

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 BEST-OF-N THROUGH THE SMOOTHING LENS: KL DIVERGENCE AND REGRET ANALYSIS

Anonymous authors

Paper under double-blind review

ABSTRACT

A simple yet effective method for inference-time alignment of generative models is Best-of- N (BoN), where N outcomes are sampled from a reference policy, evaluated using a proxy-reward model, and the highest-scoring one is selected. While prior work argues that BoN is almost optimal in reward vs KL tradeoffs, the effectiveness of BoN depends critically on the quality of the proxy-reward model used for selection. For this purpose, we study BoN through a smooth version known as Soft Best-of- N (SBoN) and develop a theoretical framework to address this gap. We analyze the scaling behaviour of BoN by providing bounds on the KL divergence between the SBoN policy and the reference policy, offering insights into how performance varies with the number of samples. We also study the regret gap, i.e., the gap between the expected true-reward under the optimal policy and the SBoN policy. Our theoretical and empirical findings show that smoothing helps SBoN mitigate reward overoptimization, especially when the quality of the proxy-reward is low.

1 INTRODUCTION

Large language models (LLMs) have transformed machine learning, achieving state-of-the-art results on a variety of tasks. Despite all advancements, LLMs can still generate undesirable outputs, such as toxic or factually incorrect responses. This has made alignment a central goal in modern LLM development (Achiam et al., 2023; Team et al., 2023).

Several post-hoc alignment methods have been proposed to address this challenge, including Reinforcement Learning from Human Feedback (RLHF) (Christiano et al., 2017; Ouyang et al., 2022), SLiC (Zhao et al., 2022), Direct Preference Optimization (Rafailov et al., 2023), controlled decoding (Mudgal et al., 2024) and Best-of- N (BoN) sampling (Stiennon et al., 2020). While these methods differ in their implementation which are ranging from training-time optimization to test-time selection, they can be viewed as approximating the solution to a KL-regularized reward maximization problem. The optimal solution to this problem is a tilted distribution over responses, which balances reward and proximity to the reference model (Yang et al., 2024).

In BoN as a test-time sampling strategy, given a prompt, N responses are sampled from the reference policy, and the one with the highest proxy-reward sample is selected. Empirically, BoN has been shown to achieve competitive or superior performance in the reward-versus-KL divergence trade-off when compared to RLHF and other alignment methods (Gao et al., 2023a; Mudgal et al., 2024) under *true-reward model*. Furthermore, under certain conditions, it asymptotically approximates the solution to the KL-regularized reward maximization objective (Yang et al., 2024). However, in practice, BoN relies on a learned proxy-reward model which is an approximation of the true-reward function, to guide this selection. As such, their effectiveness critically depends on both the *proxy-reward model* (estimation error of true-reward) and the *quality of the reference policy*.

Understanding how these two components, the quality of the proxy-reward model and the choice of reference policy, affect the alignment quality of test-time sampling algorithms is essential. There are different measures of alignment quality, including KL divergence ¹ between aligned policy and reference policy and the *regret* defined as the gap between the expected true-reward under the optimal

¹Unless stated otherwise, all KL divergences are understood to be measured between the aligned policy and the reference policy.

(tilted) policy and the alignment policy. Note that minimizing the regret gap is critical to ensuring high-quality outputs and close performance to the optimal policy. Recent work by Gao et al. (2023a) and Hilton et al. (2022) has investigated the scaling laws governing reward model optimization in both reinforcement learning (RL) and BoN settings as a function of KL divergence between aligned policy and reference policy. They empirically demonstrate that, under proxy-reward models, the improvement in expected true-reward, relative to a reference policy, scales proportionally for both RL and BoN policies.

