

CORRECTING THE MYTHOS OF KL-REGULARIZATION: DIRECT ALIGNMENT WITHOUT OVEROPTIMIZATION VIA χ^2 -PREFERENCE OPTIMIZATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Language model alignment methods, such as reinforcement learning from human feedback (RLHF), have led to impressive advances in language model capabilities. However, existing techniques are limited by a widely observed phenomenon known as *overoptimization*, where the quality of the language model degrades over the course of the alignment process. Overoptimization occurs when a language model overfits to inaccuracies in an (either explicit or implicit) offline reward model, and drifts away from preferred responses covered by the data. To discourage such distribution shift, offline alignment methods typically employ KL-regularization, but this, as we show, is too weak to prevent degradation in performance. Then, can we design an efficient algorithm that is provably robust to overoptimization?

In this paper, we advance theoretical understanding of sample-efficient offline alignment and introduce a new algorithm called χ^2 -Preference Optimization (χ PO). χ PO is a one-line change to Direct Preference Optimization (DPO; Rafailov et al. (2023)), that modifies only the logarithmic link function in the DPO objective. Despite this minimal change, χ PO implicitly implements the principle of *pessimism in the face of uncertainty* via regularization with the χ^2 -divergence—which quantifies uncertainty more effectively than KL-regularization—and provably alleviates overoptimization, achieving sample-complexity guarantees based on *single-policy concentrability*—the gold standard in offline reinforcement learning. This guarantee makes χ PO the first simple, yet general-purpose offline alignment algorithm that is provably robust to overoptimization.

1 INTRODUCTION

Large language models (LLMs) trained on unsupervised text data exhibit impressive and surprising capabilities (Brown et al., 2020; Ouyang et al., 2022; Touvron et al., 2023; OpenAI, 2023; Google, 2023), but can be difficult to control without further guidance. *Reinforcement learning from human feedback (RLHF)* and other alignment methods have emerged as a central tool to align these models to human values and elicit desired behavior (Christiano et al., 2017; Bai et al., 2022; Ouyang et al., 2022; Rafailov et al., 2023). This is achieved by treating the language model as a *policy*, and using techniques from reinforcement learning to optimize for desirable outcomes under a (explicit or implicit) reward model learned from a dataset of human-labeled responses.

Alignment methods like RLHF have led to significant advances in language model capabilities, particularly in chat domains, but existing techniques are limited by a widely observed phenomenon known as *reward overoptimization* or *reward hacking* (Michaud et al., 2020; Tien et al., 2022; Gao et al., 2023; Rafailov et al., 2024a). Since the reward model used implicitly or explicitly by these methods is an imperfect proxy for human preferences, the true quality of the language model can degrade as training proceeds, even as performance under the reward model continues to improve. Intuitively, this occurs because the language model may drift away from the manifold covered by the human-labeled data used to train the reward model and end up in a region where the reward model is inaccurate.

Overoptimization is distinct from the classical concept of overfitting because it is a causal or counterfactual phenomenon: When the human-labeled dataset does not cover all possible alternatives, the decision maker—in this case, a language model policy—cannot directly evaluate the effect of their actions. This perspective is supported by the fact that overoptimization can be mitigated by *online*

alignment techniques (Guo et al., 2024; Gao et al., 2024; Dong et al., 2024), which exploit interactive access to human or AI feedback to iteratively improve the reward model; unfortunately, gathering such feedback is costly and impractical in many settings. This raises natural theoretical questions regarding the role of overoptimization in *offline alignment*:

- Is overoptimization in offline alignment an *information-theoretic phenomenon*? This would mean that there is simply not enough information in the human-labeled (offline) preference dataset due to partial coverage, and no algorithmic intervention can avoid the overoptimization issue.
- Alternatively, is overoptimization an *algorithmic phenomenon*? This would mean that existing algorithms are not making the most of the data they have (e.g., due to optimizing the wrong objective and converging toward suboptimal solutions) and would suggest that their *sample-efficiency* can be improved, perhaps by taking more aggressive measures to avoid overfitting to the reward model.

Previous developments in the theory of offline reinforcement learning suggest that the answer may be the latter. Indeed, this literature has addressed the challenge of overoptimization—typically referred to as *distribution shift*—through the principle of *pessimism in the face of uncertainty*, which asserts that, given an offline dataset with partial coverage, a decision maker should choose their response according to the most pessimistic view of the world supported by the data. Pessimism encourages the model to avoid overfitting to the offline dataset and is supported by a rich theory offering provable robustness to overoptimization in stylized settings (Liu et al., 2020; Jin et al., 2021).

Perhaps the greatest barrier to implementing pessimism in language models is the efficient quantification of uncertainty in the offline reward, and the distillation of this information into actionable form. Most existing offline alignment methods employ KL-regularization, which penalizes the learned policy for drifting from the reference policy, but this form of uncertainty quantification is insufficient to induce pessimism (Gao et al., 2023) and is provably suboptimal in theory (Zhu et al., 2023; Song et al., 2024, see also Appendix A.1). On the other hand, offline reinforcement learning theory offers abstract pessimistic algorithms that are suitable—at least statistically—for large models (Xie et al., 2021; Uehara and Sun, 2021; Zhan et al., 2022; Chen and Jiang, 2022), but cannot be implemented directly without losing theoretical fidelity or making unrealistic modeling assumptions (Zhu et al., 2023; Zhan et al., 2023b; Li et al., 2023; Xiong et al., 2023; Liu et al., 2024; Cen et al., 2024; Fisch et al., 2024; Ji et al., 2024). Notably, the so-called “DPO+SFT” approach developed by Liu et al. (2024); Cen et al. (2024); Fisch et al. (2024) is provably suboptimal unless the language model satisfies an unrealistic convexity property (Appendix A.1). Thus we ask: *If we instead leverage the unique structure of the language modeling problem, can we develop simple, yet efficient, offline alignment methods that are certifiably robust to overoptimization?*

1.1 CONTRIBUTIONS

We introduce a new theoretical algorithm for offline alignment, χ^2 -Preference Optimization (χ PO). χ PO is simple and straightforward to implement, requiring only a single-line change to Direct Preference Optimization (Rafailov et al. (2023)), yet it is provably robust to overoptimization. Algorithmically, χ PO only differs from DPO in that we replace the usual logarithmic link function in the DPO objective with a new link function that implicitly implements pessimism via regularization with the χ^2 -divergence—a divergence that (i) plays a fundamental role in statistics due to its ability to quantify uncertainty (Tsybakov, 2008); and (ii) penalizes off-manifold behavior more effectively than KL-regularization. Statistically, we formalize robustness to overoptimization via a sample complexity guarantee based on *single-policy concentrability*—the gold standard in offline reinforcement learning—which we establish under minimal statistical and function approximation assumptions. This result implies that, in contrast to most prior work, χ PO enjoys meaningful guarantees even when the reference policy has poor coverage. Summarizing:

χ PO is the first simple, yet general-purpose algorithm for offline alignment with provable robustness to overoptimization.

The result above concerns the classical language model alignment formulation, which assumes the Bradley-Terry preference model (Christiano et al., 2017; Bai et al., 2022; Ouyang et al., 2022; Rafailov et al., 2023). Turning our attention to general preference models (Munos et al., 2023; Swamy et al., 2024; Rosset et al., 2024) where the goal is to find an approximate Nash equilibrium, we show (Appendix D) that *achieving guarantees based on single-policy concentrability is impossible*. Nonetheless, we show that an iterative variant of χ PO based on self-play achieves a sample complexity

guarantee that scales with a new local coverage condition—a condition that is stronger than single policy concentrability, but much weaker than global concentrability and the notion of unilateral concentrability introduced by Cui and Du (2022). This result provides additional evidence for the value of regularization with χ^2 -divergence for obtaining sharp sample complexity guarantees.

Technical highlights. Our analysis of χ PO leverages several new techniques. First, we show that RLHF with χ^2 -regularization is sufficient to achieve guarantees based on single-policy concentrability (Section 3.1 and Appendix C). Next, we show that a variant of the DPO reparameterization trick that combines χ^2 -regularization with KL-regularization (“mixed” χ^2 -regularization) can be used to reformulate our objective into a purely policy-based objective, in spite of the fact that χ^2 -regularization fails to satisfy certain regularity conditions found in prior work (Wang et al., 2023a). Finally, and perhaps most importantly, we use a novel analysis to show that pessimism is preserved after reparameterization. Compared to prior approaches to pessimism in offline RL (Xie et al., 2021; Uehara and Sun, 2021; Zhan et al., 2022; Chen and Jiang, 2022), χ^2 -regularization strikes a useful balance between generality and tractability, and we expect our techniques to find broader use.

2 BACKGROUND

In this section, we provide necessary background and highlight that standard algorithms in offline alignment suffer from overoptimization. We adopt standard big-oh notation, and write $f = \tilde{O}(g)$ to denote that $f = O(g \cdot \max\{1, \text{polylog}(g)\})$ and $a \lesssim b$ as shorthand for $a = O(b)$.

2.1 ALIGNMENT FROM HUMAN FEEDBACK

Following prior work (e.g., Rafailov et al. (2023); Ye et al. (2024)), we adopt a contextual bandit formulation of the alignment problem. We formalize the language model as a *policy* $\pi : \mathcal{X} \rightarrow \Delta(\mathcal{A})$ which maps a context (prompt) $x \in \mathcal{X}$ to an action (response) $a \in \mathcal{A}$ via $a \sim \pi(\cdot | x)$, and let $\rho \in \Delta(\mathcal{X})$ denote the distribution over contexts/prompts.

Offline alignment. In the offline alignment problem (Christiano et al., 2017; Bai et al., 2022; Ouyang et al., 2022), we assume access to a dataset $\mathcal{D}_{\text{pref}} = \{(x, a_+, a_-)\}$ of n prompts and labeled response pairs generated from a reference policy (language model) π_{ref} , which is typically obtained through SFT. Here, a_+ is a positive action/response and a_- is a negative action/response. Given the context/prompt $x \sim \rho$, the pair (a_+, a_-) is generated by sampling a pair (a, b) as $a \sim \pi_{\text{ref}}(\cdot | x)$ and $b \sim \pi_{\text{ref}}(\cdot | x)$, and then ordering them as (a_+, a_-) based on a binary preference $y \sim \mathbb{P}(a \succ b | b)$. We assume that preferences follow the *Bradley-Terry* model (Bradley and Terry, 1952):

$$\mathbb{P}(a \succ b | x) = \frac{\exp(r^*(x, a))}{\exp(r^*(x, a)) + \exp(r^*(x, b))}, \quad (1)$$

for an unknown reward function $r^* : \mathcal{X} \times \mathcal{A} \rightarrow [0, R_{\max}]$ for some $R_{\max} \geq 1$. From the preference dataset $\mathcal{D}_{\text{pref}}$, we aim to learn a policy $\hat{\pi}$ that has high reward in the sense that $J(\pi^*) - J(\hat{\pi}) \leq \varepsilon$ for a small $\varepsilon > 0$, where $J(\pi) := \mathbb{E}_{x \sim \rho, a \sim \pi(\cdot | x)}[r^*(x, a)]$ is the true expected reward, and π^* is any comparator policy of interest. We abbreviate $\mathbb{E}_{\pi}[\cdot] := \mathbb{E}_{x \sim \rho, a \sim \pi(\cdot | x)}[\cdot]$, and assume that $\rho(x) > 0$ for all x and $\pi_{\text{ref}}(a | x) > 0$ for all x, a without loss of generality.

Offline RLHF with KL-regularization. Classical algorithms for offline alignment (Christiano et al., 2017; Ouyang et al., 2022) are based on reinforcement learning with a *KL-regularized* reward objective, defined for a regularization parameter $\beta > 0$, via

$$J_{\beta}^{\text{KL}}(\pi) := J(\pi) - \beta \cdot D_{\text{KL}}(\pi \| \pi_{\text{ref}}) = \mathbb{E}_{\pi} \left[r^*(x, a) - \beta \log \frac{\pi(a | x)}{\pi_{\text{ref}}(a | x)} \right], \quad (2)$$

where we adopt the shorthand $D_{\text{KL}}(\pi \| \pi') = \mathbb{E}_{x \sim \rho} [D_{\text{KL}}(\pi(\cdot | x) \| \pi'(\cdot | x))]$. These methods first estimate a reward function \hat{r} from $\mathcal{D}_{\text{pref}}$ using maximum likelihood under the Bradley-Terry model:

$$\hat{r} = \operatorname{argmax}_{r \in \mathcal{R}} \sum_{(x, a_+, a_-) \in \mathcal{D}_{\text{pref}}} \log \sigma(r(a_+ | x) - r(a_- | x)), \quad (3)$$

where $\sigma(x) := \frac{\exp(x)}{1 + \exp(x)}$ is the sigmoid function and \mathcal{R} is a class of reward functions, which is typically parameterized by a neural network. Then, they apply standard policy optimization methods like PPO to optimize an estimated version of Eq. (2): $\hat{\pi} = \operatorname{argmax}_{\pi \in \Pi} \mathbb{E}_{\pi} [\hat{r}(x, a) - \beta \log \frac{\pi(a | x)}{\pi_{\text{ref}}(a | x)}]$. The regularization term in Eq. (2) is intended to encourage $\hat{\pi}$ to stay close to π_{ref} , with the hope of preventing the policy from overfitting to the potentially inaccurate reward model \hat{r} .

Direct preference optimization (DPO). χ PO is based on an alternative offline alignment approach, Direct Preference Optimization (DPO; Rafailov et al. (2023)). DPO uses the closed-form solution of the optimal KL-regularized policy under the objective Eq. (2)—which can be viewed as implicitly modeling rewards—to define a single policy optimization objective that removes the need for direct reward function estimation. Given a user specified policy class Π , DPO solves

$$\hat{\pi}_{\text{DPO}} = \operatorname{argmax}_{\pi \in \Pi} \sum_{(x, a_+, a_-) \in \mathcal{D}_{\text{pref}}} \log \left[\sigma \left(\beta \log \frac{\pi(a_+ | x)}{\pi_{\text{ref}}(a_+ | x)} - \beta \log \frac{\pi(a_- | x)}{\pi_{\text{ref}}(a_- | x)} \right) \right], \quad (4)$$

with the convention that the value of the objective is $-\infty$ if π does not satisfy $\pi \ll \pi_{\text{ref}}$.

2.2 OVEROPTIMIZATION AND INSUFFICIENCY OF KL-REGULARIZATION

Empirically, both classical RLHF and direct alignment methods like DPO have been observed to suffer from overoptimization (Gao et al., 2023; Guo et al., 2024; Rafailov et al., 2024a; Song et al., 2024), wherein model quality degrades during the optimization process as the learned policy drifts away from π_{ref} . This can be mitigated by *online alignment* techniques (Gao et al., 2024; Guo et al., 2024; Dong et al., 2024; Xie et al., 2024), which collect labeled preference data on-policy during training, but there are many settings where this is impractical or infeasible. As we will see, the overoptimization phenomena in offline alignment methods is an issue of sample-inefficiency, which can be understood through the lens of coverage coefficients developed in the theory of offline reinforcement learning (Liu et al., 2020; Jin et al., 2021; Rashidinejad et al., 2021). In particular, the performance of existing offline alignment algorithms depends on how well data covers all candidate policies, and degrades when coverage is inadequate or the number of samples is insufficiently large.

Coverage coefficients. In offline reinforcement learning theory, the sample efficiency of an algorithm refers to the number of samples required to guarantee that $J(\hat{\pi}) \approx J(\pi^*)$. It is typically quantified by a *coverage coefficient* (or concentrability coefficient) that measures the quality of the data collected by the reference π_{ref} (Farahmand et al., 2010; Xie and Jiang, 2020; Zanette et al., 2021). We will utilize the L_1 coverage coefficient, defined as $\mathcal{C}^\pi := \mathbb{E}_\pi \left[\frac{\pi(a|x)}{\pi_{\text{ref}}(a|x)} \right]$ for a policy π . *Single policy concentrability* is the gold standard for sample efficiency, and is obtained by an algorithm if, for any comparator policy π^* , the sample size required to learn $J(\hat{\pi}) \approx J(\pi^*)$ scales with \mathcal{C}^{π^*} , the coverage coefficient of π^* . This guarantees that $\hat{\pi}$ is competitive with the best policy that is sufficiently covered by offline data, and, further, is never much worse than π_{ref} itself. Single policy concentrability is typically achieved by pessimistic algorithms that penalize the valuations of candidate policies according to their uncertainty. In contrast, non-pessimistic algorithms typically suffer *all-policy concentrability*, which means that the number of samples they require scales with $\max_{\pi \in \Pi} \mathcal{C}^\pi$ (Liu et al., 2020; Jin et al., 2021; Rashidinejad et al., 2021). Unless the offline data adequately covers all candidate policies, these algorithms may learn suboptimal policies that overfit to inaccurate reward signals, which manifests as the observed phenomenon of overoptimization. As such, single policy concentrability serves as a theoretical certification that an algorithm is robust to poor data coverage, and will not overoptimize.

Pessimism in offline alignment. Zhu et al. (2023) show that PPO and DPO suffer all-policy concentrability and have sample complexities that scale with $\max_{\pi} \mathcal{C}_\infty^\pi$, for the stylized case of alignment with linearly parameterized policies where $\pi_\theta(a | x) \propto \exp(\langle \phi(x, a), \theta \rangle)$ for a known feature embedding $\phi(x, a) \in \mathbb{R}^d$ (see also Zhu et al. (2024); Song et al. (2024)). They also propose a pessimistic algorithm that achieves single policy concentrability, or $J(\pi^*) - J(\hat{\pi}) \lesssim \sqrt{\frac{\text{poly}(\mathcal{C}_\infty^{\pi^*}, d)}{n}}$ simultaneously for all π^* . While encouraging, these results are restricted to linearly parameterized policies, and cannot be directly applied to large language models. Most existing theoretical algorithms for offline alignment are similar in nature, and either place restrictive assumptions on the policy class Π (Zhu et al., 2023; Zhan et al., 2023b; Li et al., 2023; Xiong et al., 2023) or are not feasible to implement in a way that is faithful to theory (Ye et al., 2024; Ji et al., 2024).

Most relevant to our work, a series of recent papers (Liu et al., 2024; Cen et al., 2024; Fisch et al., 2024) propose implementing pessimism for general policy classes Π by solving the ‘‘DPO+SFT’’ objective

$$\operatorname{argmax}_{\pi \in \Pi} \left\{ \alpha \cdot \mathbb{E}_{\pi_{\text{ref}}} [\beta \log \pi(a | x)] + \frac{1}{n} \sum_{(x, a_+, a_-) \in \mathcal{D}_{\text{pref}}} \log \left[\sigma \left(\beta \log \frac{\pi(a_+ | x)}{\pi_{\text{ref}}(a_+ | x)} - \beta \log \frac{\pi(a_- | x)}{\pi_{\text{ref}}(a_- | x)} \right) \right] \right\}, \quad (5)$$

which augments the DPO objective (the second term) with an additional supervised fine-tuning-like (SFT) loss (the first term). While this objective is simple to apply to general policy classes, the existing single-policy concentrability guarantees for this method assume that Π satisfies restrictive *convexity* conditions which do not hold in practice for large language models. Perhaps surprisingly, we show (Appendix A.1) that without convexity, *the objective in Eq. (5) fails to achieve a single-policy concentrability guarantee.*¹ In other words, DPO+SFT is insufficient to mitigate overoptimization.

3 χ^2 -PREFERENCE OPTIMIZATION

This section presents our main algorithm, χ PO. We begin by introducing χ^2 -regularization as a general framework for mitigating overoptimization in offline alignment (Section 3.1), then derive the χ PO algorithm (Section 3.2) and finally present our main theoretical guarantee (Section 3.3).

3.1 FRAMEWORK: χ^2 -REGULARIZED REWARD OPTIMIZATION

The central algorithm design principle for our work is to (implicitly or explicitly) optimize a variant of the classical RLHF objective (Eq. (2)) that replaces KL-regularization with regularization via χ^2 -divergence, defined for a pair of probability measures \mathbb{P} and \mathbb{Q} with $\mathbb{P} \ll \mathbb{Q}$ via $D_{\chi^2}(\mathbb{P} \parallel \mathbb{Q}) := \frac{1}{2} \int \left(\frac{d\mathbb{P}}{d\mathbb{Q}} - 1 \right)^2 d\mathbb{Q}$. χ^2 -divergence is a more aggressive form of regularization than KL-divergence; we have $D_{\text{KL}}(\mathbb{P} \parallel \mathbb{Q}) \leq 2D_{\chi^2}(\mathbb{P} \parallel \mathbb{Q})$, but the converse is not true in general. We consider the following χ^2 -regularized RL objective:²

$$J_{\beta}^{\chi}(\pi) := \mathbb{E}_{\pi}[r^*(x, a)] - \beta \cdot D_{\chi^2}(\pi \parallel \pi_{\text{ref}}). \quad (6)$$

Moving to a form of regularization that penalizes deviations from π_{ref} more forcefully than KL-regularization is a natural approach to mitigating overoptimization, but an immediate concern is that this may lead to overly conservative algorithms. As we will show, however, χ^2 -divergence is better suited to the geometry of offline alignment, as it has the unique property (not shared by KL-divergence) that its value quantifies the extent to which the accuracy of a reward model \hat{r} trained under π_{ref} will transfer to a downstream policy π of interest (Lemma H.3). This implies that the χ^2 -regularized RL objective in Eq. (6) meaningfully implements a form of pessimism in the face of uncertainty, and by tuning the regularization parameter $\beta > 0$, we can keep the learned policy $\hat{\pi}$ close to π_{ref} in the “right” (uncertainty-aware) way. As such, we view optimizing χ^2 -regularized rewards, i.e., $\text{argmax}_{\pi \in \Pi} J_{\beta}^{\chi}(\pi)$ as a general principle to guide algorithm design for offline alignment (as well as offline RL more broadly), which we expect to find broader use.

We now turn our attention to the matter of how to optimize this objective. One natural approach, in the vein of classical RLHF (Christiano et al., 2017; Ouyang et al., 2022), is to estimate a reward model \hat{r} using maximum likelihood (Eq. (3)), and then use PPO or other policy optimization methods to solve

$$\hat{\pi} = \text{argmax}_{\pi \in \Pi} \mathbb{E}_{\pi} [\hat{r}(x, a)] - \beta \cdot D_{\chi^2}(\pi \parallel \pi_{\text{ref}}) = \text{argmax}_{\pi \in \Pi} \mathbb{E}_{\pi} \left[\hat{r}(x, a) - \beta \frac{\pi(a | x)}{\pi_{\text{ref}}(a | x)} \right]. \quad (7)$$

While this indeed leads to meaningful statistical guarantees (cf. Appendix C), we adopt a simpler and more direct approach inspired by DPO, which removes the need for a separate reward estimation step.

3.2 THE χ PO ALGORITHM

Our main algorithm, χ PO, is described in Algorithm 1. Given a preference dataset $\mathcal{D}_{\text{pref}}$ and policy class Π , the algorithm learns a policy $\hat{\pi}$ by solving the DPO-like optimization objective Eq. (9), which replaces the usual $\log \frac{\pi(a|x)}{\pi_{\text{ref}}(a|x)}$ terms in the original DPO objective (Eq. (4)) with a new link function:

$$\phi \left(\frac{\pi(a | x)}{\pi_{\text{ref}}(a | x)} \right) = \frac{\pi(a | x)}{\pi_{\text{ref}}(a | x)} + \log \left(\frac{\pi(a | x)}{\pi_{\text{ref}}(a | x)} \right).$$

A secondary modification is that we handle potentially unbounded density ratios by clipping to the interval $[-2R_{\text{max}}, +2R_{\text{max}}]$ via the operator $\text{clip}_R(z) = \max\{\min\{R, z\}, -R\}$. In what follows, we will show that this simple modification to DPO—that is, incorporating an additional density ratio term outside the logarithm—implicitly implements pessimism via χ^2 -regularization.

¹This finding is surprising because Xie et al. (2024) show that an *optimistic online* counterpart to Eq. (5), which negates the SFT term, enjoys online RLHF guarantees without requiring analogous convexity conditions.

²We write $D_{\chi^2}(\pi \parallel \pi_{\text{ref}}) = \mathbb{E}[D_{\chi^2}(\pi(\cdot | x) \parallel \pi_{\text{ref}}(\cdot | x))] = \mathbb{E}_{\pi} \left[\frac{\pi(a|x)}{\pi_{\text{ref}}(a|x)} \right]$.

Algorithm 1 χ^2 -Preference Optimization (χ PO)

input: Reference policy π_{ref} , preference dataset $\mathcal{D}_{\text{pref}}$, χ^2 -regularization coefficient $\beta > 0$.
 1: Define

$$\phi(z) := z + \log z. \quad (8)$$

2: Optimize χ^2 -regularized preference optimization objective:

$$\hat{\pi} \leftarrow \operatorname{argmax}_{\pi \in \Pi} \sum_{(x, a_+, a_-) \in \mathcal{D}_{\text{pref}}} \log \left[\sigma \left(\operatorname{clip}_{2, R_{\max}} \left[\beta \phi \left(\frac{\pi(a_+ | x)}{\pi_{\text{ref}}(a_+ | x)} \right) - \beta \phi \left(\frac{\pi(a_- | x)}{\pi_{\text{ref}}(a_- | x)} \right) \right] \right) \right]. \quad (9)$$

3: **return:** $\hat{\pi}$.

Algorithm derivation. Recall that DPO is derived (Rafailov et al., 2023) by observing that the optimal KL-regularized policy $\pi_{\beta; \text{KL}}^* := \operatorname{argmax}_{\pi} \{ \mathbb{E}_{\pi} [r^*(x, a)] - \beta D_{\text{KL}}(\pi \| \pi_{\text{ref}}) \}$ satisfies $r^*(x, a) = \beta \log \frac{\pi_{\beta; \text{KL}}^*(a|x)}{\pi_{\text{ref}}(a|x)} + Z_{\beta, r^*; \text{KL}}(x)$ for all $x \in \mathcal{X}$ and $a \in \mathcal{A}$ where $Z_{\beta, r^*; \text{KL}}(x)$ is a normalization constant that depends on x but not a . This facilitates reparameterizing the reward model in the maximum likelihood estimation objective (Eq. (3)) in terms of a learned policy, yielding the DPO objective in Eq. (4).

To apply a similar reparameterization trick for χ^2 -divergence, a natural starting point is an observation from Wang et al. (2023a), who show that an analogous characterization for the optimal regularized policy holds for a general class of f -divergences. For a convex function $f : \mathbb{R}_+ \rightarrow \mathbb{R}$, define the induced f -divergence by $D_f(\mathbb{P} \| \mathbb{Q}) = \int f \left(\frac{d\mathbb{P}}{d\mathbb{Q}} \right) d\mathbb{Q} = \mathbb{E}_{\mathbb{Q}} \left[f \left(\frac{d\mathbb{P}}{d\mathbb{Q}} \right) \right]$. Wang et al. (2023a) show that for any differentiable f that satisfies the technical condition $0 \notin \operatorname{dom}(f')$, the optimal f -regularized policy $\pi_{\beta; f}^* = \operatorname{argmax}_{\pi} \{ \mathbb{E}_{\pi} [r^*(x, a)] - \beta D_f(\pi \| \pi_{\text{ref}}) \}$ satisfies

$$r^*(x, a) = \beta f' \left(\frac{\pi_{\beta; f}^*(a | x)}{\pi_{\text{ref}}(a | x)} \right) + Z_{\beta, r^*; f}(x) \quad (10)$$

for a normalization constant $Z_{\beta, r^*; f}(x)$, allowing for a similar reparameterization. Informally, the condition $0 \notin \operatorname{dom}(f')$ means that $D_f(\cdot \| \pi_{\text{ref}})$ acts as a *barrier* for the positive orthant, automatically forcing $\pi_{\beta; f}^*$ to place positive probability mass on any action a for which $\pi_{\text{ref}}(a | x) > 0$.

The χ^2 -divergence is an f -divergence corresponding to $f(z) = \frac{1}{2}(z-1)^2$, but unfortunately does not satisfy the condition $0 \notin \operatorname{dom}(f')$, making Eq. (10) inapplicable. Indeed, the optimal χ^2 -regularized policy can clip action probabilities to zero in a non-smooth fashion even when $\pi_{\text{ref}}(a | x) > 0$, which means that the identity Eq. (10) does not apply. To address this issue, we augment χ^2 -regularization by considering the *mixed* χ^2 -divergence given by $f_{\chi_{\text{mix}}}(z) := \frac{1}{2}(z-1)^2 + z \log z$, which has

$$D_{f_{\chi_{\text{mix}}}}(\mathbb{P} \| \mathbb{Q}) = D_{\chi^2}(\mathbb{P} \| \mathbb{Q}) + D_{\text{KL}}(\mathbb{P} \| \mathbb{Q}).$$

In other words, we use *both* χ^2 -regularization and KL-regularization; χ^2 -regularization enforces pessimism, while KL-regularization enforces the barrier property and facilitates reparameterization. Indeed, the link function ϕ (Eq. (8)) used in χ PO has $\phi(z) := f'_{\chi_{\text{mix}}}(z) = z + \log z$, which satisfies $0 \notin \operatorname{dom}(f'_{\chi_{\text{mix}}})$, so Eq. (10) yields the reparameterization $r^*(x, a) = \beta \phi \left(\frac{\pi_{\beta; f_{\chi_{\text{mix}}}}^*(a|x)}{\pi_{\text{ref}}(a|x)} \right) + Z_{\beta, r^*; f_{\chi_{\text{mix}}}}(x)$. Substituting this identity into the maximum likelihood estimation objective (Eq. (3)) yields the χ PO algorithm.

Going forward, we define $J_{\beta, r}^{\chi_{\text{mix}}}(\pi) = \mathbb{E}_{\pi} [r(x, a)] - \beta \cdot D_{\chi^2}(\pi \| \pi_{\text{ref}}) - \beta \cdot D_{\text{KL}}(\pi \| \pi_{\text{ref}})$ for a reward function r . We use the shorthand $\pi_{\beta}^* = \operatorname{argmax}_{\pi} J_{\beta, r}^{\chi_{\text{mix}}}(\pi)$ as the optimal policy under mixed χ^2 -regularization, and abbreviate $Z_{\beta, r}(x) := Z_{\beta, r; f_{\chi_{\text{mix}}}}(x)$, so that

$$r^*(x, a) = \beta \phi \left(\frac{\pi_{\beta}^*(a | x)}{\pi_{\text{ref}}(a | x)} \right) + Z_{\beta, r}(x). \quad (11)$$

3.3 THEORETICAL GUARANTEES

To state our main sample complexity guarantee for χ PO, we begin by making standard statistical assumptions. Let the regularization parameter $\beta > 0$ in χ PO be fixed. We first make a *realizability*

assumption, which states that the policy class Π used in χ PO is sufficiently expressive to represent the optimal policy under mixed χ^2 -regularization (Eq. (11)); recall that in the context of language modeling, Π represents a class of language models with fixed architecture and varying weights.

Assumption 3.1 (Policy realizability). *The policy class Π satisfies $\pi_\beta^* \in \Pi$, where π_β^* is the optimal policy under mixed χ^2 -regularization (Eq. (11)).*

Policy realizability is a standard assumption for sample-efficient reinforcement learning (Agarwal et al., 2019; Lattimore and Szepesvári, 2020; Foster and Rakhlin, 2023), and is equivalent to reward model realizability in our setting via reparameterization. Next, our second assumption asserts that the “implicit” reward models induced by the policy class Π in χ PO have bounded range.

Assumption 3.2 (Bounded implicit rewards). *For a parameter $V_{\max} \geq R_{\max}$, it holds that for all $\pi \in \Pi$, $x \in \mathcal{X}$, and $a, b \in \mathcal{A}$, $\left| \beta \phi \left(\frac{\pi(a|x)}{\pi_{\text{ref}}(a|x)} \right) - \beta \phi \left(\frac{\pi(b|x)}{\pi_{\text{ref}}(b|x)} \right) \right| \leq V_{\max}$.*

Assumption 3.2 generalizes analogous assumptions made in the analysis of DPO-like algorithms in prior work (Rosset et al., 2024; Xie et al., 2024), and our guarantees scale polynomially with this parameter; see Section 4.3 for a detailed comparison. We emphasize that in practice, V_{\max} can be measured and directly controlled (e.g., via clipping).

Example 3.1 (Policy classes induced by reward models). A natural setting in which both Assumption 3.1 and Assumption 3.2 hold is when the policy class Π is induced by a class of bounded reward function $\mathcal{R} \subset (\mathcal{X} \times \mathcal{A} \rightarrow [0, R_{\max}])$ through the mixed- χ^2 parameterization, for $\beta > 0$:

$$\Pi_{\mathcal{R}, \beta} := \left\{ \pi(a|x) = \pi_{\text{ref}}(a|x) \cdot \phi^{-1}(\beta^{-1}(r(x, a) - Z_{\beta, r}(x))) \mid r \in \mathcal{R} \right\}. \quad (12)$$

Here, Assumption 3.1 holds whenever $r^* \in \mathcal{R}$, and Assumption 3.2 holds with $V_{\max} \leq 2R_{\max}$. \triangleleft

Finally, to quantify the coverage of the offline preference dataset $\mathcal{D}_{\text{pref}}$ generated by π_{ref} , we utilize the L_1 concentrability coefficient given by $\mathcal{C}^\pi := \mathbb{E}_\pi \left[\frac{\pi(a|x)}{\pi_{\text{ref}}(a|x)} \right]$ for a policy π . \mathcal{C}^π measures the extent to which π is covered by π_{ref} . Note that L_1 -concentrability is equivalent to χ^2 -divergence up to constant shift ($\mathcal{C}^\pi = 1 + 2D_{\chi^2}(\pi \parallel \pi_{\text{ref}})$), and that in general, we can have $\mathcal{C}^\pi \ll \left\| \frac{\pi}{\pi_{\text{ref}}} \right\|_\infty$. The following is our main sample complexity guarantee for χ PO.

Theorem 3.1 (Sample complexity bound for χ PO). *Suppose Assumptions 3.1 and 3.2 hold for some $\beta > 0$. With probability at least $1 - \delta$, χ PO (Algorithm 1) produces a policy $\hat{\pi}$ such that for all policies π^* simultaneously, we have*

$$J(\pi^*) - J(\hat{\pi}) \lesssim V_{\max} e^{2R_{\max}} \cdot \sqrt{\frac{\mathcal{C}^{\pi^*} \log(|\Pi|/\delta)}{n}} + \beta \cdot \mathcal{C}^{\pi^*} + \beta^{-1} \cdot \frac{V_{\max}^2 e^{4R_{\max}} \log(|\Pi|/\delta)}{n}. \quad (13)$$

Given any comparator policy π^* , we can choose the regularization parameter β to achieve

$$J(\pi^*) - J(\hat{\pi}) \lesssim V_{\max} e^{2R_{\max}} \cdot \sqrt{\frac{\mathcal{C}^{\pi^*} \log(|\Pi|/\delta)}{n}}. \quad (14)$$

Theorem 3.1 shows that χ PO achieves a sample complexity guarantee that scales only with the single-policy concentrability parameter \mathcal{C}^{π^*} for the comparator policy π^* , for all policies π^* simultaneously. In particular, roughly $n = O\left(\frac{\mathcal{C}^{\pi^*} \log(|\Pi|/\delta)}{\varepsilon^2}\right)$ examples are sufficient to learn a policy that is ε -suboptimal relative to π^* . As a result, χ PO is robust to overoptimization since the learned policy is as good as any π^* that is sufficiently covered by π_{ref} (in the sense that $\mathcal{C}^{\pi^*} = O(1)$), which is effectively the best one can hope for in the purely offline setting. In contrast, naive offline alignment methods like DPO have sample complexity that scales with *all-policy concentrability* (roughly, $\max_\pi \mathcal{C}^\pi$), even when the comparator policy π^* is sufficiently covered (Zhu et al., 2023; Song et al., 2024). To highlight this, in Figure 1 (see Appendix B for details) we give a concrete example in which χ PO allows the user to tune β to achieve tight statistical rates, yet no choice of β for DPO leads to comparable performance (effectively, any choice of β for DPO is either susceptible to overoptimization, or has unacceptably high bias). All prior works that achieve similar sample complexity guarantees based on single-policy concentrability are either impractical, or require more restrictive statistical assumptions on the policy class (Ye et al., 2024; Liu et al., 2024; Cen et al., 2024; Fisch et al., 2024; Ji et al., 2024).³

³A notable difference is that some of these works achieve guarantees based on tighter coverage parameters than single-policy L_1 -concentrability, which reflect the structure of the policy or reward class. This appears to be out of reach for our techniques.

