observation vector in iteration i . We will demonstrate that in the preference-feedback setting, the analogous quantity cannot always be sublinearly upper-bounded. Consider the setting defined in Section 3, with Bayesian linear regression credit assignment. Under preference feedback, we assume that the probability that one trajectory is preferred to another is fully determined by the *difference* between the trajectories’ total rewards: on iteration i , the algorithm receives a pair of observations $\mathbf{x}_{i1}, \mathbf{x}_{i2}$, with $\mathbf{x}_i := \mathbf{x}_{i2} - \mathbf{x}_{i1}$, and a preference generated according to $P(\tau_{i2} \succ \tau_{i1}) = \bar{\mathbf{r}}^T(\mathbf{x}_{i2} - \mathbf{x}_{i1}) + \frac{1}{2}$. Thus, only *differences* between compared trajectory feature vectors yield information about the rewards. Under this assumption, one can show that applying the martingale techniques yields the following variant of (65):

$$\sum_{i=1}^n \sum_{j=1}^2 \mathbf{x}_{ij}^T \left(\lambda I + \sum_{s=1}^{i-1} \mathbf{x}_s \mathbf{x}_s^T \right)^{-1} \mathbf{x}_{ij}. \quad (66)$$

This is because the expression within the matrix inverse comes from the posterior—and learning occurs with respect to the observations \mathbf{x}_i —while regret is incurred with respect to \mathbf{x}_{i1} and \mathbf{x}_{i2} ; in contrast, in the non-preference case (65), learning and regret both occur with respect to the same vectors \mathbf{x}_i .

To see that (66) does not necessarily have a sublinear upper bound, consider a deterministic MDP as a counterexample. For the regret to have a sublinear upper-bound, the probability of choosing the optimal policy must approach 1. In a fully deterministic MDP, this means that $P(\mathbf{x}_{i1} = \mathbf{x}_{i2}) \rightarrow 1$ as $i \rightarrow \infty$, and thus $P(\mathbf{x}_i = 0) \rightarrow 1$ as $i \rightarrow \infty$. Clearly, in this case, the inverted quantity in (66) acquires nonzero terms at a rate that decays in n , and so (66) does not have a sublinear upper bound.

Intuitively, to be able to upper-bound Equation (66), we would need the observations $\mathbf{x}_{i1}, \mathbf{x}_{i2}$ to contribute fully toward learning the rewards, rather than the contribution coming only from their difference. Notice that in the counterexample, even when the optimal policy is selected increasingly-often, corresponding to a low regret, (66) cannot be sublinearly bounded. In contrast, Russo and Van Roy (2016) introduces a more direct approach for quantifying the trade-off between instantaneous regret and information gained, as encapsulated by the *information ratio* defined therein; our theoretical analysis is thus based upon this framework.

D ADDITIONAL EXPERIMENTAL DETAILS

Python code for reproducing the experiments in this work is located at:
<https://github.com/ernovoseller/DuelingPosteriorSampling>.

Experiments were conducted in three simulated environments, described in Section 6: RiverSwim and random MDPs (Osband et al., 2013) and the simplified version of the Mountain Car problem described in Wirth (2017). We use a fixed episode horizon of 50 in the first two cases, while for Mountain Car, episodes have a maximum length of 500, but terminate sooner if the agent reaches the goal state. Figures 2, 3, and 4 display performance in the three environments for the five degrees of user preference noise evaluated. Experiments were run on an Ubuntu 16.04.3 machine with 32 GB of RAM and an Intel i7 processor. Some experiments were also run on an AWS server.

