# GREEDY AI ALIGNMENT METHODS ARE PROBABLY (AND PROVABLY) MORE EFFICIENT THAN YOU THINK

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Greedy AI alignment methods based on Bradley-Terry reward-model fitting and KL-regularized Alignment are remarkably effective in both offline and online preference-learning pipelines, yet existing $O(\log T)$ online and $O(1/\epsilon)$ offline KL-regularized regret rate guarantees can seem pessimistic relative to their empirical performance. We argue that this mismatch reflects a distinction between two learning targets: KL-regularized regret measures how fast we can recover the KL-regularized aligned policy, whereas we often only ask how fast we can learn the reward function that induces the correct best response. To isolate this reward-learning component, we study the traditional temperature-zero regret criterion, which evaluates only the top-ranked response at inference time. Under realizability, compactness, and a uniform best-response margin over the fitted reward class, we prove that exact reference-logged offline RLHF achieves exponentially small expected temperature-zero regret, and therefore reaches expected regret at most $\epsilon$ with $O(\log(1/\epsilon))$ offline comparisons. We also prove that exact greedy online RLHF achieves bounded $O(1)$ cumulative temperature-zero regret; the same guarantee extends to its equivalent pairwise online DPO formulation. By separating selector identification from full soft-policy recovery, our results provide a sharper theoretical explanation for the empirical efficiency of greedy alignment.

## 1 INTRODUCTION

Preference-based alignment (Christiano et al., 2017; Ouyang et al., 2022) trains AI systems to favor responses that people prefer. A platform, implicitly or explicitly, fits a *reward model* from comparisons between responses and then shifts the AI's response distribution toward higher-scoring outputs. Offline methods learn from a fixed comparison log; online methods repeatedly generate responses, collect feedback, and update the model. Both commonly use Kullback–Leibler (KL) regularization, which balances higher estimated rewards against departures from a reference model.

Recent theory establishes offline and online guarantees for such KL-regularized alignment (Zhao et al., 2024; 2025a;b). In particular, Wu et al. (2025) show that greedy online alignment achieves $O(\log T)$ cumulative KL-regularized regret without explicit optimism or pessimism. Here, *greedy* means optimizing the current estimated reward subject to KL regularization, not selecting responses deterministically; KL-regularized regret measures the gap between the learner's randomized policy and the optimal KL-regularized policy, which deliberately retains probability on lower-reward responses in order to remain close to the reference policy.

A central point of this paper is that such logarithmic KL-regularized regret guarantees need not describe how quickly the underlying reward model is being learned. KL-regularized regret asks how quickly the learner recovers an entire randomized, regularized policy. But once alignment is complete, many downstream inference-time applications do not require recovery of the entire randomized policy; they only require that the learned reward model correctly identifies which response is best (Nakano et al., 2021; Snell et al., 2024). Existing KL-regularized regret guarantees therefore leave unresolved how quickly greedy alignment learns this simpler, decision-relevant object: the identity of the best response. This distinction leaves an open question: *How large data/how long does greedy alignment need to learn a reward model accurate enough to identify which response is best?*



Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

To answer this question, we propose a traditional-style regret notion we call *temperature-zero regret*. Rather than evaluating the entire randomized policy induced by the learned reward, temperature-zero regret evaluates only the response that the learned reward ranks highest. It asks a simple question: if we use the learned reward model to choose its top response, how much reward do we lose relative to choosing the truly best response? By contrast, KL-regularized regret asks whether the entire finite-temperature response distribution has converged to the optimal regularized distribution. The former therefore measures best-response identification; the latter measures full soft-policy recovery.

This distinction matters when reward models are ultimately used in downstream tasks to rank or select outputs at inference time, where learning the exact probabilities assigned to all non-selected responses may be less important than correctly identifying the top-ranked response. Temperature-zero regret provides a direct measure of this decision-relevant component of learning. Methodologically, this criterion is closely related to simple regret in contextual bandits (Deshmukh et al., 2018; Krishnamurthy et al., 2023); our contribution is not the argmax criterion itself, but its analysis for the same greedy KL-regularized RLHF/DPO process, including adaptive online preference collection.

**Main contributions.** Under minor realistic regularity conditions, we establish two results.

- **Bounded online regret.** Standard exact greedy reinforcement learning from human feedback (RLHF) has bounded $O(1)$ expected cumulative temperature-zero regret (Theorem 4). The same guarantee holds for exact pairwise online direct preference optimization (DPO) over the reward-induced policy class (Corollary 6 in Appendix C). This bound evaluates the induced best-response rule, not the randomized responses sampled to collect feedback.
- **Exponential offline accuracy.** For exact offline RLHF trained on comparisons generated by the reference policy, expected temperature-zero regret decreases exponentially with sample size. Consequently, $O(\log(1/\epsilon))$ comparisons suffice to achieve expected regret at most $\epsilon$ (Theorem 5).

Table 1: **Greedy alignment under two evaluation criteria.** KL-regularized entries summarize prior guarantees (Wu et al., 2025); bold entries are our expected-regret guarantees (Theorems 4 and 5).

| | KL-regularized regret (Wu et al., 2025) | **Temperature-zero regret (this work)** |
|---|---|---|
| Learning target | Full soft-policy recovery | Learning reward model |
| Online cumulative regret ($T$ rounds) | $O(\log T)$ | $\boldsymbol{O(1)}$ |
| Offline one-step regret ($n$ comparisons) | $O(1/n)$ | $\boldsymbol{O(e^{-cn})}$ |
| Offline comparisons for regret $\leq \varepsilon$ | $O(1/\varepsilon)$ | $\boldsymbol{O(\log(1/\varepsilon))}$ |

A controlled simulation illustrates the distinction: one-step temperature-zero regret rapidly approaches zero and cumulative regret levels off, although preference collection remains randomized. Together, these results sharpen the interpretation of logarithmic policy-regret guarantees: *the reward model can finish learning which response to choose while the randomized policy is still being refined.*

Section 2 introduces the model and the offline and online alignment procedures. Section 3 defines temperature-zero regret and presents the theoretical results, and Section 4 reports the simulation. Related work appears in Appendix A.

## 2 Setup and greedy alignment

We study a common alignment pipeline: fit a reward model to preference feedback, then use the fitted reward to reweight a fixed reference policy.[1] The offline and online procedures we will introduce now share this pipeline. The distinction is how preference data are collected: offline learning uses a fixed reference-policy log, whereas online learning updates the sampling policy each round.

[1]For pairwise feedback, this pipeline also admits an equivalent DPO (Rafailov et al., 2023) formulation: under exact optimization over the reward-induced policy class, DPO is a reparameterization of greedy RLHF and yields the same policy iterates and temperature-zero selectors. See Appendix C.

Let $\mathcal{X}$ be a context space (prompts, queries, or tasks), and let $\mathcal{A}$ be a response space, which may be large or infinite. We do not assume that $\mathcal{A}$ is continuous; for LLMs, responses are naturally discrete token sequences, albeit from an enormous output space. Contexts are drawn independently from an unknown distribution $d_0$. A base AI model $\pi_0 : \mathcal{X} \to \Delta(\mathcal{A})$ represents the pre-alignment model, which we use as the starting point for our alignment procedure, such as a pretrained model or a supervised fine-tuned checkpoint.

Throughout, $\mathcal{A}$ is a separable metric space equipped with its Borel $\sigma$-field, and $\pi_0$ and the sampling policies are measurable probability kernels. Whenever a maximizing selector is used, we assume that the maximum is attained and that a measurable selector is available. Tie-breaking depends only on the maximizing set, so it is unchanged by adding an action-independent constant to a reward. All such selectors take values in the reference support. These are standing well-posedness conventions; compactness of a reward class alone does not imply attainment of a maximum over a noncompact action space.

Fix a slate size $K \geq 2$. Each preference observation consists of a context $x$, a slate $\mathbf{a} = (a_1, \ldots, a_K)$, and a preferred index $y \in \{1, \ldots, K\}$, drawn from an unknown choice law $P^\star(\cdot \mid x, \mathbf{a})$. The case $K = 2$ corresponds to pairwise feedback; larger $K$ allows preferences over larger slates.

### 2.1 Reward fitting and KL-regularized AI alignment

**Reward fitting** Let $\mathcal{F}$ be the learner's class of measurable reward functions $R : \mathcal{X} \times \mathcal{A} \to \mathbb{R}$. The learner associates each candidate reward with the multinomial logit preference model[2]

$$P_R(y = k \mid x, \mathbf{a}) = \frac{\exp(R(x, a_k))}{\sum_{\ell=1}^{K} \exp(R(x, a_\ell))}, \qquad k = 1, \ldots, K.$$

For $K = 2$, this is the Bradley–Terry model. Write $\mathbf{v}_R(x, \mathbf{a}) = (R(x, a_1), \ldots, R(x, a_K))$ and $\ell(\mathbf{v}, y) = \log(\sum_{k=1}^{K} e^{v_k}) - v_y$. Given a dataset $\mathcal{D}_n = \{(x_s, \mathbf{a}_s, y_s)\}_{s=1}^{n}$, the empirical preference loss and fitted reward are

$$\widehat{\mathcal{L}}_n(R) \coloneqq \frac{1}{n} \sum_{s=1}^{n} \ell\big(\mathbf{v}_R(x_s, \mathbf{a}_s), y_s\big), \quad \widehat{R}_n \in \arg\min_{R \in \mathcal{F}} \widehat{\mathcal{L}}_n(R). \tag{1}$$

The assumptions on $\mathcal{F}$ and its relation to the true feedback law $P^\star$ are stated with the theoretical results in Section 3. When the sample size or iteration index is immaterial, we suppress the subscript and write $\widehat{R}$ for a generic fitted reward.

**KL-regularized AI alignment** We then use the fitted reward $\widehat{R}$ to align the AI with KL regularization. For each context $x$, fix a constant $\eta > 0$ and assume that the fitted reward is bounded on the reference support, $\big|\widehat{R}(x, a)\big| \leq B$ for $\pi_0(\cdot \mid x)$-a.e. $a$, for some finite $B$. The policy-improvement step maximizes

$$\mathbb{E}_{A \sim \pi(\cdot \mid x)}[\widehat{R}(x, A)] - \eta^{-1} \operatorname{KL}\big(\pi(\cdot \mid x) \,\|\, \pi_0(\cdot \mid x)\big).$$

Its solution is the exponential tilt

$$\pi_{\widehat{R}}(da \mid x) = \frac{\exp\big(\eta \widehat{R}(x, a)\big)\, \pi_0(da \mid x)}{\int_{\mathcal{A}} \exp\big(\eta \widehat{R}(x, a')\big)\, \pi_0(da' \mid x)}. \tag{2}$$

which satisfies

$$e^{-2\eta B} \leq \frac{d\pi_{\widehat{R}}(\cdot \mid x)}{d\pi_0(\cdot \mid x)}(a) \leq e^{2\eta B} \qquad \pi_0(\cdot \mid x)\text{-a.s.},$$

so $\pi_{\widehat{R}}(\cdot \mid x)$ and $\pi_0(\cdot \mid x)$ are mutually absolutely continuous and therefore have the same null sets. Thus, KL-regularized AI alignment reweights responses already supported by the base AI, rather than assigning probability to response regions that the base AI rules out; see Lemma 13 in Appendix D. Both $\eta$ and the reference policy $\pi_0$ remain fixed throughout learning.

[2]The general preference extension is described in Appendix E.

### 2.2 OFFLINE AND ONLINE AI ALIGNMENT

**Offline RLHF.** The learner receives a fixed dataset collected independently across $s = 1, \dots, n$ according to

$$X_s \sim d_0, \qquad \mathbf{A}_s \sim \pi_0(\cdot \mid X_s)^{\otimes K}, \qquad Y_s \sim P^\star(\cdot \mid X_s, \mathbf{A}_s).$$

It fits $\widehat{R}_n^{\mathrm{off}}$ using (1) and forms the aligned policy $\pi_{\widehat{R}_n^{\mathrm{off}}}$ using (2). The fitted model does not influence the data in this fixed log.

**Greedy online RLHF.** Following Wu et al. (2025), initialize $\widehat{R}_0 \equiv 0$, $\pi_1 = \pi_0$, and $\mathcal{D}_0 = \emptyset$. At round $t$, collect an observation according to

$$X_t \sim d_0, \qquad \mathbf{A}_t \sim \pi_t(\cdot \mid X_t)^{\otimes K}, \qquad Y_t \sim P^\star(\cdot \mid X_t, \mathbf{A}_t).$$

Append $(X_t, \mathbf{A}_t, Y_t)$ to $\mathcal{D}_{t-1}$, refit $\widehat{R}_t$ on the accumulated dataset $\mathcal{D}_t$ using (1), and set $\pi_{t+1} = \pi_{\widehat{R}_t}$ using (2). Thus, the policy used to collect round-$t$ feedback is based on the first $t-1$ observations, while the resulting fit is used from round $t+1$. Every update remains relative to the original reference policy $\pi_0$, not the preceding policy $\pi_t$.

Both procedures produce a fitted reward and its associated randomized policy. Our analysis evaluates the response ranked highest by the fitted reward, rather than the full randomized policy. The next section formalizes this distinction through temperature-zero regret. This criterion changes only the evaluation target: online preference collection continues to use the finite-temperature policy $\pi_t$.

## 3 TEMPERATURE-ZERO REGRET ANALYSIS OF GREEDY ALIGNMENT METHODS

We now formally introduce the temperature-zero regret criterion and analyze the alignment procedures from Section 2. Throughout this section, all rewards in the fitted class are $\pi_0$-centered, that is, $\mathbb{E}_{A\sim\pi_0(\cdot|x)}[R(x, A)] = 0$ for $d_0$-almost every $x$. This is without loss of generality: subtracting, for each context $x$, an $a$-independent constant from a reward leaves the MNL choice probabilities, the KL-tilted policy, and the temperature-zero selector unchanged; see Lemma 9 of Appendix D. All envelope constants below refer to the centered class; centering a class with envelope $B$ gives an envelope of at most $2B$.

Let $\mathcal{P}$ denote the class of possible truths, where each $p \in \mathcal{P}$ induces a centered measurable reward $R_p : \mathcal{X} \times \mathcal{A} \to \mathbb{R}$, and we let $\mathcal{F}_{\mathcal{P}} := \{R_p : p \in \mathcal{P}\}$. Under truth $p$, preference feedback on a realized slate $\mathbf{a} = (a_1, \dots, a_K) \in \mathcal{A}^K$ follows $P_{R_p}(y = k \mid x, \mathbf{a}) = \exp(R_p(x, a_k))/\sum_{\ell=1}^K \exp(R_p(x, a_\ell))$, for $k = 1, \dots, K$. For $x \in \mathcal{X}$, write $S_x := \mathrm{supp}(\pi_0(\cdot \mid x))$.

### 3.1 TEMPERATURE ZERO REGRET

For each centered reward under consideration, use the standing measurable selector

$$a_R(x) \in \arg\max_{a \in S_x} R(x, a).$$

This tie-breaking convention is used throughout; for truth $p \in \mathcal{P}$, write $a_p(x) := a_{R_p}(x)$. For any centered reward $R$, define the expected *one-step temperature-zero regret* under truth $p$ by

$$\mathcal{G}_p(R) := \mathbb{E}_{X\sim d_0}\Big[R_p\big(X, a_p(X)\big) - R_p\big(X, a_R(X)\big)\Big]. \tag{3}$$

In the offline regime, the learner does not act while the fixed log is being collected. We therefore evaluate offline RLHF by $\mathbb{E}_p[\mathcal{G}_p(\widehat{R}_n^{\mathrm{off}})]$, or equivalently by the sample size needed to make this expected one-step temperature-zero regret at most $\varepsilon$, rather than by cumulative learning regret.

A learning rule $\mathsf{A}$ produces estimates $\widehat{R}_0^{\mathsf{A}}, \widehat{R}_1^{\mathsf{A}}, \widehat{R}_2^{\mathsf{A}}, \dots$, where $\widehat{R}_t^{\mathsf{A}}$ is fitted from the first $t$ online rounds and deployed on round $t+1$. Its expected *cumulative temperature-zero regret* under truth $p \in \mathcal{P}$ is

$$\mathrm{Regret}_{0,p}^{\mathsf{A}}(T) := \sum_{t=0}^{T-1} \mathbb{E}_p\big[\mathcal{G}_p(\widehat{R}_t^{\mathsf{A}})\big]. \tag{4}$$

**cf. KL-regularized regret.** Instead of temperature-zero regret used in this paper, Wu et al. (2025) uses a regret criterion called KL-regularized regret. Fix the same tilt parameter $\eta > 0$ and reference policy $\pi_0$ as in the alignment step. For a policy satisfying $\pi(\cdot \mid x) \ll \pi_0(\cdot \mid x)$, define its true KL-regularized value by

$$\mathcal{J}_{\eta,p}(\pi) \coloneqq \mathbb{E}_{X\sim d_0}\left[\mathbb{E}_{A\sim\pi(\cdot|X)}[R_p(X,A)] - \frac{1}{\eta}\,\mathrm{KL}\big(\pi(\cdot \mid X)\,\|\,\pi_0(\cdot \mid X)\big)\right]. \tag{5}$$

Under the boundedness conditions below, its maximizer is $\pi^\star_{\eta,p} \coloneqq \pi_{R_p}$, obtained from (2) using the true reward. For a bounded candidate reward $R$, define the associated *one-step KL-regularized regret* by

$$\begin{aligned}\mathcal{G}^{\mathrm{KL}}_{\eta,p}(R) &\coloneqq \mathcal{J}_{\eta,p}(\pi^\star_{\eta,p}) - \mathcal{J}_{\eta,p}(\pi_R)\\ &= \frac{1}{\eta}\mathbb{E}_{X\sim d_0}\big[\mathrm{KL}\big(\pi_R(\cdot \mid X)\,\|\,\pi^\star_{\eta,p}(\cdot \mid X)\big)\big].\end{aligned} \tag{6}$$

Indeed, applying (11) with reward $R_p$ and policy $\pi_R$ gives, at each context,

$$\int R_p(x,a)\,\pi_R(da \mid x) - \eta^{-1}\,\mathrm{KL}(\pi_R(\cdot \mid x)\|\pi_0(\cdot \mid x)) = \eta^{-1}\log Z_{R_p}(x) - \eta^{-1}\,\mathrm{KL}(\pi_R(\cdot \mid x)\|\pi^\star_{\eta,p}(\cdot \mid x)).$$

For $\pi_R = \pi^\star_{\eta,p}$, the last KL term is zero. Subtracting the two values and integrating over $d_0$ yields (6). The corresponding expected cumulative regret is

$$\mathrm{Regret}^{\mathsf{A}}_{\mathrm{KL},\eta,p}(T) \coloneqq \sum_{t=0}^{T-1}\mathbb{E}_p\Big[\mathcal{G}^{\mathrm{KL}}_{\eta,p}(\widehat{R}^{\mathsf{A}}_t)\Big], \tag{7}$$

while the offline criterion is $\mathbb{E}_p[\mathcal{G}^{\mathrm{KL}}_{\eta,p}(\widehat{R}^{\mathrm{off}}_n)]$.