Regarding the parameter V_{\max} , we observe that since the policy π_{β}^* satisfies $\left| \beta \phi \left(\frac{\pi_{\beta}^*(a|x)}{\pi_{\text{ref}}(a|x)} \right) - \beta \phi \left(\frac{\pi_{\beta}^*(b|x)}{\pi_{\text{ref}}(b|x)} \right) \right| \leq 2R_{\max}$, information-theoretically we can always achieve $V_{\max} = 2R_{\max}$ by pre-filtering the policy class Π to remove all policies for which this inequality does not hold. Since this may be non-trivial computationally, we incorporate clipping in Eq. (9), which, for precisely the above reason, we expect to improve performance. In Section 4, we discuss the role of the V_{\max} parameter and Assumption 3.2 in greater depth. See also guarantees for the χ^2 -RLHF algorithm in Appendix C, which avoid dependence on this parameter.

Tuning the regularization parameter. To achieve optimal dependence on C^{π^*} , Theorem 3.1 requires tuning $\beta > 0$ as a function of this parameter, similar to other pessimistic schemes (Liu et al., 2024). With no prior knowledge, setting $\beta \propto \sqrt{\frac{V_{\max}^2 e^{4R_{\max}} \log(|\Pi|/\delta)}{n}}$ suffices to ensure that, simultaneously for all comparator policies π^* , we have $J(\pi^*) - J(\hat{\pi}) \lesssim V_{\max} e^{2R_{\max}} \cdot \sqrt{\frac{(C^{\pi^*})^2 \log(|\mathcal{R}|/\delta)}{n}}$. This guarantee achieves a slightly worse rate than Eq. (14) but holds simultaneously for all comparator policies rather than the specific one that was used to tune β . The following result, specializing to the setting in Example 3.1, shows that there exists an optimal parameter $\beta^* > 0$ that recovers the rate in Eq. (14) and holds simultaneously for all comparator policies.

Corollary 3.1 (Sample complexity bound for χ^2 P0 with a reward model). Consider the setting in Example 3.1, where the policy class $\Pi_{\mathcal{R},\beta}$ is the set of mixed χ^2 -regularized policies induced by a reward model class \mathcal{R} with $r^* \in \mathcal{R}$ and $\beta > 0$. For any $\delta \in (0, 1)$, there exists a choice⁴ for $\beta^* > 0$ such that with probability at least $1 - \delta$, χ^2 P0 (Algorithm 1), with class $\Pi_{\mathcal{R},\beta^*}$, produces a policy $\hat{\pi}$ such that for all policies π^* simultaneously, we have $J(\pi^*) - J(\hat{\pi}) \lesssim R_{\max} e^{2R_{\max}} \cdot \sqrt{\frac{C^{\pi^*} \log(|\mathcal{R}|/\delta)}{n}}$.

4 UNDERSTANDING χ^2 P0: THE BIAS-OVEROPTIMIZATION TRADEOFF

Having derived χ^2 P0 from the mixed χ^2 -regularized RLHF objective and analyzed its performance, we now take a moment to better understand the statistical properties of the policies the algorithm learns. We focus on the tradeoff between overoptimization and bias (i.e., underoptimization) achieved by the regularization parameter $\beta > 0$, highlighting through examples how this leads to statistical benefits over naive alignment methods like DPO. See Appendix B for full discussion.

4.1 PROPERTIES OF OPTIMAL POLICY UNDER MIXED χ^2 -REGULARIZATION

We begin by deriving a (nearly) closed form solution for the optimal mixed χ^2 -regularized policy in Eq. (11); recall that we expect χ^2 P0 to converge to this policy in the limit of infinite data. We first observe that the link function $\phi(\cdot)$ is strictly increasing over \mathbb{R}_+ , and its inverse is given by $\phi^{-1}(z) = W_0(\exp(z))$; here, $W_0(y)$ denotes the Lambert W-function (Corless et al., 1996), defined for $y \geq -e^{-1}$ as the inverse of the function $x \mapsto xe^x$. Consequently, for any x , the optimal policy under mixed χ^2 -regularization satisfies

$$\pi_{\beta}^*(a | x) = \pi_{\text{ref}}(a | x) \cdot W_0(\exp(\beta^{-1}(r^*(x, a) - Z_{\beta, r^*}(x)))),$$

where $Z_{\beta, r^*}(x)$ is chosen such that $\sum_a \pi_{\beta}^*(a | x) = 1$.

Compared to KL-regularization, which leads to softmax policies that satisfy $\pi_{\beta; \text{KL}}^*(a | x) = \pi_{\text{ref}}(a | x) \cdot \exp(\beta^{-1}(r^*(x, a) - Z_{\beta, r^*; \text{KL}}(x)))$, the inverse link function $\phi^{-1}(z) = W_0(\exp(z))$ for mixed χ^2 -regularization satisfies $\phi^{-1}(z) \approx z$ for $z \geq 1$, leading to a more heavy-tailed action distribution for π_{β}^* . On the other hand, for $z \leq 1$ the inverse link behaves like the exponential function (i.e., $\phi^{-1}(z) \approx e^z$ for $z \leq 1$); see Figure 2 for an illustration, and Proposition B.1 for a formal statement. Using these properties, we derive the following upper and lower bounds on the density ratio between π_{β}^* and π_{ref} .

Proposition 4.1. For all $x \in \mathcal{X}$, $a \in \mathcal{A}$, the optimal policy π_{β}^* under mixed χ^2 -regularization satisfies

$$\exp\left(-\frac{R_{\max}}{\beta}\right) \lesssim \frac{\pi_{\beta}^*(a | x)}{\pi_{\text{ref}}(a | x)} \lesssim 1 + \frac{R_{\max}}{\beta}. \quad (15)$$

Both inequalities are tight in general (up to absolute constants).

⁴It is unclear how to select β^* in a data-driven manner, as it depends on the functionals $\pi \mapsto C^{\pi}$, $\pi \mapsto J(\pi)$.




Figure 1: The regret $J(a_0) - J(\hat{\pi})$ of χ PO and DPO for different values of n . For DPO, the error from overoptimization dominates when $\beta \leq (2 \log n)^{-1}$ (as discussed in [Appendix B.3](#)), and the error from bias dominates when $\beta > (2 \log n)^{-1}$. Taking the best choice of β for each method, DPO converges at an exponentially slower rate ($\frac{1}{\log n}$) than χ PO ($\frac{1}{\sqrt{n}}$); see [Proposition A.1](#) for formal statement and [Appendix B.3](#) for further discussion.

The upper bound in [Eq. \(15\)](#), which arises from the χ^2 term in the mixed- χ^2 objective, scales inversely with the regularization parameter β , and reflects the heavy-tailed, pessimistic behavior this regularizer induces; in contrast, the optimal policy under pure KL-regularization only satisfies $\exp\left(-\frac{R_{\max}}{\beta}\right) \lesssim \frac{\pi_{\beta; \text{KL}}^*(a|x)}{\pi_{\text{ref}}(a|x)} \lesssim \exp\left(\frac{R_{\max}}{\beta}\right)$ in general. The lower bound in [Eq. \(15\)](#) arises from the KL term in the mixed- χ^2 objective, but is not important for our analysis (outside of allowing DPO-like reparameterization).

4.2 THE BIAS-OVEROPTIMIZATION TRADEOFF

We are now well equipped to understand how χ PO modulates the tradeoff between overoptimization and bias using the regularization parameter β , and how this tradeoff compares to vanilla DPO. To showcase this, we take a reward modeling perspective, and consider the setting in which the policy class Π is induced by a given reward model class \mathcal{R} , similar to [Example 3.1](#).

Suppose we start with a reward model class $\mathcal{R} \subset (\mathcal{X} \times \mathcal{A} \rightarrow [0, R_{\max}])$ such that $r^* \in \mathcal{R}$. If we use the induced policy class

$$\Pi_{\text{DPO}, \beta} := \left\{ \pi(a|x) = \pi_{\text{ref}}(a|x) \cdot \exp(\beta^{-1}(r(x, a) - Z_{\beta, r; \text{KL}}(x))) \mid r \in \mathcal{R} \right\}, \quad (16)$$

then DPO can be interpreted as fitting a reward model \hat{r} using maximum likelihood ([Eq. \(3\)](#)) and then outputting the policy $\hat{\pi}_{\text{DPO}}(a|x) = \pi_{\text{ref}}(a|x) \cdot \exp(\beta^{-1}(\hat{r}(x, a) - Z_{\beta, \hat{r}; \text{KL}}(x)))$. Meanwhile, if we use the induced policy class

$$\Pi_{\chi\text{PO}, \beta} := \left\{ \pi(a|x) = \pi_{\text{ref}}(a|x) \cdot \phi^{-1}(\beta^{-1}(r(x, a) - Z_{\beta, r}(x))) \mid r \in \mathcal{R} \right\}, \quad (17)$$

then χ PO can be interpreted as fitting a reward model \hat{r} with the exact same maximum likelihood objective, but instead outputting the policy $\hat{\pi}_{\chi\text{PO}}(a|x) = \pi_{\text{ref}}(a|x) \cdot \phi^{-1}(\beta^{-1}(\hat{r}(x, a) - Z_{\beta, \hat{r}}(x)))$.

The policies $\hat{\pi}_{\chi\text{PO}}$ and $\hat{\pi}_{\text{DPO}}$ are induced by the same reward model \hat{r} , and both use the parameter β to balance bias and overoptimization. For both policies, large β means the policy avoids overfitting to errors in the reward model (the extreme case is $\beta \rightarrow \infty$, in which case both policies become π_{ref}), while small β means the policy has low bias, i.e., low error in the case where the model is correct in the sense that $\hat{r} = r^*$ (the extreme case is $\beta \rightarrow 0$, in which case both policies become $x \mapsto \operatorname{argmax}_{a: \pi_{\text{ref}}(a|x) > 0} \hat{r}(x, a)$). Yet, for the same choice of β , $\hat{\pi}_{\chi\text{PO}}$ is significantly more heavy-tailed than $\hat{\pi}_{\text{DPO}}$, a consequence of the pessimism induced by χ^2 -regularization; see [Figure 3](#), which plots the action distribution for both policies as a function of β .

An illustrative example. Building on the intuition above, [Figure 1](#) gives a construction in which χ PO achieves $\frac{1}{\sqrt{n}}$ regret with an appropriate choice for β , yet DPO achieves an exponentially worse rate of $\frac{1}{\log n}$, no matter how β is chosen. The intuition is that for DPO, small values of β lead to severe overfitting, yet large values of β lead to high bias. Meanwhile, χ PO avoids overoptimization using comparatively small values for β , yet has bias no worse than that of DPO, thereby striking a better tradeoff. We mention that the ‘‘DPO+SFT’’ algorithm of [Liu et al. \(2024\)](#); [Cen et al. \(2024\)](#); [Fisch et al. \(2024\)](#) also fails on the construction above; see [Proposition A.1](#) in [Appendix A.1](#) for formal details, and [Appendix B.3](#) for further discussion.

4.3 NONTRIVIALITY AND ROLE OF V_{\max} PARAMETER

We close this section by discussing the role of the V_{\max} parameter ([Assumption 3.2](#)) used in the analysis of χ PO ([Theorem 3.1](#)), motivating it using the induced policy class $\Pi_{\chi\text{PO},\beta}$ from [Section 4.2](#).

[Assumption 3.2](#) effectively implies that all policies $\pi \in \Pi$ satisfy $\|\frac{\pi}{\pi_{\text{ref}}}\|_{\infty} \lesssim \frac{V_{\max}}{\beta}$; in other words, the policy class we use in χ PO satisfies *all-policy L_{∞} -concentrability* with $\max_{\pi \in \Pi} \mathcal{C}_{\infty}^{\pi} \lesssim \frac{V_{\max}}{\beta}$. This might seem to trivialize the offline alignment problem at first glance, since it would suffice to prove a generalization guarantee based on all-policy concentrability, and then plug this bound in. We will show that this is not the case, and that this is actually an intrinsic feature of χ^2 -regularization.

In more detail, recall that for χ PO, we require the realizability assumption that $\pi_{\beta}^* \in \Pi$ ([Assumption 3.1](#)), where π_{β}^* is the optimal mixed χ^2 -regularized policy that satisfies $r^*(x, a) = \beta \phi\left(\frac{\pi_{\beta}^*(a|x)}{\pi_{\text{ref}}(a|x)}\right) + Z_{\beta, r^*}(x)$. This policy, via [Proposition B.2](#), satisfies $\|\frac{\pi_{\beta}^*}{\pi_{\text{ref}}}\|_{\infty} \lesssim \frac{R_{\max}}{\beta}$, so from a statistical perspective, we can take [Assumption 3.2](#) to hold without loss of generality by removing any policy that violates this bound. In addition, as highlighted by [Example 3.1](#), if we begin from a class of bounded reward models \mathcal{R} with $r^* \in \mathcal{R}$, [Assumption 3.2](#) holds with $V_{\max} \lesssim R_{\max}$ for the induced class $\Pi_{\chi\text{PO},\beta}$ defined in [Eq. \(17\)](#), even though knowledge of such a reward model class is a mild statistical assumption that clearly does not trivialize the learning problem.

On the other hand, for DPO, a minimal assumption is that $\pi_{\beta; \text{KL}}^* \in \Pi$ ([Xie et al., 2024](#)), where $\pi_{\beta; \text{KL}}^*$ is the optimal KL-regularized policy that satisfies $r^*(x, a) = \beta \log \frac{\pi_{\beta; \text{KL}}^*(a|x)}{\pi_{\text{ref}}(a|x)} + Z_{\beta, r^*; \text{KL}}(x)$. Unlike the optimal mixed χ^2 -regularized policy, $\pi_{\beta; \text{KL}}^*$ has $\frac{\pi_{\beta; \text{KL}}^*(a|x)}{\pi_{\text{ref}}(a|x)} \gtrsim \exp\left(\frac{R_{\max}}{\beta}\right)$. This means that it is impossible to find a policy class that simultaneously (1) realizes $\pi_{\beta; \text{KL}}^*$, and (2) satisfies all-policy concentrability with $\max_{\pi \in \Pi} \mathcal{C}_{\infty}^{\pi} \ll \exp\left(\frac{R_{\max}}{\beta}\right)$. As the bias of DPO is unacceptably large unless $\beta = \text{poly}(1/n)$ (the “small- β ” regime), this leads to vacuous guarantees.

In view of these observations, our analysis of χ PO can be interpreted as (implicitly) showing that for any bounded reward class \mathcal{R} , there exists a policy class Π (precisely, the class $\Pi_{\chi\text{PO},\beta}$ defined in [Eq. \(17\)](#)) such that the following properties hold:

1. **Bounded bias.** For all $r \in \mathcal{R}$, there exists $\pi_r \in \Pi$ such that for all π^* , $J_r(\pi^*) - J_r(\pi_r) \lesssim \beta \cdot \mathcal{C}^{\pi^*}$.
2. **Bounded overoptimization.** For all $\pi \in \Pi$, $\|\frac{\pi}{\pi_{\text{ref}}}\|_{\infty} \lesssim \frac{R_{\max}}{\beta}$.

We view this as an interesting and non-trivial contribution in its own right. We mention in passing that it is possible to analyze χ PO by proving a sample complexity guarantee based on all-policy concentrability and then using $\max_{\pi \in \Pi} \mathcal{C}_{\infty}^{\pi} \lesssim \frac{V_{\max}}{\beta}$, but this would give a loose bound relative to [Theorem 3.1](#).

5 DISCUSSION

Our work gives the first general-purpose algorithm for offline alignment with provable robustness to overoptimization and sample complexity guarantees based on single-policy concentrability. Conceptually, our results contribute to a growing body of research that highlights the statistical benefits of χ^2 -divergence for reinforcement learning ([Wang et al., 2024](#); [Gabbianelli et al., 2024](#); [Amortila et al., 2024](#)), and offer an example of fruitful interplay between reinforcement learning theory and language modeling. We expect that our analysis and algorithm design techniques will find broader use.

Natural technical directions raised by our paper include (i) developing a tight understanding of mini-max sample complexity and instance-optimality for offline alignment with general policy classes; (ii) understanding the tightest possible problem-dependent sample complexity guarantees for offline alignment with *general preference models* (in light of lower bounds in [Appendix D](#)); and (iii) extending our techniques to reinforcement learning settings beyond offline alignment (e.g., general MDPs).

Additional results. *Results deferred to the appendix for space include (i) Guarantees for RLHF with χ^2 -regularization ([Appendix C](#)), (ii) Guarantees for general preference models ([Appendix D](#)), and (iii) Experiments in language models demonstrating that χ PO mitigates overoptimization ([Appendix E](#)).*

REFERENCES

- 540
541
542 A. Agarwal, N. Jiang, and S. M. Kakade. Reinforcement learning: Theory and algorithms. <https://r1theorybook.github.io/>, 2019. Version: January 31, 2022.
543
- 544 A. Agarwal, S. Kakade, A. Krishnamurthy, and W. Sun. FLAMBE: Structural complexity and
545 representation learning of low rank MDPs. *Advances in Neural Information Processing Systems*,
546 2020.
- 547 P. Amortila, D. J. Foster, and A. Krishnamurthy. Scalable online exploration via coverability.
548 *International Conference on Machine Learning*, 2024.
549
- 550 S. Athey and S. Wager. Policy learning with observational data. *Econometrica*, 2021.
551
- 552 M. G. Azar, Z. D. Guo, B. Piot, R. Munos, M. Rowland, M. Valko, and D. Calandriello. A general
553 theoretical paradigm to understand learning from human preferences. In *International Conference*
554 *on Artificial Intelligence and Statistics*, 2024.
- 555 Y. Bai, A. Jones, K. Ndousse, A. Askell, A. Chen, N. DasSarma, D. Drain, S. Fort, D. Ganguli,
556 T. Henighan, N. Joseph, S. Kadavath, J. Kernion, T. Conerly, S. El-Showk, N. Elhage, Z. Hatfield-
557 Dodds, D. Hernandez, T. Hume, S. Johnston, S. Kravec, L. Lovitt, N. Nanda, C. Olsson, D. Amodei,
558 T. Brown, J. Clark, S. McCandlish, C. Olah, B. Mann, and J. Kaplan. Training a helpful and
559 harmless assistant with reinforcement learning from human feedback. *arXiv:2204.05862*, 2022.
560
- 561 S. Biderman, H. Schoelkopf, Q. G. Anthony, H. Bradley, K. O’Brien, E. Hallahan, M. A. Khan,
562 S. Purohit, U. S. Prashanth, E. Raff, et al. Pythia: A suite for analyzing large language models
563 across training and scaling. In *International Conference on Machine Learning*, pages 2397–2430.
564 PMLR, 2023.
- 565 R. A. Bradley and M. E. Terry. Rank analysis of incomplete block designs: I. the method of paired
566 comparisons. *Biometrika*, 1952.
567
- 568 T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam,
569 G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh,
570 D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark,
571 C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. Language models are few-shot
572 learners. In *Advances in Neural Information Processing Systems*, 2020.
- 573 S. Cen, J. Mei, K. Goshvadi, H. Dai, T. Yang, S. Yang, D. Schuurmans, Y. Chi, and B. Dai.
574 Value-incentivized preference optimization: A unified approach to online and offline RLHF.
575 *arXiv:2405.19320*, 2024.
- 576 J. D. Chang, W. Shan, O. Oertell, K. Brantley, D. Misra, J. D. Lee, and W. Sun. Dataset reset policy
577 optimization for RLHF. *arXiv:2404.08495*, 2024.
578
- 579 J. Chen and N. Jiang. Offline reinforcement learning under value and density-ratio realizability: The
580 power of gaps. In *Uncertainty in Artificial Intelligence*. PMLR, 2022.
581
- 582 X. Chen, H. Zhong, Z. Yang, Z. Wang, and L. Wang. Human-in-the-loop: Provably efficient
583 preference-based reinforcement learning with general function approximation. In *International*
584 *Conference on Machine Learning*. PMLR, 2022.
- 585 Z. Chen, Y. Deng, H. Yuan, K. Ji, and Q. Gu. Self-play fine-tuning converts weak language models
586 to strong language models. *arXiv:2401.01335*, 2024.
587
- 588 V. Chernozhukov, M. Demirer, G. Lewis, and V. Syrgkanis. Semi-parametric efficient policy learning
589 with continuous actions. *Advances in Neural Information Processing Systems*, 2019.
- 590 P. F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei. Deep reinforcement learning
591 from human preferences. *Advances in Neural Information Processing Systems*, 2017.
592
- 593 R. M. Corless, G. H. Gonnet, D. E. Hare, D. J. Jeffrey, and D. E. Knuth. On the Lambert W function.
Advances in Computational Mathematics, 1996.

- 594 T. Coste, U. Anwar, R. Kirk, and D. Krueger. Reward model ensembles help mitigate overoptimization.
595 *arXiv:2310.02743*, 2023.
- 596
- 597 Q. Cui and S. S. Du. When are offline two-player zero-sum Markov games solvable? *Advances in*
598 *Neural Information Processing Systems*, 2022.
- 599 N. Das, S. Chakraborty, A. Pacchiano, and S. R. Chowdhury. Provably sample efficient RLHF via
600 active preference optimization. *arXiv:2402.10500*, 2024.
- 601
- 602 S. A. V. de Geer. *Empirical Processes in M-Estimation*. Cambridge University Press, 2000.
- 603
- 604 H. Dong, W. Xiong, D. Goyal, Y. Zhang, W. Chow, R. Pan, S. Diao, J. Zhang, K. Shum, and T. Zhang.
605 Raft: Reward ranked finetuning for generative foundation model alignment. *arXiv:2304.06767*,
606 2023.
- 607 H. Dong, W. Xiong, B. Pang, H. Wang, H. Zhao, Y. Zhou, N. Jiang, D. Sahoo, C. Xiong, and T. Zhang.
608 RLHF workflow: From reward modeling to online RLHF. *arXiv:2405.07863*, 2024.
- 609 Y. Du, A. Winnicki, G. Dalal, S. Mannor, and R. Srikant. Exploration-driven policy optimization in
610 RLHF: Theoretical insights on efficient data utilization. *arXiv:2402.10342*, 2024.
- 611
- 612 Y. Duan, Z. Jia, and M. Wang. Minimax-optimal off-policy evaluation with linear function approxi-
613 mation. In *International Conference on Machine Learning*, 2020.
- 614 J. Duchi and H. Namkoong. Variance-based regularization with convex objectives. *Journal of*
615 *Machine Learning Research*, 2019.
- 616
- 617 M. Dudík, K. Hofmann, R. E. Schapire, A. Slivkins, and M. Zoghi. Contextual dueling bandits. In
618 *Conference on Learning Theory*, 2015.
- 619 J. Eisenstein, C. Nagpal, A. Agarwal, A. Beirami, A. D’Amour, D. Dvijotham, A. Fisch, K. Heller,
620 S. Pfohl, D. Ramachandran, P. Shaw, and J. Berant. Helping or herding? reward model ensembles
621 mitigate but do not eliminate reward hacking. *arXiv:2312.09244*, 2023.
- 622
- 623 A.-m. Farahmand, C. Szepesvári, and R. Munos. Error propagation for approximate policy and value
624 iteration. *Advances in Neural Information Processing Systems*, 2010.
- 625 A. Fisch, J. Eisenstein, V. Zayats, A. Agarwal, A. Beirami, C. Nagpal, P. Shaw, and J. Berant. Robust
626 preference optimization through reward model distillation. *arXiv:2405.19316*, 2024.
- 627
- 628 P. C. Fishburn. Probabilistic social choice based on simple voting comparisons. *The Review of*
629 *Economic Studies*, 1984.
- 630 D. J. Foster and A. Rakhlin. Foundations of reinforcement learning and interactive decision making.
631 *arXiv:2312.16730*, 2023.
- 632
- 633 G. Gabbianelli, G. Neu, and M. Papini. Importance-weighted offline learning done right. In
634 *International Conference on Algorithmic Learning Theory*, 2024.
- 635 L. Gao, J. Schulman, and J. Hilton. Scaling laws for reward model overoptimization. In *International*
636 *Conference on Machine Learning*, 2023.
- 637
- 638 Z. Gao, J. D. Chang, W. Zhan, O. Oertell, G. Swamy, K. Brantley, T. Joachims, J. A. Bagnell, J. D. Lee,
639 and W. Sun. REBEL: Reinforcement learning via regressing relative rewards. *arXiv:2404.16767*,
640 2024.
- 641 Google. Palm 2 technical report. *arXiv:2305.10403*, 2023.
- 642
- 643 S. Guo, B. Zhang, T. Liu, T. Liu, M. Khalman, F. Linares, A. Rame, T. Mesnard, Y. Zhao,
644 B. Piot, J. Ferret, and M. Blondel. Direct language model alignment from online AI feedback.
645 *arXiv:2402.04792*, 2024.
- 646
- 647 S. Huang, R. F. J. Dossa, C. Ye, J. Braga, D. Chakraborty, K. Mehta, and J. G. Araújo. Cleanrl: High-
quality single-file implementations of deep reinforcement learning algorithms. *Journal of Machine*
Learning Research, 23(274):1–18, 2022. URL <http://jmlr.org/papers/v23/21-1342.html>.

- 648 X. Ji, S. Kulkarni, M. Wang, and T. Xie. Self-play with adversarial critic: Provable and scalable
649 offline alignment for language models. *arXiv:2406.04274*, 2024.
- 650
- 651 Y. Jin, Z. Yang, and Z. Wang. Is pessimism provably efficient for offline RL? In *International
652 Conference on Machine Learning*, 2021.
- 653
- 654 N. Kallus and M. Uehara. Double reinforcement learning for efficient off-policy evaluation in markov
655 decision processes. *Journal of Machine Learning Research*, 2020.
- 656
- 657 G. H. Kramer. On a class of equilibrium conditions for majority rule. *Econometrica: Journal of the
658 Econometric Society*, 1973.
- 659
- 660 G. Kreweras. Aggregation of preference orderings. In *Mathematics and Social Sciences I: Proceedings
661 of the seminars of Menthon-Saint-Bernard, France and of Gösing, Austria*, 1965.
- 662
- 663 G. Lan. Policy mirror descent for reinforcement learning: Linear convergence, new sampling
664 complexity, and generalized problem classes. *Mathematical Programming*, 2023.
- 665
- 666 T. Lattimore and C. Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- 667
- 668 J. Lee, W. Jeon, B. Lee, J. Pineau, and K.-E. Kim. Optidice: Offline policy optimization via stationary
669 distribution correction estimation. In *International Conference on Machine Learning*, 2021.
- 670
- 671 Z. Li, Z. Yang, and M. Wang. Reinforcement learning with human feedback: Learning dynamic
672 choices via pessimism. *arXiv:2305.18438*, 2023.
- 673
- 674 T. Liu, Y. Zhao, R. Joshi, M. Khalman, M. Saleh, P. J. Liu, and J. Liu. Statistical rejection sampling
675 improves preference optimization. *arXiv:2309.06657*, 2023.
- 676
- 677 Y. Liu, A. Swaminathan, A. Agarwal, and E. Brunskill. Provably good batch off-policy reinforcement
678 learning without great exploration. *Advances in Neural Information Processing Systems*, 2020.
- 679
- 680 Z. Liu, M. Lu, S. Zhang, B. Liu, H. Guo, Y. Yang, J. Blanchet, and Z. Wang. Provably
681 mitigating overoptimization in RLHF: Your SFT loss is implicitly an adversarial regularizer.
682 *arXiv:2405.16436*, 2024.
- 683
- 684 J. Y. Ma, J. Yan, D. Jayaraman, and O. Bastani. Offline goal-conditioned reinforcement learning via
685 f -advantage regression. *Advances in Neural Information Processing Systems*, 2022a.
- 686
- 687 Y. J. Ma, A. Shen, D. Jayaraman, and O. Bastani. Smodge: Versatile offline imitation learning via
688 state occupancy matching. *arXiv:2202.02433*, 2022b.
- 689
- 690 E. J. Michaud, A. Gleave, and S. Russell. Understanding learned reward functions. *arXiv:2012.05862*,
691 2020.
- 692
- 693 T. Moskovitz, A. K. Singh, D. Strouse, T. Sandholm, R. Salakhutdinov, A. D. Dragan, and S. McAleer.
694 Confronting reward model overoptimization with constrained RLHF. *arXiv:2310.04373*, 2023.
- 695
- 696 R. Munos, M. Valko, D. Calandriello, M. G. Azar, M. Rowland, Z. D. Guo, Y. Tang, M. Geist,
697 T. Mesnard, A. Michi, M. Selvi, S. Girgin, N. Momchev, O. Bachem, D. J. Mankowitz, D. Precup,
698 and B. Piot. Nash learning from human feedback. *arXiv:2312.00886*, 2023.
- 699
- 700 E. Novoseller, Y. Wei, Y. Sui, Y. Yue, and J. Burdick. Dueling posterior sampling for preference-based
701 reinforcement learning. In *Conference on Uncertainty in Artificial Intelligence*, 2020.
- OpenAI. Gpt-4 technical report. *arXiv:2303.08774*, 2023.
- L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama,
A. Ray, J. Schulman, J. Hilton, F. Kelton, L. Miller, M. Simens, A. Askell, P. Welinder, P. Christiano,
J. Leike, and R. Lowe. Training language models to follow instructions with human feedback.
Advances in Neural Information Processing Systems, 2022.
- A. Pacchiano, A. Saha, and J. Lee. Dueling RL: Reinforcement learning with trajectory preferences.
arXiv:2111.04850, 2021.

- 702 A. Pal, D. Karkhanis, S. Dooley, M. Roberts, S. Naidu, and C. White. Smaug: Fixing failure modes
703 of preference optimisation with DPO-positive. *arXiv:2402.13228*, 2024.
- 704 R. Rafailov, A. Sharma, E. Mitchell, C. D. Manning, S. Ermon, and C. Finn. Direct preference
705 optimization: Your language model is secretly a reward model. *Advances in Neural Information*
706 *Processing Systems*, 2023.
- 707 R. Rafailov, Y. Chittepudi, R. Park, H. Sikchi, J. Hejna, B. Knox, C. Finn, and S. Niekum. Scaling laws
708 for reward model overoptimization in direct alignment algorithms. *arXiv:2406.02900*, 2024a.
- 709 R. Rafailov, J. Hejna, R. Park, and C. Finn. From r to Q^* : Your language model is secretly a
710 Q -function. *arXiv:2404.12358*, 2024b.
- 711 P. Rashidinejad, B. Zhu, C. Ma, J. Jiao, and S. Russell. Bridging offline reinforcement learning and
712 imitation learning: A tale of pessimism. *Advances in Neural Information Processing Systems*,
713 2021.
- 714 M. Rita, F. Strub, R. Chaabouni, P. Michel, E. Dupoux, and O. Pietquin. Countering reward over-
715 optimization in LLM with demonstration-guided reinforcement learning. *arXiv:2404.19409*,
716 2024.
- 717 C. Rosset, C.-A. Cheng, A. Mitra, M. Santacroce, A. Awadallah, and T. Xie. Direct Nash Optimization:
718 Teaching language models to self-improve with general preferences. *arXiv:2404.03715*, 2024.
- 719 N. Shah, S. Balakrishnan, J. Bradley, A. Parekh, K. Ramchandran, and M. Wainwright. Estimation
720 from Pairwise Comparisons: Sharp Minimax Bounds with Topology Dependence. In *International*
721 *Conference on Artificial Intelligence and Statistics*, 2015.
- 722 P. B. Simpson. On defining areas of voter choice: Professor tullock on stable voting. *The Quarterly*
723 *Journal of Economics*, 1969.
- 724 Y. Song, Y. Zhou, A. Sekhari, J. A. Bagnell, A. Krishnamurthy, and W. Sun. Hybrid RL: Using both
725 offline and online data can make RL efficient. *arXiv:2210.06718*, 2022.
- 726 Y. Song, G. Swamy, A. Singh, J. A. Bagnell, and W. Sun. Understanding preference fine-tuning
727 through the lens of coverage. *arXiv:2406.01462*, 2024.
- 728 N. Stiennon, L. Ouyang, J. Wu, D. Ziegler, R. Lowe, C. Voss, A. Radford, D. Amodei, and P. F.
729 Christiano. Learning to summarize with human feedback. *Advances in Neural Information*
730 *Processing Systems*, 33:3008–3021, 2020.
- 731 G. Swamy, C. Dann, R. Kidambi, Z. S. Wu, and A. Agarwal. A minimaximalist approach to
732 reinforcement learning from human feedback. *arXiv:2401.04056*, 2024.
- 733 F. Tajwar, A. Singh, A. Sharma, R. Rafailov, J. Schneider, T. Xie, S. Ermon, C. Finn, and A. Kumar.
734 Preference fine-tuning of LLMs should leverage suboptimal, on-policy data. *arXiv:2404.14367*,
735 2024.
- 736 Y. Tang, Z. D. Guo, Z. Zheng, D. Calandriello, R. Munos, M. Rowland, P. H. Richemond, M. Valko,
737 B. Á. Pires, and B. Piot. Generalized preference optimization: A unified approach to offline
738 alignment. *arXiv:2402.05749*, 2024.
- 739 J. Tien, J. Z.-Y. He, Z. Erickson, A. Dragan, and D. S. Brown. Causal confusion and reward
740 misidentification in preference-based reward learning. In *International Conference on Learning*
741 *Representations*, 2022.
- 742 H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra,
743 P. Bhargava, S. Bhosale, D. Bikel, L. Blecher, C. C. Ferrer, M. Chen, G. Cucurull, D. Esiobu,
744 J. Fernandes, J. Fu, W. Fu, B. Fuller, C. Gao, V. Goswami, N. Goyal, A. Hartshorn, S. Hosseini,
745 R. Hou, H. Inan, M. Kardas, V. Kerkez, M. Khabsa, I. Kloumann, A. Korenev, P. S. Koura, M.-A.
746 Lachaux, T. Lavril, J. Lee, D. Liskovitch, Y. Lu, Y. Mao, X. Martinet, T. Mihaylov, P. Mishra,
747 I. Molybog, Y. Nie, A. Poulton, J. Reizenstein, R. Rungta, K. Saladi, A. Schelten, R. Silva, E. M.
748 Smith, R. Subramanian, X. E. Tan, B. Tang, R. Taylor, A. Williams, J. X. Kuan, P. Xu, Z. Yan,
749 I. Zarov, Y. Zhang, A. Fan, M. Kambadur, S. Narang, A. Rodriguez, R. Stojnic, S. Edunov, and
750 T. Scialom. Llama 2: Open foundation and fine-tuned chat models. *arXiv:2307.09288*, 2023.

- 756 A. B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer, 2008.
757
- 758 M. Uehara and W. Sun. Pessimistic model-based offline reinforcement learning under partial coverage.
759 *arXiv:2107.06226*, 2021.
- 760 T. Van Erven and P. Harremos. Rényi divergence and kullback-leibler divergence. *IEEE Transactions*
761 *on Information Theory*, 60(7):3797–3820, 2014.
- 762 L. von Werra, Y. Belkada, L. Tunstall, E. Beeching, T. Thrush, N. Lambert, and S. Huang. Trl:
763 Transformer reinforcement learning. <https://github.com/huggingface/trl>, 2020.
- 764 C. Wang, Y. Jiang, C. Yang, H. Liu, and Y. Chen. Beyond reverse KL: Generalizing direct preference
765 optimization with diverse divergence constraints. *arXiv:2309.16240*, 2023a.
- 766 L. Wang, A. Krishnamurthy, and A. Slivkins. Oracle-efficient pessimism: Offline policy optimization
767 in contextual bandits. In *International Conference on Artificial Intelligence and Statistics*, 2024.
- 768 Y. Wang, Q. Liu, and C. Jin. Is RLHF more difficult than standard RL? *arXiv:2306.14111*, 2023b.
- 769 W. H. Wong and X. Shen. Probability inequalities for likelihood ratios and convergence rates of sieve
770 mles. *The Annals of Statistics*, 1995.
- 771 R. Wu and W. Sun. Making RL with preference-based feedback efficient via randomization.
772 *arXiv:2310.14554*, 2023.
- 773 Y. Wu, Z. Sun, H. Yuan, K. Ji, Y. Yang, and Q. Gu. Self-play preference optimization for language
774 model alignment. *arXiv:2405.00675*, 2024.
- 775 T. Xie and N. Jiang. Q* approximation schemes for batch reinforcement learning: A theoretical
776 comparison. In *Conference on Uncertainty in Artificial Intelligence*, 2020.
- 777 T. Xie, C.-A. Cheng, N. Jiang, P. Mineiro, and A. Agarwal. Bellman-consistent pessimism for offline
778 reinforcement learning. *Advances in Neural Information Processing Systems*, 2021.
- 779 T. Xie, D. J. Foster, A. Krishnamurthy, C. Rosset, A. Awadallah, and A. Rakhlin. Ex-
780 ploratory preference optimization: Harnessing implicit Q*-approximation for sample-efficient rlhf.
781 *arXiv:2405.21046*, 2024.
- 782 W. Xiong, H. Dong, C. Ye, H. Zhong, N. Jiang, and T. Zhang. Gibbs sampling from human feedback:
783 A provable KL-constrained framework for RLHF. *arXiv:2312.11456*, 2023.
- 784 Y. Xu, R. Wang, L. Yang, A. Singh, and A. Dubrawski. Preference-based reinforcement learning
785 with finite-time guarantees. *Advances in Neural Information Processing Systems*, 2020.
- 786 C. Ye, W. Xiong, Y. Zhang, N. Jiang, and T. Zhang. A theoretical analysis of Nash learning from
787 human feedback under general KL-regularized preference. *arXiv:2402.07314*, 2024.
- 788 L. Yuan, G. Cui, H. Wang, N. Ding, X. Wang, J. Deng, B. Shan, H. Chen, R. Xie, Y. Lin, Z. Liu,
789 B. Zhou, H. Peng, Z. Liu, and M. Sun. Advancing llm reasoning generalists with preference trees.
790 *arXiv:2404.02078*, 2024.
- 791 A. Zanette, M. J. Wainwright, and E. Brunskill. Provable benefits of actor-critic methods for offline
792 reinforcement learning. *Advances in Neural Information Processing Systems*, 2021.
- 793 W. Zhan, B. Huang, A. Huang, N. Jiang, and J. Lee. Offline reinforcement learning with realizability
794 and single-policy concentrability. In *Conference on Learning Theory*, 2022.
- 795 W. Zhan, S. Cen, B. Huang, Y. Chen, J. D. Lee, and Y. Chi. Policy mirror descent for regularized
796 reinforcement learning: A generalized framework with linear convergence. *SIAM Journal on*
797 *Optimization*, 2023a.
- 798 W. Zhan, M. Uehara, N. Kallus, J. D. Lee, and W. Sun. Provable offline preference-based reinforce-
799 ment learning. In *International Conference on Learning Representations*, 2023b.