We detail the ranges of hyperparameter values tested for the different DPS credit assignment models, as well as the particular hyperparameters used in the displayed performance curves. Hyperparameters were tuned by considering mean performance over 30 experiment repetitions for each parameter setting considered; we only used the least-noisy preference feedback (logistic noise, $c = 0.001$) to tune the preferences; this value of c is small enough that the preferences are close to deterministic, except that the preferences are uniformly-random in tie cases. For both Gaussian process-based credit assignment models, we use the squared exponential kernel:

$$\mathcal{K}(x_i, x_j) = \sigma_f^2 \exp \left(-\frac{1}{2} \left(\frac{x_i - x_j}{l} \right)^2 \right) + \sigma_n^2 \delta_{ij},$$

where σ_f is the signal variance, l is the kernel lengthscale, σ_n is the noise variance, and δ_{ij} is the Kronecker delta function. Please see Appendix B for definitions of the other hyperparameters in the credit assignment models. Tables 2, 3, and 4 display both the tested ranges and optimized values (those appearing in the performance curves) for each case.

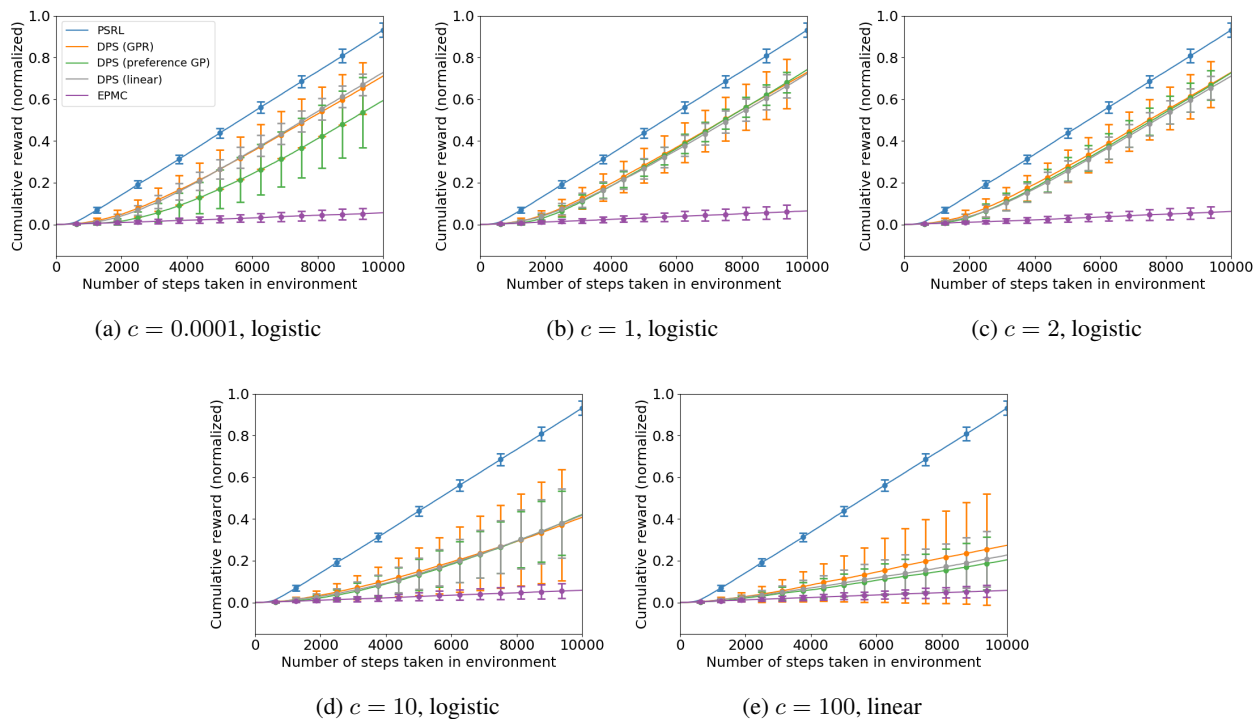


Figure 2: Empirical performance of DPS in the RiverSwim environment. Plots display mean \pm one standard deviation over 100 runs of each algorithm tested. Normalization is with respect to the total reward achieved by the optimal policy. Overall, we see that DPS performs well and is robust to the choice of credit assignment model.

