

000 001 002 003 004 005 ROBUST OPTIMIZATION FOR MITIGATING REWARD 006 HACKING WITH CORRELATED PROXIES 007 008 009

010 **Anonymous authors**
011 Paper under double-blind review
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053

ABSTRACT

Designing robust reinforcement learning (RL) agents in the presence of imperfect reward signals remains a core challenge. In practice, agents are often trained with proxy rewards that only approximate the true objective, leaving them vulnerable to reward hacking, where high proxy returns arise from unintended or exploitative behaviors. Recent work formalizes this issue using r -correlation between proxy and true rewards, but existing methods like occupancy-regularized policy optimization (ORPO) optimize against a fixed proxy and do not provide strong guarantees against broader classes of correlated proxies. In this work, we formulate reward hacking as a robust policy optimization problem over the space of all r -correlated proxy rewards. We derive a tractable max-min formulation, where the agent maximizes performance under the worst-case proxy consistent with the correlation constraint. We further show that when the reward is a linear function of known features, our approach can be adapted to incorporate this prior knowledge, yielding both improved policies and interpretable worst-case rewards. Experiments across several environments show that our algorithms consistently outperform ORPO in worst-case returns, and offer improved robustness and stability across different levels of proxy-true reward correlation. These results show that our approach provides both robustness and transparency in settings where reward design is inherently uncertain.

1 INTRODUCTION

Real-world reinforcement learning (RL) systems often struggle with reward specification: it is notoriously difficult to craft a reward function that perfectly captures the intended goals in all scenarios (Amodei et al., 2016; Ibarz et al., 2018; Stray et al., 2024). In practice, designers rely on proxy rewards that approximate the true objective (Tien et al., 2023). However, agents optimizing these imperfect proxies can lead to unintended exploitative behaviors, achieving high proxy returns while yielding poor true outcomes, a phenomenon known as reward hacking (Leike et al., 2017; Everitt et al., 2017; 2021; Koch et al., 2021). Such reward hacking behaviors are not merely hypothetical; they have led to undesirable or even catastrophic consequences in safety-critical settings (e.g., autonomous driving) (Krakovna et al., 2018; Knox et al., 2023) and are alarmingly common in real-world deployments (Kleinberg et al., 2024; Franchi et al., 2023; Milli et al., 2021; Obermeyer et al., 2019). Beyond reward hacking, interpretability and transparency of RL policies are increasingly recognized as critical requirements for real-world acceptance (Vouros, 2022; Puiutta & Veith, 2020; Iyer et al., 2018). Policymakers and practitioners in safety-critical domains require systems not only to be robust but also interpretable; they must understand which specific decision-making criteria lead to undesirable outcomes to effectively mitigate risks and ensure compliance with safety regulations (Rudin, 2019; Druce et al., 2021; Doshi-Velez & Kim, 2017). These challenges highlight the need for RL algorithms to address two fundamental challenges: robustness to uncertain or poorly-specified rewards, and interpretability to facilitate oversight and compliance by human stakeholders, especially in high-stakes, real-world environments like traffic control (Vinitsky et al., 2018), healthcare decision-making (Fox et al., 2020; Holzinger et al., 2017), and pandemic response strategies (Kompella et al., 2020).

Recent work has begun to formalize reward hacking and develop principled mitigations. (Laidlaw et al., 2025) define a proxy reward to be r -*correlated* with the true reward if it maintains a correlation coefficient $r > 0$ on state-action pairs encountered by a certain reference policy. Notably, their

054 definition permits the proxy and true reward to diverge arbitrarily in parts of the state-action space
 055 not visited by the reference policy, precisely the regions an RL agent might exploit under intensive
 056 optimization. Using this framework, reward hacking is formalized as the situation in which optimiz-
 057 ing an r -correlated proxy yields a policy with lower true reward than that of the reference policy.
 058 Building on this definition, they propose Occupancy-Regularized Policy Optimization (ORPO) as a
 059 mitigation strategy. ORPO augments the standard RL objective with a regularization term that pe-
 060 nalizes deviations between the learned policy’s occupancy measure and that of the reference policy.

061 Despite significant progress, existing solutions to reward hacking show several limitations. First,
 062 their effectiveness relies heavily on the choice of the specific proxy reward. However, designing
 063 perfect proxies is challenging, and in real-world scenarios, reward proxies are often derived heuris-
 064 tically or empirically from noisy or limited data (Jeon et al., 2020; Sadigh et al., 2017), leading to
 065 uncertainty or variability in the exact correlation with true rewards. Therefore, robustness to varia-
 066 tions in proxy rewards is crucial for dependable deployment. While the regularization method used
 067 by ORPO provides a lower bound on improvement in true reward, its guarantee on the worst-case
 068 performance against an adversarially chosen proxy is weak. Second, current methods like ORPO
 069 typically treat a reward function as a black box and learn a complex policy with no easily inter-
 070 pretable structure, making it hard to understand why the resulting policy avoids reward hacking
 071 or to trust its behavior in novel situations. Further, they cannot be easily adapted to incorporate
 072 prior knowledge of the true reward. These shortcomings underscore the need for a more robust and
 073 transparent approach to reward hacking in RL.

074 In this work, we formalize reward hacking as a robust RL problem under proxy reward uncertainty
 075 and develop new algorithms to address the above gaps. The key idea is to optimize against an
 076 adversarial proxy reward rather than trusting a single proxy. We assume the true reward could be any
 077 function that remains r -correlated with the proxy, and we train the agent to perform well against the
 078 worst-case such proxy. This approach explicitly accounts for uncertainty in proxy design and guards
 079 against unintended exploitative behaviors. Concretely, we propose a max-min formulation in which
 080 the policy chooses its strategy to maximize its guaranteed true return while an adversary minimizes
 081 the true return by selecting a reward function from the set of all r -correlated proxies. By solving
 082 this problem, the agent learns a policy that is robust to all plausible deviations of the proxy reward
 083 within the correlation bound. We derive a closed-form solution for the adversary’s worst-case reward
 084 assignment given any candidate policy, which allows efficient evaluation of the inner minimization
 085 and provides insight into how proxy reward flaws are most damaging. Building on this result, we
 086 introduce a practical algorithm for Max-Min Policy Optimization that iteratively updates the policy
 087 against this worst-case reward signal. Moreover, to improve the tractability and transparency of
 088 the inner optimization, we introduce a Linear Max-Min variant of our method. In this variant, we
 089 assume the true reward lies in a class of linear functions over known features, [an assumption that has](#)
 090 [been extensively studied in prior work on successor representations and successor features](#) (Dayan,
 091 1993; Barreto et al., 2017; 2018), and which allows us to characterize the worst-case proxy reward
 092 as a sparse linear combination of those features. While the policy itself remains parameterized by
 093 general neural networks, the learned worst-case reward function becomes interpretable in terms of
 094 its feature weights. This provides insight into which aspects of the proxy reward space the policy is
 095 robust to or vulnerable against, making it valuable for applications where understanding the failure
 096 modes of the reward design is important.

097 Finally, we empirically evaluate the proposed approaches on several challenging environments.
 098 Across all domains, our Max-Min and Linear Max-Min policies outperform ORPO in terms of
 099 worst-case reward, indicating substantially improved robustness. Moreover, under a large range
 100 of proxy-true correlation scenarios, our methods exhibit higher average reward and lower variance
 101 compared to ORPO, meaning the performance of our policies remains more consistent and reliable.
 102 These findings demonstrate the practical significance of our robust formulation, paving the way for
 103 safer and more trustworthy RL deployment in real-world applications.

104 Our main contributions can be summarized as follows: 1) We propose a novel robust RL formulation
 105 that explicitly models reward hacking as a max-min optimization problem over proxy rewards con-
 106 strained by correlation with the true rewards. 2) We develop a practical algorithm for the max-min
 107 problem, which is further extended to linear rewards with improved robustness and interpretabil-
 108 ity. 3) We provide a theoretical convergence guarantee for the max-min objective with a sample-
 109 complexity bound for the occupancy estimation. We also show that accurate occupancy estimation

108 is pivotal for robustness. 4) Experiment results demonstrate improved robustness and worst-case
 109 rewards across five real-world inspired reward hacking environments.
 110

112 2 PRELIMINARIES

114 **Reinforcement Learning.** A reinforcement learning (RL) problem can be formulated as an infinite-
 115 horizon Markov Decision Process (MDP) defined by the tuple $(\mathcal{S}, \mathcal{A}, p, \mu_0, R, \gamma)$, where \mathcal{S} and
 116 \mathcal{A} denote the state and action spaces, $p(s' | s, a)$ is the transition probability from state s to s'
 117 given action a , μ_0 is the initial state distribution and $\gamma \in [0, 1]$ is the discounted factor. The agent
 118 interacts with the environment over discrete time steps $t = 0, 1, 2, \dots$. At each time step, it selects
 119 an action $a_t \in \mathcal{A}$ based on the current state $s_t \in \mathcal{S}$ according to a policy $\pi(a | s)$, which defines a
 120 distribution over actions conditioned on the state. Upon taking action a_t , the agent receives a reward
 121 $R(s_t, a_t) \in \mathbb{R}$ and transitions to the next state s_{t+1} according to $p(s_{t+1} | s_t, a_t)$. The goal of the
 122 agent is to maximize the expected cumulative discounted return:
 123

$$J(\pi, R) = (1 - \gamma) \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t) \right], \quad (1)$$

124 where $\gamma \in [0, 1]$ is the discount factor, and the expectation is taken over trajectories generated by
 125 following policy π . We define the *state-action occupancy measure* μ_π of a policy π as: $\mu_\pi(s, a) =$
 126 $(1 - \gamma) \mathbb{E}_\pi [\sum_{t=0}^{\infty} \gamma^t \mathbb{I}\{s_t = s, a_t = a\}]$, which represents the discounted visitation frequency of
 127 each state-action pair under policy π . Using the occupancy measure, the return can be equivalently
 128 expressed as: $J(\pi, R) = \mathbb{E}_{(s, a) \sim \mu_\pi} [R(s, a)]$.
 129

130 **Correlated Proxies and Reward Hacking.** Below we give an overview of the recently proposed
 131 r -correlated proxy framework proposed in (Laidlaw et al., 2025) for detecting and mitigating reward
 132 hacking, which our work is built upon. A detailed discussion of related work on reward hacking and
 133 robust RL is given in Appendix D. In particular, they consider a setting where the agent is given a
 134 reference policy π_{ref} and a proxy reward R_{proxy} , while the true reward is hidden. They further assume
 135 that the proxy reward is *r-correlated* with the true reward under the reference policy, that is:
 136

$$\mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\left(\frac{R_{\text{proxy}} - J(\pi_{\text{ref}}, R_{\text{proxy}})}{\sigma_{R_{\text{proxy}}}} \right) \left(\frac{R_{\text{true}} - J(\pi_{\text{ref}}, R_{\text{true}})}{\sigma_{R_{\text{true}}}} \right) \right] = r, \quad (2)$$

137 where $\sigma_{R_{\text{proxy}}}^2 = \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [(R_{\text{proxy}} - J(\pi_{\text{ref}}, R_{\text{proxy}}))^2]$ and $\sigma_{R_{\text{true}}}^2 = \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [(R_{\text{true}} - J(\pi_{\text{ref}}, R_{\text{true}}))^2]$
 138 are the variances of the proxy and true rewards, respectively, under the reference policy. Reward
 139 hacking is said to occur when a policy π optimized for an r -correlated proxy reward achieves lower
 140 true reward than the reference policy: $J(\pi, R_{\text{true}}) < J(\pi_{\text{ref}}, R_{\text{true}})$. To mitigate reward hacking,
 141 they propose Occupancy-Regularized Policy Optimization (ORPO) to optimize a regularized policy
 142 objective given below, which is shown to provide a lower bound on improvement in true reward:
 143

$$\max_{\pi} J(\pi, R_{\text{proxy}}) - \lambda \sqrt{\chi^2(\mu_\pi \| \mu_{\pi_{\text{ref}}})}, \quad (3)$$

144 where $\chi^2(\mu_\pi \| \mu_{\pi_{\text{ref}}})$ denotes the χ^2 -squared divergence between the occupancy measures of π and
 145 π_{ref} , and the regularization strength $\lambda = \sigma_{R_{\text{proxy}}} \sqrt{1 - r^2}$. This encourages the learned policy to stay
 146 close to the reference distribution when the proxy reward is weakly correlated with the true reward.
 147

154 3 METHOD

155 In this section, we discuss our robust policy optimization approach for mitigating reward hacking.
 156 In contrast to regularization-based methods such as ORPO, we consider a max-min formulation
 157 that identifies a robust policy with respect to the worst-case reward across all reward functions that
 158 are r -correlated with the proxy reward. We further extend our framework to settings where the
 159 reward function is a linear combination of known features with unknown weights. Our approach
 160 effectively leverages this structural information, when known a priori, to improve both robustness
 161 and interpretability, a task that is particularly challenging for regularization-based techniques.
 162

162 3.1 MAX-MIN POLICY OPTIMIZATION
163

164 Similar to ORPO, we assume that the agent is given a proxy reward R_{proxy} and a reference policy
165 π_{ref} , while the true reward is hidden. Rather than regularizing the policy under a fixed proxy reward,
166 we consider the *entire space of rewards* $\mathcal{R}_{\text{corr}}$ that satisfy the correlation constraint with respect to a
167 known proxy reward, as defined in Equation 4:

$$168 \quad \mathcal{R}_{\text{corr}} = \left\{ R : (s, a) \rightarrow \mathbb{R} \left| \mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{R - M}{V} \cdot R_{\text{proxy}} \right] = r, J(\pi_{\text{ref}}, R) = M, \sigma_R^2 = V^2 \right. \right\}. \quad (4)$$

171 M and V denote the fixed mean and standard deviation of the reward function R under the refer-
172 ence policy π_{ref} . For simplicity, we define R_{proxy} to be the normalized proxy reward $R_{\text{proxy}}(s, a) :=$
173 $\frac{\tilde{R}_{\text{proxy}}(s, a) - J(\pi_{\text{ref}}, \tilde{R}_{\text{proxy}})}{\sigma_{\tilde{R}_{\text{proxy}}}}$, where \tilde{R}_{proxy} is the original (unnormalized) proxy reward. After normaliza-
174 tion, we have $J(\pi_{\text{ref}}, R_{\text{proxy}}) = 0$ and $\text{Var}_{\mu_{\pi_{\text{ref}}}}(R_{\text{proxy}}) = 1$, which simplifies the correlation con-
175 straint in Equation 4. The hyperparameter r controls the degree of alignment between the proxy and
176 true reward. It allows us to interpolate between strong robustness (small r) and high proxy fidelity
177 (large r), enabling a principled robustness-accuracy trade-off. We remark that it is without loss of
178 generality to consider fixed M and V , which we will further elaborate on later.

179 We propose a *worst-case optimization framework* where the policy is trained to maximize expected
180 performance under the least favorable reward within $\mathcal{R}_{\text{corr}}$. Assuming that the true reward lies some-
181 where within this set, this approach improves robustness by ensuring that the policy does not overfit
182 to any single optimistic interpretation of the proxy reward. Formally, the objective becomes:

$$183 \quad \max_{\pi} \min_{R \in \mathcal{R}_{\text{corr}}} J(\pi, R) = \max_{\pi} \min_{R \in \mathcal{R}_{\text{corr}}} \mathbb{E}_{(s, a) \sim \mu_{\pi}} [R(s, a)]. \quad (5)$$

185 However, a challenge arises: the objective $\mathbb{E}_{\mu_{\pi}} [R(s, a)]$ depends on the state-action occupancy
186 μ_{π} , whereas the constraints defining $\mathcal{R}_{\text{corr}}$ are expressed in terms of $\mu_{\pi_{\text{ref}}}$. This mismatch com-
187 plicates direct optimization. To resolve this, we apply a *change-of-measure* technique (Hu & Hong,
188 2013; Lam, 2016) to rewrite the expectation under $\mu_{\pi_{\text{ref}}}$. Specifically, let $L(s, a)$ denote the Radon-
189 Nikodym derivative: $L(s, a) = \frac{\mu_{\pi}(s, a)}{\mu_{\pi_{\text{ref}}}(s, a)}$. By definition, $L(s, a) \geq 0$ and $\mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L(s, a)] =$
190 1. Applying the change-of-measure formula, we can express the return as: $\mathbb{E}_{\mu_{\pi}} [R(s, a)] =$
191 $\int_{\mathcal{S} \times \mathcal{A}} \mu_{\pi}(s, a) R(s, a) d(s, a) = \int_{\mathcal{S} \times \mathcal{A}} \mu_{\pi_{\text{ref}}}(s, a) \frac{\mu_{\pi}(s, a)}{\mu_{\pi_{\text{ref}}}(s, a)} R(s, a) d(s, a) = \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L(s, a) R(s, a)]$.
192 Thus, both the objective and the constraints can be rewritten as expectations with respect to $\mu_{\pi_{\text{ref}}}$.

194 For notational simplicity, we will suppress variables (s, a) and write for example, L as $L(s, a)$.
195 Under this reparameterization, the inner minimization in Equation 5 can be reformulated as:

$$196 \quad \min_{R \in \mathcal{R}_{\text{corr}}} \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L \cdot R]. \quad (6)$$

198 Although the feasible set in Problem 6 is not convex due to the equality constraint on the variance, we
199 still derive an optimal solution using a Lagrangian formulation. Our approach leverages tools from
200 duality theory, commonly used in robust optimization (Delage & Ye, 2010; Goh & Sim, 2010). We
201 further justify the validity of our solution in Appendix E.2. Specifically, the Lagrangian functional
202 associated with this problem is defined as: $l_0(\lambda_1, \lambda_2, \lambda_3, R) = \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L \cdot R - \lambda_1 \frac{R - M}{V} \cdot R_{\text{proxy}} - \lambda_2 R -$
203 $\lambda_3 R^2] + \lambda_1 r + \lambda_2 M + \lambda_3 (M^2 + V^2)$, where $\lambda_1, \lambda_2, \lambda_3$ are the Lagrange multipliers corresponding to
204 the correlation constraint, mean constraint, and variance constraint, respectively. Then the original
205 problem in Equation 6 is equivalent to the following problem:

$$206 \quad \max_{\lambda_1, \lambda_2, \lambda_3} \min_{R \in \mathcal{R}_{\text{corr}}} l_0(\lambda_1, \lambda_2, \lambda_3, R). \quad (7)$$

208 We now solve the inner minimization problem in Equation 7 by finding the optimal R for fixed dual
209 variables $(\lambda_1, \lambda_2, \lambda_3)$. Taking the functional derivative of the Lagrangian l_0 with respect to $R(s, a)$
210 gives: $\frac{\partial l_0}{\partial R} = \mu_{\pi_{\text{ref}}}(s, a) [(L - \lambda_1 \frac{R_{\text{proxy}}}{V} - \lambda_2) - 2\lambda_3 R]$. When $\mu_{\pi_{\text{ref}}}(s, a) > 0$, setting the derivative
211 of the Lagrangian to zero yields the optimal adversarial reward function:

$$212 \quad R^*(s, a) = \frac{L(s, a) - \lambda_1 \frac{R_{\text{proxy}}}{V} - \lambda_2}{2\lambda_3}. \quad (8)$$

215 However, for state-action pairs where $\mu_{\pi_{\text{ref}}}(s, a) = 0$, i.e., those not visited under the reference
policy, the correlation and moment constraints become vacuous. In these regions, the adversarial

reward $R^*(s, a)$ can be driven arbitrarily poor, reflecting that no constraint prevents the adversary from assigning highly penalizing values to rarely visited or unobserved state-action pairs. Nevertheless, consider the case where $\mu_{\pi_{\text{ref}}}(s, a) > 0$, we can substitute the optimal R^* from Equation 8 into the Lagrangian l_0 and get the dual objective. After some process detailed in Appendix E.1, we get the optimal solution to problem (6), so the original max-min problem (5) reduces to:

$$\max_{\pi} r \cdot V \cdot \mathbb{E}_{\mu_{\pi}}[R_{\text{proxy}}] - V \cdot \sqrt{1 - r^2} \sqrt{\chi^2(\mu_{\pi} \parallel \mu_{\pi_{\text{ref}}}) - \mathbb{E}_{\mu_{\pi}}^2[R_{\text{proxy}}]} + M. \quad (9)$$

Thus, the final policy optimization objective becomes maximizing the proxy reward, regularized by a penalty that depends on the distributional shift between μ_{π} and $\mu_{\pi_{\text{ref}}}$ and the expectation of the current policy under proxy reward $\mathbb{E}_{\mu_{\pi}}[R_{\text{proxy}}]$, and the correlation strength r . We observe that the constants M and V do not affect the optimal policy: while they influence the absolute value of the worst-case reward for a given policy π , they only apply a linear transformation (scaling by V and shifting by M) and do not change the relative ordering of policies. Therefore, for simplicity, we set $V = 1$ and $M = 0$ in our implementation. This also provides a fair way to compare the worst-case rewards of different policies. Notice that the optimization objective in Equation 9 closely resembles the ORPO objective proposed in Equation 3. However, there are two key differences: (1) our regularization strength is $\frac{\sqrt{1-r^2}}{r}$ instead of $\sigma_{R_{\text{proxy}}} \sqrt{1-r^2}$, and (2) our penalty term is $\chi^2(\mu_{\pi} \parallel \mu_{\pi_{\text{ref}}}) - \mathbb{E}_{\mu_{\pi}}^2[R_{\text{proxy}}]$ rather than simply $\chi^2(\mu_{\pi} \parallel \mu_{\pi_{\text{ref}}})$. The proof that $\chi^2(\mu_{\pi} \parallel \mu_{\pi_{\text{ref}}}) - \mathbb{E}_{\mu_{\pi}}^2[R_{\text{proxy}}] \geq 0$ holds can be found in Appendix E.3. A detailed comparison between our policy gradient and that of ORPO is provided in Appendix E.8.

To further illustrate how our framework in Equation 9 helps prevent reward hacking, i.e., how optimizing a proxy reward can translate into an improvement in the true reward over the reference policy, as discussed in Section 2, we formalize the following theorem:

Theorem 1. Suppose that the true reward function R_{true} lies in the correlation-constrained uncertainty set $\mathcal{R}_{\text{corr}}$. Then, for any policy π such that $\mu_{\pi} \ll \mu_{\pi_{\text{ref}}}$ (i.e., $\mu_{\pi_{\text{ref}}}(s, a) = 0 \Rightarrow \mu_{\pi}(s, a) = 0$), we have

$$J(\pi, R_{\text{true}}) - J(\pi_{\text{ref}}, R_{\text{true}}) \geq r \cdot \mathbb{E}_{\mu_{\pi}}[R_{\text{proxy}}] - \sqrt{1 - r^2} \sqrt{\chi^2(\mu_{\pi} \parallel \mu_{\pi_{\text{ref}}}) - \mathbb{E}_{\mu_{\pi}}^2[R_{\text{proxy}}]}.$$

Proof can be found in Appendix E.9. From Theorem 1, we see that our objective optimizes a pessimistic lower bound on the true improvement over the reference policy. By Definition 4.2 in (Laidlaw et al., 2025), reward hacking occurs when $J(\pi, R_{\text{true}}) < J(\pi_{\text{ref}}, R_{\text{true}})$. Although this is precisely the quantity we would like to maximize, we cannot do so directly because the true reward is unobserved, and therefore we must instead optimize the max-min objective in Equation 9. Theorem 1 shows that our objective is always lower than (but anchored to) the true improvement, which explains why our framework can promote robustness against potential reward hacking: improving our surrogate objective necessarily improves a conservative bound on $J(\pi, R_{\text{true}}) - J(\pi_{\text{ref}}, R_{\text{true}})$.

Remark: Our optimization problem in Equation 5 is standard in distributionally robust optimization (DRO). However, it remains underexplored in the context of RL, with only one relevant work that considers uncertainty sets based on the first and second moments of the reward distribution (Nguyen et al., 2022). While their formulation appears similar, their results are not directly applicable to our max-min framework, and we still need to explicitly solve our formulation. We also note that under certain assumptions, the ORPO objective (Equation 3) can be reinterpreted as a special case of the max-min formulation in (Nguyen et al., 2022) (Theorem 1), providing a complementary view of the connection between these approaches. Nevertheless, our optimization objective remains structurally different. Moreover, in the pessimism offline RL setting, where distribution shift is the central challenge, the χ^2 regularization together with maxmin formulation has also been explored (Zhan et al., 2022; Huang et al., 2024) from a perspective different from ours. However, frameworks such as χ PO (Huang et al., 2024) are not applicable in our setting because they require the regularizer to be f -divergence. The square-root term in Equation 9 does not satisfy this requirement.

3.2 STRUCTURED REWARD SPACES VIA FEATURE LINEARIZATION

A natural concern with worst-case optimization is *over-conservatism*: if the reward uncertainty set $\mathcal{R}_{\text{corr}}$ is too broad, the resulting policy may become overly cautious or deviate from realistic task objectives. Additionally, the learned worst-case rewards may themselves be implausible

270 or uninterpretable. To address these issues, we introduce *structure* into the reward space by as-
 271 suming that all rewards are *linear combinations of known features*, [an assumption that has been](#)
 272 [widely adopted in prior work](#) (Dayan, 1993; Barreto et al., 2017; 2018). Specifically, we assume:
 273 $R(s, a) = \boldsymbol{\theta}^\top \phi(s, a)$, where $\phi(s, a) = [\phi_1(s, a), \phi_2(s, a), \dots, \phi_M(s, a)]^\top \in \mathbb{R}^M$ denotes a vec-
 274 tor of M known or engineered feature functions, and $\boldsymbol{\theta} = [\theta_1, \theta_2, \dots, \theta_M]^\top \in \mathbb{R}^M$ represents
 275 the uncertain feature weights. The linearization yields two key benefits: 1) **Realism and Inter-
 276 pretability:** In many real-world tasks, reward functions are naturally approximated as linear com-
 277 binations over interpretable features. For example, in a traffic control environment, features might
 278 include total commute time, vehicle speed, acceleration, and inter-vehicle headway distances. 2)
 279 **Better-Constrained Robustness:** By restricting uncertainty to structured, feature-based rewards,
 280 the worst-case optimization problem becomes more grounded and avoids pathological, unrealistic
 281 reward functions.

282 In this section, we assume that the agent is aware of the set of features but not their true weights. We
 283 show that our robust optimization framework can be naturally extended to incorporate the structure
 284 in rewards to improve robustness. In our experiments, we further demonstrate that linear rewards
 285 help interpret a policy’s performance even when it is trained without such prior knowledge. Under
 286 our assumption, the uncertainty set reduces to the set of feature weights $\boldsymbol{\theta} \in \mathbb{R}^M$ satisfying:

$$287 \quad \mathcal{R}_{\text{corr}}^{\text{lin}} = \left\{ \boldsymbol{\theta} \in \mathbb{R}^M \mid \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [\boldsymbol{\theta}^\top \phi \cdot R_{\text{proxy}}] = r, \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [\boldsymbol{\theta}^\top \phi] = 0, \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [(\boldsymbol{\theta}^\top \phi)^2] = 1 \right\}. \quad (10)$$

289 To simplify the analysis, we assume without loss of generality that the worst-case reward $R(s, a) =$
 290 $\boldsymbol{\theta}^\top \phi(s, a)$ is normalized to have zero mean and unit variance under the reference policy π_{ref} . This
 291 corresponds to setting $M = 0$ and $V = 1$, which, as shown in our earlier derivation, does not
 292 affect the resulting optimal policy. As before, R_{proxy} denotes the normalized proxy reward under
 293 π_{ref} , satisfying $\mathbb{E}_{\mu_{\pi_{\text{ref}}}} [R_{\text{proxy}}] = 0$ and $\text{Var}_{\mu_{\pi_{\text{ref}}}} [R_{\text{proxy}}] = 1$.

294 We now derive the corresponding max-min optimization under the structured reward assumption:

$$296 \quad \max_{\pi} \min_{\boldsymbol{\theta} \in \mathcal{R}_{\text{corr}}^{\text{lin}}, \boldsymbol{\theta} \geq 0} \mathbb{E}_{(s, a) \sim \mu_{\pi}} [\boldsymbol{\theta}^\top \phi(s, a)]. \quad (11)$$

298 Similar to previous steps, we introduce the Radon-Nikodym derivative $L(s, a) = \frac{\mu_{\pi}(s, a)}{\mu_{\pi_{\text{ref}}}(s, a)}$, use a
 299 change-of-measure, and define the Lagrangian functional for the inner minimization in Equation 11
 300 as: $l_1(\lambda_1, \lambda_2, \lambda_3, \boldsymbol{\theta}) = \boldsymbol{\theta}^\top \left(\sum_{(s, a)} u_{\lambda_1, \lambda_2}(s, a) \phi(s, a) \right) - \lambda_3 \boldsymbol{\theta}^\top Q \boldsymbol{\theta} + \lambda_1 r + \lambda_3$, where $u_{\lambda_1, \lambda_2} =$
 301 $\mu_{\pi} - \lambda_1 \mu_{\pi_{\text{ref}}} R_{\text{proxy}} - \lambda_2 \mu_{\pi_{\text{ref}}}$, $Q = \sum_{(s, a)} \mu_{\pi_{\text{ref}}}(s, a) \phi(s, a) \phi(s, a)^\top$. A detailed derivation can be
 302 found in Appendix E.4. Then solving the inner minimization over $\boldsymbol{\theta}$ in Equation 11 is equivalent to:
 303

$$304 \quad \max_{\lambda_1, \lambda_2, \lambda_3} \min_{\boldsymbol{\theta} \geq 0} l_1(\lambda_1, \lambda_2, \lambda_3, \boldsymbol{\theta}) = \boldsymbol{\theta}^\top \left(\sum u_{\lambda_1, \lambda_2} \phi \right) - \lambda_3 \boldsymbol{\theta}^\top Q \boldsymbol{\theta} + \lambda_1 r + \lambda_3. \quad (12)$$

306 Notice that $l_1(\lambda_1, \lambda_2, \lambda_3, \boldsymbol{\theta})$ is a convex quadratic function of $\boldsymbol{\theta}$ (assuming $\lambda_3 \leq 0$) subject to
 307 linear inequality constraints $\boldsymbol{\theta} \geq 0$. Thus, it is a standard convex quadratic program (QP) with
 308 non-negativity constraints (Boyd & Vandenberghe, 2004). However, it is not possible to derive a
 309 universal closed-form solution for the optimal $\boldsymbol{\theta}^*$ under arbitrary Q . To further simplify the problem
 310 and obtain a closed-form solution, we transform the feature vector ϕ into a whitened version $\tilde{\phi}$
 311 such that the matrix Q becomes the identity matrix I and we formally show this in Appendix E.5.
 312 Specifically, we perform a whitening transformation using the Cholesky decomposition (Boyd &
 313 Vandenberghe, 2004). Let $W = Q^{-\frac{1}{2}}$, $\tilde{\phi}(s, a) = W\phi(s, a)$, where $Q^{-\frac{1}{2}}$ denotes a matrix square
 314 root of Q^{-1} (which exists since Q is positive semi-definite and non-singular, which is detailed in
 315 Appendix E.5). Then the original problem in Equation 12 can be further simplified into:

$$317 \quad \max_{\lambda_1, \lambda_2, \lambda_3} \min_{\tilde{\boldsymbol{\theta}} \geq 0} l_1(\lambda_1, \lambda_2, \lambda_3, \tilde{\boldsymbol{\theta}}) = \tilde{\boldsymbol{\theta}}^\top \left(\sum_{(s, a)} u_{\lambda_1, \lambda_2}(s, a) \tilde{\phi}(s, a) \right) - \lambda_3 \tilde{\boldsymbol{\theta}}^\top \tilde{\boldsymbol{\theta}} + \lambda_1 r + \lambda_3. \quad (13)$$

320 where we now optimize over the parameter $\tilde{\boldsymbol{\theta}}$ using the transformed features $\tilde{\phi}$. For notational sim-
 321 plicity, we will drop the tilde and henceforth use ϕ to represent the whitened feature $\tilde{\phi}$, and $\boldsymbol{\theta}$ to
 322 represent $\tilde{\boldsymbol{\theta}}$. Then we can get a closed-form solution (we detail the steps in Appendix E.6) for optim-
 323 al $\boldsymbol{\theta}^*$ as: $\boldsymbol{\theta}^* = \max \left(0, -\frac{\sum_{(s, a)} u_{\lambda_1, \lambda_2}(s, a) \phi(s, a)}{2\lambda_3} \right)$, where the $\max(\cdot, 0)$ is applied elementwise.

324 Details for solving the outer maximization in Equation 13 can be found in Appendix E.7. After
 325 obtaining the optimal dual variables $(\lambda_1^*, \lambda_2^*, \lambda_3^*)$, we can substitute them back into the optimal θ^* to
 326 construct the worst-case reward, which is the optimal solution of the inner problem of Equation 11
 327 given π . Then we can solve the outer maximization over the policy π using standard RL algorithms.
 328

329 **ORPO with Linear Rewards.** While ORPO provides a general guarantee based on occupancy
 330 measure regularization, it does not exploit any structural assumptions about the reward function.
 331 In particular, even when the true reward is linear in a set of features, ORPO does not explicitly
 332 incorporate this structure into its policy optimization or theoretical analysis. While the lower bound
 333 (Theorem 5.1 in (Laidlaw et al., 2025)) continues to hold, it is unclear how to leverage this structure
 334 to obtain a tighter lower bound or to guide policy updates more effectively. This suggests a missed
 335 opportunity: by explicitly modeling the reward as a linear function, it becomes possible to derive
 336 stronger guarantees, interpret worst-case reward directions, and efficiently optimize against them.
 337 Our Linear Maxmin method fills this gap by parameterizing reward uncertainty directly in the space
 338 of reward weights, enabling both robustness and greater transparency.