810 W. Zhan, M. Uehara, W. Sun, and J. D. Lee. Provable reward-agnostic preference-based reinforcement
811 learning. *arXiv:2305.18505*, 2023c.

812 T. Zhang. From ϵ -entropy to KL-entropy: Analysis of minimum information complexity density
813 estimation. *The Annals of Statistics*, 2006.

814 X. Zhang, J.-F. Ton, W. Shen, H. Wang, and Y. Liu. Overcoming reward overoptimization via
815 adversarial policy optimization with lightweight uncertainty estimation. *arXiv:2403.05171*, 2024.

816 B. Zhu, M. Jordan, and J. Jiao. Principled reinforcement learning with human feedback from pairwise
817 or k-wise comparisons. In *International Conference on Machine Learning*, 2023.

818 B. Zhu, M. I. Jordan, and J. Jiao. Iterative data smoothing: Mitigating reward overfitting and
819 overoptimization in RLHF. *arXiv:2401.16335*, 2024.

820 H. Zhu and A. Zhang. Provably efficient offline goal-conditioned reinforcement learning with
821 general function approximation and single-policy concentrability. *Advances in Neural Information
822 Processing Systems*, 2024.

823 Z. Zhu, K. Lin, B. Dai, and J. Zhou. Off-policy imitation learning from observations. *Advances in
824 Neural Information Processing Systems*, 2020.

825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

Contents of Appendix

| | | |
|-----------|--|-----------|
| I | Additional Results | 17 |
| A | Additional Related Work | 17 |
| A.1 | Detailed Comparison to DPO + SFT | 19 |
| B | Detailed Discussion: χP0 and the Bias-Overoptimization Tradeoff | 21 |
| B.1 | Properties of Optimal Policy under Mixed χ^2 -Regularization | 21 |
| B.2 | The Bias-Overoptimization Tradeoff | 22 |
| B.3 | An Illustrative Example | 23 |
| B.4 | Nontriviality and Role of V_{\max} Parameter | 25 |
| C | Sample Complexity Guarantees for χ^2-RLHF | 25 |
| D | χP0 for General Preference Models | 27 |
| D.1 | Impossibility of Single-Policy Concentrability under General Preferences | 28 |
| D.2 | Iterative χ P0 for General Preferences | 29 |
| D.3 | Theoretical Analysis of Iterative χ P0 | 30 |
| E | Experiments in Offline Language Model Alignment | 32 |
| E.1 | TL;DR Summarization | 32 |
| E.2 | Experiment details | 33 |
| II | Proofs | 34 |
| F | Preliminaries | 34 |
| G | Analysis of χP0: Proof Sketch for Theorem 3.1 | 34 |
| H | Proofs for Section 3 | 36 |
| H.1 | General Version of Theorem 3.1 | 36 |
| H.2 | Proof of Theorem 3.1 | 44 |
| H.3 | Proof of Corollary 3.1 | 44 |
| I | Proofs for Appendix B | 45 |
| J | Proofs for Appendix D | 46 |
| J.1 | Proof of Theorem D.1 | 46 |
| J.2 | Proof of Theorem D.2 | 47 |
| J.3 | Proofs for Supporting Lemmas | 50 |
| K | Proofs for Appendix C | 52 |

Part I

Additional Results

A ADDITIONAL RELATED WORK

Theoretical algorithms for offline alignment. Much of prior theoretical work on offline alignment considers algorithms that are tailored to linearly parameterized policies (Zhu et al., 2023; Zhan et al., 2023b; Li et al., 2023; Xiong et al., 2023). For general policy classes, Ye et al. (2024) provide an algorithm that achieves sample complexity guarantees based on single-policy concentrability, but the algorithm requires computation of an uncertainty bonus which cannot be implemented faithfully for

918 large language models. Ji et al. (2024) provide an algorithm that achieves single-policy concentrability
 919 using self-play, but their approach requires the non-standard realizability assumption that for all
 920 $\pi \in \Pi$, there exists $\pi' \in \Pi$ such that $r(x, a) = \beta \log \frac{\pi(a|x)}{\pi'(a|x)} - Z_{\pi, \pi'}(x)$ for some function $Z_{\pi, \pi'}(x)$
 921 that depends on x , but not the action a . In addition, their algorithm is iterative, and requires solving a
 922 DPO-like objective many times (roughly $1/\varepsilon^2$ iterations are required to achieve accuracy ε). Most
 923 relevant to our work, Liu et al. (2024); Cen et al. (2024); Fisch et al. (2024) propose solving the
 924 appealingly simple DPO + SFT objective in Eq. (5). As we discuss in detail in Appendix A.1, this
 925 objective fails to achieve single-policy concentrability unless non-standard convexity assumptions on
 926 the policy class or reward model class hold.

927 A number of other works consider the *hybrid* setting for alignment where—in addition to offline
 928 preference data from π_{ref} , the algorithm has access to online feedback (Xiong et al., 2023; Gao
 929 et al., 2024; Chang et al., 2024; Song et al., 2024). While it is straightforward to achieve guarantees
 930 based on single-policy concentrability in this setting, this is a stronger feedback model than what
 931 we consider, and is not always realistic. Our work is also complementary to fully online alignment,
 932 which dispenses with coverage conditions entirely but requires active exploration (Xu et al., 2020;
 933 Novoseller et al., 2020; Pacchiano et al., 2021; Wu and Sun, 2023; Zhan et al., 2023c; Chen et al.,
 934 2022; Wang et al., 2023b; Du et al., 2024; Das et al., 2024; Ye et al., 2024; Xie et al., 2024; Cen et al.,
 935 2024).

936 **Generalizations of DPO.** Wang et al. (2023a) provide a generalization of the DPO reparameterization
 937 trick which supports general f -divergences that satisfy certain regularity conditions. Their work does
 938 not provide sample complexity guarantees or theoretical guidance on which choices of f -divergence
 939 are preferable, but our main algorithm χ P0, can be derived as a special case of their technique with a
 940 novel choice of f -divergence. Tang et al. (2024) also provide a general framework for deriving DPO
 941 variants with general loss functions, but our algorithm does not appear to be a special case of their
 942 framework.

943 **Offline reinforcement learning theory.** The theory of *offline reinforcement learning* addresses
 944 challenges similar to overoptimization, which is typically describes through the language of *distrib-*
 945 *ution shift*. Many of these works, using pessimism and related algorithmic techniques, provide
 946 guarantees that are robust to partial coverage of the data collection policy π_{ref} , which is reflected in
 947 sample complexity guarantees based on single-policy concentrability and similar coverage conditions.
 948 While this line of work provides efficient algorithms for simple (e.g., tabular or linear) settings (Liu
 949 et al., 2020; Jin et al., 2021; Rashidinejad et al., 2021), existing approaches that support general
 950 function approximation (Xie et al., 2021; Uehara and Sun, 2021; Zhan et al., 2022; Chen and Jiang,
 951 2022) cannot be implemented efficiently for language models without non-trivial modifications. See
 952 also closely related research on policy optimization and evaluation in statistics and econometrics
 953 (Athey and Wager, 2021; Chernozhukov et al., 2019; Kallus and Uehara, 2020).

954 **χ^2 -divergence in reinforcement learning.** Our work contributes to a growing body of research
 955 that uses χ^2 -divergence to derive reinforcement learning algorithms with novel statistical guarantees.⁵
 956 Notably, our work is inspired by Wang et al. (2024) (see also Gabbianelli et al. (2024)), who use a
 957 regularizer similar to χ^2 -divergence to derive single-policy concentrability guarantees for contextual
 958 bandits. Compared to the χ^2 -regularizer $\mathcal{C}^\pi = \mathbb{E}_\pi \left[\frac{\pi(a|x)}{\pi_{\text{ref}}(a|x)} \right]$ we use, their regularizer takes the form
 959 $\mathbb{E}_\pi \left[\frac{1}{\pi_{\text{ref}}(a|x)} \right]$, which is always larger. As a result of this difference, their regularizer is not suitable for
 960 large action spaces. By addressing this shortcoming, we expect our χ^2 -regularization approach to
 961 find further use in offline RL.
 962

963 Other related works include (i) Duan et al. (2020) show that χ^2 -divergence plays a fundamental
 964 role in offline RL with linear function approximation; (ii) Zhan et al. (2022) use χ^2 -regularization
 965 to provide guarantees based on single-policy concentrability for an offline RL method based on
 966 weight function learning; and (iii) Amortila et al. (2024) provide online RL algorithms that explore
 967 by directly minimizing an exploration objective based on χ^2 -divergence. We mention in passing that
 968 a number of recent empirical works apply χ^2 -regularization (Zhu et al., 2020; Lee et al., 2021; Ma
 969 et al., 2022a;b; Zhu and Zhang, 2024) to reinforcement learning in embodied domains.
 970

971 ⁵More classically, χ^2 -divergence is known to play a fundamental role in asymptotic statistics (Tsybakov,
 2008; Duchi and Namkoong, 2019).

Empirical research on offline alignment. Our work uses DPO (Rafailov et al., 2023) as a starting point. Many prior works have built upon DPO with the aim of addressing specific shortcomings, including Liu et al. (2023); Tang et al. (2024); Azar et al. (2024); Rosset et al. (2024); Chen et al. (2024); Wu et al. (2024); Tajwar et al. (2024). Closely related, there is a large body of research that attempts to understand and mitigate overoptimization in offline alignment from a purely empirical perspective (Michaud et al., 2020; Tien et al., 2022; Coste et al., 2023; Dong et al., 2023; Eisenstein et al., 2023; Gao et al., 2023; Moskovitz et al., 2023; Pal et al., 2024; Rita et al., 2024; Rafailov et al., 2024a; Zhang et al., 2024).

A.1 DETAILED COMPARISON TO DPO + SFT

In this section, we give additional background on the suboptimality of the DPO + SFT objective in Eq. (5). Let $\beta > 0$ be the KL-regularization parameter and $\alpha > 0$ be an optimism parameter. Consider the setting in which $\Pi = \{\pi_r(a | x) = \pi_{\text{ref}}(a | x) \exp(\beta^{-1}(r(x, a) - Z_r(x))) \mid r \in \mathcal{R}\}$ for a given reward function class $\mathcal{R} \subset (\mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R})$. Liu et al. (2024); Cen et al. (2024); Fisch et al. (2024) propose solving (variants of) the objective

$$\hat{\pi}_{\max\text{-min}} = \operatorname{argmax}_{\pi} \min_{r \in \mathcal{R}} \left\{ \alpha \left(\mathbb{E}_{x \sim \rho, a \sim \pi(\cdot | x), b \sim \pi_{\text{ref}}(\cdot | x)} [r(a) - r(b)] - \beta D_{\text{KL}}(\pi \parallel \pi_{\text{ref}}) \right) + \mathcal{L}(r) \right\}, \quad (18)$$

where the max ranges over the space of all policies, and where $\mathcal{L}(r) := -\frac{1}{n} \sum_{(x, a_+, a_-) \in \mathcal{D}_{\text{pref}}} \log \sigma[r(x, a_+) - r(x, a_-)]$ is the negative log-likelihood under the Bradley-Terry model. Liu et al. (2024) show that for general policy classes, this algorithm attains sample complexity guarantees scaling with single-policy concentrability; Cen et al. (2024) provide similar results for the special case of linearly parameterized policies.

The objective in Eq. (18) is non-trivial to implement for language models. To derive the DPO + SFT objective in Eq. (5), Liu et al. (2024) observe that if \mathcal{R} is convex, the minimax theorem implies that the objective value in Eq. (18) is equivalent to the value for the min-max objective

$$\min_{r \in \mathcal{R}} \max_{\pi} \left\{ \alpha \left(\mathbb{E}_{x \sim \rho, a \sim \pi(\cdot | x), b \sim \pi_{\text{ref}}(\cdot | x)} [r(a) - r(b)] - \beta D_{\text{KL}}(\pi \parallel \pi_{\text{ref}}) \right) + \mathcal{L}(r) \right\}. \quad (19)$$

This leads to a natural algorithmic strategy adopted by (Liu et al., 2024; Cen et al., 2024; Fisch et al., 2024): Let $\hat{r}_{\min\text{-max}}$ be the minimizing reward function in Eq. (19) and let $\pi_{\hat{r}_{\min\text{-max}}}$ —the optimal policy in the KL-regularized MDP with reward function $\hat{r}_{\min\text{-max}}$ —be the final policy returned by the algorithm. After standard manipulations, one can then show that $\pi_{\hat{r}_{\min\text{-max}}}$ is equivalent to

$$\operatorname{argmax}_{\pi \in \Pi} \left\{ \alpha \cdot \mathbb{E}_{\pi_{\text{ref}}} [\beta \log \pi(a | x)] + \frac{1}{n} \sum_{(x, a_+, a_-) \in \mathcal{D}_{\text{pref}}} \log \left[\sigma \left(\beta \log \frac{\pi(a_+ | x)}{\pi_{\text{ref}}(a_+ | x)} - \beta \log \frac{\pi(a_- | x)}{\pi_{\text{ref}}(a_- | x)} \right) \right] \right\}. \quad (20)$$

We call this policy $\hat{\pi}_{\text{DPO+SFT}}$. The sample complexity analyses for the $\hat{\pi}_{\text{DPO+SFT}}$ policy (Eq. (20)) in (Liu et al., 2024; Cen et al., 2024) rely on showing that the objective value in Eq. (19) is equivalent to the value in Eq. (18), which is not guaranteed to hold if \mathcal{R} is non-convex (e.g., if \mathcal{R} is a class of neural networks).⁶ Indeed, the following proposition shows that, for non-convex reward classes \mathcal{R} , the DPO + SFT objective in Eq. (20) fails to achieve a statistical guarantee based on single-policy concentrability, even when Eq. (18) succeeds.

Proposition A.1. *Let $n \in \mathbb{N}$ with $n \geq 2$ be given. There exists a reward class \mathcal{R} with $|\mathcal{R}| = 2$, a problem instance (ρ, r) satisfying realizability ($r \in \mathcal{R}$) and $r \in [0, 1]$, a data collection policy π_{ref} , and universal constants $c_1 \in (0, 1)$ and $c_2, c_3 > 0$ such that the following hold:*

1. *There exists a policy $\tilde{\pi}$ such that $\|\tilde{\pi}/\pi_{\text{ref}}\|_{\infty} \leq 2$; yet*
2. *For any $\beta \leq (2 \log(n))^{-1}$ and $\alpha \geq 0$, the minimax policy $\hat{\pi}_{\min\text{-max}}$ (Eq. (19)) and DPO+SFT policy $\hat{\pi}_{\text{DPO+SFT}}$ (Eq. (20)) derived from a dataset $\mathcal{D}_{\text{pref}}$ of n samples from π_{ref} incur suboptimality*

$$J(\tilde{\pi}) - J(\hat{\pi}_{\text{DPO+SFT}}) = J(\tilde{\pi}) - J(\hat{\pi}_{\min\text{-max}}) \geq c_2,$$

with probability at least c_1 .

⁶Precisely, Liu et al. (2024) provide guarantees for $\hat{\pi}_{\max\text{-min}}$ with general reward class \mathcal{R} and establish equivalence of $\hat{\pi}_{\max\text{-min}}$ and $\hat{\pi}_{\min\text{-max}}$ when \mathcal{R} is convex, while Cen et al. (2024) consider linear function approximation, which yields the required convexity.

3. For any $\beta \geq (2 \log(n))^{-1}$ and $\alpha \geq 0$, the minimax policy $\hat{\pi}_{\min\text{-max}}$ (Eq. (19)) and DPO+SFT policy $\hat{\pi}_{\text{DPO+SFT}}$ (Eq. (20)) derived from a dataset $\mathcal{D}_{\text{pref}}$ of n samples from π_{ref} incur suboptimality

$$J(\hat{\pi}) - J(\hat{\pi}_{\text{DPO+SFT}}) = J(\hat{\pi}) - J(\hat{\pi}_{\min\text{-max}}) \geq \frac{c_3}{\log(n)},$$

with probability at least c_1 .

On the other hand, we observe that for the instance in Proposition A.1, χP0 (via Theorem 3.1) with $\beta \propto 1/\sqrt{n}$ and the class $\Pi = \{\pi(a | x) = \pi_{\text{ref}}(a | x) \cdot \phi^{-1}(\beta^{-1}(r(x, a) - Z_r(x))) \mid r \in \mathcal{R}\}$ achieves

$$J(\hat{\pi}) - J(\hat{\pi}_{\text{DPO+SFT}}) \lesssim \sqrt{\frac{(\mathcal{C}_{\hat{\pi}})^2}{n}} \lesssim \sqrt{\frac{1}{n}},$$

highlighting the fact that χP0 meaningfully adapts to single-policy concentrability even when the technical conditions required by DPO+SFT do not hold; see also Appendix B. We find this conclusion to be somewhat surprising, as Xie et al. (2024) show that an optimistic counterpart to Eq. (20), which negates the SFT term, enjoys strong guarantees for online alignment with general policy classes without requiring convexity conditions.

Although our construction does not establish inconsistency in the $\beta \geq (2 \log(n))^{-1}$ regime, in general, DPO+SFT will incur $O(\beta)$ bias if one aims to compete with the optimal policy. Due to restriction that β must be rather large, this results in an exponentially slower rate of convergence than χP0 .

Proof of Proposition A.1. Let $n \in \mathbb{N}$ with $n \geq 2$ be given. We consider a problem instance with $\mathcal{X} = \{x_1, x_2\}$ and $\mathcal{A} = \{a_0, a_1, a_2, a_3\}$, so that $|\mathcal{A}| = 4$. We define a reward class with two reward functions $\mathcal{R} := \{r_1, r_2\}$ as follows. For $i \in \{1, 2\}$:

$$\begin{aligned} r_i(x_1, a_0) &= \zeta, & \text{and} & & r_i(x_1, a_1) &= r_i(x_1, a_2) &= r_i(x_1, a_3) &= 0 \\ r_i(x_2, a_0) &= 1/2, & r_i(x_2, a_i) &= 1, & \text{and} & r_i(x_2, a_j) &= 0 & \forall j \neq i. \end{aligned}$$

Here $\zeta \in [0, 1]$ will be chosen at the end of the proof. The context distribution is $\rho = \text{unif}(\mathcal{X})$, and we define π_{ref} for each $x_i \in \{x_1, x_2\}$ via

$$\pi_{\text{ref}}(a_0 | x_i) = 1/2, \quad \pi_{\text{ref}}(a_1 | x_i) = \pi_{\text{ref}}(a_2 | x_i) = 1/(2n), \quad \text{and} \quad \pi_{\text{ref}}(a_3 | x_i) = (n-2)/(2n).$$

Let r_1 be the true reward function. Recall that $\mathcal{D}_{\text{pref}} = \{(x, a_+, a_-)\}$ consists of n tuples (x, a_+, a_-) obtained by sampling $x \sim \rho$ and a pair of actions $(a, b) \sim \pi_{\text{ref}}$ and labeling them as (a_+, a_-) via the Bradley-Terry model in Eq. (1) with reward r_1 . Define a ‘‘bad’’ event under this process:

$$\mathcal{E} := \{\text{No tuples in } \mathcal{D}_{\text{pref}} \text{ contain } a_1 \text{ or } a_2\}.$$

We can lower bound the probability of \mathcal{E} as follows:

$$\begin{aligned} \mathbb{P}[\mathcal{E}^c] &\leq \mathbb{P}[a_1 \text{ in } \mathcal{D}_{\text{pref}}] + \mathbb{P}[a_2 \text{ in } \mathcal{D}_{\text{pref}}] \\ &= 2(1 - (1 - 1/2n)^n) \leq 2(1 - e^{-1/2}(1 - 1/(4n))) \leq 2(1 - 7e^{-1/2}/8) \leq 0.94, \end{aligned}$$

where the first inequality uses that $(1 - x/n)^n \geq e^{-x}(1 - x^2/n)$ for $n \geq 1$ and $|x| < n$. We conclude that

$$\mathbb{P}[\mathcal{E}] \geq 0.06 =: c_1.$$

Let $\mathcal{L}(r; \mathcal{D}_{\text{pref}}) := -\frac{1}{n} \sum_{(x, a_+, a_-) \in \mathcal{D}_{\text{pref}}} \log \sigma[r(x, a_+) - r(x, a_-)]$ denote the DPO loss. Observe that conditioned on \mathcal{E} , we have that $\mathcal{L}(r_1; \mathcal{D}_{\text{pref}}) = \mathcal{L}(r_2; \mathcal{D}_{\text{pref}})$. Noting that

$$\max_{\pi} \{\mathbb{E}_{\pi} [r] - \mathbb{E}_{\pi_{\text{ref}}} [r] - \beta D_{\text{KL}}(\pi \parallel \pi_{\text{ref}})\} = \mathbb{E}_{\pi_r} [r] - \mathbb{E}_{\pi_{\text{ref}}} [r] - \beta D_{\text{KL}}(\pi_r \parallel \pi_{\text{ref}}),$$

is the same for both $r \in \mathcal{R}$, we see that both r_1 and r_2 optimize the minimax objective in Eq. (19). Thus, breaking ties adversarially, we can choose $\hat{\pi}_{\min\text{-max}} = \pi_{r_2}$ under \mathcal{E} for all values of $\beta > 0$ and $\alpha \geq 0$. By the equivalence between the minimax objective in Eq. (19) and the DPO+SFT objective in Eq. (20) (Liu et al., 2024; Cen et al., 2024; Fisch et al., 2024), for $\Pi = \{\pi_{r_1}, \pi_{r_2}\}$, we can choose $\hat{\pi}_{\text{DPO+SFT}} = \pi_{r_2}$ in Eq. (20) under \mathcal{E} . Indeed, under \mathcal{E} , the DPO+SFT objective is equivalent to $\arg\max_{\pi \in \Pi} \mathbb{E}_{\pi_{\text{ref}}} [\log \pi(a)]$, and π_{r_1} and π_{r_2} have the same value for this objective.

To conclude we choose $\tilde{\pi}(\cdot) = a_0$, which has $\|\tilde{\pi}/\pi_{\text{ref}}\|_{\infty} = 2$. It remains to calculate the suboptimality gap.

$$J(\tilde{\pi}) - J(\hat{\pi}_{\text{DPO+SFT}}) = J(\tilde{\pi}) - J(\hat{\pi}_{\text{min-max}}) = J(\tilde{\pi}) - J(\pi_{r_2})$$

under \mathcal{E} . Note that $J(\tilde{\pi}) = \zeta/2 + 1/4$. We decompose the reward for π_{r_2} on instance r_1 into two components, corresponding to the two contexts x_1, x_2 :

$$\begin{aligned} J(\pi_{r_2}) &= \frac{1}{2}(\mathbb{E}_{a \sim \pi_{r_2}}[r_1(x_1, a)] + \mathbb{E}_{a \sim \pi_{r_2}}[r_1(x_2, a)]) =: \frac{1}{2}(J_1(\beta) + J_2(\beta)) \\ J_1(\beta) &= \frac{r_1(x_1, a_0)\pi_{\text{ref}}(a_0 | x_1) \exp(r_2(x_1, a_0)/\beta)}{Z(r_2, x_1)} = \frac{\zeta/2 \exp(\zeta/\beta)}{1/2 \exp(\zeta/\beta) + 1/2} \\ J_2(\beta) &= \frac{r_1(x_2, a_0)\pi_{\text{ref}}(a_0 | x_2) \exp(r_2(x_2, a_0)/\beta) + r_1(x_1, a_1)\pi_{\text{ref}}(a_1 | x_2) \exp(r_2(x_2, a_1)/\beta)}{Z(r_2, x_2)} \\ &= \frac{1/4e^{1/2\beta} + 1/(2n)}{1/2e^{1/2\beta} + e^{1/\beta}/(2n) + (n-1)/(2n)}, \end{aligned}$$

where $Z(r_2, x) := \sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a | x) \exp(r_2(x, a)/\beta)$.

We first consider the small β regime. Here we use the upper bound $J_1(\beta) \leq \zeta$ and focus on $J_2(\beta)$. Note that $J_2(\beta)$ is increasing with β for $\beta \leq 1/(2 \log(n))$. In particular, if we consider $\beta = 1/(c \log(n))$ for $c \geq 2$, then the expression above is equal to

$$J_2(\beta) = \frac{n^{c/2}/4 + 1/(2n)}{n^{c/2}/2 + n^{c-1}/2 + (n-1)/(2n)} \leq \frac{n^{c/2}/4 + 1/(2n)}{n^{c/2} + (n-1)/(2n)} \leq 1/4 + \frac{1}{2n^{c/2+1}} \leq 3/8,$$

where the last inequality holds when $c \geq 2$ and $n \geq 2$. We set $c = 2$, so that as long as $n \geq 2$, $J(\pi_{r_2}) \leq \frac{3}{8}$. Thus, the suboptimality is

$$J(\tilde{\pi}) - J(\pi_{r_2}) \geq \frac{\zeta}{2} + \frac{1}{4} - \left(\frac{\zeta}{2} + \frac{3}{16} \right) \geq \frac{1}{16} =: c_2.$$

Next consider the regime where $\beta \geq 1/(2 \log(n))$. Analogously to before, note that $J_2(\beta) \leq 1/2$. On the other hand, $J_1(\beta)$ is monotonically decreasing with β , so using $\beta \geq 1/(2 \log(n))$ we obtain the bound

$$J_1(\beta) \leq \frac{\zeta \exp(2\zeta \log(n))}{\exp(2\zeta \log(n)) + 1} = \zeta \cdot \frac{n^{2\zeta}}{n^{2\zeta} + 1}.$$

So in this case, the suboptimality is

$$J(\tilde{\pi}) - J(\pi_{r_2}) \geq \frac{\zeta}{2} \cdot \left(1 - \frac{n^{2\zeta}}{n^{2\zeta} + 1} \right) \geq \frac{\zeta}{4} \cdot \frac{1}{n^{2\zeta}} = \frac{\log(2)}{16 \log(n)},$$

if we set $\zeta = \log(2)/(2 \log(n))$ which is in $[0, 1]$ under the assumption that $n \geq 2$. \square

B DETAILED DISCUSSION: χ^0 PO AND THE BIAS-OVEROPTIMIZATION TRADEOFF

Having derived χ^0 PO from the mixed χ^2 -regularized RLHF objective and analyzed its performance, we now take a moment to better understand the statistical properties of the policies the algorithm learns. We focus on the tradeoff between overoptimization and bias (i.e., underoptimization) achieved by the regularization parameter $\beta > 0$, highlighting through examples how this leads to statistical benefits over naive alignment methods like DPO.

B.1 PROPERTIES OF OPTIMAL POLICY UNDER MIXED χ^2 -REGULARIZATION

We begin by deriving a (nearly) closed form solution for the optimal mixed χ^2 -regularized policy in Eq. (11); recall that we expect χ^0 PO to converge to this policy in the limit of infinite data.

We first observe that the link function $\phi(\cdot)$ is strictly increasing over \mathbb{R}_+ , and its inverse is given by $\phi^{-1}(z) = W_0(\exp(z))$; here, $W_0(y)$ denotes the Lambert W-function (Corless et al., 1996), defined




Figure 2: Behavior of the mixed χ^2 -regularization link function $\phi_{\chi^2\text{PO}}(z) = z + \log z$ and inverse $\phi_{\chi^2\text{PO}}^{-1}(z) = W_0(\exp(z))$, compared to the KL-regularization link function $\phi_{\text{DPO}}(z) = \log z$ and inverse $\phi_{\text{DPO}}^{-1}(z) = \exp(z)$. $\phi_{\chi^2\text{PO}}^{-1}(z) \approx z$ for $z \geq 1$, leading to favorable heavy-tailed, pessimistic behavior.

for $y \geq -e^{-1}$ as the inverse of the function $x \mapsto xe^x$. Consequently, for any x , the optimal policy under mixed χ^2 -regularization satisfies

$$\pi_{\beta}^*(a | x) = \pi_{\text{ref}}(a | x) \cdot W_0(\exp(\beta^{-1}(r^*(x, a) - Z_{\beta, r^*}(x)))),$$

where $Z_{\beta, r^*}(x)$ is chosen such that $\sum_a \pi_{\beta}^*(a | x) = 1$. We can better understand how this policy behaves using the following simple upper and lower bounds on the inverse link function $\phi^{-1}(z) = W_0(\exp(z))$.

Proposition B.1. *The link function $\phi(z) = z + \log z$ is strictly increasing over $(0, \infty)$, and its inverse $\phi^{-1}(z) = W_0(\exp(z))$ is strictly increasing over $(-\infty, \infty)$. The inverse link function ϕ^{-1} satisfies*

$$\frac{z}{2} \leq \phi^{-1}(z) \leq z \quad \forall z \in [1, \infty), \quad \text{and} \quad e^{z-e} \leq \phi^{-1}(z) \leq e^z \quad \forall z \in (-\infty, 1].$$

Compared to KL-regularization, which leads to softmax policies that satisfy $\pi_{\beta; \text{KL}}^*(a | x) = \pi_{\text{ref}}(a | x) \cdot \exp(\beta^{-1}(r^*(x, a) - Z_{\beta, r^*; \text{KL}}(x)))$, we see that the inverse link function $\phi^{-1}(z) = W_0(\exp(z))$ for mixed χ^2 -regularization satisfies $\phi^{-1}(z) \approx z$ for $z \geq 1$, leading to a more heavy-tailed action distribution for π_{β}^* . On the other hand, for $z \leq 1$ the inverse link behaves like the exponential function (i.e., $\phi^{-1}(z) \approx e^z$ for $z \leq 1$); see Figure 2 for an illustration. Using these properties, we can derive the following upper and lower bounds on the density ratio between π_{β}^* and π_{ref} .

Proposition B.2 (Proposition 4.1 restated). *For all $x \in \mathcal{X}$ and $a \in \mathcal{A}$, the optimal policy π_{β}^* under mixed χ^2 -regularization satisfies*

$$\exp\left(-\frac{R_{\max}}{\beta}\right) \lesssim \frac{\pi_{\beta}^*(a | x)}{\pi_{\text{ref}}(a | x)} \lesssim 1 + \frac{R_{\max}}{\beta}. \quad (21)$$

Both inequalities are tight in general (up to absolute constants).

The upper bound in Eq. (21), which arises from the χ^2 term in the mixed- χ^2 objective, scales inversely with the regularization parameter β , and reflects the heavy-tailed, pessimistic behavior this regularizer induces; in contrast, the optimal policy under pure KL-regularization only satisfies

$$\exp\left(-\frac{R_{\max}}{\beta}\right) \lesssim \frac{\pi_{\beta; \text{KL}}^*(a | x)}{\pi_{\text{ref}}(a | x)} \lesssim \exp\left(\frac{R_{\max}}{\beta}\right) \quad (22)$$

in general. The lower bound in Eq. (21) arises from the KL term in the mixed- χ^2 objective, but is not important for our analysis (outside of allowing for DPO-like reparameterization).

B.2 THE BIAS-OVEROPTIMIZATION TRADEOFF

We are now well equipped to understand how $\chi^2\text{PO}$ modulates the tradeoff between overoptimization and bias using the regularization parameter β , and how this tradeoff compares to vanilla DPO. To showcase this, we take a reward modeling perspective, and consider the setting in which the policy class Π is induced by a given reward model class \mathcal{R} , similar to Example 3.1.




Figure 3: Action probabilities for policies learned by χ PO and DPO on the example from Appendix B.3, under the “bad” event \mathcal{E} in which the true reward model is $r^* = r_1$ but the estimated reward model is $\hat{r} = r_2$ ($n = 10$). Here, $r^*(a_{\text{good}}) = 1$ and $r^*(a_{\text{bad}}) = 0$, but $\hat{r}(a_{\text{good}}) = 0$ and $\hat{r}(a_{\text{bad}}) = 1$; both reward functions have $r^*(a_0) = \hat{r}(a_0) = 1/2$, and the goal is to compete with a comparator policy that deterministically plays a_0 .

Overoptimization. The DPO policy is greedier with respect to the incorrect reward model and places much larger mass on the bad action a_{bad} for all $\beta \in (0, \frac{1}{2 \log n}]$ (Right). As a result, the DPO policy places much smaller mass on the baseline action a_0 , suffering significantly more overoptimization error compared to χ PO (Left; see also Figure 1).

Bias. Compared to DPO, χ PO has a higher probability of taking the optimal action a_{good} , and thus has lower bias (Center). As a result, it strikes a better bias-overoptimization tradeoff than DPO, and is competitive with respect to the comparator a_0 even when DPO fails to converge.

Suppose we start with a reward model class $\mathcal{R} \subset (\mathcal{X} \times \mathcal{A} \rightarrow [0, R_{\max}])$ such that $r^* \in \mathcal{R}$. If we use the induced policy class

$$\Pi_{\text{DPO}, \beta} := \{ \pi(a | x) = \pi_{\text{ref}}(a | x) \cdot \exp(\beta^{-1}(r(x, a) - Z_{\beta, r; \text{KL}}(x))) \mid r \in \mathcal{R} \}, \quad (23)$$

then DPO can be interpreted as fitting a reward model \hat{r} using maximum likelihood (Eq. (3)) and then outputting the policy $\hat{\pi}_{\text{DPO}}(a | x) = \pi_{\text{ref}}(a | x) \cdot \exp(\beta^{-1}(\hat{r}(x, a) - Z_{\beta, \hat{r}; \text{KL}}(x)))$. Meanwhile, if we use the induced policy class

$$\Pi_{\chi\text{PO}, \beta} := \{ \pi(a | x) = \pi_{\text{ref}}(a | x) \cdot \phi^{-1}(\beta^{-1}(r(x, a) - Z_{\beta, r}(x))) \mid r \in \mathcal{R} \}, \quad (24)$$

then χ PO can be interpreted as fitting a reward model \hat{r} with the exact same maximum likelihood objective, but instead outputting the policy $\hat{\pi}_{\chi\text{PO}}(a | x) = \pi_{\text{ref}}(a | x) \cdot \phi^{-1}(\beta^{-1}(\hat{r}(x, a) - Z_{\beta, \hat{r}}(x)))$.

The policies $\hat{\pi}_{\chi\text{PO}}$ and $\hat{\pi}_{\text{DPO}}$ are induced by the same reward model \hat{r} , and both use the parameter β to balance bias and overoptimization. For both policies, large β means the policy avoids overfitting to errors in the reward model (the extreme case is $\beta \rightarrow \infty$, in which case both policies become π_{ref}), while small β means the policy has low *bias*, i.e., low error in the case where the model is correct in the sense that $\hat{r} = r^*$ (the extreme case is $\beta \rightarrow 0$, in which case both policies become $x \mapsto \operatorname{argmax}_{a: \pi_{\text{ref}}(a|x) > 0} \hat{r}(x, a)$). Yet, for the same choice of β , $\hat{\pi}_{\chi\text{PO}}$ is significantly more heavy-tailed than $\hat{\pi}_{\text{DPO}}$, a consequence of the pessimism induced by χ^2 -regularization; see Figure 3, which plots the action distribution for both policies as a function of β .

B.3 AN ILLUSTRATIVE EXAMPLE

We now give a concrete example in which χ PO allows the user to tune β to achieve tight statistical rates, yet no choice of β for DPO leads to comparable performance (effectively, any choice of β is either susceptible to overoptimization, or has unacceptably high bias). This illustrates the favorable tradeoff between bias and overoptimization achieved by χ PO.