The dynamics model, meanwhile, has a Dirichlet prior and posterior. Not to assume domain knowledge that differentiates the state-action pairs, we set all prior parameters of the Dirichlet model to be equal; for the RiverSwim and Random MDP environments, the prior is set to 1 for each state-action, creating a uniform prior over all dynamics models. This is a reasonable choice with small numbers of states and actions, so we do not optimize over different values. For the Mountain Car problem, smaller prior values perform better because they favor sparse dynamics distributions. For this environment, we test prior parameters ranging from 0.0001 to 1, and found 0.0005 to be the best-performing value among those tested.

The EMPC algorithm (Wirth and Furnkranz, 2013b) has two hyperparameter values, α and η . We optimize both of these jointly via a grid search over values of $(0.1, 0.2, \dots, 0.9)$, with 100 repetitions of each pair of values. The best-performing hyperparameter values (i.e. those achieving the highest total reward) are displayed in Table 1; these are the hyperparameter values depicted in the performance curve plots.

Finally, Figures 5, 6, and 7 illustrate how DPS’s performance varies as the hyperparameters are modified over a set of representative values from the ranges that we tested. These plots demonstrate that DPS is largely robust across many choices of model hyperparameters.

Table 1: Hyperparameters for the EPMC baseline algorithm. Each table element shows the best-performing α/η values for the corresponding simulation domain and noise parameter.

NOISE	LOGISTIC, 10	LOGISTIC, 2	LOGISTIC, 1	LOGISTIC, 0.0001	LINEAR
RiverSwim	0.1/0.8	0.3/0.7	0.1/0.2	0.8/0.8	0.3/0.1
Random MDPs	0.2/0.2	0.7/0.7	0.6/0.4	0.2/0.8	0.7/0.1

NOISE	LOGISTIC, 100	LOGISTIC, 20	LOGISTIC, 10	LOGISTIC, 0.0001	LINEAR
MountainCar	0.1/0.8	0.1/0.7	0.1/0.6	0.1/0.4	0.2/0.5

Table 2: Credit assignment hyperparameters for the RiverSwim Environment

MODEL	HYPERPARAMETER	RANGE TESTED	OPTIMIZED VALUE
Bayesian linear regression	σ	[0.05, 5]	0.5
	λ	[0.01, 10]	0.1
GP regression	σ_f^2	[0.001, 0.5]	0.1
	l	[0, 0] ([state, action])	0
	σ_n^2	[0.0001, 0.1]	0.001
GP preference (special case: Bayesian logistic regression)	$\lambda = \sigma_f^2 + \sigma_n^2$	[0.1, 30]	1
	α	[0.01, 1]	1
GP preference (varying c)	c	[1, 13]	N/A
	σ_f^2	[1]	
	l	[0, 0] ([state, action])	
	σ_n^2	[0.001]	
	α	[1]	

Table 3: Credit assignment hyperparameters for the Random MDP Environment

MODEL	HYPERPARAMETER	RANGE TESTED	OPTIMIZED VALUE
Bayesian linear regression	σ	[0.05, 5]	0.1
	λ	[0.01, 20]	10
GP regression	σ_f^2	[0.001, 1]	0.05
	l	[0, 0] ([state, action])	0
	σ_n^2	[0.0001, 0.1]	0.0005
GP preference (special case: Bayesian logistic regression)	$\lambda = \sigma_f^2 + \sigma_n^2$	[1, 15]	0.1
	α	[0.01, 1]	0.01
GP preference (varying c)	c	[0.0001, 1000]	N/A
	σ_f^2	[1]	
	l	[0, 0] ([state, action])	
	σ_n^2	[0.03]	
	α	[1]	