339 3.3 OCCUPANCY ESTIMATION AND CONVERGENCE

340 A core step in both our algorithms and ORPO is to estimate the Radon-Nikodym derivative $L(s, a)$.
 341 To this end, following prior works (Laidlaw et al., 2025; Kang et al., 2018; Ho & Ermon, 2016), we
 342 fit a discriminator network $d_\phi(s, a)$ with $L_\phi(s, a) = \exp d_\phi(s, a)$. We learn ϕ by minimizing:
 343

$$344 \phi = \arg \min_{\phi} \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [\log(1 + e^{d_\phi(s, a)})] + \mathbb{E}_{\mu_\pi} [\log(1 + e^{-d_\phi(s, a)})]. \quad (14)$$

345 It is known that the optimal discriminator satisfies $d^*(s, a) = \log \frac{\mu_\pi(s, a)}{\mu_{\pi_{\text{ref}}}(s, a)}$ and we estimate $L_\phi(s, a)$
 346 as $\tilde{L}_\phi(s, a) = \exp \tilde{d}_\phi(s, a)$ with $\tilde{d}_\phi(s, a) \approx d^*(s, a)$. As discussed in Section 3.1, if the policy π
 347 visits state-action pairs that the reference policy π_{ref} rarely or never visits, the adversarial reward can
 348 be arbitrarily poor. In theory, the estimated $\tilde{L}(s, a)$ is expected to grow arbitrarily large in this case,
 349 which should discourage the learned policy from exploiting such regions. However, we observe
 350 empirically (Section 4.2) that the ORPO policy still visits some of these low-coverage regions under
 351 π_{ref} . This is because in the original ORPO implementation, the discriminator is not fully optimized
 352 during policy learning. Specifically, the discriminator receives only a small number of gradient
 353 updates per RL iteration, resulting in underfitting and inaccurate estimates of the Radon-Nikodym
 354 derivative $\tilde{L}(s, a)$. To address this, we substantially increase the number of gradient updates per
 355 iteration and carefully tune the learning rate. Our goal is to strike a practical balance between
 356 training time and discriminator quality, which we discuss in Appendix F.1. We further show that the following theorem, which establishes that the discriminator estimation achieves a sample complexity
 357 of $\mathcal{O}(n^{-1/4})$, where n denotes the sample size.
 358

359 **Theorem 2** (Occupancy ratio L_ϕ error bound). Under assumptions, let $\tilde{L} := e^{\tilde{d}}$ be the empirical
 360 estimation and $L^* = e^{d^*}$ be the true ratio. Then, with probability at least $1 - \delta$,
 361

$$362 \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [|\tilde{L} - L^*|] \leq C \left(\gamma' + \left(\frac{\log(M/\delta)}{n} \right)^{1/2} \right)^{1/2}.$$

363 where C , γ' and M are some constants.

364 The full argument is presented in Appendix G.1, where we adopt the optimistic cover notion (Definition 3)
 365 from (Huang et al., 2023) as a technical tool and establish the new concentration analysis
 366 as well as the resulting complexity bounds specific for estimating the loss in Equation 14.

367 To compute the final objective for our Max-Min policy in Equation 9, we estimate the χ^2 divergence,
 368 the normalized proxy reward R_{proxy} , and the first and second moments $\mathbb{E}_{\mu_\pi}[R_{\text{proxy}}]$ and $\mathbb{E}_{\mu_\pi}^2[R_{\text{proxy}}]$.
 369 These components together define the robust optimization objective used to update the policy. A
 370 simplified Max-Min policy optimization procedure is outlined in Algorithm 1. We provide detailed
 371 descriptions of each estimation step, as well as the complete algorithmic implementation for both
 372 Max-Min and Linear Max-Min in Appendices F.2 and F.3. We further obtain a convergence bound
 373 of $\mathcal{O}(1/T + 1/N + n^{-1/4})$ for our Max-Min algorithm, by viewing (9) as maximizing a general
 374 utility considered in (Zhang et al., 2022; Barakat et al., 2024). Here T is the number of iterations
 375 and N is the batch size for policy update. Detailed proofs and the convergence analysis for Linear
 376 Max-Min are in Appendix G.

378 **Algorithm 1** Max-Min Policy Optimization (Simplified)

379
 380 1: Initialize policy parameters θ
 381 2: Initialize reference policy π_{ref} and collect trajectories
 382 3: Estimate mean and variance of the proxy reward under π_{ref}
 383 4: **for** each iteration **do**
 384 5: Collect trajectories from current policy π_θ
 385 6: Normalize the proxy rewards for state-action pairs in the collected trajectories
 386 7: Estimate the expected proxy reward and its second moment under the current policy
 387 8: Estimate the discriminator using Equation 14 and χ^2 divergence between μ_π and $\mu_{\pi_{\text{ref}}}$
 388 9: Update the policy using PPO to maximize the Max-Min objective in Equation 9
 389 10: **end for**
 390
 391

392 **4 EXPERIMENT**
 393394 **4.1 EXPERIMENT SETUP**
 395

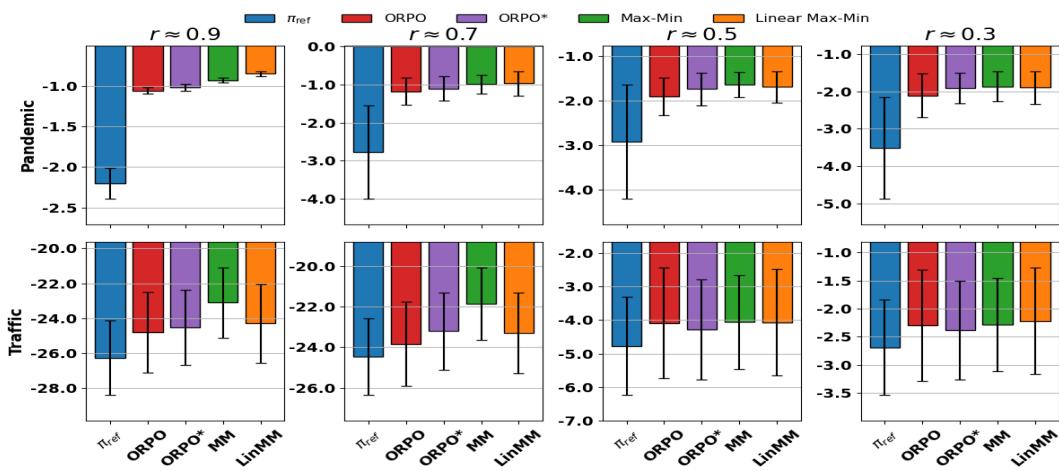
396 We evaluate our method across five realistic benchmark environments: *Traffic*, *Pandemic*, *Glucose*
 397 *Monitoring*, *Tomato Watering* and *GridWorld* and *RLHF*. These environments were originally proposed
 398 in (Pan et al., 2022; Leike et al., 2017) and represent diverse forms of proxy reward hacking, includ-
 399 ing misweighting, ontological mismatch, and scope misalignment (Pan et al., 2022). A detailed de-
 400 scription of the environments and their respective reward structures is provided in Appendix F.4. In
 401 each of the five environments, we train policies using both our Max-Min and Linear Max-Min
 402 algorithms for 5 random seeds. For baselines, we compare against the ORPO policy. To isolate the
 403 impact of discriminator training, we also include an ablation: ORPO*, where we train the ORPO
 404 policy using the same full discriminator training schedule as in our algorithms. This variant shares
 405 the same architecture and optimization settings as the original ORPO, differing only in the extent of
 406 discriminator training. Including this baseline allows us to evaluate the specific contribution of dis-
 407 criminator optimization to policy robustness. **For the RLHF environment, we additionally include**
 408 **the Ensemble baseline (Eisenstein et al., 2023), a reward-centric approach designed to mitigate**
 409 **reward hacking in RLHF.** We include more detailed experimental settings in Appendix F.5 and a
 410 discussion of training time and complexity of all algorithms in Appendix F.6.

411 Since the correlation r may only be approximately estimated and there is currently no principled
 412 method for selecting its optimal value. We adopted a similar approach used by ORPO (Laidlaw
 413 et al., 2025). For each environment, we first performed a grid search over several different values
 414 of r , and for each fixed r , we trained the policy using our algorithm. We then selected the r value
 415 that leads to the policy with the best expected worst-case return (detailed in Appendix H.2), which
 416 is 0.3 for Traffic, 0.7 for Pandemic, 0.9 for Glucose, 0.4 for Tomato, and 0.4 for RLHF. Results on
 417 all searched r can be found in Appendix H.5. Notice that ORPO selects the optimal r that yields the
 418 best expected return under the true reward, which is infeasible in practice when the true reward is
 419 unknown during training. **On the other hand, when the exact correlation r is unknown, our approach**
 420 **also raises a concern about how to interpret which worst-case reward is actually meaningful.** We
 421 include a detailed discussion about how to choose r in practice in Appendix I.

422 As for evaluation metrics, we report both the expected proxy and true rewards, along with the ex-
 423 pected worst-case reward as described in Section 3.1. Note that some policies may visit state-action
 424 pairs that are not covered by the reference policy π_{ref} . In such cases, we exclude those trajectories
 425 and report the occupancy measure of the unseen state-action pairs. Additionally, we evaluate each
 426 policy using two variants of the expected linear worst-case reward introduced in Section 3.2. The
 427 first uses only the features present in the proxy reward, while the second variant, denoted *Linear*
 428 *Worst**, leverages features from the true reward, some of which remain unseen during training. This
 429 setup mimics a more realistic real-world scenario in which the true reward function may depend on
 430 features not explicitly modeled at training time. Comparing performance under this setting allows
 431 us to assess the robustness of each policy to unseen or misaligned reward structures. All rewards
 432 are normalized with respect to the reference policy π_{ref} to ensure a consistent scale across metrics,
 433 enabling fair and meaningful comparisons. Note that all worst-case rewards are reported using the
 434 fixed correlation level r specified during training.

432
 433 Table 1: Evaluation results on Traffic, Pandemic, Glucose, and RLHF environments. All policies are trained
 434 using **only the proxy reward**. In Traffic, the proxy reward is based on *vel, accel, headway* (1, 1, 0.1), while the
 435 true reward uses *commute, accel, headway* (1, 1, 0.1). In Pandemic, the proxy reward includes *infection, lower*
 436 *stage, smooth changes* (10, 0.1, 0.01), while the true reward additionally includes *political* with weight 10 after
 437 *infection*. In Glucose, the proxy uses *expected patient cost*, and the true reward uses *magni_bg*. In RLHF, the
 438 proxy uses a 70M LLM, and the true reward uses a 8B LLM. We report θ in the same order as feature weights.
 439 **Occ** denotes total occupancy over state-action pairs unseen by π_{ref} , where discriminator outputs infinity.

| Env | Traffic | | | | | |
|----------------|-----------------------|-------------|--------------------|-------------------------------|-------------------------------------|--------------------|
| Method | True | Proxy | Worst | Linear Worst (θ) | Linear Worst* (θ) | Occ ↓ |
| ORPO | 16.91±0.12 | 3.41±0.13 | -1.96e+04±0.02e+04 | -0.69±0.01 (0.71, 0.21, 0.69) | -0.83±0.02 (0.63, 0.12, 0.97) | 3.82e-04 ±0.13e-04 |
| ORPO* | 10.26±0.09 | 1.35±0.09 | -1.35e+04±0.02e+04 | -0.44±0.02 (0.46, 0.18, 0.86) | -0.45±0.01 (0.58, 0.06, 0.81) | 1.84e-04 ±0.07e-04 |
| Max-Min | 12.70±0.06 | 3.63±0.09 | -268.31±4.14 | -0.06±0.01 (0.01, 0.02, 0.96) | -0.06±0.01 (0.001, 0.02, 0.99) | 0.00±0.00 |
| Linear Max-Min | 16.46±0.10 | 2.40±0.11 | -1.19e+04±0.01e+04 | 0.20±0.01 (0.64, 0.07, 0.76) | -0.12±0.01 (0.91, 0.01, 0.67) | 0.00±0.00 |
| Env | Pandemic | | | | | |
| Method | True | Proxy | Worst | Linear Worst (θ) | Linear Worst* (θ) | |
| ORPO | -1.04±0.21 | 1.75±0.19 | -5.31e+06±0.01e+06 | -2.41±0.02 (0.23, 0.95, 0.17) | -2.65±0.02 (0.02, 0.95, 0.92, 0.08) | |
| ORPO* | 1.18±0.19 | 1.18±0.19 | -4.46e+06±0.03e+06 | -1.36±0.01 (0.25, 0.97, 0.13) | -1.36±0.01 (0.25, 0, 0.97, 0.13) | |
| Max-Min | 1.25±0.18 | 1.25±0.18 | -63.29±3.35 | -1.11±0.01 (0.14, 0.99, 0.01) | -1.11±0.01 (0.14, 0, 0.99, 0.01) | |
| Linear Max-Min | 3.65±0.11 | 7.60±0.13 | -6.82e+05±0.01e+05 | 0.65±0.01 (0.001, 0.23, 0.02) | -0.17±0.02 (0.01, 0.97, 0.22, 0.09) | |
| Env | Glucose | | | | | |
| Method | True($\times 10^3$) | Proxy | Worst | True | Proxy | Worst |
| ORPO | 6.0±0.1 | 100.48±0.54 | -27.54±0.32 | 8.30±1.07 | 0.63±0.21 | -1.84±0.03 |
| ORPO* | 6.3±0.2 | 116.36±0.56 | -8.79±0.27 | N/A | N/A | N/A |
| Max-Min | 6.3±0.1 | 102.66±0.58 | -1.71±0.25 | 5.38±0.92 | 0.84±0.11 | -0.10±0.01 |
| Ensemble | N/A | N/A | N/A | 2.31±1.23 | 1.26±0.11 | -1.70±0.04 |



450
 451 Figure 1: Mean reward and standard deviation under sampled θ and true reward features at different proxy-true
 452 reward correlation levels r for the Traffic and Pandemic environments. Our methods (Max-Min and Linear
 453 Max-Min) yield more stable and higher average performance across all choices of r .
 454
 455
 456
 457
 458
 459
 460
 461
 462
 463
 464

4.2 RESULTS

471
 472 **Worst-Case Performance.** Table 1 presents the evaluation results on the Traffic, Pandemic, Glu-
 473 cose, and RLHF environments. Additional results for Tomato are provided in Appendix H. Note
 474 that we omit the Linear Max-Min policy from the Glucose and RLHF environments for the fol-
 475 lowing reasons. In the Glucose and RLHF environment, both the proxy and true rewards used in
 476 prior work (Laidlaw et al., 2025; Baker et al., 2025) are based on a single feature, making the lin-
 477 ear reward formulation trivial. Although the original Glucose simulator provides multiple candidate
 478 features related to patient health, selecting an appropriate feature combination without prior knowl-
 479 edge of clinical intent is nontrivial. Therefore, in both settings, we report only the results for the
 480 Max-Min policy alongside the baselines.

481 Our Max-Min and Linear Max-Min policies achieve better expected worst-case performance
 482 under both general and linear adversarial rewards, while remaining competitive with baselines in
 483 terms of expected true and proxy rewards. Notably, the Max-Min policy attains the highest expected
 484 worst-case return, followed by Linear Max-Min. Conversely, Linear Max-Min yields the
 485 highest expected linear worst-case reward, followed by Max-Min, demonstrating the robustness
 486 of both approaches under worst-case scenarios. For the Linear Worst* evaluation, which uses re-
 487 ward features unseen during training, we observe minimal degradation in Max-Min policy’s perfor-

486 mance, indicating its strong robustness to feature variation. In contrast, the performance of Linear
 487 Max–Min declines in this case, suggesting its advantage diminishes when prior assumptions about
 488 feature structure are inaccurate. We also find that ORPO* exhibits better worst-case performance
 489 than the original ORPO. In particular, training the discriminator more thoroughly significantly re-
 490 duces the occupancy of state-action pairs that are not visited by the reference policy, indicating that
 491 more accurate estimation of the Radon–Nikodym derivative leads to improved policy robustness.
 492 Notably, in the Pandemic and Glucose environment, we observe no such unvisited state-action pairs,
 493 and the discriminator outputs remain small across all policies. This could be due to either the dis-
 494 criminator network not being fully optimized or its inability to capture rare events that fall outside
 495 the support of π_{ref} . Developing more reliable techniques for handling such rare or unseen state-
 496 action pairs remains an open direction for future work. We also report the adversarial weight vectors
 497 θ for each policy. These weights reveal which features are most vulnerable to proxy exploitation
 498 under the learned policy and can be used to diagnose and revise the proxy reward function, thereby
 499 improving robustness. This highlights the interpretability benefits of our framework. Moreover,
 500 several patterns emerge from the results, which is detailed in Appendix H.2. **We also notice that the**
 501 **Ensemble baseline in the RLHF setting achieves only limited improvement in expected true return**
 502 **over the reference policy and attains a lower expected worst-case return than our method.** These
 503 results indicate that using reward ensembles alone is insufficient to effectively mitigate reward hack-
 504 ing compared to our approach. However, such reward-centric methods, including InfoRM (Miao
 505 et al., 2024) and RRM (Liu et al., 2024), can be easily integrated into our framework. In particular,
 506 these approaches can be used to construct a stronger proxy reward, which can then be plugged into
 507 our method to further improve performance.

508 **Robustness Across Correlation Levels.** To further assess the robustness of each policy across
 509 a broader range of proxy–true correlation scenarios, we also compute the Linear Worst* for each
 510 policy under varying r values. Specifically, for each r , we sample 1000 vectors θ such that $\theta \in \mathcal{R}_{\text{corr}}^{\text{lin}}$,
 511 and report the average return and variance achieved by each policy over these sampled rewards.
 512 Importantly, the variation in r is applied **only during evaluation**; all policies are fixed and trained
 513 using the specific r values reported in Appendix F.5. Unlike evaluations that only consider several
 514 reward functions, this approach evaluates policy performance across the entire reward set $\mathcal{R}_{\text{corr}}^{\text{lin}}$,
 515 providing a more comprehensive measure of robustness and better reflecting real-world scenarios
 516 where the true reward and correlation r are unknown.

517 Figure 1 shows the average reward and variance achieved by each method under different levels of
 518 proxy–true reward correlation r . As expected, the reference policy π_{ref} (blue) performs the worst
 519 across all correlation levels in both environments. In Traffic, its variance is relatively small, sug-
 520 gesting consistently poor but stable behavior. In contrast, variance is highest in the Pandemic
 521 environment, indicating increased policy fragility. Notably, ORPO* (purple) consistently achieves lower
 522 variance than ORPO (red) across both environments and outperforms it in terms of average reward at
 523 $r \approx 0.9$ and $r \approx 0.7$ in Traffic, and across nearly all r values in Pandemic. This underscores the im-
 524 portance of accurate discriminator training for improving both stability and robustness. Max–Min
 525 (green) demonstrates the highest average reward and lowest variance across a wide range of r val-
 526 ues in both environments, showing strong resilience to reward misspecification. While Linear
 527 Max–Min (orange) achieves the best performance at specific correlation levels, particularly $r \approx 0.3$
 528 in Traffic and $r \approx 0.7$ – 0.9 in Pandemic. As r decreases and the proxy becomes less informative, dif-
 529 ferences in average reward among methods shrink, while variance increases. These results highlight
 530 the significance of variance control in low-correlation regimes and demonstrate that Max–Min and
 531 Linear Max–Min offer robust and stable performance under high uncertainty.

532 5 CONCLUSION

533 In this work, we formalize reward hacking as a robust optimization problem and introduce both
 534 a Max–Min formulation with a closed-form adversarial reward and a Linear Max–Min variant that
 535 further improves interpretability and tractability. We develop efficient algorithms and empirically
 536 validate that both Max–Min and Linear Max–Min policies achieve stronger worst-case performance
 537 and improved stability compared to prior baselines such as ORPO across diverse environments. We
 538 further discuss limitations and broader impacts of our method in Appendices B and C.

540 REFERENCES
541

542 Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In
543 *International Conference on Machine learning (ICML)*, 2004.

544 Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Con-
545 crete problems in ai safety. *arXiv preprint arXiv:1606.06565*, 2016.

546 J Andrew Bagnell, Andrew Y Ng, and Jeff G Schneider. Solving uncertain markov decision pro-
547 cesses. 2001.

548 Bowen Baker, Joost Huizinga, Leo Gao, Zehao Dou, Melody Y Guan, Aleksander Madry, Wojciech
549 Zaremba, Jakub Pachocki, and David Farhi. Monitoring reasoning models for misbehavior and
550 the risks of promoting obfuscation. *arXiv preprint arXiv:2503.11926*, 2025.

551 Anas Barakat, Souradip Chakraborty, Peihong Yu, Pratap Tokekar, and Amrit Singh Bedi. Towards
552 scalable general utility reinforcement learning: Occupancy approximation, sample complexity
553 and global optimality. *arXiv preprint arXiv:2410.04108*, 2024.

554 André Barreto, Will Dabney, Rémi Munos, Jonathan J Hunt, Tom Schaul, Hado P van Hasselt,
555 and David Silver. Successor features for transfer in reinforcement learning. *Advances in neural
556 information processing systems*, 30, 2017.

557 Andre Barreto, Diana Borsa, John Quan, Tom Schaul, David Silver, Matteo Hessel, Daniel
558 Mankowitz, Augustin Zidek, and Remi Munos. Transfer in deep reinforcement learning using
559 successor features and generalised policy improvement. In *International Conference on Machine
560 Learning*, pp. 501–510. PMLR, 2018.

561 Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric
562 Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, et al.
563 Pythia: A suite for analyzing large language models across training and scaling. In *International
564 Conference on Machine Learning*, pp. 2397–2430. PMLR, 2023.

565 Serena Booth, W Bradley Knox, Julie Shah, Scott Niekum, Peter Stone, and Alessandro Allievi.
566 The perils of trial-and-error reward design: misdesign through overfitting and invalid task spec-
567 ifications. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp.
568 5920–5929, 2023.

569 Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

570 Ronen Brafman, Giuseppe De Giacomo, and Fabio Patrizi. Ltlf/ldlf non-markovian rewards. In
571 *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.

572 Alberto Camacho, Oscar Chen, Scott Sanner, and Sheila McIlraith. Non-markovian rewards ex-
573 pressed in ltl: guiding search via reward shaping. In *Proceedings of the International Symposium
574 on Combinatorial Search*, volume 8, pp. 159–160, 2017.

575 Thomas Coste, Usman Anwar, Robert Kirk, and David Krueger. Reward model ensembles help
576 mitigate overoptimization. In *NeurIPS 2023 Workshop on Instruction Tuning and Instruction
577 Following*.

578 Peter Dayan. Improving generalization for temporal difference learning: The successor representa-
579 tion. *Neural computation*, 5(4):613–624, 1993.

580 Erick Delage and Yinyu Ye. Distributionally robust optimization under moment uncertainty with
581 application to data-driven problems. *Operations research*, 58(3):595–612, 2010.

582 Esther Derman, Matthieu Geist, and Shie Mannor. Twice regularized mdps and the equivalence
583 between robustness and regularization. *Advances in Neural Information Processing Systems*, 34:
584 22274–22287, 2021.

585 Dotan Di Castro, Aviv Tamar, and Shie Mannor. Policy gradients with variance related risk criteria.
586 In *International Conference on Machine learning (ICML)*, 2012.

594 Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning.
 595 *stat*, 1050:2, 2017.

596

597 Jeff Druce, Michael Harradon, and James Tittle. Explainable artificial intelligence (xai) for in-
 598 creasing user trust in deep reinforcement learning driven autonomous systems. *arXiv preprint*
 599 *arXiv:2106.03775*, 2021.

600 Yann Dubois, Chen Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos
 601 Guestrin, Percy S Liang, and Tatsunori B Hashimoto. Alpacafarm: A simulation framework for
 602 methods that learn from human feedback. *Advances in Neural Information Processing Systems*,
 603 36:30039–30069, 2023.

604

605 Jacob Eisenstein, Chirag Nagpal, Alekh Agarwal, Ahmad Beirami, Alex D’Amour, DJ Dvijotham,
 606 Adam Fisch, Katherine Heller, Stephen Pfohl, Deepak Ramachandran, et al. Helping or herd-
 607 ing? reward model ensembles mitigate but do not eliminate reward hacking. *arXiv preprint*
 608 *arXiv:2312.09244*, 2023.

609 Tom Everitt, Victoria Krakovna, Laurent Orseau, Marcus Hutter, and Shane Legg. Reinforcement
 610 learning with a corrupted reward channel. In *International Joint Conference on Artificial Intelli-
 611 gence (IJCAI)*, 2017.

612 Tom Everitt, Marcus Hutter, Ramana Kumar, and Victoria Krakovna. Reward tampering problems
 613 and solutions in reinforcement learning: A causal influence diagram perspective. *Synthese*, 198
 614 (Suppl 27):6435–6467, 2021.

615

616 Benjamin Eysenbach and Sergey Levine. Maximum entropy rl (provably) solves some robust rl
 617 problems. In *International Conference on Learning Representations*.

618 Alan Fern, Sriraam Natarajan, Kshitij Judah, and Prasad Tadepalli. A decision-theoretic model of
 619 assistance. *Journal of Artificial Intelligence Research*, 50:71–104, 2014.

620

621 Ian Fox, Joyce Lee, Rodica Pop-Busui, and Jenna Wiens. Deep reinforcement learning for closed-
 622 loop blood glucose control. In *Machine Learning for Healthcare Conference*, pp. 508–536.
 623 PMLR, 2020.

624 Matt Franchi, JD Zamfirescu-Pereira, Wendy Ju, and Emma Pierson. Detecting disparities in police
 625 deployments using dashcam data. In *Proceedings of the 2023 ACM Conference on Fairness,
 626 Accountability, and Transparency*, pp. 534–544, 2023.

627

628 Clement A Gehring. Approximate linear successor representation. *Reinforcement Learning Decision
 629 Making*, 2015.

630

631 Joel Goh and Melvyn Sim. Distributionally robust optimization and its tractable approximations.
Operations research, 58(4-part-1):902–917, 2010.

632

633 Charles AE Goodhart and CAE Goodhart. *Problems of monetary management: the UK experience*.
 634 Springer, 1984.

635

636 Vineet Goyal and Julien Grand-Clement. Robust markov decision process: Beyond rectangularity.
arXiv preprint arXiv:1811.00215, 2018.

637

638 Julien Grand-Clément and Christian Kroer. Scalable first-order methods for robust mdps. In *Pro-
 639 ceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 12086–12094, 2021.

640

641 Dylan Hadfield-Menell, Smitha Milli, Pieter Abbeel, Stuart J Russell, and Anca Dragan. Inverse
 642 reward design. In *Advances in Neural Information Processing Systems*, 2017.

643

644 Yuval Noah Harari. *Nexus: A brief history of information networks from the stone age to AI*. Signal,
 2024.

645 Alexander Havrilla, Maksym Zhuravinskyi, Duy Phung, Aman Tiwari, Jonathan Tow, Stella Bider-
 646 man, Quentin Anthony, and Louis Castricato. trlx: A framework for large scale reinforcement
 647 learning from human feedback. In *Proceedings of the 2023 Conference on Empirical Methods in
 Natural Language Processing*, pp. 8578–8595, 2023.

648 Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Advances in Neural*
 649 *Information Processing Systems*, 2016.

650

651 Andreas Holzinger, Chris Biemann, Constantinos S Pattichis, and Douglas B Kell. What do we need
 652 to build explainable ai systems for the medical domain? *arXiv preprint arXiv:1712.09923*, 2017.

653

654 Zhaolin Hu and L Jeff Hong. Kullback-leibler divergence constrained distributionally robust opti-
 655 mization. *Available at Optimization Online*, 1(2):9, 2013.

656

657 Audrey Huang, Jinglin Chen, and Nan Jiang. Reinforcement learning in low-rank mdps with density
 658 features. In *International Conference on Machine Learning*, pp. 13710–13752. PMLR, 2023.

659

660 Audrey Huang, Wenhao Zhan, Tengyang Xie, Jason D Lee, Wen Sun, Akshay Krishnamurthy, and
 661 Dylan J Foster. Correcting the mythos of kl-regularization: Direct alignment without overopti-
 662 mization via chi-squared preference optimization. *arXiv preprint arXiv:2407.13399*, 2024.

663

664 Borja Ibarz, Jan Leike, Tobias Pohlen, Geoffrey Irving, Shane Legg, and Dario Amodei. Reward
 665 learning from human preferences and demonstrations in atari. In *Advances in Neural Information*
 666 *Processing Systems*, 2018.

667

668 Rodrigo Toro Icarte, Toryn Q Klassen, Richard Valenzano, and Sheila A McIlraith. Reward ma-
 669 chines: Exploiting reward function structure in reinforcement learning. *Journal of Artificial In-*
 670 *telligence Research*, 73:173–208, 2022.

671

672 Hamish Ivison, Yizhong Wang, Jiacheng Liu, Zeqiu Wu, Valentina Pyatkin, Nathan Lambert,
 673 Noah A Smith, Yejin Choi, and Hanna Hajishirzi. Unpacking dpo and ppo: Disentangling best
 674 practices for learning from preference feedback. *Advances in neural information processing sys-*
 675 *tems*, 37:36602–36633, 2024.

676

677 Garud Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–
 678 280, 2005.

679

680 Rahul Iyer, Yuezhang Li, Huao Li, Michael Lewis, Ramitha Sundar, and Katia Sycara. Transparency
 681 and explanation in deep reinforcement learning neural networks. In *Proceedings of the 2018*
 682 *AAAI/ACM Conference on AI, Ethics, and Society*, pp. 144–150, 2018.

683

684 Hong Jun Jeon, Smitha Milli, and Anca Dragan. Reward-rational (implicit) choice: A unifying
 685 formalism for reward learning. In *Advances in Neural Information Processing Systems*, 2020.

686

687 Minqi Jiang, Edward Grefenstette, and Tim Rocktäschel. Prioritized level replay. In *International*
 688 *Conference on Machine Learning*, pp. 4940–4950. PMLR, 2021.

689

690 Chi Jin, Praneeth Netrapalli, and Michael Jordan. What is local optimality in nonconvex-nonconcave
 691 minimax optimization? In *International conference on machine learning*, pp. 4880–4889. PMLR,
 692 2020.

693

694 Bingyi Kang, Zequn Jie, and Jiashi Feng. Policy optimization with demonstrations. In *International*
 695 *Conference on Machine Learning (ICML)*, 2018.

696

697 David L Kaufman and Andrew J Schaefer. Robust modified policy iteration. *INFORMS Journal on*
 698 *Computing*, 25(3):396–410, 2013.

699

700 Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. The challenge of understanding what
 701 users want: Inconsistent preferences and engagement optimization. *Management Science*, 70(9):
 6336–6355, 2024.

702

703 W Bradley Knox and James MacGlashan. How to specify reinforcement learning objectives. In
 704 *Finding the Frame: An RLC Workshop for Examining Conceptual Frameworks*, 2024.

705

706 W Bradley Knox, Alessandro Allievi, Holger Banzhaf, Felix Schmitt, and Peter Stone. Reward
 707 (mis) design for autonomous driving. *Artificial Intelligence*, 316:103829, 2023.

708

709 Jack Koch, Lauro Langosco, Jacob Pfau, James Le, and Lee Sharkey. Objective robustness in deep
 710 reinforcement learning. *arXiv preprint arXiv:2105.14111*, 2, 2021.

702 V Kompella, R Capobianco, S Jong, J Browne, S Fox, L Meyers, P Wurman, P Stone, et al. Reinforcement learning for optimization of covid-19 mitigation policies. In *CEUR WORKSHOP PROCEEDINGS*, volume 2884. CEUR-WS, 2020.

703

704

705 Victoria Krakovna. Specification gaming examples in ai, April 2018. URL <https://vkrakovna.wordpress.com/2018/04/02/specification-gaming-examples-in-ai/>. Blog post.

706

707

708

709 Victoria Krakovna. Classifying specification problems as variants of goodhart’s law, August 2019. URL <https://vkrakovna.wordpress.com/2019/08/19/classifying-specification-problems-as-variants-of-goodharts-law/>. Blog post.

710

711

712

713 Victoria Krakovna, Laurent Orseau, Ramana Kumar, Miljan Martic, and Shane Legg. Penalizing side effects using stepwise relative reachability. *arXiv preprint arXiv:1806.01186*, 2018.

714

715

716 Cassidy Laidlaw, Eli Bronstein, Timothy Guo, Dylan Feng, Lukas Berglund, Justin Svegliato, Stuart Russell, and Anca Dragan. Scalably solving assistance games. In *ICLR 2025 Workshop on Bidirectional Human-AI Alignment*, 2024.

717

718

719 Cassidy Laidlaw, Shivam Singhal, and Anca Dragan. Correlated proxies: A new definition and improved mitigation for reward hacking. In *International Conference on Learning Representations*, 2025.

720

721

722

723 Henry Lam. Robust sensitivity analysis for stochastic systems. *Mathematics of Operations Research*, 41(4):1248–1275, 2016.

724

725

726 Jan Leike, Miljan Martic, Victoria Krakovna, Pedro Ortega, Tom Everitt, Ryan Lefrancq, Laurent Orseau, and Shane Legg. AI safety gridworlds. *arXiv preprint arXiv:1711.09883*, 2017.

727

728 Eric Liang, Richard Liaw, Robert Nishihara, Philipp Moritz, Roy Fox, Ken Goldberg, Joseph E Gonzalez, Michael I Jordan, and Ion Stoica. Rllib: Abstractions for distributed reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2018.

729

730

731 Fei Liu et al. Learning to summarize from human feedback. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 583–592, 2020.

732

733

734 Tianqi Liu, Wei Xiong, Jie Ren, Lichang Chen, Junru Wu, Rishabh Joshi, Yang Gao, Jiaming Shen, Zhen Qin, Tianhe Yu, et al. Rrm: Robust reward model training mitigates reward hacking. *arXiv preprint arXiv:2409.13156*, 2024.

735

736

737 Chiara Dalla Man, Francesco Micheletto, Dayu Lv, Marc Breton, Boris Kovatchev, and Claudio Cobelli. The uva/padova type 1 diabetes simulator: new features. *Journal of diabetes science and technology*, 8(1):26–34, 2014.

738

739

740

741 Jeremy McMahan, Giovanni Artiglio, and Qiaomin Xie. Roping in uncertainty: Robustness and regularization in markov games. In *International Conference on Machine Learning*, pp. 35267–35295. PMLR, 2024.

742

743

744 Yuchun Miao, Sen Zhang, Liang Ding, Rong Bao, Lefei Zhang, and Dacheng Tao. Inform: Mitigating reward hacking in rlhf via information-theoretic reward modeling. *Advances in Neural Information Processing Systems*, 37:134387–134429, 2024.

745

746

747

748 Smitha Milli, Luca Belli, and Moritz Hardt. From optimizing engagement to measuring value. In *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, pp. 714–722, 2021.

749

750

751 Nathan Monette, Alistair Letcher, Michael Beukman, Matthew T Jackson, Alexander Rutherford, Alexander D Goldie, and Jakob N Foerster. An optimisation framework for unsupervised environment design. *arXiv preprint arXiv:2505.20659*, 2025.

752

753

754 Jorge J. Moré. The levenberg-marquardt algorithm: Implementation and theory. Technical Report ANL-80-20, Argonne National Laboratory, Argonne, IL, 1978. Lecture Notes in Mathematics, vol. 630.

755

756 Tong Mu, Alec Helyar, Johannes Heidecke, Joshua Achiam, Andrea Vallone, Ian Kivlichan, Molly
 757 Lin, Alex Beutel, John Schulman, and Lilian Weng. Rule based rewards for language model
 758 safety. *Advances in Neural Information Processing Systems*, 37:108877–108901, 2024.

759

760 Calarina Muslimani, Kerrick Johnstonbaugh, Suyog Chandramouli, Serena Booth, W Bradley Knox,
 761 and Matthew E Taylor. Towards improving reward design in rl: A reward alignment metric for rl
 762 practitioners. *arXiv preprint arXiv:2503.05996*, 2025.

763 Hoang Nam Nguyen, Abdel Lisser, and Vikas Vikram Singh. Distributionally robust chance-
 764 constrained markov decision processes. *arXiv preprint arXiv:2212.08126*, 2022.

765

766 Arnab Nilim and Laurent El Ghaoui. Robust control of markov decision processes with uncertain
 767 transition matrices. *Operations Research*, 53(5):780–798, 2005.