Let $n \in \mathbb{N}$ with $n \geq 2$ be given. We consider a problem instance with $\mathcal{X} = \{\emptyset\}$ and $\mathcal{A} = \{a_0, a_1, a_2, a_3\}$. We define π_{ref} via

$$\pi_{\text{ref}}(a_0) = \frac{1}{2}, \quad \pi_{\text{ref}}(a_1) = \pi_{\text{ref}}(a_2) = \frac{1}{2n}, \quad \text{and} \quad \pi_{\text{ref}}(a_3) = \frac{n-2}{2n}.$$

We define a reward class with two reward functions $\mathcal{R} := \{r_1, r_2\}$ as follows. For $i \in \{1, 2\}$:

$$r_i(a_0) = 1/2, \quad r_i(a_i) = 1, \quad r_i(a_j) = 0, \quad \forall j \neq i.$$

Let $\beta > 0$ be fixed. To compare χ PO and DPO, we consider their behavior when invoked with the induced policy classes $\Pi_{\chi\text{PO}, \beta}$ and $\Pi_{\text{DPO}, \beta}$ defined above. Recall that with this choice, the two algorithms can be interpreted as fitting a reward model \hat{r} using maximum likelihood (Eq. (3)) and returning the policies $\hat{\pi}_{\chi\text{PO}}(a | x) = \pi_{\text{ref}}(a | x) \cdot \phi^{-1}(\beta^{-1}(\hat{r}(x, a) - Z_{\beta, \hat{r}}(x)))$ and $\hat{\pi}_{\text{DPO}}(a | x) = \pi_{\text{ref}}(a | x) \cdot \exp(\beta^{-1}(\hat{r}(x, a) - Z_{\beta, \hat{r}; \text{KL}}(x)))$, respectively.

Suppose that r_1 is the true reward function. It is hopeless (information-theoretically) to compete with the unconstrained optimal action a_1 , as we are in a sample-starved regime where $C^{a_1} = 2n$ (in the language of Eq. (13)). Indeed, one can show (see proof of Proposition A.1 in Appendix A) that with constant probability, none of the examples in the offline dataset $\mathcal{D}_{\text{pref}}$ contain actions a_1 or a_2 . Under this event, which we denote by \mathcal{E} , the value for the maximum likelihood objective in Eq. (3) is identical for r_1 and r_2 , so we may obtain $\hat{r} = r_2$ (due to adversarial tie-breaking). However, in spite of the fact that the policies $\hat{\pi}_{\chi\text{PO}}$ and $\hat{\pi}_{\text{DPO}}$ are induced by the same (incorrect) reward function $\hat{r} = r_2$, they produce very different action distributions, as highlighted in Figure 3.

To understand this, note that even in the sample-starved regime, we can still hope to compete with the ‘‘baseline’’ action a_0 ; Figure 1 shows that χ PO has low regret against this action, while DPO has high regret. In particular, since $C^{a_0} = 2$, Theorem 3.1 (Eq. (13)) implies that χ PO achieves

$$J(a_0) - J(\hat{\pi}_{\chi\text{PO}}) \lesssim \sqrt{\frac{1}{n}} + \beta + \beta^{-1} \frac{1}{n},$$

and setting $\beta \propto \sqrt{\frac{1}{n}}$ leads to $J(a_0) - J(\hat{\pi}_{\chi\text{PO}}) \lesssim \sqrt{\frac{1}{n}}$. This is a consequence of the pessimistic, heavy-tailed nature of $\hat{\pi}_{\chi\text{PO}}$ (cf. Proposition B.2), which places no more than β^{-1}/n probability mass on the (incorrect) greedy action a_2 for $\hat{r} = r_2$, thereby correctly capturing the inherent uncertainty in the reward for this action.

On the other hand, it is straightforward to show that for all possible values $\beta \leq (2 \log n)^{-1}$, the DPO policy $\hat{\pi}_{\text{DPO}}$ has regret

$$J(a_0) - J(\hat{\pi}_{\text{DPO}}) \geq \frac{1}{2} \left(1 - \frac{1}{1 + \frac{1}{n} e^{\frac{1}{2}} + (1 - \frac{1}{n}) e^{-\frac{1}{2\beta}}} \right) - \frac{1}{2n} \geq \Omega(1)$$

whenever $n \geq 2$. This is because when $\beta \leq (2 \log n)^{-1}$, $\hat{\pi}_{\text{DPO}}$ assigns excessively high probability to the incorrect greedy action a_2 , an instance of overoptimization. Meanwhile, larger choices for β lead to excessively large bias in general (see Appendix A.1 for a more sophisticated construction which extends this lower bound to all possible β). In other words, as illustrated in Figure 1, no choice of β gives a favorable tradeoff between overoptimization and bias.

To summarize, for DPO, large values of β are required to avoid overfitting to the reward function, incurring high bias. Meanwhile, χ PO avoids overoptimization using comparatively small values for β , yet has bias no worse than that of DPO, thereby striking a better tradeoff. We mention that the ‘‘DPO+SFT’’ algorithm of Liu et al. (2024); Cen et al. (2024); Fisch et al. (2024) also fails on the construction above; see Proposition A.1 in Appendix A.1 for details.

Remark B.1 (DPO decreases probabilities of preferred and rejected responses). *Various recent works have noted an empirical phenomenon in which DPO decreases the probabilities for both preferred and rejected responses throughout training (Yuan et al., 2024; Pal et al., 2024; Rafailov et al., 2024b). Interestingly, we observe that the example above exhibits this phenomenon. Notably, if $\beta < (2 \log n)^{-1}$, then under the event \mathcal{E} in which the offline dataset $\mathcal{D}_{\text{pref}}$ does not contain the actions*

a_1 or a_2 (so that $\hat{r} = r_2$), we observe that $\hat{\pi}_{\text{DPO}}(a_0) = \frac{\frac{1}{2} e^{\frac{1}{2\beta}}}{\frac{1}{2} e^{\frac{1}{2\beta}} + \frac{1}{2n} e^{\frac{1}{\beta}} + \frac{n-1}{2n}} < \frac{1}{2} = \pi_{\text{ref}}(a_0)$, and for

all $i > 2$, $\hat{\pi}_{\text{DPO}}(a_i) = \frac{\frac{1}{2n}}{\frac{1}{2} e^{\frac{1}{2\beta}} + \frac{1}{2n} e^{\frac{1}{\beta}} + \frac{n-1}{2n}} < \frac{1}{2n} = \pi_{\text{ref}}(a_i)$. We conclude that for all $a \in \mathcal{D}_{\text{pref}}$,

$$\hat{\pi}_{\text{DPO}}(a) < \pi_{\text{ref}}(a).$$

We emphasize that this behavior arises due to the use of function approximation. When the reward class \mathcal{R} (equivalently, the policy class $\Pi_{\text{DPO},\beta}$) is restricted, the algorithm can aggressively (and incorrectly) extrapolate rewards for actions outside the dataset and, in doing so, inadvertently decrease the probabilities for preferred responses in the dataset. Meanwhile, in the same parameter range, χPO satisfies (see Figure 3)

$$\widehat{\pi}_{\chi\text{PO}}(a_0) > \pi_{\text{ref}}(a_0),$$

highlighting that pessimism can mitigate this phenomenon.

B.4 NONTRIVIALITY AND ROLE OF V_{\max} PARAMETER

To close this section, we discuss the role of the V_{\max} parameter (Assumption 3.2) used in the analysis of χPO (Theorem 3.1) in depth, motivating it from the perspective of the induced policy class $\Pi_{\chi\text{PO},\beta}$ from Section 4.2.

Assumption 3.2 effectively implies that all policies $\pi \in \Pi$ satisfy $\|\frac{\pi}{\pi_{\text{ref}}}\|_{\infty} \lesssim \frac{V_{\max}}{\beta}$; in other words, the policy class we use in χPO satisfies *all-policy L_{∞} -concentrability* with $\max_{\pi \in \Pi} \mathcal{C}_{\infty}^{\pi} \lesssim \frac{V_{\max}}{\beta}$. This might seem to trivialize the offline alignment problem at first glance, since it would suffice to prove a generalization guarantee based on all-policy concentrability, and then plug this bound in. We will show that this is not the case, and that this is actually an intrinsic feature of χ^2 -regularization.

In more detail, recall that for χPO , we require the realizability assumption that $\pi_{\beta}^* \in \Pi$ (Assumption 3.1), where π_{β}^* is the optimal mixed χ^2 -regularized policy that satisfies $r^*(x, a) = \beta \phi\left(\frac{\pi_{\beta}^*(a|x)}{\pi_{\text{ref}}(a|x)}\right) + Z_{\beta, r^*}(x)$. This policy, via Proposition B.2, satisfies $\|\frac{\pi_{\beta}^*}{\pi_{\text{ref}}}\|_{\infty} \lesssim \frac{R_{\max}}{\beta}$, so from a statistical perspective, we can take Assumption 3.2 to hold without loss of generality by removing any policy that violates this bound. In addition, as highlighted by Example 3.1, if we begin from a class of bounded reward models \mathcal{R} with $r^* \in \mathcal{R}$, Assumption 3.2 holds with $V_{\max} \lesssim R_{\max}$ for the induced class $\Pi_{\chi\text{PO},\beta}$ defined in Eq. (24), even though knowledge of such a reward model class is a mild statistical assumption that clearly does not trivialize the learning problem.

On the other hand, for DPO, a minimal assumption is that $\pi_{\beta; \text{KL}}^* \in \Pi$ (Xie et al., 2024), where $\pi_{\beta; \text{KL}}^*$ is the optimal KL-regularized policy that satisfies $r^*(x, a) = \beta \log \frac{\pi_{\beta; \text{KL}}^*(a|x)}{\pi_{\text{ref}}(a|x)} + Z_{\beta, r^*; \text{KL}}(x)$. Unlike the optimal mixed χ^2 -regularized policy, $\pi_{\beta; \text{KL}}^*$ has $\frac{\pi_{\beta; \text{KL}}^*(a|x)}{\pi_{\text{ref}}(a|x)} \gtrsim \exp\left(\frac{R_{\max}}{\beta}\right)$. This means that it is impossible to find a policy class that simultaneously (1) realizes $\pi_{\beta; \text{KL}}^*$, and (2) satisfies all-policy concentrability with $\max_{\pi \in \Pi} \mathcal{C}_{\infty}^{\pi} \ll \exp\left(\frac{R_{\max}}{\beta}\right)$. As the bias of DPO is unacceptably large unless $\beta = \text{poly}(1/n)$ (the “small- β ” regime), this leads to vacuous guarantees.

In view of these observations, our analysis of χPO can be interpreted as (implicitly) showing that for any bounded reward class \mathcal{R} , there exists a policy class Π (precisely, the class $\Pi_{\chi\text{PO},\beta}$ defined in Eq. (24)) such that the following properties hold:

1. **Bounded bias.** For every $r \in \mathcal{R}$, there exists $\pi_r \in \Pi$ such that for all policies π^* , $J_r(\pi^*) - J_r(\pi_r) \lesssim \beta \cdot \mathcal{C}^{\pi^*}$.
2. **Bounded overoptimization.** For all $\pi \in \Pi$, $\|\frac{\pi}{\pi_{\text{ref}}}\|_{\infty} \lesssim \frac{R_{\max}}{\beta}$.

We view this as an interesting and non-trivial contribution in its own right. We mention in passing that while it is indeed possible to analyze χPO by first proving a sample complexity guarantee based on all-policy concentrability and then using that $\max_{\pi \in \Pi} \mathcal{C}_{\infty}^{\pi} \lesssim \frac{V_{\max}}{\beta}$, this would lead to a loose bound relative to Theorem 3.1.

C SAMPLE COMPLEXITY GUARANTEES FOR χ^2 -RLHF

The χ^2 -regularization framework we consider (Section 3.1) can be used to derive algorithms beyond just χPO , and we expect it to find broader use. To highlight this, in this section we analyze the algorithm that directly optimizes a variant of the χ^2 -regularized RLHF objective in Eq. (6); this can be accomplished via policy optimization methods such as PPO, in the vein of classical RLHF approaches to offline alignment (Christiano et al., 2017; Bai et al., 2022; Ouyang et al., 2022; von Werra et al., 2020). As we will show, a benefit of directly optimizing the RLHF objective is that it

Algorithm 2 χ^2 -RLHF

input: Reference policy π_{ref} , preference dataset $\mathcal{D}_{\text{pref}}$, unlabeled context dataset \mathcal{D}_x , χ^2 -regularization coefficient $\beta > 0$, smoothing parameter $\eta \geq 0$.

1: **Estimate reward model via maximum likelihood:**

$$\hat{r} \leftarrow \operatorname{argmax}_{r \in \mathcal{R}} \sum_{(x, a_+, a_-) \in \mathcal{D}_{\text{pref}}} \log [\sigma (r(x, a_+) - r(x, a_-))]. \quad (26)$$

2: Define χ^2 -regularized RLHF objective:

$$\hat{J}_{\beta, \eta}(\pi) := \frac{1}{n_x} \sum_{x \in \mathcal{D}_x} \left(\mathbb{E}_{a \sim \pi(\cdot | x)} [\hat{r}(x, a)] - \beta \sum_a \frac{\pi^2(a | x)}{\pi_{\text{ref}}(a | x) + \eta \pi(a | x)} \right).$$

3: **Policy optimization:** Compute $\hat{\pi} \in \Pi$ such that

$$\hat{J}_{\beta, \eta}(\hat{\pi}) \geq \max_{\pi \in \Pi} \hat{J}_{\beta, \eta}(\pi) - \varepsilon_{\text{opt}}.$$

4: **return:** $\hat{\pi}$.

allows us to provide guarantees that avoid dependence on the V_{max} parameter in [Theorem 3.1](#), which may lead to improvement when Π includes policies with very large or very small density ratios $\frac{\pi}{\pi_{\text{ref}}}$.

Algorithm. Our algorithm, χ^2 -RLHF is displayed in [Algorithm 2](#). At the population level, the algorithm aims to optimize a variant of [Eq. \(7\)](#) that incorporates a small but important modification that allows us to avoid dependencies on $\frac{\pi}{\pi_{\text{ref}}}$. Given *smoothing parameter* $\eta > 0$, define the *smoothed χ^2 -divergence* $D_{\chi^2; \eta}(\pi \parallel \pi_{\text{ref}}) := \mathbb{E}_{\pi} \left[\frac{\pi(a | x)}{\pi_{\text{ref}}(a | x) + \eta \pi(a | x)} \right]$. We aim to find

$$\begin{aligned} \operatorname{argmax}_{\pi} J_{\beta, \eta}(\pi) &:= \mathbb{E}_{\pi} [r^*(x, a)] - \beta D_{\chi^2; \eta}(\pi \parallel \pi_{\text{ref}}) \\ &= \operatorname{argmax}_{\pi} \mathbb{E}_{\pi} \left[r^*(x, a) - \beta \frac{\pi(a | x)}{\pi_{\text{ref}}(a | x) + \eta \pi(a | x)} \right]. \end{aligned} \quad (25)$$

The smoothing parameter η effectively clips the policy ratio in $D_{\chi^2; \eta}(\pi \parallel \pi_{\text{ref}})$ where $\pi_{\text{ref}}(a | x) \ll \eta \pi(a | x)$; $D_{\chi^2}(\cdot \parallel \cdot)$ corresponds to the special (non-clipped) case where $\eta = 0$. In particular, clipping ensures a uniform bound of the form $D_{\chi^2; \eta}(\pi \parallel \pi_{\text{ref}}) \leq \eta^{-1}$, whereas the best bound we can hope for with the unclipped χ^2 -divergence is $D_{\chi^2}(\pi \parallel \pi_{\text{ref}}) = \mathbb{E}_{\pi} \left[\frac{\pi(a | x)}{\pi_{\text{ref}}(a | x)} \right] \leq \mathcal{C}_{\infty}^{\pi}$. For this reason, smoothing will allow us to obtain guarantees that avoid dependence on all-policy concentrability or parameters similar to V_{max} .

To optimize [Eq. \(25\)](#), [Algorithm 2](#) takes two datasets as input, along with a user-specified reward model class \mathcal{R} and policy class Π . The first dataset, $\mathcal{D}_{\text{pref}}$, is labeled with human preferences, and is used to learn a reward model \hat{r} via maximum likelihood estimation in [Line 1](#). The second, \mathcal{D}_x , contains *only unlabeled contexts* sampled from ρ , and is utilized in [Line 3](#) to learn a policy that approximately maximizes an empirical version of [Eq. \(25\)](#). Importantly, because [Line 3](#) involves an empirical expectation over only contexts, it is a purely computational problem that we can solve using algorithms like PPO; we allow for tolerance ε_{opt} in [Line 3](#) to accommodate optimization error from such algorithms. By using unlabeled contexts in [Line 3](#), we can obtain tighter guarantees when \mathcal{D}_x is large. This is often the case in practice, where unlabeled contexts are cheap to obtain, but preferences can be expensive to query.

Theoretical guarantees. To analyze χ^2 -RLHF, we make similar assumptions to those utilized in [Theorem 3.1](#) for χ PO. Since χ^2 -RLHF utilizes separate reward and policy classes, we require realizability conditions for both. Namely, \mathcal{R} must be able to express the true reward function r^* , and Π must include the optimal policy for the regularized RLHF objective in [Eq. \(25\)](#).

Assumption C.1. *The reward function class satisfies $r^* \in \mathcal{R}$, and is bounded so that $r(x, a) \in [0, R_{\text{max}}]$ for all $r \in \mathcal{R}$ and $(x, a) \in \mathcal{X} \times \mathcal{A}$.*

Assumption C.2. *The policy class Π satisfies $\pi_{\beta, \eta}^* \in \Pi$, where $\pi_{\beta, \eta}^*$ is the optimal policy for [Eq. \(25\)](#).*

Below is our main sample complexity guarantee for χ^2 -RLHF. While it is stated for a fixed, β -dependent smoothing parameter for compactness, the most general version of this result ([Theorem K.1](#)) applies to general η .

Theorem C.1. *Let $\beta > 0$ be given, and suppose [Assumptions C.1](#) and [C.2](#) hold any $\eta \in \left[0, \frac{\beta}{8R_{\max}}\right]$. With probability at least $1 - \delta$, χ^2 -RLHF ([Algorithm 2](#)) produces a policy $\hat{\pi}$ such that for all policies π^* simultaneously, we have*

$$J(\pi^*) - J(\hat{\pi}) \lesssim R_{\max} e^{2R_{\max}} \cdot \sqrt{\frac{C^{\pi^*} \log(|\mathcal{R}|/\delta)}{n}} + \beta \cdot C^{\pi^*} + \beta^{-1} \cdot \frac{R_{\max}^2 e^{4R_{\max}} \log(|\mathcal{R}|/\delta)}{n} + R_{\max} \sqrt{\frac{\log(|\Pi|/\delta)}{n_x}} + \varepsilon_{\text{opt}}.$$

In particular, given any comparator policy π^* , we can choose the regularization parameter β to achieve

$$J(\pi^*) - J(\hat{\pi}) \lesssim R_{\max} e^{2R_{\max}} \cdot \sqrt{\frac{C^{\pi^*} \log(|\mathcal{R}|/\delta)}{n}} + R_{\max} \sqrt{\frac{\log(|\Pi|/\delta)}{n_x}} + \varepsilon_{\text{opt}}. \quad (27)$$

Above, we see that χ^2 -RLHF, like χ PO, has sample complexity that scales only with the single-policy concentrability coefficient C^{π^*} , and holds for all comparator policies π^* simultaneously. Since the choice of β induces a similar bias-overoptimization tradeoff in the first statement of [Theorem C.1](#) as it did in [Theorem 3.1](#) for χ PO, we focus our discussion on the guarantee for a tuned choice of β ([Eq. \(27\)](#)). The first term in [Eq. \(27\)](#) accounts for the reward estimation error ([Line 1](#)) and scales with C^{π^*} ; as before, this accounts for how well rewards estimated from π_{ref} transfer to other candidate policies. The second term in [Eq. \(27\)](#) accounts for the statistical error from sampled contexts used in [Line 3](#) for policy optimization. In particular, it is possible to drive this term to be much smaller than the first by using a larger unlabeled context dataset, which is typically far cheaper to acquire.

Computationally efficiency. [Theorem C.1](#) bounds the sample complexity of χ^2 -RLHF under the assumption that we can solve [Line 3](#) up to ε_{opt} -accuracy. This is a purely computational problem, and in practice it can be solved using policy gradient methods such as PPO.

Comparison to χ PO. Unlike χ PO ([Theorem 3.1](#)), [Theorem C.1](#) has no dependence on the parameter V_{\max} or quantities such as $\frac{\pi}{\pi_{\text{ref}}} \leq \max_{\pi} C_{\infty}^{\pi}$. We primarily attribute this to the fact that χ^2 -RLHF uses an explicit reward function class \mathcal{R} , and normalizing or clipping it to the reward range R_{\max} is both natural and routinely done in practice ([Shah et al., 2015](#); [Christiano et al., 2017](#); [Ouyang et al., 2022](#)). In comparison, the implicit reward models induced by the policy class Π in χ PO can have larger range, and clipping the policy class in χ PO directly, e.g., so that $|\beta\phi(\frac{\pi}{\pi_{\text{ref}}})|$ is bounded, is misguided, because the policy class may lose realizability ([Assumption 3.1](#)). This is because $r^*(x, a) = \beta\phi\left(\frac{\pi_{\beta}^*(a|x)}{\pi_{\text{ref}}(a|x)}\right) + Z_{\beta, r^*}(x)$, and the normalization factor Z_{β, r^*} cannot be reasonably accounted for when clipping Π . While the V_{\max} ([Assumption 3.2](#)) parameter involves pairs of action probabilities, and thereby sidesteps the normalization constant issue, it may not always be practical to modify Π so that V_{\max} is bounded, since this would require checking all pairs of each policy’s action probabilities.

However, using an explicit reward function class alone is not enough. As discussed previously, when we move from implicit to explicit χ^2 -regularization, incorporating the smoothing parameter η in [Eq. \(25\)](#) is essential to avoid statistical errors due to policies with large density ratios when we approximate the χ^2 -regularizer with empirical data. A careful choice of $\eta = \beta/R_{\max}$ in [Theorem C.1](#) balances the benefits of clipping against the bias it introduces. Without smoothing (i.e., $\eta = 0$), a guarantee that depends on $\max_{\pi} C_{\infty}^{\pi}$ for χ^2 -RLHF would be unavoidable, since the sample complexity must scale with the range of the problem, which grows with the magnitude of the regularizer. See [Corollary K.2](#) in [Appendix K](#) for a guarantee in the case where $\eta = 0$, which highlights this.

D χ PO FOR GENERAL PREFERENCE MODELS

All of our results so far concern the Bradley-Terry model ([Eq. \(1\)](#)), which, as highlighted in prior work, is somewhat restrictive. Thus, in this section, we turn our attention to offline alignment under a

1458 *general preference model* which does not assume transitivity (Munos et al., 2023; Wang et al., 2023b;
1459 Swamy et al., 2024; Rosset et al., 2024; Ye et al., 2024). The setup is the same as Section 2, but we
1460 assume that for a given context x and pair of actions (a, b) , the preference $y \in \{0, 1\}$ is generated via
1461 a Bernoulli Distribution

$$1462 \quad y \sim \text{Ber}(\mathcal{P}^*(a \succ b \mid x)), \quad (28)$$

1463 where $\mathcal{P}^*(a \succ b \mid x) \in [0, 1]$ is a general preference distribution. For a pair of policies π, π' , let
1464 $\mathcal{P}^*(\pi \succ \pi') := \mathbb{E}_{x \sim \rho}[\mathcal{P}^*(\pi(x) \succ \pi'(x) \mid x)]$. Following Wang et al. (2023b); Munos et al. (2023);
1465 Swamy et al. (2024), we consider the *minimax winner* (Kreweras, 1965; Simpson, 1969; Kramer,
1466 1973; Fishburn, 1984) or *von Neumann winner* (Dudík et al., 2015) as a solution concept:

$$1467 \quad \pi_{\text{MW}} := \operatorname{argmax}_{\pi \in \Pi} \min_{\pi' \in \Pi} \mathcal{P}^*(\pi \succ \pi').$$

1470 It will be useful to slightly reparameterize this formulation by introducing the preference function
1471 $\ell^*(x, a, b) := 2\mathcal{P}^*(a \succ b \mid x) - 1$. Note that for any well-defined preference model, we have
1472 $\mathcal{P}^*(a \succ b \mid x) + \mathcal{P}^*(b \succ a \mid x) = 1$ for all x, a, b , which indicates that ℓ^* satisfies skew symmetry:
1473

$$1474 \quad \ell^*(x, a, a) = 0, \quad \ell^*(x, a, b) + \ell^*(x, b, a) = 0, \quad \forall x \in \mathcal{X}, a, b \in \mathcal{A}.$$

1475 Furthermore, the minimax winner above is equivalent to

$$1476 \quad \pi_{\text{MW}} := \operatorname{argmax}_{\pi \in \Pi} \min_{\pi' \in \Pi} \ell^*(\pi, \pi'), \quad (29)$$

1477 where $\ell^*(\pi, \pi') := \mathbb{E}_{x \sim \rho, a \sim \pi(x), b \sim \pi'(x)}[\ell^*(x, a, b)]$. Concretely, our goal is to use the logged
1478 preference data $\mathcal{D}_{\text{pref}} = \{(x, a_+, a_-)\}$ (with (a_+, a_-) labeled according to Eq. (28)) to compute a
1479 policy $\hat{\pi}$ that is an ε -approximate minimax winner, in the sense that

$$1480 \quad \text{DG}(\hat{\pi}) := \max_{\pi \in \Pi} \ell^*(\pi, \hat{\pi}) - \min_{\pi \in \Pi} \ell^*(\hat{\pi}, \pi) \leq \varepsilon. \quad (30)$$

1484 D.1 IMPOSSIBILITY OF SINGLE-POLICY CONCENTRABILITY UNDER GENERAL PREFERENCES

1485 While the general preference framework above is more powerful than the Bradley-Terry model, we
1486 now show that there is a statistical cost for this generality. In particular, our first result in this section
1487 shows that in contrast to the Bradley-Terry model, it is not possible to achieve sample complexity
1488 guarantees that scale with single-policy concentrability under general preferences, even when the
1489 learner has access to a small class of preference models \mathcal{P} that contains the true preference model \mathcal{P}
1490 (i.e., $\mathcal{P}^* \in \mathcal{P}$).

1491 **Theorem D.1** (Impossibility of single-policy concentrability under general preferences). *There*
1492 *exists two problem instances $\theta_1 = (\rho, \mathcal{P}_1^*, \Pi)$ and $\theta_2 = (\rho, \mathcal{P}_2^*, \Pi)$ differing only in their ground*
1493 *truth preference model, a data collection policy π_{ref} , and a preference model class $\mathcal{P} = \{\mathcal{P}_1^*, \mathcal{P}_2^*\}$*
1494 *with $|\mathcal{P}| = 2$ such that the following hold:*

- 1495 1. *For both instances, the single-policy L_∞ -concentrability coefficient for a minimax winner is*
1496 *bounded: $\min_{\pi_{\text{MW}}} \mathcal{C}_\infty^{\pi_{\text{MW}}} \leq 2$.⁷*
- 1497 2. *For any $n \in \mathbb{N}$ and any algorithm Alg which derives a policy $\hat{\pi}$ from a dataset $\mathcal{D}_{\text{pref}}$ of n samples,*
1498 *there exists an instance $\theta \in \{\theta_1, \theta_2\}$ such that π_{ref} incurs constant suboptimality:*

$$1500 \quad \min_{\text{Alg}} \max_{i \in \{1, 2\}} \mathbb{E}_{\mathcal{D}_{\text{pref}} \sim \theta_i} [\text{DG}(\text{Alg}(\mathcal{D}_{\text{pref}}); \theta_i)] \geq \frac{1}{8},$$

1501 where $\text{DG}(\pi; \theta)$ is the duality gap for policy π on instance θ .

1502 This lower bound is inspired by similar results in the literature on offline RL in two-player zero-sum
1503 Markov games (Cui and Du, 2022). However, the lower bound constructions in Cui and Du (2022)
1504 cannot be directly applied as-is, because they do not satisfy the skew-symmetry property required
1505 by the general preference alignment framework. Our lower bound highlights that even under skew-
1506 symmetry, it is impossible to achieve single-policy concentrability for offline learning in two-player
1507 zero-sum games.

1508 ⁷In general, the minimax winner may not be unique. We compete against the minimax winner with the *best*
1509 possible single-policy concentrability coefficient.

Algorithm 3 Iterative χ PO for General Preferences

- 1: **Input:** labeled preference dataset $\mathcal{D}_{\text{pref}}$, preference model class \mathcal{L} , regularization coefficient β , stepsize η , total number of iterations T .
- 2: **Initialize:** $\pi^1 = \pi_{\text{ref}}$.
- 3: Learn a preference model $\hat{\ell}$ via least-squares regression:

$$\hat{\ell} = \operatorname{argmin}_{\ell \in \mathcal{L}} \sum_{(x, a_+, a_-) \in \mathcal{D}_{\text{pref}}} (\ell(x, a_+, a_-) - 1)^2.$$

- 4: Collect m samples $\mathcal{D}_x = \{(x, a, b)\}$ where each sample is drawn i.i.d. from $x \sim \rho, a \sim \pi_{\text{ref}}(x), b \sim \pi_{\text{ref}}(x)$.
- 5: **for** $t = 1, \dots, T$ **do**
- 6: Sample $b_t \sim \pi^t(x)$ and let $\hat{r}^t(x, a) = \hat{\ell}(x, a, b_t)$ for all $x \in \mathcal{X}, a \in \mathcal{A}$.
- 7: Compute

$$\pi^{t+1} = \operatorname{argmin}_{\pi \in \Pi} \sum_{(x, a, b) \in \mathcal{D}_x} \left(\operatorname{clip}_4 \left(f_{\pi, \pi^t}^{\beta, \eta}(x, a, b) \right) - (\hat{r}^t(x, a) - \hat{r}^t(x, b)) \right)^2, \quad (32)$$

where $f_{\pi, \pi^t}^{\beta, \eta}(x, a, b)$ is defined in Eq. (31).

- 8: **Output:** $\hat{\pi} = \operatorname{unif}(\{\pi^t\}_{t=1}^T)$.

D.2 ITERATIVE χ PO FOR GENERAL PREFERENCES

In spite of the hardness in the prequel, we now show that an iterative variant of χ PO—based on self-play—can learn a near-optimal minimax winner under the general preference model under a new local coverage condition—a condition that is stronger than the single policy concentrability but much weaker than global/all-policy concentrability and the notion of unilateral concentrability introduced by Cui and Du (2022).

Our algorithm, Iterative χ PO, is described in Algorithm 3, and consists of two main steps.

Preference model estimation via least squares regression on $\mathcal{D}_{\text{pref}}$. We first (Line 3) learn a preference model from the offline preference dataset $\mathcal{D}_{\text{pref}}$. We assume access to a preference function class \mathcal{L} which is realizable in the sense that $\ell^* \in \mathcal{L}$ and where all $\ell \in \mathcal{L}$ satisfy skew-symmetry, and we will estimate ℓ^* rather than \mathcal{P}^* . We perform least-squares regression on $\mathcal{D}_{\text{pref}}$ with \mathcal{L} to learn ℓ^* :

$$\hat{\ell} = \operatorname{argmin}_{\ell \in \mathcal{L}} \sum_{(x, a_+, a_-) \in \mathcal{D}_{\text{pref}}} (\ell(x, a_+, a_-) - 1)^2.$$

Policy optimization with iterative χ PO update. Given the estimated model $\hat{\ell}$, we compute an approximate minimax winner using an iterative regression scheme inspired by Gao et al. (2024). We proceed in T iterations (Line 5), where at each iteration t , we define an iteration-dependent reward function $\bar{r}^t(x, a)$ based on the current policy π^t as

$$\bar{r}^t(x, a) = \mathbb{E}_{b \sim \pi^t(x)} [\hat{\ell}(x, a, b)], \quad \forall x \in \mathcal{X}, a \in \mathcal{A}.$$

Then, for all $\pi, \pi' \in \Pi$, we define a policy-dependent predictor $f_{\pi, \pi'}^{\beta, \eta}(x, a, b)$, whose motivation will be described in detail momentarily, as follows:

$$\begin{aligned} f_{\pi, \pi'}^{\beta, \eta}(x, a, b) := & \left(1 + \frac{1}{\eta} \right) \cdot \left(\beta \phi \left(\frac{\pi(a|x)}{\pi_{\text{ref}}(a|x)} \right) - \beta \phi \left(\frac{\pi(b|x)}{\pi_{\text{ref}}(b|x)} \right) \right) \\ & - \frac{1}{\eta} \left(\beta \phi \left(\frac{\pi'(a|x)}{\pi_{\text{ref}}(a|x)} \right) - \beta \phi \left(\frac{\pi'(b|x)}{\pi_{\text{ref}}(b|x)} \right) \right) \end{aligned} \quad (31)$$

Using $f_{\pi, \pi^t}^{\beta, \eta}(x, a, b)$ as a policy-parameterized regression function, we (Line 7) compute the next policy π^{t+1} by solving a least-squares regression problem in which the Bayes optimal solution is the *relative reward* $\bar{r}^t(x, a) - \bar{r}^t(x, b)$ for iteration t .

Let us now explain the intuition behind the predictor $f_{\pi, \pi'}^{\beta, \eta}(x, a, b)$. Suppose that the regression step in [Line 7](#) learns a predictor that can perfectly model the relative reward, i.e.,

$$\forall x, a, b, \quad f_{\pi^{t+1}, \pi^t}^{\beta, \eta}(x, a, b) = \bar{r}^t(x, a) - \bar{r}^t(x, b),$$

In this case, we can show that the returned policy π^{t+1} is the optimal policy for the following mixed χ^2 -regularized RL objective:

$$\pi^{t+1}(x) = \operatorname{argmax}_{p \in \Delta(\mathcal{X})} \left\{ \mathbb{E}_{a \sim p} [\bar{r}^t(x, a)] - \beta D_{f_{\chi_{\text{mix}}}}(p \parallel \pi_{\text{ref}}(x)) - \frac{\beta}{\eta} B_x(p, \pi^t) \right\}, \quad \forall x \in \mathcal{X}, \quad (33)$$

where $B_x(p, \pi^t)$ is the Bregman divergence induced by the regularizer $p \mapsto D_{f_{\chi_{\text{mix}}}}(p \parallel \pi_{\text{ref}}(x))$, i.e., $B_x(p, q) := D_{f_{\chi_{\text{mix}}}}(p \parallel \pi_{\text{ref}}(x)) - D_{f_{\chi_{\text{mix}}}}(q \parallel \pi_{\text{ref}}(x)) - \langle \nabla D_{f_{\chi_{\text{mix}}}}(q \parallel \pi_{\text{ref}}(x)), p - q \rangle$, $\forall x \in \mathcal{X}$.

Thus, the algorithm can be understood as running mirror descent on the iteration-dependent loss function $-\bar{r}^t$, with $p \mapsto D_{f_{\chi_{\text{mix}}}}(p \parallel \pi_{\text{ref}}(x))$ as a per-context regularizer. This technique draws inspiration from [Chang et al. \(2024\)](#), in which the authors apply a similar regularized mirror descent algorithm to learn the optimal policy for the *reward-based* setting. The motivation for using mixed- χ^2 regularization is exactly the same as in χPO : we want to ensure that $\frac{\pi^{t+1}(a|x)}{\pi_{\text{ref}}(a|x)} \leq 1 + \frac{1}{\beta}$, thereby mitigating overoptimization.

D.3 THEORETICAL ANALYSIS OF ITERATIVE χPO

We now present our main theoretical guarantees for Iterative χPO . We begin by stating a number of statistical assumptions. We first assume that the preference model class contains the ground truth preference function ℓ^* .

Assumption D.1 (Preference function realizability). *The model class \mathcal{L} satisfies $\ell^* \in \mathcal{L}$ where ℓ^* is the ground truth preference function.*

In addition, since [Algorithm 3](#) iteratively applies an χPO update, we require that a policy realizability assumption analogous to [Assumption 3.1](#) holds for each of the sub-problems in [Eq. \(33\)](#). Concretely, we make the following assumption.

Assumption D.2 (Policy realizability for general preferences). *For any policy $\pi \in \Pi$ and $\ell \in \mathcal{L}$, the policy class Π contains the minimizer of the following regularized RL objective:*

$$\bar{\pi}(x; \ell, \pi) := \operatorname{argmax}_{p \in \Delta(\mathcal{X})} \left\{ \mathbb{E}_{a \sim p, b \sim \pi(x)} [\ell(x, a, b)] - \beta D_{f_{\chi_{\text{mix}}}}(p \parallel \pi_{\text{ref}}(x)) - \frac{\beta}{\eta} B_x(p, \pi) \right\}, \quad \forall x \in \mathcal{X}.$$

Finally, we require that the implicit reward functions in [Eq. \(32\)](#) are bounded, analogous to [Assumption 3.2](#).