The two criteria evaluate different objects produced by the same fit. Temperature-zero regret vanishes once $a_R(X)$ is optimal under $R_p$ almost surely, whereas KL-regularized regret vanishes only when $\pi_R(\cdot \mid X) = \pi^\star_{\eta,p}(\cdot \mid X)$ for $d_0$-almost every context. Thus, a fitted reward can identify the correct best response while still inducing an imperfect finite-temperature distribution. The distinction is between *best-response identification* and *full soft-policy recovery*, not between learning error and an unavoidable randomization cost.

### 3.2 Model class

For a function on the reference support, write $\|R\|_{\infty,\mathrm{supp}(\pi_0)} \coloneqq \sup_{x\in\mathcal{X}}\sup_{a\in S_x}|R(x,a)|$. We use this supportwise supremum norm, rather than an essential supremum over actions, because the evaluation rule maximizes over all of $S_x$. Functions that agree on the reference support are identified when this norm is used. Assume that the class of possible truths $\mathcal{P}$ is nonempty and compact. We impose the following regularity conditions.

(C1) **Compact continuous reward class.** The induced reward class $\mathcal{F}_{\mathcal{P}}$ is compact under $\|\cdot\|_{\infty,\mathrm{supp}(\pi_0)}$. Moreover, for every $R \in \mathcal{F}_{\mathcal{P}}$, the map $a \mapsto R(x,a)$ is continuous on $S_x$ for $d_0$-a.e. $x \in \mathcal{X}$.

(C2) **Reward gap on $\mathcal{P}$.** There exists $\Delta^{\mathcal{P}}_{\min} > 0$ such that for every truth $p \in \mathcal{P}$, $a_p(X)$ is unique and

$$R_p\big(X, a_p(X)\big) - \sup_{a\in S_X,\ a\neq a_p(X)} R_p(X,a) \geq \Delta^{\mathcal{P}}_{\min} \qquad d_0\text{-a.s.}$$

Define also $\Delta^{\mathcal{P}}_{\max} \coloneqq \sup_{p\in\mathcal{P}}\sup_{x\in\mathcal{X}}\sup_{a\in S_x}\big(R_p(x,a_p(x)) - R_p(x,a)\big)$.

Condition (C1) is a compactness-based support-local regularity condition that subsumes the finite-class and bounded-covering-number assumptions commonly used in recent online alignment analyses (Xiong et al., 2024; Ye et al., 2024; Wu et al., 2025). Continuity in Condition (C1) is mild for standard parametric reward classes, including bounded-parameter linear and neural reward families with continuous activations; see Appendix D.1. Condition (C2) is a uniform best-response margin assumption: it restricts attention to settings in which the best response is separated from its supported alternatives by a behaviorally meaningful difference, rather than an arbitrarily small difference that users may not reliably distinguish. In the intended LLM setting, $\mathcal{A}$ is a discrete set of token sequences

rather than a continuous action manifold, so Condition (C2) should be read as a separation condition on supported candidate responses.[3]

For any centered reward $R$, define the truth-centered population loss under truth $p$ by

$$\mathcal{L}_p(R) \coloneqq \mathbb{E}\big[\ell(\mathbf{v}_R, Y)\big], \tag{8}$$

where $X \sim d_0$, $\mathbf{A} = (A_1, \dots, A_K) \sim \pi_0(\cdot \mid X)^{\otimes K}$, $Y \sim P_{R_p}(\cdot \mid X, \mathbf{A})$, and $\mathbf{v}_R \coloneqq (R(X, A_1), \dots, R(X, A_K))$.

For the rest of this section, we instantiate the ERM learners from Section 2 with $\mathcal{F} = \mathcal{F}_\mathcal{P}$. In particular, Condition (C2) applies to every member of the fitted class, not only to the realized truth. For data whose slate coordinates lie in the reference support, Lemma 10 gives

$$|\widehat{\mathcal{L}}_n(R) - \widehat{\mathcal{L}}_n(R')| \le 2\|R - R'\|_{\infty,\mathrm{supp}(\pi_0)}.$$

The empirical objective is therefore continuous on the compact class $\mathcal{F}_\mathcal{P}$ and attains its minimum for every $n \ge 1$. We fix a measurable choice of minimizer. Existence of a minimizer and measurability of its selection are distinct issues; the latter is part of our standing convention.

### 3.3 Bounded cumulative regret in online RLHF

Condition (C1) implies the finite uniform envelope $B_\mathcal{P} \coloneqq \sup_{R \in \mathcal{F}_\mathcal{P}} \|R\|_{\infty,\mathrm{supp}(\pi_0)} < \infty$, and hence $\Delta^\mathcal{P}_{\max} \le 2B_\mathcal{P}$. Since every reward in $\mathcal{F}_\mathcal{P}$ is bounded by $B_\mathcal{P}$ on $\mathrm{supp}(\pi_0)$, Lemma 13 in Appendix D implies $\beta^{-1} \le \frac{d\pi_{t+1}(\cdot|x)}{d\pi_0(\cdot|x)}(a) \le \beta, \beta \coloneqq e^{2\eta B_\mathcal{P}}$, for $d_0$-a.e. $x$ and $\pi_0(\cdot \mid x)$-a.s. $a$.

**Lemma 1** (Automatic isolation of positive regret on $\mathcal{F}_\mathcal{P}$)**.** *Assume (C1)–(C2). Then there exists a constant*

$$\varepsilon^\mathcal{P}_{\mathrm{iso}} > 0$$

*such that for every $p, q \in \mathcal{P}$,*

$$\mathcal{G}_p(R_q) \in \{0\} \cup [\varepsilon^\mathcal{P}_{\mathrm{iso}}, \infty).$$

**Lemma 2** (Zero truth-centered loss identifies the reward on the reference support)**.** *Assume (C1). Then for every truth $p \in \mathcal{P}$ and every $R \in \mathcal{F}_\mathcal{P}$,*

$$\mathcal{L}_p(R) = \mathcal{L}_p(R_p) \quad \Longrightarrow \quad R(x,a) = R_p(x,a) \qquad \forall a \in S_x, \quad d_0\text{-a.e. } x.$$

*In particular,*

$$\mathcal{G}_p(R) = 0.$$

**Lemma 3** (Automatic verification of the truth-centered loss gap)**.** *Assume (C1)–(C2). Fix any $\varepsilon_0 > 0$. Then there exists a finite constant*

$$\gamma_{\mathcal{P},\varepsilon_0} > 0$$

*such that for every pair $p, q \in \mathcal{P}$,*

$$\mathcal{G}_p(R_q) \ge \varepsilon_0 \quad \Longrightarrow \quad \mathcal{L}_p(R_q) - \mathcal{L}_p(R_p) \ge \gamma_{\mathcal{P},\varepsilon_0}.$$

**Theorem 4** (Bounded $O(1)$ regret for online RLHF)**.** *Assume (C1)–(C2). Then*

$$\sup_{p \in \mathcal{P}} \sup_{T \ge 1} \mathrm{Regret}^{\mathrm{ERM}}_{0,p}(T) < \infty.$$

*More precisely, set $\varepsilon_0 \coloneqq \varepsilon^\mathcal{P}_{\mathrm{iso}}$, $\gamma \coloneqq \beta^{-K}\gamma_{\mathcal{P},\varepsilon_0}$, $N_\gamma \coloneqq \mathcal{N}(\mathcal{F}_\mathcal{P}, \gamma/32, \|\cdot\|_{\infty,\mathrm{supp}(\pi_0)})$, and $c_\gamma \coloneqq \gamma^2/[128(\log K + 2B_\mathcal{P})^2]$, where $\mathcal{N}(\mathcal{F}, \varepsilon, \|\cdot\|)$ denotes the smallest cardinality of an $\varepsilon$-net*

[3] This interpretation is consistent with models of imprecise valuation and indifference thresholds, under which sufficiently similar alternatives can be difficult to distinguish (Woodford, 2020; Cantillo et al., 2010). It is also natural under Bradley–Terry: a reward difference of $\Delta$ implies preference probability $\sigma(\Delta)$, which approaches $1/2$ as $\Delta \to 0$. Thus, near-zero reward gaps correspond to nearly indistinguishable preference signals.

*with centers in $\mathcal{F}$, under the indicated norm. Then*

$$\sup_{p\in\mathcal{P}}\sup_{T\geq 1}\mathrm{Regret}_{0,p}^{\mathrm{ERM}}(T)\leq\Delta_{\max}^{\mathcal{P}}\left(1+\left\lceil\frac{1}{c_\gamma}\log(2N_\gamma)\right\rceil+\frac{1}{e^{c_\gamma}-1}\right).$$

### 3.4 Logarithmic accuracy complexity in offline RLHF

The offline result is cleaner than the online result because the logged slates are already sampled from $\pi_0(\cdot\mid X)^{\otimes K}$, which is exactly the slate distribution used in the truth-centered population loss $\mathcal{L}_p$. Therefore the online likelihood-ratio transfer factor $\beta^{-K}$ is not needed.

The fixed-scale concentration statement underlying the theorem is Lemma 20 in Appendix D. We state the logarithmic expected-regret consequence as the main offline result.

**Theorem 5** (Offline logarithmic accuracy complexity)**.** *Assume (C1)–(C2). For each truth $p\in\mathcal{P}$, let the offline data be generated as in Section 2.2 with $P^\star=P_{R_p}$, and let $\widehat{R}_n^{\mathrm{off}}$ be an exact offline ERM over $\mathcal{F}_{\mathcal{P}}$. Set $\varepsilon_0:=\varepsilon_{\mathrm{iso}}^{\mathcal{P}}$, $\gamma_0:=\gamma_{\mathcal{P},\varepsilon_{\mathrm{iso}}^{\mathcal{P}}}$, $N_0^{\mathrm{off}}:=\mathcal{N}(\mathcal{F}_{\mathcal{P}},\gamma_0/32,\|\cdot\|_{\infty,\mathrm{supp}(\pi_0)})$, and $c_0^{\mathrm{off}}:=\gamma_0^2/[128(\log K+2B_{\mathcal{P}})^2]$. Then, for every $n\geq 1$,*

$$\sup_{p\in\mathcal{P}}\mathbb{E}_p\left[\mathcal{G}_p(\widehat{R}_n^{\mathrm{off}})\right]\leq 2\Delta_{\max}^{\mathcal{P}}N_0^{\mathrm{off}}\exp\{-c_0^{\mathrm{off}}n\}.$$

*If $\Delta_{\max}^{\mathcal{P}}=0$, the expected regret is zero for every sample size. Otherwise, for any $\varepsilon>0$, exact reference-logged offline RLHF achieves expected temperature-zero regret at most $\varepsilon$ for every integer sample size satisfying*

$$n\geq\max\left\{1,\left\lceil\frac{1}{c_0^{\mathrm{off}}}\log\left(\frac{2\Delta_{\max}^{\mathcal{P}}N_0^{\mathrm{off}}}{\varepsilon}\right)\right\rceil\right\}.$$

*For a fixed model class, the offline accuracy complexity is therefore $O(\log(1/\varepsilon))$ as $\varepsilon\downarrow 0$.*

## 4 Experiments

To isolate the paper's main point, we use the linear BT testbed of Wu et al. (2025) while changing the evaluation target: instead of evaluating full soft-policy recovery, we evaluate the temperature-zero action induced by the learned reward estimate. This directly illustrates the paper's claim that continued KL-regularized regret need not imply continued error in best-response identification.

### 4.1 Shared testbed and evaluation

Both experiments use the same finite-action linear BT model, so their results concern the same best-response identification task. To mirror the positive-gap requirement in Condition (C2) in the simulation, we first draw a fixed bank $\mathcal{X}_{\mathrm{sim}}$ of 20,000 contexts independently from $\mathrm{Unif}([0,1]^5)$ and define the simulation context distribution $d_0$ to be uniform on this finite bank. A fixed set of six actions $\mathcal{A}=\{a_1,\dots,a_6\}\subset[0,1]^5$ and a ground-truth matrix $W^\star\in[0,1]^{5\times 5}$ define the reward

$$R^\star(x,a)=x^\top W^\star a.$$

The reference policy $\pi_0$ is uniform over the six actions. Given a context and a pair $(a_1,a_2)$, the first response is preferred with probability $\sigma(R^\star(x,a_1)-R^\star(x,a_2))$, where $\sigma(z)=(1+e^{-z})^{-1}$. We fit a reward of the same bilinear form by numerical BT maximum likelihood with box constraints $\widehat{W}\in[0,1]^{5\times 5}$, starting from $\widehat{W}_0=0$.

Our common evaluation target is the temperature-zero regret of the fitted selector. For a fitted matrix $\widehat{W}$, define $\widehat{a}(x)=\arg\max_{a\in\mathcal{A}}x^\top\widehat{W}a$ and evaluate

$$\mathbb{E}_X[R^\star(X,a^\star(X))-R^\star(X,\widehat{a}(X))],$$

where $a^\star(x)$ maximizes the true reward. We estimate this expectation on fresh draws from $d_0$, independent of the preference data, and report means and pointwise standard errors across independent

repeats. The online experiment additionally accumulates these one-step regrets over rounds; the offline experiment relates them to the size of a fixed reference-policy log. The selected problem instance has minimum true top-two gap $0.2703$ over the entire support $\mathcal{X}_{\mathrm{sim}}$, so the experimental truth satisfies the positive-gap condition used in Condition (C2). This verifies the margin for the selected truth, not Condition (C2) for the entire fitted box-constrained class: that class contains the zero matrix, which ties all actions. Moreover, the numerical optimizer is not an exact ERM oracle. The experiments therefore illustrate the qualitative prediction rather than verify all the assumptions of the theorems. Appendix B gives the instance-selection procedure and the separate sampling and optimization settings.

### 4.2 Greedy online alignment

The online experiment asks whether best-response identification can stabilize while feedback collection remains randomized. We initialize $\pi_1 = \pi_0$. At round $t$, both responses are sampled independently from the current KL-tilted policy

$$\pi_t(a \mid x_t) \propto \pi_0(a) \exp\big(\eta \widehat{R}_{t-1}(x_t, a)\big), \qquad \widehat{R}_{t-1}(x, a) = x^\top \widehat{W}_{t-1} a,$$

exactly as in the online procedure of Section 2.2. After observing the BT preference label, we refit the reward on all comparisons collected so far. We use $\eta \in \{1, 2, 3\}$, a horizon of $T = 200$, and $50$ independent trajectories per value of $\eta$. Each fitted selector is evaluated on a fresh batch of $4096$ contexts at every iteration.

Figure 1 illustrates the same qualitative pattern, uniformly across all three regularization levels: the one-step temperature-zero regret falls rapidly to essentially zero well before round 100, and the cumulative regret curves flatten after roughly the first 100 rounds. This is consistent with the bounded $O(1)$ regret behavior predicted by Theorem 4: once the learned reward estimate induces the correct top-ranked action in almost all contexts, the remaining temperature-zero regret becomes zero, so the cumulative curve plateaus rather than continuing to grow throughout the horizon. Note that the online DPO curves are visually indistinguishable from the online RLHF curves at each $\eta$, exactly as predicted by the policy-space equivalence in the pairwise setting discussed earlier.

It is useful to compare this with the results reported in Wu et al. (2025), which studies greedy online RLHF under KL-regularized regret. Their theory yields logarithmic KL-regularized cumulative regret. The comparison highlights our measurement point: the best-response selector can stabilize even while KL-regularized regret continues to evaluate discrepancies in the full randomized policy.

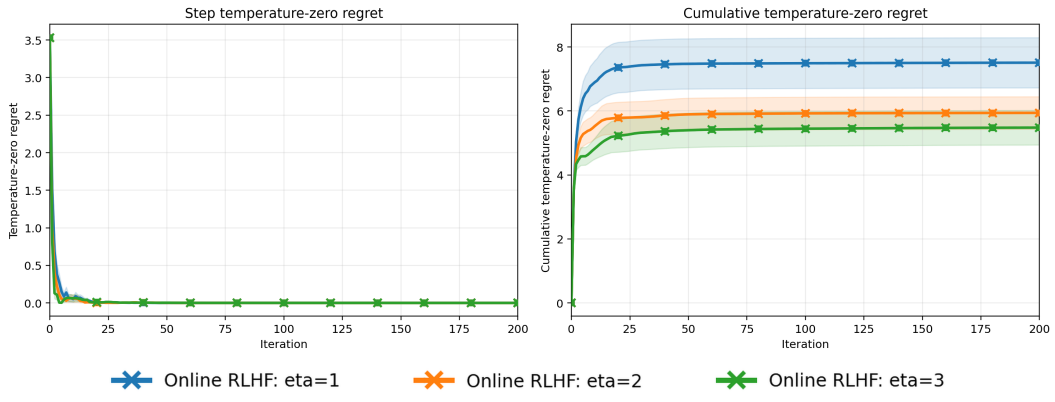


Figure 1: Greedy online RLHF in the finite-action linear Bradley–Terry experiment of Wu et al. (2025), evaluated using temperature-zero regret. Curves show the mean over $50$ repeated runs, and shaded bands denote the standard error. For each $\eta \in \{1, 2, 3\}$, one-step temperature-zero regret (left) rapidly approaches zero and cumulative temperature-zero regret (right) levels off.