768 Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. Dissecting racial bias
 769 in an algorithm used to manage the health of populations. *Science*, 366(6464):447–453, 2019.

770

771 Alexander Pan, Kush Bhatia, and Jacob Steinhardt. The effects of reward misspecification: Mapping
 772 and mitigating misaligned models. In *International Conference on Learning Representations*,
 773 2022.

774 A Paszke. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint
 775 arXiv:1912.01703*, 2019.

776

777 Erika Puiutta and Eric MSP Veith. Explainable reinforcement learning: A survey. In *International
 778 cross-domain conference for machine learning and knowledge extraction*, pp. 77–95. Springer,
 779 2020.

780 Alexandre Ramé, Nino Vieillard, Léonard Hussonot, Robert Dadashi, Geoffrey Cideron, Olivier
 781 Bachem, and Johan Ferret. Warm: On the benefits of weight averaged reward models. *arXiv
 782 preprint arXiv:2401.12187*, 2024.

783

784 Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and
 785 use interpretable models instead. *Nature machine intelligence*, 1(5):206–215, 2019.

786

787 Karim Abdel Sadek, Matthew Farrugia-Roberts, Usman Anwar, Hannah Erlebach, Christian
 788 Schroeder de Witt, David Krueger, and Michael Dennis. Mitigating goal misgeneralization
 789 via minimax regret. *arXiv preprint arXiv:2507.03068*, 2025.

790

791 Dorsa Sadigh, Anca Dragan, Shankar Sastry, and Sanjit Seshia. Active preference-based learning of
 792 reward functions. In *Robotics: Science and Systems*, 2017.

793

794 Andreas Schlaginhaufen and Maryam Kamgarpour. Identifiability and generalizability in con-
 795 strained inverse reinforcement learning. In *International Conference on Machine Learning
 (ICML)*, 2023.

796

797 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
 798 optimization algorithms. In *arXiv preprint arXiv:1707.06347*, 2017.

799

800 Rohin Shah, Pedro Freire, Neel Alex, Rachel Freedman, Dmitrii Krasheninnikov, Lawrence Chan,
 801 Michael D Dennis, Pieter Abbeel, Anca Dragan, and Stuart Russell. Benefits of assistance over
 802 reward learning. 2020.

803

804 Joar Skalse, Nikolaus Howe, Dmitrii Krasheninnikov, and David Krueger. Defining and character-
 805 izing reward gaming. *Advances in Neural Information Processing Systems*, 35:9460–9471, 2022.

806

807 Garry M Steil. Algorithms for a closed-loop artificial pancreas: the case for proportional-integral-
 808 derivative control. *Journal of diabetes science and technology*, 7(6):1621–1631, 2013.

809

810 Jonathan Stray, Alon Halevy, Parisa Assar, Dylan Hadfield-Menell, Craig Boutilier, Amar Ashar,
 811 Chloe Bakalar, Lex Beattie, Michael Ekstrand, Claire Leibowicz, et al. Building human values
 812 into recommender systems: An interdisciplinary synthesis. *ACM Transactions on Recommender
 813 Systems*, 2(3):1–57, 2024.

810 Richard S Sutton and Andrew G Barto. *Reinforcement Learning: An Introduction*. MIT press, 2018.
 811

812 Jeremy Tien, Jerry Zhi-Yang He, Zackory Erickson, Anca D Dragan, and Daniel S Brown. Causal
 813 confusion and reward misidentification in preference-based reward learning. In *International
 814 Conference on Learning Representations*, 2023.

815 Martin Treiber, Ansgar Hennecke, and Dirk Helbing. Congested traffic states in empirical observa-
 816 tions and microscopic simulations. *Physical Review E*, 62(2):1805, 2000.
 817

818 Eugene Vinitsky, Aboudy Kreidieh, Luc Le Flem, Nishant Kheterpal, Kathy Jang, Cathy Wu,
 819 Fangyu Wu, Richard Liaw, Eric Liang, and Alexandre M Bayen. Benchmarks for reinforcement
 820 learning in mixed-autonomy traffic. In *Conference on Robot Learning*, pp. 399–409. PMLR,
 821 2018.

822 Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau,
 823 Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der
 824 Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nel-
 825 son, Eric Jones, Robert Kern, Eric Larson, C. J. Carey, İlhan Polat, Yu Feng, Eric W. Moore,
 826 Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero,
 827 Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mul-
 828 bregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental algorithms for scientific computing in
 829 python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.

830 George A Vouros. Explainable deep reinforcement learning: state of the art and challenges. *ACM
 831 Computing Surveys*, 55(5):1–39, 2022.
 832

833 Cathy Wu, Abdul Rahman Kreidieh, Kanaad Parvate, Eugene Vinitsky, and Alexandre M Bayen.
 834 Flow: A modular learning framework for mixed autonomy traffic. *IEEE Transactions on Robotics*,
 835 38(2):1270–1286, 2021.

836 Tengyang Xie, Bo Liu, Yangyang Xu, Mohammad Ghavamzadeh, Yinlam Chow, Daoming Lyu, and
 837 Daesub Yoon. A block coordinate ascent algorithm for mean-variance optimization. In *Advances
 838 in Neural Information Processing Systems*, 2018.

839 Rui Yu, Shenghua Wan, Yucen Wang, Chen-Xiao Gao, Le Gan, Zongzhang Zhang, and De-
 840 Chuan Zhan. Reward models in deep reinforcement learning: A survey. *arXiv preprint
 841 arXiv:2506.15421*, 2025.
 842

843 Wenhao Zhan, Baihe Huang, Audrey Huang, Nan Jiang, and Jason Lee. Offline reinforcement
 844 learning with realizability and single-policy concentrability. In *Conference on Learning Theory*,
 845 pp. 2730–2775. PMLR, 2022.

846 Junyu Zhang, Amrit Singh Bedi, Mengdi Wang, and Alec Koppel. Multi-agent reinforcement learn-
 847 ing with general utilities via decentralized shadow reward actor-critic. In *Proceedings of the AAAI
 848 Conference on Artificial Intelligence*, volume 36, pp. 9031–9039, 2022.
 849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

| | | |
|-----|--|-----------|
| 864 | APPENDIX | |
| 865 | | |
| 866 | | |
| 867 | A LLM Usage | 19 |
| 868 | | |
| 869 | B Limitations and Future Work | 19 |
| 870 | | |
| 871 | | |
| 872 | C Broader Impacts | 20 |
| 873 | | |
| 874 | D Related Work | 20 |
| 875 | | |
| 876 | D.1 Reward Hacking | 20 |
| 877 | D.2 Reward Modeling in Reinforcement Learning | 21 |
| 878 | D.3 Robust Reinforcement Learning | 21 |
| 879 | D.4 Successor Representations in Reinforcement Learning | 22 |
| 880 | | |
| 881 | | |
| 882 | E Proofs and Additional Theoretical Results | 22 |
| 883 | | |
| 884 | E.1 Solve the Max-Min Objective | 22 |
| 885 | E.2 Proof of Optimality | 26 |
| 886 | E.3 Proof that $\chi^2(\mu_\pi \ \mu_{\pi_{\text{ref}}}) \geq \mathbb{E}_{\mu_\pi}^2[R_{\text{proxy}}]$ | 28 |
| 887 | E.4 Derive Lagrangian Functional for Linear Max-Min Objective | 29 |
| 888 | E.5 Proof for Whitening Transformation | 30 |
| 889 | E.6 Derive Optimal Primal Variable for Linear Max-Min Objective | 31 |
| 890 | E.7 Solve the Dual Objective for Linear Max-Min Objective | 32 |
| 891 | E.8 Policy Gradient Derivation | 33 |
| 892 | E.9 Proof of Theorem 1 | 35 |
| 893 | | |
| 894 | | |
| 895 | | |
| 896 | | |
| 897 | F Additional Implementation Details | 35 |
| 898 | | |
| 899 | F.1 Training Discriminator Network | 35 |
| 900 | F.2 Derivation of Max-Min Policy Optimization | 37 |
| 901 | F.3 Derivation of Linear Max-Min Policy Optimization | 38 |
| 902 | F.4 Environment Description and Reward Hacking Types | 40 |
| 903 | F.5 Additional Experiment Setup | 42 |
| 904 | F.6 Training Time and Complexity | 44 |
| 905 | | |
| 906 | | |
| 907 | | |
| 908 | G Convergence Analysis | 45 |
| 909 | | |
| 910 | G.1 Sample Complexity of Occupancy Measure Estimation | 45 |
| 911 | G.2 Guarantees for Max-Min with Occupancy Measure Approximation | 51 |
| 912 | G.3 Convergences for Linear Max-Min algorithm | 52 |
| 913 | | |
| 914 | | |
| 915 | H Additional Experiment Results | 53 |
| 916 | | |
| 917 | H.1 Feature Weights in Linear Max-Min Optimization during Training | 53 |
| | H.2 Additional Worst-Case Performance Results | 54 |

| | | |
|-----|---|-----------|
| 918 | H.3 Additional Results for Robustness Across Correlation Levels | 56 |
| 919 | | |
| 920 | H.4 Additonal Unnormalized Results | 57 |
| 921 | | |
| 922 | H.5 Results for all r | 57 |
| 923 | | |
| 924 | I How to choose r in practice? | 58 |
| 925 | | |
| 926 | | |
| 927 | | |
| 928 | | |
| 929 | | |
| 930 | | |
| 931 | | |
| 932 | | |
| 933 | | |
| 934 | | |
| 935 | | |
| 936 | | |
| 937 | | |
| 938 | | |
| 939 | | |
| 940 | | |
| 941 | | |
| 942 | | |
| 943 | | |
| 944 | | |
| 945 | | |
| 946 | | |
| 947 | | |
| 948 | | |
| 949 | | |
| 950 | | |
| 951 | | |
| 952 | | |
| 953 | | |
| 954 | | |
| 955 | | |
| 956 | | |
| 957 | | |
| 958 | | |
| 959 | | |
| 960 | | |
| 961 | | |
| 962 | | |
| 963 | | |
| 964 | | |
| 965 | | |
| 966 | | |
| 967 | | |
| 968 | | |
| 969 | | |
| 970 | | |
| 971 | | |

972 **A LLM USAGE**
973974 In this work, we used ChatGPT to improve grammar, wording, and paragraph flow throughout the
975 paper after completing an initial draft. We also used ChatGPT’s research capabilities to help sur-
976 face potentially relevant prior work for the related work and introduction sections. All references
977 were independently verified by the authors. No algorithms, proofs, or experimental results were
978 generated by ChatGPT, and no proprietary or sensitive data were shared with the tool. All technical
979 contributions and analyses are solely the authors’ work.
980981 **B LIMITATIONS AND FUTURE WORK**
982983 Despite the effectiveness of our framework, several limitations remain. First, although our Max-Min
984 formulation can be extended naturally beyond linear reward structures, incorporating more expres-
985 sive representations such as neural networks makes the inner optimization problem significantly
986 harder. In such cases, the inner minimization may no longer admit a closed-form solution, necessi-
987 tating iterative training between the policy and adversary. This increases computational complexity
988 and undermines the efficiency advantages of our current formulation. Developing scalable solutions
989 for general reward representations remains an open direction.
990991 Second, in our Linear Max-Min algorithm, computing Q results in $O(d^2)$ space and $O(d^3)$ time
992 complexity. While low-rank approximations could potentially reduce computational cost, such
993 methods often discard small eigenvalues. However, in our setting, these small eigenvalues be-
994 come critical due to the inversion in the whitening step, and removing them may severely distort
995 the worst-case reward direction. Therefore, naive low-rank approximations may not be applicable
996 in our setting, and we emphasize the need for principled, scalable extensions when applying our
method to settings with very high-dimensional feature spaces.
997998 Third, for complex environments, constructing effective features for the reward function is often
999 challenging without prior domain knowledge. For example, in the Glucose environment, a large
1000 number of health-related indices are provided. However, without medical expertise or knowledge of
1001 glucose monitoring, it is difficult to determine which combination of indices best captures patient
1002 health or blood glucose trends. Using arbitrarily selected features in such cases can lead to proxy
1003 rewards that exhibit little or no correlation with the true reward. While our max-min formulation
1004 can still offer robustness under such misspecification, the resulting policy is nevertheless expected to
1005 perform poorly due to the fundamental misalignment between the proxy and true objectives. There-
1006 fore, designing meaningful reward features remains a fundamental and unresolved challenge, and
1007 we will include this as a limitation of our method in the main text. Moreover, in some environ-
1008 ments, such as Tomato, the reward function is not explicitly feature-based. Although our general
1009 max-min algorithm still applies in this setting, incorporating non feature-based reward structure into
the uncertainty set remains an open problem.
10101011 Fourth, like ORPO, our framework assumes access to a fixed proxy reward, a reference policy,
1012 and a pre-specified correlation parameter r , all provided offline. This setup limits the ability of
1013 the algorithm to adapt or refine its reward model based on new information. However, we observe
1014 that the adversarial rewards generated by our method, particularly the structured linear ones, can
1015 serve as diagnostic tools to identify vulnerable reward features. These insights could be leveraged to
1016 guide human-in-the-loop refinement or adaptive querying of stronger feedback models (e.g., large
language models). Extending our framework to close the loop between diagnostic robustness and
reward learning is an exciting direction for future work.
10171018 Fifth, while our experimental results demonstrate that the proposed method improves robustness
1019 across a range of proxy-true reward correlation levels, an alternative and perhaps more direct strategy
1020 would be to train the policy against multiple proxy rewards sampled at varying levels of correlation
1021 r . In principle, optimizing the average performance across a diverse set of proxies could yield a
1022 policy that is robust to a wider distribution of potential reward misspecifications. However, this
1023 approach presents several practical challenges. First, there is a trade-off between computational cost
1024 and coverage: sampling too few proxies may fail to represent the full space of plausible reward
1025 deviations, while sampling many proxies significantly increases training time. Second, efficiently
generating reward functions that satisfy a fixed correlation constraint with the proxy reward becomes
non-trivial in high-dimensional or continuous state-action spaces. Designing scalable and effective
1026

1026 reward sampling mechanisms (such as leveraging diffusion models) under correlation constraints
 1027 remains an open problem and a promising direction for future research.
 1028

1029 C BROADER IMPACTS

1030 Designing reward functions that faithfully reflect designer intent remains a fundamental challenge
 1031 in deploying reinforcement learning (RL) systems in the real world. When reward misspecification
 1032 occurs, agents can behave in undesirable or even dangerous ways. Our work addresses this issue
 1033 by proposing a robust policy optimization framework that explicitly accounts for uncertainty in the
 1034 reward function, improving worst-case performance across a range of plausible reward proxies. This
 1035 approach has the potential to increase the safety and reliability of RL systems in safety-critical ap-
 1036 plications such as healthcare, autonomous driving, and digital infrastructure, where poorly specified
 1037 incentives can lead to unintended consequences. In addition to robustness, our linear variant con-
 1038 tributes to policy interpretability by yielding explicit weightings over features that can be inspected
 1039 and audited. This can help practitioners identify vulnerable components in their reward specification
 1040 and make better-informed decisions when refining proxies. However, while our method is primarily
 1041 intended to prevent reward exploitation, one could conceivably use adversarial reward modeling to
 1042 stress-test or attack policies. We believe the benefits of improved safety and robustness outweigh
 1043 this risk, especially when combined with interpretability. Overall, this work contributes to the safe
 1044 and trustworthy deployment of RL by equipping practitioners with more robust and explainable
 1045 optimization tools.
 1046

1047 D RELATED WORK

1048 D.1 REWARD HACKING

1049 Early work in AI safety underscored the pitfalls of optimizing an imperfect proxy reward. Amodei
 1050 et al. (Amodei et al., 2016) famously illustrate how an agent can “game” its reward function: for
 1051 example, a cleaning robot rewarded for not seeing any messes might simply close its cameras or
 1052 create messes to clean up, maximizing the proxy reward while betraying the designer’s intent. Other
 1053 examples of such reward hacking include an agent in a racing game that spins in circles to collect
 1054 points instead of completing the race (Skalse et al., 2022), social media recommendation systems
 1055 that promote emotionally extreme content to increase engagement (Harari, 2024), and Large Lan-
 1056 guage Models (LLMs) that generate trivial or hard-coded solutions to pass unit tests rather than
 1057 producing general, correct code (Baker et al., 2025). Krakovna et al. (Krakovna, 2018) have cata-
 1058 logued many such failure cases across diverse domains. Several studies have analyzed the causes of
 1059 reward hacking (Amodei et al., 2016; Krakovna, 2019; Skalse et al., 2022), often interpreting it as
 1060 a manifestation of Goodhart’s Law (Goodhart & Goodhart, 1984): when a proxy metric becomes a
 1061 target for optimization, it ceases to be a good measure. In reinforcement learning, this risk is partic-
 1062 ularly acute because agents can exploit even small imperfections in the reward specification. Pan et
 1063 al. (Pan et al., 2022) further propose a taxonomy of proxy reward misspecification into three types:
 1064 *misweighting*, *ontological*, and *scope* errors.
 1065

1066 To mitigate such risks, several reward-centric methods have been proposed (Hadfield-Menell et al.,
 1067 2017; Ramé et al., 2024). Inverse Reward Design (Hadfield-Menell et al., 2017) aims to infer the in-
 1068 tended true objective from a given proxy and its training context, helping agents generalize without
 1069 exploiting flawed signals. Recent work by Rame et al. (Ramé et al., 2024) averages the parameters
 1070 of multiple reward models to smooth out idiosyncratic errors, reduce the impact of individual proxy
 1071 biases, and demonstrate reduced reward hacking on held-out tests. Another line of defense focuses
 1072 on regularizing policy behavior to reduce sensitivity to reward flaws. Common approaches include
 1073 penalizing divergence from a reference policy using KL-regularization (Liu et al., 2020). Recent
 1074 research by Laidlaw et al. (Laidlaw et al., 2025) proposes Occupancy-Regularized Policy Optimiza-
 1075 tion (ORPO), which applies a χ^2 penalty on the state-action distribution to constrain deviation from
 1076 a baseline policy and reduce exploitative behaviors. Another complementary paradigm is assistance
 1077 games (Fern et al., 2014; Shah et al., 2020), in which human users remain actively involved and
 1078 the agent’s actions complement the user’s to achieve optimal joint performance. Assistance games
 1079 can mitigate reward hacking by removing incentives for deception since the agent’s performance

1080 depends on the human’s latent (true) reward. Recent work has developed scalable assistance-game
 1081 approaches in practice (Laidlaw et al., 2024).
 1082

1083 Overall, existing approaches either attempt to correct the reward specification, regularize against a
 1084 fixed proxy, or explicitly involve human interaction. In contrast, our method trains policies against an
 1085 entire set of plausible proxy rewards, those that remain sufficiently correlated with the true reward,
 1086 offering robustness to a broader range of misspecifications. Moreover, we show that this robust
 1087 training objective can be reformulated as an equivalent regularized optimization problem, providing
 1088 both theoretical and practical benefits.
 1089

1090 D.2 REWARD MODELING IN REINFORCEMENT LEARNING

1091 In standard RL benchmarks, the reward is usually assumed to be given, but real-world applications
 1092 rarely offer a well-defined reward signal upfront. Therefore, designing an effective reward function,
 1093 often referred to as reward modeling, is a critical yet challenging aspect of RL (Skalse et al., 2022;
 1094 Booth et al., 2023; Knox & MacGlashan, 2024). Moreover, evaluating whether a designed reward
 1095 truly captures the designer’s intent is non-trivial. Recent work (Muslimani et al., 2025) has proposed
 1096 to measure reward alignment via a “trajectory alignment coefficient,” which quantifies how closely
 1097 the rankings of trajectories induced by a given reward match a human stakeholder’s preferences.
 1098 Such efforts underscore the importance of conceptual frameworks that go beyond treating the reward
 1099 as a black box, instead focusing on principled reward design and evaluation.

1100 To incorporate domain knowledge and improve interpretability, researchers have explored structured
 1101 or rule-based reward modeling frameworks (Icarte et al., 2022; Brafman et al., 2018; Camacho et al.,
 1102 2017). One prominent example is the use of reward machines (Icarte et al., 2022) that explicitly rep-
 1103 resent the reward function’s logic. A reward machine exposes the internal structure of the reward
 1104 (e.g. different sub-goal phases or conditions) to the agent, enabling techniques like automated re-
 1105 ward shaping and task decomposition for more sample-efficient learning. Even before the advent
 1106 of reward machines, prior works had leveraged logical task specifications to design rewards. For
 1107 instance, translating Linear Temporal Logic(LTL) formulas into automata and rewarding the agent
 1108 upon reaching designated accepting states (Camacho et al., 2017; Brafman et al., 2018). By defin-
 1109 ing rewards through such rules or logical templates, the intended behavior is encoded transparently,
 1110 making the reward function more interpretable. Along similar lines, many approaches assume a
 1111 structured parametric form for the reward function itself to aid transparency (Yu et al., 2025; Mu
 1112 et al., 2024). In particular, it is common to model the reward as a linear combination of feature
 1113 functions, a simplification used in inverse RL and preference-based reward learning to make reward
 1114 inference tractable and explainable (Yu et al., 2025). Recent work in RL from human feedback also
 1115 implements rule-based reward signals as linear models over interpretable features (Mu et al., 2024).
 1116 Our approach follows this tradition: by assuming the reward is linear in a set of human-interpretable
 1117 features, we improve the interpretability of the learned policies and reveal which feature components
 1118 are robust or vulnerable.

1119 D.3 ROBUST REINFORCEMENT LEARNING

1120 Our work is also related to robust reinforcement learning, where the agent assumes the reward
 1121 function (and/or transition dynamics) lies within a given uncertainty set, and it seeks to maximize
 1122 performance against the worst-case realization from that set. This can be formulated as a zero-sum
 1123 dynamic game between the agent and an adversary who selects the most adverse reward or dynamics;
 1124 solving the robust MDP thus involves a challenging max-min optimization (Iyengar, 2005; Nilm
 1125 & El Ghaoui, 2005). To alleviate the computational complexity, early works in this vein rely on
 1126 a rectangularity assumption that is crucial for traceability. Thus, classical robust RL formulation
 1127 typically considers rectangular uncertainty sets on rewards or transition probabilities, which lead to
 1128 conservative solutions but permit efficient algorithms such as robust value iteration (Bagnell et al.,
 1129 2001; Grand-Clément & Kroer, 2021) or modified policy iteration (MPI) (Kaufman & Schaefer,
 1130 2013).

1131 Recent theoretical work has revealed an intimate connection between adversarial robustness and pol-
 1132 icy regularization in the context of rectangular uncertainty sets. Several researchers have shown that
 1133 solving a robust MDP is equivalent to solving a certain regularized RL problem (Derman et al., 2021;
 Eysenbach & Levine; McMahan et al., 2024). In particular, the worst-case effect of the adversary

can often be captured via an additional penalty term in the agent’s objective. Derman et al. (Derman et al., 2021) prove that any entropy- or L^2 -regularized MDP can be interpreted as a robust MDP with uncertain rewards – in fact, a regularized MDP is a special case of a reward-robust MDP. Their analysis establishes a duality between a max-min reward-robust objective and a single-agent maximization of expected reward plus a regularization term. Eysenbach and Levine (Eysenbach & Levine) show that the optimal policy from a maximum entropy RL formulation is provably robust to some adversarial reward perturbations. More recently, these insights have been extended and formalized for general MDPs and even multi-agent settings. McMahan et al. (Mcmahan et al., 2024) study robust Markov games with (s, a) -rectangular uncertainty, and they prove that computing a robust equilibrium is polynomial-time equivalent to computing an equilibrium in a corresponding regularized game. In their framework, the added regularization term is exactly the support function of the uncertainty set, effectively the dual representation of the adversary’s worst-case reward selection. This means that for common uncertainty sets (e.g., those inducing entropy or ℓ_p -norm regularizers), one can replace the inner minimization over rewards with an explicit regularization term in the objective.

The setting in our work departs from the above literature by considering **non-rectangular** reward uncertainty. In particular, we assume a correlation-constrained uncertainty set for the reward function, meaning that the adversary’s permissible deviations in reward are coupled across states. This structure can mitigate the conservativeness of the worst-case solution (the adversary cannot simultaneously push all state rewards to their extreme worst values) (Goyal & Grand-Clement, 2018), but it also means that the neat robustness-regularization duality from the rectangular-case no longer applies and the robust optimization must be solved (or approximated) directly. In summary, our work tackles a form of reward uncertainty which lies beyond the scope of existing robustness-as-regularization analysis.

D.4 SUCCESSOR REPRESENTATIONS IN REINFORCEMENT LEARNING

The linear reward assumption and the use of discounted feature expectations are also closely related to the literature on successor representations/successor features. Successor representations and successor features represent values as inner products between reward weights and discounted occupancies or features of future states and actions. It was introduced as a generalization of the value function (Dayan, 1993). This idea was later generalized by (Gehring, 2015; Barreto et al., 2017; 2018) to handle high-dimensional, continuous state spaces and to use the method for transfer learning. Specifically, (Barreto et al., 2017) formalize this idea into the successor features (SF) framework for transfer learning, assuming that tasks share dynamics but differ only in their reward functions parameterized as linear combinations of features. This yields a value function representation that effectively decouples the environment’s transition dynamics from the reward parameters. (Barreto et al., 2018) further extend successor features to deep reinforcement learning and introduce generalized policy improvement over multiple tasks, demonstrating effective transfer by reusing learned successor features across a family of related tasks.

E PROOFS AND ADDITIONAL THEORETICAL RESULTS

E.1 SOLVE THE MAX-MIN OBJECTIVE

In this section, we show the complete proof for solving the following max-min problem:

$$\max_{\pi} \min_{R \in \mathcal{R}_{\text{corr}}} J(\pi, R) = \max_{\pi} \min_{R \in \mathcal{R}_{\text{corr}}} \mathbb{E}_{(s, a) \sim \mu_{\pi}} [R(s, a)]. \quad (15)$$

where $\mathcal{R}_{\text{corr}}$ is the entire space of rewards that satisfy the correlation constraint with respect to a known proxy reward, as defined below:

$$\begin{aligned} \mathcal{R}_{\text{corr}} &= \left\{ R : (s, a) \rightarrow \mathbb{R} \left| \mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{R - M}{V} \cdot R_{\text{proxy}} \right] = r, J(\pi_{\text{ref}}, R) = M, \sigma_R^2 = V^2 \right. \right\} \quad (16) \\ &= \left\{ R : (s, a) \rightarrow \mathbb{R} \left| \mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{R - M}{V} \cdot R_{\text{proxy}} \right] = r, \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [R] = M, \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [R^2] = V^2 + M^2 \right. \right\}. \end{aligned}$$

1188 M and V denote the fixed mean and standard deviation of the reward function R under the reference
 1189 policy π_{ref} . R_{proxy} is the normalized proxy reward
 1190

$$1191 \quad R_{\text{proxy}}(s, a) := \frac{\tilde{R}_{\text{proxy}}(s, a) - J(\pi_{\text{ref}}, \tilde{R}_{\text{proxy}})}{\sigma_{\tilde{R}_{\text{proxy}}}}$$

$$1192$$

$$1193$$

1194 where \tilde{R}_{proxy} is the original (unnormalized) proxy reward. After normalization, we have
 1195 $J(\pi_{\text{ref}}, R_{\text{proxy}}) = 0$ and $\text{Var}_{\mu_{\pi_{\text{ref}}}}(R_{\text{proxy}}) = 1$.
 1196

1197 To solve the challenge that the objective $\mathbb{E}_{\mu_{\pi}}[R(s, a)]$ depends on the state-action occupancy μ_{π} ,
 1198 whereas the constraints defining $\mathcal{R}_{\text{corr}}$ are expressed in terms of $\mu_{\pi_{\text{ref}}}$. We apply a *change-of-measure*
 1199 technique (Hu & Hong, 2013; Lam, 2016) to rewrite the expectation under $\mu_{\pi_{\text{ref}}}$. Specifically, let
 1200 $L(s, a)$ denote the Radon-Nikodym derivative:
 1201

$$1201 \quad L(s, a) = \frac{\mu_{\pi}(s, a)}{\mu_{\pi_{\text{ref}}}(s, a)}$$

$$1202$$

1203 By definition, $L(s, a) \geq 0$ and $\mathbb{E}_{\mu_{\pi_{\text{ref}}}}[L(s, a)] = 1$. Applying the change-of-measure formula, we
 1204 can express the return as:
 1205

$$1207 \quad \mathbb{E}_{\mu_{\pi}}[R(s, a)] = \int_{\mathcal{S} \times \mathcal{A}} \mu_{\pi}(s, a) R(s, a) d(s, a)$$

$$1208$$

$$1209 \quad = \int_{\mathcal{S} \times \mathcal{A}} \mu_{\pi_{\text{ref}}}(s, a) \frac{\mu_{\pi}(s, a)}{\mu_{\pi_{\text{ref}}}(s, a)} R(s, a) d(s, a)$$

$$1210$$

$$1211 \quad = \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[L(s, a) R(s, a)]$$

$$1212$$

$$1213$$

1214 Thus, both the objective and the constraints can be rewritten as expectations with respect to the
 1215 reference distribution $\mu_{\pi_{\text{ref}}}$. For notational simplicity, we will suppress the variables (s, a) when
 1216 necessary. Under this reparameterization, the max-min objective in Equation 15 can be reformulated
 1217 as:
 1218

$$1218 \quad \max_{\pi} \min_{R \in \mathcal{R}_{\text{corr}}} \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[L \cdot R]. \quad (17)$$

$$1219$$

1220 **Solve Inner Minimization Problem.** The Lagrangian functional associated with the inner mini-
 1221 mization problem of 17 is defined as:
 1222

$$1224 \quad l_0(\lambda_1, \lambda_2, \lambda_3, R) = \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[L \cdot R - \lambda_1 \frac{R - M}{V} \cdot R_{\text{proxy}} - \lambda_2 R - \lambda_3 R^2] + \lambda_1 r + \lambda_2 M + \lambda_3 (M^2 + V^2)$$

$$1225$$

$$1226$$

1227 where $\lambda_1, \lambda_2, \lambda_3$ are the Lagrange multipliers corresponding to the correlation constraint, mean con-
 1228 straint, and variance constraint, respectively. Then the inner minimization problem in Equation 17
 1229 is equivalent to the following problem:
 1230

$$1230 \quad \max_{\lambda_1, \lambda_2, \lambda_3} \min_{R \in \mathcal{R}_{\text{corr}}} l_0(\lambda_1, \lambda_2, \lambda_3, R) \quad (18)$$

$$1231$$

1232 We now solve the inner minimization problem in Equation 18 by finding the optimal R for fixed dual
 1233 variables $(\lambda_1, \lambda_2, \lambda_3)$. Taking the functional derivative of the Lagrangian l_0 with respect to $R(s, a)$
 1234 gives:
 1235

$$1236 \quad \frac{\partial l_0}{\partial R} = \mu_{\pi_{\text{ref}}}(s, a) [(L - \lambda_1 \frac{R_{\text{proxy}}(s, a)}{V} - \lambda_2) - 2\lambda_3 R]$$

$$1237$$

1238 When $\mu_{\pi_{\text{ref}}}(s, a) > 0$, setting the derivative of the Lagrangian to zero yields the optimal adversarial
 1239 reward function:
 1240

$$1241 \quad R^*(s, a) = \frac{L(s, a) - \lambda_1 \frac{R_{\text{proxy}}(s, a)}{V} - \lambda_2}{2\lambda_3} \quad (19)$$

1242 However, for state-action pairs where $\mu_{\pi_{\text{ref}}}(s, a) = 0$, i.e., those not visited under the reference
 1243 policy, the correlation and moment constraints become vacuous. In these regions, the adversarial
 1244 reward $R^*(s, a)$ can be driven arbitrarily poor, reflecting that no constraint prevents the adversary
 1245 from assigning highly penalizing values to rarely visited or unobserved state-action pairs. Never-
 1246 theless, consider the case where $\mu_{\pi_{\text{ref}}}(s, a) > 0$, after substituting the optimal R^* from Equation 19
 1247 into the Lagrangian l_0 in Equation 18 and simplifying, we obtain the following dual objective:
 1248

$$\max_{\lambda_1, \lambda_2, \lambda_3} l_0(\lambda_1, \lambda_2, \lambda_3, R^*) = \mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{(L(s, a) - \lambda_1 \frac{R_{\text{proxy}}(s, a)}{V} - \lambda_2)^2}{4\lambda_3} \right] + \lambda_1 r + \lambda_2 M + \lambda_3 (M^2 + V^2) \quad (20)$$

1251

1252 We now compute the gradients of the dual objective with respect to the dual variables:
 1253

$$\begin{aligned} \frac{\partial l_0}{\partial \lambda_1} &= -\mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{(L - \lambda_1 \frac{R_{\text{proxy}}(s, a)}{V} - \lambda_2) \frac{R_{\text{proxy}}(s, a)}{V}}{2\lambda_3} \right] + r \\ \frac{\partial l_0}{\partial \lambda_2} &= -\mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{L - \lambda_1 \frac{R_{\text{proxy}}(s, a)}{V} - \lambda_2}{2\lambda_3} \right] + M \\ \frac{\partial l_0}{\partial \lambda_3} &= -\mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{(L - \lambda_1 \frac{R_{\text{proxy}}(s, a)}{V} - \lambda_2)^2}{4\lambda_3^2} \right] + M^2 + V^2 \end{aligned} \quad (21)$$

1261

1262 Setting these gradients to zero yields the system of equations:
 1263

$$\mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{(L - \lambda_1 \frac{R_{\text{proxy}}(s, a)}{V} - \lambda_2) \frac{R_{\text{proxy}}(s, a)}{V}}{2\lambda_3} \right] = r, \quad (22)$$

$$\mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{L - \lambda_1 \frac{R_{\text{proxy}}(s, a)}{V} - \lambda_2}{2\lambda_3} \right] = M, \quad (23)$$

$$\mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{(L - \lambda_1 \frac{R_{\text{proxy}}(s, a)}{V} - \lambda_2)^2}{4\lambda_3^2} \right] = M^2 + V^2. \quad (24)$$

1272

1273 Expanding and simplifying each condition:
 1274

Solving for λ_2 : Starting with Equation 23,

$$\mathbb{E}_{\mu_{\pi_{\text{ref}}}}[L] - \lambda_1 \mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{R_{\text{proxy}}(s, a)}{V} \right] - \lambda_2 = 2\lambda_3 M$$

1275 Recall from our normalization that $\mathbb{E}_{\mu_{\pi_{\text{ref}}}}[R_{\text{proxy}}] = 0$. Thus,
 1276

$$\lambda_2 = \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[L] - 2\lambda_3 M$$

1277 Since $L(s, a) = \frac{\mu_{\pi_{\text{ref}}}(s, a)}{\mu_{\pi_{\text{ref}}}(s, a)}$, and using properties of Radon-Nikodym derivatives, we have:
 1278

$$\mathbb{E}_{\mu_{\pi_{\text{ref}}}}[L] = 1$$

1279 Thus, we find:
 1280

$$\boxed{\lambda_2 = 1 - 2\lambda_3 M}$$

1281 **Solving for λ_1 :** Substituting $\lambda_2 = 1 - 2\lambda_3 M$ into Equation 22,
 1282

$$\begin{aligned} 2r\lambda_3 &= \mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\left(L - \lambda_1 \frac{R_{\text{proxy}}}{V} - 1 + 2\lambda_3 M \right) \frac{R_{\text{proxy}}}{V} \right] \\ &= \mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[L \cdot \frac{R_{\text{proxy}}}{V} \right] - \lambda_1 \mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{R_{\text{proxy}}^2}{V^2} \right] - (1 - 2\lambda_3 M) \mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{R_{\text{proxy}}}{V} \right] \end{aligned}$$

1283 Again, using normalization, $\mathbb{E}_{\mu_{\pi_{\text{ref}}}}[R_{\text{proxy}}] = 0$ and $\mathbb{E}_{\mu_{\pi_{\text{ref}}}}[R_{\text{proxy}}^2] = 1$, so we get:
 1284

$$2r\lambda_3 = \mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[L \cdot \frac{R_{\text{proxy}}}{V} \right] - \frac{\lambda_1}{V^2}$$

1285 which rearranges to:
 1286

$$\boxed{\lambda_1 = V \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[L \cdot R_{\text{proxy}}] - 2r\lambda_3 V^2}$$

1296 **Solving for λ_3 :** Substituting $\lambda_2 = 1 - 2\lambda_3 M$ into Equation 24,

$$\begin{aligned} 1298 \quad 4\lambda_3^2(M^2 + V^2) &= \mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[(L - \lambda_1 \frac{R_{\text{proxy}}}{V} - 1 + 2\lambda_3 M)^2 \right] \\ 1299 &= \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L^2] + \lambda_1^2 \mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{R_{\text{proxy}}^2}{V^2} \right] + (1 - 2\lambda_3 M)^2 - 2\lambda_1 \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L \cdot \frac{R_{\text{proxy}}}{V}] \\ 1300 &\quad - 2(1 - 2\lambda_3 M) \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L] + 2\lambda_1(1 - 2\lambda_3 M) \mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{R_{\text{proxy}}}{V} \right] \\ 1301 &\quad - 2(1 - 2\lambda_3 M) \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L] + 2\lambda_1(1 - 2\lambda_3 M) \mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{R_{\text{proxy}}}{V} \right] \end{aligned}$$

1304 Again, using normalization ($\mathbb{E}[R_{\text{proxy}}] = 0$, $\mathbb{E}[R_{\text{proxy}}^2] = 1$, $\mathbb{E}[L] = 1$), this simplifies to:

$$\begin{aligned} 1306 \quad 4\lambda_3^2(M^2 + V^2) &= \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L^2] + \frac{\lambda_1^2}{V^2} - 2\lambda_1 \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L \cdot \frac{R_{\text{proxy}}}{V}] + 4\lambda_3^2 M^2 - 1 \\ 1307 & \\ 1308 \quad 4\lambda_3^2 V^2 &= \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L^2] + \frac{\lambda_1^2}{V^2} - 2\lambda_1 \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L \cdot \frac{R_{\text{proxy}}}{V}] - 1 \\ 1309 & \end{aligned}$$

1310 Now substitute $\lambda_1 = V \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L \cdot R_{\text{proxy}}] - 2r\lambda_3 V^2$ into this expression. After rearrangement and 1311 simplification, we obtain:

$$1313 \quad 4\lambda_3^2(1 - r^2)V^2 = \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L^2] - \mathbb{E}_{\mu_{\pi_{\text{ref}}}}^2 [L \cdot R_{\text{proxy}}] - 1$$

1314 Thus,

$$1315 \quad \lambda_3 = \pm \frac{1}{2} \frac{\sqrt{\mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L^2] - \mathbb{E}_{\mu_{\pi_{\text{ref}}}}^2 [L \cdot R_{\text{proxy}}] - 1}}{V \sqrt{1 - r^2}}$$

1318 We argue that $\lambda_3 < 0$ yields the optimal dual variable. To determine which root maximizes the 1319 above dual objective in Equation 20, we compute the second derivative from Equation 21:

$$1321 \quad \frac{\partial^2 l_0}{\partial \lambda_3^2} = \mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{(L - \lambda_1 \frac{R_{\text{proxy}}(s,a)}{V} - \lambda_2)^2}{2\lambda_3^3} \right]$$

1324 Since the numerator is always non-negative and when $\lambda_3 < 0$, we have $\frac{\partial^2 l_0}{\partial \lambda_3^2} < 0$, which implies 1325 the dual objective is concave in λ_3 around this root. Thus, selecting the negative root yields a local 1326 maximum of the dual objective.