Assumption D.3 (Bounded implicit rewards for general preferences). *For a parameter $V_{\max} \geq 2$, it holds that for all $\pi, \pi' \in \Pi$, $x \in \mathcal{X}$, and $a, b \in \mathcal{A}$,*

$$\left| f_{\pi, \pi'}^{\beta, \eta}(x, a, b) \right| \leq V_{\max}. \quad (34)$$

Our main guarantee for [Algorithm 3](#) is as follows.

Theorem D.2. *Fix any $\delta \in (0, 1]$. Suppose [Algorithm 3](#) is invoked with $T = \frac{mn}{nV_{\max}^2 + m}$, $\beta = \frac{1}{\sqrt{T}}$, and $\eta = \frac{1}{T}$. Then under [Assumption D.1](#), [Assumption D.2](#) and [Assumption D.3](#), we have that probability at least $1 - \delta$,*

$$\text{DG}(\hat{\pi}) \lesssim \min_{C \geq 1} \left\{ \text{subopt}(\hat{\pi}, C) + C \left(\frac{V_{\max} \log(|\Pi|/\delta)}{\sqrt{m}} + \frac{\log(|\Pi||\mathcal{L}|/\delta)}{\sqrt{n}} \right) \right\},$$

where $\text{subopt}(\hat{\pi}, C) := \max_{\pi \in \Pi} \ell^*(\pi, \hat{\pi}) - \max_{\pi \in \Pi_C} \ell^*(\pi, \hat{\pi})$ and $\Pi_C := \{\pi : \max_{x \in \mathcal{X}} D_{\chi^2}(\pi(x) \parallel \pi_{\text{ref}}(x)) \leq C\}$. In particular, if we define the unilateral concentrability coefficient as

$$C_{\text{uni}} := \max_{\pi \in \Pi, x \in \mathcal{X}, a, b \in \mathcal{A}} \frac{\pi(a|x)\pi_{\text{MW}}(b|x)}{\pi_{\text{ref}}(a|x)\pi_{\text{ref}}(b|x)},$$

then the bound above implies that

$$\text{DG}(\hat{\pi}) \lesssim C_{\text{uni}} \cdot \left(\frac{V_{\max} \log(|\Pi|/\delta)}{\sqrt{m}} + \frac{\log(|\Pi||\mathcal{L}|/\delta)}{\sqrt{n}} \right).$$

The first result gives a tradeoff between the statistical error and the approximation error $\text{subopt}(\hat{\pi}, C)$, which is modulated by the parameter C . This tradeoff is analogous to, but more subtle, than the one for χPO in the reward-based setting. In the reward-based setting, χPO has low regret to the best policy covered π_{ref} . In the general preference setting, [Algorithm 3](#) has small duality gap if, for any policy, there is an approximate best response that is covered by π_{ref} (this implies that $\text{subopt}(\hat{\pi}, C)$ is small for small C). Crucially, [Algorithm 3](#) does not require that all policies are covered by π_{ref} , which is a distinctive feature of mixed χ^2 -regularization and reflects the algorithm's robustness to overoptimization.

The second result concerns the setting where all policies are covered by π_{ref} and is easier to interpret. Indeed, if all $\pi \in \Pi$ satisfy $D_{\chi^2}(\pi \parallel \pi_{\text{ref}}) \leq C^*$, then $\text{subopt}(\hat{\pi}, C^*) = 0$, which implies that we can learn an ε -approximate minimizer using $\tilde{O}(C^*/\varepsilon^2)$ samples. Thus, we obtain a guarantee based on unilateral concentrability ([Cui and Du, 2022](#)), which is a stronger condition, i.e., we always have $\max_{\pi} D_{\chi^2}(\pi \parallel \pi_{\text{ref}}) \leq C_{\text{uni}}$. However, per the above discussion, the first part of [Theorem D.2](#) is stronger than results based on unilateral concentrability and hints at a new notion of coverage for the general preference setting. Lastly, we remark that the parameter V_{max} only affects $\sqrt{1/m}$ in [Theorem D.2](#), so this dependence on this parameter can be mitigated using unlabeled data.

[Theorem D.2](#) is closely related to recent work of [Ye et al. \(2024\)](#), which uses pessimism to learn a regularized minimax winner, and achieves polynomial sample complexity with a concentrability assumption similar to [Theorem D.2](#). However, there are two key differences. First, their learning objective is the KL-regularized minimax winner, while we study the unregularized objective and use χ^2 -regularization. More importantly, their theoretical algorithm is computationally inefficient as it constructs an explicit confidence set for the preference model and performs max-min-style policy optimization. In contrast, our algorithm only requires solving standard supervised learning problems.

Proof sketch. We now sketch the proof of [Theorem D.2](#). For a parameter $C \geq 1$, let $\tilde{\pi}_C$ be the best response to $\hat{\pi}$, restricted to the set Π_C :

$$\tilde{\pi}_C = \operatorname{argmax}_{\pi \in \Pi_C} \mathbb{E}_{x \sim \rho, a \sim \pi(x), b \sim \hat{\pi}(x)} [\ell^*(x, a, b)].$$

Let $\tilde{\pi} = \tilde{\pi}_{\infty}$ be the global best response to $\hat{\pi}$. Recall that $\bar{r}^t(x, a) := \mathbb{E}_{b \sim \hat{\pi}^t(x)} [\hat{\ell}(x, a, b)]$, and let $g_x(p) := \beta D_{f_{\chi^{\text{mix}}}}(p \parallel \pi_{\text{ref}}(x))$ denote the (context-wise) mixed χ^2 -divergence. Due to the skew symmetry of ℓ^* , we only need to bound $\ell^*(\tilde{\pi}, \hat{\pi})$, which can be decomposed into the following terms:

$$\begin{aligned} \ell^*(\tilde{\pi}, \hat{\pi}) &= \text{subopt}(\hat{\pi}, C) + \underbrace{\frac{1}{T} \sum_{t=1}^T (\hat{r}^t(\tilde{\pi}_C) - \hat{r}^t(\pi^t))}_{(1)} + \underbrace{\frac{1}{T} \sum_{t=1}^T (\ell^*(\tilde{\pi}_C, \pi^t) - \hat{\ell}(\tilde{\pi}_C, \pi^t))}_{(2)} \\ &\quad + \underbrace{\frac{1}{T} \sum_{t=1}^T (\bar{r}^t(\tilde{\pi}_C) - \hat{r}^t(\tilde{\pi}_C))}_{(3)} + \underbrace{\frac{1}{T} \sum_{t=1}^T (\hat{r}^t(\pi^t) - \bar{r}^t(\pi^t))}_{(4)}, \end{aligned}$$

where we denote $\bar{r}^t(\pi) := \mathbb{E}_{x \sim \rho, a \sim \pi(x)} [\bar{r}^t(x, a)]$. The decomposition utilizes the fact that $\bar{r}^t(\pi^t) = 0$ and $\bar{r}^t(\tilde{\pi}_C) = \hat{\ell}(\tilde{\pi}_C, \pi^t)$. This implies that we only need to bound terms (1), (2), (3), and (4) to upper bound the gap of $\hat{\pi}$. Note that term (2) is the estimation error of $\hat{\ell}$ under the distribution $x \sim \rho, a \sim \tilde{\pi}_C(x), b \sim \pi^t(x)$, and we can bound it by using similar techniques to the proof of [Theorem 3.1](#). Terms (3) and (4) can be bounded by standard martingale concentration arguments. The challenge of the analysis lies in controlling term (1), because [Algorithm 3](#) applies the χPO update (with function approximation) rather than the idealized per-state policy mirror descent update, which prevents us from utilizing existing techniques in the literature ([Zhan et al., 2023a](#); [Lan, 2023](#)). We solve this difficulty by using the following key lemma:

Lemma D.1. For any $C \geq 0$, we have that for all policies $\pi \in \Pi_C$,

$$\frac{1}{T} \sum_{t=1}^T (\hat{r}^t(\pi) - \hat{r}^t(\pi^t)) \lesssim \frac{\beta C}{\eta T} + \beta C + \frac{\eta}{\beta} - \frac{1}{T} \sum_{t=1}^{T+1} \mathbb{E}_{x \sim \rho} [g_x(\pi^t)]$$

$$+ \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{x \sim \rho} [\langle \hat{r}^t(x, \cdot) - G^t(\pi^{t+1}, x, \cdot), \pi(x) - \pi^{t+1}(x) \rangle],$$

where $G^t(\pi, x, a) := \beta \left((1 + \frac{1}{\eta}) \phi \left(\frac{\pi(a|x)}{\pi_{\text{ref}}(a|x)} \right) - \frac{1}{\eta} \phi \left(\frac{\pi^t(a|x)}{\pi_{\text{ref}}(a|x)} \right) \right)$.

With this lemma, we only need to bound the quantity $\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{x \sim \rho} [\langle \hat{r}^t(x, \cdot) - G^t(\pi^{t+1}, x, \cdot), \pi(x) - \pi^{t+1}(x) \rangle]$ in order to bound term (1). However, this term can be upper bounded by $\mathbb{E}_{x \sim \rho, a \sim \pi(x), b \sim \pi_{\text{ref}}(x)} \left[\left| \hat{r}^t(x, a) - \hat{r}^t(x, b) - f_{\pi^{t+1}, \pi^t}^{\beta, \eta}(x, a, b) \right| \right]$ and $\mathbb{E}_{x \sim \rho, a \sim \pi^{t+1}(x), b \sim \pi_{\text{ref}}(x)} \left[\left| \hat{r}^t(x, a) - \hat{r}^t(x, b) - f_{\pi^{t+1}, \pi^t}^{\beta, \eta}(x, a, b) \right| \right]$, which can then be bounded by applying similar techniques to those in the proof of [Lemma H.3](#). This concludes the proof.

E EXPERIMENTS IN OFFLINE LANGUAGE MODEL ALIGNMENT

E.1 TL;DR SUMMARIZATION

We perform preliminary evaluations of χ PO for offline language model alignment on the TL;DR dataset ([Stiennon et al., 2020](#)), and DPO serves as our baseline of comparison. The reference policy π_{ref} is the Pythia-1b model ([Biderman et al., 2023](#)) pre-trained on SFT data (cleanr1/eleutherai-pythia-1b-deduped__sft__tldr from [Huang et al. \(2022\)](#)), and performance is measured via winrate against a baseline, as judged by GPT-4o. All parameters that are not algorithm-specific, such as the learning rate, are shared by both χ PO and DPO in order to ensure a fair comparison (see [Appendix E.2](#) for details).

Table 1: Winrate on TL;DR Summarization of the model learned by χ PO and DPO, over several choices of the number of training epochs and of the regularization parameter β . Standard error over the 3 seeds is also reported.

| β | Epochs | χ PO winrate (%) | DPO winrate (%) |
|---------|--------|-----------------------|-----------------|
| 0.05 | 1 | 56.5 ± 1.3 | 55.8 ± 2.1 |
| | 2 | 56.1 ± 0.6 | 50.3 ± 0.8 |
| | 4 | 48.0 ± 1.6 | 38.0 ± 0.7 |
| 0.005 | 1 | 50.6 ± 1.6 | 14.7 ± 3.9 |
| | 2 | 52.8 ± 2.3 | 3.4 ± 1.5 |
| | 4 | 51.6 ± 0.8 | 0.5 ± 0.2 |

In [Table 1](#) we display the winrates of χ PO and DPO over several choices of training epochs, as well as regularization parameter β . The winrate corresponds to the final checkpoint learned each algorithm for each set of hyperparameters. We consider $\beta = 0.05$ and 1 epoch of training to be a standard setup for DPO ([Gao et al., 2024](#); [Guo et al., 2024](#); [Rafailov et al., 2024a](#)), and, as we are particularly concerned with regimes where overoptimization is of concern, we additionally analyze performance when epochs are increased, and/or β is decreased (corresponding to less regularization).

Over all choices of β and epochs, χ PO achieves a higher average winrate than DPO. While the difference is not significant for $\beta = 0.05$ and 1 training epoch, the performance gap grows significantly as the number of training epochs increases, demonstrating the robustness of χ PO to overoptimization in offline alignment. Further, while DPO degrades completely for $\beta = 0.005$, χ PO is robust over two orders of magnitude of β , reinforcing trends seen earlier in [Figure 1](#) and the more favorable bias-overoptimization tradeoff from our theoretical analysis.

In addition, χ PO exhibits better performance and robustness longitudinally throughout training, as shown in [Appendix E.1](#). While DPO peaks early with high variance around 0.5 epochs and degrades thereafter, χ PO continues to improve gently then plateaus over the last epoch. Further, for the same regularization parameter β , the χ PO policy has significantly lower KL-divergence relative to π_{ref} , demonstrating that the χ^2 -regularization is both a stronger regularizer and one that effectively mitigates overoptimization.

These observations are encouraging, and we are evaluating χ PO in a more thorough manner, e.g. with more seeds, on models of different sizes, on more datasets. In future work, we plan to investigate




Figure 4: (left) TL;DR Summarization winrate recorded longitudinally over 2 epochs of training every 250 steps. Shaded area displays ± 1 standard error over 3 seeds. At 1 epoch χ PO already obtains better performance, and continues to improve over the course of training, while DPO degrades over time. (right) KL divergence $D_{\text{KL}}(\hat{\pi} \parallel \pi_{\text{ref}})$ averaged over 2 of the seeds. For the same β , χ PO constrains the learned policy to be significantly closer to π_{ref} , thereby striking a better bias-variance tradeoff.

questions such as the effect of different β , γ , and other parameters on the KL divergence and overall performance.

E.2 EXPERIMENT DETAILS

Dataset. For training we use `trl-internal-testing/tldr-preference-trl-style`, with 92.9K train samples and 83.8K validation samples.

Training details. Our implementation of χ PO is built upon the DPO trainer from Transformer Reinforcement Learning (TRL) (von Werra et al., 2020). χ PO comes with strong robustness and theoretical properties, but the policy ratios can sometimes challenge stability in training. In practice, we have observed that better stability and performance can be achieved by utilizing the (more general form) link function $\tilde{\phi}(z) := \exp(\text{clip}_{[-88, 20]}(\alpha \cdot z)) + \gamma \cdot \log z$ in Algorithm 1, and performing a small grid search over additional parameters $\alpha = \{\frac{1}{4}, 1\}$ and $\gamma = \{0.1, 1\}$ for a fixed β . We briefly discuss each parameter in turn. The mixing parameter γ controls the relative ratios of KL- and χ^2 -regularization, our analysis in Appendix H.1 shows that Theorem 3.1 holds more generally for $\gamma \in (0, 1]$ (see Theorem H.1). Next, ignoring the clip for a second, $\alpha \in (0, 1]$ in $\tilde{\phi}$ implements regularization with the $(1 + \alpha)$ -divergence (or Renyi divergence), which is an f -divergence that is stronger than KL-regularization but weaker than χ^2 -regularization (Van Erven and Harremos, 2014), and also carries single-policy concentrability guarantees (although with a slower-rate dependence on sample size n). For example $\alpha = \frac{1}{4}$ corresponds to a link function with $(\frac{\pi}{\pi_{\text{ref}}})^{1/4}$, which easier to optimize than $\frac{\pi}{\pi_{\text{ref}}}$, which is recovered by $\alpha = 1$. Lastly, we provide some additional explanation for the clip portion of the link function. We observed that `torch.exp` is prone to underflow when $\log \frac{\pi}{\pi_{\text{ref}}}$ is very negative, and clipping the upper range to 20 can help reduce numerical instabilities. Clipping in such a manner is supported by our analysis in Proposition 4.1, which shows that $\frac{\pi^*}{\pi_{\text{ref}}} \leq 1 + \frac{R_{\text{max}}}{\beta}$ (though technically we do not know what R_{max} is, we can guess a reasonable upper bound). The parameters for all experiments are displayed in Table 2.

Generation details. For winrate evaluation, we use greedy decoding (temperature = 0) for all algorithms. For computation of the KL divergence, we sample from the model with temperature 1. The maximum prompt length is 512, and the maximum response length is 200. We use the standard generation prompt "TL;DR:" (Gao et al., 2024).

Evaluation of performance. The performance of each algorithm is measured against via winrate against responses in the SFT dataset, as measured by GPT-4o (global standard). The winrate is computed on a subset of 512 prompts from the SFT validation set (`trl-internal-testing/tldr-preference-sft-trl-style`), and the order of the model and reference responses are randomized each round.

Table 2: Parameter settings in TL;DR summarization

| Algorithm | Parameters |
|------------------------------|---|
| DPO | batch size: 64 learning rate: 1e-6 scheduler: cosine optimizer: adamw |
| χ PO | batch size: 64 clip range: [-88, 20] learning rate: 1e-6 scheduler: cosine optimizer: adamw |
| $\beta = 0.05$, 1 epoch | $\alpha : 1.25, \gamma : 1.0$ |
| $\beta = 0.05$, 2 epochs | $\alpha : 2.00, \gamma : 1.0$ |
| $\beta = 0.05$, 4 epochs | $\alpha : 1.25, \gamma : 0.1$ |
| $\beta = 0.005$, all epochs | $\alpha : 1.25, \gamma : 0.1$ |

Part II

Proofs

F PRELIMINARIES

Recall that for a pair of probability measures \mathbb{P} and \mathbb{Q} with a common dominating measure ω , Hellinger distance is defined via

$$D_{\text{H}}^2(\mathbb{P}, \mathbb{Q}) = \int \left(\sqrt{\frac{d\mathbb{P}}{d\omega}} - \sqrt{\frac{d\mathbb{Q}}{d\omega}} \right)^2 d\omega. \quad (35)$$

Lemma F.1 (MLE for conditional density estimation (e.g., Wong and Shen (1995); de Geer (2000); Zhang (2006); Agarwal et al. (2020))). *Consider a conditional density $p^* : \mathcal{X} \rightarrow \Delta(\mathcal{Y})$, where \mathcal{X} is the instance space and \mathcal{Y} is the target space. Let $\mathcal{D} = \{(x^i, y^i)\}_{i=1}^n$ be a dataset in which (x^i, y^i) are drawn i.i.d. as $x^i \sim \rho \in \Delta(\mathcal{X})$ and $y^i \sim p^*(y | x)$. Suppose we have a finite function class \mathcal{P} such that $p^* \in \mathcal{P}$, where $p(\cdot | x) \in \Delta(\mathcal{Y})$ for all $p \in \mathcal{P}$ and $x \in \mathcal{X}$. Define the maximum likelihood estimator*

$$\hat{p} := \operatorname{argmax}_{p \in \mathcal{P}} \sum_{(x, y) \in \mathcal{D}} \log p(y | x).$$

Then with probability at least $1 - \delta$,

$$\mathbb{E}_{x \sim \rho} [D_{\text{H}}^2(\hat{p}(\cdot | x), p^*(\cdot | x))] \leq \frac{2 \log(|\mathcal{P}| \delta^{-1})}{n}.$$

G ANALYSIS OF χ PO: PROOF SKETCH FOR THEOREM 3.1

In this section, we sketch the proof of the main guarantee for χ PO, [Theorem 3.1](#), with the full proof deferred to [Appendix H](#). A central object in the proof is the *implicit* reward model induced by the χ PO policy $\hat{\pi}$, which we define via

$$\hat{r}(x, a) := \beta \phi \left(\frac{\hat{\pi}(a | x)}{\pi_{\text{ref}}(a | x)} \right). \quad (36)$$

As we will show, this reward model is a natural bridge between χ PO and the corresponding mixed χ^2 -regularized RLHF objective in [Section 3.1](#), and allows us to view χ PO from a reward-based perspective. In particular, note that if we analogously define an induced reward model class $\mathcal{R}_{\Pi} := \{r(x, a) = \beta \phi \left(\frac{\pi(a | x)}{\pi_{\text{ref}}(a | x)} \right) : \pi \in \Pi\}$, then [Line 2](#) of χ PO can be viewed as performing maximum

likelihood estimation over this class (in the sense of Eq. (3)) under the Bradley-Terry model. Under Assumption 3.1, \mathcal{R}_Π realizes the true reward function r up to an action-independent shift. As a result, if we define $\Delta^r(x, a, b) := r(x, a) - r(x, b)$, then using a fairly standard generalization bound for maximum likelihood estimation (e.g., Wong and Shen (1995); Zhang (2006); de Geer (2000); see Lemma H.1), we can show that

$$\varepsilon_{\text{stat}}^2 := \mathbb{E}_{x \sim \rho, a \sim \pi_{\text{ref}}, b \sim \pi_{\text{ref}}} \left[\left| \Delta^{\hat{r}}(x, a, b) - \Delta^{r^*}(x, a, b) \right|^2 \right] \leq O \left(V_{\text{max}} e^{2R_{\text{max}}} \cdot \frac{\log(|\Pi|/\delta)}{n} \right). \quad (37)$$

In other words, the estimated reward model \hat{r} is accurate under the action distribution induced by π_{ref} . However, \hat{r} may still be inaccurate for policies that select different actions from π_{ref} , raising concerns of overoptimization. To address this issue, we use the following lemma, which shows that χ^2 -divergence bounds the extent to which the accuracy of a reward model \hat{r} trained under π_{ref} will transfer to a downstream policy π of interest; this will motivate our use of χ^2 -regularization.

Lemma G.1 (Informal version of Lemma H.3). *For any policy $\pi : \mathcal{X} \rightarrow \Delta(\mathcal{A})$, it holds that*

$$\mathbb{E}_{x \sim \rho, a \sim \pi(\cdot|x), b \sim \pi_{\text{ref}}(\cdot|x)} \left[\left| \Delta^{\hat{r}}(x, a, b) - \Delta^{r^*}(x, a, b) \right| \right] \lesssim \sqrt{(1 + D_{\chi^2}(\pi \parallel \pi_{\text{ref}})) \cdot \varepsilon_{\text{stat}}^2}.$$

Going forward, let us abbreviate $\mathbb{E}_{\pi, \pi_{\text{ref}}}[\cdot] = \mathbb{E}_{x \sim \rho, a \sim \pi(\cdot|x), b \sim \pi_{\text{ref}}(\cdot|x)}[\cdot]$. Let π^* be an arbitrary policy. Noting that $\mathcal{C}^\pi = 1 + 2D_{\chi^2}(\pi \parallel \pi_{\text{ref}})$ and that

$$J(\pi^*) - J(\hat{\pi}) \lesssim \mathbb{E}_{\pi^*, \pi_{\text{ref}}} \left[\left| \Delta^{\hat{r}}(x, a, b) - \Delta^{r^*}(x, a, b) \right| \right] + \mathbb{E}_{\hat{\pi}, \pi_{\text{ref}}} \left[\left| \Delta^{\hat{r}}(x, a, b) - \Delta^{r^*}(x, a, b) \right| \right],$$

it follows immediately from Lemma G.1 that χ PO obtains a crude guarantee scaling with all-policy concentrability, i.e. $J(\pi^*) - J(\hat{\pi}) \lesssim \sqrt{(\mathcal{C}^{\pi^*} + \mathcal{C}^{\hat{\pi}}) \varepsilon_{\text{stat}}^2} \leq \sqrt{(\mathcal{C}^{\pi^*} + \max_{\pi \in \Pi} \mathcal{C}^\pi) \varepsilon_{\text{stat}}^2}$. This inequality is tight for non-pessimistic algorithms like DPO, which reflects their sensitivity to overoptimization. To obtain the improved guarantee for χ PO in Theorem 3.1, which scales only with single-policy concentrability \mathcal{C}^{π^*} , the crux of the remaining proof will be to show that χ PO implicitly implements pessimism via mixed χ^2 -regularization. For this, we appeal to the following central technical lemma, which we expect to find broader use.

Lemma G.2 (Informal version of Lemma H.2). *Let f be a convex function with $\text{dom}(f) = \mathbb{R}_+$ that is differentiable over its domain. Given any parameter $\beta > 0$ and policy $\bar{\pi} : \mathcal{X} \rightarrow \Delta(\mathcal{A})$ with $\bar{\pi}(a | x) \in \text{dom}(f')$ for all x, a , define the reward model $\bar{r}(x, a) = \beta f' \left(\frac{\bar{\pi}(a|x)}{\pi_{\text{ref}}(a|x)} \right)$. Then*

$$\bar{\pi} \in \underset{\pi}{\text{argmax}} \mathbb{E}_\pi[\bar{r}(x, a)] - \beta \cdot D_f(\pi \parallel \pi_{\text{ref}}).$$

Under Assumption 3.2 we have $\hat{\pi} \in \text{dom}(f'_{\chi_{\text{mix}}})$. Then recalling that $\hat{r}(x, a) := \beta \phi \left(\frac{\hat{\pi}(a|x)}{\pi_{\text{ref}}(a|x)} \right) = \beta f'_{\chi_{\text{mix}}} \left(\frac{\hat{\pi}(a|x)}{\pi_{\text{ref}}(a|x)} \right)$ and that $f_{\chi_{\text{mix}}}$ is convex, Lemma G.2 implies that the policy $\hat{\pi}$ produced by χ PO satisfies

$$\hat{\pi} \in \underset{\pi \in \Pi}{\text{argmax}} J_{\beta, \hat{r}}^{\chi_{\text{mix}}}(\pi) := \mathbb{E}_\pi[\hat{r}] - \beta D_{\chi^2}(\pi \parallel \pi_{\text{ref}}) - \beta D_{\text{KL}}(\pi \parallel \pi_{\text{ref}}). \quad (38)$$

In other words,

The χ PO policy $\hat{\pi}$ optimizes the mixed χ^2 -regularized RLHF objective under its own implicit reward model.

This formally justifies the claim that χ PO implicitly implements pessimism via χ^2 -regularization. With this result in hand, we are now ready to prove Theorem 3.1. Let π^* be an arbitrary policy. Since $J_{\beta, \hat{r}}^{\chi_{\text{mix}}}(\hat{\pi}) \geq J_{\beta, \hat{r}}^{\chi_{\text{mix}}}(\pi^*)$ by Eq. (38), we can decompose the regret $J(\pi^*) - J(\hat{\pi})$ as

$$\begin{aligned} J(\pi^*) - J(\hat{\pi}) &\leq J(\pi^*) - J_{\beta, \hat{r}}^{\chi_{\text{mix}}}(\pi^*) + J_{\beta, \hat{r}}^{\chi_{\text{mix}}}(\hat{\pi}) - J(\hat{\pi}) \\ &= \underbrace{J(\pi^*) - J(\pi_{\text{ref}}) - J_{\beta, \hat{r}}^{\chi_{\text{mix}}}(\pi^*) + J_{\beta, \hat{r}}^{\chi_{\text{mix}}}(\pi_{\text{ref}})}_{\text{(I)}} + \underbrace{J_{\beta, \hat{r}}^{\chi_{\text{mix}}}(\hat{\pi}) - J_{\beta, \hat{r}}^{\chi_{\text{mix}}}(\pi_{\text{ref}}) - J(\hat{\pi}) + J(\pi_{\text{ref}})}_{\text{(II)}}. \end{aligned}$$

In the second line, we have added or subtracted the baselines $J(\pi_{\text{ref}})$ and $J_{\beta, \hat{r}}^{\chi_{\text{mix}}}(\pi_{\text{ref}})$ to center the objectives with the performance of the reference policy. Up to statistical errors, the first term

(I) corresponds to error from how much $J_{\beta, \hat{r}}^{\chi^2 \text{mix}}(\pi^*)$ *underestimates* the return of π^* (bias), and the second term (II) corresponds to error from how much $J_{\beta, \hat{r}}^{\chi^2 \text{mix}}(\hat{\pi})$ *overestimates* the return of $\hat{\pi}$ (overoptimization). As we will see shortly, these two sources of error are directly controlled (in opposing ways) by the strength of the regularization parameter β in Eq. (38).

First, expanding the definition of $J_{\beta, \hat{r}}^{\chi^2 \text{mix}}(\pi^*)$ and centering the returns using the reference policies, we have

$$\begin{aligned} \text{(I)} &= J(\pi^*) - J_{\beta, \hat{r}}^{\chi^2 \text{mix}}(\pi^*) - J(\pi_{\text{ref}}) + J_{\beta, \hat{r}}^{\chi^2 \text{mix}}(\pi_{\text{ref}}) \\ &= \mathbb{E}_{\pi^*}[r^*(x, a)] - \mathbb{E}_{\pi^*}[\hat{r}(x, a)] + \beta D_{\chi^2}(\pi^* \parallel \pi_{\text{ref}}) + \beta D_{\text{KL}}(\pi^* \parallel \pi_{\text{ref}}) - \mathbb{E}_{\hat{\pi}}[r^*(x, a)] + \mathbb{E}_{\pi_{\text{ref}}}[\hat{r}(x, a)] \\ &= \mathbb{E}_{\pi^*, \pi_{\text{ref}}}[\Delta^{r^*}(x, a, b) - \Delta^{\hat{r}}(x, a, b)] + \beta D_{\chi^2}(\pi^* \parallel \pi_{\text{ref}}) + \beta D_{\text{KL}}(\pi^* \parallel \pi_{\text{ref}}) \\ &\leq \sqrt{(1 + D_{\chi^2}(\pi^* \parallel \pi_{\text{ref}})) \cdot \varepsilon_{\text{stat}}^2} + \underbrace{\beta \cdot D_{\chi^2}(\pi^* \parallel \pi_{\text{ref}})}_{\text{bias}}. \end{aligned}$$

Above, we have used that $D_{\text{KL}}(\pi \parallel \pi_{\text{ref}}) \leq D_{\chi^2}(\pi \parallel \pi_{\text{ref}})$ for any policy π , along with the bound on reward estimation error from Lemma G.1. Next, expanding $J_{\beta, \hat{r}}^{\chi^2 \text{mix}}(\hat{\pi})$ and centering the returns in a similar fashion,

$$\begin{aligned} \text{(II)} &= J_{\beta, \hat{r}}^{\chi^2 \text{mix}}(\hat{\pi}) - J(\hat{\pi}) - J_{\beta, \hat{r}}^{\chi^2 \text{mix}}(\pi_{\text{ref}}) + J(\pi_{\text{ref}}) \\ &= \mathbb{E}_{\hat{\pi}, \pi_{\text{ref}}}[\Delta^{\hat{r}}(x, a, b) - \Delta^{r^*}(x, a, b)] - \beta D_{\chi^2}(\hat{\pi} \parallel \pi_{\text{ref}}) - \beta D_{\text{KL}}(\hat{\pi} \parallel \pi_{\text{ref}}) \\ &\leq \sqrt{(1 + D_{\chi^2}(\hat{\pi} \parallel \pi_{\text{ref}})) \cdot \varepsilon_{\text{stat}}^2} - \beta \cdot D_{\chi^2}(\hat{\pi} \parallel \pi_{\text{ref}}) \\ &\lesssim \varepsilon_{\text{stat}} + \underbrace{\beta^{-1} \varepsilon_{\text{stat}}^2}_{\text{overoptimization error}}. \end{aligned}$$

Above, the first inequality uses $D_{\text{KL}}(\pi \parallel \pi_{\text{ref}}) \geq 0$ and Lemma G.1, while the second inequality uses AM-GM. Critically, by using χ^2 -regularization, we are able to cancel the on-policy error term $\sqrt{(1 + D_{\chi^2}(\hat{\pi} \parallel \pi_{\text{ref}})) \cdot \varepsilon_{\text{stat}}^2}$ that arises from change-of-measure, leading to a modest $\beta^{-1} \varepsilon_{\text{stat}}^2$ penalty for overoptimization.

Combining these results, and recalling that $\mathcal{C}^\pi = 1 + 2D_{\chi^2}(\pi \parallel \pi_{\text{ref}})$, we conclude that

$$J(\pi^*) - J(\hat{\pi}) \lesssim \sqrt{\mathcal{C}^{\pi^*} \cdot \varepsilon_{\text{stat}}^2} + \underbrace{\beta \cdot \mathcal{C}^{\pi^*}}_{\text{bias}} + \underbrace{\beta^{-1} \cdot \varepsilon_{\text{stat}}^2}_{\text{overoptimization error}}.$$

The bias and overoptimization errors above arise from how well our chosen uncertainty quantifier, $\beta D_{\chi^2}(\pi \parallel \pi_{\text{ref}})$, accounts for the on-policy statistical error $\sqrt{(1 + D_{\chi^2}(\pi \parallel \pi_{\text{ref}})) \cdot \varepsilon_{\text{stat}}^2}$ arising from Lemma G.1; this is controlled by the magnitude of the regularization parameter β . When β is too large, the uncertainty quantifier is overly pessimistic about the quality of the reward model \hat{r} under π^* , which increases the *bias* of χ^2 PO. In contrast, the *overoptimization error* increases when β is too small. In this regime, $\hat{\pi}$ overfits to \hat{r} because the regularizer under-evaluates the statistical error of the learned policy. In order to obtain tight statistical rates, the choice of regularization parameter β must carefully balance its opposing effects on bias and overoptimization error. For a fixed π^* , choosing $\beta \propto (\varepsilon_{\text{stat}}^2 / \mathcal{C}^{\pi^*})^{1/2}$ results in the second claim in Theorem 3.1.

H PROOFS FOR SECTION 3

This section is organized as follows. First, in Appendix H.1, we analyze a more general version of χ^2 PO that mixes KL-regularization with χ^2 -regularization using a mixing parameter $\gamma \in (0, 1]$, and present its sample complexity guarantee in Theorem H.1. χ^2 PO is a special case with $\gamma = 1$, and Appendix H.2 shows (with a one-line proof) that Theorem 3.1 follows directly from Theorem H.1 with this parameter choice.

H.1 GENERAL VERSION OF THEOREM 3.1

As previously described at the end of Section 3.3, χ^2 PO can be applied in a more general form where the KL-regularization is mixed with χ^2 -regularization using a weight parameter $\gamma \in (0, 1]$. In this

section, we analyze the sample complexity for this form of the algorithm, of which χ PO is a special case with $\gamma = 1$, which directly leads to the guarantee in [Theorem 3.1](#).

Concretely, given regularization parameter $\beta > 0$ and weight parameter $\gamma \in (0, 1]$, we aim to solve the mixed χ^2 -regularized objective

$$\operatorname{argmax}_{\pi: \mathcal{X} \rightarrow \Delta(\mathcal{A})} J_{\beta, \gamma}^{\chi^{\text{mix}}}(\pi) := \mathbb{E}_{\pi}[r^*(x, a)] - \beta \cdot D_{\chi^2}(\pi \parallel \pi_{\text{ref}}) - \beta\gamma \cdot D_{\text{KL}}(\pi \parallel \pi_{\text{ref}}). \quad (39)$$

The regularization term $D_{\chi^2}(\pi \parallel \pi_{\text{ref}}) + \gamma \cdot D_{\text{KL}}(\pi \parallel \pi_{\text{ref}}) = D_{f_{\chi^{\text{mix}}, \gamma}}(\pi \parallel \pi_{\text{ref}})$ is an f -divergence induced by the function $f_{\chi^{\text{mix}}, \gamma}(z) := \frac{1}{2}(z-1)^2 + \gamma z \log z$. Correspondingly, we replace the link function $\phi(\cdot)$ in χ PO with

$$\phi_{\gamma}(z) := z + \gamma \log(z),$$

and output the policy

$$\hat{\pi} \leftarrow \operatorname{argmax}_{\pi \in \Pi} \sum_{(x, a_+, a_-) \in \mathcal{D}_{\text{pref}}} \log \left[\sigma \left(\operatorname{clip}_{2R_{\max}} \left[\beta \phi_{\gamma} \left(\frac{\pi(a_+ | x)}{\pi_{\text{ref}}(a_+ | x)} \right) - \beta \phi_{\gamma} \left(\frac{\pi(a_- | x)}{\pi_{\text{ref}}(a_- | x)} \right) \right] \right) \right]. \quad (40)$$

To give a sample complexity guarantee for [Eq. \(40\)](#), we require that Π can express the optimal regularized policy for the objective $J_{\beta, \gamma}^{\chi^{\text{mix}}}$ in [Eq. \(39\)](#). This generalizes [Assumption 3.1](#) for χ PO, which corresponds to the special case where $\gamma = 1$.

Assumption H.1 (Policy realizability). *The policy class Π satisfies $\pi_{\beta, \gamma}^* \in \Pi$, where $\pi_{\beta, \gamma}^*$ is the optimal policy under mixed χ^2 -regularization ([Eq. \(11\)](#)).*

We also assert that, analogous to [Assumption 3.2](#), the ‘‘implicit’’ reward models induced by the policy class Π and the link function ϕ_{γ} have bounded range.