### 4.3 Offline AI alignment

The offline experiment asks whether the same best-response identification task can be learned rapidly from a fixed reference-policy log. Unlike the online experiment, responses are sampled independently

from the reference policy $\pi_0$, so the distribution of collected comparisons does not depend on the fitted reward. We repeated the experiment 200 times; in each repeat, we evaluate all fitted selectors on the same independent batch of 16,384 contexts. Figure 2 reports the resulting mean temperature-zero regret for $n \leq 1000$.

Figure 2 illustrates the qualitative pattern predicted by Theorem 5: temperature-zero regret falls rapidly toward zero as the size of the fixed comparison log increases. To compare the predicted accuracy-complexity forms, we fit exponential decay $Ae^{-cn}$ and inverse decay $A/n$ to the observed mean regret. These correspond, respectively, to $O(\log(1/\epsilon))$ and $O(1/\epsilon)$ sample complexity for reaching target regret $\epsilon$. Over $n \leq 1000$, the exponential specification achieves log-space $R^2 = 0.875$, compared with 0.656 for inverse decay, and reduces log-space RMSE from 0.793 to 0.478, a 39.7% reduction. Thus, over this finite-sample range, the observed decline in temperature-zero regret is more closely described by exponential than inverse decay under the chosen fitting criterion.

It is useful to compare this pattern with the standard inverse accuracy dependence associated with KL-regularized offline alignment guarantees (Wu et al., 2025). Theorem 5 predicts a sharper result for the decision-relevant task studied here: because temperature-zero regret only requires identifying the correct top-ranked response, expected regret can decrease exponentially with the number of reference-policy comparisons, yielding $O(\log(1/\epsilon))$ rather than $O(1/\epsilon)$ accuracy complexity. The simulation is consistent with this distinction: the fitted reward can identify the best response rapidly even though recovering the full finite-temperature aligned policy is a more demanding objective.

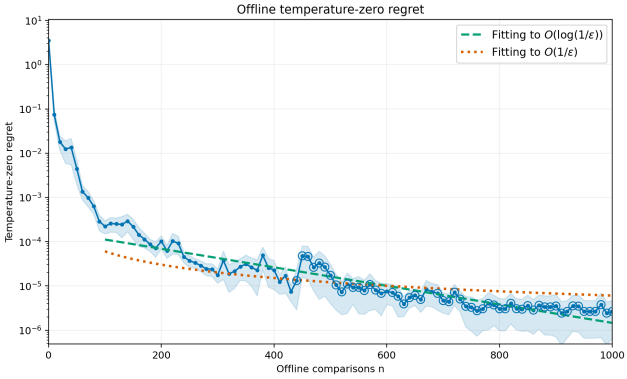


| Fitted complexity | Log-space $R^2 \uparrow$ | Log-space RMSE $\downarrow$ |
|---|---|---|
| $O(\log(1/\epsilon))$ | **0.875** | **0.478** |
| $O(1/\epsilon)$ | 0.656 | 0.793 |

Figure 2: Offline temperature-zero regret and fitted accuracy complexity. The blue curve is the mean over 200 independent offline logs; shading denotes one standard error across repeats. Open circles mark sample sizes with fewer than 10 positive-regret repeats.

## 5 Conclusion

We show that the apparent efficiency of greedy AI alignment is much less mysterious once we distinguish two learning problems that standard KL-regularized regret combines: learning a reward model and recovering the entire randomized aligned policy. Under realizability, compactness, and a uniform best-response margin over the fitted reward class, offline RLHF identifies the optimal response with exponentially decreasing expected temperature-zero regret, while greedy online RLHF incurs only bounded cumulative temperature-zero regret. Thus, the reward model can effectively finish the decision-relevant task of learning what to choose even while the KL-regularized policy continues to adjust probabilities over non-selected responses. Existing logarithmic policy-regret guarantees therefore do not show that greedy alignment must remain slow at learning the best response; they partly reflect the more demanding objective of recovering the full soft policy. What can look like the unreasonable effectiveness of greedy alignment is, under the decision-relevant metric studied here, a consequence of best-response identification being substantially easier than full policy recovery.

### AI use statement

Generative AI tools were used to assist with manuscript organization and editing, expansion and checking of mathematical proofs, LaTeX formatting, figure-legend editing, and document preparation.

## REFERENCES


Víctor Cantillo, Johanna Amaya, and J de D Ortúzar. Thresholds and indifference in stated choice surveys. *Transportation Research Part B: Methodological*, 44(6):753–763, 2010.

Yurong Chen, Yu He, Michael I. Jordan, and Fan Yao. How sampling shapes LLM alignment: From one-shot optima to iterative dynamics, 2026. URL https://arxiv.org/abs/2602.12180.

Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.

Aniket Anand Deshmukh, Srinagesh Sharma, James W. Cutler, Mark Moldwin, and Clayton Scott. Simple regret minimization for contextual bandits. *arXiv preprint arXiv:1810.07371*, 2018.

Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Remi Munos, Mark Rowland, Michal Valko, and Daniele Calandriello. A general theoretical paradigm to understand learning from human preferences. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pages 4447–4455. PMLR, 2024. URL https://proceedings.mlr.press/v238/gheshlaghi-azar24a.html.

Shangmin Guo, Biao Zhang, Tianlin Liu, Tianqi Liu, Misha Khalman, Felipe Llinares, Alexandre Rame, Thomas Mesnard, Yao Zhao, Bilal Piot, et al. Direct language model alignment from online ai feedback. *arXiv preprint arXiv:2402.04792*, 2024.

Audrey Huang, Adam Block, Qinghua Liu, Nan Jiang, Akshay Krishnamurthy, and Dylan J. Foster. Is best-of-$n$ the best of them? coverage, scaling, and optimality in inference-time alignment. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, 2025. URL https://arxiv.org/abs/2503.21878.

Juno Kim, Jihun Yun, Jason D. Lee, and Kwang-Sung Jun. Coverage improvement and fast convergence of on-policy preference learning, 2026. URL https://arxiv.org/abs/2601.08421.

Sanath Kumar Krishnamurthy, Ruohan Zhan, Susan Athey, and Emma Brunskill. Proportional response: Contextual bandits for simple and cumulative regret minimization. *arXiv preprint arXiv:2307.02108*, 2023.

Junghyun Lee, Minju Hong, Kwang-Sung Jun, Chulhee Yun, and Se-Young Yun. Regularized online RLHF with generalized bilinear preferences, 2026. URL https://arxiv.org/abs/2602.23116.

Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*, 2021.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *Advances in Neural Information Processing Systems*, volume 36, 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/a85b405ed65c6477a4fe8302b5e06ce7-Paper-Conference.pdf.

Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters, 2024. *URL https://arxiv. org/abs/2408.03314*, 20:5, 2024.

Yuanhao Wang, Qinghua Liu, and Chi Jin. Is RLHF more difficult than standard RL? a theoretical perspective. In *Advances in Neural Information Processing Systems*, volume 36, 2023. URL https://papers.nips.cc/paper_files/paper/2023/hash/efb9629755e598c4f261c44aeb6fde5e-Paper-Conference.pdf.

Michael Woodford. Modeling imprecision in perception, valuation, and choice. *Annual Review of Economics*, 12(1):579–601, 2020.

Di Wu, Chengshuai Shi, Jing Yang, and Cong Shen. Greedy sampling is provably efficient for RLHF. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=jQH0gdwsuT.

Wei Xiong, Hanze Dong, Chenlu Ye, Ziqi Wang, Han Zhong, Heng Ji, Nan Jiang, and Tong Zhang. Iterative preference learning from human feedback: Bridging theory and practice for RLHF under KL-constraint. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 54715–54754. PMLR, 2024. URL https://proceedings.mlr.press/v235/xiong24a.html.

Yichong Xu, Ruosong Wang, Lin F. Yang, Aarti Singh, and Artur Dubrawski. Preference-based reinforcement learning with finite-time guarantees. In *Advances in Neural Information Processing Systems*, volume 33, 2020. URL https://papers.nips.cc/paper_files/paper/2020/hash/d9d3837ee7981e8c064774da6cdd98bf-Paper.pdf.

Chenlu Ye, Wei Xiong, Yuheng Zhang, Hanze Dong, Nan Jiang, and Tong Zhang. Online iterative reinforcement learning from human feedback with general preference model. In *Advances in Neural Information Processing Systems*, volume 37, pages 81773–81807, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/94d13c2401fe119e57ba325b6fe526e0-Abstract-Conference.html.

Heyang Zhao, Chenlu Ye, Quanquan Gu, and Tong Zhang. Sharp analysis for KL-regularized contextual bandits and RLHF, 2024. URL https://arxiv.org/abs/2411.04625.

Heyang Zhao, Chenlu Ye, Wei Xiong, Quanquan Gu, and Tong Zhang. Logarithmic regret for online KL-regularized reinforcement learning, 2025a. URL https://arxiv.org/abs/2502.07460.

Qingyue Zhao, Kaixuan Ji, Heyang Zhao, Tong Zhang, and Quanquan Gu. Nearly optimal sample complexity of offline KL-regularized contextual bandits under single-policy concentrability, 2025b. URL https://arxiv.org/abs/2502.06051v1.

Banghua Zhu, Michael Jordan, and Jiantao Jiao. Principled reinforcement learning with human feedback from pairwise or $K$-wise comparisons. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 43037–43067. PMLR, 2023. URL https://proceedings.mlr.press/v202/zhu23f.html.

## A Related works

**Preference-based alignment.** Early theory for RLHF established statistical guarantees for learning from pairwise and $K$-wise comparisons and clarified the role of reward-model estimation in preference-based control (Zhu et al., 2023); adjacent work in preference-based RL studied finite-time guarantees and compared the intrinsic difficulty of RLHF to standard reward-based RL (Xu et al., 2020; Wang et al., 2023). A complementary unifying view is given by ΨPO, which treats RLHF, DPO, and identity-based objectives as special cases of a broader preference-learning framework and makes explicit when scalar-reward reductions are or are not faithful to the underlying preference problem (Gheshlaghi Azar et al., 2024).

**KL-regularized iterative RLHF.** A recent theory line studies the same KL-regularized contextual-bandit / RLHF objective considered here: Xiong et al. (2024) analyze offline, online, and hybrid iterative preference learning under KL constraints, and Ye et al. (2024) extend the online analysis beyond Bradley–Terry to general preference oracles. Subsequent sharp analyses show that KL regularization can fundamentally change the local statistical difficulty, yielding fast rates in offline and online settings under suitable assumptions (Zhao et al., 2024; 2025a;b).

**Greedy and on-policy preference learning.** Most closely related, Wu et al. (2025) show that greedy sampling itself can attain the same $O(\log T)$ order as confidence-based methods in RLHF, including under general preference models, so explicit optimism or pessimism is not always necessary. Kim et al. (2026), Chen et al. (2026), and Lee et al. (2026) refine this picture by analyzing coverage improvement in on-policy DPO, the effect of sampling and reference choices on iterative dynamics, and greedy learning under generalized bilinear or potentially intransitive preference models with strongly convex regularization.

**Simple-regret and deployment-centric objectives.** Methodologically, our temperature-zero criterion is closer to simple-regret evaluation than to standard finite-temperature policy regret, because it scores the final selected action rather than the full stochastic policy used during data collection. This objective has been studied in contextual bandits both in pure-exploration / exploitation formulations and in more recent algorithms that explicitly trade off cumulative and simple regret (Deshmukh et al., 2018; Krishnamurthy et al., 2023). Our contribution is therefore not to introduce simple regret itself, but to characterize this decision-centric target for the same greedy KL-regularized preference-learning process used in RLHF/DPO, including adaptive online data collection, and to contrast its rate with full soft-policy recovery. Recent inference-time alignment work similarly studies the top-of-ranking decision more directly, for example by analyzing best-of-$N$ scaling, coverage, and reward hacking for response selection with a fixed reward model (Huang et al., 2025).

## B Experiment reproducibility details

This appendix records the shared problem instance and the separate online and offline runs in Section 4. The online artifacts correspond to running the attached paper's code.py file:

```
python Experiments/code.py \
  --output-dir \
  Experiments/horizon200_gapforced_0p2_seed3500_rep50 \
  --seed 3500 \
  --dimension 5 \
  --num-actions 6 \
  --horizon 200 \
  --repeats 50 \
  --eval-contexts 4096 \
  --mle-maxiter 50 \
  --mle-ftol 1e-9 \
  --min-probe-gap 0.2 \
  --gap-probe-contexts 20000 \
  --problem-search-limit 1000
```

**Problem-instance selection.** Using the base seed 3500, the script first draws a fixed bank $\mathcal{X}_{\mathrm{sim}}$ of 20,000 contexts from $\mathrm{Unif}([0,1]^5)$. This bank is the support of the simulation context distribution $d_0$, from which training and evaluation contexts are sampled uniformly with replacement. The script then scans candidate problem seeds $3500, 3501, \ldots$ until it finds an instance whose minimum true top-two reward gap over this entire support is at least 0.2. For the run used in Figure 1, the accepted candidate is the 47-th one, with chosen problem seed 3546. Its minimum support-wide gap is 0.2703, and the mean support-wide gap is 1.9018. Thus the simulated truth has a uniform positive best-response gap on the support of $d_0$.

**Online data collection and fitting.** The experiment uses $k = 5$-dimensional contexts, $|\mathcal{A}| = 6$ actions, horizon $T = 200$, and regularization coefficients $\eta \in \{1, 2, 3\}$. Each context $x_t$ is drawn independently from the finite-support distribution $d_0$ defined above. Both actions are sampled independently from the current KL-tilted policy, exactly as in the theoretical online procedure, and the binary preference is then drawn from the Bradley–Terry probability

$$\sigma\big(x_t^\top W^\star (a_{t,1} - a_{t,2})\big).$$

After each round, the estimate $\widehat{W}_t$ is refit on all comparisons collected so far by bounded L-BFGS-B maximum likelihood on the 25-dimensional parameter vector, initialized from the previous iterate, with box constraints $[0, 1]$, `maxiter`$= 50$, and `ftol`$= 10^{-9}$.

**Evaluation and uncertainty reporting.** For each $\eta$, we run 50 independent trajectories. At every iteration $t$, the one-step temperature-zero regret is estimated on a fresh Monte Carlo batch of 4096 evaluation contexts sampled from $d_0$, independently of the online trajectory. Figure 1 plots the mean across the 50 trajectories, and the shaded bands are pointwise standard errors, computed as the sample standard deviation divided by $\sqrt{50}$. Appendix C.3 presents the equivalent pairwise policy-space interpretation of these same trajectories.

**Offline data collection, fitting, and evaluation.** The offline run uses the same saved problem instance and base simulation seed 4500. For each of 200 repeats, we generate 5000 independent reference-policy comparisons and fit only the prefixes $n = 0, 10, \ldots, 2000$; generating the full log preserves the random streams used in the extended simulation. Both actions are sampled independently and uniformly, and preferences follow the same BT law as above. Each fit uses bounded L-BFGS-B with `maxiter`$= 1000$, `ftol`$= 10^{-12}$, and `gtol`$= 10^{-8}$, initialized from the preceding fitted prefix. At $n = 0$, the reward matrix is zero. Evaluation uses 16,384 independent draws from $d_0$ per repeat, shared across prefixes. All 40,000 nonempty-prefix optimizations reported convergence. The run is reproduced from the project root by

```
python offline_experiment.py \
  --seed 4500 --repeats 200 --eval-contexts 16384 \
  --sample-step 10 --max-samples 2000 --data-horizon 5000 \
  --output-dir Experiments/offline_dense_n0_2000_step10_seed4500_rep200
```

Figure 2 uses the first 1000 comparisons and fits both decay curves on $100 \le n \le 1000$. Their fitted forms are $1.81386 \times 10^{-4} \exp(-0.00481934n)$ and $0.00605738/n$. The reported scores are computed in natural-log regret on these same 91 points; all their mean regrets are positive, and no positive-repeat threshold is applied to either fit.

# C Online DPO as an equivalent pairwise formulation

This appendix collects the exact pairwise DPO formulation, its bounded-regret corollary, and the corresponding simulation overlay. The optimization is over the policies induced by the same reward class used in the main text; no separate regret argument is needed.

## C.1 Equivalent policy-space view: greedy online DPO

When the feedback is pairwise ($K = 2$), the same greedy loop can be written directly in policy coordinates in the style of direct preference optimization (DPO) (Guo et al., 2024). For the regular

reward class $\mathcal{F} = \mathcal{F}_{\mathcal{P}}$ used in the theoretical results, define the induced policy class

$$\Pi_{\mathcal{F}}^{\mathrm{DPO}} := \{\pi_R : R \in \mathcal{F}\},$$

where $\pi_R$ is the KL-tilted policy from (2).

Given pairwise data

$$(x_s, a_{s,1}, a_{s,2}, y_s), \qquad y_s \in \{1, 2\},$$

define the exact online DPO empirical loss of a policy $\pi \in \Pi_{\mathcal{F}}^{\mathrm{DPO}}$ by

$$\widehat{\mathcal{L}}_t^{\mathrm{DPO}}(\pi) := \frac{1}{t}\sum_{s=1}^{t} -\log\sigma\left(\frac{1}{\eta}\left[\log\frac{d\pi(\cdot \mid x_s)}{d\pi_0(\cdot \mid x_s)}(a_{s,y_s}) - \log\frac{d\pi(\cdot \mid x_s)}{d\pi_0(\cdot \mid x_s)}(a_{s,3-y_s})\right]\right), \tag{9}$$

where $\sigma(u) := (1 + e^{-u})^{-1}$. Initialize $\widehat{\pi}_0 = \pi_0$, collect pairwise feedback under the same online sampling rule as in the main text, and refit using all accumulated comparisons. The exact online DPO update is

$$\widehat{\pi}_t \in \operatorname*{arg\,min}_{\pi \in \Pi_{\mathcal{F}}^{\mathrm{DPO}}} \widehat{\mathcal{L}}_t^{\mathrm{DPO}}(\pi), \qquad \pi_{t+1} \leftarrow \widehat{\pi}_t. \tag{10}$$

Under exact optimization over this induced policy class, this is a reparameterization of the reward-space update. We evaluate density ratios using the canonical version $d\pi_R/d\pi_0 = e^{\eta R}/Z_R$ on $S_x$, where $Z_R(x) = \int e^{\eta R(x,a')}\,\pi_0(da' \mid x)$. Thus, at every supported action, this version satisfies

$$\frac{1}{\eta}\log\frac{d\pi_R(\cdot \mid x)}{d\pi_0(\cdot \mid x)}(a) = R(x,a) - \frac{1}{\eta}\log\int_{\mathcal{A}} e^{\eta R(x,a')}\,\pi_0(da' \mid x).$$

Hence the normalization term cancels in pairwise differences, and for every realized example $(x_s, a_{s,1}, a_{s,2}, y_s)$,

$$\frac{1}{\eta}\left[\log\frac{d\pi_R(\cdot \mid x_s)}{d\pi_0(\cdot \mid x_s)}(a_{s,y_s}) - \log\frac{d\pi_R(\cdot \mid x_s)}{d\pi_0(\cdot \mid x_s)}(a_{s,3-y_s})\right] = R(x_s, a_{s,y_s}) - R(x_s, a_{s,3-y_s}).$$

Therefore

$$-\log\sigma\big(R(x_s, a_{s,y_s}) - R(x_s, a_{s,3-y_s})\big) = \log\Big(e^{R(x_s,a_{s,1})} + e^{R(x_s,a_{s,2})}\Big) - R(x_s, a_{s,y_s}),$$

so

$$\widehat{\mathcal{L}}_t^{\mathrm{DPO}}(\pi_R) = \widehat{\mathcal{L}}_t(R) \qquad \text{when } K = 2.$$

**Scope of the equivalence.** The equality of empirical losses holds for every realized comparison history. With common data and a matched measurable choice of exact minimizer, the two formulations also have identical policy iterates and temperature-zero selectors. This is an equivalence between exact optimization problems over matched classes, not a claim that arbitrary approximate implementations return identical iterates.