1327 Recognizing that $\mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L^2] - 1$ corresponds to the χ^2 divergence between the occupancy measures:

$$1329 \quad \chi^2(\mu_{\pi} \parallel \mu_{\pi_{\text{ref}}}) = \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L^2] - 1$$

1330 and noting that:

$$1331 \quad \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L \cdot R_{\text{proxy}}] = \mathbb{E}_{\mu_{\pi}} [R_{\text{proxy}}]$$

1332 we can express the solution for λ_3 as:

$$1334 \quad \boxed{\lambda_3 = - \frac{\sqrt{\chi^2(\mu_{\pi} \parallel \mu_{\pi_{\text{ref}}}) - \mathbb{E}_{\mu_{\pi}}^2 [R_{\text{proxy}}]}}{2V \sqrt{1 - r^2}}.} \quad (25)$$

1337 **Solve Outer Maximization Problem.** Now that we have solved for the optimal primal variable 1338 R and dual variables λ_1 , λ_2 , and λ_3 , we plug them back into the original max-min objective in 1339 Equation 17:

$$1341 \quad \max_{\pi} \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L \cdot R] = \max_{\pi} \mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[L(s, a) \cdot \frac{L(s, a) - \lambda_1 \frac{R_{\text{proxy}}(s,a)}{V} - \lambda_2}{2\lambda_3} \right] \quad (26)$$

1344 Using the earlier substitutions:

$$1345 \quad \lambda_1 = V \cdot \mathbb{E}_{\mu_{\pi}} [R_{\text{proxy}}] - 2r\lambda_3 V^2,$$

$$1346 \quad \lambda_2 = 1 - 2\lambda_3 M,$$

$$1348 \quad \lambda_3 = - \frac{1}{2} \cdot \frac{\sqrt{\chi^2(\mu_{\pi} \parallel \mu_{\pi_{\text{ref}}}) - \mathbb{E}_{\mu_{\pi}}^2 [R_{\text{proxy}}]}}{V \sqrt{1 - r^2}},$$

1350

We simplify the expression:

1351

$$\max_{\pi} \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[LR] = \max_{\pi} \frac{1}{2\lambda_3} \left(\mathbb{E}_{\mu_{\pi_{\text{ref}}}}[L^2] - \lambda_1 \mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[L \cdot \frac{R_{\text{proxy}}}{V} \right] - \lambda_2 \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[L] \right)$$

1354

1355

Recall the identities:

1356

$$\begin{aligned} \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[L^2] &= \chi^2(\mu_{\pi} \parallel \mu_{\pi_{\text{ref}}}) + 1, \\ \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[L \cdot R_{\text{proxy}}] &= \mathbb{E}_{\mu_{\pi}}[R_{\text{proxy}}], \\ \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[L] &= 1, \end{aligned}$$

1360

We substitute these and get:

1361

$$\max_{\pi} \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[LR] = \frac{1}{2\lambda_3} \left(\chi^2(\mu_{\pi} \parallel \mu_{\pi_{\text{ref}}}) + 1 - \lambda_1 \cdot \frac{\mathbb{E}_{\mu_{\pi}}[R_{\text{proxy}}]}{V} - \lambda_2 \right)$$

1364

1365

Now substitute the expressions for λ_1 and λ_2 :

1366

$$\begin{aligned} \max_{\pi} \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[LR] &= \max_{\pi} \frac{1}{2\lambda_3} \left(\chi^2 + 1 - (\mathbb{E}_{\mu_{\pi}}[R_{\text{proxy}}] - 2r\lambda_3 V) \cdot \frac{\mathbb{E}_{\mu_{\pi}}[R_{\text{proxy}}]}{V} - (1 - 2\lambda_3 M) \right) \\ &= \max_{\pi} \frac{1}{2\lambda_3} \left(\chi^2 - \frac{\mathbb{E}_{\mu_{\pi}}^2[R_{\text{proxy}}]}{V} + 2r\lambda_3 \mathbb{E}_{\mu_{\pi}}[R_{\text{proxy}}] + 2\lambda_3 M \right) \end{aligned}$$

1371

1372

Now cancel out $2\lambda_3$ in numerator and denominator:

1373

$$\max_{\pi} \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[LR] = \max_{\pi} \mathbb{E}_{\mu_{\pi}}[R_{\text{proxy}}] \cdot r - \frac{1}{2\lambda_3} \cdot \left(\frac{\mathbb{E}_{\mu_{\pi}}^2[R_{\text{proxy}}]}{V} - \chi^2 \right) + M$$

1376

1377

Now plug in the expression for λ_3 :

1378

1379

1380

$$\lambda_3 = -\frac{1}{2} \cdot \frac{\sqrt{\chi^2 - \mathbb{E}_{\mu_{\pi}}^2[R_{\text{proxy}}]}}{V\sqrt{1-r^2}}$$

1381

1382

This gives the final outer problem for the original max-min objective in Equation 15:

1383

1384

1385

$$\boxed{\max_{\pi} r \cdot V \cdot \mathbb{E}_{\mu_{\pi}}[R_{\text{proxy}}] - V \cdot \sqrt{1-r^2} \cdot \sqrt{\chi^2(\mu_{\pi} \parallel \mu_{\pi_{\text{ref}}}) - \mathbb{E}_{\mu_{\pi}}^2[R_{\text{proxy}}]} + M} \quad (27)$$

1386

1387

E.2 PROOF OF OPTIMALITY

1388

1389

Recall the inner minimization problem of our max-min objective in Equation 17:

1390

1391

$$\min_{R \in \mathcal{R}_{\text{corr}}} \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[L \cdot R]$$

1392

1393

where $L = \mu_{\pi}(s, a) / \mu_{\pi_{\text{ref}}}(s, a)$ is treated as fixed, and the feasible set is:

1394

1395

1396

1397

1398

$$\begin{aligned} \mathcal{R}_{\text{corr}} = \left\{ R : (s, a) \rightarrow \mathbb{R} \mid \begin{array}{l} \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[(R - M) \cdot R_{\text{proxy}}] = r \cdot V, \\ \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[R] = M, \quad \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[R^2] = V^2 + M^2 \end{array} \right\} \end{aligned}$$

1399

1400

1401

1402

1403

The feasible region is not convex due to the *quadratic equality* constraint $\mathbb{E}_{\mu_{\pi_{\text{ref}}}}[R^2] = V^2 + M^2$. This defines the boundary of an L^2 ball (a hypersphere) in function space, which is not convex. Therefore, traditional convex programming tools and strong duality do not directly apply.

However, we still claim that the resulting R^* derived in Appendix E.1 is globally optimal. This is supported by the following facts:

1404
 1405 **Stationarity.** When considering R^* for any fixed dual variables $\lambda_1, \lambda_2, \lambda_3$, we are looking at the
 1406 inner minimization problem in Equation 18 as follows:
 1407

$$\min_{R \in \mathcal{R}_{\text{corr}}} l_0(\lambda_1, \lambda_2, \lambda_3, R)$$

1408 The term with $R(s, a)$ in l_0 is:
 1409

$$\mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L \cdot R - \lambda_1 \frac{R - M}{V} \cdot R_{\text{proxy}} - \lambda_2 R - \lambda_3 R^2]$$

1410 For this quadratic in R to have a minimum (since it is a minimization problem for R), the coefficient
 1411 of R^2 must be positive. In our case, the coefficient is $-\lambda_3$. Therefore, for the minimization problem
 1412 to be well-posed and have a finite minimum, we must have $\lambda_3 < 0$. This condition ensures that the
 1413 quadratic term in R is a concave upward parabola, which means that a minimum exists. Moreover, in
 1414 Appendix E.1, we explicitly state that $R^*(s, a)$ is derived by setting the derivative of the Lagrangian
 1415 function l_0 in Equation 18 to zero with respect to $R(s, a)$. Thus, $R^*(s, a)$ is indeed the optimal value
 1416 for the minimization problem for fixed $\lambda_1, \lambda_2, \lambda_3$ where $\lambda_3 < 0$. The Stationarity in this context
 1417 implies that R^* lies within the domain where the Lagrangian is well-defined and differentiable,
 1418 which it does.
 1419

1420
 1421 **Feasibility.** We also argue that the closed-form primal solution $R^*(\lambda^*)$, where λ^* denotes the
 1422 optimal dual solution, is feasible in the original sense, that is, it satisfies the three equality constraints
 1423 in the feasible set $\mathcal{R}_{\text{corr}}$. Specifically, as detailed in Appendix E.1, we substitute R^* back into the
 1424 dual objective l_0 and compute the gradient with respect to each dual variable. We then solve:
 1425

$$\frac{\partial l_0(\lambda_1, \lambda_2, \lambda_3, R^*(\lambda_1, \lambda_2, \lambda_3))}{\partial \lambda_i} = 0, \quad \text{for } i = 1, 2, 3,$$

1426 to find the optimal values $\lambda_1^*, \lambda_2^*, \lambda_3^*$.
 1427

1428 By the chain rule, we have:
 1429

$$\frac{\partial l_0(\lambda_1, \lambda_2, \lambda_3, R^*(\lambda_1, \lambda_2, \lambda_3))}{\partial \lambda_i} = \left\langle \frac{\partial l_0}{\partial R}, \frac{\partial R^*}{\partial \lambda_i} \right\rangle + \frac{\partial l_0}{\partial \lambda_i},$$

1430 where the first term vanishes because R^* is chosen to minimize l_0 for fixed λ (i.e., $\partial l_0 / \partial R = 0$ at
 1431 R^*). Therefore, the derivative simplifies to:
 1432

$$\frac{\partial l_0(\lambda_1, \lambda_2, \lambda_3, R^*)}{\partial \lambda_i} = \frac{\partial l_0}{\partial \lambda_i}.$$

1433 Setting these derivatives to zero yields:
 1434

$$\begin{aligned} \frac{\partial l_0}{\partial \lambda_1} &= -\mathbb{E}_{\mu_{\pi_{\text{ref}}}} [(R^* - M)R_{\text{proxy}}] + rV = 0, \\ \frac{\partial l_0}{\partial \lambda_2} &= -\mathbb{E}_{\mu_{\pi_{\text{ref}}}} [R^*] + M = 0, \\ \frac{\partial l_0}{\partial \lambda_3} &= -\mathbb{E}_{\mu_{\pi_{\text{ref}}}} [(R^*)^2] + V^2 + M^2 = 0, \end{aligned}$$

1435 which exactly recover the original feasibility constraints. Hence, the solution $R^*(\lambda^*)$ is feasible by
 1436 construction.
 1437

1438 **Therefore, R^* satisfies both stationarity and feasibility. In general, stationarity and feasibility are**
 1439 **not sufficient for global optimality when the feasible set is nonconvex. In our case, however, global**
 1440 **optimality does hold, relying on the specific structure of the inner problem.**

1441 Recall the inner minimization problem discussed above, and we work in the Hilbert space $\mathcal{H} =$
 1442 $L^2(\mu_{\pi_{\text{ref}}})$. Using the normalization assumptions: $\mathbb{E}_{\mu_{\pi_{\text{ref}}}} [R_{\text{proxy}}] = 0$ and $\mathbb{E}_{\mu_{\pi_{\text{ref}}}} [R_{\text{proxy}}^2] = 1$, the
 1443 constraints can be rewritten as inner products in \mathcal{H} :
 1444

- 1445 • $\langle R, \mathbf{1} \rangle = M$ (mean constraint)
- 1446 • $\langle R, R_{\text{proxy}} \rangle = rV$ (correlation constraint)

1458 • $\|R\|_2^2 = V^2 + M^2$ (norm constraint).

1459

1460 Let $\{e_0, e_1, e_2, \dots\}$ be an orthonormal basis of \mathcal{H} , where

1461

- e_0 is proportional to the constant function $\mathbf{1}$
- $e_1 = R_{\text{proxy}}$,
- and $\{e_k\}_{k \geq 2}$, spans the orthogonal complement of $\text{span}\{\mathbf{1}, R_{\text{proxy}}\}$.

1462

1463 e_0 and e_1 is orthonormal because $\mathbb{E}_{\mu_{\pi_{\text{ref}}}}[R_{\text{proxy}}] = 0$. Expanding

1464

1465

$$R = \alpha_0 e_0 + \alpha_1 e_1 + \sum_{k \geq 2} \alpha_k e_k$$

1466

1467 Notice that the mean constraint and correlation constraints uniquely fix α_0 and α_1 . The norm constraint then forces:

1468

$$\sum_k \alpha_k = \rho^2$$

1469

1470 for some constant radius $\rho > 0$. Hence the remaining degrees of freedom lie on a sphere in the
1471 subspace orthogonal to $\mathbf{1}$ and R_{proxy} . This is to say, although $\mathcal{R}_{\text{corr}}$ is not convex in the ambient
1472 space, it is a spherical manifold (the boundary of an L^2 -ball intersected with an affine subspace),
1473 which is compact and smooth. Moreover, the objective is linear in R :

1474

1475

$$\mathbb{E}_{\mu_{\pi_{\text{ref}}}}[L \cdot R] = \langle L, R \rangle = \text{const} + \langle L', R' \rangle$$

1476

1477 where L' is the projection of $L = \frac{\mu_{\pi}}{\mu_{\pi_{\text{ref}}}}$ onto the subspace spanned by $\{e_k\}_{k \geq 2}$ and $R' =$
1478 $\sum_{k \geq 2} \alpha_k e_k$. Therefore the optimization reduces to

1479

1480

$$\min_{\|R'\|_2 = \rho} \langle L', R' \rangle$$

1481

1482 This is simply minimizing a linear function over a Euclidean sphere. In this setting, it is well-known
1483 that the only stationary points of a linear functional on a sphere are its global maximum and global
1484 minimum. There are no other local minima or saddle points. Thus, on this particular nonconvex
1485 feasible set, **any feasible stationary point is automatically a global optimizer**.

1486

1487 In summary, our previous analysis shows that:

1488

1. For fixed $(\lambda_1, \lambda_2, \lambda_3)$ with $\lambda_3 < 0$, the Lagrangian is a strictly convex quadratic in R , so
1489 its stationary point $R^*(\lambda)$ is the unique global minimizer of the inner problem with those
1490 multipliers.
2. Solving the dual and enforcing feasibility recovers the specific choice of multipliers λ^* for
1491 which $R^*(\lambda^*)$ lies on the sphere defined by the norm constraint.
3. Because the reduced problem is linear over a sphere, this feasible stationary point $R^*(\lambda^*)$
1492 must be the global minimizer of the original inner problem.

1493

1501 E.3 PROOF THAT $\chi^2(\mu_{\pi} \parallel \mu_{\pi_{\text{ref}}}) \geq \mathbb{E}_{\mu_{\pi}}^2[R_{\text{proxy}}]$

1502

1503 To ensure that the inner term of the square root in Equation 25 remains non-negative, we need to
1504 show that

1505

$$\chi^2(\mu_{\pi} \parallel \mu_{\pi_{\text{ref}}}) \geq \mathbb{E}_{\mu_{\pi}}^2[R_{\text{proxy}}]$$

1506

1507 **Proof.** Recall that

1508

$$\mathbb{E}_{\mu_{\pi}}[R_{\text{proxy}}] = \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[L \cdot R_{\text{proxy}}],$$

1509

1510 where $L(s, a) = \frac{\mu_{\pi}(s, a)}{\mu_{\pi_{\text{ref}}}(s, a)}$ is the Radon-Nikodym derivative. Since R_{proxy} is normalized to have zero
1511 mean under $\mu_{\pi_{\text{ref}}}$, we have:

1512

$$\mathbb{E}_{\mu_{\pi_{\text{ref}}}}[R_{\text{proxy}}] = 0.$$

1512 Thus,

1513

$$\begin{aligned} 1514 \mathbb{E}_{\mu_\pi}[R_{\text{proxy}}] &= \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[R_{\text{proxy}}(s, a)(L(s, a) - 1)] \\ 1515 &= \sum_{(s, a)} R_{\text{proxy}}(s, a)\mu_{\pi_{\text{ref}}}(s, a)(L(s, a) - 1) \\ 1516 \\ 1517 \end{aligned}$$

1518 Applying the Cauchy-Schwarz inequality:

1519

$$\left(\sum_{(s, a)} R_{\text{proxy}}(s, a)\mu_{\pi_{\text{ref}}}(s, a)(L(s, a) - 1) \right)^2 \leq \left(\sum_{(s, a)} \mu_{\pi_{\text{ref}}}(s, a)R_{\text{proxy}}^2(s, a) \right) \left(\sum_{(s, a)} \mu_{\pi_{\text{ref}}}(s, a)(L(s, a) - 1)^2 \right)$$

1520

1521 By the assumptions: $\mathbb{E}_{\mu_{\pi_{\text{ref}}}}[R_{\text{proxy}}^2] = 1$, $\sum_{(s, a)} \mu_{\pi_{\text{ref}}}(s, a)(L(s, a) - 1)^2 = \chi^2(\mu_\pi \parallel \mu_{\pi_{\text{ref}}})$.

1522 We obtain:

1523

$$\mathbb{E}_{\mu_\pi}^2[R_{\text{proxy}}] \leq \chi^2(\mu_\pi \parallel \mu_{\pi_{\text{ref}}})$$

1524

1525 as desired.

1526

1527 E.4 DERIVE LAGRANGIAN FUNCTIONAL FOR LINEAR MAX-MIN OBJECTIVE

1528 Recall that our max-min optimization under the structured reward assumption is as follows:

1529

$$\max_{\pi} \min_{\boldsymbol{\theta} \in \mathcal{R}_{\text{corr}}^{\text{lin}}, \boldsymbol{\theta} \geq 0} \mathbb{E}_{(s, a) \sim \mu_\pi} [\boldsymbol{\theta}^\top \boldsymbol{\phi}(s, a)]. \quad (28)$$

1530

1531 where $\mathcal{R}_{\text{corr}}^{\text{lin}}$ is the uncertainty set defined as follow:

1532

1533

$$\mathcal{R}_{\text{corr}}^{\text{lin}} = \left\{ \boldsymbol{\theta} \in \mathbb{R}^M \mid \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[\boldsymbol{\theta}^\top \boldsymbol{\phi} \cdot R_{\text{proxy}}] = r, \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[\boldsymbol{\theta}^\top \boldsymbol{\phi}] = 0, \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[(\boldsymbol{\theta}^\top \boldsymbol{\phi})^2] = 1 \right\}. \quad (29)$$

1534

1535 We assume without loss of generality that the worst-case reward $R(s, a) = \boldsymbol{\theta}^\top \boldsymbol{\phi}(s, a)$ is normalized
1536 to have zero mean and unit variance under the reference policy π_{ref} . This corresponds to setting $M =$
1537 0 and $V = 1$, which, as shown in our earlier derivation, does not affect the resulting optimal policy.
1538 As before, R_{proxy} denotes the normalized proxy reward under π_{ref} , satisfying $\mathbb{E}_{\mu_{\pi_{\text{ref}}}}[R_{\text{proxy}}] = 0$ and
1539 $\text{Var}_{\mu_{\pi_{\text{ref}}}}[R_{\text{proxy}}] = 1$.

1540

1541 Similar to previous steps, we introduce the Radon-Nikodym derivative

1542

$$L(s, a) = \frac{\mu_\pi(s, a)}{\mu_{\pi_{\text{ref}}}(s, a)}$$

1543

1544

1545 We use a change-of-measure, and define the Lagrangian functional for the inner minimization in
1546 Equation 28 as:

1547

$$\begin{aligned} 1548 l_1(\lambda_1, \lambda_2, \lambda_3, \boldsymbol{\theta}) &= \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[L \cdot \boldsymbol{\theta}^\top \boldsymbol{\phi}] - \lambda_1 \left(\mathbb{E}_{\mu_{\pi_{\text{ref}}}}[R_{\text{proxy}} \cdot \boldsymbol{\theta}^\top \boldsymbol{\phi}] - r \right) \\ 1549 &\quad - \lambda_2 \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[\boldsymbol{\theta}^\top \boldsymbol{\phi}] - \lambda_3 \left(\mathbb{E}_{\mu_{\pi_{\text{ref}}}}[(\boldsymbol{\theta}^\top \boldsymbol{\phi})^2] - 1 \right) \\ 1550 &= \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[(L - \lambda_1 R_{\text{proxy}} - \lambda_2) \boldsymbol{\theta}^\top \boldsymbol{\phi}] - \lambda_3 \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[(\boldsymbol{\theta}^\top \boldsymbol{\phi})^2] + \lambda_1 r + \lambda_3 \\ 1551 &= \sum_{(s, a)} \mu_{\pi_{\text{ref}}}(s, a) [(L - \lambda_1 R_{\text{proxy}} - \lambda_2) \boldsymbol{\theta}^\top \boldsymbol{\phi} - (\boldsymbol{\theta}^\top \boldsymbol{\phi})^2] + \lambda_1 r + \lambda_3 \end{aligned}$$

1552

1553

1554 Define the following terms for simplicity:

1555

$$\begin{aligned} 1556 v(s, a) &= \mu_\pi(s, a) \\ 1557 D(s, a) &= \mu_{\pi_{\text{ref}}}(s, a) \cdot R_{\text{proxy}}(s, a) \\ 1558 C(s, a) &= \mu_{\pi_{\text{ref}}}(s, a) \\ 1559 u_{\lambda_1, \lambda_2}(s, a) &= v(s, a) - \lambda_1 D(s, a) - \lambda_2 C(s, a) \end{aligned}$$

1560

1566 Then the Lagrangian function simplifies to:
 1567
 1568 $l_1(\lambda_1, \lambda_2, \lambda_3, \theta) = \sum_{(s,a)} [u_{\lambda_1, \lambda_2}(s, a)\theta^\top \phi(s, a) - \lambda_3 C(s, a)(\theta^\top \phi(s, a))^2] + \lambda_1 r + \lambda_3$
 1569
 1570 $= \theta^\top \left(\sum_{(s,a)} u_{\lambda_1, \lambda_2}(s, a)\phi(s, a) \right) - \lambda_3 \theta^\top \left(\sum_{(s,a)} C(s, a)\phi(s, a)\phi(s, a)^\top \right) \theta + \lambda_1 r + \lambda_3$
 1571
 1572
 1573

1574 where we expand the quadratic term:
 1575

$$\begin{aligned} 1576 \sum_{(s,a)} C(s, a)(\theta^\top \phi(s, a))^2 &= \sum_{(s,a)} C(s, a)(\phi(s, a)^\top \theta)^2 \\ 1577 &= \sum_{(s,a)} C(s, a)\theta^\top \phi(s, a)\phi(s, a)^\top \theta \\ 1578 &= \theta^\top \left(\sum_{(s,a)} C(s, a)\phi(s, a)\phi(s, a)^\top \right) \theta \\ 1579 & \\ 1580 & \\ 1581 & \\ 1582 & \\ 1583 & \\ 1584 & \end{aligned}$$

1585 Let

$$Q = \sum_{(s,a)} C(s, a)\phi(s, a)\phi(s, a)^\top \quad (30)$$

1586 then we can write the Lagrangian function as:
 1587

$$1588 l_1(\lambda_1, \lambda_2, \lambda_3, \theta) = \theta^\top \left(\sum_{(s,a)} u_{\lambda_1, \lambda_2}(s, a)\phi(s, a) \right) - \lambda_3 \theta^\top Q \theta + \lambda_1 r + \lambda_3$$

1589
 1590 And the inner minimization problem in Equation 28 becomes:
 1591
 1592

$$1593 \max_{\lambda_1, \lambda_2, \lambda_3} \min_{\theta \geq 0} l_1(\lambda_1, \lambda_2, \lambda_3, R) \quad (31)$$

1594 E.5 PROOF FOR WHITENING TRANSFORMATION

1595 To simplify the problem associated with the Lagrangian function above, we transform the feature
 1596 vector ϕ into a whitened version $\tilde{\phi}$ such that the matrix Q as defined in Equation 30 becomes the
 1597 identity matrix I . Specifically, we perform a whitening transformation using the Cholesky decom-
 1598 position (Boyd & Vandenberghe, 2004). Let
 1599

$$1600 W = Q^{-\frac{1}{2}}, \quad \tilde{\phi}(s, a) = W\phi(s, a)$$

1601 where $Q^{-\frac{1}{2}}$ denotes a matrix square root of Q^{-1} . Then we have:
 1602

$$\begin{aligned} 1603 \sum_{(s,a)} C(s, a)\tilde{\phi}(s, a)\tilde{\phi}(s, a)^\top &= \sum_{(s,a)} C(s, a)(W\phi(s, a))(W\phi(s, a))^\top \\ 1604 &= \sum_{(s,a)} C(s, a)W\phi(s, a)\phi(s, a)^\top W^\top \\ 1605 &= W \left(\sum_{(s,a)} C(s, a)\phi(s, a)\phi(s, a)^\top \right) W^\top \\ 1606 &= W Q W^\top \\ 1607 &= Q^{-\frac{1}{2}} Q Q^{-\frac{1}{2}} \\ 1608 &= I \end{aligned}$$

1609 as desired.
 1610
 1611
 1612
 1613
 1614
 1615
 1616
 1617
 1618
 1619

1620 Note that the whitening step requires Q to be invertible so that $Q^{-1/2}$ (and hence Q^{-1}) exists.
 1621 It holds when Q is positive semi-definite and non-singular. Q is positive semi-definite since it is a
 1622 sum of outer products $\phi(s, a)\phi(s, a)^\top$ weighted by non-negative coefficients (occupancy measure of
 1623 $\pi_{\text{ref}} \geq 0$). For Q to be non-singular, it is necessary that the span of $\{\phi(s, a) : \mu_{\pi_{\text{ref}}}(s, a) > 0\}$ covers
 1624 \mathbb{R}^n , i.e., the features associated with state-action pairs visited by π_{ref} must span the full feature space.
 1625 To achieve these conditions, the reference policy should visit a diverse and representative subset of
 1626 the state-action space with non-trivial occupancy. This is more likely when π_{ref} is derived from ei-
 1627 ther expert demonstrations that exhibit rich behavior or from stochastic or exploratory policies (e.g.,
 1628 entropy-regularized policies or policies trained with exploration bonuses). Moreover, the feature
 1629 mapping $\phi(s, a)$ must exhibit sufficient variation across the visited state-action pairs. This typically
 1630 holds when ϕ encodes task-relevant dynamics (e.g., learned embeddings or expressive hand-crafted
 1631 features) and when π_{ref} does not collapse to trivial or deterministic behavior. In our experiments
 1632 (Appendix F.4), the reference policies for the Traffic and Pandemic environments are trained via
 1633 behavioral cloning on large, diverse trajectories generated by human experts or hand-crafted con-
 1634 trollers. The feature representations used in these environments, such as velocity, acceleration, and
 1635 headway in Traffic, and infection level, disease stage, and smooth transitions in Pandemic, encode
 1636 meaningful task-relevant dynamics. These demonstrations cover a wide range of task-relevant be-
 1637 haviors, and the induced occupancy over state-action pairs spans a high-dimensional subspace of the
 1638 feature space. We empirically verified that the resulting Q matrices in our experiments are full-rank
 1639 and numerically well-conditioned. Though ensuring sufficient coverage of the feature space by the
 1640 reference policy is generally challenging in practice.

1641 E.6 DERIVE OPTIMAL PRIMAL VARIABLE FOR LINEAR MAX-MIN OBJECTIVE

1642 After whitening transformation as discussed in Appendix E.5, the problem in Equation 31 becomes:
 1643

$$1644 \max_{\lambda_1, \lambda_2, \lambda_3} \min_{\tilde{\theta} \geq 0} l_1(\lambda_1, \lambda_2, \lambda_3, \tilde{\theta}) = \tilde{\theta}^\top \left(\sum_{(s, a)} u_{\lambda_1, \lambda_2}(s, a) \tilde{\phi}(s, a) \right) - \lambda_3 \tilde{\theta}^\top \tilde{\theta} + \lambda_1 r + \lambda_3. \quad (32)$$

1644 where we now optimize over the parameter $\tilde{\theta}$ using the transformed features $\tilde{\phi}$. For notational
 1645 simplicity, we will drop the tilde and henceforth use ϕ to represent the whitened feature $\tilde{\phi}$, and θ to
 1646 represent the whitened weights $\tilde{\theta}$.
 1647

1648 **Separable Structure.** In the whitened feature space, the objective becomes separable across co-
 1649 ordinates of θ . Thus, the inner minimization problem in Equation 32 decouples into M independent
 1650 one-dimensional convex minimization problems, one for each feature coordinate $i \in \{1, 2, \dots, M\}$:
 1651

$$1652 \min_{\theta_i \geq 0} \left(\sum_{(s, a)} u_{\lambda_1, \lambda_2}(s, a) \phi_i(s, a) \right) \theta_i - \lambda_3 \theta_i^2$$

1653 Let us solve the i -th subproblem. Assuming $\lambda_3 < 0$, the objective is a convex quadratic function in
 1654 θ_i (an upward-opening parabola). The unconstrained minimum occurs at:
 1655

$$1656 \theta_i^* = -\frac{\sum_{(s, a)} u_{\lambda_1, \lambda_2}(s, a) \phi_i(s, a)}{2\lambda_3}$$

1657 Considering the constraint $\theta_i \geq 0$, we have two cases:
 1658

- 1659 • If the unconstrained minimum $\theta_i^* \geq 0$, then it is also the solution to the constrained prob-
 1660 lem.
- 1661 • If $\theta_i^* < 0$, then the constrained minimum occurs at the boundary $\theta_i = 0$.