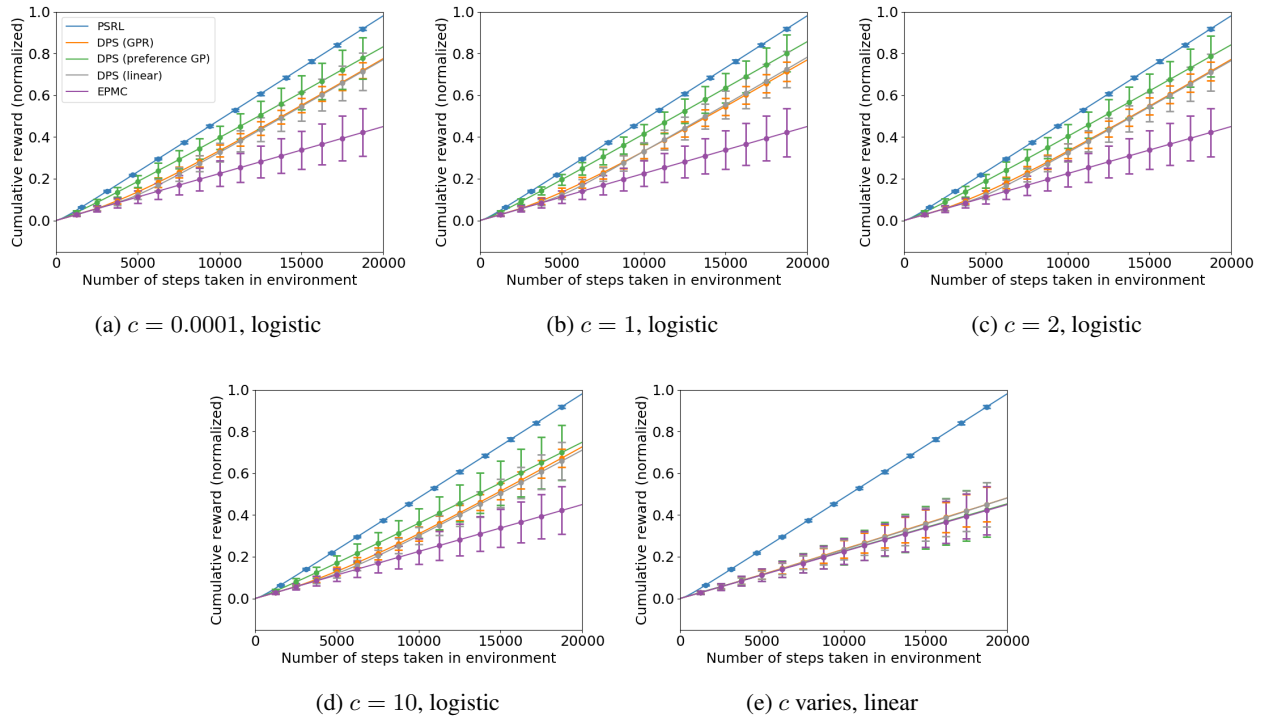


Figure 3: Empirical performance of DPS in the Random MDP environment. Plots display mean +/- one standard deviation over 100 runs of each algorithm tested. Normalization is with respect to the total reward achieved by the optimal policy. Overall, we see that DPS performs well and is robust to the choice of credit assignment model.

Table 4: Credit assignment hyperparameters for the Mountain Car Environment

MODEL	HYPERPARAMETER	RANGE TESTED	OPTIMIZED VALUE
Bayesian linear regression	σ	[0.001, 30]	10
	λ	[0.001, 10]	1
GP regression	σ_f^2	[0.0001, 10]	0.01
	l	$[x, x, 0], x \in [1, 3]$ ([position, velocity, action])	$x = 2$
	σ_n^2	[1e-7, 0.01]	1e-5
GP preference (special case: Bayesian logistic regression)	$\lambda = \sigma_f^2 + \sigma_n^2$	[0.0001, 10]	0.0001
	α	[0.0001, 1]	0.01
GP preference (varying c)	c	[10, 10000]	N/A
	σ_f^2	[1]	
	l	[2, 2, 0] ([position, velocity, action])	
	σ_n^2	[0.001]	
	α	[1]	