### C.2 Bounded temperature-zero regret as a corollary

Specialize to pairwise feedback $K = 2$ and set $\mathcal{F} = \mathcal{F}_{\mathcal{P}}$. Write

$$\Pi_{\mathcal{P}}^{\mathrm{DPO}} := \Pi_{\mathcal{F}_{\mathcal{P}}}^{\mathrm{DPO}}, \qquad \Pi_{\mathcal{P}}^{\mathrm{DPO},0} := \Pi_{\mathcal{P}}^{\mathrm{DPO}} \cup \{\pi_0\}.$$

For $\pi \in \Pi_{\mathcal{P}}^{\mathrm{DPO}}$, choose any $R \in \mathcal{F}_{\mathcal{P}}$ such that $\pi = \pi_R$ and set $a_\pi(x) := a_R(x)$. This is well defined $d_0$-a.s. because any two such rewards differ on $S_x$ only by an action-independent constant, as shown in the proof below. For the initial policy $\pi_0$, use $a_{\pi_0}(x) := a_0(x)$, with the same measurable tie-breaking convention as in the main text. Define

$$\mathcal{G}_p^{\mathrm{DPO}}(\pi) := \mathbb{E}_{X \sim d_0}\left[R_p(X, a_p(X)) - R_p(X, a_\pi(X))\right].$$

Let $(\widehat{\pi}_t)_{t \geq 0}$ be the exact online DPO iterates from (10), initialized at $\widehat{\pi}_0 = \pi_0$. Use a measurable representing-reward rule, so that for $t \geq 1$, $\widehat{\pi}_t = \pi_{\widehat{R}_t^{\mathrm{DPO}}}$ with $\widehat{R}_t^{\mathrm{DPO}} \in \mathcal{F}_{\mathcal{P}}$. A matched measurable reward-space ERM selection, as in the preceding subsection, provides such a rule.

**Corollary 6** (Bounded temperature-zero regret for online DPO)**.** *Assume (C1)–(C2), and specialize to pairwise feedback $K = 2$. Then*

$$\sup_{p\in\mathcal{P}}\sup_{T\geq 1}\sum_{t=0}^{T-1}\mathbb{E}_p\big[\mathcal{G}_p^{\mathrm{DPO}}(\widehat{\pi}_t)\big] < \infty.$$

*More precisely, with the same constants $\varepsilon_0, \gamma_{\mathcal{P},\varepsilon_0}, \gamma, N_\gamma, c_\gamma$ as in Theorem 4, specialized to $K = 2$,*

$$\sup_{p\in\mathcal{P}}\sup_{T\geq 1}\sum_{t=0}^{T-1}\mathbb{E}_p\big[\mathcal{G}_p^{\mathrm{DPO}}(\widehat{\pi}_t)\big] \leq \Delta_{\max}^{\mathcal{P}}\left(1+\left\lceil\frac{1}{c_\gamma}\log(2N_\gamma)\right\rceil+\frac{1}{e^{c_\gamma}-1}\right).$$

*Proof of Corollary 6.* Work on the common full-measure context set supplied by Lemma 15. Suppose $R, R' \in \mathcal{F}_{\mathcal{P}}$ induce the same policy at a context in this set. Their densities with respect to $\pi_0(\cdot \mid x)$ agree almost surely, and therefore

$$\frac{e^{\eta R(x,a)}}{Z_R(x)} = \frac{e^{\eta R'(x,a)}}{Z_{R'}(x)} \quad \text{for } \pi_0(\cdot \mid x)\text{-almost every } a.$$

Both sides are positive, so taking logarithms and rearranging gives

$$R(x,a) - R'(x,a) = \frac{\log Z_R(x) - \log Z_{R'}(x)}{\eta} \quad \pi_0(\cdot \mid x)\text{-almost surely.}$$

The right-hand side is constant in $a$. Integrating and using centering of both rewards shows that this constant is zero. Their continuous action sections then agree on all of $S_x$, by Lemma 14. Thus the induced selector is independent of the representing reward, up to a $d_0$-null set. The same argument with $R' = 0$ shows that if $\pi_0$ already belongs to the induced policy class, its representing reward is zero on the support at almost every context. Its selector then agrees with the initial selector $a_0$ under the common tie-breaking convention. The definition on the union with $\{\pi_0\}$ is therefore consistent.

For any supported pair and observed preferred index $y \in \{1, 2\}$, the canonical densities satisfy

$$\begin{aligned}\frac{1}{\eta}\left[\log\frac{d\pi_R}{d\pi_0}(a_y \mid x) - \log\frac{d\pi_R}{d\pi_0}(a_{3-y} \mid x)\right] &= \frac{1}{\eta}\big[\eta R(x,a_y) - \log Z_R(x) - \eta R(x,a_{3-y}) + \log Z_R(x)\big] \\ &= R(x,a_y) - R(x,a_{3-y}).\end{aligned}$$

Moreover,

$$\begin{aligned}-\log\sigma(R(x,a_y) - R(x,a_{3-y})) &= \log\Big(1 + e^{R(x,a_{3-y})-R(x,a_y)}\Big) \\ &= \log\Big(e^{R(x,a_y)} + e^{R(x,a_{3-y})}\Big) - R(x,a_y) \\ &= \ell(\mathbf{v}_R(x,(a_1,a_2)), y).\end{aligned}$$

Summing over observations proves $\widehat{\mathcal{L}}_t^{\mathrm{DPO}}(\pi_R) = \widehat{\mathcal{L}}_t(R)$ simultaneously over the fitted class on an almost-sure set of histories. The common context set ensures that this statement is not affected by candidate-dependent exceptional contexts.

It follows in both directions that the minimizer sets correspond:

$$\operatorname*{arg\,min}_{\pi\in\Pi_{\mathcal{P}}^{\mathrm{DPO}}}\widehat{\mathcal{L}}_t^{\mathrm{DPO}}(\pi) = \left\{\pi_R : R \in \operatorname*{arg\,min}_{U\in\mathcal{F}_{\mathcal{P}}}\widehat{\mathcal{L}}_t(U)\right\}.$$

Indeed, a reward minimizer has no policy competitor with a smaller loss, because every policy competitor has a representing reward with the same loss. Conversely, if an exact policy minimizer has a representing reward $R$ that is not an empirical reward minimizer, a reward with strictly smaller empirical loss would induce a policy with strictly smaller DPO loss, a contradiction.

Consequently the measurable representing rewards $\widehat{R}_t^{\mathrm{DPO}}$ form an exact ERM sequence. Their sampling policies are $\pi_1 = \pi_0$ and $\pi_{t+1} = \pi_{\widehat{R}_t^{\mathrm{DPO}}}$, so this sequence satisfies the online RLHF assumptions. Also, including the initialization,

$$\mathcal{G}_p^{\mathrm{DPO}}(\widehat{\pi}_t) = \mathcal{G}_p(\widehat{R}_t^{\mathrm{DPO}}), \qquad \widehat{R}_0^{\mathrm{DPO}} = 0.$$

Applying Theorem 4 to this exact ERM rule, with $K = 2$, gives the stated bound after summing over $t$. If the same minimizer choices and sampling randomness are used in both formulations, equality of their policies holds initially and, by the loss identity, is preserved after every common observation; induction then gives identical trajectories. Such matched choices are needed for equality of particular trajectories, but the regret bound holds for any measurable representing-reward rule satisfying the exact optimization requirement. □

### C.3 Simulation overlay and interpretation

Figure 3 retains the paired RLHF–DPO presentation of the simulation used in Section 4. In this pairwise implementation, the DPO overlay uses exactly the same trajectory data as the RLHF curves. The reported maximum absolute RLHF-versus-DPO difference is therefore zero for both the one-step and cumulative curves at every $\eta$, by construction. The overlay illustrates the equivalent policy-space representation; it is not a separate optimization experiment or an independent empirical comparison. Figure 1 in the main text labels each shared trajectory only as online RLHF.

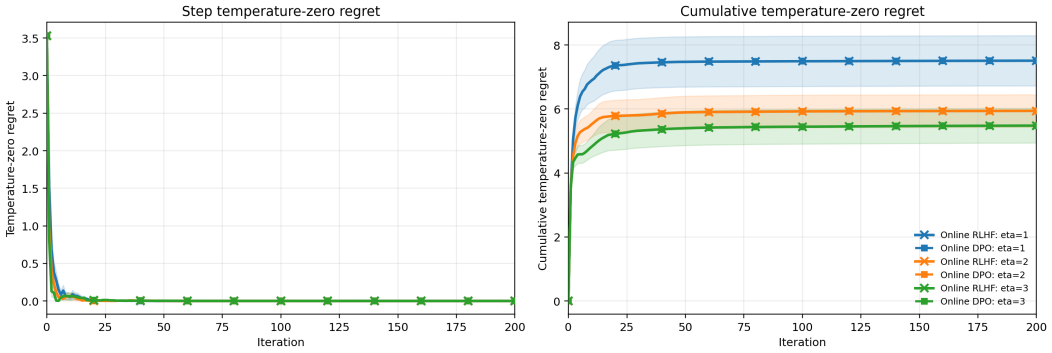


Figure 3: Online RLHF and its equivalent pairwise DPO representation for the simulation in Section 4. Both labels use the same trajectory data, so the curves coincide by construction at each $\eta \in \{1, 2, 3\}$. Curves show the mean over 50 repeated runs, and shaded bands denote the standard error. The left panel shows one-step temperature-zero regret; the right panel shows cumulative temperature-zero regret.

## D Proofs for Section 3

We first record examples of reward families satisfying Condition (C1).

### D.1 Linear and neural reward families

Write

$$S_x := \mathrm{supp}(\pi_0(\cdot \mid x)), \qquad S_{\pi_0} := \{(x, a) \in \mathcal{X} \times \mathcal{A} : \ a \in S_x\}.$$

**Proposition 7** (Centered linear reward classes satisfy Condition (C1))**.** *Let $\Theta \subset \mathbb{R}^d$ be compact. Let $\phi : \mathcal{X} \times \mathcal{A} \to \mathbb{R}^d$ be measurable and satisfy*

$$M_\phi := \sup_{(x,a) \in S_{\pi_0}} \|\phi(x, a)\|_2 < \infty.$$

*Assume also that for $d_0$-a.e. $x$, the map $a \mapsto \phi(x, a)$ is continuous on $S_x$. For $\theta \in \Theta$, define the centered linear reward*

$$R_\theta(x, a) := \theta^\top \phi(x, a) - \mathbb{E}_{A \sim \pi_0(\cdot|x)}[\theta^\top \phi(x, A)].$$

*Then the class $\mathcal{F}_\Theta^{\mathrm{lin}} := \{R_\theta : \theta \in \Theta\}$ is compact under $\|\cdot\|_{\infty,\mathrm{supp}(\pi_0)}$, and for every $\theta \in \Theta$, the map $a \mapsto R_\theta(x, a)$ is continuous on $S_x$ for $d_0$-a.e. $x$.*

*Proof.* Define the reference feature mean

$$\bar{\phi}(x) \coloneqq \int_{\mathcal{A}} \phi(x,a)\,\pi_0(da \mid x).$$

Each coordinate is integrable because $\|\phi(x,a)\|_2 \le M_\phi$ on the reference support, which has full reference measure by Lemma 14. The kernel integral is measurable, and

$$\|\bar{\phi}(x)\|_2 \le \int_{\mathcal{A}} \|\phi(x,a)\|_2\,\pi_0(da \mid x) \le M_\phi.$$

Consequently,

$$R_\theta(x,a) = \theta^\top\big(\phi(x,a) - \bar{\phi}(x)\big), \qquad \int R_\theta(x,a)\,\pi_0(da \mid x) = 0.$$

For every context at which the action section of $\phi$ is continuous, the action section of $R_\theta$ is continuous, since the vector $\bar{\phi}(x)$ does not depend on $a$.

For $\theta, \tilde{\theta} \in \Theta$ and $(x,a) \in S_{\pi_0}$, the Cauchy–Schwarz and triangle inequalities give

$$\begin{aligned} |R_\theta(x,a) - R_{\tilde{\theta}}(x,a)| &= \Big|(\theta - \tilde{\theta})^\top(\phi(x,a) - \bar{\phi}(x))\Big| \\ &\le \|\theta - \tilde{\theta}\|_2\big(\|\phi(x,a)\|_2 + \|\bar{\phi}(x)\|_2\big) \\ &\le 2M_\phi\|\theta - \tilde{\theta}\|_2. \end{aligned}$$

Taking the supremum over $S_{\pi_0}$ proves continuity of the map $\theta \mapsto R_\theta$ into the supportwise supremum norm. To verify compactness directly, let $(R_{\theta_n})_{n\ge1}$ be any sequence in the class. Compactness of $\Theta$ supplies a subsequence $\theta_{n_j} \to \theta \in \Theta$. The displayed bound then implies $\|R_{\theta_{n_j}} - R_\theta\|_{\infty,\mathrm{supp}(\pi_0)} \to 0$. Thus every sequence has a subsequence converging to a member of the class. Sequential compactness is equivalent to compactness in a metric space, which proves the claim. □

**Proposition 8** (Centered neural reward classes satisfy Condition (C1)). *Let $\Theta \subset \mathbb{R}^m$ be compact. For each $\theta \in \Theta$, let $f_\theta : S_{\pi_0} \to \mathbb{R}$ be a bounded measurable function. Assume that: (i) for $d_0$-a.e. $x$ and every $\theta \in \Theta$, the map $a \mapsto f_\theta(x,a)$ is continuous on $S_x$; and (ii) the parameter-to-function map $\theta \mapsto f_\theta$ is continuous from $\Theta$ to $\ell_\infty(S_{\pi_0})$. Define*

$$R_\theta(x,a) \coloneqq f_\theta(x,a) - \mathbb{E}_{A\sim\pi_0(\cdot|x)}[f_\theta(x,A)].$$

*Then $\mathcal{F}^{\mathrm{nn}}_\Theta \coloneqq \{R_\theta : \theta \in \Theta\}$ is compact under $\|\cdot\|_{\infty,\mathrm{supp}(\pi_0)}$, and every action section is continuous on $S_x$ for $d_0$-a.e. $x$.*

*Proof.* For a bounded measurable function $f$ on $S_{\pi_0}$, define

$$(Cf)(x,a) \coloneqq f(x,a) - \int f(x,b)\,\pi_0(db \mid x).$$

The integral is finite, its dependence on $x$ is measurable, and $\int (Cf)(x,a)\,\pi_0(da \mid x) = 0$. If the action section of $f$ is continuous, then the action section of $Cf$ is continuous because the subtracted term is constant in $a$. For any two bounded functions $f, g$,

$$\begin{aligned} |(Cf)(x,a) - (Cg)(x,a)| &\le |f(x,a) - g(x,a)| + \int |f(x,b) - g(x,b)|\,\pi_0(db \mid x) \\ &\le 2\|f - g\|_{\infty,\mathrm{supp}(\pi_0)}. \end{aligned}$$

Hence $C$ is a continuous linear map, with operator norm at most 2. By the assumed continuity of $\theta \mapsto f_\theta$, its composition $\theta \mapsto Cf_\theta = R_\theta$ is continuous.

For completeness, if $(R_{\theta_n})_{n\ge1}$ is a sequence in the centered class, compactness of $\Theta$ gives a subsequence $\theta_{n_j} \to \theta \in \Theta$. Assumption (ii) and the preceding bound yield

$$\|R_{\theta_{n_j}} - R_\theta\|_{\infty,\mathrm{supp}(\pi_0)} \le 2\|f_{\theta_{n_j}} - f_\theta\|_{\infty,\mathrm{supp}(\pi_0)} \longrightarrow 0.$$

The limit belongs to the centered class, so that class is compact. Assumption (i) and the first paragraph give the claimed continuity of its action sections. □

### D.2 Basic technical lemmas

**Lemma 9** ($\pi_0$-centering is w.l.o.g. for MNL likelihoods)**.** *Fix a measurable reward $R : \mathcal{X} \times \mathcal{A} \to \mathbb{R}$. At every context under consideration, assume $\int |R(x,a)|\, \pi_0(da \mid x) < \infty$; for the policy identity, also assume $0 < \int e^{\eta R(x,a)}\, \pi_0(da \mid x) < \infty$. These requirements hold for bounded rewards. Define its $\pi_0$-centered version*

$$\bar{R}(x,a) \coloneqq R(x,a) - m_R(x), \qquad m_R(x) \coloneqq \mathbb{E}_{a' \sim \pi_0(\cdot|x)}[R(x,a')].$$

*Then for every $x$ and every slate $\mathbf{a}$: (i) $P_R(\cdot \mid x, \mathbf{a}) = P_{\bar{R}}(\cdot \mid x, \mathbf{a})$; (ii) $\pi_R(\cdot \mid x) = \pi_{\bar{R}}(\cdot \mid x)$ for KL tilts; and (iii)* $\arg\max_{a \in S_x} R(x,a) = \arg\max_{a \in S_x} \bar{R}(x,a)$.