1662 Thus, the final optimal θ_i^* is:
 1663

$$1664 \theta_i^* = \max \left(0, -\frac{\sum_{(s, a)} u_{\lambda_1, \lambda_2}(s, a) \phi_i(s, a)}{2\lambda_3} \right)$$

1674 Collecting across all i , we express the final optimal solution θ^* as:
 1675

$$1676 \theta^* = \max \left(0, -\frac{\sum_{(s,a)} u_{\lambda_1, \lambda_2}(s, a) \phi(s, a)}{2\lambda_3} \right) \quad (33)$$

1677
 1678

1679 where the $\max(\cdot, 0)$ is applied elementwise.
 1680

1681 E.7 SOLVE THE DUAL OBJECTIVE FOR LINEAR MAX-MIN OBJECTIVE

1682

1683 Let the outer objective in Equation 32 be:
 1684

$$1685 g(\lambda_1, \lambda_2, \lambda_3) = l_1(\lambda_1, \lambda_2, \lambda_3, \theta^*)$$

1686

1687 Then we want to solve the following dual objective:
 1688

$$\max_{\lambda_1, \lambda_2, \lambda_3} g(\lambda_1, \lambda_2, \lambda_3) \quad (34)$$

1689

1690 Let

$$1691 q_j(\lambda_1, \lambda_2) = \sum_{(s,a)} (v(s, a) - \lambda_1 D(s, a) - \lambda_2 C(s, a)) \phi_j(s, a)$$

1692
 1693

1694 denote the linear coefficient for each feature $j \in \{1, \dots, M\}$.
 1695

1696 The optimal θ_j^* is:

$$1697 \theta_j^*(\lambda) = \max \left(0, \frac{q_j(\lambda_1, \lambda_2)}{2\lambda_3} \right)$$

1698
 1699

1700 Now, we compute the gradients:
 1701

1702 Gradient with respect to λ_1 :

$$1703 \frac{\partial g}{\partial \lambda_1}(\lambda) = \frac{\partial l_1}{\partial \lambda_1}(\lambda, \theta^*(\lambda))$$

1704

$$1705 = \sum_{(s,a)} (-D(s, a)(\theta^{*T} \phi(s, a))) + r$$

1706

$$1707 = r - \sum_{j=1}^M D_{\phi,j} \cdot \theta_j^*(\lambda)$$

1708

1709

1710

1711 where

$$1713 D_{\phi,j} = \sum_{(s,a)} D(s, a) \phi_j(s, a)$$

1714

1715

1716 Gradient with respect to λ_2 :

$$1717 \frac{\partial g}{\partial \lambda_2}(\lambda) = \frac{\partial l_1}{\partial \lambda_2}(\lambda, \theta^*(\lambda))$$

1718

$$1719 = \sum_{(s,a)} (-C(s, a)(\theta^{*T} \phi(s, a)))$$

1720

$$1721 = - \sum_{j=1}^M C_{\phi,j} \cdot \theta_j^*(\lambda)$$

1722

1723

1724

1725 where

$$1727 C_{\phi,j} = \sum_{(s,a)} C(s, a) \phi_j(s, a)$$

1728

1728 **Gradient with respect to λ_3 :**

$$\begin{aligned} \frac{\partial g}{\partial \lambda_3}(\lambda) &= \frac{\partial l_1}{\partial \lambda_3}(\lambda, \boldsymbol{\theta}^*(\lambda)) \\ &= \sum_{(s,a)} (-C(s,a)(\boldsymbol{\theta}^{*T}\boldsymbol{\phi}(s,a))^2) + 1 \\ &= 1 - \sum_{j=1}^M (\theta_j^*(\lambda))^2 \end{aligned}$$

1737 where we use the whitening assumption $\sum_{(s,a)} C(s,a)\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top = I$.

1739 Thus, the full gradients are:

$$\begin{aligned} \frac{\partial g}{\partial \lambda_1}(\lambda) &= r - \sum_{j=1}^M D_{\phi,j} \cdot \max\left(0, \frac{q_j(\lambda_1, \lambda_2)}{2\lambda_3}\right) \\ \frac{\partial g}{\partial \lambda_2}(\lambda) &= - \sum_{j=1}^M C_{\phi,j} \cdot \max\left(0, \frac{q_j(\lambda_1, \lambda_2)}{2\lambda_3}\right) \\ \frac{\partial g}{\partial \lambda_3}(\lambda) &= 1 - \sum_{j=1}^M \left(\max\left(0, \frac{q_j(\lambda_1, \lambda_2)}{2\lambda_3}\right)\right)^2 \end{aligned}$$

1751 We can solve for the optimal dual variables $(\lambda_1, \lambda_2, \lambda_3)$ using standard first-order optimization
1752 methods. Since $g(\lambda)$ is concave (under the condition $\lambda_3 < 0$), optimization is well-behaved and
1753 converges reliably. After obtaining the optimal primal variables $\boldsymbol{\theta}^*$ and dual variables $(\lambda_1^*, \lambda_2^*, \lambda_3^*)$,
1754 we can substitute them back into Equation 28 and solve the outer maximization over the policy π
1755 using standard reinforcement learning algorithms, such as PPO (Schulman et al., 2017).

1757 E.8 POLICY GRADIENT DERIVATION

1759 We now derive the gradient of the robust objective equation 27 with respect to the policy parameters
1760 θ . Recall that the robust objective is:

$$\begin{aligned} \mathcal{J}(\mu_{\pi_\theta}) &= r \cdot V \cdot \mathbb{E}_{\mu_{\pi_\theta}}[R_{\text{proxy}}] - V \cdot \sqrt{1-r^2} \cdot \sqrt{\chi^2(\mu_\pi \parallel \mu_{\pi_{\text{ref}}}) - \left(\mathbb{E}_{\mu_{\pi_\theta}}[R_{\text{proxy}}]\right)^2} + M \\ &= \mathbb{E}_{\mu_{\pi_\theta}}[R_{\text{proxy}}] - \frac{\sqrt{1-r^2}}{r} \sqrt{\chi^2(\mu_{\pi_\theta} \parallel \mu_{\pi_{\text{ref}}}) - \left(\mathbb{E}_{\mu_{\pi_\theta}}[R_{\text{proxy}}]\right)^2} \end{aligned}$$

1766 where we set $M = 0$ and $V = 1$ without loss of generality. We also divide the entire objective by r ,
1767 which is assumed to be positive ($r > 0$), so this rescaling preserves the optimization direction and
1768 does not affect the final policy solution. The χ^2 divergence is defined as:

$$\chi^2(\mu_{\pi_\theta} \parallel \mu_{\pi_{\text{ref}}}) = \sum_{(s,a)} \frac{\mu_{\pi_\theta}(s,a)^2}{\mu_{\pi_{\text{ref}}}(s,a)} - 1$$

1772 Applying the chain rule, we compute:

$$\nabla_\theta \mathcal{J} = \nabla_\theta \mathbb{E}_{\mu_{\pi_\theta}}[R_{\text{proxy}}] - \frac{\sqrt{1-r^2}}{r} \nabla_\theta \left(\sqrt{h(\mu_{\pi_\theta})} \right) \quad (35)$$

1776 where we define:

$$h(\mu_{\pi_\theta}) = \chi^2(\mu_{\pi_\theta} \parallel \mu_{\pi_{\text{ref}}}) - \left(\mathbb{E}_{\mu_{\pi_\theta}}[R_{\text{proxy}}]\right)^2$$

1779 Using the chain rule again:

$$\nabla_\theta \sqrt{h(\mu_{\pi_\theta})} = \frac{1}{2\sqrt{h(\mu_{\pi_\theta})}} \nabla_\theta h(\mu_{\pi_\theta})$$

1782 Now compute $\nabla_\theta h(\mu_{\pi_\theta})$:
 1783

$$\begin{aligned} 1784 \nabla_\theta h(\mu_{\pi_\theta}) &= \nabla_\theta \chi^2(\mu_{\pi_\theta} \| \mu_{\pi_{\text{ref}}}) - \nabla_\theta \left(\mathbb{E}_{\mu_{\pi_\theta}} [R_{\text{proxy}}]^2 \right) \\ 1785 &= \nabla_\theta \left(\sum_{(s,a)} \frac{\mu_{\pi_\theta}(s,a)^2}{\mu_{\pi_{\text{ref}}}(s,a)} - 1 \right) - 2\mathbb{E}_{\mu_{\pi_\theta}} [R_{\text{proxy}}] \nabla_\theta \mathbb{E}_{\mu_{\pi_\theta}} [R_{\text{proxy}}] \\ 1786 \\ 1787 \\ 1788 \end{aligned}$$

1789 The individual terms are:
 1790

$$\begin{aligned} 1791 \nabla_\theta \left(\sum_{(s,a)} \frac{\mu_{\pi_\theta}(s,a)^2}{\mu_{\pi_{\text{ref}}}(s,a)} - 1 \right) &= 2 \sum_{(s,a)} \frac{\mu_{\pi_\theta}(s,a)}{\mu_{\pi_{\text{ref}}}(s,a)} \nabla_\theta \mu_{\pi_\theta}(s,a) \\ 1792 \\ 1793 \\ 1794 \nabla_\theta \mathbb{E}_{\mu_{\pi_\theta}} [R_{\text{proxy}}] &= \sum_{(s,a)} \nabla_\theta \mu_{\pi_\theta}(s,a) R_{\text{proxy}}(s,a) \\ 1795 \\ 1796 \end{aligned}$$

1797 Thus:
 1798

$$\nabla_\theta h(\mu_{\pi_\theta}) = \sum_{(s,a)} \nabla_\theta \mu_{\pi_\theta}(s,a) \left(2 \frac{\mu_{\pi_\theta}(s,a)}{\mu_{\pi_{\text{ref}}}(s,a)} - 2\mathbb{E}_{\mu_{\pi_\theta}} [R_{\text{proxy}}] R_{\text{proxy}}(s,a) \right)$$

1801 Then we compute $\nabla_\theta \mathbb{E}_{\mu_{\pi_\theta}} [R_{\text{proxy}}]$:
 1802

$$\begin{aligned} 1803 \nabla_\theta \mathbb{E}_{\mu_{\pi_\theta}} [R_{\text{proxy}}] &= \nabla_\theta \sum_{(s,a)} \mu_{\pi_\theta}(s,a) R_{\text{proxy}}(s,a) \\ 1804 &= \sum_{(s,a)} \nabla_\theta \mu_{\pi_\theta}(s,a) R_{\text{proxy}}(s,a) \\ 1805 \\ 1806 \\ 1807 \end{aligned}$$

1808 Put them together, we get the final gradient in Equation 35 as:
 1809

$$\begin{aligned} 1810 \nabla_\theta \mathcal{J} &= \nabla_\theta \mathbb{E}_{\mu_{\pi_\theta}} [R_{\text{proxy}}] - \frac{\sqrt{1-r^2}}{r} \nabla_\theta \left(\sqrt{h(\mu_{\pi_\theta})} \right) \\ 1811 &= \sum_{(s,a)} \nabla_\theta \mu_{\pi_\theta}(s,a) R_{\text{proxy}}(s,a) - \frac{\sqrt{1-r^2}}{r} \frac{1}{2\sqrt{h(\mu_{\pi_\theta})}} \nabla_\theta h(\mu_{\pi_\theta}) \\ 1812 \\ 1813 &= \sum_{(s,a)} \nabla_\theta \mu_{\pi_\theta}(s,a) R_{\text{proxy}}(s,a) \\ 1814 \\ 1815 &\quad - \frac{\sqrt{1-r^2}}{r} \frac{1}{2\sqrt{h(\mu_{\pi_\theta})}} \sum_{(s,a)} \nabla_\theta \mu_{\pi_\theta}(s,a) \left(2 \frac{\mu_{\pi_\theta}(s,a)}{\mu_{\pi_{\text{ref}}}(s,a)} - 2\mathbb{E}_{\mu_{\pi_\theta}} [R_{\text{proxy}}] R_{\text{proxy}}(s,a) \right) \\ 1816 \\ 1817 &= \sum_{(s,a)} \nabla_\theta \mu_{\pi_\theta}(s,a) \left[R_{\text{proxy}} - \frac{\sqrt{1-r^2}}{r} \frac{1}{\sqrt{h(\mu_{\pi_\theta})}} \left(\frac{\mu_{\pi_\theta}(s,a)}{\mu_{\pi_{\text{ref}}}(s,a)} - \mathbb{E}_{\mu_{\pi_\theta}} [R_{\text{proxy}}] R_{\text{proxy}}(s,a) \right) \right] \\ 1818 \\ 1819 \\ 1820 \\ 1821 \\ 1822 \\ 1823 \\ 1824 \\ 1825 \\ 1826 \\ 1827 \end{aligned} \tag{36}$$

The full policy gradient for the ORPO algorithm, as presented in Appendix B of (Laidlaw et al., 2025), is given by:

$$\sum_{(s,a)} \nabla_\theta \mu_{\pi_\theta}(s,a) \left[R_{\text{proxy}}(s,a) - \frac{\lambda}{\sqrt{\chi^2(\mu_{\pi_\theta} \| \mu_{\pi_{\text{ref}}})}} \cdot \frac{\mu_{\pi_\theta}(s,a)}{\mu_{\pi_{\text{ref}}}(s,a)} \right]. \tag{37}$$

1828 **Interpretation.** The policy gradient consists of two terms:
 1829

- 1830 • A standard term encouraging the policy to increase $R_{\text{proxy}}(s,a)$.
 1831
- 1832 • A correction term that penalizes deviations from the reference occupancy $\mu_{\pi_{\text{ref}}}$, while also
 1833 adjusting for alignment with the proxy reward.
 1834

1836 This correction enforces robustness to potential reward hacking by optimizing against adversarially
 1837 misaligned interpretations of the proxy reward.
 1838

1839 Notice that our derived policy gradient in Equation 36 shares structural similarities with ORPO but
 1840 is rooted in a formal robust optimization framework. Unlike ORPO, our formulation introduces
 1841 an additional correction term involving both the occupancy ratio and the expected proxy reward,
 1842 capturing how the proxy is aligned with the current policy’s behavior. This structure more explicitly
 1843 penalizes the combination of distributional shift and proxy overoptimization, discouraging policies
 1844 from exploiting proxy-specific artifacts. Both methods share the goal of improving robustness, but
 1845 our approach is derived from first principles by directly optimizing for worst-case performance over
 1846 a correlation-constrained uncertainty set.
 1847

E.9 PROOF OF THEOREM 1

1849 *Proof.* For any reward function R , define the performance difference

$$\Delta J(\pi, R) := J(\pi, R) - J(\pi_{\text{ref}}, R).$$

1850 By definition of the correlated uncertainty set, our distributionally robust objective considers

$$F(\pi) := \min_{R \in \mathcal{R}_{\text{corr}}} \Delta J(\pi, R).$$

1851 Under the assumptions on the correlation, mean, and variance of rewards in $\mathcal{R}_{\text{corr}}$, Equation 27
 1852 provides a closed-form expression for this inner minimum. In particular, for any policy π with
 1853 $\mu_\pi \ll \mu_{\pi_{\text{ref}}}$, Equation 27 gives

$$F(\pi) = r \cdot \mathbb{E}_{\mu_\pi}[R_{\text{proxy}}] - \sqrt{1 - r^2} \sqrt{\chi^2(\mu_\pi \parallel \mu_{\pi_{\text{ref}}}) - \mathbb{E}_{\mu_\pi}^2[R_{\text{proxy}}]}.$$

1854 Now assume that the true reward R_{true} lies in $\mathcal{R}_{\text{corr}}$. Since R_{true} is one feasible element of the
 1855 uncertainty set, we must have

$$F(\pi) = \min_{R \in \mathcal{R}_{\text{corr}}} \Delta J(\pi, R) \leq \Delta J(\pi, R_{\text{true}}) = J(\pi, R_{\text{true}}) - J(\pi_{\text{ref}}, R_{\text{true}}).$$

1856 Rearranging yields

$$J(\pi, R_{\text{true}}) - J(\pi_{\text{ref}}, R_{\text{true}}) \geq F(\pi),$$

1857 and substituting the explicit form of $F(\pi)$ from Equation 27 gives the claimed inequality. \square

F ADDITIONAL IMPLEMENTATION DETAILS

F.1 TRAINING DISCRIMINATOR NETWORK

1858 A core step in our Max-Min optimization algorithm and ORPO is to estimate the Radon-Nikodym
 1859 derivative $L(s, a)$, which is critical for computing the χ^2 divergence, as detailed in Appendix F.2.
 1860 To this end, we follow prior works (Laidlaw et al., 2025; Kang et al., 2018; Ho & Ermon, 2016)
 1861 and train a discriminator network. Specifically, we sample a batch of trajectories $D_{\pi_{\text{ref}}}$ from the
 1862 reference policy π_{ref} and another batch D_π from the current policy π . The batch sizes used for each
 1863 are specified in Table 2. And then we use a discriminator architecture identical to that in (Laidlaw
 1864 et al., 2025), denoted by $d_\phi(s, a)$, which is optimized according to:

$$\begin{aligned} \phi &= \arg \min_{\phi} \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [\log(1 + e^{d_\phi(s, a)})] + \mathbb{E}_{\mu_\pi} [\log(1 + e^{-d_\phi(s, a)})] \\ &\approx \arg \min_{\phi} \mathbb{E}_{D_{\pi_{\text{ref}}}} [\log(1 + e^{d_\phi(s, a)})] + \mathbb{E}_{D_\pi} [\log(1 + e^{-d_\phi(s, a)})] \end{aligned} \quad (38)$$

1866 It is known that the optimal discriminator satisfies $d^*(s, a) = \log \frac{\mu_\pi(s, a)}{\mu_{\pi_{\text{ref}}}(s, a)}$ and we estimate $L(s, a)$
 1867 as $\tilde{L}(s, a) = e^{d_\phi(s, a)}$ with $d_\phi(s, a) \approx d^*(s, a)$. However, in the original ORPO implementation¹,
 1868

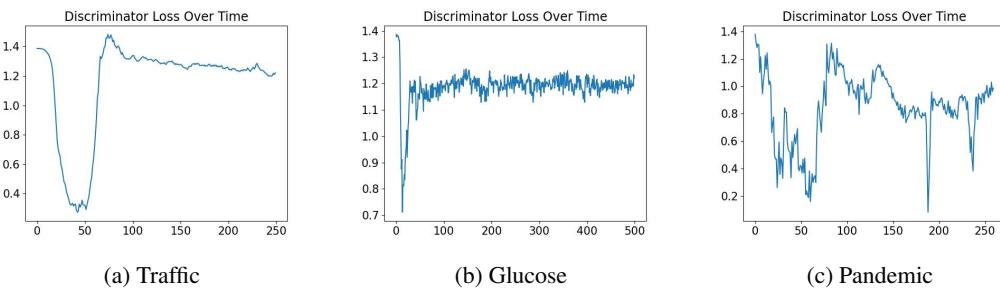
1¹<https://github.com/cassidylaidlaw/orpo/tree/main>

1890 the discriminator is not fully optimized during policy learning. Specifically, the discriminator re-
 1891 ceives only a small number of gradient updates per reinforcement learning iteration, resulting in
 1892 underfitting and inaccurate estimates of the Radon-Nikodym derivative $L(s, a)$.
 1893

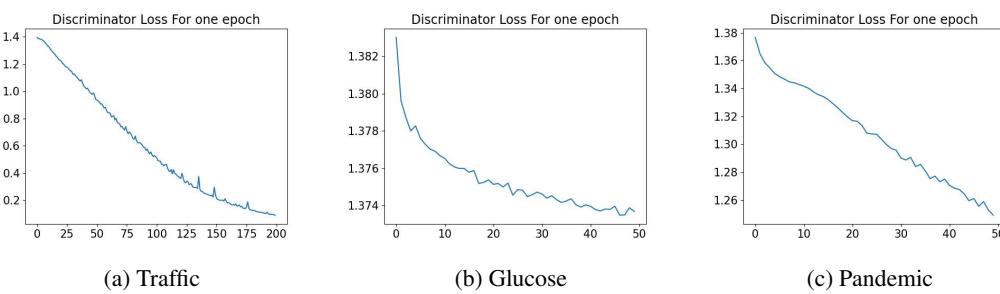
1894 This undertraining is evident in Figure 2, which shows the discriminator loss across RL iterations.
 1895 The loss remains nearly constant (e.g., around 1.4 in the Traffic environment, which is the initial loss
 1896 value as shown in Figure 3a), indicating that the discriminator is not learning effectively. This limits
 1897 its ability to distinguish between π and π_{ref} , especially for state-action pairs where their occupancy
 1898 distributions diverge.
 1899

1900 To address this, we substantially increase the number of gradient updates per iteration and carefully
 1901 tune the learning rate. Our goal is to strike a practical balance between training time and discrimi-
 1902 nator quality: while fully training the discriminator to convergence each iteration is computationally
 1903 expensive, insufficient training leads to inaccurate divergence estimates and unstable optimization.
 1904

1905 Figure 3 shows that in our implementation, the discriminator loss consistently decreases within each
 1906 iteration, e.g., from an initial value around 1.4 to below 0.2 in the Traffic environment, indicating
 1907 effective optimization and more accurate occupancy-ratio estimation. In the Glucose and Pandemic
 1908 environments, however, we observe that training the discriminator for too long leads to slower con-
 1909 vergence and little improvement in loss. In these cases, we apply early stopping to limit training
 1910 time. The specific training schedules are provided in Table 2.
 1911



1920 Figure 2: Discriminator loss across RL iterations in the **original ORPO implementation**. The loss
 1921 stays flat and high (~ 1.4 for the traffic environment), indicating the discriminator is not adequately
 1922 trained. This undermines the accuracy of the estimated occupancy ratios.
 1923



1936 Figure 3: Discriminator loss over training steps *within each RL iteration* in our implementation. The
 1937 loss decreases rapidly from its initial value (e.g., 1.4 to values near 0.2 in the Traffic environment),
 1938 indicating successful training and improved accuracy of occupancy-ratio estimates.
 1939

1940 As for the discriminator network architecture, we follow the same structure described in (Laidlaw
 1941 et al., 2025). For each environment, we employ a fully connected neural network with two hidden
 1942 layers, each consisting of 256 units and ReLU activations. Table 2 summarizes the hyperparameters
 1943 used for discriminator training across different environments.
 1944

1944 Table 2: Hyperparameters used for discriminator network training across different environments.
1945

| Hyperparameter | Traffic | Glucose | Pandemic |
|--------------------------|--------------------|--------------------|--------------------|
| Learning rate | 5×10^{-3} | 1×10^{-2} | 5×10^{-4} |
| SGD epochs per iteration | 200 | 20 | 15 |
| Batch size | 40000 | 100000 | 3860 |
| SGD minibatch size | 16384 | 1024 | 64 |

1953 F.2 DERIVATION OF MAX-MIN POLICY OPTIMIZATION

1954 Using the estimated $d_\phi(s, a)$ from trained discriminator as discussed in Appendix F.1, we can compute the χ^2 divergence via:

1955
$$\chi^2(\mu_\pi \| \mu_{\pi_{\text{ref}}}) = \mathbb{E}_{\mu_\pi} \left[\frac{\mu_\pi(s, a)}{\mu_{\pi_{\text{ref}}}(s, a)} - 1 \right] \approx \mathbb{E}_{D_\pi} \left[e^{d_\phi(s, a)} - 1 \right]. \quad (39)$$

1960 For environments where both state and action spaces are discrete, we directly estimate the occupancy
1961 measure via empirical sampling. Specifically, given the same batch of trajectories D collected from
1962 policy π (as used for training the discriminator), we approximate the discounted occupancy measure
1963 as follows (Schlaginhaufen & Kamgarpour, 2023; Abbeel & Ng, 2004):
1964

1965
$$\tilde{\mu}_\pi^D(s, a) := (1 - \gamma) \frac{1}{N} \sum_{i=1}^N \sum_{t=0}^T \gamma^t \mathbb{I}\{s_t^i = s, a_t^i = a\}, \quad (40)$$

1968 where $\mathbb{I}\{\cdot\}$ is the indicator function. Using this empirical estimate, we can compute the Radon-
1969 Nikodym derivative and χ^2 divergence without training a discriminator.1970 In our formulation, we assume that the proxy reward is normalized with respect to the reference
1971 policy π_{ref} . To achieve this, we reuse the same batch of trajectories $D_{\pi_{\text{ref}}}$ sampled from π_{ref} to
1972 estimate the expected return $\tilde{J}(\pi_{\text{ref}}, R_{\text{proxy}})$ using:
1973

1974
$$\tilde{J}(\pi_{\text{ref}}, R_{\text{proxy}}) = (1 - \gamma) \frac{1}{N} \sum_{i=1}^N R_{\text{proxy}}(\tau^{(i)}) \quad (41)$$

1977 where each $\tau^{(i)} = (s_0, a_0, s_1, a_1, \dots, s_T) \sim D_{\pi_{\text{ref}}}$ is a trajectory sampled from π_{ref} , N is the number
1978 of sampled trajectories, and $R_{\text{proxy}}(\tau^{(i)}) = \sum_{t=0}^T \gamma^t R_{\text{proxy}}(s_t^{(i)}, a_t^{(i)})$. This estimation is unbiased
1979 when using trajectories generated by the policy π . To estimate the empirical variance of the proxy
1980 reward, we use:
1981

1982
$$\tilde{\sigma}_{R_{\text{proxy}}}^2 = \tilde{\mathbb{E}}_{\mu_{\pi_{\text{ref}}}}^2 [R_{\text{proxy}}] - \tilde{\mathbb{E}}_{\mu_{\pi_{\text{ref}}}} [R_{\text{proxy}}]^2 \quad (42)$$

1983 We estimate $\tilde{\mathbb{E}}_{\mu_{\pi_{\text{ref}}}} [R_{\text{proxy}}^2]$ using:

1984
$$\tilde{\mathbb{E}}_{\mu_{\pi_{\text{ref}}}} [R_{\text{proxy}}^2] = (1 - \gamma) \frac{1}{N} \sum_{i=1}^N R_{\text{proxy}}^2(\tau^{(i)})$$

1988 where $R_{\text{proxy}}^2(\tau^{(i)}) = \sum_{t=0}^T \gamma^t R_{\text{proxy}}^2(s_t^{(i)}, a_t^{(i)})$. However, estimating $\tilde{\mathbb{E}}_{\mu_{\pi_{\text{ref}}}}^2 [R_{\text{proxy}}]$ directly from a
1989 single batch introduces bias, because the square of an empirical mean is not an unbiased estimator of
1990 the square of the true mean. To obtain an unbiased estimate of $\tilde{\mathbb{E}}_{\mu_{\pi_{\text{ref}}}}^2 [R_{\text{proxy}}]$, we apply the double-
1991 sampling technique (Di Castro et al., 2012; Xie et al., 2018). Specifically, we independently sample
1992 another batch of trajectories, denoted $D_{\pi_{\text{ref}}}^*$, from the reference policy π_{ref} , and compute:
1993

1994
$$\tilde{\mathbb{E}}_{\mu_{\pi_{\text{ref}}}}^2 [R_{\text{proxy}}] = \tilde{\mathbb{E}}_{\mu_{\pi_{\text{ref}}}}^{D_{\pi_{\text{ref}}}} [R_{\text{proxy}}] \times \tilde{\mathbb{E}}_{\mu_{\pi_{\text{ref}}}}^{D_{\pi_{\text{ref}}}} [R_{\text{proxy}}]$$

1995 where $\tilde{\mathbb{E}}_{\mu_{\pi_{\text{ref}}}}^{D_{\pi_{\text{ref}}}} [R_{\text{proxy}}]$ and $\tilde{\mathbb{E}}_{\mu_{\pi_{\text{ref}}}}^{D_{\pi_{\text{ref}}}} [R_{\text{proxy}}]$ denote the empirical returns computed from the two inde-
1996 pendent batches using Equation 41. This ensures an unbiased estimation of $\tilde{\mathbb{E}}_{\mu_{\pi}}^2 [R_{\text{proxy}}]$, which is
1997 critical for correctly computing the regularization term in the objective 27.

1998
1999We then normalize the proxy reward for each state-action pair in D_π as:2000
2001
2002

$$R_{\text{proxy}}^{\text{norm}}(s, a) = \frac{R_{\text{proxy}}(s, a) - \tilde{J}(\pi_{\text{ref}}, R_{\text{proxy}})}{\tilde{\sigma}_{R_{\text{proxy}}}} \quad (43)$$

2003
2004For notational simplicity, we will continue to use R_{proxy} to denote the normalized proxy reward throughout the remainder of this section.2005
2006We use the same batch of sampled trajectories D_π from current policy π to estimate $\mathbb{E}_{\mu_\pi}[R_{\text{proxy}}]$ using:2007
2008
2009

$$\tilde{\mathbb{E}}_{\mu_\pi}[R_{\text{proxy}}] = \tilde{J}(\pi, R_{\text{proxy}}) = (1 - \gamma) \frac{1}{N} \sum_{i=1}^N R_{\text{proxy}}(\bar{\tau}^{(i)}) \quad (44)$$

2010
2011
2012where each $\bar{\tau}^{(i)} = (s_0, a_0, s_1, a_1, \dots, s_T) \sim D_\pi$ is a trajectory sampled from π . To estimate $\tilde{\mathbb{E}}_{\mu_\pi}^2[R_{\text{proxy}}]$, we apply the same double-sampling technique. Specifically, we independently sample another batches of trajectories, denoted D_π^* , from the current policy π , and compute:2013
2014
2015

$$\tilde{\mathbb{E}}_{\mu_\pi}^2[R_{\text{proxy}}] = \tilde{\mathbb{E}}_{\mu_\pi}^{D_\pi}[R_{\text{proxy}}] \times \tilde{\mathbb{E}}_{\mu_\pi}^{D_\pi^*}[R_{\text{proxy}}], \quad (45)$$

2016
2017where $\tilde{\mathbb{E}}_{\mu_\pi}^{D_\pi}[R_{\text{proxy}}]$ and $\tilde{\mathbb{E}}_{\mu_\pi}^{D_\pi^*}[R_{\text{proxy}}]$ denote the empirical returns computed from the two independent batches using Equation 44.2018
2019

Putting all the steps together, the maxmin algorithm is in Algorithm 2:

2020
2021**Algorithm 2** Max-Min Policy Optimization2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037

```

1: Initialize policy parameters  $\theta$ 
2: Initialize reference policy  $\pi_{\text{ref}}$  and collect trajectories  $D_{\pi_{\text{ref}}}$ 
3: Estimate  $J(\pi_{\text{ref}}, R_{\text{proxy}})$  using Equation 41 and  $\sigma_{R_{\text{proxy}}}^2$  using Equation 42
4: for each iteration do
5:   Sample trajectories  $D_\pi$  from current policy  $\pi_\theta$ 
6:   if discrete environment then
7:     Estimate occupancy measure using Equation 40
8:   else
9:     Train discriminator  $d_\phi$  by minimizing Equation 38
10:  end if
11:  Estimate  $\chi^2$  divergence using Equation 39
12:  Normalize proxy reward for each state-action pair in  $D_\pi$  using Equation 43
13:  Estimate proxy reward expectation  $\mathbb{E}_{\mu_\pi}[R_{\text{proxy}}]$  using Equation 44
14:  Estimate  $\mathbb{E}_{\mu_\pi}^2[R_{\text{proxy}}]$  via double-sampling using Equation 45
15:  Update policy  $\pi_\theta$  using PPO to maximize robust objective in Equation 27
16: end for

```

2038

F.3 DERIVATION OF LINEAR MAX-MIN POLICY OPTIMIZATION2039
2040
2041As for the Linear Max-Min optimization problem, following the discussion in E.4–E.7, we first need to estimate Q :2042
2043

$$Q = \sum_{(s, a)} C(s, a) \phi(s, a) \phi(s, a)^\top$$

2044
2045
2046
2047
2048
2049where $C(s, a) = \mu_{\pi_{\text{ref}}}(s, a)$, $\phi(s, a)$ is a vector of known feature functions sampled under the current policy π . To recover the reference occupancy μ_{ref} , define

$$\bar{d}_\phi(s, a) = \log \frac{\mu_{\pi_{\text{ref}}}(s, a)}{\mu_\pi(s, a)}$$

2050
2051we can rewrite Q via importance sampling (Sutton & Barto, 2018):

$$Q = \mathbb{E}_{\mu_\pi} \left[e^{\bar{d}_\phi(s, a)} \phi(s, a) \phi(s, a)^\top \right]$$

2052 Note that $\bar{d}_\phi(s, a)$ differs slightly from $d_\phi(s, a) = \log \frac{\mu_\pi(s, a)}{\mu_{\pi_{\text{ref}}}(s, a)}$ used in the Max-Min optimization
 2053 algorithm discussed in Appendix F.2, where the numerator and denominator are reversed. To es-
 2054 timate $\bar{d}_\phi(s, a)$, we again train a similar discriminator network as described in Appendix F.1 by
 2055 minimizing:
 2056

$$\begin{aligned} 2057 \phi &= \arg \min_{\phi} \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [\log(1 + e^{-d_\phi(s, a)})] + \mathbb{E}_{\mu_\pi} [\log(1 + e^{d_\phi(s, a)})] \\ 2058 \\ 2059 &\approx \arg \min_{\phi} \mathbb{E}_{D_{\pi_{\text{ref}}}} [\log(1 + e^{-d_\phi(s, a)})] + \mathbb{E}_{D_\pi} [\log(1 + e^{d_\phi(s, a)})] \end{aligned} \quad (46)$$

2061 At optimality, the discriminator satisfies:

$$2063 \bar{d}^*(s, a) = \log \frac{\mu_{\pi_{\text{ref}}}(s, a)}{\mu_\pi(s, a)} \\ 2064$$

2065 And we use $\bar{d}_\phi(s, a) \approx \bar{d}^*(s, a)$. We then estimate Q :
 2066

$$2067 \tilde{Q} = (1 - \gamma) \mathbb{E}_{D_\pi} \left[\sum_{t=0}^{\infty} \gamma^t e^{\bar{d}_\phi(s_t, a_t)} \phi(s_t, a_t) \phi(s_t, a_t)^\top \right] \quad (47)$$

2070 We then perform feature whitening by applying a linear transformation:
 2071

$$2072 \tilde{\phi}(s, a) = \tilde{W} \phi(s, a) \\ 2073$$

2074 where $\tilde{W} = \tilde{Q}^{-1/2}$ is the matrix square root inverse of \tilde{Q} .
 2075

2076 All subsequent quantities are computed using the transformed features $\tilde{\phi}$. As before, we also nor-
 2077 malize the proxy reward for each state-action pair in D_π using Equation 43.