Assumption H.2 (Bounded implicit rewards). *For a parameter $V_{\max} \geq R_{\max}$, it holds that for all $\pi \in \Pi$, $x \in \mathcal{X}$, and $a, b \in \mathcal{A}$,*

$$\left| \beta \phi_{\gamma} \left(\frac{\pi(a | x)}{\pi_{\text{ref}}(a | x)} \right) - \beta \phi_{\gamma} \left(\frac{\pi(b | x)}{\pi_{\text{ref}}(b | x)} \right) \right| \leq V_{\max}. \quad (41)$$

We now state the sample complexity guarantee for the policy learned in [Eq. \(40\)](#). The first bound applies to general $\beta > 0$ and $\gamma \in (0, 1]$, while in the second we obtain a tight statistical rate by choosing the parameter β as a function of the comparator policy π^* .

Theorem H.1 (General version of [Theorem 3.1](#)). *Suppose [Assumptions H.1](#) and [H.2](#) hold for some $\beta > 0$ and $\gamma \in (0, 1]$. With probability at least $1 - \delta$, the variant of χ PO in [Eq. \(40\)](#) produces a policy $\hat{\pi}$ such that for all policies π^* simultaneously, we have*

$$J(\pi^*) - J(\hat{\pi}) \leq 32V_{\max}e^{2R_{\max}} \cdot \sqrt{\frac{2C^{\pi^*} \log(|\Pi|/\delta)}{n}} + \beta(1 + \gamma) \cdot \frac{C^{\pi^*}}{2} + \beta^{-1} \cdot \frac{256V_{\max}^2 e^{4R_{\max}} \log(|\Pi|/\delta)}{n}.$$

In particular, given any comparator policy π^ , we can choose $\beta = 32V_{\max}e^{2R_{\max}} \sqrt{\frac{2 \log(|\Pi|/\delta)}{nC^{\pi^*}}}$ to achieve*

$$J(\pi^*) - J(\hat{\pi}) \leq (64 + 4\gamma)V_{\max}e^{2R_{\max}} \cdot \sqrt{\frac{C^{\pi^*} \log(|\Pi|/\delta)}{n}}.$$

The bias-overoptimization tradeoffs induced by the choice of β in [Theorem H.1](#) are identical to those for [Theorem 3.1](#) (and described there). Let us briefly discuss the influence of γ on the sample complexity. We first observe that choice of $\gamma \in (0, 1]$ changes the bound by only a small multiplicative factor, which implies that γ can be arbitrarily small as long as it is positive. For the analysis, this is natural because the KL-divergence is dominated by the χ^2 -divergence, and, as discussed in [Section 3.2](#), KL-regularization is only needed to enable the DPO-style reparameterization trick for [Eq. \(40\)](#) (in particular, the χ^2 -RLHF algorithm in [Appendix C](#), which does not require reparameterization, obtains similar guarantees using pure χ^2 -regularization). It is worth noting, however, that the γ parameter

can implicitly influence the magnitude of V_{\max} , as well as the policy realizability condition. As such, practical consequences of this hyperparameter choice may not be fully captured by [Theorem H.1](#).

Proof of [Theorem H.1](#). Recall that the link function ϕ_γ induces a correspondence between policies in the class Π and the implicit reward functions they induce (or, equivalently, between policies and the Bradley-Terry preference models they express). Our proof centers around the implicit reward model induced by the learned policy $\hat{\pi}$,

$$\hat{r}(x, a) := \beta \cdot \phi_\gamma \left(\frac{\hat{\pi}(a | x)}{\pi_{\text{ref}}(a | x)} \right),$$

which will allow us to move between the χ PO objective ([Eq. \(40\)](#)) and the RLHF objective ([Eq. \(39\)](#)). In particular, we establish two key facts, which together show that [Eq. \(40\)](#) implicitly solves [Eq. \(39\)](#):

1. ([Lemma H.3](#)) The reward model \hat{r} is an accurate estimate of r^* on the distribution of π_{ref} . Moreover, we can transfer this guarantee to the distribution of any policy π by paying a multiplicative $(1 + 2D_{\chi^2}(\pi \| \pi_{\text{ref}}))$ -factor.
2. ([Lemma H.2](#)) $\hat{\pi}$ maximizes the RLHF objective in [Eq. \(39\)](#) with reward model \hat{r} , namely,

$$\hat{\pi} = \operatorname{argmax}_{\pi \in \Pi} \mathbb{E}_\pi[\hat{r}(x, a)] - \beta \cdot D_{\chi^2}(\pi \| \pi_{\text{ref}}) - \beta\gamma \cdot D_{\text{KL}}(\pi \| \pi_{\text{ref}}). \quad (42)$$

Establishing these relationships enables us to analyze the χ PO policy $\hat{\pi}$ defined in [Eq. \(40\)](#) through the RLHF formulation in [Eq. \(42\)](#), allowing us to appeal to pessimism-based arguments to show that χ PO is insensitive to overoptimization error that might otherwise be encountered when learning a policy from off-policy data.

Implicit reward model \hat{r} . The χ PO objective in [Eq. \(40\)](#) is equivalent to maximum likelihood estimation with the Bradley-Terry preference model over the induced reward function class

$$\mathcal{R}_\Pi := \left\{ r(x, a) = \beta \cdot \phi_\gamma \left(\frac{\pi(a | x)}{\pi_{\text{ref}}(a | x)} \right) : \pi \in \Pi \right\}.$$

Then, since $\hat{\pi}$ is the maximizer in [Eq. \(40\)](#), we can equivalently write

$$\hat{r} = \operatorname{argmax}_{r \in \mathcal{R}_\Pi} \sum_{(x, a_+, a_-) \in \mathcal{D}_{\text{pref}}} \log \sigma(\text{clip}_{2R_{\max}}[r(a_+ | x) - r(a_- | x)]). \quad (43)$$

The following lemma, which builds on a standard MLE generalization bound ([Lemma F.1](#)) bounds the error of \hat{r} under the action distribution induced by π_{ref} . Recall that we use $\mathbb{E}_{\pi, \pi'}[\cdot]$ as shorthand for $\mathbb{E}_{x \sim \rho, a \sim \pi(\cdot | x), b \sim \pi'(\cdot | x)}[\cdot]$.

Lemma H.1. *Suppose [Assumption H.1](#) holds. Then with probability at least $1 - \delta$, the policy $\hat{\pi}$ output by [Eq. \(40\)](#) satisfies*

$$\varepsilon_{\text{Stat}}^2 =: \mathbb{E}_{\pi_{\text{ref}}, \pi_{\text{ref}}} \left[\left(\text{clip}_{2R_{\max}}[\hat{r}(x, a) - \hat{r}(x, b)] - \text{clip}_{2R_{\max}}[r^*(x, a) - r^*(x, b)] \right)^2 \right] \leq \frac{128R_{\max}^2 e^{4R_{\max}} \log(|\Pi|/\delta)}{n}.$$

[Lemma H.1](#), along with all further supporting lemmas, is proven in the sequel. This result measures the error of \hat{r} using the clipped differences of rewards for pairs of actions (x, a, b) drawn from π_{ref} . Clipping the range of the implicit/explicit reward functions to $2R_{\max}$ ensures that the statistical error does not depend on V_{\max} . One minor but important detail in the proof is showing that [Assumption H.1](#) implies \mathcal{R}_Π includes the true reward function r^* up to an action-independent shift, so that the true preference model is realizable.

Implicit RLHF policy optimization. Having established the accuracy of \hat{r} , we now show that [Eq. \(40\)](#) finds the optimal policy to the RLHF objective in [Eq. \(42\)](#) when \hat{r} is used as the reward model, i.e.,

$$\hat{\pi} = \operatorname{argmax}_{\pi \in \Pi} J_{\beta, \gamma, \hat{r}}^{\chi_{\text{mix}}}(\pi) := \mathbb{E}_\pi[\hat{r}(x, a)] - \beta \cdot D_{\chi^2}(\pi \| \pi_{\text{ref}}) - \beta\gamma \cdot D_{\text{KL}}(\pi \| \pi_{\text{ref}}). \quad (44)$$

This is a direct consequence of the result in [Lemma H.2](#), which shows that an analogous property holds for general f -divergences. In particular, for any convex function f and policy π , the policy π is itself the optimal solution to the f -divergence-regularized RLHF objective under the implicit reward model induced by π with the link function f' .

Lemma H.2. Let $f : (0, \infty) \rightarrow \mathbb{R}$ be a convex function with $f(1) = 0$. Further, f is differentiable almost everywhere and $0 \notin \text{dom}(f')$, where we define $f'(0) := \lim_{x \downarrow 0} \frac{f(x) - f(0)}{x}$ and $f(0) := \lim_{x \downarrow 0} f(x)$. Given any parameter $\beta > 0$ and valid policy $\bar{\pi} : \mathcal{X} \rightarrow \Delta(\mathcal{A})$, with $\bar{\pi}(a | x) \in \text{dom}(f')$ for all (x, a) , let $\bar{r}(x, a) = \beta f' \left(\frac{\bar{\pi}(a|x)}{\bar{\pi}_{\text{ref}}(a|x)} \right)$ be the implicit reward model. Then

$$\bar{\pi} \in \operatorname{argmax}_{\pi: \mathcal{X} \rightarrow \Delta(\mathcal{A})} \mathbb{E}_{\pi}[\bar{r}(x, a)] - \beta D_f(\pi \| \pi_{\text{ref}}).$$

Since $f'_{\chi^{\text{mix}, \gamma}} = \phi_{\gamma} = x + \gamma \log x$ for $\gamma > 0$, clearly $0 \notin \text{dom}(\phi_{\gamma})$. Further, under [Assumption H.2](#), $\pi(a | x) > 0$ for all $\pi \in \Pi$ (otherwise V_{max} would be undefined), thus $\pi(a | x) \in \text{dom}(\phi_{\gamma})$ for all (x, a) . The claim in [Eq. \(44\)](#) then directly follows.

Estimation error translation. To proceed, we will use condition on [Lemma H.1](#) and use the event in this lemma to relate the estimated RLHF objective in [Eq. \(42\)](#) to the “true” RLHF objective that replaces \hat{r} with r^* . An immediate challenge is that the RLHF objective in [Eq. \(42\)](#) must evaluate $\mathbb{E}_{\pi}[\hat{r}(x, a)]$ for all $\pi \in \Pi$, and accuracy under π_{ref} does not immediately imply that \hat{r} is accurate for other policies. The following bound quantifies the effects of this distribution shift using the χ^2 -divergence, and expresses how the estimation guarantee for \hat{r} in [Lemma H.1](#) transfers to other policies π of interest.

Lemma H.3. Suppose [Assumption 3.1](#) holds. Then for any $\pi : \mathcal{X} \rightarrow \Delta(\mathcal{A})$, under the event in [Lemma H.1](#), we have

$$\mathbb{E}_{\pi, \pi_{\text{ref}}} [|\hat{r}(x, a) - \hat{r}(x, b) - (r^*(x, a) - r^*(x, b))|] \leq \frac{2V_{\text{max}}}{R_{\text{max}}} \cdot \sqrt{(1 + 2D_{\chi^2}(\pi \| \pi_{\text{ref}})) \cdot \varepsilon_{\text{stat}}^2},$$

where $\varepsilon_{\text{stat}}^2$ is the off-policy estimation error defined in [Lemma H.1](#).

It is worth noting that [Lemma H.3](#) bounds the *unclipped* on-policy estimation error (on the LHS) in terms of the *clipped* off-policy estimation error, and in making this translation we pay for V_{max} . As we will see shortly, working with the unclipped \hat{r} object is necessary for showing that [Eq. \(40\)](#) implicitly optimizes [Eq. \(42\)](#).

Pessimism-based regret decomposition. Equipped with the preceding lemmas, we can now bound the regret for χ PO. We decompose the regret using the RLHF objective $J_{\beta, \gamma, \hat{r}}^{\chi^{\text{mix}}}(\pi^*)$ defined in [Eq. \(44\)](#). Fixing an arbitrary comparator policy π^* , we have

$$\begin{aligned} J(\pi^*) - J(\hat{\pi}) &= \mathbb{E}_{\pi^*}[r^*(x, a)] - \mathbb{E}_{\hat{\pi}}[r^*(x, a)] \\ &= \mathbb{E}_{\pi^*}[r^*(x, a)] - J_{\beta, \gamma, \hat{r}}^{\chi^{\text{mix}}}(\pi^*) + J_{\beta, \gamma, \hat{r}}^{\chi^{\text{mix}}}(\pi^*) - \mathbb{E}_{\hat{\pi}}[r^*(x, a)] \\ &\leq \mathbb{E}_{\pi^*}[r^*(x, a)] - J_{\beta, \gamma, \hat{r}}^{\chi^{\text{mix}}}(\pi^*) + J_{\beta, \gamma, \hat{r}}^{\chi^{\text{mix}}}(\hat{\pi}) - \mathbb{E}_{\hat{\pi}}[r^*(x, a)], \end{aligned}$$

where the last inequality uses the optimality of $\hat{\pi}$ for [Eq. \(44\)](#).

Expanding the expression for $J_{\beta, \gamma, \hat{r}}^{\chi^{\text{mix}}}$, we can further bound this by

$$\begin{aligned} J(\pi^*) - J(\hat{\pi}) &\leq \mathbb{E}_{\pi^*}[r^*(x, a) - \hat{r}(x, a)] + \beta D_{\chi^2}(\pi^* \| \pi_{\text{ref}}) + \beta \gamma D_{\text{KL}}(\pi^* \| \pi_{\text{ref}}) \\ &\quad + \mathbb{E}_{\hat{\pi}}[\hat{r}(x, a) - r^*(x, a)] - \beta D_{\chi^2}(\hat{\pi} \| \pi_{\text{ref}}) - \beta \gamma D_{\text{KL}}(\hat{\pi} \| \pi_{\text{ref}}) \\ &\leq \mathbb{E}_{\pi^*}[r^*(x, a) - \hat{r}(x, a)] + \beta(1 + \gamma) D_{\chi^2}(\pi^* \| \pi_{\text{ref}}) \\ &\quad + \mathbb{E}_{\hat{\pi}}[\hat{r}(x, a) - r^*(x, a)] - \beta D_{\chi^2}(\hat{\pi} \| \pi_{\text{ref}}). \end{aligned} \tag{45}$$

In the last line, we use the fact that $0 \leq D_{\text{KL}}(\pi \| \pi_{\text{ref}}) \leq D_{\chi^2}(\pi \| \pi_{\text{ref}})$ for any policy π to consolidate the f -divergence terms. Specifically, this allows us to eliminate $D_{\text{KL}}(\hat{\pi} \| \pi_{\text{ref}})$, and combine $D_{\text{KL}}(\pi^* \| \pi_{\text{ref}})$ and $D_{\chi^2}(\pi^* \| \pi_{\text{ref}})$.

In order to bound the reward estimation error terms in [Eq. \(45\)](#) using the guarantee we have previously established ([Lemma H.3](#)), we first center them using the return under the reference policy:

$$\begin{aligned} &\mathbb{E}_{\pi^*}[r^*(x, a) - \hat{r}(x, a)] + \mathbb{E}_{\hat{\pi}}[\hat{r}(x, a) - r^*(x, a)] \\ &= \mathbb{E}_{\pi^*, \pi_{\text{ref}}}[r^*(x, a) - \hat{r}(x, a) - r^*(x, b) + \hat{r}(x, b)] + \mathbb{E}_{\hat{\pi}, \pi_{\text{ref}}}[\hat{r}(x, a) - r^*(x, a) - \hat{r}(x, b) + r^*(x, b)] \\ &= \mathbb{E}_{\pi^*, \pi_{\text{ref}}}[\Delta^*(x, a, b) - \hat{\Delta}(x, a, b)] + \mathbb{E}_{\hat{\pi}, \pi_{\text{ref}}}[\hat{\Delta}(x, a, b) - \Delta^*(x, a, b)], \end{aligned}$$

where $\Delta^*(x, a, b) := r^*(x, a) - r^*(x, b)$ and $\widehat{\Delta}(x, a, b) := \widehat{r}(x, a) - \widehat{r}(x, b)$. Substituting this identity back into the regret decomposition in Eq. (45), we apply Lemma H.3 with $\varepsilon_{\text{stat}}^2 := 128R_{\text{max}}^2 e^{4R_{\text{max}}} \frac{\log(|\Pi|/\delta)}{n}$ (from Lemma H.1) to obtain

$$\begin{aligned}
J(\pi^*) - J(\widehat{\pi}) &\leq \mathbb{E}_{\pi^*, \pi_{\text{ref}}} \left[\Delta^*(x, a, b) - \widehat{\Delta}(x, a, b) \right] + \beta(1 + \gamma) D_{\chi^2}(\pi^* \parallel \pi_{\text{ref}}) \\
&\quad + \mathbb{E}_{\widehat{\pi}, \pi_{\text{ref}}} \left[\widehat{\Delta}(x, a, b) - \Delta^*(x, a, b) \right] - \beta D_{\chi^2}(\widehat{\pi} \parallel \pi_{\text{ref}}) \\
&\leq \frac{2V_{\text{max}}}{R_{\text{max}}} \sqrt{(1 + 2D_{\chi^2}(\pi^* \parallel \pi_{\text{ref}})) \cdot \varepsilon_{\text{stat}}^2} + \beta(1 + \gamma) D_{\chi^2}(\pi^* \parallel \pi_{\text{ref}}) \\
&\quad + \frac{2V_{\text{max}}}{R_{\text{max}}} \sqrt{(1 + 2D_{\chi^2}(\widehat{\pi} \parallel \pi_{\text{ref}})) \cdot \varepsilon_{\text{stat}}^2} - \beta D_{\chi^2}(\widehat{\pi} \parallel \pi_{\text{ref}}) \\
&= \frac{2V_{\text{max}}}{R_{\text{max}}} \sqrt{\mathcal{C}^{\pi^*} \cdot \varepsilon_{\text{stat}}^2} + \frac{\beta(1 + \gamma)}{2} \cdot (\mathcal{C}^{\pi^*} - 1) + \frac{2V_{\text{max}}}{R_{\text{max}}} \sqrt{\mathcal{C}^{\widehat{\pi}} \cdot \varepsilon_{\text{stat}}^2} - \frac{\beta}{2} \cdot (\mathcal{C}^{\widehat{\pi}} - 1) \\
&\leq \frac{2V_{\text{max}}}{R_{\text{max}}} \sqrt{\mathcal{C}^{\pi^*} \cdot \varepsilon_{\text{stat}}^2} + \frac{\beta(1 + \gamma)}{2} \cdot \mathcal{C}^{\pi^*} + \frac{2V_{\text{max}}}{R_{\text{max}}} \sqrt{\mathcal{C}^{\widehat{\pi}} \cdot \varepsilon_{\text{stat}}^2} - \frac{\beta}{2} \cdot \mathcal{C}^{\widehat{\pi}},
\end{aligned}$$

since $\mathcal{C}^\pi = 1 + 2D_{\chi^2}(\pi \parallel \pi_{\text{ref}})$, or equivalently $D_{\chi^2}(\pi \parallel \pi_{\text{ref}}) = \frac{1}{2}(\mathcal{C}^\pi - 1)$. Lastly, we use the AM-GM inequality to upper bound

$$\frac{2V_{\text{max}}}{R_{\text{max}}} \sqrt{\mathcal{C}^{\widehat{\pi}} \cdot \varepsilon_{\text{stat}}^2} \leq \frac{2V_{\text{max}}^2 \varepsilon_{\text{stat}}^2}{R_{\text{max}}^2 \beta} + \frac{\beta \mathcal{C}^{\widehat{\pi}}}{2},$$

allowing us to conclude that

$$J(\pi^*) - J(\widehat{\pi}) \leq \frac{2V_{\text{max}}}{R_{\text{max}}} \sqrt{\mathcal{C}^{\pi^*} \cdot \varepsilon_{\text{stat}}^2} + \frac{\beta(1 + \gamma)}{2} \cdot \mathcal{C}^{\pi^*} + 2\beta^{-1} \cdot \frac{V_{\text{max}}^2 \varepsilon_{\text{stat}}^2}{R_{\text{max}}^2}.$$

Plugging in the expression for $\varepsilon_{\text{stat}}^2$ results in the first statement of Theorem H.1.

Choosing β for tight rates. For the second statement, given a comparator policy π^* , choosing $\beta = \frac{2V_{\text{max}}}{R_{\text{max}}} \sqrt{\frac{\varepsilon_{\text{stat}}^2}{\mathcal{C}^{\pi^*}}}$ gives

$$\begin{aligned}
J(\pi^*) - J(\widehat{\pi}) &\leq \frac{2V_{\text{max}}}{R_{\text{max}}} \sqrt{\mathcal{C}^{\pi^*} \cdot \varepsilon_{\text{stat}}^2} + (1 + \gamma) \frac{V_{\text{max}}}{R_{\text{max}}} \sqrt{\mathcal{C}^{\pi^*} \cdot \varepsilon_{\text{stat}}^2} + \frac{V_{\text{max}}}{R_{\text{max}}} \sqrt{\mathcal{C}^{\pi^*} \cdot \varepsilon_{\text{stat}}^2} \\
&= (4 + \gamma) \frac{V_{\text{max}}}{R_{\text{max}}} \sqrt{\mathcal{C}^{\pi^*} \cdot \varepsilon_{\text{stat}}^2}.
\end{aligned}$$

□

H.1.1 PROOFS FOR SUPPORTING LEMMAS

Proof of Lemma H.1. Recall the reward-based MLE objective in Eq. (43),

$$\widehat{r} = \operatorname{argmax}_{r \in \mathcal{R}_{\Pi}} \sum_{(x, a_+, a_-) \in \mathcal{D}_{\text{pref}}} \log \sigma(\operatorname{clip}_{2R_{\text{max}}}[r(x, a_+) - r(x, a_-)]).$$

To leverage standard generalization bounds for MLE, we re-interpret this objective as maximum likelihood over a class of preference distributions under the Bradley-Terry model. For a reward function r , define for all $y \in \{+1, -1\}$ and $(x, a, b) \in \mathcal{X} \times \mathcal{A} \times \mathcal{A}$ its induced preference distribution:

$$P_r(y|x, a, b) = \mathbb{I}\{y = +1\} \cdot \sigma(\operatorname{clip}_{2R_{\text{max}}}[r(x, a) - r(x, b)]) + \mathbb{I}\{y = -1\} \cdot \sigma(\operatorname{clip}_{2R_{\text{max}}}[r(x, b) - r(x, a)]).$$

Consider the a class of preference models induced by \mathcal{R}_{Π} under this definition, $\mathcal{P}_{\Pi} := \{P_r : r \in \mathcal{R}_{\Pi}\}$. We can equivalently write that

$$P_{\widehat{r}} = \operatorname{argmax}_{p \in \mathcal{P}_{\Pi}} \sum_{(x, a_+, a_-) \in \mathcal{D}_{\text{pref}}} \log p(+1 | x, a_+, a_-),$$

or, interpreting each tuple (x, a_+, a_-) in $\mathcal{D}_{\text{pref}}$ as being induced by a tuple (x, a, \tilde{a}, y) in which $(a_+, a_-) = (a, \tilde{a})$ if $y = +1$ and $(a_+, a_-) = (\tilde{a}, a)$ if $y = -1$,

$$P_{\widehat{r}} = \operatorname{argmax}_{p \in \mathcal{P}_{\Pi}} \sum_{(x, a, \tilde{a}, y) \in \mathcal{D}_{\text{pref}}} \log p(y | x, a, \tilde{a}).$$

Next, we show that $P_{r^*} \in \mathcal{P}_\Pi$, ie., the induced preference model class realizes the true distribution. For $\pi_{\beta, \gamma}^*$, define the reward model

$$\tilde{r}^*(x, a) = \phi_\gamma \left(\frac{\pi_{\beta, \gamma}^*(a | x)}{\pi_{\text{ref}}(a | x)} \right),$$

which is equivalent to r^* up to an action-independent shift, namely, the normalization factor $\lambda_{\beta, \gamma}^*$ in [Lemma H.4](#). Since $\pi_{\beta, \gamma}^* \in \Pi$ under [Assumption H.1](#), we have $\tilde{r}^* \in \mathcal{R}_\Pi$, and for all $(x, a, b) \in \mathcal{X} \times \mathcal{A} \times \mathcal{A}$, it holds that

$$\text{clip}_{2R_{\max}}[\tilde{r}^*(x, a) - \tilde{r}^*(x, b)] = \text{clip}_{2R_{\max}}[r^*(x, a) - r^*(x, b)] = r^*(x, a) - r^*(x, b).$$

The first equality is because action-independent shift between \tilde{r}^* and r^* is cancelled out when taking the difference of rewards, and the second equality is because, by assumption, $r^* \in [0, R_{\max}]$. As a result, the reward difference is bounded in the same range and never clipped.

From this we conclude that $P_{\tilde{r}^*} = P_{r^*} \in \mathcal{P}_\Pi$, and realizability is satisfied. Further, it is easy to see that \mathcal{P}_Π contains only valid distributions. Thus, having satisfied the necessary preconditions, we can invoke [Lemma F.1](#), which guarantees that with probability at least $1 - \delta$, we have

$$\mathbb{E}_{\pi_{\text{ref}}, \pi_{\text{ref}}} [D_{\text{H}}^2(P_{\tilde{r}}(\cdot | x, a, b), P_{r^*}(\cdot | x, a, b))] \leq \frac{2 \log(|\Pi|/\delta)}{n}.$$

To conclude, we extract a bound on reward estimation error from this Hellinger distance bound by using [Lemma H.5](#) with $R = V = 2R_{\max}$, giving

$$\begin{aligned} & \mathbb{E}_{\pi_{\text{ref}}, \pi_{\text{ref}}} \left[\left(\text{clip}_{2R_{\max}}[\hat{r}(x, a) - \hat{r}(x, b)] - \text{clip}_{2R_{\max}}[r^*(x, a) - r^*(x, b)] \right)^2 \right] \\ & \leq 64e^{4R_{\max}} R_{\max}^2 \cdot \mathbb{E}_{\pi_{\text{ref}}, \pi_{\text{ref}}} [D_{\text{H}}^2(P_{\hat{r}}(\cdot | x, a, b), P_{r^*}(\cdot | x, a, b))] \\ & \leq 128e^{4R_{\max}} R_{\max}^2 \cdot \frac{\log(|\Pi|/\delta)}{n}. \end{aligned}$$

□

Proof of [Lemma H.2](#).

First we rewrite the objective as a minimization problem,

$$\begin{aligned} & \underset{\pi}{\text{argmin}} \quad -\mathbb{E}_\pi[\bar{r}(x, a)] + \beta D_f(\pi \| \pi_{\text{ref}}) \\ & \text{s.t.} \quad \rho(x) \sum_a \pi(a | x) = \rho(x) \quad \forall x, \\ & \quad \quad \rho(x) \pi(a | x) \geq 0 \quad \forall x, a. \end{aligned}$$

Here, π is the primal variable, and denote the dual variables as $\lambda : \mathcal{X} \rightarrow \mathbb{R}$ and $\alpha : \mathcal{X} \times \mathcal{A} \rightarrow [0, \infty)$, which correspond to the first and second constraints, respectively. The Lagrangian form is then

$$\mathcal{L}(\pi, \lambda, \alpha) = -\mathbb{E}_\pi[\bar{r}(x, a)] + \beta D_f(\pi \| \pi_{\text{ref}}) + \sum_x \rho(x) \lambda(x) \left(\sum_a \pi(a | x) - 1 \right) - \sum_x \rho(x) \sum_a \alpha(x, a) \pi(a | x).$$

Slater's condition holds since $\bar{\pi}$ itself is a strictly feasible solution, and the objective is convex in $\pi(a | x)$. Then if (π, λ, α) satisfy the KKT conditions, they are the optimal primal and dual variables, which, overloading notation, we denote as $(\pi^*, \lambda^*, \alpha^*)$.

We will demonstrate that setting $\pi^* = \bar{\pi}$, $\lambda^* = 0$, and $\alpha^* = 0$ satisfies the KKT conditions. First, we observe that the proposed solutions are primal and dual feasible. Further, we have $\bar{\pi} > 0$ since $0 \notin \text{dom}(f')$ and $\bar{\pi}(a | x) \in \text{dom}(f')$. As a result, $\rho(x) \alpha^*(x, a) \pi(a | x) = 0$ for all x, a , and complementary slackness is satisfied. Lastly, for stationarity,

$$\begin{aligned} \frac{\partial \mathcal{L}(\pi, \lambda, \alpha)}{\partial \pi(a | x)} &= \rho(x) \left(-\bar{r}(x, a) + \beta f' \left(\frac{\bar{\pi}(a | x)}{\pi_{\text{ref}}(a | x)} \right) + \lambda^*(x) - \alpha^*(x, a) \right) \\ &= \rho(x) \left(-\bar{r}(x, a) + \beta f' \left(\frac{\bar{\pi}(a | x)}{\pi_{\text{ref}}(a | x)} \right) \right) \end{aligned}$$

$$\begin{aligned}
&= \rho(x) \left(-\beta f' \left(\frac{\bar{\pi}(a|x)}{\pi_{\text{ref}}(a|x)} \right) + \beta f' \left(\frac{\bar{\pi}(a|x)}{\pi_{\text{ref}}(a|x)} \right) \right) \\
&= 0,
\end{aligned}$$

where in the second line we substitute $\lambda^* = 0$ and $\alpha^* = 0$, and in third line we have utilized the definition of $\bar{r}(x, a)$ from the lemma statement. \square

Proof of Lemma H.3. For a pair of policies π, π' and $p \geq 1$, we define the norm $\|\cdot\|_{p, \pi \times \pi'} := (\mathbb{E}_{\rho, a \sim \pi, b \sim \pi'} [|\cdot|^p])^{1/p}$. In addition, for notational compactness, we abbreviate $\widehat{\Delta}(x, a, b) := \widehat{r}(x, a) - \widehat{r}(x, b)$, and $\Delta^*(x, a, b) := r^*(x, a) - r^*(x, b)$.

Recall that our goal is to bound the (unclipped) reward estimation error under π using the (clipped) reward estimation error π_{ref} . We begin by decomposing

$$\begin{aligned}
\|\Delta^* - \widehat{\Delta}\|_{1, \pi \times \pi_{\text{ref}}} &= \|\Delta^* - \text{clip}_{2R_{\text{max}}}[\widehat{\Delta}] + \text{clip}_{2R_{\text{max}}}[\widehat{\Delta}] - \widehat{\Delta}\|_{1, \pi \times \pi_{\text{ref}}} \\
&\leq \|\Delta^* - \text{clip}_{2R_{\text{max}}}[\widehat{\Delta}]\|_{1, \pi \times \pi_{\text{ref}}} + \left\| \left(\text{clip}_{2R_{\text{max}}}[\widehat{\Delta}] - \widehat{\Delta} \right) \cdot \mathbb{I} \left[\text{clip}_{2R_{\text{max}}}[\widehat{\Delta}] \neq \widehat{\Delta} \right] \right\|_{1, \pi \times \pi_{\text{ref}}} \\
&\leq \underbrace{\|\Delta^* - \text{clip}_{2R_{\text{max}}}[\widehat{\Delta}]\|_{1, \pi \times \pi_{\text{ref}}}}_{\text{(I) clipped on-policy estimation error}} + \underbrace{V_{\text{max}} \cdot \mathbb{P}_{\pi, \pi_{\text{ref}}} \left(\text{clip}_{2R_{\text{max}}}[\widehat{\Delta}] \neq \widehat{\Delta} \right)}_{\text{(II) bias from clipping}}.
\end{aligned}$$

This splits our bound into two terms. The first is the on-policy error of the clipped reward differences, and can be directly bounded by Lemma H.1 using a standard change-of-measure argument. The second expresses the error of translating the clipped estimates to the unclipped ones in our target bound. For the first term, using Cauchy-Schwarz gives

$$\begin{aligned}
\text{(I)} &= \|\Delta^* - \text{clip}_{2R_{\text{max}}}[\widehat{\Delta}]\|_{1, \pi \times \pi_{\text{ref}}} \leq \sqrt{\mathcal{C}^\pi \cdot \|\Delta^* - \text{clip}_{2R_{\text{max}}}[\widehat{\Delta}]\|_{2, \pi_{\text{ref}} \times \pi_{\text{ref}}}^2} \\
&= \sqrt{\mathcal{C}^\pi \cdot \|\text{clip}_{2R_{\text{max}}}[\Delta^*] - \text{clip}_{2R_{\text{max}}}[\widehat{\Delta}]\|_{2, \pi_{\text{ref}} \times \pi_{\text{ref}}}^2},
\end{aligned}$$

where the last equality uses that $\Delta^* \in [-R_{\text{max}}, R_{\text{max}}]$.

Next, for the second term, we again use Cauchy-Schwarz to change measure onto the offline distribution,

$$\text{(II)} = V_{\text{max}} \cdot \mathbb{P}_{\pi \times \pi_{\text{ref}}} \left(\text{clip}_{2R_{\text{max}}}[\widehat{\Delta}] \neq \widehat{\Delta} \right) \leq V_{\text{max}} \cdot \sqrt{\mathcal{C}^\pi \cdot \mathbb{P}_{\pi_{\text{ref}}, \pi_{\text{ref}}} \left(\text{clip}_{2R_{\text{max}}}[\widehat{\Delta}] \neq \widehat{\Delta} \right)}.$$

Further, using Markov's inequality along with the fact that $\Delta^* \in [-R_{\text{max}}, R_{\text{max}}]$,

$$\begin{aligned}
\mathbb{P}_{\pi_{\text{ref}}, \pi_{\text{ref}}} \left(\text{clip}_{2R_{\text{max}}}[\widehat{\Delta}] \neq \widehat{\Delta} \right) &\leq \mathbb{P}_{\pi_{\text{ref}}, \pi_{\text{ref}}} \left(\left| \text{clip}_{2R_{\text{max}}}[\widehat{\Delta}] \right| = 2R_{\text{max}} \right) \\
&\leq \mathbb{P}_{\pi_{\text{ref}}, \pi_{\text{ref}}} \left(\left| \text{clip}_{2R_{\text{max}}}[\widehat{\Delta}] - \text{clip}_{2R_{\text{max}}}[\Delta^*] \right| \geq R_{\text{max}} \right) \\
&\leq \frac{1}{R_{\text{max}}^2} \left\| \text{clip}_{2R_{\text{max}}}[\widehat{\Delta}] - \text{clip}_{2R_{\text{max}}}[\Delta^*] \right\|_{2, \pi_{\text{ref}} \times \pi_{\text{ref}}}^2.
\end{aligned}$$

Combining inequalities, we obtain

$$\begin{aligned}
\|\Delta^* - \widehat{\Delta}\|_{1, \pi \times \pi_{\text{ref}}} &\leq \left(1 + \frac{V_{\text{max}}}{R_{\text{max}}} \right) \sqrt{\mathcal{C}^\pi \cdot \left\| \text{clip}_{2R_{\text{max}}}[\widehat{\Delta}] - \text{clip}_{2R_{\text{max}}}[\Delta^*] \right\|_{2, \pi_{\text{ref}} \times \pi_{\text{ref}}}^2} \\
&= \left(1 + \frac{V_{\text{max}}}{R_{\text{max}}} \right) \sqrt{(1 + 2D_{\chi^2}(\pi \parallel \pi_{\text{ref}})) \cdot \varepsilon_{\text{stat}}^2} \\
&\leq \frac{2V_{\text{max}}}{R_{\text{max}}} \sqrt{(1 + 2D_{\chi^2}(\pi \parallel \pi_{\text{ref}})) \cdot \varepsilon_{\text{stat}}^2}.
\end{aligned}$$

In the second line we have used $\mathcal{C}^\pi = 1 + 2D_{\chi^2}(\pi \parallel \pi_{\text{ref}})$ and the definition of $\varepsilon_{\text{stat}}^2$ from Lemma H.1, and in the last line we use $V_{\text{max}} \geq R_{\text{max}}$. \square

Lemma H.4. When $\pi_{\text{ref}}(a | x) > 0$ for all $x \in \mathcal{X}$, the optimal policy $\pi_{\beta, \gamma}^*$ for Eq. (39) satisfies

$$r^*(x, a) = \phi_\gamma \left(\frac{\pi_{\beta, \gamma}^*(a | x)}{\pi_{\text{ref}}(a | x)} \right) + \lambda_{\beta, \gamma}^*(x),$$

where $\lambda_{\beta, \gamma}^*$ is an optimal dual variable that normalizes $\pi_{\beta, \gamma}^*$.

Proof of Lemma H.4. It is easy to see that strong duality holds for Eq. (39), since it is convex and strictly feasible (e.g., for the policy π_{ref}). Thus, the KKT conditions give the optimal primal and dual solutions.