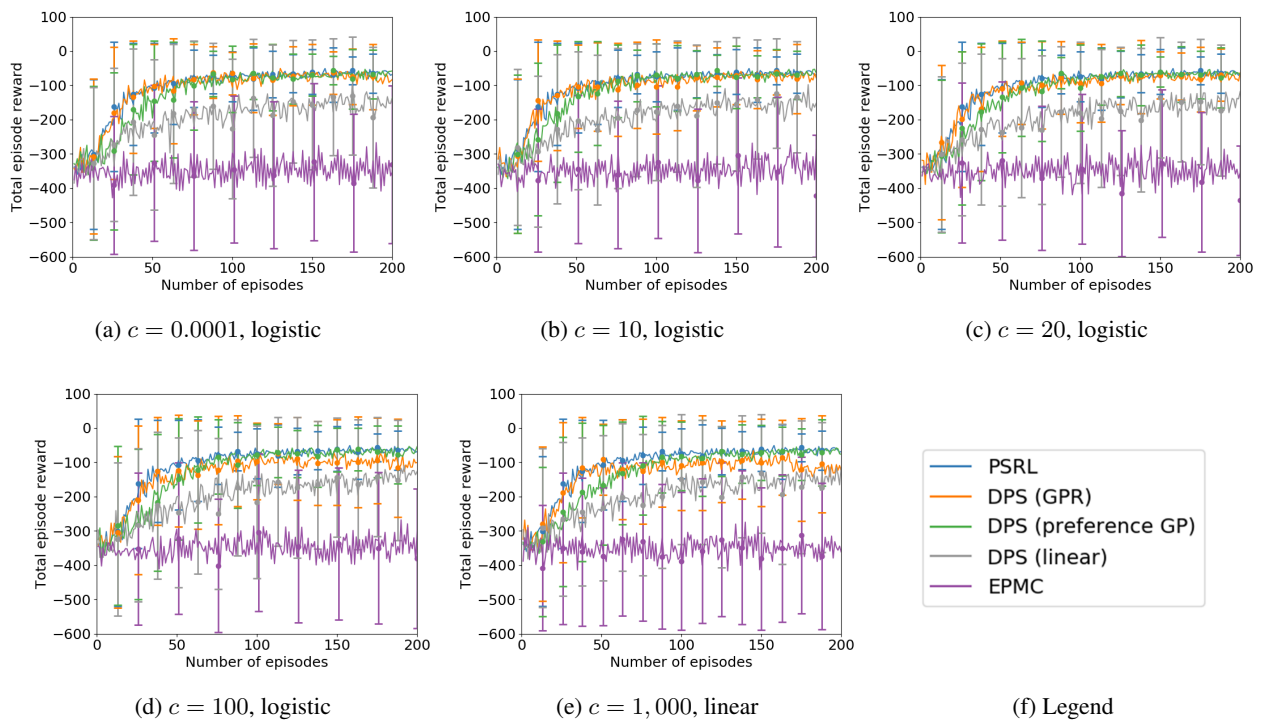





Figure 4: Empirical performance of DPS in the Mountain Car environment. Plots display mean \pm one standard deviation over 100 runs of each algorithm tested. Overall, we see that DPS performs well and is robust to the choice of credit assignment model.