2078 After whitening, we need to estimate the gradient of each dual variables $(\lambda_1, \lambda_2, \lambda_3)$ as derived in
 2079 Appendix E.7.
 2080

2081 **Estimating $C_{\phi, j}$ and $D_{\phi, j}$.** Recall that
 2082

$$2083 C_{\phi, j} = \sum_{(s, a)} C(s, a) \phi_j(s, a) \\ 2084$$

2085 which can be rewritten via importance sampling as:
 2086

$$2087 C_{\phi, j} = \mathbb{E}_{\mu_\pi} \left[e^{\bar{d}_\phi(s, a)} \tilde{\phi}_j(s, a) \right] \\ 2088$$

2089 and then can be approximated via:
 2090

$$2091 \tilde{C}_{\phi, j} = (1 - \gamma) \mathbb{E}_{D_\pi} \left[\sum_{t=0}^{\infty} \gamma^t e^{\bar{d}_\phi(s_t, a_t)} \tilde{\phi}_j(s_t, a_t) \right] \\ 2092$$

2093 Similarly, recall that
 2094

$$2095 D_{\phi, j} = \sum_{(s, a)} D(s, a) \phi_j(s, a) \\ 2096$$

2097 where $D(s, a) = \mu_{\pi_{\text{ref}}}(s, a) \cdot R_{\text{proxy}}(s, a)$. Using importance sampling, we can write:
 2098

$$2100 D_{\phi, j} = \mathbb{E}_{\mu_\pi} \left[e^{\bar{d}_\phi(s, a)} R_{\text{proxy}}(s, a) \tilde{\phi}_j(s, a) \right] \\ 2101$$

2102 and can be approximated in practice by:
 2103

$$2104 \tilde{D}_{\phi, j} = (1 - \gamma) \mathbb{E}_{D_\pi} \left[\sum_{t=0}^{\infty} \gamma^t e^{\bar{d}_\phi(s_t, a_t)} R_{\text{proxy}}(s_t, a_t) \tilde{\phi}_j(s_t, a_t) \right] \\ 2105$$

2106 **Estimating $q_j(\lambda_1, \lambda_2)$.** Recall that
 2107

$$\begin{aligned} q_j(\lambda_1, \lambda_2) &= \sum_{(s,a)} (v(s,a) - \lambda_1 D(s,a) - \lambda_2 C(s,a)) \phi_j(s,a) \\ &= \sum_{(s,a)} v(s,a) \phi_j(s,a) - \lambda_1 \sum_{(s,a)} D(s,a) \phi_j(s,a) - \lambda_2 \sum_{(s,a)} C(s,a) \phi_j(s,a) \\ &= \mathbb{E}_{\mu_\pi}[\phi_j(s,a)] - \lambda_1 D_{\phi,j} - \lambda_2 C_{\phi,j} \end{aligned}$$

2114 where $v(s,a) = \mu_\pi(s,a)$ and $\mathbb{E}_{\mu_\pi}[\phi_j(s,a)]$ is the discounted feature expectation under the policy
 2115 π . We can estimate the first term using:
 2116

$$\tilde{\mathbb{E}}_{\mu_\pi}[\phi_j(s,a)] = (1 - \gamma) \frac{1}{N} \sum_{i=1}^N \tilde{\phi}_j(\bar{\tau}^i)$$

2120 where each $\bar{\tau}^{(i)} = (s_0, a_0, s_1, a_1, \dots, s_T) \sim D_\pi$ is a trajectory sampled from π , and $\tilde{\phi}_j(\bar{\tau}^{(i)}) =$
 2121 $\sum_{t=0}^T \gamma^t \tilde{\phi}_j(s_t^{(i)}, a_t^{(i)})$. Given the above estimates, we can finally compute:
 2122

$$\tilde{q}_j(\lambda_1, \lambda_2) = \tilde{\mathbb{E}}_{\mu_\pi}[\phi_j(s,a)] - \lambda_1 \tilde{D}_{\phi,j} - \lambda_2 \tilde{C}_{\phi,j}$$

2125 With the above estimation, we can compute the gradient and solve for the optimal dual variables
 2126 $\lambda = (\lambda_1, \lambda_2, \lambda_3)$ using the Levenberg-Marquardt algorithm (Moré, 1978), a damped least-squares
 2127 method designed for solving nonlinear systems of equations. Specifically, we use the root solver in
 2128 SciPy (Virtanen et al., 2020) to find the stationary point of the gradient $\nabla_\lambda g(\lambda) = 0$. We initialize
 2129 the optimization with $\lambda_1 = 0$, $\lambda_2 = 0$, and $\lambda_3 = -1$, and enforce $\lambda_3 < 0$ throughout training
 2130 to ensure concavity of the dual objective $g(\lambda)$. To enforce the non-negativity constraint $\theta \geq 0$
 2131 as required by the analytical form in Equation 33, we manually clip each θ_i to ensure it remains
 2132 non-negative. Future work may explore alternative solvers better suited to constrained optimization.
 2133

Recall that the optimal primal variable θ_j^* is:

$$\theta_j^*(\lambda) = \max \left(0, \frac{q_j(\lambda_1, \lambda_2)}{2\lambda_3} \right)$$

2137 After optimizing for the dual variables, we can substitute the optimal $(\lambda_1^*, \lambda_2^*, \lambda_3^*)$ back into the
 2138 above equation and get:
 2139

$$\tilde{\theta}_j^*(\lambda) = \max \left(0, \frac{\tilde{q}_j(\lambda_1^*, \lambda_2^*)}{2\lambda_3^*} \right)$$

2142 for all features. Then we can substitute the optimal $\tilde{\theta}^*$ back in the robust reward objective in Equation
 2143 28 and train the policy π to maximize the outer problem using the standard reinforcement
 2144 learning algorithm proximal policy optimization (PPO) (Schulman et al., 2017).
 2145

Putting all the steps together, the linear maxmin algorithm is in Algorithm 3:
 2146

2147 F.4 ENVIRONMENT DESCRIPTION AND REWARD HACKING TYPES

2149 **Traffic.** This environment simulates a highway merging scenario, adapted from (Pan et al., 2022;
 2150 Wu et al., 2021; Vinitsky et al., 2018), where a group of autonomous vehicles (AVs) controlled
 2151 by an RL agent must merge into human-driven traffic. Each AV observes its own state (position
 2152 and velocity) and those of nearby vehicles, and outputs continuous acceleration actions. The true
 2153 reward is designed to ensure smooth and efficient traffic flow, encouraging low commute times
 2154 and gentle accelerations. The reference policy π_{ref} is a behavioral cloning (BC) policy trained on
 2155 demonstrations generated by the Intelligent Driver Model (IDM) (Treiber et al., 2000).
 2156

2157 **Pandemic.** Based on the PandemicSimulator (Kompella et al., 2020), this environment models
 2158 infection dynamics using an extended SEIR model. At each timestep, the agent selects a lockdown
 2159 policy to control the spread of disease while minimizing societal costs. The true reward balances
 2160 infection severity, political disruption, and policy smoothness over time. The reference policy is
 2161 trained via behavioral cloning on a mixture of realistic and hand-crafted policy trajectories.
 2162

2160 **Algorithm 3** Linear Max-Min Policy Optimization

```

2161 1: Initialize policy parameters  $\theta$ 
2162 2: Initialize reference policy  $\pi_{\text{ref}}$  and collect trajectories  $D_{\pi_{\text{ref}}}$ 
2163 3: Estimate  $J(\pi_{\text{ref}}, R_{\text{proxy}})$  using Equation 41 and  $\sigma_{R_{\text{proxy}}}^2$  using Equation 42
2164 4: for each iteration do
2165 5:   Sample trajectories  $D_{\pi}$  from current policy  $\pi_{\theta}$ 
2166 6:   if discrete environment then
2167 7:     Estimate occupancy measure using Equation 40
2168 8:   else
2169 9:     Train discriminator  $d_{\phi}$  by minimizing Equation 46
2170 10:  end if
2171 11:  Normalize proxy reward for each state-action pair in  $D_{\pi}$  using Equation 43
2172 12:  Estimate  $\tilde{Q}$  using Equation 47
2173 13:  Compute feature transformation  $\tilde{W} = \tilde{Q}^{-1/2}$  and transform features  $\tilde{\phi}(s, a) = \tilde{W}\phi(s, a)$ 
2174 14:  Estimate  $\tilde{C}_{\phi, j}$ ,  $\tilde{D}_{\phi, j}$  using transformed features
2175 15:  Compute  $\tilde{q}_j(\lambda_1, \lambda_2)$  for all features
2176 16:  Solve for optimal dual variables  $(\lambda_1, \lambda_2, \lambda_3)$  by maximizing the dual objective (Equation 34)
2177 17:  Compute the optimal primal variable  $\tilde{\theta}_j^*$  for all features
2178 18:  Update policy  $\pi_{\theta}$  using PPO to maximize the robust objective in Equation 28
2179 19: end for
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213

```

Glucose Monitoring. This environment uses the SimGlucose simulator (Man et al., 2014; Fox et al., 2020), where an RL agent administers insulin doses to a simulated patient with Type 1 Diabetes. The goal is to maintain safe blood glucose levels and minimize long-term health risk. The reference policy is trained via behavioral cloning using data generated by a PID controller with clinically tuned parameters (Steil, 2013). Proxy rewards in this setting often reflect surrogate objectives such as treatment cost or patient burden.

Tomato Watering GridWorld. This environment presents a simple spatial grid where the agent waters tomato plants. The true reward corresponds to the number of tomatoes correctly watered. However, the proxy reward includes an artificially high bonus at a specific grid location (a “sprinkler state”), which causes the agent to overfit by remaining in that region despite little actual benefit to overall tomato growth. The reference policy follows (Laidlaw et al., 2025), with 10% random actions added to allow for policy improvement.

RLHF. This environment builds on prior work (Laidlaw et al., 2025; Coste et al.) that studies overoptimization of LLM-based reward models. The proxy reward function is derived from a reward model fine-tuned on the AlpacaFarm preference dataset (Dubois et al., 2023), using the Pythia-70M model (Biderman et al., 2023), a comparatively small model. For the true reward signal, we adopt the Llama 3 Tulu V2 8B reward model released by AI2 (Ivison et al., 2024). The reference policy corresponds to the supervised fine-tuned (SFT) model from (Laidlaw et al., 2025; Coste et al.), which was trained on the AlpacaFarm SFT dataset using Pythia-1.4B.

Types of Reward Hacking. We adopt the taxonomy proposed in (Pan et al., 2022) to classify the kinds of proxy reward misalignments that lead to reward hacking. Our selected environments span all three major categories:

- **Misweighting:** The proxy reward includes all relevant objectives but uses incorrect relative weights. Our Linear Max-Min method specifically seeks the most adversarial weighting in this space.
- **Ontological:** The proxy captures the correct high-level goal using different or incomplete features. In the **Traffic** environment, the true reward combines commute time, acceleration, and headway, whereas the proxy replaces commute time with velocity. In the **Pandemic** environment, the true reward penalizes infections, political cost, lower stage changes, and non-smooth policies, while the proxy omits the political cost entirely. Similarly, in **Glucose**, the proxy reward includes a high bonus for staying in a specific grid cell (the “sprinkler state”), which causes the agent to overfit to that state despite little actual benefit to overall tomato growth.

2214
 2215 **close**, the proxy reward only considers the expected patient costs while the true reward only
 2216 measures the health risk.
 2217
 2218 • **Scope:** The proxy evaluates behavior over a limited domain. In the **Tomato** environment,
 2219 the true reward reflects the number of tomatoes successfully watered. However, the proxy
 2220 introduces a large bonus at a specific state (the sprinkler), incentivizing the agent to pursue
 2221 this location at the expense of fulfilling the intended watering task. In the **RLHF** environ-
 2222 ment, the proxy reward is produced by a comparatively small model with limited evaluative
 2223 capacity, whereas the true reward is derived from a much larger, stronger model. Conse-
 2224 quently, the proxy reward provides a less reliable evaluation signal.

2224 F.5 ADDITIONAL EXPERIMENT SETUP

2226
 2227 **Non-LLM Experiments.** For the policy networks, we follow the architectures described in (Laid-
 2228 law et al., 2025). In the Pandemic, Traffic, and Tomato environments, we use fully connected neural
 2229 networks with 2 layers of 128 units, 4 layers of 512 units, and 4 layers of 512 units, respectively.
 2230 For the Glucose environment, we employ a three-layer LSTM network, where each LSTM layer has
 2231 64 units. We use the pre-trained policies provided in the ORPO repository² as the reference policies
 2232 π_{ref} . We initialize the policy network with the corresponding pre-trained checkpoint for the Traffic,
 2233 Glucose, and Pandemic environments, and initialize a random policy for the Tomato environment.
 2234 Table 3 summarizes the hyperparameters used for PPO training across all models and environments.

2235 Table 3: Hyperparameters used for PPO training across different environments.

| 2237 Hyperparameter | 2238 Traffic | 2239 Glucose | 2240 Pandemic | 2241 Tomato |
|---|---------------------|---------------------|----------------------|--------------------|
| 2239 Training iterations | 250 | 500 | 260 | 500 |
| 2240 Batch size | 40000 | 100000 | 3860 | 3000 |
| 2241 Optimizer | Adam | Adam | Adam | Adam |
| 2242 Learning rate | 5×10^{-5} | 1×10^{-5} | 0.0003 | 1×10^{-3} |
| 2243 Gradient clipping | N/A | 10 | 10 | 0.1 |
| 2244 Discount factor (γ) | 0.99 | 0.99 | 0.99 | 0.99 |
| 2245 Random seed | 0 | 0 | 0 | 0 |
| 2246 GAE coefficient (λ) | 0.97 | 0.98 | 0.95 | 0.98 |
| 2247 Entropy coefficient (start) | 0.01 | 0.01 | 0.1 | 0.01 |
| 2248 Entropy coefficient (end) | 0.01 | 0.01 | 0.01 | 0.01 |
| 2249 KL target | 0.02 | 1×10^{-3} | 0.01 | 1×10^{-3} |
| 2250 Value function loss clipping | 10000 | 100 | 20 | 10 |
| 2251 Value function loss coefficient | 0.5 | 0.0001 | 0.5 | 0.1 |
| 2252 Share value function layers | True | True | True | False |

2253
 2254 As for the reward used during training and evaluation, we follow the same setup as ORPO (Laidlaw
 2255 et al., 2025). All policies are trained using **only the proxy reward**, while both the true and proxy
 2256 rewards are used for evaluation.

2257 In the **Traffic** environment, the proxy reward is a weighted combination of *velocity*, *acceleration*,
 2258 and *headway*, with weights 1, 1, and 0.1, respectively. The true reward, on the other hand, uses
 2259 *commute time*, *acceleration*, and *headway*, also weighted 1, 1, and 0.1.

2260 In the **Pandemic** environment, the proxy reward is composed of *infection summary absolute*, *lower*
 2261 *stage*, and *smooth stage changes*, with weights 10, 0.1, and 0.01. The true reward adds a *political*
 2262 component to these three features and is weighted with 10.

2263 For the **Glucose** environment, the proxy reward includes only one feature: *expected patient cost*.
 2264 The true reward is based on *magnitude*, which measures the health risk of the patient.

2265
 2266 ²https://github.com/cassidylaidlaw/orpo/tree/main/data/base_policy_checkpoints

In the **Tomato** environment, the true reward counts the number of *watered tomato*. The proxy reward adds a large bonus at a specific state (*sprinkler*), incentivizing the agent to reach that location regardless of its impact on the primary task.

LLM Experiments. We adopt the common formulation for RLHF as a *contextual bandit* problem, where the environment is modeled as a Markov Decision Process (MDP) with a discount factor $\gamma = 0$. In this setup, the return of a policy π under a given reward function R simplifies to:

$$J(\pi, R) = \mathbb{E}_{s \sim \mu_0, a \sim \pi} [R(s, a)],$$

where μ_0 denotes the distribution over initial states. In the context of RLHF, each state corresponds to a prompt sampled from a dataset, and the action is the model’s generated response. The reward is then computed based on this prompt–response pair.

Under this contextual bandit assumption, the χ^2 divergence between occupancy measures reduces to the divergence between action distributions conditioned on prompts:

$$\chi^2(\mu_\pi \| \mu_{\pi_{\text{ref}}}) = \mathbb{E}_{s \sim \mu_0} [\chi^2(\pi(\cdot|s) \| \pi_{\text{ref}}(\cdot|s))],$$

as established in Lemma A.6 of (Laidlaw et al., 2025). This allows us to avoid discriminator-based estimation of occupancy ratios in this setting. Instead, we directly estimate the χ^2 divergence using the following estimator:

$$\tilde{\chi}^2(\mu_\pi \| \mu_{\pi_{\text{ref}}}) = \mathbb{E}_{s \sim \mu_0, a \sim \pi} \left[\frac{\pi(a|s)}{\pi_{\text{ref}}(a|s)} + \frac{\pi_{\text{ref}}(a|s)}{\pi(a|s)} - 2 \right],$$

as proposed in (Laidlaw et al., 2025).

For policy optimization, we apply the same Max-Min training algorithm described in Appendix F.2, adapting it to the contextual bandit structure without discriminator training. We evaluate our approach in the RLHF setting using a setup consistent with prior work (Laidlaw et al., 2025; Coste et al.). The proxy reward function is derived from a reward model fine-tuned on the AlpacaFarm preference dataset (Dubois et al., 2023), using the Pythia-70M model (Biderman et al., 2023). For the true reward signal, we adopt the Llama 3 Tulu V2 8B reward model released by AI2 (Ivison et al., 2024). The reference policy corresponds to the supervised fine-tuned (SFT) model from (Laidlaw et al., 2025; Coste et al.), which was trained on the AlpacaFarm SFT dataset using Pythia-1.4B. All policy evaluations, both proxy and true rewards, are conducted on the same set of prompts used in (Laidlaw et al., 2025).

To further strengthen our experimental evaluation regarding the RLHF setting, we also compare against the Reward Ensemble method (referred to as Ensemble for brevity) (Eisenstein et al., 2023). Specifically, we adopt their *finetune ensembles* setting: we fine-tune five reward models on the AlpacaFarm preference dataset, all initialized from the same pre-trained Pythia-70M model but using five different random seeds, and aggregate their outputs using the mean rule. This setup is directly comparable to our RLHF configuration for ORPO and our methods, where both ORPO and our approach use a single fine-tuned Pythia-70M reward model.

Selection of r . For our Max-Min and Linear Max-Min policy optimization algorithms, the correlation parameter r serves as an additional hyperparameter. In practice, as with ORPO (Laidlaw et al., 2025), r may only be approximately estimated, and there is currently no principled method for selecting its optimal value. To address this, we perform a grid search over $r \in \{0.1, 0.2, \dots, 0.9\}$ for each environment and measure the resulting Max-Min and Linear Max-Min policy expected returns under the worst-case or linear worst-case reward. Results on all searched r can be found in Appendix H.5. Additional analysis of how different training values of r affect robustness under varying evaluation r values is provided in Appendix H.2. Unless otherwise noted, we use the following r values for training and evaluation throughout our experiments: $r = 0.3$ for **Traffic**, $r = 0.7$ for **Pandemic**, $r = 0.9$ for **Glucose**, $r = 0.4$ for **Tomato**, and $r = 0.4$ for RLHF. As for ORPO policy, we trained with occupancy-measure χ^2 regularization, using the official implementation from (Laidlaw et al., 2025). All hyperparameters are set as recommended to ensure optimal performance. The ORPO* shares the exact same setting as the ORPO policy with the full discriminator training schedule as in our algorithms.

Evaluation of the worst-case performance. Theoretically, in the absence of structural constraints on the reward function, as opposed to the case of linear rewards, the worst-case reward of a policy

in state-action pairs unvisited by π_{ref} can be arbitrarily negative without violating the correlation constraint. However, assigning extremely negative values is impractical in real-world scenarios due to domain constraints. Moreover, doing so would render all policies with at least one unseen state-action pair equally poor in terms of worst-case reward, obscuring meaningful comparisons. To address this, we define a minimal feasible reward value R_{\min} and assign it to all unseen state-action pairs. The actual expected worst-case reward (**Worst***) is thus calculated as:

$$\sum_{(s,a):\mu_{\pi_{\text{ref}}}(s,a)>0} \mu_{\pi}(s,a)R_{\text{worst}} + \sum_{(s',a'):\mu_{\pi_{\text{ref}}}(s',a')=0} \mu_{\pi}(s',a')R_{\min}$$

where the first part is derived from the adversarial reward function given by our inner minimization solution, and the second part applies to state-action pairs unvisited by π_{ref} .

In practice, however, environments like Traffic, Pandemic, and Glucose are continuous with large state-action spaces, making it difficult to reliably estimate $\mu_{\pi}(s,a)$ and $\mu_{\pi_{\text{ref}}}(s,a)$ from a limited number of trajectories. As a result, identifying unvisited or low-density regions in these environments is far more ambiguous. Therefore, for these continuous environments, we rely on the output of the discriminator as a signal for detecting unseen state-action pairs. Specifically, if the discriminator outputs infinity (or diverges numerically) for a given state-action, we treat this as an indication that the state-action was never visited by the reference policy π_{ref} . We approximate the total occupancy (**Occ**) over such state-action pairs by computing their frequency in the sampled trajectories, and use the expected worst-case reward (**Worst**) of a policy π over the remaining state-action pairs as the default worst-case performance metric: $\sum_{(s,a):d_{\phi}(s,a)<\infty} \mu_{\pi}(s,a)R_{\text{worst}}(s,a)$. In contrast, for the discrete Tomato environment, we directly estimate the occupancy measure by sampling state-action pairs and then compute **Worst*** accordingly. Further details on this procedure are provided in Appendix H.2.

To compare the worst-case performance of different policies, we sample 200 trajectories in the **Traffic** and **Glucose** environments, 20 trajectories in **Pandemic**, 1000 trajectories in **Tomato**, and 8 answers per prompt in **RLHF** to estimate the worst-case performance of a policy.

Evaluation of policy robustness. To evaluate robustness across different correlation levels, we uniformly sample candidate vectors θ , where each component θ_i is drawn from the interval $[0, 1]$. We use the same number of trajectories sampled from the reference policy π_{ref} to determine whether it satisfies the correlation constraint:

$$\mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\frac{\theta^{\top} \phi - M}{V} \cdot R_{\text{proxy}} \right] = r,$$

where M and V denote the mean and standard deviation of $\theta^{\top} \phi$ under the reference policy.

Note: For our worst-case performance evaluation, we explicitly normalize the reward to have zero mean and unit variance under the reference policy (enforcing $M = 0$ and $V = 1$). In contrast, for the robustness evaluation across correlation levels, we do not apply such normalization. This allows us to report the average reward under each θ in its original scale, reflecting variability more comparable to the original true reward landscape.

F.6 TRAINING TIME AND COMPLEXITY

Table 4 reports the total training time for each algorithm across different environments. All experiments were conducted on a single NVIDIA RTX 4090 GPU (24GB memory) and a 13th Gen Intel Core i9-13900KF CPU (32 threads). We implemented all methods in Python 3.9 using PyTorch 2.6.0 (Paszke, 2019), RLlib (Liang et al., 2018) and trlX (Havrilla et al., 2023).

The training times and memory usages across different environments are summarized in Table 4 and Table 5. Since the training durations for ORPO*, Max-Min, and Linear Max-Min differ by less than one hour in each setting, and the memory usages differ by less than 10MB, we group them together for brevity. Since the RLHF environment does not require training a discriminator, the training times and memory usages for ORPO and our Max-Min are identical. We therefore exclude RLHF from the runtime analysis. As shown in Table 4, all three methods require more training time compared to the original ORPO implementation. The increased training time primarily results from additional gradient steps used to more thoroughly train the discriminator network. Specifically, the

2376 per-iteration training time is approximately 2.5 minutes for Traffic, 4.6 minutes for Pandemic, and
 2377 8.9 minutes for Glucose. This leads to a total training time increase from roughly 7 hours to 37 hours
 2378 in Glucose. However, the added cost is environment-dependent and remains moderate in simpler
 2379 settings, for example, increasing from 5 to 10 hours in Traffic. **In contrast, the memory footprint**
 2380 **of our methods is very close to that of ORPO: the peak CPU memory usage differs by less than**
 2381 **30–50MB across environments (within a few percent of ORPO in all cases).** Overall, these results
 2382 indicate that our methods introduce a modest runtime overhead and negligible memory overhead,
 2383 achieving a practical trade-off between computational cost and the improved quality of divergence
 2384 estimation.

2385 Table 4: Approximate training time for each algorithm across different environments.
 2386

| Algorithm | Traffic | Glucose | Pandemic | Tomato |
|----------------------------------|---------|---------|----------|--------|
| ORPO | ≈5h | ≈7h | ≈14h | ≈1h |
| ORPO* / Max-Min / Linear Max-Min | ≈10h | ≈37h | ≈19h | ≈1h |

2387 Table 5: Approximate memory usage for each algorithm across different environments.
 2388

| Algorithm | Traffic | Glucose | Pandemic | Tomato |
|----------------------------------|---------|---------|----------|---------|
| ORPO | ≈1679MB | ≈1662MB | ≈1813MB | ≈2148MB |
| ORPO* / Max-Min / Linear Max-Min | ≈1706MB | ≈1674MB | ≈1864MB | ≈1903MB |

2392 **Complexity.** At first glance, regularization-based approaches like ORPO may appear more
 2393 computationally efficient than max-min optimization, which often involves iterative procedures to solve
 2394 both inner and outer objectives. However, in practice, ORPO requires repeatedly estimating the χ^2
 2395 divergence between policy distributions during each policy update step. This estimation is done
 2396 by training a discriminator network, which itself involves multiple optimization steps per iteration.
 2397 In contrast, our Max-Min formulation admits a closed-form solution for the inner minimization
 2398 over reward functions. This allows us to avoid iterative solving in the inner loop entirely. For the
 2399 Linear Max-Min variant, although a closed-form expression for the dual variables is not available,
 2400 the corresponding dual optimization problem is smooth and well-posed, and can be solved
 2401 efficiently using standard gradient-based methods. Therefore, despite the max-min structure, our
 2402 method does not incur higher practical complexity compared to ORPO. In fact, both approaches rely
 2403 on discriminator-based divergence estimation and perform comparable amounts of computation per
 2404 iteration. The main difference lies in the structure of the objective, not in the asymptotic or em-
 2405 pirical complexity. In summary, ORPO does not inherently enjoy a complexity advantage over our
 2406 Max-Min or Linear Max-Min algorithms.
 2407

2408

G CONVERGENCE ANALYSIS

2409 In this section, we study the convergence of our Max-Min and Linear Max-Min algorithms. As both
 2410 methods rely on accurately estimating the occupancy measure, we begin by analyzing the sample
 2411 complexity of this estimation via the discriminator described in Appendix F.1.
 2412

2413

G.1 SAMPLE COMPLEXITY OF OCCUPANCY MEASURE ESTIMATION

2414 In this section, we analyze the sample complexity of estimating the occupancy measure via the
 2415 discriminator described in Appendix F.1. Our argument adapts techniques from Huang et al. (2023);
 2416 Barakat et al. (2024) to our setting. To start with, we define the following notations for convenience.
 2417 Let $x = (s, a)$ range over the space \mathcal{X} of all possible state-action pairs. We consider two reference
 2418 distributions on \mathcal{X} : $\mu_{\pi_{\text{ref}}}$ and μ_{π} . The (true) density ratio of μ_{π} with respect to $\mu_{\pi_{\text{ref}}}$ is
 2419

$$2420 L^*(x) := \frac{\mu_{\pi}(x)}{\mu_{\pi_{\text{ref}}}(x)} \quad \text{on the support of } \mu_{\pi_{\text{ref}}},$$

2430 and its log-ratio is $d^*(x) := \log L^*(x)$.
 2431

2432 We work with a parametric log-ratio class $\mathcal{D} = \{d_\phi : \mathcal{X} \rightarrow \mathbb{R}\}$ and the induced ratio class $\mathcal{L} =$
 2433 $\{L_\phi := e^{d_\phi}\}$. Following Equation 38 in Appendix F.1, we learn d_ϕ by minimizing the following
 2434 loss:

$$2435 \mathcal{R}(d) := \frac{1}{2} \mathbb{E}_{x \sim \mu_{\pi_{\text{ref}}}} [\log(1 + e^{d(x)})] + \frac{1}{2} \mathbb{E}_{x \sim \mu_\pi} [\log(1 + e^{-d(x)})]. \quad (48)$$

2436 Given n_{ref} i.i.d. samples $\{x_i^{\text{ref}}\}_{i=1}^{n_{\text{ref}}} \sim \mu_{\pi_{\text{ref}}}$ and n_π i.i.d. samples $\{x_j^\pi\}_{j=1}^{n_\pi} \sim \mu_\pi$, which are inde-
 2437 pendent, we can minimize the empirical loss as follows in practice:

$$2438 \widehat{\mathcal{R}}(d) := \frac{1}{2} \cdot \frac{1}{n_{\text{ref}}} \sum_{i=1}^{n_{\text{ref}}} \log(1 + e^{d(x_i^{\text{ref}})}) + \frac{1}{2} \cdot \frac{1}{n_\pi} \sum_{j=1}^{n_\pi} \log(1 + e^{-d(x_j^\pi)}). \quad (49)$$

2439 Let the true loss minimizer be
 2440

$$2441 d^* \in \arg \min_{d \in \mathcal{D}} \mathcal{R}(d), \quad L^* := e^{d^*}.$$

2442 Let the empirical loss minimizer be
 2443

$$2444 \widehat{d} \in \arg \min_{d \in \mathcal{D}} \widehat{\mathcal{R}}(d), \quad \widehat{L} := e^{\widehat{d}}.$$

2445 For convenience, we also define the mixture distribution $\mu_{\text{mix}} := \frac{1}{2} \mu_{\pi_{\text{ref}}} + \frac{1}{2} \mu_\pi$, which will be used
 2446 later. We make the following assumption throughout the analysis.

2447 **Assumption 1** (Modeling, boundedness and cover). *The following conditions hold throughout the
 2448 analysis:*

- 2449 1. **Realizability.** *The true log-ratio belongs to the model class: $d^* \in \mathcal{D}$ (equivalently, $L^* \in$
 2450 \mathcal{L}).*
- 2451 2. **Bounded.** *There exist constants $0 < \alpha \leq \beta < \infty$ such that*

$$2452 \alpha \leq L_\phi(x) \leq \beta \quad \text{for all } x \in \mathcal{X}, \phi,$$

2453 and hence $\alpha \leq L^*(x) \leq \beta$ as well. Equivalently, $d_\phi(x) \in [\log \alpha, \log \beta]$.

- 2454 3. **L_1 optimistic cover (Definition 3 of (Huang et al., 2023)).** *There exists a finite set $\overline{\mathcal{L}} \subset$
 2455 $(0, \infty)^\mathcal{X}$ with cardinality $|\overline{\mathcal{L}}| = M$ and a scale $\gamma > 0$ such that for every $L \in \mathcal{L}$ there is
 2456 $\overline{L} \in \overline{\mathcal{L}}$ with*

$$2457 \overline{L}(x) \geq L(x) \quad \text{for all } x, \quad \mathbb{E}_{x \sim \mu_{\text{ref}}} [|\overline{L}(x) - L(x)|] \leq \gamma, \quad \alpha \leq \overline{L}(x) \leq \beta.$$

2458 We denote $\overline{\mathcal{D}} := \{\overline{d} := \log \overline{L} : \overline{L} \in \overline{\mathcal{L}}\}$.

2459 Assumption 1 collects the conditions used throughout our analysis. First, *realizability* is standard
 2460 in likelihood-based occupancy estimation and allows us to control the estimation error through the
 2461 complexity of the parametric class rather than the size of \mathcal{X} . In practice, a sufficiently expres-
 2462 sive neural discriminator makes this assumption reasonable. Second, *boundedness* guarantees well-
 2463 posedness on the support of $\mu_{\pi_{\text{ref}}}$ and prevents divisions by zero. It can be enforced by restricting
 2464 attention to the support of $\mu_{\pi_{\text{ref}}}$ or by applying ratio clipping during training. Finally, the *L_1 opti-
 2465 mistic cover* (adopted from Definition 3 of Huang et al. (2023)) is the technical device that enables
 2466 uniform concentration and converts control of the loss in Equation 48 into an L_1 error with clean
 2467 constants. We instantiate this cover for our discriminator class later in the proof.

2468 We begin by stating some auxiliary lemmas that formalize the structural claims used later.

2469 **Lemma 1** (Strong convexity of $\mathcal{R}(d)$). *Let Assumption 1 hold true. Define*

$$2470 \lambda := \frac{\min\{\alpha, \beta\}}{(1 + \max\{\alpha, \beta\})^2} > 0.$$

2471 Then for any measurable $d \in \mathcal{D}$ and for the unique minimizer d^* to Equation 48, we have

$$2472 \mathcal{R}(d) - \mathcal{R}(d^*) \geq \frac{\lambda}{2} \mathbb{E}_{x \sim \mu_{\text{mix}}} [(d(x) - d^*(x))^2].$$

2484 *Proof.* Fix x and define the pointwise loss as:

$$2486 \quad r_x(d) := (1 - \eta(x)) \log(1 + e^d) + \eta(x) \log(1 + e^{-d}), \quad \eta(x) := \frac{\mu_\pi(x)}{\mu_\pi(x) + \mu_{\pi_{\text{ref}}}(x)}.$$

2488 Its derivatives w.r.t. the scalar d are $r'_x(d) = \sigma(d) - \eta(x)$ and $r''_x(d) = \sigma(d)(1 - \sigma(d)) > 0$, where
2489 $\sigma(d) = \frac{e^d}{1+e^d}$. We notice that $r''_x(d)$ is independent of $\eta(x)$. Therefore, at every x , the pointwise loss
2490 r_x is strictly convex in d .

2492 Now we estimate the lower bound for $r''_x(d)$. On the range $d \in [\log \alpha, \log \beta]$ (boundedness from
2493 Assumption 1), let $y = e^d \in [\alpha, \beta]$; then

$$2495 \quad r''_x(d) = \sigma(d)(1 - \sigma(d)) = \frac{y}{(1+y)^2} =: f(y)$$

2497 Let's consider the monotonicity of $f(y)$, $f'(y) = \frac{1-y}{(1-y)^3}$, so f increases on $(0, 1]$ and decreases on
2498 $[1, \infty)$. Therefore,

$$2499 \quad \min_{y \in [\alpha, \beta]} f(y) = \min \left\{ \frac{\alpha}{(1+\alpha)^2}, \frac{\beta}{(1+\beta)^2} \right\}.$$

2501 We can consider a slightly more conservative but simpler bound:

$$2503 \quad \lambda := \frac{\min\{\alpha, \beta\}}{(1 + \max\{\alpha, \beta\})^2} \leq \min_{t \in [\alpha, \beta]} f(y),$$

2505 Thus, for all x and all $d \in [\log \alpha, \log \beta]$, we have $r''_x(d) \geq \lambda$. Notice that strong convexity (with
2506 parameter λ) of a C^2 univariate function g satisfies:

$$2508 \quad g(u) \geq g(v) + g'(v)(u - v) + \frac{\lambda}{2}(u - v)^2 \quad \text{for all } u, v.$$

2510 Applying this with $g(\cdot) = r_x(\cdot)$, $u = d(x)$, and $v = d^*(x)$, where d^* is the unique pointwise
2511 minimizer (so $r'_x(d^*(x)) = 0$). We get for every x ,

$$2513 \quad r_x(d(x)) - r_x(d^*(x)) \geq \frac{\lambda}{2}(d(x) - d^*(x))^2.$$

2515 Recall that $\mu_{\text{mix}} = \frac{1}{2}\mu_{\pi_{\text{ref}}} + \frac{1}{2}\mu_\pi$ and

$$2516 \quad \mathcal{R}(d) = \frac{1}{2}\mathbb{E}_{x \sim \mu_{\pi_{\text{ref}}}}[\log(1 + e^{d(x)})] + \frac{1}{2}\mathbb{E}_{x \sim \mu_\pi}[\log(1 + e^{-d(x)})] = \mathbb{E}_{x \sim \mu_{\text{mix}}} [r_x(d)]$$

2518 Taking expectation with respect to the mixture μ_{mix} gives

$$2520 \quad \mathcal{R}(d) - \mathcal{R}(d^*) = \mathbb{E}_{x \sim \mu_{\text{mix}}}[r_x(d(x)) - r_x(d^*(x))] \geq \frac{\lambda}{2}\mathbb{E}_{x \sim \mu_{\text{mix}}}[(d(x) - d^*(x))^2],$$

2522 which is the desired inequality. \square

2524 We now establish three Lipschitz bounds that will be used repeatedly in the analysis.

2525 **Lemma 2** (Lipschitz bounds). *Let $L_+ := \frac{\beta}{1+\beta}$ and $L_- := \frac{1}{1+\alpha}$. For all $d, \tilde{d} \in [\log \alpha, \log \beta]$, the
2526 following hold:*

- 2528 1. $|\log(1 + e^d) - \log(1 + e^{\tilde{d}})| \leq L_+ |d - \tilde{d}|$.
- 2529 2. $|\log(1 + e^{-d}) - \log(1 + e^{-\tilde{d}})| \leq L_- |d - \tilde{d}|$.
- 2531 3. $|e^d - e^{\tilde{d}}| \leq \beta |d - \tilde{d}|$.