Since Eq. (39) is constrained optimization problem (over valid policies), we first define the dual variables. Below, $\lambda : \mathcal{X} \rightarrow \mathbb{R}$ corresponds to the equality constraint that $\sum_a \pi(a | x) = 1$ for all $x \in \mathcal{X}$, and $\alpha : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}_{\geq 0}$ corresponds to the inequality constraint that $\pi(a | x) \geq 0$ for all $(x, a) \in \mathcal{X} \times \mathcal{A}$. After converting Eq. (39) from maximization to minimization, we write Eq. (39) in Lagrangian form as

$$\mathcal{L}(\pi, \lambda, \alpha) = -\mathbb{E}_\pi[r^*(x, a)] + \beta D_{f_{\text{mix}, \gamma}}(\pi \| \pi_{\text{ref}}) + \sum_x \rho(x) \lambda(x) \left(\sum_a \pi(a | x) - 1 \right) - \sum_x \rho(x) \sum_a \alpha(x, a) \pi(a | x),$$

since multiplying each of the solutions by $\rho(x)$ does not affect the value of the saddle-point problem. We denote the optimal primal variable as $\pi_{\beta, \gamma}^*$, and optimal dual variables as $(\lambda_{\beta, \gamma}^*, \alpha_{\beta, \gamma}^*)$.

From stationarity, the optimal primal and dual variables satisfy

$$r^*(x, a) = \phi_\gamma \left(\frac{\pi_{\beta, \gamma}^*(a | x)}{\pi_{\text{ref}}(a | x)} \right) + \lambda_{\beta, \gamma}^*(x) - \alpha_{\beta, \gamma}^*(x, a).$$

Next, for a function g let g^{-1} denote its left inverse, such that $g^{-1}(g(x)) = x$. Because ϕ_γ is injective (see proof of Lemma H.2), it has a left inverse $(\phi_\gamma)^{-1}$, and we can write

$$\pi_{\beta, \gamma}^*(a | x) = \pi_{\text{ref}}(a | x) \cdot (\phi_\gamma)^{-1}(r^*(x, a) - \lambda_{\beta, \gamma}^*(x) + \alpha_{\beta, \gamma}^*(x, a)).$$

Because $\phi_\gamma(z) = z + \gamma \log(z)$, $0 \notin \text{dom}(\phi_\gamma)$, and therefore $0 \notin \text{range}((\phi_\gamma)^{-1})$. Then from the above expression, we observe that $\pi_{\beta, \gamma}^*(a | x) > 0$ since $\pi_{\text{ref}}(a | x) > 0$. It immediately follows that $\alpha_{\beta, \gamma}^*(x, a) = 0$ for all (x, a) from complementary slackness, which states that the optimal solutions satisfy $\pi_{\beta, \gamma}^*(a | x) \cdot \alpha_{\beta, \gamma}^*(x, a) = 0$ for all x, a . This allows us to reduce the expression for r^* to the stated result, that is,

$$r^*(x, a) = \phi_\gamma \left(\frac{\pi_{\beta, \gamma}^*(a | x)}{\pi_{\text{ref}}(a | x)} \right) + \lambda_{\beta, \gamma}^*(x).$$

□

Lemma H.5. For $z \in [-R, R]$ and $z' \in [-V, V]$ where $V \geq R \geq 1$, we have

$$|z - z'| \leq 4e^{2R}V \cdot |\sigma(z) - \sigma(z')|.$$

Additionally, if we define the distribution $P_z(y) = \mathbb{I}\{y = +1\}\sigma(z) + \mathbb{I}\{y = -1\}\sigma(-z)$ for $y \in \{-1, +1\}$ and define $P_{z'}$ analogously, then

$$|z - z'| \leq 4e^{2R}V \cdot D_{\text{H}}(P_z, P_{z'}).$$

Proof of Lemma H.5. We begin with the first statement, and write

$$|z - z'| = \frac{|z - z'|}{|\sigma(z) - \sigma(z')|} \cdot |\sigma(z) - \sigma(z')|.$$

Since $\sigma(z') \in (0, 1)$ but $z' \in [-V, V]$, it can be observed that the slope $\frac{|z - z'|}{|\sigma(z) - \sigma(z')|}$ is smallest where $z \approx z'$, and increases as we move away from this region in either direction. To better intuit the scaling of the slope in terms of V , we expand $|\sigma(z) - \sigma(z')|$ in the denominator to write

$$|z - z'| = \frac{|z - z'|(1 + e^z)(1 + e^{z'})}{|e^z - e^{z'}|} \cdot |\sigma(z) - \sigma(z')|.$$

This indicates that the slope should scale linearly (not exponentially) with the range of z' . For example, as $z' \rightarrow \infty$, $(1 + e^{z'})/|e^z - e^{z'}| = O(1)$.

To make this intuition precise, we split into two cases. First, whenever $e^{z'} \geq \frac{e^{R+z}+1}{e^{R-1}}$ or $e^{z'} \leq \frac{e^{R+z}-1}{e^{R+1}}$ (this constitutes the range where “ $z' \approx z$ ”), we have $1 + e^{z'} \leq e^R |e^z - e^{z'}|$. Then in this region,

$$|z - z'| = \frac{|z - z'| (1 + e^z)(1 + e^{z'})}{|e^z - e^{z'}|} |\sigma(z) - \sigma(z')| \leq 2V(1 + e^R)e^R \cdot |\sigma(z) - \sigma(z')|.$$

Next, for $e^{z'} \in [\frac{e^{R+z}-1}{e^{R+1}}, \frac{e^{R+z}+1}{e^{R-1}}]$, we apply the mean value theorem. Since $\sigma'(x) = e^x(1 + e^{-x})^{-2}$,

$$\begin{aligned} \frac{|z - z'|}{|\sigma(z) - \sigma(z')|} &\leq \sup_{\tilde{z} \in [\min\{z, z'\}, \max\{z, z'\}]} e^{\tilde{z}}(1 + e^{-\tilde{z}})^{-2} \\ &\leq \sup_{e^{\tilde{z}} \in [\frac{e^{R+z}-1}{e^{R+1}}, \frac{e^{R+z}+1}{e^{R-1}}]} e^{\tilde{z}}(1 + e^{-\tilde{z}})^{-2} \\ &\leq 4e^R. \end{aligned}$$

In the second inequality, we use the fact that $e^{z'}, e^z \in [\frac{e^{R+z}-1}{e^{R+1}}, \frac{e^{R+z}+1}{e^{R-1}}]$, and in the third inequality we use the fact that $\sigma'(x)$ is increasing in x , and that $|z| \leq R$. Combining the inequalities for the two regions of $e^{z'}$ gives the result.

For the second statement, we use the fact that

$$2D_{\text{H}}^2(P_z, P_{z'}) \geq \sum_{y \in \{+1, -1\}} \frac{(P_z(y) - P_{z'}(y))^2}{P_z(y) + P_{z'}(y)}.$$

As a result,

$$\sum_{y \in \{+1, -1\}} (P_z(y) - P_{z'}(y))^2 \leq 4D_{\text{H}}^2(P_z, P_{z'}).$$

Since $P_z(y) = 1 - P_z(-y)$ and $P_z(+1) = \sigma(z)$,

$$\sum_{y \in \{+1, -1\}} (P_z(y) - P_{z'}(y))^2 = 2(\sigma(z) - \sigma(z'))^2,$$

and therefore $(\sigma(z) - \sigma(z'))^2 \leq 2D_{\text{H}}^2(P_z, P_{z'})$. The result follows from taking the square root of both sides and combining with the first statement in the lemma. \square

H.2 PROOF OF THEOREM 3.1

Proof of Theorem 3.1. The policy optimization in Line 2 of Algorithm 1 is a special case of Eq. (40) with $\gamma = 1$. As a result, Theorem 3.1 follows directly from Theorem H.1 when instantiated with $\gamma = 1$. \square

H.3 PROOF OF COROLLARY 3.1

Proof of Corollary 3.1. Recall that for any $\beta > 0$, Theorem 3.1 (Eq. (13)) with the policy class $\Pi_{\mathcal{R}}$ ensures that with probability at least $1 - \delta$, for all π^* ,

$$J(\pi^*) - J(\hat{\pi}) \leq c_1 R_{\max} e^{2R_{\max}} \cdot \sqrt{\frac{\mathcal{C}^{\pi^*} \log(|\mathcal{R}|/\delta)}{n}} + c_2 \beta \mathcal{C}^{\pi^*} + c_3 \beta^{-1} \frac{R_{\max}^2 e^{4R_{\max}} \log(|\mathcal{R}|/\delta)}{n} \quad (46)$$

for absolute constants $c_1, c_2, c_3 > 0$. Let us invoke this result with

$$\beta^* = \operatorname{argmax}_{\beta > 0} \max_{\pi^*} \left\{ J(\pi^*) - c_1 R_{\max} e^{2R_{\max}} \cdot \sqrt{\frac{\mathcal{C}^{\pi^*} \log(|\mathcal{R}|/\delta)}{n}} - c_2 \beta \mathcal{C}^{\pi^*} - c_3 \beta^{-1} \frac{R_{\max}^2 e^{4R_{\max}} \log(|\mathcal{R}|/\delta)}{n} \right\}.$$

Then Eq. (46) implies that

$$\max_{\pi^*} \left\{ J(\pi^*) - c_1 R_{\max} e^{2R_{\max}} \cdot \sqrt{\frac{\mathcal{C}^{\pi^*} \log(|\mathcal{R}|/\delta)}{n}} - c_2 \beta^* \mathcal{C}^{\pi^*} - c_3 (\beta^*)^{-1} \frac{R_{\max}^2 e^{4R_{\max}} \log(|\mathcal{R}|/\delta)}{n} \right\} - J(\hat{\pi}) \leq 0,$$

so that by the definition of β^* ,

$$\max_{\beta > 0} \max_{\pi^*} \left\{ J(\pi^*) - c_1 R_{\max} e^{2R_{\max}} \cdot \sqrt{\frac{\mathcal{C}^{\pi^*} \log(|\mathcal{R}|/\delta)}{n}} - c_2 \beta \mathcal{C}^{\pi^*} - c_3 \beta^{-1} \frac{R_{\max}^2 e^{4R_{\max}} \log(|\mathcal{R}|/\delta)}{n} \right\} - J(\hat{\pi}) \leq 0,$$

or equivalently

$$J(\pi^*) - J(\hat{\pi}) \leq c_1 R_{\max} e^{2R_{\max}} \cdot \sqrt{\frac{\mathcal{C}^{\pi^*} \log(|\mathcal{R}|/\delta)}{n}} + c_2 \beta \mathcal{C}^{\pi^*} + c_3 \beta^{-1} \frac{R_{\max}^2 e^{4R_{\max}} \log(|\mathcal{R}|/\delta)}{n} \quad \forall \pi^*, \forall \beta > 0.$$

It follows that for all comparator policies π^* , we have

$$J(\pi^*) - J(\hat{\pi}) \lesssim R_{\max} e^{2R_{\max}} \cdot \sqrt{\frac{\mathcal{C}^{\pi^*} \log(|\mathcal{R}|/\delta)}{n}}$$

by choosing $\beta \propto \sqrt{\frac{R_{\max}^2 e^{4R_{\max}} \log(|\mathcal{R}|/\delta)}{\mathcal{C}^{\pi^*} n}}$ above. □

I PROOFS FOR APPENDIX B

Proof of Proposition B.1. To see that ϕ and ϕ^{-1} are strictly increasing, we note that $\phi'(z) = 1 + \frac{1}{z} > 0$ for all $z > 0$.

We now bound the inverse function ϕ^{-1} . We will use the fact that $z \mapsto W_0(z)$ is increasing over $z \geq 0$ throughout. We first consider the regime where $z \geq 1$. Since $W_0(\cdot)$ is increasing, we have that $\phi^{-1}(z) = W_0(e^z) \leq z$ if and only if $e^z \leq z e^z$, which is clearly true for $z \geq 1$. On the other hand, for $c > 0$ we have $\phi^{-1}(z) = W_0(e^z) \geq c \cdot z$ if and only if $e^z \geq c z e^{cz}$; setting $c = 1/2$ is clearly sufficient.

We now consider the regime where $z \leq 1$. Here, we see that $\phi^{-1}(z) = W(e^z) \leq e^z$ if and only if $e^z \leq e^z e^{e^z}$, which holds for all $z \in \mathbb{R}$. On the other hand have that $\phi^{-1}(z) = W(e^z) \geq e^{-e} e^z$ if and only if $e^z \geq e^{-e} e^z e^{-e e^z}$. Since $z \leq 1$, we have

$$e^{-e} e^z e^{-e e^z} \leq e^{-e} e^z e^{e^z} \leq e^{-e} e^z e^e = e^z,$$

which establishes the result. □

Proof of Proposition B.2. Recall that the optimal policy satisfies

$$r(x, a) = \beta \phi \left(\frac{\pi_{\beta}^*(a | x)}{\pi_{\text{ref}}(a | x)} \right) + Z_{\beta, r}(x), \quad (47)$$

where $Z_{\beta, r}(x)$ is a normalization constant chosen such that $\pi_{\beta}^*(\cdot | x)$ is a valid probability distribution.

We begin by bounding $Z_{\beta, r}(x)$. We will use that $r(x, a) \in [0, R_{\max}]$. Let $x \in \mathcal{X}$ be fixed. By averaging Eq. (47) over $a \sim \pi_{\beta}^*(x)$, we have

$$\mathbb{E}_{a \sim \pi_{\beta}^*(x)} [r(x, a)] = \beta \mathbb{E}_{a \sim \pi_{\beta}^*(x)} \left[\frac{\pi_{\beta}^*(a | x)}{\pi_{\text{ref}}(a | x)} \right] + \beta D_{\text{KL}}(\pi_{\beta}^* \| \pi_{\text{ref}}) + Z_{\beta, r}(x) \geq Z_{\beta, r}(x),$$

so $Z_{\beta, r}(x) \leq R_{\max}$. On the other hand, averaging over $a \sim \pi_{\text{ref}}(x)$, we have

$$\mathbb{E}_{a \sim \pi_{\beta}^*(x)} [r(x, a)] = \beta \mathbb{E}_{a \sim \pi_{\text{ref}}(x)} \left[\frac{\pi_{\beta}^*(a | x)}{\pi_{\text{ref}}(a | x)} \right] - \beta D_{\text{KL}}(\pi_{\text{ref}} \| \pi_{\beta}^*) + Z_{\beta, r}(x)$$

2430

$$\leq \beta + Z_{\beta,r}(x),$$

2431

so $Z_{\beta,r}(x) \geq -\beta$.

2432

2433

Having established that $Z_{\beta,r}(x) \in [-\beta, R_{\max}]$, we will use that $\phi\left(\frac{\pi_{\beta}^*(a|x)}{\pi_{\text{ref}}(a|x)}\right) = \beta^{-1}(r(x, a) - Z_{\beta,r}(x))$, so that our bound on $Z_{\beta,r}$ implies that

2434

2435

2436

2437

$$-\beta^{-1}R_{\max} \leq \phi\left(\frac{\pi_{\beta}^*(a|x)}{\pi_{\text{ref}}(a|x)}\right) \leq 1 + \beta^{-1}R_{\max},$$

2438

or, since ϕ^{-1} is increasing,

2439

2440

2441

$$e^{-e} \cdot e^{-\beta^{-1}R_{\max}} \leq \phi^{-1}(-\beta^{-1}R_{\max}) \leq \frac{\pi_{\beta}^*(a|x)}{\pi_{\text{ref}}(a|x)} \leq \phi^{-1}(1 + \beta^{-1}R_{\max}) \leq 1 + \beta^{-1}R_{\max},$$

2442

where we have used that $\phi^{-1}(z) \leq z$ for $z \geq 1$ and $\phi^{-1}(z) \geq e^{z-e}$ for $z \leq 1$ (by [Proposition B.1](#)).

2443

2444

2445

2446

J PROOFS FOR APPENDIX D

2447

2448

J.1 PROOF OF [THEOREM D.1](#)

2449

2450

2451

Proof of [Theorem D.1](#). We consider a family of instances in which there is a single context (prompt) $\mathcal{X} = \{\emptyset\}$ and four actions (responses) $\mathcal{A} = \{a, b, c, d\}$. We consider the reference policy π_{ref} given by

2452

2453

2454

$$\pi_{\text{ref}}(a' | x) = \begin{cases} \frac{1}{C}, & \text{if } a' = a \text{ or } a' = b, \\ 1 - \frac{2}{C}, & \text{if } a' = c. \end{cases}$$

2455

We consider a preference model class $\mathcal{P} = \{\mathcal{P}^1, \mathcal{P}^2\}$ in which

2456

2457

2458

2459

2460

2461

2462

$$\mathcal{P}^i(a^0 \succ a^1 | x) = (1 + \ell^i(x, a^0, a^1))/2$$

for a function $\ell^i(x, a^0, a^1) \in [-1, +1]$. The functions ℓ^1 and ℓ^2 are defined as follows (we omit the dependence on x , since there is a single context):

$$\ell^1(a^0, a^1) = \ell^2(a^0, a^1) = 0, \quad \forall a^0 \in \mathcal{A}, a^1 \in \{a, b, c\},$$

$$\ell^1(a, d) = 0, \quad \ell^1(b, d) = -1, \quad \ell^1(c, d) = 1$$

$$\ell^2(a, d) = -1, \quad \ell^2(b, d) = 0, \quad \ell^2(c, d) = -1.$$

2463

2464

2465

Note that both functions are skew-symmetric in the sense that $\ell(x, a', a') = 0$ and $\ell(x, a^0, a^1) + \ell(x, a^1, a^0) = 0$ for all $x \in \mathcal{X}$ and $a^0, a^1 \in \mathcal{A}$.

2466

2467

It is straightforward to see that the deterministic policies $\pi_{\text{MW}}^1(x) = a$ and $\pi_{\text{MW}}^2(x) = b$ are minimax winners for ℓ^1 and ℓ^2 respectively. Observe that for both policies, we have

2468

$$\mathcal{C}_{\infty}^{\pi_{\text{MW}}^1} = \mathcal{C}_{\infty}^{\pi_{\text{MW}}^2} = C.$$

2469

2470

To proceed, we compute duality gap an arbitrary policy π under \mathcal{P}^1 and \mathcal{P}^2 . Let $\text{DG}(\pi; \mathcal{P})$ denote the value of $\text{DG}(\pi)$ when \mathcal{P} is the true preference model. Then we have:

2471

2472

2473

2474

2475

2476

2477

Therefore we know

$$\text{DG}(\pi; \mathcal{P}^1) = 2 \max_{q \in \Delta(\mathcal{A})} q(d)(\pi(b) - \pi(c)) - \pi(d)(q(b) - q(c))$$

2478

2479

Following similar computations, we have

2480

2481

2482

2483

$$\text{DG}(\pi; \mathcal{P}^2) = 2 \max_{q \in \Delta(\mathcal{A})} q(d)(\pi(a) + \pi(c)) - \pi(d)(q(a) + q(c)).$$

We aim to show that for all policies π , $\text{DG}(\pi; \mathcal{P}^1) + \text{DG}(\pi; \mathcal{P}^2) \geq \frac{1}{2}$. To do so, we consider two cases. Going forward, we will use that $\text{DG}(\pi; \mathcal{P}^i) \geq 0$.

2484 **Case (1):** $\pi(a) + \pi(c) \geq \frac{1}{2}$. In this case, we have $\text{DG}(\pi; \mathcal{P}^2) \geq \frac{1}{2}$, and thus $\text{DG}(\pi; \mathcal{P}^1) +$
 2485 $\text{DG}(\pi; \mathcal{P}^2) \geq \frac{1}{2}$.

2487 **Case (2):** $\pi(a) + \pi(c) < \frac{1}{4}$. In this case, let $\theta := \pi(b) - \pi(c)$. Then we have $\text{DG}(\pi; \mathcal{P}^1) \geq$
 2488 $2 \max\{\theta, \pi(d)\}$. We observe that $\theta + \pi(d) = \pi(b) + \pi(d) - \pi(c) > \frac{3}{4} - \frac{1}{4} = \frac{1}{2}$. This implies that
 2489 $\text{DG}(\pi; \mathcal{P}^1) > \frac{1}{2}$, and thus $\text{DG}(\pi; \mathcal{P}^1) + \text{DG}(\pi; \mathcal{P}^2) \geq \frac{1}{2}$.

2491 Having established that all π satisfy $\text{DG}(\pi; \mathcal{P}^1) + \text{DG}(\pi; \mathcal{P}^2) \geq \frac{1}{2}$ we can apply the Le Cam two-
 2492 point method (specifically, the variant based on the Bretagnolle-Huber inequality (e.g., Theorem 14.2
 2493 in Lattimore and Szepesvári (2020))), which leads to the following inequality

$$2494 \inf_{\text{Alg}} \sup_{\mathcal{P} \in \mathcal{P}} \mathbb{E}_{\mathcal{D}_{\text{pref}}} [\text{DG}(\hat{\pi}; \mathcal{P})] \geq \frac{1}{8} \exp(-n \cdot D_{\text{KL}}(\rho \otimes \pi_{\text{ref}} \otimes \pi_{\text{ref}} \otimes \mathcal{P}^1 \parallel \rho \otimes \pi_{\text{ref}} \otimes \pi_{\text{ref}} \otimes \mathcal{P}^2)).$$

2496 It can be observed that $D_{\text{KL}}(\rho \otimes \pi_{\text{ref}} \otimes \pi_{\text{ref}} \otimes \mathcal{P}^1 \parallel \rho \otimes \pi_{\text{ref}} \otimes \pi_{\text{ref}} \otimes \mathcal{P}^2) = 0$, since $\ell^1(a^0, a^1) =$
 2497 $\ell^2(a^0, a^1) = 0$ for all $a^0, a^1 \in \{a, b, c\}$, and π_{ref} is supported on $\{a, b, c\}$. We conclude that any
 2498 policy derived from $\mathcal{D}_{\text{pref}}$ must have

$$2500 \mathbb{E}[\text{DG}(\hat{\pi}; \mathcal{P}^i)] \geq \frac{1}{8}$$

2501 for some i . □

2504 J.2 PROOF OF THEOREM D.2

2505 **Proof of Theorem D.2.** Let $\tilde{\pi}$ be the global best response of $\hat{\pi}$:

$$2507 \tilde{\pi} = \operatorname{argmax}_{\pi \in \Pi} \mathbb{E}_{x \sim \rho, a \sim \pi(x), b \sim \hat{\pi}(x)} [\ell^*(x, a, b)],$$

2509 and let $\tilde{\pi}_C$ be the best response within Π_C of $\hat{\pi}$ where $C \geq 1$ (recall that $\Pi_C := \{\pi :$
 2510 $\max_{x \in \mathcal{X}} D_{\chi^2}(\pi(x) \parallel \pi_{\text{ref}}(x)) \leq C\}$ denotes the set of policies with bounded χ^2 -divergence w.r.t.
 2511 π_{ref}):

$$2512 \tilde{\pi}_C = \operatorname{argmax}_{\pi \in \Pi_C} \mathbb{E}_{x \sim \rho, a \sim \pi(x), b \sim \hat{\pi}(x)} [\ell^*(x, a, b)].$$

2514 Recall that $\bar{r}^t(x, a) := \mathbb{E}_{b \sim \pi^t(x)} [\hat{\ell}(x, a, b)]$. Then we know

$$2516 \ell^*(\tilde{\pi}, \hat{\pi}) = \underbrace{\text{subopt}(\hat{\pi}, C)}_{(1)} + \underbrace{\frac{1}{T} \sum_{t=1}^T (\bar{r}^t(\tilde{\pi}_C) - \bar{r}^t(\pi^t))}_{(2)} + \underbrace{\frac{1}{T} \sum_{t=1}^T (\ell^*(\tilde{\pi}_C, \pi^t) - \hat{\ell}(\tilde{\pi}_C, \pi^t))}_{(3)}$$

$$2518 + \underbrace{\frac{1}{T} \sum_{t=1}^T (\bar{r}^t(\tilde{\pi}_C) - \hat{r}^t(\tilde{\pi}_C))}_{(4)} + \underbrace{\frac{1}{T} \sum_{t=1}^T (\hat{r}^t(\pi^t) - \bar{r}^t(\pi^t))}_{(4)}, \quad (48)$$

2524 where $r(\pi) := \mathbb{E}_{x \sim \rho, a \sim \pi(x)} [r(x, a)]$. The decomposition utilizes the fact that $\bar{r}^t(\pi^t) = 0$ and
 2525 $\bar{r}^t(\tilde{\pi}_C) = \hat{\ell}(\tilde{\pi}_C, \pi^t)$. This implies that we only need to bound term (1)(2)(3)(4) in Eq. (48) to upper
 2526 bound the gap of $\hat{\pi}$.

2527 **Bounding term (1).** Let $g_x(p)$ to denote the mixed divergence $\beta D_{f_{\text{mix}}}(p(x) \parallel \pi_{\text{ref}}(x))$. Then we
 2528 have the following guarantee on regularized policy mirror descent (formal version of Lemma D.1):

2530 **Lemma J.1.** For any $C \geq 0$, we have for all policy $\pi \in \Pi_C$ that

$$2531 \frac{1}{T} \sum_{t=1}^T (\hat{r}^t(\pi) - \hat{r}^t(\pi^t)) \leq \frac{2\beta C}{\eta T} + 2\beta C - \frac{1}{T} \sum_{t=1}^{T+1} \mathbb{E}_{x \sim \rho} [g_x(\pi^t)]$$

$$2533 + \frac{\eta}{2\beta} + \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{x \sim \rho} [\langle \hat{r}^t(x, \cdot) - G^t(\pi^{t+1}, x, \cdot), \pi(x) - \pi^{t+1}(x) \rangle],$$

2534 where $G^t(\pi, x, a) := \beta \left((1 + \frac{1}{\eta}) \phi \left(\frac{\pi(a|x)}{\pi_{\text{ref}}(a|x)} \right) - \frac{1}{\eta} \phi \left(\frac{\pi^t(a|x)}{\pi_{\text{ref}}(a|x)} \right) \right)$ for all $\pi \in \Pi, x \in \mathcal{X}, a \in \mathcal{A}$.

To simplify writing, we use $\bar{\pi}^{t+1}$ to denote the minimizer of the following regularized RL objective:

$$\bar{\pi}^{t+1}(x) := \arg \min_{p \in \Delta(\mathcal{X})} \langle -\hat{r}^t(x, \cdot), p \rangle + \beta D_{f_{\chi_{\text{mix}}}}(p \| \pi_{\text{ref}}(x)) + \frac{\beta}{\eta} B_x(p, \pi^t), \quad \forall x \in \mathcal{X}.$$

Then [Assumption D.2](#) indicates that $\bar{\pi}^{t+1} \in \Pi$ for all $t \in [T]$. In addition, by introducing Lagrangian multipliers into the above optimization problem and following similar arguments in the proof of [Lemma H.4](#), we know

$$f_{\bar{\pi}^{t+1}, \pi^t}^{\beta, \eta}(x, a, b) - (\hat{r}^t(x, a) - \hat{r}^t(x, b)) = 0, \quad \forall x \in \mathcal{X}, a, b \in \mathcal{A}. \quad (49)$$

Recall that by definition $f_{\pi, \pi^t}^{\beta, \eta}(x, a, b) = G^t(\pi, x, a) - G^t(\pi, x, b)$ for all policies $\pi \in \Pi$. This implies that we have

$$\begin{aligned} & \mathbb{E}_{x \sim \rho} [\langle \hat{r}^t(x, \cdot) - G^t(\pi^{t+1}, x, \cdot), \pi(x) - \pi^{t+1}(x) \rangle] \\ &= \mathbb{E}_{x \sim \rho} [\langle \hat{r}^t(x, \cdot) - G^t(\pi^{t+1}, x, \cdot), \pi(x) - \pi_{\text{ref}}(x) \rangle] + \mathbb{E}_{x \sim \rho} [\langle \hat{r}^t(x, \cdot) - G^t(\pi^{t+1}, x, \cdot), \pi_{\text{ref}}(x) - \pi^{t+1}(x) \rangle] \\ &= \underbrace{(f_{\bar{\pi}^{t+1}, \pi^t}^{\beta, \eta} - f_{\pi^{t+1}, \pi^t}^{\beta, \eta})(\rho, \pi, \pi_{\text{ref}})}_{(5)} + \underbrace{(f_{\pi^{t+1}, \pi^t}^{\beta, \eta} - f_{\bar{\pi}^{t+1}, \pi^t}^{\beta, \eta})(\rho, \pi^{t+1}, \pi_{\text{ref}})}_{(6)}, \end{aligned}$$

where we use $f(\rho, \pi, \pi')$ to denote the expectation $\mathbb{E}_{x \sim \rho, a \sim \pi(x), b \sim \pi'(x)}[f(x, a, b)]$ and the last step utilizes [Eq. \(49\)](#). Therefore, to bound term (1), we need to bound term (5) and (6) respectively. To simplify writing, we define $L(\pi, \pi', \pi'')$ as follows:

$$L(\pi, \pi', \pi'') := \mathbb{E}_{x \sim \rho, a \sim \pi_{\text{ref}}(x), b \sim \pi_{\text{ref}}(x)} \left[\left(\text{clip}_4(f_{\pi, \pi''}^{\beta, \eta}(x, a, b)) - \text{clip}_4(f_{\pi', \pi''}^{\beta, \eta}(x, a, b)) \right)^2 \right],$$

Note that we have the following guarantee of least squares regression from the literature ([Lemma 15 in Song et al. \(2022\)](#))

Lemma J.2 (least squares regression). *Let $\{(y_i, z_i)\}_{i=1}^K$ be a dataset of K points where each point are independently sampled from $y_i \sim \mu$ and $z_i \sim p(\cdot | y_i) := h^*(y_i) + \varepsilon_i$. Let $\mathcal{H} : \mathcal{Y} \rightarrow [-R, R]$ be a real valued functions where $h^* \in \mathcal{H}$ and $R > 0$. Then if $\{\varepsilon_i\}_{i=1}^K$ are independent random variables such that $\mathbb{E}[z_i | y_i] = h^*(y_i)$, the least squares solution $\hat{h} = \arg \min_{h \in \mathcal{H}} \sum_{i=1}^K (h(y_i) - z_i)^2$ satisfies with probability at least $1 - \delta$ that*

$$\mathbb{E}_{x \sim \mu} [(\hat{h}(y) - h^*(y))^2] \lesssim \frac{R^2 \log(|\mathcal{H}|/\delta)}{K}.$$

The proof of the above lemma is omitted. Applying [Lemma J.2](#) to the least squares solution π^{t+1} , we have the following concentration lemma:

Lemma J.3 (concentration in optimization). *Suppose [Assumption D.2](#) and [Assumption D.3](#) hold. Then with probability at least $1 - \delta/4$, we have for all policy $t \in [T]$ that*

$$L(\pi^{t+1}, \bar{\pi}^{t+1}, \pi^t) \leq \frac{C_{\text{con}} \log(|\Pi|/\delta)}{m} := \varepsilon_{\text{md}}^2,$$

where $C_{\text{con}} > 0$ is a universal constant.

In the following discussion, we use \mathcal{E}_1 to denote the event in [Lemma J.3](#). Then under \mathcal{E}_1 , by following the same arguments in the proof of [Lemma H.3](#), we have the following bound on $\|f_{\bar{\pi}^{t+1}, \pi^t}^{\beta, \eta} - f_{\pi^{t+1}, \pi^t}^{\beta, \eta}\|_{1, \pi \times \pi_{\text{ref}}}$:

$$\|f_{\bar{\pi}^{t+1}, \pi^t}^{\beta, \eta} - f_{\pi^{t+1}, \pi^t}^{\beta, \eta}\|_{1, \pi \times \pi_{\text{ref}}} \leq V_{\text{max}} \sqrt{(1 + 2D_{\chi^2}(\pi \| \pi_{\text{ref}})) \varepsilon_{\text{md}}^2}, \quad \forall \pi \in \Pi, t \in [T]. \quad (50)$$

Therefore, with [Eq. \(50\)](#) we know that conditioned on \mathcal{E}_1 , for any policy $\pi \in \Pi_C$ we have

$$(5) \leq V_{\text{max}} \sqrt{3C\varepsilon_{\text{md}}^2}, \quad (6) \leq V_{\text{max}} \sqrt{(1 + 2D_{\chi^2}(\pi^{t+1} \| \pi_{\text{ref}})) \varepsilon_{\text{md}}^2} \leq \frac{V_{\text{max}}^2 \varepsilon_{\text{md}}^2}{\beta} + \frac{1}{2} \mathbb{E}_{x \sim \rho} [g_x(\pi^{t+1})] + V_{\text{max}} \varepsilon_{\text{md}},$$

where we use AM-GM inequality in the last step, the definition of $g_x(\pi) := \beta D_{f_{\chi_{\text{mix}}}}(\pi(\cdot|x) \| \pi_{\text{ref}}(\cdot|x))$, and $D_{f_{\chi_{\text{mix}}}}(p(x) \| \pi_{\text{ref}}(x)) \geq D_{\chi^2}(p(x) \| \pi_{\text{ref}}(x))$ since KL is non-negative

In summary, conditioned on \mathcal{E}_1 , we have

$$(1) \leq \frac{2\beta C}{\eta T} + 2\beta C - \frac{1}{2T} \sum_{t=1}^{T+1} \mathbb{E}_{x \sim \rho} [g_x(\pi^t)] + \frac{\eta}{2\beta} + V_{\max} \sqrt{4C\varepsilon_{\text{md}}^2} + \frac{V_{\max}^2 \varepsilon_{\text{md}}^2}{\beta}. \quad (51)$$

Bounding term (2). From Cauchy-Schwartz's inequality, we have

$$\begin{aligned} & \ell^*(\tilde{\pi}_C, \pi^t) - \widehat{\ell}(\tilde{\pi}_C, \pi^t) \\ & \leq \sqrt{\mathbb{E}_{x \sim \rho, a \sim \pi_{\text{ref}}(x), b \sim \pi_{\text{ref}}(x)} [(\ell^*(x, a, b) - \widehat{\ell}(x, a, b))^2] (1 + 2D_{\chi^2}(\rho \otimes \tilde{\pi}_C \otimes \pi^t \| \rho \otimes \pi_{\text{ref}} \otimes \pi_{\text{ref}}))}, \end{aligned}$$

where $\rho \otimes \pi_1 \otimes \pi_2$ denotes the joint distribution of (x, a, b) where $x \sim \rho, a \sim \pi_1(x), b \sim \pi_2(x)$ for all $\pi_1, \pi_2 \in \Pi$. Applying the guarantee of least squares regression (Lemma J.2) to the least squares solution $\widehat{\ell}$, we have under Assumption D.1, with probability at least $1 - \delta/4$, the following event holds:

$$\mathbb{E}_{x \sim \rho, y^0 \sim \pi_{\text{ref}}(x), y^1 \sim \pi_{\text{ref}}(x)} \left[(\widehat{\ell}(x, y^0, y^1) - \ell^*(x, y^0, y^1))^2 \right] \leq O\left(\frac{\ln(|\mathcal{L}|/\delta)}{n}\right) := \varepsilon_{\text{general}}^2. \quad (52)$$

Denote the event in Eq. (52) by \mathcal{E}_2 . On the other hand, we can obtain that:

$$\begin{aligned} 1 + 2D_{\chi^2}(\rho \otimes \tilde{\pi}_C \otimes \pi^t \| \rho \otimes \pi_{\text{ref}} \otimes \pi_{\text{ref}}) &= \sum_x \rho(x) \sum_a \frac{(\tilde{\pi}_C(a|x))^2}{\pi_{\text{ref}}(a|x)} \sum_b \frac{(\pi^t(b|x))^2}{\pi_{\text{ref}}(b|x)} \\ &= \sum_x \rho(x) (1 + 2D_{\chi^2}(\tilde{\pi}_C(x) \| \pi_{\text{ref}}(x))) (1 + 2D_{\chi^2}(\pi^t(x) \| \pi_{\text{ref}}(x))) \\ &\leq 6C (\mathbb{E}_{x \sim \rho} [D_{\chi^2}(\pi^t(x) \| \pi_{\text{ref}}(x))] + 1) \end{aligned}$$

where the last step is due to $\tilde{\pi}_C \in \Pi_C$. Therefore, conditioned on \mathcal{E}_2 , we have

$$\begin{aligned} \ell^*(\tilde{\pi}_C, \pi^t) - \widehat{\ell}(\tilde{\pi}_C, \pi^t) &\leq \sqrt{6C \mathbb{E}_{x \sim \rho} [D_{\chi^2}(\pi^t(x) \| \pi_{\text{ref}}(x))] \varepsilon_{\text{general}}^2} + \sqrt{6C \varepsilon_{\text{general}}^2} \\ &\leq \frac{1}{2} \mathbb{E}_{x \sim \rho} [g_x(\pi^t)] + \frac{3C \varepsilon_{\text{general}}^2}{\beta} + \sqrt{6C \varepsilon_{\text{general}}^2}. \end{aligned}$$

In summary, we have

$$\frac{1}{T} \sum_{t=1}^T \ell^*(\tilde{\pi}_C, \pi^t) - \widehat{\ell}(\tilde{\pi}_C, \pi^t) \leq \frac{1}{2T} \sum_{t=1}^T \mathbb{E}_{x \sim \rho} [g_x(\pi^t)] + \frac{3C \varepsilon_{\text{general}}^2}{\beta} + \sqrt{6C \varepsilon_{\text{general}}^2}. \quad (53)$$