*Proof.* Fix a context satisfying the stated integrability conditions, and write $m = m_R(x)$. Since $m$ is finite,

$$\int \bar{R}(x,a)\, \pi_0(da \mid x) = \int R(x,a)\, \pi_0(da \mid x) - m \int 1\, \pi_0(da \mid x) = 0.$$

For every slate $\mathbf{a} = (a_1, \ldots, a_K)$ and every index $k$,

$$P_{\bar{R}}(k \mid x, \mathbf{a}) = \frac{e^{R(x,a_k)-m}}{\sum_{\ell=1}^K e^{R(x,a_\ell)-m}} = \frac{e^{-m} e^{R(x,a_k)}}{e^{-m} \sum_{\ell=1}^K e^{R(x,a_\ell)}} = P_R(k \mid x, \mathbf{a}).$$

Thus the choice probabilities, and hence their negative log-likelihoods, are unchanged.

Let $Z_R(x) = \int e^{\eta R(x,a)}\, \pi_0(da \mid x)$. Then

$$Z_{\bar{R}}(x) = \int e^{\eta(R(x,a)-m)}\, \pi_0(da \mid x) = e^{-\eta m} Z_R(x).$$

Substitution into the density formula gives

$$\frac{d\pi_{\bar{R}}(\cdot \mid x)}{d\pi_0(\cdot \mid x)}(a) = \frac{e^{\eta R(x,a) - \eta m}}{e^{-\eta m} Z_R(x)} = \frac{d\pi_R(\cdot \mid x)}{d\pi_0(\cdot \mid x)}(a),$$

so the two policy measures are equal. Finally, for any $a, b \in S_x$, $\bar{R}(x,a) \geq \bar{R}(x,b)$ holds if and only if $R(x,a) \geq R(x,b)$. The maximizing sets therefore agree, and the common tie-breaking convention gives the same selector. □

**Lemma 10** (Boundedness and Lipschitzness of the MNL log-loss)**.** *Let $\ell(\mathbf{v}, y) = \log(\sum_{k=1}^K e^{v_k}) - v_y$ be the MNL negative log-likelihood. If $\max_k v_k - \min_k v_k \leq 2B$, then for every $y$,*

$$0 \leq \ell(\mathbf{v}, y) \leq \log K + 2B.$$

*Moreover, for any $\mathbf{v}, \mathbf{v}' \in \mathbb{R}^K$ and any $y$,*

$$|\ell(\mathbf{v}, y) - \ell(\mathbf{v}', y)| \leq 2\|\mathbf{v} - \mathbf{v}'\|_\infty.$$

*Proof.* For every $y$,

$$\ell(\mathbf{v}, y) = \log\left(\sum_{k=1}^K e^{v_k - v_y}\right).$$

The summand with $k = y$ equals 1, so the expression is nonnegative. Writing $m = \min_k v_k$ and $M = \max_k v_k$, every summand is at most $e^{M-m}$. Therefore

$$\ell(\mathbf{v}, y) \leq \log\big(K e^{M-m}\big) = \log K + M - m \leq \log K + 2B.$$

To prove the Lipschitz bound without a differentiability argument, set $d = \|\mathbf{v} - \mathbf{v}'\|_\infty$. For each $k$, $v'_k - d \leq v_k \leq v'_k + d$, whence

$$e^{-d} \sum_k e^{v'_k} \leq \sum_k e^{v_k} \leq e^d \sum_k e^{v'_k}.$$

Taking logarithms gives

$$\left| \log \sum_k e^{v_k} - \log \sum_k e^{v'_k} \right| \leq d.$$

Since $|v_y - v'_y| \le d$, the triangle inequality now yields

$$|\ell(\mathbf{v}, y) - \ell(\mathbf{v}', y)| \le d + d = 2\|\mathbf{v} - \mathbf{v}'\|_\infty.$$

□

**Lemma 11** (One-step excess loss equals choice-model KL)**.** *For fixed $(x, \mathbf{a})$, truth reward $R^\star$, and candidate reward $R$, let $p^\star = P_{R^\star}(\cdot \mid x, \mathbf{a})$ and $p = P_R(\cdot \mid x, \mathbf{a})$. Then*

$$\mathbb{E}_{y\sim p^\star}\big[\ell(\mathbf{v}_R, y) - \ell(\mathbf{v}_{R^\star}, y)\big] = \mathrm{KL}(p^\star \| p).$$

*Proof.* Every softmax coordinate is strictly positive. By the definition of $\ell$, $\ell(\mathbf{v}_R, k) = -\log p_k$ and $\ell(\mathbf{v}_{R^\star}, k) = -\log p^\star_k$. Consequently,

$$\begin{aligned}\mathbb{E}_{Y\sim p^\star}[\ell(\mathbf{v}_R, Y) - \ell(\mathbf{v}_{R^\star}, Y)] &= \sum_{k=1}^K p^\star_k(-\log p_k + \log p^\star_k) \\ &= \sum_{k=1}^K p^\star_k \log \frac{p^\star_k}{p_k} = \mathrm{KL}(p^\star \| p).\end{aligned}$$

We will also use nonnegativity and its equality case. The scalar inequality $-\log z \ge 1 - z$, with equality exactly when $z = 1$, gives

$$\mathrm{KL}(p^\star \| p) \ge \sum_{k=1}^K p^\star_k \left(1 - \frac{p_k}{p^\star_k}\right) = \sum_{k=1}^K (p^\star_k - p_k) = 0.$$

Because all weights $p^\star_k$ are positive, equality holds if and only if $p_k = p^\star_k$ for every $k$. □

**Lemma 12** (Slate expectation domination)**.** *Fix $x$ and a distribution $\pi(\cdot \mid x) \ll \pi_0(\cdot \mid x)$ whose Radon–Nikodym derivative $r_x = d\pi(\cdot \mid x)/d\pi_0(\cdot \mid x)$ satisfies $r_x \ge \beta^{-1}$ $\pi_0(\cdot \mid x)$-almost surely. Let $\mathbf{A} \sim \pi(\cdot \mid x)^{\otimes K}$ and $\mathbf{A}_0 \sim \pi_0(\cdot \mid x)^{\otimes K}$. Then for every nonnegative measurable $g$,*

$$\mathbb{E}[g(\mathbf{A})] \ge \beta^{-K}\mathbb{E}[g(\mathbf{A}_0)].$$

*Proof.* Write $\mu = \pi_0(\cdot \mid x)$, $\nu = \pi(\cdot \mid x)$, and $r = d\nu/d\mu$. For every measurable rectangle $E_1 \times \cdots \times E_K$, iterated integration gives

$$\int_{E_1\times\cdots\times E_K} \prod_{k=1}^K r(a_k)\, \mu^{\otimes K}(d\mathbf{a}) = \prod_{k=1}^K \int_{E_k} r(a)\, \mu(da) = \prod_{k=1}^K \nu(E_k).$$

The measure with density $\prod_{k=1}^K r(a_k)$ relative to $\mu^{\otimes K}$ is thus $\nu^{\otimes K}$, since finite measures agreeing on measurable rectangles agree on the product $\sigma$-field. Outside a $\mu^{\otimes K}$-null set, every coordinate satisfies $r(a_k) \ge \beta^{-1}$, and hence the product is at least $\beta^{-K}$. For a nonnegative measurable $g$, including one with infinite expectation, it follows that

$$\mathbb{E}[g(\mathbf{A})] = \int g(\mathbf{a}) \prod_{k=1}^K r(a_k)\, \mu^{\otimes K}(d\mathbf{a}) \ge \beta^{-K} \int g(\mathbf{a})\, \mu^{\otimes K}(d\mathbf{a}) = \beta^{-K}\mathbb{E}[g(\mathbf{A}_0)].$$

□

**Lemma 13** (KL-tilt optimality and likelihood-ratio bounds)**.** *Fix $\eta > 0$. Let $R$ satisfy $|R(x, a)| \le B$ for $\pi_0(\cdot \mid x)$-a.e. $a$. Then the KL tilt $\pi_R(a \mid x) \propto \pi_0(a \mid x)e^{\eta R(x,a)}$ satisfies*

$$e^{-2\eta B} \le \frac{d\pi_R(\cdot \mid x)}{d\pi_0(\cdot \mid x)}(a) \le e^{2\eta B} \qquad \pi_0(\cdot \mid x)\text{-a.s.}$$

*Moreover, $\pi_R(\cdot \mid x)$ is the unique maximizer, over $\pi \ll \pi_0(\cdot \mid x)$, of $\int R(x, a)\, \pi(da) - \eta^{-1}\,\mathrm{KL}(\pi \| \pi_0(\cdot \mid x))$. Writing $Z_R(x) = \int e^{\eta R(x,a)}\, \pi_0(da \mid x)$, its value is $\eta^{-1} \log Z_R(x)$,*

*and for every such $\pi$,*

$$\int R(x,a)\,\pi(da) - \eta^{-1}\,\mathrm{KL}(\pi\|\pi_0(\cdot \mid x)) = \eta^{-1}\log Z_R(x) - \eta^{-1}\,\mathrm{KL}(\pi\|\pi_R(\cdot \mid x)). \tag{11}$$

*The identity is interpreted in the extended real sense when a KL divergence is infinite.*

*Proof.* Fix the context and abbreviate $\mu = \pi_0(\cdot \mid x)$ and $Z = Z_R(x)$. Exponentiating the reward bound and integrating against the probability measure $\mu$ gives

$$e^{-\eta B} \le e^{\eta R(x,a)} \le e^{\eta B} \quad \mu\text{-a.s.}, \qquad e^{-\eta B} \le Z \le e^{\eta B}.$$

Thus $Z$ is finite and strictly positive. The function $r(a) = e^{\eta R(x,a)}/Z$ integrates to one and satisfies

$$\frac{e^{-\eta B}}{e^{\eta B}} \le r(a) \le \frac{e^{\eta B}}{e^{-\eta B}},$$

which proves the stated density bounds. In particular, $r > 0$ $\mu$-almost surely, so $\pi_R$ and $\mu$ have the same null sets.

Let $\pi \ll \mu$ and put $f = d\pi/d\mu$. Since $r > 0$, $d\pi/d\pi_R = f/r$, and $\log r = \eta R(x,\cdot) - \log Z$. The latter function is bounded. If $\mathrm{KL}(\pi\|\mu) < \infty$, substitution into the definition of KL therefore yields

$$\begin{aligned}
\mathrm{KL}(\pi\|\pi_R) &= \int \log \frac{f(a)}{r(a)}\,\pi(da) \\
&= \int \log f(a)\,\pi(da) - \eta \int R(x,a)\,\pi(da) + \log Z \\
&= \mathrm{KL}(\pi\|\mu) - \eta \int R(x,a)\,\pi(da) + \log Z.
\end{aligned}$$

The rearranged equality is (11). If $\mathrm{KL}(\pi\|\mu) = \infty$, subtracting the bounded function $\log r$ leaves the KL divergence infinite; both objective expressions in (11) then equal $-\infty$. Here the negative part of $\int \log f\, d\pi = \int f \log f\, d\mu$ is finite because $-u \log u \le e^{-1}$ on $0 \le u \le 1$.

Finally, for a probability measure $\nu$ and $\pi \ll \nu$ with density $g$, the inequality $u \log u - u + 1 \ge 0$ gives

$$\mathrm{KL}(\pi\|\nu) = \int (g \log g - g + 1)\, d\nu \ge 0,$$

with equality only when $g = 1$ $\nu$-almost surely. Applying this fact with $\nu = \pi_R$ to (11) shows that the maximum is $\eta^{-1} \log Z$ and is attained only at $\pi = \pi_R$. □

**Lemma 14** (Supportwise upgrade of almost-sure equality)**.** *Let $(\mathcal{A}, d)$ be a separable metric space, let $\mu$ be a Borel probability measure on $\mathcal{A}$, and let $S = \mathrm{supp}(\mu)$. If continuous functions $f, g : S \to \mathbb{R}$ satisfy $f = g$ $\mu$-a.s., then $f = g$ on $S$.*

*Proof.* A separable metric space has a countable base of open sets; fix one and denote it by $(U_j)_{j\ge1}$. By the definition of support, a point $a \notin S$ has an open neighborhood of $\mu$-measure zero. A base element containing $a$ and contained in that neighborhood also has measure zero. Conversely, every point of a zero-measure open set lies outside the support. Hence

$$\mathcal{A} \setminus S = \bigcup_{j:\,\mu(U_j)=0} U_j, \qquad \mu(\mathcal{A} \setminus S) = 0.$$

The last equality uses countable subadditivity.

Set $h = f - g$, which is continuous on $S$. Suppose for a contradiction that $h(a_0) \ne 0$ for some $a_0 \in S$, and put $\delta = |h(a_0)|/2 > 0$. Relative continuity gives an open set $V \subseteq \mathcal{A}$ containing $a_0$ such that $|h(a) - h(a_0)| < \delta$ for all $a \in V \cap S$. The reverse triangle inequality implies $|h(a)| > \delta$ on $V \cap S$. Since $a_0 \in S$, $\mu(V) > 0$; since $\mu(\mathcal{A} \setminus S) = 0$, $\mu(V \cap S) = \mu(V) > 0$. Thus $h$ is nonzero on a set of positive measure, contradicting $h = 0$ $\mu$-almost surely. No such $a_0$ exists, and $f = g$ on all of $S$. □

**Lemma 15** (A common regularity set for the compact reward class)**.** *Under Condition (C1) and the centering convention, there is a measurable set $\mathcal{X}_0$ of $d_0$-measure one on which every reward in $\mathcal{F}_\mathcal{P}$ is centered and has a continuous action section. Under Condition (C2), $\mathcal{X}_0$ can also be chosen so that every reward in the class has its stated unique maximizer and margin on this same set.*

*Proof.* The compact metric class has a countable dense subset $\{R_j : j \geq 1\}$. For each $j$, choose a full-measure measurable set on which $R_j$ is centered and its action section is continuous; when Condition (C2) is assumed, also impose its margin property. The countable intersection of these sets, denoted by $\mathcal{X}_0$, has measure one.

Fix $x \in \mathcal{X}_0$ and $R \in \mathcal{F}_\mathcal{P}$. Choose a sequence $R_{j_n} \to R$ in the supportwise supremum norm. Its action sections at $x$ converge uniformly on $S_x$. The uniform limit of continuous functions is continuous, so the action section of $R$ is continuous. Also,

$$\left|\int R(x,a)\,\pi_0(da \mid x)\right| = \left|\int (R - R_{j_n})(x,a)\,\pi_0(da \mid x)\right| \leq \|R - R_{j_n}\|_{\infty,\mathrm{supp}(\pi_0)} \longrightarrow 0,$$

which proves centering.

Suppose in addition that Condition (C2) holds. Since the approximating sequence is Cauchy, for some $n_0$ and all $n, m \geq n_0$, $\|R_{j_n} - R_{j_m}\|_{\infty,\mathrm{supp}(\pi_0)} < \Delta^\mathcal{P}_{\min}/2$. At the fixed context $x$, the pointwise comparison used in Lemma 17 then shows that these rewards all have the same maximizing action, say $a_\star$: for any other supported action, its reward disadvantage under $R_{j_n}$ can decrease by at most twice that norm when the reward changes to $R_{j_m}$, and therefore remains strictly positive. For every $a \in S_x \setminus \{a_\star\}$ and $n \geq n_0$,

$$R_{j_n}(x,a_\star) - R_{j_n}(x,a) \geq \Delta^\mathcal{P}_{\min}.$$

Taking the limit in this inequality gives $R(x,a_\star) - R(x,a) \geq \Delta^\mathcal{P}_{\min}$. This holds for every such $a$, so $a_\star$ is the unique maximizer of $R$ with the same margin. Since $R$ was arbitrary, the common set has all the asserted properties. □

### D.3 Proofs of structural lemmas and offline/online results

*Proof of Lemma 2.* Fix $p \in \mathcal{P}$ and $R \in \mathcal{F}_\mathcal{P}$, and define

$$K_{p,R}(x,\mathbf{a}) := \mathrm{KL}\big(P_{R_p}(\cdot \mid x,\mathbf{a}) \,\big\|\, P_R(\cdot \mid x,\mathbf{a})\big)\,.$$

Let $\nu_0$ be the reference context–slate law

$$\nu_0(dx, d\mathbf{a}) = d_0(dx)\,\pi_0(\cdot \mid x)^{\otimes K}(d\mathbf{a}).$$

Conditioning first on $(X, \mathbf{A})$ and using Lemma 11 gives

$$\mathcal{L}_p(R) - \mathcal{L}_p(R_p) = \int K_{p,R}(x,\mathbf{a})\,\nu_0(dx, d\mathbf{a}). \tag{12}$$

The integrand is measurable, nonnegative, and bounded by $\log K + 2B_\mathcal{P}$: its conditional excess-loss representation is at most the conditional candidate loss, which has this upper bound. Thus all the expectations in the identity are finite.

Suppose $\mathcal{L}_p(R) = \mathcal{L}_p(R_p)$. For every integer $m \geq 1$, nonnegativity gives

$$0 = \int K_{p,R}\,d\nu_0 \geq \frac{1}{m}\,\nu_0\{K_{p,R} \geq 1/m\}.$$

Hence each set on the right has measure zero. Their countable union is $\{K_{p,R} > 0\}$, so $K_{p,R} = 0$ $\nu_0$-almost surely. Integrating its indicator first over the slate distribution shows that, for $d_0$-almost every $x$, $K_{p,R}(x,\mathbf{a}) = 0$ for $\pi_0(\cdot \mid x)^{\otimes K}$-almost every slate. The equality case in Lemma 11 then gives equality of the two choice laws on those slates.