2534 *Proof.* (1) Define $f_+(u) = \log(1 + e^u)$. Then $f'_+(u) = \sigma(u) = \frac{e^u}{1+e^u}$. On $u \in [\log \alpha, \log \beta]$ we
2535 have $e^u \in [\alpha, \beta]$, hence

$$2537 \quad |f'_+(u)| = \frac{e^u}{1+e^u} \leq \sup_{y \in [\alpha, \beta]} \frac{y}{1+y} = \frac{\beta}{1+\beta} = L_+.$$

2538 By the mean value theorem, $|f_+(d) - f_+(\tilde{d})| \leq L_+ |d - \tilde{d}|$.
 2539

2540 (2) Define $f_-(u) = \log(1 + e^{-u})$. Then $f'_-(u) = -\sigma(-u) = -\frac{1}{1+e^u}$. For $u \in [\log \alpha, \log \beta]$,

2541

$$2542 |f'_-(u)| = \frac{1}{1+e^u} \leq \sup_{y \in [\alpha, \beta]} \frac{1}{1+y} = \frac{1}{1+\alpha} = L_-.$$

2543

2544 Again by the mean value theorem, $|f_-(d) - f_-(\tilde{d})| \leq L_- |d - \tilde{d}|$.
 2545

2546 (3) For $g(u) = e^u$ we have $g'(u) = e^u$. On $[\log \alpha, \log \beta]$, $e^u \leq \beta$, hence $|g'(u)| \leq \beta$. The mean
 2547 value theorem yields $|e^d - e^{\tilde{d}}| \leq \beta |d - \tilde{d}|$. \square
 2548

2549 **Lemma 3** (Uniform deviation over the finite cover). *Let Assumption 1 hold true. Define*

2550

$$B := \frac{1}{2} \log(1 + \beta) + \frac{1}{2} \log\left(1 + \frac{1}{\alpha}\right).$$

2551

2552 Let $\bar{\mathcal{D}}$ be a finite cover (L_1 optimistic cover from Assumption 1) with cardinality $|\bar{\mathcal{D}}| = M$. Define
 2553 $n := \min\{n_{\text{ref}}, n_{\pi}\}$ and $\eta := \sqrt{\frac{\log(M/\delta)}{n}}$ for any $\delta \in (0, 1)$. Then, with probability at least $1 - \delta$,

2554

$$\sup_{\bar{d} \in \bar{\mathcal{D}}} |\hat{\mathcal{R}}(\bar{d}) - \mathcal{R}(\bar{d})| \leq 2B\eta. \quad (50)$$

2555

2556

2557 *Proof.* Fix $\bar{d} \in \bar{\mathcal{D}}$. Define

2558

$$\Delta_{\text{ref}}(\bar{d}) := \sum_{i=1}^{n_{\text{ref}}} [\log(1 + e^{\bar{d}(x)})] - \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [\log(1 + e^{\bar{d}(x)})],$$

2559

2560

$$\Delta_{\pi}(\bar{d}) := \sum_{i=1}^{n_{\pi}} [\log(1 + e^{-\bar{d}(x)})] - \mathbb{E}_{\mu_{\pi}} [\log(1 + e^{-\bar{d}(x)})].$$

2561

2562 Then

2563

$$\hat{\mathcal{R}}(\bar{d}) - \mathcal{R}(\bar{d}) = \frac{1}{2} \Delta_{\text{ref}}(\bar{d}) + \frac{1}{2} \Delta_{\pi}(\bar{d}), \Rightarrow |\hat{\mathcal{R}}(\bar{d}) - \mathcal{R}(\bar{d})| \leq \frac{1}{2} |\Delta_{\text{ref}}(\bar{d})| + \frac{1}{2} |\Delta_{\pi}(\bar{d})|.$$

2564

2565 By the boundedness of Assumption 1, we have $d(x) \in [\log \alpha, \log \beta]$, each summand satisfies

2566

$$0 \leq \log(1 + e^{\bar{d}(x)}) \leq \log(1 + \beta), \quad 0 \leq \log(1 + e^{-\bar{d}(x)}) \leq \log\left(1 + \frac{1}{\alpha}\right).$$

2567

2568 Hence, by the Hoeffding's inequality, for any $t > 0$,

2569

$$\mathbb{P}(|\Delta_{\text{ref}}(\bar{d})| \geq t) \leq 2 \exp\left(-\frac{2n_{\text{ref}}t^2}{\log(1 + \beta)^2}\right), \quad \mathbb{P}(|\Delta_{\pi}(\bar{d})| \geq t) \leq 2 \exp\left(-\frac{2n_{\pi}t^2}{\log(1 + \frac{1}{\alpha})^2}\right).$$

2570

2571 Choose

2572

$$t_{\text{ref}} := \log(1 + \beta) \sqrt{\frac{\log(2M/\delta)}{2n_{\text{ref}}}}, \quad t_{\pi} := \log\left(1 + \frac{1}{\alpha}\right) \sqrt{\frac{\log(2M/\delta)}{2n_{\pi}}}.$$

2573

2574 Then $\mathbb{P}(|\Delta_{\text{ref}}(\bar{d})| \geq t_{\text{ref}}) \leq \delta/M$ and $\mathbb{P}(|\Delta_{\pi}(\bar{d})| \geq t_{\pi}) \leq \delta/M$. Taking a union bound over all
 2575 $\bar{d} \in \bar{\mathcal{D}}$ yields, with probability at least $1 - \delta$,

2576

$$\sup_{\bar{d} \in \bar{\mathcal{D}}} |\hat{\mathcal{R}}(\bar{d}) - \mathcal{R}(\bar{d})| \leq \frac{1}{2} t_{\text{ref}} + \frac{1}{2} t_{\pi}.$$

2577

2578 Finally, since $n = \min\{n_{\text{ref}}, n_{\pi}\}$, we have

2579

$$\sqrt{\frac{\log(2M/\delta)}{2n_{\text{ref}}}} \leq \sqrt{\frac{\log(M/\delta)}{n}}, \quad \sqrt{\frac{\log(2M/\delta)}{2n_{\pi}}} \leq \sqrt{\frac{\log(M/\delta)}{n}},$$

2580

2581 up to benign constant factors that we absorb into the front constant. Using the definition of B and
 2582 setting $\eta = \sqrt{\log(M/\delta)/n}$ gives

2583

$$\sup_{\bar{d} \in \bar{\mathcal{D}}} |\hat{\mathcal{R}}(\bar{d}) - \mathcal{R}(\bar{d})| \leq 2B\eta,$$

2584

2585 which is Equation 50. \square

We now prove the transfer bounds that move deviations on the cover element \bar{d} back to an arbitrary d , measured either in the original loss or the empirical loss. These inequalities will let us relate risk differences to L_1 discrepancies between ratio functions.

Lemma 4 (Transfer bounds from a cover element to an arbitrary discriminator). *Let Assumption 1 hold true. Let $L_+ := \frac{\beta}{1+\beta}$, $L_- := \frac{1}{1+\alpha}$, and define*

$$C_\Delta := \frac{L_+ + \beta L_-}{2\alpha}.$$

Then:

$$|\mathcal{R}(\bar{d}) - \mathcal{R}(d)| \leq C_\Delta \gamma, \quad (51)$$

$$|\hat{\mathcal{R}}(\bar{d}) - \hat{\mathcal{R}}(d)| \leq C_\Delta \gamma. \quad (52)$$

Proof. Start from the definition

$$\mathcal{R}(d) = \frac{1}{2} \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [\log(1 + e^d)] + \frac{1}{2} \mathbb{E}_{\mu_\pi} [\log(1 + e^{-d})].$$

By the triangle inequality,

$$|\mathcal{R}(\bar{d}) - \mathcal{R}(d)| \leq \frac{1}{2} \mathbb{E}_{\mu_{\pi_{\text{ref}}}} |\log(1 + e^{\bar{d}}) - \log(1 + e^d)| + \frac{1}{2} \mathbb{E}_{\mu_\pi} |\log(1 + e^{-\bar{d}}) - \log(1 + e^{-d})|. \quad (53)$$

Recall the Lipschitz bounds derived in Lemma 2:

$$|\log(1 + e^u) - \log(1 + e^v)| \leq L_+ |u - v|, \quad |\log(1 + e^{-u}) - \log(1 + e^{-v})| \leq L_- |u - v|$$

yield from Equation 53

$$|\mathcal{R}(\bar{d}) - \mathcal{R}(d)| \leq \frac{1}{2} L_+ \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [|\bar{d} - d|] + \frac{1}{2} L_- \mathbb{E}_{\mu_\pi} [|\bar{d} - d|]. \quad (54)$$

Since $\mu_\pi = L \cdot \mu_{\pi_{\text{ref}}}$ and $L \leq \beta$ by boundedness from Assumption 1,

$$\mathbb{E}_{\mu_\pi} [|\bar{d} - d|] = \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [L |\bar{d} - d|] \leq \beta \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [|\bar{d} - d|]. \quad (55)$$

Plug Equation 55 back into Equation 54:

$$|\mathcal{R}(\bar{d}) - \mathcal{R}(d)| \leq \frac{L_+ + \beta L_-}{2} \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [|\bar{d} - d|]. \quad (56)$$

Next, we need to convert $\mathbb{E}_{\mu_{\pi_{\text{ref}}}} [|\bar{d} - d|]$ to $\mathbb{E}_{\mu_{\pi_{\text{ref}}}} [|\bar{L} - L|]$. Use the mean value theorem for \log on $[\alpha, \beta]$,

$$|\bar{d} - d| = |\log \bar{L} - \log L| = \frac{1}{\xi} |\bar{L} - L| \leq \frac{1}{\alpha} |\bar{L} - L|, \quad \xi \text{ between } \bar{L} \text{ and } L.$$

Therefore, we have

$$\mathbb{E}_{\mu_{\pi_{\text{ref}}}} [|\bar{d} - d|] \leq \frac{1}{\alpha} \mathbb{E}_{\mu_{\pi_{\text{ref}}}} [|\bar{L} - L|]. \quad (57)$$

Combining Equation 56 and Equation 57:

$$|\mathcal{R}(\bar{d}) - \mathcal{R}(d)| \leq C_\Delta \mathbb{E}_{x \sim \mu_{\pi_{\text{ref}}}} [|\bar{L}(x) - L(x)|], \quad C_\Delta = (L_+ + \beta L_-)/(2\alpha).$$

Finally, according to the L_1 optimistic cover of Assumption 1, we have

$$\mathbb{E}_{\mu_{\pi_{\text{ref}}}} [|\bar{L} - L|] \leq \gamma$$

We get

$$|\mathcal{R}(\bar{d}) - \mathcal{R}(d)| \leq C_\Delta \gamma.$$

The same derivation holds if we replace expectations by empirical averages (sample means). Every inequality we used above (triangle inequality and the Lipschitz bounds) is pointwise and hence holds averaging over a finite sample instead of the distribution. Concretely,

$$|\hat{\mathcal{R}}(\bar{d}) - \hat{\mathcal{R}}(d)| \leq C_\Delta \sum_{i=1}^{n_{\text{ref}}} [|\bar{L}(x) - L(x)|] \leq C_\Delta \gamma$$

where the last inequality uses the empirical L_1 closeness of \bar{L} and L .

□

We now control the excess loss of the empirical minimizer by combining the uniform deviation bound over the optimistic cover with the transfer inequalities.

Lemma 5 (Excess-loss bound for the empirical minimizer). *Let Assumption 1 hold true. Let $B = \frac{1}{2} \log(1+\beta) + \frac{1}{2} \log(1+1/\alpha)$, $C_\Delta = (L_+ + \beta L_-)/(2\alpha)$ with $L_+ = \beta/(1+\beta)$ and $L_- = 1/(1+\alpha)$, and $\eta = \sqrt{\log(M/\delta)/n}$ where $n = \min\{n_{\text{ref}}, n_\pi\}$ defined as before. Then, with probability at least $1 - \delta$,*

$$\mathcal{R}(\hat{d}) - \mathcal{R}(d^*) \leq 3C_\Delta \gamma + 4B\eta. \quad (58)$$

Proof. Since $\hat{d} \in \arg \min_{d \in \mathcal{D}} \hat{\mathcal{R}}(d)$, we have:

$$\hat{\mathcal{R}}(\hat{d}) \leq \hat{\mathcal{R}}(d) \quad \text{for all } d,$$

in particular for $d = d^*$ and for $d = \bar{d}^*$ (the cover of d^* , $\bar{d}^* = \log \bar{L}^*$). Start with a standard add-subtract trick:

$$\mathcal{R}(\hat{d}) - \mathcal{R}(d^*) = (\mathcal{R}(\hat{d}) - \hat{\mathcal{R}}(\hat{d})) + \underbrace{(\hat{\mathcal{R}}(\hat{d}) - \hat{\mathcal{R}}(\bar{d}^*))}_{\leq 0} + (\hat{\mathcal{R}}(\bar{d}^*) - \mathcal{R}(\bar{d}^*)) + (\mathcal{R}(\bar{d}^*) - \mathcal{R}(d^*)), \quad (59)$$

For the first difference in Equation 59, insert the cover element $\bar{d} = \log \bar{L}$ of \hat{d} and apply the transfer bounds (Lemma 4) and the uniform deviation bound over the finite cover (Lemma 3):

$$\begin{aligned} \mathcal{R}(\hat{d}) - \hat{\mathcal{R}}(\hat{d}) &= (\mathcal{R}(\hat{d}) - \mathcal{R}(\bar{d})) + (\mathcal{R}(\bar{d}) - \hat{\mathcal{R}}(\bar{d})) + (\hat{\mathcal{R}}(\bar{d}) - \hat{\mathcal{R}}(\hat{d})) \\ &\leq |\mathcal{R}(\hat{d}) - \mathcal{R}(\bar{d})| + |\mathcal{R}(\bar{d}) - \hat{\mathcal{R}}(\bar{d})| + |\hat{\mathcal{R}}(\bar{d}) - \hat{\mathcal{R}}(\hat{d})| \\ &\leq C_\Delta \gamma + 2B\eta + C_\Delta \gamma. \end{aligned}$$

The first term uses the transfer bounds from \hat{d} to its cover \bar{d} . The middle term uses the uniform deviation bound on the finite cover. It applies directly because $\bar{d} \in \bar{\mathcal{D}}$. The third term uses the empirical transfer bounds. Thus, we have

$$\mathcal{R}(\hat{d}) - \hat{\mathcal{R}}(\hat{d}) \leq 2C_\Delta \gamma + 2B\eta. \quad (60)$$

Returning to Equation 59, the middle term is nonpositive by optimality, and the remaining two terms are bounded by the same two lemmas:

$$|\hat{\mathcal{R}}(\bar{d}^*) - \mathcal{R}(\bar{d}^*)| \leq 2B\eta, \quad |\mathcal{R}(\bar{d}^*) - \mathcal{R}(d^*)| \leq C_\Delta \gamma.$$

Combining with Equation 60 yields

$$\mathcal{R}(\hat{d}) - \mathcal{R}(d^*) \leq (2C_\Delta \gamma + 2B\eta) + 2B\eta + C_\Delta \gamma = 3C_\Delta \gamma + 4B\eta,$$

which is Equation 58. \square

Finally, we derive the occupancy ratio error bound.

Theorem 3 (Occupancy ratio L_1 error bound). *Let Assumption 1 hold true. Let*

$$B := \frac{1}{2} \log(1 + \beta) + \frac{1}{2} \log(1 + 1/\alpha), \quad L_+ := \frac{\beta}{1 + \beta}, \quad L_- := \frac{1}{1 + \alpha},$$

$$C_\Delta := \frac{L_+ + \beta L_-}{2\alpha}, \quad \lambda := \frac{\min\{\alpha, \beta\}}{(1 + \max\{\alpha, \beta\})^2} > 0,$$

and $n := \min\{n_{\text{ref}}, n_\pi\}$, $\eta := \sqrt{\log(M/\delta)/n}$ for any $\delta \in (0, 1)$. Let $\hat{d} \in \arg \min_{d \in \mathcal{D}} \hat{\mathcal{R}}(d)$ be the empirical minimizer, $\hat{L} := e^{\hat{d}}$, and $L^* = e^{d^*}$ be the true ratio. Then, with probability at least $1 - \delta$,

$$\mathbb{E}_{x \sim \mu_{\pi_{\text{ref}}}} [|\hat{L}(x) - L^*(x)|] \leq \beta \sqrt{\frac{4}{\lambda}} \sqrt{3C_\Delta \gamma + 4B\eta}. \quad (61)$$

2700 *Proof.* By Lemma 5 (excess-risk bound for the empirical minimizer),
 2701

$$2702 \quad \mathcal{R}(\hat{d}) - \mathcal{R}(d^*) \leq 3C_\Delta \gamma + 4B\eta \quad \text{with probability at least } 1 - \delta.$$

2703 From Lemma 1 we have,
 2704

$$2705 \quad \mathbb{E}_{x \sim \mu_{\text{mix}}}[(\hat{d}(x) - d^*(x))^2] \leq \frac{2}{\lambda} (\mathcal{R}(\hat{d}) - \mathcal{R}(d^*)) \leq \frac{2}{\lambda} (3C_\Delta \gamma + 4B\eta),$$

2707 where $\mu_{\text{mix}} = \frac{1}{2}\mu_{\pi_{\text{ref}}} + \frac{1}{2}\mu_\pi$. Since $\mu_{\text{mix}} \geq \frac{1}{2}\mu_{\pi_{\text{ref}}}$, we have $\mathbb{E}_{\mu_{\pi_{\text{ref}}}}[\cdot] \leq 2\mathbb{E}_{\mu_{\text{mix}}}[\cdot]$, and
 2708

$$2709 \quad \mathbb{E}_{x \sim \mu_{\pi_{\text{ref}}}}[(\hat{d}(x) - d^*(x))^2] \leq \frac{4}{\lambda} (3C_\Delta \gamma + 4B\eta).$$

2711 Finally, by the third point of Lemma 2 (the exponential map is β -Lipschitz on $[\log \alpha, \log \beta]$) and
 2712 Cauchy–Schwarz,
 2713

$$2714 \quad \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[|\hat{L} - L^*|] \leq \beta \mathbb{E}_{\mu_{\pi_{\text{ref}}}}[|\hat{d} - d^*|] \leq \beta \sqrt{\mathbb{E}_{\mu_{\pi_{\text{ref}}}}[(\hat{d} - d^*)^2]} \leq \beta \sqrt{\frac{4}{\lambda} \sqrt{3C_\Delta \gamma + 4B\eta}},$$

2716 which is Equation 61. \square
 2717

2718 G.2 GUARANTEES FOR MAX-MIN WITH OCCUPANCY MEASURE APPROXIMATION

2719 In this section, we establish convergence guarantees for our Max-Min Algorithm 2. Our analysis
 2720 follows the Reinforcement Learning with General Utility (RLGU) (Zhang et al., 2022; Barakat
 2721 et al., 2024), where given a utility function $F(\cdot)$, $\theta \mapsto F(\mu_{\pi_\theta})$ over the policy-induced occupancy
 2722 measure μ_{π_θ} , the goal of RLGU is to find a policy π_θ^* such that $\pi_\theta^* \in \arg \max_\theta F(\mu_{\pi_\theta})$. In RLGU,
 2723 there is no reward function. Instead, we can view $\nabla_\theta F(\mu_\theta)$ as a pseudo-reward depending on the
 2724 unknown occupancy measure induced by the policy. The procedure for solving RLGU follows
 2725 three steps: (i) estimate the occupancy μ_{π_θ} (e.g., by MLE), (ii) form the pseudo-reward from this
 2726 estimate, and (iii) update the policy. Our Max-Min algorithm mirrors this pipeline. Specifically,
 2727 we first estimate the occupancy ratio by training a discriminator. Then we construct the worst-case
 2728 reward using Equation 19 from the estimation. Finally, we perform a policy update. Consequently,
 2729 the general RLGU sample complexity guarantees apply to our algorithm after replacing the pseudo-
 2730 reward $\nabla_\theta F(\mu_\theta)$ with our worst-case reward and substituting their occupancy-estimation error with
 2731 our occupancy-ratio error obtained above. We formalize this correspondence and state the resulting
 2732 bounds below.

2733 For each iteration of our Max-Min algorithm $t = 1, 2, \dots, T$, let the pseudo-reward $r_t(s, a)$ is defined
 2734 as in Equation 19. Let's define

$$2735 \quad F(\mu_t) := \frac{1}{4\lambda_3} \int \frac{\mu_t(s, a)^2}{\mu_{\pi_{\text{ref}}}(s, a)} d(s, a) - \int c(s, a) \mu_t(s, a) d(s, a), \quad (62)$$

2736 where $c(s, a) := \frac{1}{2\lambda_3} (\lambda_1 \frac{R_{\text{proxy}}(s, a)}{V} + \lambda_2)$. By construction, we have $r_t(s, a) = \nabla_\mu F(\mu_t)(s, a)$,
 2737 which means the utility gradient in μ is exactly the pseudo-reward. Let's $\hat{\mu}_t := \hat{L}_t \cdot \mu_{\pi_{\text{ref}}}$ be
 2738 the occupancy estimator, $\hat{r}_t(s, a) := \nabla_\mu F(\hat{\mu}_t)(s, a)$ be the estimated pseudo-reward and $F^* \in$
 2739 $\max_\theta F(\mu_{\pi_\theta})$ be the maximum.

2740 We next introduce some assumptions that are required for our results, which are adapted
 2741 from (Barakat et al., 2024).

2742 **Assumption 2** (Policy parametrization, Assumption 6 from (Barakat et al., 2024)). *For every*
 2743 $(s, a) \in \mathcal{S} \times \mathcal{A}$ *and every* $\theta \in \mathbb{R}^d$, *the policy has full support, i.e.* $\pi_\theta(a | s) > 0$. *Moreover, the*
 2744 *mapping* $\theta \mapsto \pi_\theta(a | s)$ *is continuously differentiable, and the score function* $\theta \mapsto \nabla_\theta \log \pi_\theta(a | s)$
 2745 *is uniformly bounded:*

$$2746 \quad \|\nabla_\theta \log \pi_\theta(a | s)\| \leq l_\psi \quad \text{for some constant } l_\psi > 0 \text{ and all } (s, a), \theta.$$

2747 This assumption typically holds in practice, for instance, with the standard softmax policy parameterization.
 2748 Next, we make a smoothness assumption on the utility function, which is crucial in
 2749 deriving the final convergence bound. We also verify that the defined utility function in Equation 62
 2750 satisfies the smoothness assumption.

2754 **Assumption 3** (General utility smoothness, Assumption 7 from (Barakat et al., 2024)). *For utility*
 2755 *function $F(\cdot), \theta \mapsto F(\mu_{\pi_\theta})$, there exist constants $L_\mu > 0$ such that for all $\mu_1, \mu_2 \in \mathcal{X}$,*
 2756

$$2757 \quad \|\nabla_\mu F(\mu_1)\|_2 \leq \ell_\mu \quad \text{and} \quad \|\nabla_\mu F(\mu_1) - \nabla_\mu F(\mu_2)\|_2 \leq L_\mu \|\mu_1 - \mu_2\|_2.$$

2759 Notice that Assumption 3 holds in our setting since Hessian is diagonal with entries at most
 2760 $\nabla_\mu^2 F(\mu)(s, a) = 1/(2\lambda_3 \mu_{\pi_{\text{ref}}}(s, a))$. Thus, if $\mu_{\pi_{\text{ref}}}(s, a) \geq \rho_{\min} > 0$ on the support of all (s, a) ,
 2761 which we assume it holds, then we have that $\nabla_\mu F(\mu)$ is L_μ -Lipschitz with
 2762

$$2763 \quad \|\nabla_\mu F(\mu) - \nabla_\mu F(\mu')\|_2 \leq L_\mu \|\mu - \mu'\|_2, \quad L_\mu = \frac{1}{2|\lambda_3|} \rho_{\min}^{-1},$$

2765 Under Assumptions 2 and 3, the utility function $\theta \mapsto F(\mu_{\pi_\theta})$ is L_θ -smooth. Using these properties,
 2766 our Max-Min algorithm admits the following first-order stationarity guarantee:
 2767

2768 **Theorem 4** (Guarantee for the Max-Min update). *Assume Assumptions 2 and 3 hold. Let N be*
 2769 *the batch size for estimating the policy gradient at each iteration, α be the stepsizes satisfying*
 2770 $\alpha_t \leq 1/(2L_\theta)$, $K_{\text{conv}} := \|\mu_{\pi_{\text{ref}}}\|_\infty (\beta - \alpha)$ and
 2771

$$2772 \quad \varepsilon_L := \beta \sqrt{\frac{4}{\lambda} \sqrt{3C_\Delta \gamma + 4B\eta}} \geq \mathbb{E}_{x \sim \mu_{\pi_{\text{ref}}}} [|\widehat{L}(x) - L^*(x)|]$$

2774 Then we have:

$$2776 \quad \mathbb{E} [\|\nabla_\theta F(\mu_{\pi_{\theta_\tau}})\|^2] \leq \frac{16 (F^* - \mathbb{E}[F(\mu_{\pi_{\theta_1}})])}{\alpha T} + \frac{C_1}{N} + C_2 K_{\text{conv}} \varepsilon_L, \quad (63)$$

2778 where τ is drawn uniformly from $\{1, \dots, T\}$, expectation is w.r.t all randomness (in (θ_t) and τ),
 2779 $C_1 = \frac{8l_\mu^2 l_\psi^2}{(1-\gamma')^4}$ and $C_2 = \frac{8l_\psi^2 L_\mu^2}{(1-\gamma')^4}$ with γ' be the discount factor in RL.
 2780

2782 *Proof.* Since we already verified that Assumptions 2 and 3 hold in our setting, according to Theorem 8 in (Barakat et al., 2024)), we directly have
 2783

$$2785 \quad \frac{1}{T} \sum_{t=1}^T \mathbb{E} [\|\nabla_\theta F(\mu_{\pi_{\theta_\tau}})\|^2] \leq \frac{16 (F^* - \mathbb{E}[F(\mu_{\pi_{\theta_1}})])}{\alpha T} + \frac{C_1}{N} + C_2 \frac{1}{T} \sum_{t=1}^T \mathbb{E} [\|\widehat{\mu}_t - \mu_t\|_2^2].$$

2788 To control the last term via the ratio error derived in Theorem 3, note that with $\widehat{\mu}_t = \widehat{L}_t \mu_{\pi_{\text{ref}}}$ and
 2789 $\mu_t = L_t \mu_{\pi_{\text{ref}}}$,
 2790

$$2791 \quad \|\widehat{\mu}_t - \mu_t\|_2^2 = \sum_{s,a} \mu_{\pi_{\text{ref}}}(s, a)^2 (\widehat{L}_t(s, a) - L_t(s, a))^2 \leq \|\mu_{\pi_{\text{ref}}}\|_\infty (\beta - \alpha) \sum_{s,a} \mu_{\pi_{\text{ref}}}(s, a) |\widehat{L}_t - L_t|,$$

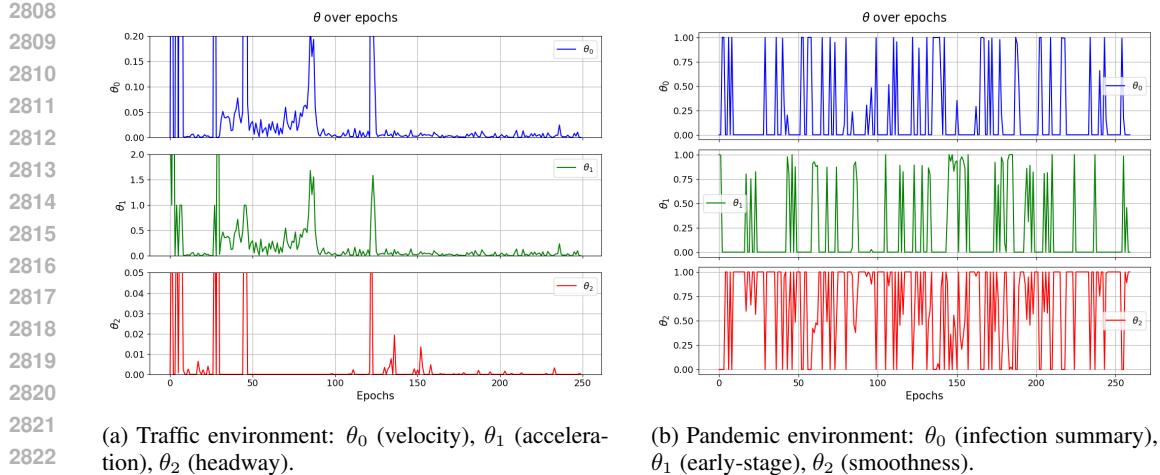
2793 using $x^2 \leq (\beta - \alpha)|x|$ for $|x| \leq \beta - \alpha$ and $\mu_{\pi_{\text{ref}}}^2 \leq \|\mu_{\pi_{\text{ref}}}\|_\infty \mu_{\pi_{\text{ref}}}$. Therefore,
 2794

$$2795 \quad \mathbb{E} [\|\widehat{\mu}_t - \mu_t\|_2^2] \leq K_{\text{conv}} \mathbb{E}_{x \sim \mu_{\pi_{\text{ref}}}} [|\widehat{L}_t(x) - L_t(x)|] \leq K_{\text{conv}} \varepsilon_L.$$

2797 Averaging over t and drawing τ uniformly from $\{1, \dots, T\}$ yields Equation 63. \square
 2798

2799 G.3 CONVERGENCES FOR LINEAR MAX-MIN ALGORITHM

2801 For our Linear Max-Min algorithm, it is challenging to derive a convergence bound directly. How-
 2802 ever, as discussed in Appendix E.6, the inner optimization problem, i.e., finding the worst-case
 2803 reward for a given policy, admits a **globally optimal closed-form solution** under our formulation
 2804 in the tabular setting. Therefore, for any given policy π , we have access to an oracle that outputs
 2805 the optimal worst-case reward R^* , and our Linear Max-Min algorithm can be viewed as alternating
 2806 between gradient ascent on π and the optimal minimization on R^* . As shown in Section 4 of (Jin
 2807 et al., 2020), our algorithm converges, and the resulting policy π corresponds to an approximate
 2808 stationary point of the outer optimization problem.



(a) Traffic environment: θ_0 (velocity), θ_1 (acceleration), θ_2 (headway). (b) Pandemic environment: θ_0 (infection summary), θ_1 (early-stage), θ_2 (smoothness).

Figure 4: Evolution of adversarial reward weights θ over training epochs for different environments using the Linear Maxmin method.

H ADDITIONAL EXPERIMENT RESULTS

H.1 FEATURE WEIGHTS IN LINEAR MAX-MIN OPTIMIZATION DURING TRAINING

Figure 4 visualizes the evolution of each component of the linear worst-case reward weight vector θ during training in the Traffic and Pandemic environments. We observe distinct behaviors in the dynamics of θ across tasks.

In the Traffic environment (Figure 4a), we observe that the three θ parameters vary significantly in scale. Specifically, θ_1 (acceleration) exhibits the largest magnitude, ranging from 0 to 2, while θ_2 (headway) has the smallest scale, ranging from 0 to 0.05. This highlights how the linear max-min algorithm assigns different levels of penalization to each feature. Moreover, we also observe distinct phases in the dynamics of θ over the course of training. In the early epochs (<50), all components, especially θ_1 (acceleration) and θ_2 (headway), exhibit high-frequency fluctuations. At this point, the dual optimization problem is not yet well-conditioned, and the adversarial reward is highly sensitive to small changes in occupancy or feature values. As training progresses (~epochs 100–250), the parameters begin to stabilize. Most notably, θ_2 (headway) converges close to zero and remains suppressed, indicating that the worst-case reward does not emphasize this feature. This may suggest that headway is less harmful under adversarial reweighting compared to others (velocity or accel) or is already well aligned with the reference policy π_{ref} . Meanwhile, θ_1 (acceleration) consistently exhibits higher values and sharper spikes than the other components. This indicates that acceleration plays a dominant role in the adversarial reward, likely because policies that optimize for the proxy reward tend to exploit aggressive acceleration patterns that diverge significantly from the behavior of π_{ref} . In contrast, θ_0 (velocity) remains small and relatively stable throughout training, suggesting that speed alone is not strongly penalized under adversarial interpretations.

Overall, the observed pattern reflects the interpretability and sparsity benefits of the linear max-min formulation. The model is able to selectively emphasize features that are most vulnerable to reward hacking, while suppressing those that are either irrelevant or well-aligned. This structured behavior supports the practical value of using linearly parameterized worst-case rewards to improve policy robustness.