Bounding term (3). Recall that $\widehat{r}^t(x, a) = \widehat{\ell}(x, a, b_t)$ where $b_t \sim \pi^t(x)$ is an unbiased estimator of \bar{r}^t . Fix any policy $\pi \in \Pi$, then from Azuma-Hoeffding's inequality, we have with probability at least $1 - \delta'$ that

$$\left| \sum_{t=1}^T \widehat{r}^t(\pi) - \sum_{t=1}^T \bar{r}^t(\pi) \right| \lesssim \sqrt{T \log(1/\delta')}.$$

By union bound, with probability at least $1 - \delta/4$ we have that for all $\pi \in \Pi$:

$$\left| \sum_{t=1}^T \widehat{r}^t(\pi) - \sum_{t=1}^T \bar{r}^t(\pi) \right| \lesssim \sqrt{T \log(|\Pi|/\delta)}.$$

Therefore, specifically for $\tilde{\pi}_C$, we have

$$(3) \lesssim \sqrt{\frac{\log(|\Pi|/\delta)}{T}}. \quad (54)$$

2646 **Bounding term (4).** From Azuma-Hoeffding’s inequality, we have with probability at least $1 - \delta/4$
 2647 that

$$2648 \left| \sum_{t=1}^T \hat{r}^t(\pi^t) - \sum_{t=1}^T \bar{r}^t(\pi^t) \right| \lesssim \sqrt{T \log(1/\delta')}.$$

2651 Therefore, we have

$$2652 (4) \lesssim \sqrt{\frac{\log(1/\delta)}{T}}. \quad (55)$$

2655 **Putting everything together.** Substituting Eq. (51)(53)(54)(55) into (48), we have with probability
 2656 at least $1 - \delta$ that

$$2657 \ell^*(\tilde{\pi}, \hat{\pi}) \lesssim \text{subopt}(\hat{\pi}, C) + \frac{C\beta}{\eta T} + C\beta + \frac{\eta}{\beta} + V_{\max} \sqrt{C\varepsilon_{\text{md}}^2} + \frac{V_{\max}^2 \varepsilon_{\text{md}}^2}{2\beta}$$

$$2660 + \frac{C\varepsilon_{\text{general}}^2}{\beta} + \sqrt{C\varepsilon_{\text{general}}^2} + \sqrt{\frac{\log \frac{|\Pi|}{\delta}}{T}}.$$

2662 By selecting

$$2663 T = \frac{mn}{nV_{\max}^2 + m}, \quad \beta = \frac{1}{\sqrt{T}}, \quad \eta = \frac{1}{T},$$

2666 we have with probability at least $1 - \delta$ that

$$2667 \ell^*(\tilde{\pi}, \hat{\pi}) \lesssim \text{subopt}(\hat{\pi}, C) + C \left(\frac{V_{\max} \log(|\Pi|/\delta)}{\sqrt{m}} + \frac{\log(|\Pi||\mathcal{L}|/\delta)}{\sqrt{n}} \right)$$

2670 Note that due to the skew symmetry of ℓ^* , we have:

$$2671 \min_{\pi \in \Pi} \mathbb{E}_{x \sim \rho, a \sim \hat{\pi}(x), b \sim \pi(x)} [\ell^*(x, a, b)] = - \max_{\pi \in \Pi} \mathbb{E}_{x \sim \rho, a \sim \pi(x), b \sim \hat{\pi}(x)} [\ell^*(x, a, b)] = -\ell^*(\tilde{\pi}, \hat{\pi}).$$

2673 This implies that $\text{DG}(\hat{\pi}) \leq 2\ell^*(\tilde{\pi}, \hat{\pi})$, which concludes our proof. \square

2675 J.3 PROOFS FOR SUPPORTING LEMMAS

2676 **Proof of Lemma J.1.** First for all $t \in [T]$, $s \in \mathcal{S}$ and any policy $\pi \in \Pi_C$, we have

$$2677 \langle \eta \hat{r}^t(x), \pi(x) - \pi^t(x) \rangle + \eta g_x(\pi^t) - \eta g_x(\pi)$$

$$2678 = \langle \eta \hat{r}^t(x) - (1 + \eta) \nabla g_x(\pi^{t+1}) + \nabla g_x(\pi^t), \pi(x) - \pi^{t+1}(x) \rangle$$

$$2679 + \underbrace{\langle \nabla g_x(\pi^{t+1}) - \nabla g_x(\pi^t), \pi(x) - \pi^{t+1}(x) \rangle}_{(7)} + \underbrace{\langle \eta \hat{r}^t(x), \pi^{t+1}(x) - \pi^t(x) \rangle}_{(8)}$$

$$2680 + \underbrace{\langle \eta \nabla g_x(\pi^{t+1}), \pi(x) - \pi^{t+1}(x) \rangle + \eta g_x(\pi^t) - \eta g_x(\pi)}_{(9)},$$

2686 Note that we have

$$2687 \langle \eta \hat{r}^t(x) - (1 + \eta) \nabla g_x(\pi^{t+1}) + \nabla g_x(\pi^t), \pi(x) - \pi^{t+1}(x) \rangle = \eta \langle \hat{r}^t(x, \cdot) - G^t(\pi^{t+1}, x, \cdot), \pi(x) - \pi^{t+1}(x) \rangle$$

2688 Next we bound the term (7)(8)(9) respectively.

2689 **Bounding term (7).** Note that we have the following three point lemma:

2691 **Lemma J.4** (three point lemma). *For any $p_1, p_2, p_3 : \mathcal{X} \mapsto \Delta(\mathcal{Y})$, we have for all $x \in \mathcal{X}$*

$$2692 \frac{1}{\beta} \langle \nabla g_x(p_1) - \nabla g_x(p_2), p_3(x) - p_1(x) \rangle = B_x(p_3, p_2) - B_x(p_3, p_1) - B_x(p_1, p_2).$$

2695 **Proof.** By definition, we know

$$2696 \beta B_x(p, p') = g_x(p) - g_x(p') - \langle \nabla g_x(p'), p - p' \rangle.$$

2697 Substitute the definition into Lemma J.4 and we can prove the lemma. \square

2698 From Lemma J.4, we can rewrite (7) as follows:

$$2699 (7) = \beta (B_x(\pi, \pi^t) - B_x(\pi, \pi^{t+1}) - B_x(\pi^{t+1}, \pi^t)).$$

Bounding term (8). From Cauchy-Schwartz inequality, we have

$$(8) \leq \sum_{a \in \mathcal{A}} \frac{\beta(\pi^{t+1}(a|x) - \pi^t(a|x))^2}{2\pi_{\text{ref}}(a|x)} + \frac{\pi_{\text{ref}}(a|x)\eta^2(\hat{r}^t(x, a))^2}{2\beta} \leq \beta B_x(\pi^{t+1}, \pi^t) + \frac{\eta^2}{2\beta},$$

where the last step comes from the definition of B_x .

Bounding term (9). Since g_x is convex, we know

$$\langle \eta \nabla g_x(\pi^{t+1}), \pi - \pi^{t+1} \rangle \leq \eta g_x(\pi) - \eta g_x(\pi^{t+1}).$$

This implies that

$$(3) \leq \eta (g_x(\pi^t) - g_x(\pi^{t+1})).$$

In summary, for all $t \in [T]$, $s \in \mathcal{S}$ and any policy $\pi \in \Pi_C$, we have

$$\begin{aligned} & \langle \eta \hat{r}^t(x), \pi(x) - \pi^t(x) \rangle + \eta g_x(\pi^t) - \eta g_x(\pi) \leq \beta (B_x(\pi, \pi^t) - B_x(\pi, \pi^{t+1})) \\ & + \eta (g_x(\pi^t) - g_x(\pi^{t+1})) + \frac{\eta^2}{2\beta} + \eta \langle \hat{r}^t(x, \cdot) - G^t(\pi^{t+1}, x, \cdot), \pi(x) - \pi^{t+1}(x) \rangle. \end{aligned}$$

This implies that for any policy $\pi \in \Pi_C$:

$$\begin{aligned} \sum_{t=1}^T (\hat{r}^t(\pi) - \hat{r}^t(\pi^t)) & \leq T \mathbb{E}_{x \sim \rho} [g_x(\pi)] - \sum_{t=1}^{T+1} \mathbb{E}_{x \sim \rho} [g_x(\pi^t)] + \frac{\beta}{\eta} \mathbb{E}_{x \sim \rho} [B_x(\pi, \pi^1)] + \frac{\eta T}{2\beta} \\ & + \sum_{t=1}^T \mathbb{E}_{x \sim \rho} [\langle \hat{r}^t(x, \cdot) - G^t(\pi^{t+1}, x, \cdot), \pi(x) - \pi^{t+1}(x) \rangle] \\ & \leq 2TC\beta - \sum_{t=1}^{T+1} \mathbb{E}_{x \sim \rho} [g_x(\pi^t)] + \frac{2C\beta}{\eta} + \frac{\eta T}{2\beta} \\ & + \sum_{t=1}^T \mathbb{E}_{x \sim \rho} [\langle \hat{r}^t(x, \cdot) - G^t(\pi^{t+1}, x, \cdot), \pi(x) - \pi^{t+1}(x) \rangle] \end{aligned}$$

Here the last step uses the fact that $B_x(\cdot, \pi_{\text{ref}}) = \frac{1}{\beta} g_x(\cdot)$ and $\pi \in \Pi_C$. This concludes our proof. \square

Proof of Lemma J.3. Let $\hat{L}(\pi, \pi', \pi'')$ denote the empirical squared loss:

$$\hat{L}(\pi, \pi', \pi'') := \sum_{(\bar{x}, \bar{a}, \bar{b})} \left(\text{clip}_4(f_{\pi, \pi'}^{\beta, \eta}(\bar{x}, \bar{a}, \bar{b})) - \text{clip}_4(f_{\pi', \pi''}^{\beta, \eta}(\bar{x}, \bar{a}, \bar{b})) \right)^2.$$

Fix any $\pi', \pi'' \in \Pi$ and consider the following LSR problems:

$$\pi(\pi', \pi'') := \operatorname{argmin}_{\pi \in \Pi} \hat{L}(\pi, \pi', \pi'').$$

Then from Lemma J.2, we know with probability at least $1 - \delta'$ that

$$L(\pi(\pi', \pi''), \pi', \pi'') \lesssim \frac{\log(|\Pi|/\delta')}{M}.$$

Therefore, by union bound, we know with probability at least $1 - \delta'$ that for all $\pi', \pi'' \in \Pi$:

$$L(\pi(\pi', \pi''), \pi', \pi'') \lesssim \frac{\log(|\Pi|/\delta')}{M}.$$

The proof is concluded by noticing that $\pi^{t+1} = \operatorname{argmin}_{\pi \in \Pi} \hat{L}(\pi, \bar{\pi}^{t+1}, \pi^t)$ under Assumption D.2. \square

2754 K PROOFS FOR APPENDIX C

2755 The section contains the proofs for the main guarantee χ^2 -RLHF in [Appendix C \(Theorem C.1\)](#). We
 2756 first prove two results, [Theorem K.1](#) and [Corollary K.1](#), which correspond to exact (i.e., including
 2757 precise constants) versions of the two statements in [Theorem C.1](#). We also analyze χ^2 -RLHF with
 2758 $\eta = 0$ in [Corollary K.2](#).
 2759

2760 Throughout this section, we make use of the following η -smoothed version of the L_1 concentrability
 2761 coefficient:

$$2762 \mathcal{C}_\eta^\pi := \mathbb{E}_\pi \left[\frac{\pi(a | x)}{\pi_{\text{ref}}(a | x) + \eta\pi(a | x)} \right].$$

2765 It is easy to see that for any $\eta \geq 0$ we have $\mathcal{C}_\eta^\pi \leq \mathcal{C}^\pi$, as well as $\mathcal{C}_\eta^\pi \leq \eta^{-1}$.

2766 **Theorem K.1** (General regret bound for [Algorithm 2](#)). *Suppose [Assumption C.1](#) and [Assumption C.2](#)*
 2767 *hold for parameters $\beta > 0$ and $\eta \in [0, \frac{\beta}{8R_{\max}}]$. Then with probability at least $1 - \delta$, the policy $\hat{\pi}$*
 2768 *produced by χ^2 -RLHF ([Algorithm 2](#)) satisfies*
 2769

$$2770 J(\pi^*) - J(\hat{\pi}) \leq 2\sqrt{\mathcal{C}_{\eta^*}^{\pi^*} \cdot \varepsilon_{\text{stat}}^2} + 2\beta \cdot \mathcal{C}_{\eta^*}^{\pi^*} + 4\beta^{-1} \cdot \varepsilon_{\text{stat}}^2$$

$$2771 + 4\beta \cdot \left(\min\{\mathcal{C}_\infty^{\pi^*}, \eta^{-1}\} + \min\left\{\max_{\pi \in \Pi} \mathcal{C}_\infty^\pi, \eta^{-1}\right\} \right) \varepsilon_x^2 + 2R_{\max}\varepsilon_x.$$

2774 where $\varepsilon_{\text{stat}}^2 = \frac{32R_{\max}^2 e^{4R_{\max}} \log(3|\mathcal{R}|/\delta)}{n}$ and $\varepsilon_x = \sqrt{\frac{\log(3|\Pi|/\delta)}{2n_x}}$.
 2775

2776 The following results are immediate consequences of [Theorem K.1](#).
 2777

2778 **Corollary K.1** (Smoothed χ^2 -regularization). Given π^* , let $\eta = \frac{\beta}{8R_{\max}}$ and $\beta =$
 2779 $2\sqrt{\frac{32R_{\max}^2 e^{4R_{\max}} \log(3|\mathcal{R}|/\delta)}{n\mathcal{C}^{\pi^*}}}$. Then under the preconditions of [Theorem K.1](#), with probability at least
 2780 $1 - \delta$, the policy $\hat{\pi}$ produced by χ^2 -RLHF ([Algorithm 2](#)) satisfies
 2781

$$2782 J(\pi^*) - J(\hat{\pi}) \leq 20R_{\max}e^{2R_{\max}}\sqrt{\frac{2\mathcal{C}^{\pi^*} \log(3|\mathcal{R}|/\delta)}{n}} + R_{\max}\sqrt{\frac{2\log(3|\Pi|/\delta)}{n_x}} + \frac{32R_{\max} \log(3|\Pi|/\delta)}{n_x}.$$

2785 **Corollary K.2** (Non-smoothed χ^2 -regularization). Given π^* , let $\eta = 0$ and $\beta =$
 2786 $2\sqrt{\frac{32R_{\max}^2 e^{4R_{\max}} \log(3|\mathcal{R}|/\delta)}{n\mathcal{C}^{\pi^*}}}$. Then under the preconditions of [Theorem K.1](#), with probability at least
 2787 $1 - \delta$, the policy $\hat{\pi}$ produced by χ^2 -RLHF ([Algorithm 2](#)) satisfies
 2788

$$2789 J(\pi^*) - J(\hat{\pi}) \leq 20R_{\max}e^{2R_{\max}}\sqrt{\frac{2\mathcal{C}^{\pi^*} \log(3|\mathcal{R}|/\delta)}{n}} + R_{\max}\sqrt{\frac{2\log(3|\Pi|/\delta)}{n_x}}$$

$$2790 + 32\left(\mathcal{C}_\infty^{\pi^*} + \max_{\pi \in \Pi} \mathcal{C}_\infty^\pi\right) \cdot \frac{\log(3|\Pi|/\delta)}{n_x} \cdot \sqrt{\frac{2\log(3|\mathcal{R}|/\delta)}{n}}.$$

2795 **Proof of [Theorem K.1](#)**. The proof follows largely the same lines of analyses as the proof of
 2796 [Theorem H.1](#). One difference is that in [Algorithm 2](#), we approximate the RLHF objective using
 2797 contexts are sampled from \mathcal{D}_x , so we require additional concentration arguments to show that the
 2798 empirical objective approximates its population counterpart.
 2799

2800 **Basic concentration results.** We begin by stating the two concentration inequalities, which, given
 2801 the reward model \hat{r} produced in [Eq. \(26\)](#), bound the error between $\hat{J}_{\beta, \eta}^{\hat{r}}$ and its the population version
 2802 $J_{\beta, \eta}^{\hat{r}}$.
 2803

2804 We will handle the return and regularization terms separately, which will later allow us to obtain
 2805 tighter bounds. Define

$$2806 \hat{J}(\pi) := \frac{1}{n_x} \sum_{x \in \mathcal{D}_x} \mathbb{E}_\pi[\hat{r}(x, a) | x],$$

and

$$\widehat{\mathcal{C}}_\eta^\pi(\pi) := \frac{1}{n_x} \sum_{x \in \mathcal{D}_x} \mathbb{E}_\pi \left[\sum_a \frac{\pi^2(a|x)}{\pi_{\text{ref}}(a|x) + \eta\pi(a|x)} \mid x \right],$$

so that $\widehat{J}_{\beta,\eta}^\pi(\pi) = \widehat{J}(\pi) - \beta\widehat{\mathcal{C}}_\eta^\pi(\pi)$.

Fix $\delta' \in (0, 1]$, which we will specify at the end of this proof. Since $\max_x \mathbb{E}_\pi[\widehat{r}(x, a) \mid x] \leq R_{\max}$, a straightforward application of Hoeffding's inequality guarantees that with probability at most $1 - \delta'$, for all $\pi \in \Pi$ we have that

$$\left| \widehat{J}(\pi) - \mathbb{E}_\pi[\widehat{r}(x, a)] \right| \leq R_{\max} \sqrt{\frac{\log(2|\Pi|/\delta')}{2n_x}}. \quad (56)$$

Next, we consider the regularization term. Since $\sum_a \frac{\pi^2(a|x)}{\pi_{\text{ref}}(a|x) + \eta\pi(a|x)} \leq \min\{\mathcal{C}_\infty^\pi, \eta^{-1}\}$ for any $x \in \mathcal{X}$, we use Bernstein's inequality to derive the following result.

Lemma K.1. *With probability at least $1 - \delta$, for any $\pi \in \Pi$, we have*

$$\left| \widehat{\mathcal{C}}_\eta^\pi - \mathcal{C}_\eta^\pi \right| \leq \frac{\mathcal{C}_\eta^\pi}{2} + \frac{2 \min\{\mathcal{C}_\infty^\pi, \eta^{-1}\} \log(2|\Pi|/\delta)}{n_x}.$$

Define $\varepsilon_x := \sqrt{\frac{\log(2|\Pi|/\delta')}{2n_x}}$. The above lemma implies that for all $\pi \in \Pi$, we have

$$\widehat{\mathcal{C}}_\eta^\pi \leq \frac{3\mathcal{C}_\eta^\pi}{2} + 4 \min\{\mathcal{C}_\infty^\pi, \eta^{-1}\} \cdot \varepsilon_x^2, \quad \text{and} \quad \widehat{\mathcal{C}}_\eta^\pi \geq \frac{\mathcal{C}_\eta^\pi}{2} - 4 \min\{\mathcal{C}_\infty^\pi, \eta^{-1}\} \cdot \varepsilon_x^2.$$

Together with Eq. (56), this implies that for all $\pi \in \Pi$,

$$\widehat{J}_{\beta,\eta}^\pi(\pi) = \widehat{J}(\pi) - \beta\widehat{\mathcal{C}}_\eta^\pi \leq \mathbb{E}_\pi[\widehat{r}(x, a)] - \frac{\beta\mathcal{C}_\eta^\pi}{2} + 4\beta \min\{\mathcal{C}_\infty^\pi, \eta^{-1}\} \varepsilon_x^2 + R_{\max} \varepsilon_x, \quad (57)$$

and

$$\widehat{J}_{\beta,\eta}^\pi(\pi) = \widehat{J}(\pi) - \beta\widehat{\mathcal{C}}_\eta^\pi \geq \mathbb{E}_\pi[\widehat{r}(x, a)] - \frac{3\beta\mathcal{C}_\eta^\pi}{2} - 4\beta \min\{\mathcal{C}_\infty^\pi, \eta^{-1}\} \varepsilon_x^2 - R_{\max} \varepsilon_x. \quad (58)$$

Estimation error bounds. Next, we state the following off- and on-policy reward estimation error bounds for the reward model \widehat{r} , analogous to Lemma H.1 and Lemma H.3 for χPO .

Lemma K.2. *Suppose Assumption C.1 holds. Then with probability at least $1 - \delta$, the reward model \widehat{r} learned in Eq. (26) satisfies*

$$\varepsilon_{\text{stat}}^2 =: \mathbb{E}_{\pi_{\text{ref}}, \pi_{\text{ref}}} \left[((\widehat{r}(x, a) - \widehat{r}(x, b)) - (r^*(x, a) - r^*(x, b)))^2 \right] \leq \frac{32R_{\max}^2 e^{4R_{\max}} \log(|\Pi|/\delta)}{n}.$$

Lemma K.3. *Under the event in Lemma K.2, we have that for all $\pi : \mathcal{X} \rightarrow \Delta(\mathcal{A})$,*

$$\mathbb{E}_{\pi, \pi_{\text{ref}}} [|(\widehat{r}(x, a) - \widehat{r}(x, b)) - (r^*(x, a) - r^*(x, b))|] \leq 2\sqrt{\mathcal{C}_\eta^\pi \varepsilon_{\text{stat}}^2} + 2\mathcal{C}_\eta^\pi R_{\max} \eta,$$

where $\varepsilon_{\text{stat}}^2$ is defined in Lemma K.2.

Regret decomposition. Equipped with these concentration and estimation error bounds, we now bound the regret of Algorithm 2 using a pessimism-based analysis similar to the proof of Theorem H.1. Condition on the events in Eq. (56), Lemma K.1, and Lemma K.2, which hold together with probability at least $1 - 3\delta'$. We decompose the regret of $\widehat{\pi}$ using $\widehat{J}_{\beta,\eta}^\pi$, then leverage the inequalities in Eq. (57) and Eq. (58):

$$\begin{aligned} J(\pi^*) - J(\widehat{\pi}) &= J(\pi^*) - \widehat{J}_{\beta,\eta}^\pi(\pi^*) + \widehat{J}_{\beta,\eta}^\pi(\pi^*) - J(\widehat{\pi}) \\ &\leq J(\pi^*) - \widehat{J}_{\beta,\eta}^\pi(\pi^*) + \widehat{J}_{\beta,\eta}^\pi(\widehat{\pi}) - J(\widehat{\pi}) \end{aligned}$$

$$\begin{aligned}
&\leq J(\pi^*) - \mathbb{E}_{\pi^*}[\widehat{r}(x, a)] + \frac{3\beta\mathcal{C}_\eta^{\pi^*}}{2} + 4\beta \min\{\mathcal{C}_\infty^{\pi^*}, \eta^{-1}\}\varepsilon_x^2 + R_{\max}\varepsilon_x \\
&\quad + \mathbb{E}_{\widehat{\pi}}[\widehat{r}(x, a)] - \frac{\beta\mathcal{C}_\eta^{\widehat{\pi}}}{2} + 4\beta \min\{\mathcal{C}_\infty^{\widehat{\pi}}, \eta^{-1}\}\varepsilon_x^2 + R_{\max}\varepsilon_x - J(\widehat{\pi}) \\
&= \mathbb{E}_{\pi^*, \pi_{\text{ref}}}[\Delta^*(x, a, b) - \widehat{\Delta}(x, a, b)] + \frac{3\beta\mathcal{C}_\eta^{\pi^*}}{2} + \mathbb{E}_{\widehat{\pi}, \pi_{\text{ref}}}[\widehat{\Delta}(x, a, b) - \Delta^*(x, a, b)] - \frac{\beta\mathcal{C}_\eta^{\widehat{\pi}}}{2} \\
&\quad + 4\beta\varepsilon_x^2 \left(\min\{\mathcal{C}_\infty^{\pi^*}, \eta^{-1}\} + \min\{\mathcal{C}_\infty^{\widehat{\pi}}, \eta^{-1}\} \right) + 2R_{\max}\varepsilon_x.
\end{aligned}$$

In the last line above, we have introduced the notation $\Delta^*(x, a, b) = r^*(x, a) - r^*(x, b)$ and $\widehat{\Delta}(x, a, b) = \widehat{r}(x, a) - \widehat{r}(x, b)$, and centered the returns. Next, applying [Lemma K.3](#) to bound the reward estimation error above, we have

$$\begin{aligned}
J(\pi^*) - J(\widehat{\pi}) &\leq 2\sqrt{\mathcal{C}_\eta^{\pi^*}\varepsilon_{\text{stat}}^2} + 2\eta R_{\max}\mathcal{C}_\eta^{\pi^*} + \frac{3\beta\mathcal{C}_\eta^{\pi^*}}{2} \\
&\quad + 2\sqrt{\mathcal{C}_\eta^{\widehat{\pi}}\varepsilon_{\text{stat}}^2} + 2\eta R_{\max}\mathcal{C}_\eta^{\widehat{\pi}} - \frac{\beta\mathcal{C}_\eta^{\widehat{\pi}}}{2} \\
&\quad + 4\beta\varepsilon_x^2 \left(\min\{\mathcal{C}_\infty^{\pi^*}, \eta^{-1}\} + \min\{\mathcal{C}_\infty^{\widehat{\pi}}, \eta^{-1}\} \right) + 2R_{\max}\varepsilon_x.
\end{aligned}$$

Applying the AM-GM inequality to $2\sqrt{\mathcal{C}_\eta^{\widehat{\pi}}\varepsilon_{\text{stat}}^2}$ for $\eta \in \left[0, \frac{\beta}{4R_{\max}}\right]$, we have

$$\begin{aligned}
2\sqrt{\mathcal{C}_\eta^{\widehat{\pi}}\varepsilon_{\text{stat}}^2} &= \sqrt{(\beta - 4\eta R_{\max})\mathcal{C}_\eta^{\widehat{\pi}} \cdot \frac{4\varepsilon_{\text{stat}}^2}{(\beta - 4\eta R_{\max})}} \\
&\leq \frac{\beta\mathcal{C}_\eta^{\widehat{\pi}}}{2} - 2\eta R_{\max}\mathcal{C}_\eta^{\widehat{\pi}} + \frac{2\varepsilon_{\text{stat}}^2}{\beta - 4\eta R_{\max}} \\
&\leq \frac{\beta\mathcal{C}_\eta^{\widehat{\pi}}}{2} - 2\eta R_{\max}\mathcal{C}_\eta^{\widehat{\pi}} + \frac{4\varepsilon_{\text{stat}}^2}{\beta},
\end{aligned}$$

where in the last line we use the fact that $\eta \leq \frac{\beta}{8R_{\max}}$ so $4\eta R_{\max} \leq \frac{\beta}{2}$. Then plugging this back into our regret decomposition cancels out the $\mathcal{C}_\eta^{\widehat{\pi}}$ terms to give

$$\begin{aligned}
J(\pi^*) - J(\widehat{\pi}) &\leq 2\sqrt{\mathcal{C}_\eta^{\pi^*}\varepsilon_{\text{stat}}^2} + 2\eta R_{\max}\mathcal{C}_\eta^{\pi^*} + \frac{3\beta\mathcal{C}_\eta^{\pi^*}}{2} + \frac{4\varepsilon_{\text{stat}}^2}{\beta} \\
&\quad + 4\beta\varepsilon_x^2 \left(\min\{\mathcal{C}_\infty^{\pi^*}, \eta^{-1}\} + \min\{\mathcal{C}_\infty^{\widehat{\pi}}, \eta^{-1}\} \right) + 2R_{\max}\varepsilon_x \\
&\leq 2\sqrt{\mathcal{C}_\eta^{\pi^*}\varepsilon_{\text{stat}}^2} + 2\beta\mathcal{C}_\eta^{\pi^*} + \frac{4\varepsilon_{\text{stat}}^2}{\beta} \\
&\quad + 4\beta\varepsilon_x^2 \left(\min\{\mathcal{C}_\infty^{\pi^*}, \eta^{-1}\} + \min\{\mathcal{C}_\infty^{\widehat{\pi}}, \eta^{-1}\} \right) + 2R_{\max}\varepsilon_x,
\end{aligned}$$

where in the last line we consolidate $\mathcal{C}_\eta^{\pi^*}$ terms by again using $4\eta R_{\max} \leq \frac{\beta}{2}$. Plugging in $\delta' = \delta/3$ and the values for $\varepsilon_{\text{stat}}^2$ and ε_x results in the theorem statement. \square

Proof of Corollary K.1. When $\eta = \frac{\beta}{8R_{\max}}$, [Theorem K.1](#) states that

$$\begin{aligned}
J(\pi^*) - J(\widehat{\pi}) &\leq 2\sqrt{\mathcal{C}_\eta^{\pi^*}\varepsilon_{\text{stat}}^2} + 2\beta\mathcal{C}_\eta^{\pi^*} + \frac{4\varepsilon_{\text{stat}}^2}{\beta} + 4\beta\varepsilon_x^2 \cdot \left(\min\{\mathcal{C}_\infty^{\pi^*}, \eta^{-1}\} + \min\left\{\max_{\pi \in \Pi} \mathcal{C}_\infty^\pi, \eta^{-1}\right\} \right) + 2R_{\max}\varepsilon_x \\
&\leq 2\sqrt{\mathcal{C}_\eta^{\pi^*}\varepsilon_{\text{stat}}^2} + 2\beta\mathcal{C}_\eta^{\pi^*} + \frac{4\varepsilon_{\text{stat}}^2}{\beta} + 8\beta\varepsilon_x^2 \cdot \eta^{-1} + 2R_{\max}\varepsilon_x \\
&= 2\sqrt{\mathcal{C}_\eta^{\pi^*}\varepsilon_{\text{stat}}^2} + 2\beta\mathcal{C}_\eta^{\pi^*} + \frac{4\varepsilon_{\text{stat}}^2}{\beta} + 64R_{\max}\varepsilon_x^2 + 2R_{\max}\varepsilon_x.
\end{aligned}$$

2916 Setting $\beta = 2\sqrt{\frac{\varepsilon_{\text{stat}}^2}{\mathcal{C}^{\pi^*}}}$, we obtain

$$2917 \quad J(\pi^*) - J(\hat{\pi}) \leq 5\sqrt{\mathcal{C}_\eta^{\pi^*} \varepsilon_{\text{stat}}^2} + 64R_{\max}\varepsilon_x^2 + 2R_{\max}\varepsilon_x.$$

□

2922 **Proof of Corollary K.2.** When $\eta = 0$, [Theorem K.1](#) states that

$$2923 \quad J(\pi^*) - J(\hat{\pi}) \leq 2\sqrt{\mathcal{C}^{\pi^*} \varepsilon_{\text{stat}}^2} + 2\beta\mathcal{C}^{\pi^*} + \frac{4\varepsilon_{\text{stat}}^2}{\beta} + 4\beta\varepsilon_x^2 \cdot \left(\mathcal{C}_\infty^{\pi^*} + \max_{\pi \in \Pi} \mathcal{C}_\infty^\pi \right) + 2R_{\max}\varepsilon_x$$

2926 Setting $\beta = 2\sqrt{\frac{\varepsilon_{\text{stat}}^2}{\mathcal{C}^{\pi^*}}}$, we obtain

$$2927 \quad J(\pi^*) - J(\hat{\pi}) \leq 5\sqrt{\mathcal{C}^{\pi^*} \varepsilon_{\text{stat}}^2} + 8\varepsilon_{\text{stat}}\varepsilon_x^2 \cdot \left(\mathcal{C}_\infty^{\pi^*} + \max_{\pi \in \Pi} \mathcal{C}_\infty^\pi \right) + 2R_{\max}\varepsilon_x.$$

□

2933 **Proof of Lemma K.2.** We use similar reasoning and notation to the proof of [Lemma H.1](#). Since $r^* \in \mathcal{R}$ under [Assumption C.1](#), [Lemma F.1](#) guarantees that with probability at least $1 - \delta$ we have

$$2934 \quad \mathbb{E}_{\pi_{\text{ref}}, \pi_{\text{ref}}} [D_{\text{H}}^2(P_{\hat{r}}(\cdot | x, a, b), P_{r^*}(\cdot | x, a, b))] \leq \frac{2\log(|\mathcal{R}|/\delta)}{n}.$$

2939 Since $|r(x, a) - r(x, b)| \leq R_{\max}$ for all $r \in \mathcal{R}$ under [Assumption C.1](#), we then apply [Lemma H.5](#) with $R = V = R_{\max}$.

$$2940 \quad \begin{aligned} & \mathbb{E}_{\pi_{\text{ref}}, \pi_{\text{ref}}} \left[(\hat{r}(x, a) - \hat{r}(x, b) - (r^*(x, a) - r^*(x, b)))^2 \right] \\ & \leq 16e^{4R_{\max}} R_{\max}^2 \cdot \mathbb{E}_{\pi_{\text{ref}}, \pi_{\text{ref}}} [D_{\text{H}}^2(P_{\hat{r}}(\cdot | x, a, b), P_{r^*}(\cdot | x, a, b))] \\ & \leq 32e^{4R_{\max}} R_{\max}^2 \cdot \frac{\log(|\mathcal{R}|/\delta)}{n}. \end{aligned}$$

□

2948 **Proof of Lemma K.3.** Abbreviate $\Delta^*(x, a, b) = r^*(x, a) - r^*(x, b)$, and $\hat{\Delta}(x, a, b) = \hat{r}(x, a) - \hat{r}(x, b)$. For a pair of policies π, π' and $p \geq 1$, we define the norm $\|\cdot\|_{p, \pi \times \pi'} := (\mathbb{E}_{\rho, a \sim \pi, b \sim \pi'} [|\cdot|^p])^{1/p}$, so that $\mathbb{E}_{\pi, \pi_{\text{ref}}} [\|\Delta^*(x, a, b) - \hat{\Delta}(x, a, b)\|] = \|\Delta^* - \hat{\Delta}\|_{1, \pi \times \pi_{\text{ref}}}$. Then via Cauchy-Schwarz,

$$2953 \quad \begin{aligned} & \|\Delta^* - \hat{\Delta}\|_{1, \pi \times \pi_{\text{ref}}} \leq \sqrt{\mathbb{E}_\rho \left[\sum_{a, b} \frac{\pi^2(a | x) \pi_{\text{ref}}^2(b | x)}{(\pi_{\text{ref}}(a | x) + \eta\pi(a | x)) \pi_{\text{ref}}(b | x)} \right]} \\ & \quad \cdot \sqrt{\mathbb{E}_\rho \left[\sum_{a, b} (\pi_{\text{ref}}(a | x) + \eta\pi(a | x)) \pi_{\text{ref}}(b | x) (\Delta^*(x, a, b) - \hat{\Delta}(x, a, b))^2 \right]} \\ & = \sqrt{\mathcal{C}_\eta^\pi \cdot \left(\|\Delta^* - \hat{\Delta}\|_{2, \pi_{\text{ref}} \times \pi_{\text{ref}}}^2 + \eta \|\Delta^* - \hat{\Delta}\|_{2, \pi \times \pi_{\text{ref}}}^2 \right)} \\ & \leq \sqrt{\mathcal{C}_\eta^\pi \cdot \|\Delta^* - \hat{\Delta}\|_{2, \pi_{\text{ref}} \times \pi_{\text{ref}}}^2} + \sqrt{2\eta R_{\max} \mathcal{C}_\eta^\pi \cdot \|\Delta^* - \hat{\Delta}\|_{1, \pi \times \pi_{\text{ref}}}}. \end{aligned}$$

2967 Applying the AM-GM inequality to the second term, we obtain

$$2968 \quad \|\Delta^* - \hat{\Delta}\|_{1, \pi \times \pi_{\text{ref}}} \leq \sqrt{\mathcal{C}_\eta^\pi \cdot \|\Delta^* - \hat{\Delta}\|_{2, \pi_{\text{ref}} \times \pi_{\text{ref}}}^2} + \eta R_{\max} \mathcal{C}_\eta^\pi + \frac{1}{2} \|\Delta^* - \hat{\Delta}\|_{1, \pi \times \pi_{\text{ref}}}.$$

2970
2971
2972
2973
2974
2975
2976
2977
2978
2979
2980
2981
2982
2983
2984
2985
2986
2987
2988
2989
2990
2991
2992
2993
2994
2995
2996
2997
2998
2999
3000
3001
3002
3003
3004
3005
3006
3007
3008
3009
3010
3011
3012
3013
3014
3015
3016
3017
3018
3019
3020
3021
3022
3023

Rearranging,

$$\left\| \Delta^* - \widehat{\Delta} \right\|_{1, \pi \times \pi_{\text{ref}}} \leq 2\sqrt{C_{\eta}^{\pi} \cdot \left\| \Delta^* - \widehat{\Delta} \right\|_{2, \pi_{\text{ref}} \times \pi_{\text{ref}}}^2} + 2\eta R_{\max} C_{\eta}^{\pi}.$$

□