Fix such a context at which both rewards are centered and have continuous action sections. Write $\mu = \pi_0(\cdot \mid x)$ and $h(a) = R(x,a) - R_p(x,a)$. Since $K \geq 2$, equality of the first and second choice odds implies

$$\begin{aligned} e^{R(x,a_1)-R(x,a_2)} &= \frac{P_R(1 \mid x,\mathbf{a})}{P_R(2 \mid x,\mathbf{a})} \\ &= \frac{P_{R_p}(1 \mid x,\mathbf{a})}{P_{R_p}(2 \mid x,\mathbf{a})} = e^{R_p(x,a_1)-R_p(x,a_2)}. \end{aligned}$$

Taking logarithms yields $h(a_1) = h(a_2)$ for $\mu^{\otimes K}$-almost every slate. The set on which this equality fails depends only on $(a_1, a_2)$. Its $K$-fold product measure is therefore its two-fold product measure, because integrating the remaining $K - 2$ coordinates contributes a factor of one. Consequently, if $A, B$ are independent with law $\mu$, then $h(A) = h(B)$ almost surely.

The boundedness of the rewards permits the following second-moment calculation:

$$\begin{aligned} 0 &= \mathbb{E}[(h(A) - h(B))^2] \\ &= \mathbb{E}[h(A)^2] + \mathbb{E}[h(B)^2] - 2\mathbb{E}[h(A)h(B)] \\ &= 2\mathbb{E}[h(A)^2] - 2\big(\mathbb{E}[h(A)]\big)^2 = 2\,\mathrm{Var}(h(A)). \end{aligned}$$

Independence justifies the factorization of the cross term. Centering of both rewards gives $\mathbb{E}[h(A)] = 0$, so $\mathbb{E}[h(A)^2] = 0$ and $h(A) = 0$ almost surely. Lemma 14, applied to the two continuous action sections, now yields $R(x, a) = R_p(x, a)$ for every $a \in S_x$.

This conclusion holds for $d_0$-almost every context. At each such context, a maximizer of $R$ is a maximizer of $R_p$, so $R_p(x, a_R(x)) = R_p(x, a_p(x))$. Integrating this zero pointwise regret proves $\mathcal{G}_p(R) = 0$. This last conclusion does not require uniqueness of the maximizer. □

**Lemma 16** (Expected one-step regret versus disagreement mass)**.** *Assume (C2). Then for every $p, q \in \mathcal{P}$,*

$$\Delta^{\mathcal{P}}_{\min}\, d_0\{a_q \neq a_p\} \leq \mathcal{G}_p(R_q) \leq \Delta^{\mathcal{P}}_{\max}\, d_0\{a_q \neq a_p\}.$$

*Proof.* Let $D_{p,q} = \{x : a_q(x) \neq a_p(x)\}$, a measurable set by the selector convention. Outside $D_{p,q}$, the selected actions coincide and their reward difference is zero. At every context in $D_{p,q}$ where Condition (C2) holds, $a_q(x)$ is a supported action different from the unique maximizer $a_p(x)$. Hence

$$\Delta^{\mathcal{P}}_{\min} \leq R_p(x, a_p(x)) - R_p(x, a_q(x)) \leq \Delta^{\mathcal{P}}_{\max}.$$

Combining the two cases gives the almost-sure inequalities

$$\Delta^{\mathcal{P}}_{\min} \mathbf{1}_{D_{p,q}}(X) \leq R_p(X, a_p(X)) - R_p(X, a_q(X)) \leq \Delta^{\mathcal{P}}_{\max} \mathbf{1}_{D_{p,q}}(X).$$

The expectation of the middle expression is $\mathcal{G}_p(R_q)$, and $\mathbb{E}[\mathbf{1}_{D_{p,q}}(X)] = d_0(D_{p,q})$. Taking expectations therefore proves both bounds. □

**Lemma 17** (Local stability of the temperature-zero selector)**.** *Assume (C2). If $p, q \in \mathcal{P}$ satisfy*

$$\|R_q - R_p\|_{\infty, \mathrm{supp}(\pi_0)} < \frac{\Delta^{\mathcal{P}}_{\min}}{2},$$

*then $a_q(X) = a_p(X)$ $d_0$-a.s.*

*Proof.* Put $d = \|R_q - R_p\|_{\infty, \mathrm{supp}(\pi_0)}$, so $2d < \Delta^{\mathcal{P}}_{\min}$. Fix a context at which the gap condition for $R_p$ holds. For any $a \in S_x$ different from $a_p(x)$,

$$\begin{aligned} R_q(x, a_p(x)) - R_q(x, a) &= R_p(x, a_p(x)) - R_p(x, a) \\ &\quad + (R_q - R_p)(x, a_p(x)) - (R_q - R_p)(x, a) \\ &\geq \Delta^{\mathcal{P}}_{\min} - |(R_q - R_p)(x, a_p(x))| - |(R_q - R_p)(x, a)| \\ &\geq \Delta^{\mathcal{P}}_{\min} - 2d > 0. \end{aligned}$$

Thus $a_p(x)$ is also the unique maximizer of $R_q(x, \cdot)$, and $a_q(x) = a_p(x)$. The set of contexts used in this argument has $d_0$-measure one, which proves the assertion. □

*Proof of Lemma 1.* Say that two rewards belong to the same selector class when their selectors agree $d_0$-almost surely. Almost-sure equality is reflexive, symmetric, and transitive, so this is an equivalence relation. Let $r = \Delta^{\mathcal{P}}_{\min}/8$. Compactness of $\mathcal{F}_{\mathcal{P}}$ gives a finite $r$-net with centers $R_{p_1}, \ldots, R_{p_M}$ in the class. For any $R_p \in \mathcal{F}_{\mathcal{P}}$, some center satisfies

$$\|R_p - R_{p_i}\|_{\infty, \mathrm{supp}(\pi_0)} \leq r < \Delta^{\mathcal{P}}_{\min}/2.$$

Lemma 17 implies that $R_p$ and $R_{p_i}$ have the same selector class. There are therefore at most $M$ distinct selector classes.

Choose one representative selector $a^{(1)}, \dots, a^{(m)}$ from each of the distinct classes. If $m = 1$, every pair of rewards has the same selector almost surely, so all the regrets in the lemma are zero and we may set $\varepsilon^{\mathcal{P}}_{\mathrm{iso}} = 1$. Suppose now that $m \geq 2$. For $i \neq j$, define

$$\delta_{ij} \coloneqq d_0\{x : a^{(i)}(x) \neq a^{(j)}(x)\}.$$

The definition of distinct equivalence classes gives $\delta_{ij} > 0$. There are finitely many pairs, and hence

$$\delta_{\mathcal{P}} \coloneqq \min_{1 \leq i < j \leq m} \delta_{ij} > 0.$$

Changing a representative on a null set does not change any of these disagreement probabilities.

Set $\varepsilon^{\mathcal{P}}_{\mathrm{iso}} = \Delta^{\mathcal{P}}_{\min}\delta_{\mathcal{P}}$. For $p, q \in \mathcal{P}$ in the same selector class, the regret is zero. For $p, q$ in distinct classes, their disagreement mass is at least $\delta_{\mathcal{P}}$, so Lemma 16 gives

$$\mathcal{G}_p(R_q) \geq \Delta^{\mathcal{P}}_{\min} d_0\{a_p \neq a_q\} \geq \Delta^{\mathcal{P}}_{\min}\delta_{\mathcal{P}} = \varepsilon^{\mathcal{P}}_{\mathrm{iso}}.$$

These two cases establish the claimed isolation. □

**Lemma 18** (Continuity of the truth-centered regret and risk maps)**.** *Assume (C1)–(C2). Then the maps*

$$(R^\star, R) \longmapsto \mathcal{G}_{R^\star}(R) \qquad \text{and} \qquad (R^\star, R) \longmapsto \mathcal{L}_{R^\star}(R) - \mathcal{L}_{R^\star}(R^\star)$$

*are continuous on $\mathcal{F}^2_{\mathcal{P}}$.*

*Proof.* For a reward $U \in \mathcal{F}_{\mathcal{P}}$, let $a_U$ denote its selector, and use $\mathcal{G}_U$ and $\mathcal{L}_U$ for the regret and population loss with $U$ as the truth. Consider a sequence $(U_n, V_n) \to (U, V)$ in $\mathcal{F}_{\mathcal{P}} \times \mathcal{F}_{\mathcal{P}}$. For all sufficiently large $n$, both

$$\|U_n - U\|_{\infty,\mathrm{supp}(\pi_0)} < \Delta^{\mathcal{P}}_{\min}/2, \qquad \|V_n - V\|_{\infty,\mathrm{supp}(\pi_0)} < \Delta^{\mathcal{P}}_{\min}/2.$$

Lemma 17 then gives $a_{U_n} = a_U$ and $a_{V_n} = a_V$ almost surely. Taking a countable intersection over these $n$ supplies a common full-measure set of contexts. On that set,

$$\begin{aligned} &U_n(x, a_{U_n}(x)) - U_n(x, a_{V_n}(x)) - U(x, a_U(x)) + U(x, a_V(x)) \\ &\quad = (U_n - U)(x, a_U(x)) - (U_n - U)(x, a_V(x)). \end{aligned}$$

It follows that

$$|\mathcal{G}_{U_n}(V_n) - \mathcal{G}_U(V)| \leq 2\|U_n - U\|_{\infty,\mathrm{supp}(\pi_0)} \longrightarrow 0.$$

This proves continuity of the regret map.

For the excess-risk map, use the fixed reference context–slate law $\nu_0(dx, d\mathbf{a}) = d_0(dx)\pi_0(\cdot \mid x)^{\otimes K}(d\mathbf{a})$. By Lemma 11,

$$\mathcal{L}_{U_n}(V_n) - \mathcal{L}_{U_n}(U_n) = \int \mathrm{KL}(P_{U_n}(\cdot \mid x, \mathbf{a}) \,\|\, P_{V_n}(\cdot \mid x, \mathbf{a}))\ \nu_0(dx, d\mathbf{a}).$$

For every reference-supported slate, the two reward vectors converge coordinatewise to those associated with $U$ and $V$. Softmax is continuous and has strictly positive coordinates. Therefore the finite sum defining the displayed KL integrand converges to the corresponding integrand for $(U, V)$. Each integrand lies between zero and $\ell_{\max} = \log K + 2B_{\mathcal{P}}$, by its conditional excess-loss representation and Lemma 10. The constant function $\ell_{\max}$ is integrable under the probability measure $\nu_0$. The dominated convergence theorem now gives

$$\mathcal{L}_{U_n}(V_n) - \mathcal{L}_{U_n}(U_n) \longrightarrow \mathcal{L}_U(V) - \mathcal{L}_U(U).$$

Both domains are metric spaces, so the sequential continuity just proved is equivalent to continuity. □

*Proof of Lemma 3.* Define

$$\mathcal{K}_{\varepsilon_0} \coloneqq \{(U, V) \in \mathcal{F}_{\mathcal{P}}^2 : \mathcal{G}_U(V) \geq \varepsilon_0\}.$$

By Lemma 18, this set is the inverse image of the closed interval $[\varepsilon_0, \infty)$ under a continuous map. It is therefore closed in $\mathcal{F}_{\mathcal{P}}^2$. The latter space is compact as the product of two compact spaces, so $\mathcal{K}_{\varepsilon_0}$ is compact.

If $\mathcal{K}_{\varepsilon_0}$ is empty, there are no pairs satisfying the antecedent of the lemma, and we may take $\gamma_{\mathcal{P},\varepsilon_0} = 1$. Otherwise, consider the continuous, nonnegative function

$$\Psi(U, V) \coloneqq \mathcal{L}_U(V) - \mathcal{L}_U(U).$$

Its continuity follows from Lemma 18, and its nonnegativity follows from Lemma 11. Because $\mathcal{K}_{\varepsilon_0}$ is nonempty and compact, $\Psi$ attains its minimum at some $(U_0, V_0) \in \mathcal{K}_{\varepsilon_0}$. If this minimum were zero, Lemma 2 would give $\mathcal{G}_{U_0}(V_0) = 0$. Membership in $\mathcal{K}_{\varepsilon_0}$, however, gives $\mathcal{G}_{U_0}(V_0) \geq \varepsilon_0 > 0$, a contradiction. Thus the attained minimum is strictly positive.

Take

$$\gamma_{\mathcal{P},\varepsilon_0} \coloneqq \min\left\{1, \min_{(U,V)\in\mathcal{K}_{\varepsilon_0}} \Psi(U, V)\right\}.$$

This is a finite, strictly positive number, and every pair with $\mathcal{G}_U(V) \geq \varepsilon_0$ has $\Psi(U, V) \geq \gamma_{\mathcal{P},\varepsilon_0}$. Substituting $U = R_p$ and $V = R_q$ proves the claim uniformly over $p, q$. No continuity assumption on the indexing map $p \mapsto R_p$ is needed: the compactness argument takes place directly in the reward class. □

**Predictable risks and measurability.** For the concentration arguments, $\mathscr{H}_s$ is the history through round $s$, including any randomization used by the learner before the next observation. The sampling rule is measurable with respect to the preceding history. For a fixed candidate reward, write

$$\ell_s(R) \coloneqq \ell(\mathbf{v}_R(X_s, \mathbf{A}_s), Y_s), \qquad m_s(R) \coloneqq \mathbb{E}[\ell_s(R) \mid \mathscr{H}_{s-1}].$$

We define $m_s(R)$, simultaneously for all $R$, by integration against the specified conditional observation kernel. This fixes versions for which the Lipschitz inequalities below hold simultaneously on the reward class. In particular, evaluating the random function $m_s(\cdot)$ at a fitted reward does not mean that the fitted reward is held fixed inside an ordinary conditional expectation.

**Lemma 19** (Deviation tail bound via an $\epsilon$-net)**.** *Assume $\mathcal{F}$ is separable and bounded in $\|\cdot\|_{\infty,\mathrm{supp}(\pi_0)}$ by $B$, and that at each round every slate coordinate is sampled from a distribution absolutely continuous with respect to $\pi_0(\cdot \mid x)$. Let*

$$\mathcal{L}_t(R) = \frac{1}{t}\sum_{s=1}^{t} \mathbb{E}[\ell_s(R) \mid \mathscr{H}_{s-1}], \qquad b_t = \sup_{R\in\mathcal{F}} |\widehat{\mathcal{L}}_t(R) - \mathcal{L}_t(R)|.$$

*If $\mathcal{C}_\epsilon \subseteq \mathcal{F}$ is a finite $\epsilon$-net of $\mathcal{F}$, with centers in the class, then for every $u > 0$,*

$$\mathbb{P}(b_t \geq u + 4\epsilon) \leq 2|\mathcal{C}_\epsilon| \exp\left(-\frac{tu^2}{2(\log K + 2B)^2}\right).$$

*Proof.* Set $L = \log K + 2B > 0$, fix a deterministic $R \in \mathcal{F}$, and let $Z_s(R) = \ell_s(R) - m_s(R)$. Absolute continuity of the slate coordinates and full reference measure of the support imply that $0 \leq \ell_s(R) \leq L$ almost surely. Conditional expectation preserves these bounds, so $0 \leq m_s(R) \leq L$, $|Z_s(R)| \leq L$, and $\mathbb{E}[Z_s(R) \mid \mathscr{H}_{s-1}] = 0$.

We derive the martingale tail bound explicitly. For $z \in [-L, L]$ and $\lambda \in \mathbb{R}$, convexity of the exponential function gives the chord bound

$$e^{\lambda z} \leq \frac{L + z}{2L} e^{\lambda L} + \frac{L - z}{2L} e^{-\lambda L}.$$

Taking conditional expectations with $z = Z_s(R)$ yields

$$\mathbb{E}[e^{\lambda Z_s(R)} \mid \mathscr{H}_{s-1}] \leq \frac{e^{\lambda L} + e^{-\lambda L}}{2} = \cosh(\lambda L) \leq e^{\lambda^2 L^2/2}.$$

The last inequality follows from the power series: for every integer $j \geq 0$, $(2j)! \geq 2^j j!$, and hence $\sum_{j\geq 0} v^{2j}/(2j)! \leq \sum_{j\geq 0}(v^2/2)^j/j!$. Since $\sum_{s=1}^{t-1} Z_s(R)$ is $\mathscr{H}_{t-1}$-measurable, the tower property gives

$$\mathbb{E}\exp\left(\lambda \sum_{s=1}^{t} Z_s(R)\right) = \mathbb{E}\Big[e^{\lambda \sum_{s=1}^{t-1} Z_s(R)} \mathbb{E}[e^{\lambda Z_t(R)} \mid \mathscr{H}_{t-1}]\Big]$$
$$\leq e^{\lambda^2 L^2/2} \mathbb{E}\exp\left(\lambda \sum_{s=1}^{t-1} Z_s(R)\right).$$

Iterating from $t$ down to zero gives $\mathbb{E}e^{\lambda \sum_{s=1}^{t} Z_s(R)} \leq e^{t\lambda^2 L^2/2}$. For $u > 0$ and $\lambda > 0$, Markov's inequality therefore implies

$$\mathbb{P}\left(\frac{1}{t}\sum_{s=1}^{t} Z_s(R) \geq u\right) \leq \exp\left(-\lambda t u + \frac{t\lambda^2 L^2}{2}\right).$$

Choosing $\lambda = u/L^2$ gives $e^{-tu^2/(2L^2)}$. Applying the same argument to $-Z_s(R)$, and adding the two one-sided bounds, gives

$$\mathbb{P}\Big(|\widehat{\mathcal{L}}_t(R) - \mathcal{L}_t(R)| \geq u\Big) \leq 2\exp\left(-\frac{tu^2}{2L^2}\right). \tag{13}$$

No independence across rounds was used.

For the passage from the net to the class, let $R \in \mathcal{F}$ and choose $R^\circ \in \mathcal{C}_\epsilon$ with $\|R - R^\circ\|_{\infty,\mathrm{supp}(\pi_0)} \leq \epsilon$. Lemma 10 gives $|\ell_s(R) - \ell_s(R^\circ)| \leq 2\epsilon$ on supported observations. Integration against the conditional observation kernel also gives $|m_s(R) - m_s(R^\circ)| \leq 2\epsilon$. Averaging these inequalities separately yields

$$|\widehat{\mathcal{L}}_t(R) - \widehat{\mathcal{L}}_t(R^\circ)| \leq 2\epsilon, \qquad |\mathcal{L}_t(R) - \mathcal{L}_t(R^\circ)| \leq 2\epsilon.$$

The triangle inequality then gives

$$|\widehat{\mathcal{L}}_t(R) - \mathcal{L}_t(R)| \leq |\widehat{\mathcal{L}}_t(R^\circ) - \mathcal{L}_t(R^\circ)| + 4\epsilon.$$

Taking the supremum over $R$ proves

$$b_t \leq \max_{R^\circ \in \mathcal{C}_\epsilon} |\widehat{\mathcal{L}}_t(R^\circ) - \mathcal{L}_t(R^\circ)| + 4\epsilon. \tag{14}$$

Both risk functions are Lipschitz in $R$. Thus, if $D$ is a countable dense subset of the separable class, the supremum defining $b_t$ equals the supremum over $D$: approximate any $R$ by a sequence in $D$ and use continuity. Consequently $b_t$ is measurable.