In the Pandemic environment (Figure 4b), unlike the Traffic environment, where θ converged to a sparse and interpretable solution, we observe high variability across all components throughout training. In particular, we find the following pattern:

1. **Persistent fluctuations.** All three components exhibit frequent oscillations over the course of 260 epochs. This ongoing instability suggests that the adversarial reward continually

2862 adapts as the policy changes, likely due to the environment’s temporal sensitivity and complex dynamics.
 2863

2864 2. **θ_2 (smoothness) remains active.** The smoothness-related component θ_2 is frequently non-
 2865 zero and relatively stable compared to the others. This indicates that the worst-case reward
 2866 consistently emphasizes penalizing erratic or unstable responses in the infection trajectory
 2867 — a behavior often neglected by naive proxy metrics.
 2868

2869 3. **θ_1 (early-stage transitions) is highly volatile.** The component associated with early in-
 2870 fection stage changes spikes intermittently. This suggests that early-stage mismanagement
 2871 is a recurring vulnerability in the learned policy that the adversarial reward seeks to exploit.
 2872

2873 4. **θ_0 (overall infection) activates intermittently.** Although θ_0 sometimes spikes, it does not
 2874 dominate the adversarial reward. This may indicate that the learned policy already accounts
 2875 for infection magnitude reasonably well, or that smoothness and early-stage control offer
 2876 more leverage for reward hacking under the proxy constraint.
 2877

2878 Overall, this pattern highlights that in more dynamic and temporally complex environments like
 2879 Pandemic, the worst-case reward remains non-sparse and adapts to different policy weaknesses
 2880 throughout training. In contrast to the Traffic environment, adversarial emphasis here is broader
 2881 and more reactive.
 2882

2883 H.2 ADDITIONAL WORST-CASE PERFORMANCE RESULTS

2884 Table 6: Evaluation results on Traffic, Pandemic, Glucose, and RLHF environments. All policies are trained
 2885 using **only the proxy reward**. In Traffic, the proxy reward is based on *vel*, *accel*, *headway* (1, 1, 0.1), while the
 2886 true reward uses *commute*, *accel*, *headway* (1, 1, 0.1). In Pandemic, the proxy reward includes *infection*, *lower*
 2887 *stage*, *smooth changes* (10, 0.1, 0.01), while the true reward additionally includes *political* with weight 10 after
 2888 *infection*. In Glucose, the proxy uses *expected patient cost*, and the true reward uses *magni_bg*. In RLHF, the
 2889 proxy uses a 70M LLM, and the true reward uses a 8B LLM. We report θ in the same order as feature weights.
 2890 **Occ** denotes total occupancy over state-action pairs unseen by π_{ref} , where discriminator outputs infinity.
 2891

| Env | | Traffic | | | | |
|----------------|--|-----------------------|-------------|--------------------|----------------------------------|-------------------------------------|
| Method | | True | Proxy | Worst | Linear Worst (θ) | Linear Worst* (θ) |
| ORPO | | 16.91±0.12 | 3.41±0.13 | -1.96e+04±0.02e+04 | -0.69±0.01 (0.71, 0.21, 0.69) | -0.83±0.02 (0.63, 0.12, 0.97) |
| ORPO* | | 10.26±0.09 | 1.35±0.09 | -1.35e+04±0.02e+04 | -0.44±0.02 (0.46, 0.18, 0.86) | -0.45±0.01 (0.58, 0.06, 0.81) |
| Max-Min | | 12.70±0.06 | 3.63±0.09 | -268.31±4.14 | -0.06±0.01 (0.01, 0.02, 0.96) | -0.06±0.01 (0.001, 0.02, 0.99) |
| Linear Max-Min | | 16.46±0.10 | 2.40±0.11 | -1.19e+04±0.01e+04 | 0.20±0.01 (0.64, 0.07, 0.76) | -0.12±0.01 (0.91, 0.01, 0.67) |
| Env | | Pandemic | | | | |
| Method | | True | Proxy | Worst | Linear Worst (θ) | Linear Worst* (θ) |
| ORPO | | -1.04±0.21 | 1.75±0.19 | -5.31e+06±0.01e+06 | -2.41±0.02 (0.23, 0.95, 0.17) | -2.65±0.02 (0.02, 0.95, 0.92, 0.08) |
| ORPO* | | 1.18±0.19 | 1.68±0.19 | -4.46e+06±0.03e+06 | -1.36±0.01 (0.25, 0, 0.97, 0.13) | -1.36±0.01 (0.25, 0, 0.97, 0.13) |
| Max-Min | | 1.25±0.18 | 1.25±0.18 | -63.29±3.35 | -1.11±0.01 (0.14, 0.99, 0.01) | -1.11±0.01 (0.14, 0, 0.99, 0.01) |
| Linear Max-Min | | 3.65±0.11 | 7.60±0.13 | -6.82e+05±0.01e+05 | 0.65±0.01 (0.001, 0.23, 0.02) | -0.17±0.02 (0.01, 0.97, 0.22, 0.09) |
| Env | | Glucose | | | | |
| Method | | True($\times 10^3$) | Proxy | Worst | True | Proxy |
| ORPO | | 6.0±0.1 | 100.48±0.54 | -27.54±0.32 | 8.30 ± 1.07 | 0.63±0.21 |
| ORPO* | | 6.3±0.2 | 116.36±0.56 | -8.79±0.27 | N/A | N/A |
| Max-Min | | 6.3±0.1 | 102.66±0.58 | -1.71±0.25 | 5.38 ± 0.92 | 0.84±0.11 |

2900 **Adversarial Weight Analysis.** In Table 6, we also report the adversarial weight vectors θ for
 2901 each policy. These weights reveal which features are most vulnerable to proxy exploitation under
 2902 the learned policy and can be used to diagnose and revise the proxy reward function, thereby im-
 2903 proving robustness. This highlights the interpretability benefits of our framework. Moreover, several
 2904 patterns emerge from the results. In the Traffic environment, first, we observe a clear dominance of
 2905 the headway feature, with all methods assigning it the highest weight. This suggests that headway
 2906 is the most critical component exposed to reward hacking under correlation constraints. Second,
 2907 the acceleration feature is consistently downweighted across all methods. This indicates that ac-
 2908 celeration may be less prone to exploitation or already well aligned with the reference policy. Third,
 2909 the velocity feature is moderately emphasized by Linear Max-Min and ORPO (e.g., 0.64 and
 2910 0.71), while Max-Min nearly suppresses it (0.01). This contrast suggests that Linear Max-Min
 2911 anticipates some vulnerability from velocity deviations, while Max-Min focuses almost entirely on
 2912 headway. In the Pandemic environment, first, both ORPO* and Max-Min assign zero weight to
 2913 the political feature. This occurs because the expected feature value under their policies is exactly
 2914 zero, making the correlation constraint inactive for that dimension. Interestingly, this feature plays
 2915 a significant role in the adversarial rewards for both ORPO and Linear Max-Min, with their cor-
 2916 responding θ assigning non-negligible weight to it (e.g., 0.95 and 0.97 respectively). This suggests

2916 that these policies expose themselves to vulnerability in feature dimensions that are entirely ignored
 2917 by Max–Min and ORPO*. Second, the lower stage feature consistently receives the highest weight
 2918 across all methods, indicating it is the most sensitive component under proxy misalignment.
 2919

2920 Table 7: Evaluation results on Tomato environments. All policies are trained using **only the proxy**
 2921 **reward**. In Tomato, the proxy includes *number of watered tomatoes* plus a bonus at a specific state
 2922 (sprinkler), while the true reward only measures *watered tomatoes*. **Occ** in the Tomato environment
 2923 denotes total occupancy over state-action pairs unseen by π_{ref} , based on 1000 sampled trajectories.
 2924 **Worst** refers to the expected worst-case reward computed while excluding those unseen state-action
 2925 pairs. **Worst*** denotes the actual expected worst-case reward, while R_{\min} represents the minimum
 2926 possible reward of any state-action pair. All rewards are normalized according to the reference
 2927 policy π_{ref} .
 2928

| Method | Tomato | | | | |
|---------|-----------|-----------|------------|-------------------|-----------------------------|
| | True | Proxy | Worst | Occ ↓ | Worst* |
| ORPO | 6.28±0.22 | 6.83±0.28 | -1.51±0.09 | 2.50e-04±0.63e-04 | -1.51+ R_{\min} ·2.50e-04 |
| ORPO* | 4.00±0.18 | 3.98±0.23 | -1.09±0.10 | 3.09e-05±0.59e-05 | -1.09+ R_{\min} ·3.09e-05 |
| Max–Min | 4.56±0.20 | 4.68±0.25 | -1.37±0.06 | 1.01e-05±0.43e-05 | -1.37+ R_{\min} ·1.01e-05 |

2932
 2933 **Worst-Case Performance in Tomato Environment.** Table 7 reports worst-case performance re-
 2934 sults for the Tomato environment. We omit the Linear Max–Min policy from these experiments
 2935 for the following reasons. In the Tomato environment, the reward structure is difficult to express in
 2936 a clean feature-based form suitable for linear modeling. Therefore, we report only the results for the
 2937 Max–Min policy alongside the baselines.
 2938

2939 The results for the Tomato environments exhibit trends similar to those observed in other environ-
 2940 ments (Section 4.2). In particular, ORPO* appears to outperform others in the Tomato environment
 2941 in terms of worst-case performance. Recall that these results are reported under **Worst**, the expected
 2942 worst-case reward restricted to state-action pairs observed under π_{ref} . Since the Tomato environment
 2943 is discrete, we can explicitly identify which state-action pairs are unseen through sampling, enabling
 2944 clearer interpretation of their physical meaning as well as the evaluation of the actual worst-case per-
 2945 formance **Worst***. The latter corresponds to the **Worst** value plus the product of the occupancy in
 2946 unseen regions (**Occ**) and R_{\min} . Because Max–Min exhibits the lowest occupancy among all meth-
 2947 ods, it demonstrates greater robustness under varying assumptions about R_{\min} .
 2948

2949 Nevertheless, ORPO* still shows marked improvement over ORPO, both in worst-case return and in
 2950 reducing occupancy over unseen state-action pairs. As previously noted, in the Glucose environ-
 2951 ment, the discriminator fails to detect any state-action pairs missed by the reference policy. This
 2952 reinforces our earlier concern that the current discriminator training procedures may have limited
 2953 capacity to identify rare or out-of-distribution events.
 2954

2955 **Impact of Correlation Parameter Selection on Robustness.** In this section, we present addi-
 2956 tional experiment results to examine how the proxy-true reward correlation parameter r used during
 2957 training affects the robustness under varying evaluation r .
 2958

2959 Table 8: Evaluation of robustness in the Tomato environment across different training-time corre-
 2960 lation levels r . **Occ** denotes total occupancy over state-action pairs unseen by π_{ref} , based on 1000
 2961 sampled trajectories. **Worst** refers to the expected worst-case reward computed while excluding
 2962 those unseen state-action pairs.
 2963

| r | Occ | Worst ($r = 0.1$) | Worst ($r = 0.4$) | Worst ($r = 0.7$) | Worst ($r = 0.9$) |
|-----|-------------------|---------------------|---------------------|---------------------|---------------------|
| 0.1 | 1.36e-03±0.12e-03 | -1.34±0.05 | -1.12±0.04 | -0.74±0.03 | -0.27±0.02 |
| 0.4 | 1.01e-05±0.43e-05 | -1.66±0.07 | -1.37±0.06 | -0.70±0.03 | -0.05±0.02 |
| 0.7 | 1.05e-02±0.10e-02 | -2.10±0.08 | -1.82±0.07 | -1.33±0.06 | -0.66±0.04 |
| 0.9 | 1.29e-05±0.41e-05 | -9.10±0.20 | -8.92±0.18 | -7.60±0.15 | -5.49±0.12 |

2967 Table 8 and Table 9 report the robustness evaluation results in the **Tomato** and **Traffic** environ-
 2968 ments under different training-time values of the correlation parameter r . Several consistent patterns
 2969 emerge across both environments.
 2970

2970 Table 9: Evaluation of robustness in the Traffic environment across different training-time corre-
 2971 lation levels r . **Occ** denotes total occupancy over state-action pairs unseen by π_{ref} , based on 200
 2972 sampled trajectories. **Worst** refers to the expected worst-case reward computed while excluding
 2973 those unseen state-action pairs.

| 2975 r | 2976 Occ | 2977 Worst ($r = 0.1$) | 2978 Worst ($r = 0.3$) | 2979 Worst ($r = 0.5$) | 2980 Worst ($r = 0.9$) |
|----------------------------|-------------------------|--|--|--|--|
| 0.3 | 0.00 \pm 0.00 | -2794.63 \pm 42.10 | -268.31 \pm 4.14 | -82.07 \pm 2.04 | -22.03 \pm 0.88 |
| 0.5 | 0.00 \pm 0.00 | -7.71e+04 \pm 1.20e+03 | -1.95e+04 \pm 3.10e+02 | -6168.40 \pm 124.75 | -1350.22 \pm 27.95 |
| 0.9 | 9.66e-05 \pm 1.84e-05 | -3.01e+05 \pm 6.05e+03 | -9.51e+04 \pm 1.89e+03 | -2.73e+04 \pm 5.45e+02 | -9.33e+03 \pm 1.88e+02 |

2981 First, for any fixed policy (i.e., fixed training r), we observe that the expected worst-case reward
 2982 monotonically increases as the evaluation r increases. This aligns with intuition: higher correlation
 2983 levels correspond to smaller uncertainty sets over rewards, meaning the worst-case reward functions
 2984 are less adversarial. In contrast, low r values expand the reward uncertainty set, allowing more
 2985 pathological or implausible reward functions, and thus lead to more pessimistic evaluations. **How-2986**
 2987 **ever, this does not hold universally. For a fixed policy, the expected worst-case reward in Equation 27**
 2988 **is monotone with respect to r only when the policy has a positive expected proxy return (which is**
 2989 **the case here). If the policy’s expected proxy return is negative, this monotonicity condition fails.**

2990 Second, we find that training with a moderate correlation level, particularly around $r = 0.3$ to 0.4 ,
 2991 yields better robustness across a wide range of evaluation r values. In contrast, training with overly
 2992 small (e.g., $r = 0.1$ for Tomato) or large (e.g., $r = 0.9$ for Tomato and Traffic) correlation levels
 2993 degrades robustness. A small r leads to overly conservative training, anticipating extreme forms of
 2994 reward hacking and thus hurting general performance. On the other hand, a high r overly trusts the
 2995 proxy reward and fails to hedge against potential deviations, resulting in poor worst-case behavior
 2996 under reward misspecification. This trade-off highlights that intermediate values of r may strike a
 2997 better balance between conservativeness and optimism, enabling the policy to generalize to a broader
 2998 and more plausible spectrum of reward functions. Therefore, in the absence of prior knowledge of
 2999 r , starting with a moderate r is a practical heuristic.

2999 H.3 ADDITIONAL RESULTS FOR ROBUSTNESS ACROSS CORRELATION LEVELS

3000 As discussed previously (Appendix H.2), we do not include linear worst-case evaluation for the
 3001 Glucose and Tomato environments. Consequently, we cannot perform a uniform search over θ as
 3002 we do for the Traffic and Pandemic environments (Appendix F.5). As noted in Appendix B, it is
 3003 generally difficult, and often infeasible, to sample a full reward function over all state-action pairs,
 3004 particularly in high-dimensional or continuous environments, such as the Glucose environment. To
 3005 approximate this process for the Tomato environment, which is a discrete environment, we instead
 3006 sample 1000 trajectories using the reference policy π_{ref} . We then restrict the search to the visited
 3007 state-action pairs. For each such pair, we perturb the original proxy reward by adding Gaussian noise
 3008 with zero mean and variance sampled uniformly from the interval $[0.001, 1]$. We then check whether
 3009 the resulting perturbed reward \tilde{R} satisfies the constraint $\tilde{R} \in \mathcal{R}_{\text{corr}}$. As in previous evaluations, we
 3010 do not explicitly constrain M and V . We sample 20 perturbed reward functions and then use them
 3011 to evaluate each policy. **Note:** Some policies, such as ORPO, may visit state-action pairs that are not
 3012 included in the sampled set from π_{ref} . For these unseen state-action pairs, the proxy-true correlation
 3013 constraint does not apply, as no corresponding reference data is available. In such cases, we default
 3014 to using the original proxy reward to evaluate those portions of the trajectory.

3015 We emphasize that this procedure is neither optimal nor efficient, and is employed solely for eval-
 3016 uation purposes in the Tomato environment. Designing more principled and scalable methods for
 3017 reward sampling under correlation constraints remains an important direction for future work.

3018 Figure 5 presents the average reward and standard deviation across varying correlation levels r for
 3019 the Tomato environment. As expected, the reference policy π_{ref} (blue) consistently underperforms
 3020 across all values of r , though it exhibits the lowest variance—indicating stable yet suboptimal be-
 3021 havior. Interestingly, ORPO* (purple) performs worse than ORPO (red) throughout, suggesting that
 3022 improving the accuracy of occupancy measure estimation does not necessarily enhance robustness
 3023 in this environment. In contrast, our Max-Min method (green) achieves the highest average reward
 across all correlation levels, highlighting its better robustness under reward uncertainty.

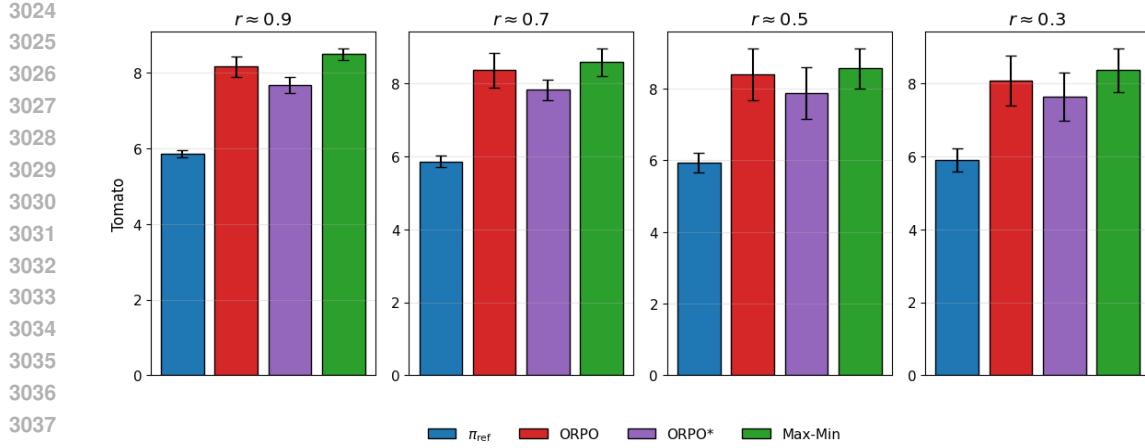


Figure 5: Mean reward and standard deviation under sampled reward functions at different proxy-true reward correlation levels r for the Tomato environment. Our methods (Max-Min) yield higher average performance across all choices of r .

H.4 ADDITIONAL UNNORMALIZED RESULTS

To ensure a fair comparison with prior work (Laidlaw et al., 2025), which reports results in the unnormalized reward scale, we also include the raw (unnormalized) expected proxy and true rewards. However, for worst-case reward metrics, it is nontrivial to reverse the normalization transformation, as our formulation explicitly constrains the reward to have zero mean and unit variance under the reference policy. Therefore, we omit worst-case results in the unnormalized setting.

Table 10: Unnormalized performance comparison across all environments.

| Method | Traffic | | Pandemic | | Glucose | | Tomato | | RLHF | |
|-------------------------|---------------------|--------------------|-------------------|-------------------|------------------------|--------------------|-----------------|-----------------|------------------|------------------|
| | True | Proxy | True | Proxy | True ($\times 10^3$) | Proxy | True | Proxy | True | Proxy |
| π_{ref} | -1004.33 \pm 0.00 | 1474.30 \pm 0.00 | -12.01 \pm 0.00 | -12.01 \pm 0.00 | -79.7 \pm 0.0 | -117.75 \pm 0.00 | 5.96 \pm 0.00 | 6.37 \pm 0.00 | 15.97 \pm 0.00 | -0.29 \pm 0.00 |
| ORPO | -666.13 \pm 2.34 | 1542.57 \pm 2.62 | -12.84 \pm 0.17 | -10.61 \pm 0.15 | -49.7 \pm 0.6 | -67.51 \pm 0.27 | 9.10 \pm 0.11 | 9.10 \pm 0.11 | 16.51 \pm 0.07 | -0.23 \pm 0.02 |
| ORPO* | -799.02 \pm 1.77 | 1501.29 \pm 1.83 | -11.06 \pm 0.15 | -11.06 \pm 0.15 | -48.3 \pm 1.0 | -59.57 \pm 0.28 | 7.96 \pm 0.09 | 7.96 \pm 0.09 | N/A | N/A |
| Max-Min | -750.32 \pm 1.28 | 1546.86 \pm 1.82 | -11.01 \pm 0.14 | -11.01 \pm 0.14 | -48.2 \pm 0.5 | -66.42 \pm 0.29 | 8.24 \pm 0.10 | 8.24 \pm 0.10 | 16.32 \pm 0.06 | -0.21 \pm 0.01 |
| Linear Max-Min Ensemble | -675.12 \pm 2.04 | 1522.34 \pm 2.19 | -9.90 \pm 0.09 | -5.93 \pm 0.10 | N/A | N/A | N/A | N/A | N/A | N/A |

Table 10 presents the unnormalized performance results across all environments. We observe that both our Max-Min and Linear Max-Min policies achieve comparable performance to ORPO on most tasks. Interestingly, the ORPO* variant (with a fully trained discriminator) outperforms the original ORPO in some environments (e.g., Pandemic and Glucose), but performs worse in others, such as Traffic and Tomato. While our earlier analysis (Section 4.2) shows that better discriminator training generally improves worst-case robustness, these results suggest that accurate discriminator estimation does not always translate to improved performance for every specific reward function. Understanding the nuanced effects of discriminator optimization on various reward metrics is beyond the scope of this paper and remains an important direction for future research.

H.5 RESULTS FOR ALL r

Here, we report the results of a uniform grid search over $r \in [0.1, 0.9]$ for our Max-Min algorithm across all training-time correlation levels r on the Traffic, Tomato, Pandemic and Glucose environments in Tables 11, 12, 13 and 14. Each table presents the mean and standard deviation of the expected true reward, expected worst-case reward, and occupancy measure achieved by the learned policy over five random seeds. Here, the “true reward” refers to the original (unnormalized) true reward, while the “worst reward” refers to the normalized reward. We also include the corresponding ORPO results for a fair comparison.

Note that the original ORPO algorithm performs a grid search over λ , where $\lambda = \sigma_{R_{\text{proxy}}} \sqrt{1 - r^2}$ and $\sigma_{R_{\text{proxy}}}$ is the standard deviation of R_{proxy} under π_{ref} , which is generally unknown and must be

estimated. Thus, ORPO effectively searches over λ , whereas our method searches directly over r . For a fair comparison, we estimate $\sigma_{R_{\text{proxy}}}$ and use it to map each value of r in our grid to a corresponding λ for ORPO. However, across all environments we find that the resulting λ values occupy a much narrower scale than r : while r spans the full range from 0.1 to 0.9, the induced λ values are confined to a small interval (e.g., approximately 0.021–0.05 in **Tomato** and 0.035–0.08 in **Pandemic**). As a consequence, the ORPO policies change only marginally across the mapped r values, and their expected true and worst-case returns appear similar in the tables. In practice, ORPO would need to search over a broader range of λ values. By contrast, the expected true and worst-case returns for our Max-Min method vary meaningfully across the full span of r . This highlights that, in practice, our method and ORPO naturally operate on different hyperparameter scales when tuning their respective robustness parameters.

Table 11: Evaluation in the Tomato environment across training-time correlation levels r for ORPO and Max–Min. λ denotes ORPO’s coefficient, with $\lambda = \sigma_{R_{\text{proxy}}} \sqrt{1 - r^2}$. We use $\sigma_{R_{\text{proxy}}} = 0.05$ in this environment, consistent with the ORPO setting. **Occ** denotes the total occupancy over state-action pairs unseen by π_{ref} . **Worst** refers to the expected worst-case reward computed under the training r while excluding those unseen state-action pairs.

| r | ORPO | | | Max–Min | | | |
|-----|-----------|-----------|------------|-------------------|-----------|------------|-------------------|
| | λ | True | Worst | Occ | True | Worst | Occ |
| 0.1 | 0.050 | 0.46±0.14 | -6.08±0.07 | 8.22e-03±0.27e-03 | 0.13±0.15 | -1.34±0.05 | 1.36e-03±0.12e-03 |
| 0.2 | 0.049 | 0.66±0.03 | -6.88±0.08 | 2.8e-03±0.14e-03 | 7.79±0.10 | -0.72±0.05 | 2.18e-03±0.25e-03 |
| 0.3 | 0.048 | 0.70±0.03 | -7.55±0.04 | 1.87e-03±0.16e-03 | 7.68±0.11 | -0.96±0.04 | 1.85e-03±0.16e-03 |
| 0.4 | 0.046 | 0.16±0.02 | -9.20±0.06 | 1.51e-03±0.13e-03 | 8.24±0.10 | -1.37±0.06 | 1.01e-05±0.43e-05 |
| 0.5 | 0.043 | 0.52±0.08 | -6.53±0.05 | 0.028±0.0010 | 7.38±0.12 | -1.21±0.04 | 2.15e-03±0.37e-03 |
| 0.6 | 0.040 | 0.51±0.07 | -7.44±0.06 | 0.028±0.0011 | 6.65±0.17 | -1.22±0.06 | 1.18e-03±0.32e-03 |
| 0.7 | 0.035 | 0.84±0.12 | -5.85±0.07 | 0.027±0.0010 | 0.16±0.08 | -1.33±0.06 | 1.05e-02±0.10e-02 |
| 0.8 | 0.030 | 0.16±0.03 | -7.01±0.06 | 2.88e-03±0.19e-03 | 1.02±0.13 | -2.86±0.04 | 5.77e-04±0.35e-04 |
| 0.9 | 0.021 | 0.11±0.09 | -7.30±0.08 | 0.032±0.0009 | 0.37±0.13 | -5.49±0.12 | 1.29e-05±0.41e-05 |

Table 12: Evaluation in the Traffic environment across training-time correlation levels r for ORPO and Max–Min. λ denotes ORPO’s coefficient, with $\lambda = \sigma_{R_{\text{proxy}}} \sqrt{1 - r^2}$. We use $\sigma_{R_{\text{proxy}}} = 2e - 4$ in this environment, consistent with the ORPO setting. **Occ** denotes the total occupancy over state-action pairs unseen by π_{ref} . **Worst** refers to the expected worst-case reward computed under the training r while excluding those unseen state-action pairs.

| r | ORPO | | | Max–Min | | | |
|-----|-----------|---------------|--------------------|-------------------|---------------|--------------------|-------------------|
| | λ | True | Worst | Occ | True | Worst | Occ |
| 0.1 | 1.99e-4 | -1063.75±1.08 | 1.20e+03±0.02e+03 | 5.53e-03±0.03e-03 | -1428.21±5.36 | -3.43e+04±0.66e+04 | 2.29e-04±0.36e-04 |
| 0.2 | 1.95e-4 | -775.79±1.15 | 4.62e+04±0.04e+04 | 8.71e-04±0.03e-04 | -1312.67±9.14 | -2.83e+04±0.71e+04 | 1.49e-04±0.42e-04 |
| 0.3 | 1.91e-4 | -689.31±1.12 | -5.13e+04±0.04e+04 | 3.98e-04±0.02e-04 | -750.32±1.28 | -268.31±4.14 | 0.00±0.00 |
| 0.4 | 1.83e-4 | -1109.41±1.59 | -1.84e+03±0.04e+03 | 4.94e-03±0.05e-03 | -732.86±1.19 | -314.14±5.37 | 0.00±0.00 |
| 0.5 | 1.73e-4 | -673.45±1.36 | -1.51e+04±0.03e+04 | 6.85e-04±0.01e-04 | -1034.42±2.32 | -6168.40±124.75 | 0.00±0.00 |
| 0.6 | 1.60e-4 | -768.04±1.55 | -4.65e+04±0.04e+04 | 5.43e-03±0.02e-03 | -1322.74±3.92 | -6.73e-03±0.28e+03 | 4.50e-05±1.27e-05 |
| 0.7 | 1.43e-4 | -816.62±1.03 | -4.93e+04±0.04e+04 | 7.60e-03±0.03e-03 | -1398.63±2.73 | -3.82e+04±1.74e+04 | 2.38e-04±0.35e-04 |
| 0.8 | 1.20e-4 | -782.01±1.08 | -8.74e+04±0.09e+04 | 5.97e-03±0.02e-03 | -1359.23±2.08 | -4.94e+04±1.45e+04 | 3.48e-04±0.25e-04 |
| 0.9 | 8.72e-5 | -669.89±1.01 | -1.39e+04±0.04e+04 | 4.41e-03±0.03e-03 | -1337.41±2.45 | -9.33e+03±1.88e+02 | 9.66e-05±1.84e-05 |

I HOW TO CHOOSE r IN PRACTICE?

When r is unknown, both our method and ORPO lack a principled mechanism for selecting an appropriate value. Besides the simple heuristics derived from our experiments as discussed in Appendix H.2, we outline two potential approaches to this important problem below.

Statistical inference of r . If we have access to the true reward on a subset of state-action pairs, or if such labels can be acquired through active learning, we can estimate r using the definition:

$$\mathbb{E}_{\mu_{\pi_{\text{ref}}}} \left[\left(\frac{R_{\text{proxy}} - J(\pi_{\text{ref}}, R_{\text{proxy}})}{\sigma_{R_{\text{proxy}}}} \right) \left(\frac{R_{\text{true}} - J(\pi_{\text{ref}}, R_{\text{true}})}{\sigma_{R_{\text{true}}}} \right) \right] = r, \quad (64)$$

Table 13: Evaluation in the Pandemic environment across training-time correlation levels r for ORPO and Max–Min. λ denotes ORPO’s coefficient, with $\lambda = \sigma_{R_{\text{proxy}}} \sqrt{1 - r^2}$. We use $\sigma_{R_{\text{proxy}}} = 0.08$ in this environment, consistent with the ORPO setting. **Worst** refers to the expected worst-case reward computed under the training r while excluding those unseen state-action pairs.

| r | ORPO | | | Max–Min | |
|-----|-----------|-------------|--------------------|-------------|--------------------|
| | λ | True | Worst | True | Worst |
| 0.1 | 0.080 | -12.22±0.14 | -7.29e+06±0.05e+06 | -18.48±0.37 | -7.44e+04±0.19e+04 |
| 0.2 | 0.078 | -11.77±0.11 | -1.70e+07±0.10e+07 | -16.31±0.49 | -6.03e+04±0.12e+04 |
| 0.3 | 0.076 | -12.49±0.20 | -2.49e+06±0.05e+06 | -19.27±0.38 | -7.04e+04±0.05e+04 |
| 0.4 | 0.073 | -12.17±0.17 | -1.25e+06±0.05e+06 | -19.21±0.15 | -2.03e+03±0.16e+03 |
| 0.5 | 0.069 | -12.26±0.24 | -1.07e+06±0.04e+06 | -13.75±0.14 | -2.58e+03±0.15e+03 |
| 0.6 | 0.064 | -12.08±0.28 | -2.65e+06±0.09e+06 | -13.15±0.24 | -104.00±0.22 |
| 0.7 | 0.057 | -11.45±0.23 | -2.92e+05±0.10e+05 | -11.01±0.14 | -63.29±3.35 |
| 0.8 | 0.048 | -12.22±0.14 | -9.37e+05±0.06e+05 | -11.20±0.22 | -123.65±0.15 |
| 0.9 | 0.035 | -12.02±0.20 | -3.29e+04±0.09e+04 | -11.05±0.13 | -77.20±2.08 |

Table 14: Evaluation in the Glucose environment across training-time correlation levels r for ORPO and Max–Min. λ denotes ORPO’s coefficient, with $\lambda = \sigma_{R_{\text{proxy}}} \sqrt{1 - r^2}$. We use $\sigma_{R_{\text{proxy}}} = 0.05$ in this environment, consistent with the ORPO setting. **Worst** refers to the expected worst-case reward computed under the training r while excluding those unseen state-action pairs.

| r | ORPO | | | Max–Min | |
|-----|-----------|-----------------------|--------------|-----------------------|--------------|
| | λ | True($\times 10^3$) | Worst | True($\times 10^3$) | Worst |
| 0.1 | 0.050 | -90.2±0.7 | -350.94±0.37 | -169.3±0.6 | -317.97±0.12 |
| 0.2 | 0.049 | -88.1±0.8 | -199.26±0.59 | -150.2±0.6 | -304.15±0.34 |
| 0.3 | 0.048 | -79.1±0.6 | -225.71±0.45 | -118.3±0.4 | -139.47±0.46 |
| 0.4 | 0.046 | -72.2±0.4 | -206.40±0.27 | -113.5±0.8 | -123.11±0.32 |
| 0.5 | 0.043 | -94.4±0.4 | -215.00±0.27 | -125.1±0.7 | -126.41±0.43 |
| 0.6 | 0.040 | -68.0±0.9 | -266.50±0.47 | -95.9±0.8 | -84.67±0.17 |
| 0.7 | 0.035 | -71.6±0.5 | -314.48±0.23 | -51.7±0.8 | -18.84±0.43 |
| 0.8 | 0.030 | -53.5±0.6 | -227.79±0.28 | -33.3±0.4 | -11.25±0.27 |
| 0.9 | 0.021 | -50.9±0.5 | -255.07±0.18 | -48.2±0.5 | -1.71±0.25 |

In fact, Equation 64 defines the Pearson correlation coefficient r between the true reward R_{true} and the proxy reward R_{proxy} under the occupancy measure $\mu_{\pi_{\text{ref}}}$. Given a batch of n state-action pairs $\{(s_i, a_i)\}_{i=1}^n$ sampled from π_{ref} for which we have both $R_{\text{true}}^{(i)}$ and $R_{\text{proxy}}^{(i)}$, we can estimate this correlation using the sample correlation coefficient:

$$\hat{r} = \frac{\sum_{i=1}^n (R_{\text{true}}^{(i)} - \bar{R}_{\text{true}})(R_{\text{proxy}}^{(i)} - \bar{R}_{\text{proxy}})}{\sqrt{\sum_{i=1}^n (R_{\text{true}}^{(i)} - \bar{R}_{\text{true}})^2} \cdot \sqrt{\sum_{i=1}^n (R_{\text{proxy}}^{(i)} - \bar{R}_{\text{proxy}})^2}}$$

We can then use **Fisher’s z-transformation** to compute the confidence intervals for r . After getting this bounded range, we can plug this bound into our framework to define a tighter reward uncertainty set. For example, we can use r_{lower} for more pessimistic robustness. Or we can redefine the correlation constraint in Equation 64 to be bounded by both r_{lower} and r_{upper} . The optimal solution under this new constraint can be similarly obtained using the approach in the paper.

A min–max regret approach. A more principled approach to addressing the uncertainty in r may come from a regret-based perspective. Let $J_r(\pi)$ denote the worst-case return for a given policy π under a specific correlation level r , i.e., $J_r(\pi) = \min_{R \in R_{\text{corr}}(r)} J(\pi, R)$. The regret can then be defined as $\text{Reg}(\pi, r) = \max_{\pi^*} J_r(\pi^*) - J_r(\pi)$, which quantifies the performance gap between the optimal policy under r and the current policy. With this formulation, a robust objective

3186 can be expressed as $\min_{\pi} \max_r \text{Reg}(\pi, r)$, aiming to find a policy that minimizes the worst-case
3187 regret across all possible values of r . This framework enables us to train policies that are robust to
3188 uncertainty in the correlation parameter r . We think this is a promising future direction, especially
3189 for cases where r may be misspecified during training. As studied in (Sadek et al., 2025), minimax-
3190 regret may provide strong robustness guarantees under distribution shifts for r . In such settings,
3191 methods like Prioritized Level Replay (Jiang et al., 2021) and recent progress in (Monette et al.,
3192 2025) could be adapted to solve the problem by sampling multiple r and solving Equation 27 in
3193 our paper. We should note that the reason these frameworks are potentially applicable is that our
3194 formulation admits a closed-form solution for the inner minimization. However, the main challenge
3195 lies in estimating the occupancy measure. An interesting direction for future work is to investigate
3196 whether policy gradients can be approximated without explicitly occupancy estimation.
3197
3198
3199
3200
3201
3202
3203
3204
3205
3206
3207
3208
3209
3210
3211
3212
3213
3214
3215
3216
3217
3218
3219
3220
3221
3222
3223
3224
3225
3226
3227
3228
3229
3230
3231
3232
3233
3234
3235
3236
3237
3238
3239