(a) Bayesian linear regression



(b) Gaussian process regression




(c) Gaussian process preference model (special case: Bayesian logistic regression)




(d) Gaussian process preference model (varying c)

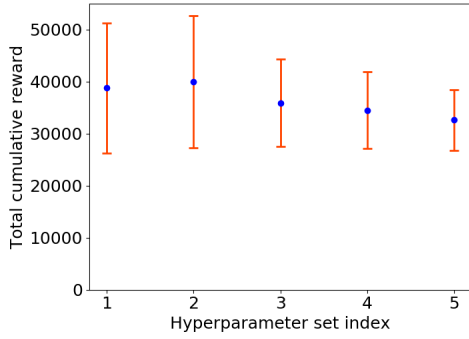
Figure 5: Empirical performance of DPS in the RiverSwim environment for different hyperparameter combinations. Plots display mean \pm one standard deviation over 30 runs of each algorithm tested with logistic user noise and $c = 0.001$. Overall, we see that DPS is robust to the choice of hyperparameters. The hyperparameter values depicted in each plot are (from left to right): for Bayesian linear regression, $(\sigma, \lambda) = \{(0.5, 0.1), (0.5, 10), (0.1, 0.1), (0.1, 10), (1, 0.1)\}$; for GP regression, $(\sigma_f^2, \sigma_n^2) = \{(0.1, 0.001), (0.1, 0.1), (0.01, 0.001), (0.001, 0.0001), (0.5, 0.1)\}$; for Bayesian logistic regression (special case of the GP preference model), $(\lambda, \alpha) = \{(1, 1), (30, 1), (20, 0.5), (1, 0.5), (30, 0.1)\}$; and additionally for the GP preference model, $c \in \{0.5, 1, 2, 5, 13\}$. See Table 2 for the values of any hyperparameters not specifically mentioned here.




(a) Bayesian linear regression



(b) Gaussian process regression




(c) Gaussian process preference model (special case: Bayesian logistic regression)




(d) Gaussian process preference model (varying c)


Figure 6: Empirical performance of DPS in the Random MDP environment for different hyperparameter combinations. Plots display mean \pm one standard deviation over 30 runs of each algorithm tested with logistic user noise and $c = 0.001$. Overall, we see that DPS is robust to the choice of hyperparameters. The hyperparameter values depicted in each plot are (from left to right): for Bayesian linear regression, $(\sigma, \lambda) = \{(0.1, 10), (0.1, 0.1), (0.05, 0.01), (0.5, 20), (1, 10)\}$; for GP regression, $(\sigma_f^2, \sigma_n^2) = \{(0.05, 0.0005), (0.001, 0.0001), (0.05, 0.1), (0.001, 0.0005), (1, 0.1)\}$; for Bayesian logistic regression (special case of the GP preference model), $(\lambda, \alpha) = \{(0.1, 0.01), (1, 0.01), (0.1, 1), (30, 0.1), (5, 0.5)\}$; and additionally for the GP preference model, $c \in \{1, 10, 15, 19, 100\}$. See Table 3 for the values of any hyperparameters not specifically mentioned here.




(a) Bayesian linear regression



(b) Gaussian process regression



(c) Gaussian process preference model (special case: Bayesian logistic regression)



(d) Gaussian process preference model (varying c)

Figure 7: Empirical performance of DPS in the Mountain Car environment for different hyperparameter combinations. Plots display mean \pm one standard deviation over 30 runs of each algorithm tested with logistic user noise and $c = 0.001$. Overall, we see that DPS is robust to the choice of hyperparameters. The hyperparameter values depicted in each plot are (from left to right): for Bayesian linear regression, $(\sigma, \lambda) = \{(10, 1), (10, 10), (30, 0.001), (0.001, 10), (0.1, 0.1)\}$; for GP regression, $(\sigma_f^2, l, \sigma_n^2) = \{(0.01, 2, 10^{-5}), (0.01, 1, 10^{-5}), (0.1, 2, 0.01), (1, 2, 0.001), (0.001, 3, 10^{-6})\}$; for Bayesian logistic regression (special case of the GP preference model), $(\lambda, \alpha) = \{(0.0001, 0.01), (0.1, 0.01), (0.0001, 0.0001), (0.001, 0.0001), (0.001, 0.01)\}$; and additionally for the GP preference model, $c \in \{10, 300, 400, 700, 1000\}$. See Table 4 for the values of any hyperparameters not specifically mentioned here.