By (14), the event $\{b_t \geq u + 4\epsilon\}$ is contained in the union, over $R^\circ \in \mathcal{C}_\epsilon$, of the events $\{|\widehat{\mathcal{L}}_t(R^\circ) - \mathcal{L}_t(R^\circ)| \geq u\}$. The union bound and (13) therefore give

$$\mathbb{P}(b_t \geq u + 4\epsilon) \leq 2|\mathcal{C}_\epsilon| \exp\left(-\frac{tu^2}{2L^2}\right),$$

as asserted. □

#### D.3.1 Offline exact ERM fixed-scale control

**Lemma 20** (Offline exact ERM: exponential fixed-scale control)**.** *Assume (C1)–(C2). Fix any $\varepsilon_0 > 0$, and let $\gamma_{\mathcal{P},\varepsilon_0} > 0$ be the fixed-scale loss-gap constant from Lemma 3. Consider offline data generated under truth $p \in \mathcal{P}$ by*

$$X_s \sim d_0, \qquad \mathbf{A}_s \sim \pi_0(\cdot \mid X_s)^{\otimes K}, \qquad Y_s \sim P_{R_p}(\cdot \mid X_s, \mathbf{A}_s), \qquad s = 1, \ldots, n,$$

*independently across $s$. Let*

$$\widehat{R}_n^{\mathrm{off}} \in \operatorname*{arg\,min}_{R \in \mathcal{F}_{\mathcal{P}}} \frac{1}{n}\sum_{s=1}^{n} \ell\big(\mathbf{v}_R(X_s, \mathbf{A}_s), Y_s\big)$$

*be an exact offline ERM. Define*

$$\ell_{\max} \coloneqq \log K + 2B_{\mathcal{P}}, \qquad N_{\varepsilon_0}^{\mathrm{off}} \coloneqq \mathcal{N}\big(\mathcal{F}_{\mathcal{P}}, \gamma_{\mathcal{P},\varepsilon_0}/32, \|\cdot\|_{\infty,\mathrm{supp}(\pi_0)}\big),$$

*and*

$$c_{\varepsilon_0}^{\mathrm{off}} \coloneqq \frac{\gamma_{\mathcal{P},\varepsilon_0}^2}{128(\log K + 2B_{\mathcal{P}})^2}.$$

*Then, for every $n \geq 1$,*

$$\sup_{p\in\mathcal{P}} \mathbb{P}_p\Big(\mathcal{G}_p(\widehat{R}_n^{\mathrm{off}}) \geq \varepsilon_0\Big) \leq 2N_{\varepsilon_0}^{\mathrm{off}} \exp\{-c_{\varepsilon_0}^{\mathrm{off}} n\}.$$

*Proof.* Fix $p \in \mathcal{P}$, and abbreviate $\gamma_0 = \gamma_{\mathcal{P},\varepsilon_0}$ and $L = \log K + 2B_{\mathcal{P}}$. Let

$$\widehat{L}_n(R) = \frac{1}{n}\sum_{s=1}^{n} \ell_s(R), \qquad b_n = \sup_{R\in\mathcal{F}_{\mathcal{P}}} |\widehat{L}_n(R) - \mathcal{L}_p(R)|.$$

The offline observations are independent and each has the reference context–slate distribution and truth choice law used to define $\mathcal{L}_p$. Thus, for every fixed $R$, its predictable loss is $m_s(R) = \mathcal{L}_p(R)$ at every round. This equality holds simultaneously in $R$ when both sides are defined by integration against that common observation kernel. Consequently $b_n$ is exactly the uniform deviation appearing in Lemma 19. The class is compact, hence separable, and has finite covering numbers. The hypotheses of that lemma are therefore satisfied.

Because $R_p$ is a feasible competitor for the exact ERM, $\widehat{L}_n(\widehat{R}_n^{\mathrm{off}}) \leq \widehat{L}_n(R_p)$. On every sample path where the deviation is defined,

$$\begin{aligned}\mathcal{L}_p(\widehat{R}_n^{\mathrm{off}}) - \mathcal{L}_p(R_p) &= \big[\mathcal{L}_p(\widehat{R}_n^{\mathrm{off}}) - \widehat{L}_n(\widehat{R}_n^{\mathrm{off}})\big] + \big[\widehat{L}_n(\widehat{R}_n^{\mathrm{off}}) - \widehat{L}_n(R_p)\big] \\ &\quad + \big[\widehat{L}_n(R_p) - \mathcal{L}_p(R_p)\big] \\ &\leq b_n + 0 + b_n = 2b_n. \end{aligned} \tag{15}$$

The use of the random fitted reward in this inequality is justified by the supremum over the entire class in the definition of $b_n$.

Let $E_n = \{\mathcal{G}_p(\widehat{R}_n^{\mathrm{off}}) \geq \varepsilon_0\}$. On $E_n$, Lemma 3 gives $\mathcal{L}_p(\widehat{R}_n^{\mathrm{off}}) - \mathcal{L}_p(R_p) \geq \gamma_0$. Combining this with (15) yields

$$E_n \subseteq \{b_n \geq \gamma_0/2\} \subseteq \{b_n \geq \gamma_0/4\}. \tag{16}$$

The weaker threshold matches the net radius used below.

Choose an in-class net $\mathcal{C}$ of radius $r = \gamma_0/32$ and cardinality $N_{\varepsilon_0}^{\mathrm{off}}$, and set $u = \gamma_0/8$. The definition of the covering number permits such a net; all its centers obey the same reward envelope. Since $u + 4r = \gamma_0/8 + \gamma_0/8 = \gamma_0/4$, the fixed-reward bound (13) and a union bound give

$$\mathbb{P}_p\left(\max_{R\in\mathcal{C}} |\widehat{L}_n(R) - \mathcal{L}_p(R)| \geq \gamma_0/8\right) \leq 2N_{\varepsilon_0}^{\mathrm{off}} \exp\left(-\frac{n\gamma_0^2}{128L^2}\right). \tag{17}$$

The net extension (14) gives $b_n \leq \max_{R\in\mathcal{C}} |\widehat{L}_n(R) - \mathcal{L}_p(R)| + 4r$. Thus $b_n \geq \gamma_0/4$ implies that this maximum is at least $\gamma_0/4 - 4r = \gamma_0/8$, and hence

$$\begin{aligned}\mathbb{P}_p(b_n \geq \gamma_0/4) &\leq 2N_{\varepsilon_0}^{\mathrm{off}} \exp\left(-\frac{n(\gamma_0/8)^2}{2L^2}\right) \\ &= 2N_{\varepsilon_0}^{\mathrm{off}} \exp\left(-\frac{n\gamma_0^2}{128L^2}\right) = 2N_{\varepsilon_0}^{\mathrm{off}} e^{-c_{\varepsilon_0}^{\mathrm{off}} n}. \end{aligned} \tag{18}$$

By (16), the same upper bound applies to $\mathbb{P}_p(E_n)$. The envelope, covering number, and loss-gap constant were chosen for the whole class, independently of $p$. Taking the supremum over $p \in \mathcal{P}$ proves the lemma. □

*Proof of Theorem 5.* Fix $p \in \mathcal{P}$ and $n \geq 1$. The fitted reward belongs to $\mathcal{F}_{\mathcal{P}}$, so Lemma 1 ensures that its regret is either zero or at least $\varepsilon_{\mathrm{iso}}^{\mathcal{P}}$. Since any supported selector has regret at most $\Delta_{\max}^{\mathcal{P}}$,

$$\mathcal{G}_p(\widehat{R}_n^{\mathrm{off}}) \leq \Delta_{\max}^{\mathcal{P}}\, \mathbf{1}\{\mathcal{G}_p(\widehat{R}_n^{\mathrm{off}}) \geq \varepsilon_{\mathrm{iso}}^{\mathcal{P}}\}.$$

Taking expectations and applying Lemma 20 at $\varepsilon_0 = \varepsilon^{\mathcal{P}}_{\mathrm{iso}}$ gives

$$\mathbb{E}_p[\mathcal{G}_p(\widehat{R}^{\mathrm{off}}_n)] \le \Delta^{\mathcal{P}}_{\max}\mathbb{P}_p\Big(\mathcal{G}_p(\widehat{R}^{\mathrm{off}}_n) \ge \varepsilon^{\mathcal{P}}_{\mathrm{iso}}\Big) \le 2\Delta^{\mathcal{P}}_{\max}N^{\mathrm{off}}_0 e^{-c^{\mathrm{off}}_0 n}.$$

The constants are uniform in $p$, which permits taking the supremum over $\mathcal{P}$.

If $\Delta^{\mathcal{P}}_{\max} = 0$, the regret is identically zero, so no logarithm of this constant is needed. Otherwise set $A = 2\Delta^{\mathcal{P}}_{\max}N^{\mathrm{off}}_0 > 0$ and $c = c^{\mathrm{off}}_0 > 0$. For $\varepsilon > 0$, monotonicity of the logarithm and positivity of $c$ give the equivalences

$$Ae^{-cn} \le \varepsilon \iff \log A - cn \le \log\varepsilon \iff n \ge c^{-1}\log(A/\varepsilon).$$

Rounding this last threshold upward and imposing $n \ge 1$ yields the sample-size condition in the statement. For the fixed class, $c^{-1}\log(A/\varepsilon) = c^{-1}\log(1/\varepsilon) + c^{-1}\log A$, so the required sample size has the asserted logarithmic order. □

**Lemma 21** (Fixed-scale control from a truth-centered loss gap)**.** *Assume (C1)–(C2), and fix $\varepsilon_0 > 0$. Let $\gamma_{\mathcal{P},\varepsilon_0} > 0$ be any witness from Lemma 3. Define*

$$\gamma = \beta^{-K}\gamma_{\mathcal{P},\varepsilon_0}, \qquad N_\gamma = \mathcal{N}\big(\mathcal{F}_{\mathcal{P}}, \gamma/32, \|\cdot\|_{\infty,\mathrm{supp}(\pi_0)}\big), \qquad c_\gamma = \frac{\gamma^2}{128(\log K + 2B_{\mathcal{P}})^2}.$$

*Then, for every truth $p \in \mathcal{P}$ and every $t \ge 1$,*

$$\mathbb{P}_p(\mathcal{G}_p(\widehat{R}_t) \ge \varepsilon_0) \le 2N_\gamma e^{-c_\gamma t}.$$

*Proof.* Fix a truth $p$, and let $\mathscr{H}_s$ and $m_s(R)$ have the predictable-risk meanings specified above. Put

$$\widehat{L}_t(R) = \frac{1}{t}\sum_{s=1}^{t}\ell_s(R), \qquad \bar{L}_t(R) = \frac{1}{t}\sum_{s=1}^{t}m_s(R), \qquad b_t = \sup_{R\in\mathcal{F}_{\mathcal{P}}}|\widehat{L}_t(R) - \bar{L}_t(R)|.$$

The function $\bar{L}_t$ is generally random, because the sampling policies depend on the past. It is not identified with the reference population loss $\mathcal{L}_p$.

At round $s = 1$, the sampling policy is $\pi_0$, so its density relative to $\pi_0$ is one. At every round $s \ge 2$, the policy is $\pi_{\widehat{R}_{s-1}}$, and the fitted reward belongs to $\mathcal{F}_{\mathcal{P}}$. Lemma 13 therefore gives, conditionally on the preceding history,

$$\frac{d\pi_s(\cdot \mid x)}{d\pi_0(\cdot \mid x)}(a) \ge \beta^{-1}, \qquad \beta = e^{2\eta B_{\mathcal{P}}} \ge 1,$$

for reference-almost every action and for $d_0$-almost every context. This also holds at $s = 1$.

For an arbitrary candidate $R \in \mathcal{F}_{\mathcal{P}}$, define $K_{p,R}$ as in (12). Integrating first over the preference label, whose conditional law is $P_{R_p}$, and then over the context and slate gives

$$\begin{aligned} m_s(R) - m_s(R_p) &= \int_{\mathcal{X}}\int_{\mathcal{A}^K} K_{p,R}(x,\mathbf{a})\,\pi_s(\cdot \mid x)^{\otimes K}(d\mathbf{a})\,d_0(dx) \\ &\ge \beta^{-K}\int_{\mathcal{X}}\int_{\mathcal{A}^K} K_{p,R}(x,\mathbf{a})\,\pi_0(\cdot \mid x)^{\otimes K}(d\mathbf{a})\,d_0(dx) \\ &= \beta^{-K}\big(\mathcal{L}_p(R) - \mathcal{L}_p(R_p)\big). \end{aligned} \tag{19}$$

The inequality is Lemma 12, applied to the nonnegative function $K_{p,R}$. The last equality is (12). The conditional risks here are explicit kernel integrals, so the displayed comparison holds simultaneously for all candidates at each admissible history. Averaging (19) over $s = 1, \dots, t$ yields

$$\bar{L}_t(R) - \bar{L}_t(R_p) \ge \beta^{-K}\big(\mathcal{L}_p(R) - \mathcal{L}_p(R_p)\big) \qquad (R \in \mathcal{F}_{\mathcal{P}}). \tag{20}$$

Since the inequality is uniform in the candidate, it can now be evaluated at the data-dependent $\widehat{R}_t$. On the event $E_t = \{\mathcal{G}_p(\widehat{R}_t) \ge \varepsilon_0\}$, Lemma 3 and (20) imply

$$\bar{L}_t(\widehat{R}_t) - \bar{L}_t(R_p) \ge \beta^{-K}\gamma_{\mathcal{P},\varepsilon_0} = \gamma.$$

On the other hand, exact empirical minimization gives $\widehat{L}_t(\widehat{R}_t) \le \widehat{L}_t(R_p)$. Therefore, on every sample path,

$$\begin{aligned}\bar{L}_t(\widehat{R}_t) - \bar{L}_t(R_p) &= \big[\bar{L}_t(\widehat{R}_t) - \widehat{L}_t(\widehat{R}_t)\big] + \big[\widehat{L}_t(\widehat{R}_t) - \widehat{L}_t(R_p)\big] \\ &\quad + \big[\widehat{L}_t(R_p) - \bar{L}_t(R_p)\big] \\ &\le b_t + 0 + b_t = 2b_t.\end{aligned}$$

It follows that

$$E_t \subseteq \{b_t \ge \gamma/2\} \subseteq \{b_t \ge \gamma/4\}. \tag{21}$$

The substitution is justified by uniformity in the reward argument; no predictability of $\widehat{R}_t$ is required.

The reward class is compact and hence separable. Its envelope is $B_{\mathcal{P}}$, and every online slate coordinate is absolutely continuous with respect to the reference policy. Choose an in-class net of radius $r = \gamma/32$ and size $N_\gamma$, and put $u = \gamma/8$. All constants are independent of the chosen truth. Since $u + 4r = \gamma/4$, (21) and Lemma 19 give

$$\begin{aligned}\mathbb{P}_p(E_t) &\le \mathbb{P}_p(b_t \ge \gamma/4) \\ &\le 2N_\gamma \exp\left(-\frac{t(\gamma/8)^2}{2(\log K + 2B_{\mathcal{P}})^2}\right) \\ &= 2N_\gamma \exp\left(-\frac{t\gamma^2}{128(\log K + 2B_{\mathcal{P}})^2}\right) = 2N_\gamma e^{-c_\gamma t}.\end{aligned}$$

□

*Proof of Theorem 4.* Fix $p \in \mathcal{P}$. The initialization $\widehat{R}_0 = 0$ need not belong to $\mathcal{F}_{\mathcal{P}}$, but its selector is supported, so $0 \le \mathcal{G}_p(\widehat{R}_0) \le \Delta^{\mathcal{P}}_{\max}$. For every $t \ge 1$, the exact ERM does belong to $\mathcal{F}_{\mathcal{P}}$. Lemma 1 and the definition $\varepsilon_0 = \varepsilon^{\mathcal{P}}_{\mathrm{iso}}$ therefore give, pathwise,

$$0 \le \mathcal{G}_p(\widehat{R}_t) \le \Delta^{\mathcal{P}}_{\max} \mathbf{1}\{\mathcal{G}_p(\widehat{R}_t) \ge \varepsilon_0\}.$$

The fixed-scale estimate in Lemma 21 implies

$$\mathbb{E}_p[\mathcal{G}_p(\widehat{R}_t)] \le \Delta^{\mathcal{P}}_{\max} \min\{1, 2N_\gamma e^{-c_\gamma t}\}.$$

The minimum with one is valid because the intervening quantity is a probability. Consequently, for every $T \ge 1$,

$$\mathrm{Regret}^{\mathrm{ERM}}_{0,p}(T) \le \Delta^{\mathcal{P}}_{\max}\left(1 + \sum_{t=1}^{T-1} \min\{1, 2N_\gamma e^{-c_\gamma t}\}\right). \tag{22}$$

Let $m = \lceil c_\gamma^{-1} \log(2N_\gamma) \rceil$. The class is nonempty and compact, so $1 \le N_\gamma < \infty$, and the positive loss gap gives $c_\gamma > 0$. In particular, $m$ is finite and $2N_\gamma e^{-c_\gamma m} \le 1$. The summands with $1 \le t \le m$ are each at most one. For $t = m + j$, $j \ge 1$,

$$\min\{1, 2N_\gamma e^{-c_\gamma(m+j)}\} \le 2N_\gamma e^{-c_\gamma m} e^{-c_\gamma j} \le e^{-c_\gamma j}.$$

All terms are nonnegative, so extending the finite sum in (22) to an infinite sum gives

$$\begin{aligned}\sum_{t=1}^{T-1} \min\{1, 2N_\gamma e^{-c_\gamma t}\} &\le m + \sum_{j=1}^{\infty} e^{-c_\gamma j} \\ &= m + \frac{e^{-c_\gamma}}{1 - e^{-c_\gamma}} = m + \frac{1}{e^{c_\gamma} - 1}.\end{aligned}$$

Substitution into (22) gives the stated bound. Its right-hand side depends only on the model class, $K$, and $\eta$, not on $p$ or $T$. Taking both suprema proves the theorem. □

# E BOUNDED TEMPERATURE-ZERO REGRET UNDER THE GENERAL PREFERENCE MODEL

The main results use MNL/Bradley–Terry feedback, where the reward itself defines the selected action. Here we make explicit the finite-action, finite-class extension based on the Nash score of a KL-regularized preference game. The game is the general-preference formulation of Wu et al. (2025); the data-collection rule below uses two independent draws from the current policy, as in this paper. Preference probabilities equal to zero or one are allowed.

Let $\mathcal{A}$ be finite, let $\eta > 0$, and let $\mathfrak{P}_{\mathrm{GP}}$ be a nonempty finite class of measurable kernels $P : \mathcal{X} \times \mathcal{A} \times \mathcal{A} \to [0,1]$ satisfying

$$P(x,a,b) + P(x,b,a) = 1.$$

The true kernel belongs to this class. Write $S_x = \{a \in \mathcal{A} : \pi_0(a \mid x) > 0\}$, and use a fixed ordering of $\mathcal{A}$ to break all selector ties. For distributions $\mu, \nu$ on $S_x$, define

$$J_{P,x}(\mu,\nu) = \sum_{a,b \in S_x} \mu(a)\nu(b)P(x,a,b) - \eta^{-1}\,\mathrm{KL}(\mu\|\pi_0(\cdot \mid x)) + \eta^{-1}\,\mathrm{KL}(\nu\|\pi_0(\cdot \mid x)). \tag{23}$$

The first argument is maximized and the second is minimized. All the KL divergences in this finite simplex are finite, with the convention $0 \log 0 = 0$.

**Lemma 22** (The Nash policy and its reference coverage)**.** *For every $P \in \mathfrak{P}_{\mathrm{GP}}$, the game* (23) *has a unique symmetric Nash policy $\pi_P$. It is measurable in the context and satisfies*

$$\phi_P(x,a) = \sum_{b \in S_x} \pi_P(b \mid x)P(x,a,b), \qquad \pi_P(a \mid x) = \frac{\pi_0(a \mid x)e^{\eta\phi_P(x,a)}}{\sum_{c \in S_x} \pi_0(c \mid x)e^{\eta\phi_P(x,c)}}. \tag{24}$$

*For all supported actions,*

$$0 \le \phi_P(x,a) \le 1, \qquad e^{-\eta} \le \frac{\pi_P(a \mid x)}{\pi_0(a \mid x)} \le e^{\eta}.$$

*Proof.* Fix $x$ and abbreviate $\rho = \pi_0(\cdot \mid x)$. For a probability vector $\nu$ on the full finite action set, put

$$u_\nu(a) = \sum_{b \in \mathcal{A}} \nu(b)P(x,a,b), \qquad T(\nu)(a) = \frac{\rho(a)e^{\eta u_\nu(a)}}{\sum_c \rho(c)e^{\eta u_\nu(c)}}.$$

Here $u_\nu(a) \in [0,1]$. The denominator lies in $[1, e^\eta]$, so $T$ is continuous on the finite-dimensional probability simplex and takes values in that same simplex, with support contained in $S_x$. The simplex is nonempty, compact, and convex. Brouwer's fixed-point theorem therefore supplies a vector $v = T(v)$. It is supported on $S_x$, and in fact $v(a) > 0$ for each $a \in S_x$.

By Lemma 13, applied with reward $u_v$, $T(v)$ is the unique maximizer of $\sum_a \mu(a)u_v(a) - \eta^{-1}\,\mathrm{KL}(\mu\|\rho)$. The remaining term in $J_{P,x}(\mu,v)$ is independent of $\mu$. Thus $v = T(v)$ maximizes $J_{P,x}(\mu,v)$, and the variational identity gives

$$J_{P,x}(\mu,v) = J_{P,x}(v,v) - \eta^{-1}\,\mathrm{KL}(\mu\|v). \tag{25}$$

The kernel symmetry and cancellation of the regularizers imply

$$J_{P,x}(\mu,\nu) + J_{P,x}(\nu,\mu) = \sum_{a,b} \mu(a)\nu(b)\big(P(x,a,b) + P(x,b,a)\big) = 1.$$

In particular, $J_{P,x}(v,v) = 1/2$. It follows that for all $\mu, \nu$ on $S_x$,

$$J_{P,x}(\mu,v) \le \tfrac{1}{2}, \qquad J_{P,x}(v,\nu) = 1 - J_{P,x}(\nu,v) \ge \tfrac{1}{2}.$$

These are precisely the saddle inequalities, so $(v,v)$ is a symmetric Nash equilibrium.

If $w$ is another symmetric Nash policy, it must maximize $J_{P,x}(\mu,w)$; uniqueness of the KL-regularized best response therefore gives $w = T(w)$. Apply (25) to the two fixed points to obtain

$$J_{P,x}(w,v) = \tfrac{1}{2} - \eta^{-1}\,\mathrm{KL}(w\|v), \qquad J_{P,x}(v,w) = \tfrac{1}{2} - \eta^{-1}\,\mathrm{KL}(v\|w).$$

Adding these identities and using the symmetry identity above gives $\mathrm{KL}(w\|v) + \mathrm{KL}(v\|w) = 0$. Nonnegativity and its equality case imply $v = w$. Define $\pi_P(\cdot \mid x) = v$. The fixed-point formula is (24). Its numerator factor $e^{\eta\phi_P(x,a)}$ and normalizer both lie in $[1, e^\eta]$, which gives the claimed density bounds.

To check measurability as $x$ varies, let $(q_j)_{j\geq 1}$ be a countable dense subset of the full action simplex and restore the notation $T_x$. For every fixed $q_j$, the residual $r_j(x) = \|T_x(q_j) - q_j\|_2$ is measurable. For each integer $n \geq 1$, let $j_n(x)$ be the least $j$ with $r_j(x) < 1/n$. Such an index exists by continuity of the residual and existence of a fixed point. It is measurable since

$$\{x : j_n(x) = j\} = \{x : r_j(x) < 1/n\} \cap \bigcap_{i<j} \{x : r_i(x) \geq 1/n\}.$$

Thus $q_{j_n(x)}$ is a measurable vector-valued function of $x$. For a fixed context, every convergent subsequence of these vectors has residual zero at its limit, by continuity of $T_x$. The unique fixed point is $\pi_P(\cdot \mid x)$, so compactness of the simplex implies that the entire sequence converges to it: a subsequence staying a positive distance away would have a further convergent subsequence with a different limit. Hence $\pi_P$ is the pointwise limit of measurable vectors and is measurable. □

Define the temperature-zero selector and regret by

$$a_P(x) \in \arg\max_{a\in S_x} \phi_P(x, a), \qquad \mathcal{G}_P^{\mathrm{GP}}(Q) = \mathbb{E}_{X\sim d_0}\big[\phi_P(X, a_P(X)) - \phi_P(X, a_Q(X))\big].$$

These selectors are measurable because the action set is finite, the scores and support indicators are measurable, and ties are resolved by a fixed ordering. Since every score lies in $[0, 1]$, $0 \leq \mathcal{G}_P^{\mathrm{GP}}(Q) \leq 1$. For $p, q \in [0, 1]$, use the squared Hellinger convention

$$h^2(p, q) \coloneqq \tfrac{1}{2}\left[(\sqrt{p} - \sqrt{q})^2 + (\sqrt{1-p} - \sqrt{1-q})^2\right] = 1 - \sqrt{pq} - \sqrt{(1-p)(1-q)}.$$

Thus $0 \leq h^2 \leq 1$, and $h^2(p, q) = 0$ exactly when $p = q$. Define the reference-comparison separation

$$H_0^2(P, Q) \coloneqq \mathbb{E}_{X\sim d_0} \sum_{a,b\in S_X} \pi_0(a \mid X)\pi_0(b \mid X) h^2(P(X, a, b), Q(X, a, b)). \tag{26}$$

**Lemma 23** (Reference identifiability of the Nash selector)**.** *In the finite-action setting, $H_0^2(P, Q) = 0$ implies $\mathcal{G}_P^{\mathrm{GP}}(Q) = 0$. Consequently, in a finite model class, positive-regret pairs have a uniform strictly positive reference Hellinger separation whenever any such pairs exist.*

*Proof.* The integrand in (26) is nonnegative. If its expectation is zero, then it is zero for $d_0$-almost every context. At each such context the sum is finite, and the coefficient $\pi_0(a \mid x)\pi_0(b \mid x)$ is strictly positive for every $a, b \in S_x$. Therefore every summand has $h^2(P(x, a, b), Q(x, a, b)) = 0$, and hence $P(x, a, b) = Q(x, a, b)$ for all supported pairs.

The two regularized games at this context are identical on the reference-supported simplex. Uniqueness in Lemma 22 gives $\pi_P = \pi_Q$ there. Substitution into the score formula shows $\phi_P(x, a) = \phi_Q(x, a)$ for every supported action. Thus any maximizer selected under $Q$ is also a maximizer under $P$, and the pointwise regret is zero. Integrating gives $\mathcal{G}_P^{\mathrm{GP}}(Q) = 0$.

Its contrapositive says that every positive-regret pair has $H_0^2(P, Q) > 0$. There are finitely many ordered pairs in a finite class, so the minimum of these strictly positive numbers, over a nonempty set of positive-regret pairs, is strictly positive. □

We now specify the greedy MLE-to-Nash procedure. Choose any fixed $\widehat{P}_0 \in \mathfrak{P}_{\mathrm{GP}}$ for the initial selector, and initialize the sampling policy at $\pi_1 = \pi_0$. Under truth $P$, round $s$ is generated conditionally on the preceding history by

$$X_s \sim d_0, \qquad (A_s, B_s) \sim \pi_s(\cdot \mid X_s)^{\otimes 2}, \qquad Y_s \sim \mathrm{Bernoulli}(P(X_s, A_s, B_s)).$$

Here $Y_s = 1$ means that the first action is preferred. For a candidate $Q$, define

$$f_Q(y \mid x, a, b) = \begin{cases} Q(x, a, b), & y = 1, \\ 1 - Q(x, a, b), & y = 0, \end{cases} \qquad \Lambda_t(Q) = \prod_{s=1}^{t} f_Q(Y_s \mid X_s, A_s, B_s).$$

Choose a measurable exact maximizer

$$\widehat{P}_t \in \arg\max_{Q \in \mathfrak{P}_{\mathrm{GP}}} \Lambda_t(Q), \qquad \pi_{t+1} = \pi_{\widehat{P}_t}.$$

A fixed ordering of the finite class provides a measurable maximizer, including in the presence of ties. The product formulation avoids undefined logarithms when a candidate assigns zero probability to an observed label. It is equivalent to minimizing the extended-valued negative log-likelihood.

**Theorem 24** (Finite-class general-preference guarantee)**.** *Let $M = |\mathfrak{P}_{\mathrm{GP}}|$, and define*

$$\mathcal{B} = \{(P, Q) \in \mathfrak{P}_{\mathrm{GP}}^2 : \mathcal{G}_P^{\mathrm{GP}}(Q) > 0\}.$$

*If $\mathcal{B} \neq \varnothing$, set $h_\star = \min_{(P,Q)\in\mathcal{B}} H_0^2(P, Q) > 0$; if $\mathcal{B} = \varnothing$, set $h_\star = 1$. Let $\kappa = e^{-2\eta} h_\star > 0$. Then, for every truth $P \in \mathfrak{P}_{\mathrm{GP}}$ and $t \geq 1$,*

$$\mathbb{P}_P\Big(\mathcal{G}_P^{\mathrm{GP}}(\widehat{P}_t) > 0\Big) \leq (M-1)e^{-\kappa t}, \qquad \mathbb{E}_P[\mathcal{G}_P^{\mathrm{GP}}(\widehat{P}_t)] \leq \min\{1, (M-1)e^{-\kappa t}\}.$$

*In particular,*

$$\sup_{P\in\mathfrak{P}_{\mathrm{GP}}} \sup_{T\geq 1} \sum_{t=0}^{T-1} \mathbb{E}_P[\mathcal{G}_P^{\mathrm{GP}}(\widehat{P}_t)] \leq 1 + \frac{M-1}{e^\kappa - 1} < \infty.$$

*Proof.* If $\mathcal{B}$ is empty, every candidate has zero regret under every truth, so all the displayed claims hold. Otherwise Lemma 23 justifies $h_\star > 0$. Fix a truth $P$, and consider any candidate $Q$ for which $\mathcal{G}_P^{\mathrm{GP}}(Q) > 0$. Let $\mathscr{H}_s$ be the history through round $s$, and abbreviate $Z_s = (X_s, A_s, B_s)$.

At round one the sampling density relative to the reference is one. At later rounds, Lemma 22 gives $\pi_s(a \mid x) \geq e^{-\eta}\pi_0(a \mid x)$ for every supported action. Thus, at every round and conditional on every admissible preceding history, the pair probabilities are at least $e^{-2\eta}$ times their reference counterparts. Since $h^2$ is nonnegative, integrating this product bound yields

$$\begin{aligned}
&\mathbb{E}_P[h^2(P(Z_s), Q(Z_s)) \mid \mathscr{H}_{s-1}] \\
&\quad \geq e^{-2\eta}\mathbb{E}_{X\sim d_0} \sum_{a,b\in S_X} \pi_0(a \mid X)\pi_0(b \mid X) h^2(P(X, a, b), Q(X, a, b)) \\
&\quad = e^{-2\eta} H_0^2(P, Q) \geq e^{-2\eta} h_\star = \kappa.
\end{aligned} \tag{27}$$

For each observation, the true conditional probability of its realized label is positive almost surely: conditionally on $Z_s$, the probability of seeing a label assigned probability zero is the sum of its zero probabilities. A countable union over rounds shows that $\Lambda_t(P) > 0$ simultaneously for all finite $t$, almost surely. On this event define

$$W_t(Q) = \sqrt{\frac{\Lambda_t(Q)}{\Lambda_t(P)}}, \qquad W_0(Q) = 1.$$

The values on the null complement may be set arbitrarily. Conditionally on the preceding history and the current pair, direct summation over the two possible labels gives

$$\begin{aligned}
&\mathbb{E}_P\left[\sqrt{\frac{f_Q(Y_s \mid Z_s)}{f_P(Y_s \mid Z_s)}} \,\middle|\, \mathscr{H}_{s-1}, Z_s\right] \\
&\quad = \sum_{y:\, f_P(y|Z_s)>0} \sqrt{f_P(y \mid Z_s) f_Q(y \mid Z_s)} \\
&\quad = \sqrt{P(Z_s)Q(Z_s)} + \sqrt{(1 - P(Z_s))(1 - Q(Z_s))} = 1 - h^2(P(Z_s), Q(Z_s)).
\end{aligned}$$

A term with zero true probability contributes zero to the affinity; this is why the formula remains valid when preference probabilities are at the boundary of $[0, 1]$.

The likelihood product identity gives $W_s(Q) = W_{s-1}(Q)\sqrt{f_Q(Y_s \mid Z_s)/f_P(Y_s \mid Z_s)}$. Since $W_{s-1}(Q)$ is nonnegative and $\mathscr{H}_{s-1}$-measurable, the tower property and (27) imply

$$\begin{aligned}
\mathbb{E}_P[W_s(Q) \mid \mathscr{H}_{s-1}] &= W_{s-1}(Q)\left(1 - \mathbb{E}_P[h^2(P(Z_s), Q(Z_s)) \mid \mathscr{H}_{s-1}]\right) \\
&\leq (1 - \kappa) W_{s-1}(Q).
\end{aligned}$$

Iterating expectations gives

$$\mathbb{E}_P[W_t(Q)] \le (1-\kappa)^t \le e^{-\kappa t}.$$

Here $0 < \kappa < 1$, since $h_\star \le 1$ and $\eta > 0$, and $1 - \kappa \le e^{-\kappa}$. In particular these expectations are finite. The sampling policies need not be independent of the past: the contraction was established conditionally on that past. The likelihood comparison uses the same realized sampling rule for both candidates; it does not replace the observed policies by policies that would have been generated under a different candidate.

Exact maximum likelihood and feasibility of the truth imply $\Lambda_t(\widehat{P}_t) \ge \Lambda_t(P)$. Hence $\{\widehat{P}_t = Q\} \subseteq \{W_t(Q) \ge 1\}$, and Markov's inequality gives

$$\mathbb{P}_P(\widehat{P}_t = Q) \le \mathbb{P}_P(W_t(Q) \ge 1) \le \mathbb{E}_P[W_t(Q)] \le e^{-\kappa t}.$$

Only candidates with positive regret need to be included. There are at most $M - 1$ such candidates, since the truth has zero regret under itself. A union bound therefore yields

$$\mathbb{P}_P\Big(\mathcal{G}_P^{\mathrm{GP}}(\widehat{P}_t) > 0\Big) \le \sum_{Q:\, \mathcal{G}_P^{\mathrm{GP}}(Q) > 0} \mathbb{P}_P(\widehat{P}_t = Q) \le (M-1)e^{-\kappa t}.$$

Since $0 \le \mathcal{G}_P^{\mathrm{GP}}(Q) \le 1$, its expectation is at most the probability of positive regret and is also at most one. This proves the two one-round bounds.

The initialization contributes at most one. Extending the remaining nonnegative sum to all positive times and evaluating the geometric series gives

$$\begin{aligned}\sum_{t=0}^{T-1} \mathbb{E}_P[\mathcal{G}_P^{\mathrm{GP}}(\widehat{P}_t)] &\le 1 + (M-1)\sum_{t=1}^{\infty} e^{-\kappa t} \\ &= 1 + (M-1)\frac{e^{-\kappa}}{1-e^{-\kappa}} = 1 + \frac{M-1}{e^{\kappa}-1}.\end{aligned}$$

The constants $M$, $h_\star$, and $\kappa$ do not depend on the realized truth or horizon. Taking both suprema completes the proof. □