While recent work analyzes BoN under the idealized settings where there is no discrepancy between the proxy-reward and the true-reward (Yang et al., 2024; Beirami et al., 2024; Mroueh, 2024; Huang et al., 2025), our work relaxes this assumption to study the interplay between the reward discrepancy measured through regret and the KL-divergence. We present a theoretical study of **Soft Best-of- N** (SBoN), a smoothed variant of BoN recently introduced by Mayrink Verdun et al. (2025); Jinnai et al. (2024). Unlike BoN, SBoN draws the final response *probabilistically* from the N candidates, yielding a policy that is tunable with a temperature parameter. Our analysis centres on two metrics:

- (a) the Kullback-Leibler divergence between SBoN policy (under the *true* reward or proxy-reward model) and the reference policy, and
- (b) the *regret*, i.e. the expected true-reward gap between optimal policy and SBoN policy.

We show how these results specialize to the BoN (as a limit of SBoN for the temperature goes to infinity) and quantify the estimation error incurred by using a *proxy* reward model instead of the true-reward. Finally, we characterize regimes in which SBoN attains *lower* regret bound than BoN when we use the proxy-reward model. Our main contributions are:

- We derive finite-sample upper bounds for KL divergence between the SBoN policy and reference policy, and upper and lower bounds for the regret gap of the SBoN policy, and we extend these bounds to BoN. These bounds reveal how the number of responses N , proxy-reward model quality and reference policy model affect performance.
- We quantify cases where SBoN performs better than BoN under overoptimization scenario where the proxy-reward model is used instead of the true-reward model.
- We provide experimental validation using various proxy-reward models to demonstrate SBoN’s advantages in the overoptimization scenario. Furthermore, we provide numerical experiments to evaluate our bounds.

2 RELATED WORKS

In this section, we discuss related works on BoN, the theoretical foundation of (Soft) BoN and overoptimization. More related works for the theoretical foundation of RLHF and smoothing of maximum are provided in the Appendix (App B).

Best-of- N : Despite many recent advancements in alignment, a simple, popular, and well-performing method continues to be the BoN policy (Nakano et al., 2021; Stiennon et al., 2020; Beirami et al., 2024). In fact, Gao et al. (2023b); Mudgal et al. (2024); Eisenstein et al. show that BoN consistently achieves compelling win rate-KL tradeoff curves, often outperforming KL-regularized reinforcement learning and other more complex alignment strategies. LLaMA 2 (Touvron et al., 2023) leverages BoN outputs as teacher signals to further finetune the base model. Mudgal et al. (2024) extend BoN through Q-learning to block-wise BoN decoding. This empirical effectiveness has also inspired research into distilling BoN behaviour into standalone models (Amini et al., 2025; Sessa et al., 2024; Qiu et al., 2024). Hughes et al. (2024) utilize BoN as an effective method for jailbreaking, while BoN is also commonly used as a strong baseline for scaling inference-time compute (Brown et al., 2024; Snell et al., 2024). Given the broad success of BoN, we are motivated to theoretically investigate the BoN policies and the effect of the proxy-reward model (reward hacking) and the quality of the reference policy.

Theoretical Foundation of (Soft) BoN: KL divergence of BoN is studied in (Beirami et al., 2024; Mroueh, 2024) via information theoretical tools where the KL divergence of BoN sampling from the reference distribution is bounded by $\log(N) - (N - 1)/N$. Scaling laws governing reward as a function of KL divergence is empirically studied by Gao et al. (2023b) and theoretically formalized by Mroueh (2024). Furthermore, the asymptotic case and the equivalence of BoN to the KL-constrained reinforcement learning solution are studied by Yang et al. (2024) under the assumption that the BoN policy is a smooth approximation of the optimal policy.

108 tion of access to optimal reward. Gui et al. (2024) further characterized the win rate–KL gap in
 109 the asymptotic regime where a model assigns extremely low likelihoods to successful completions.
 110 Furthermore, Sun et al. (2024) accelerated BoN using speculative rejection sampling. The regret
 111 of BoN under some assumptions is studied in (Huang et al., 2025). The convergence rate of the
 112 SBoN policy to the optimal tilted policy has been analyzed by Mayrink Verdun et al. (2025). Ad-
 113 ditionally, Geuter et al. (2025) investigate a variant of SBoN that incorporates speculative samples
 114 from a small auxiliary model, providing both theoretical and empirical insights. However, the regret
 115 gap and KL divergence of SBoN under overoptimization scenario remain largely unexplored in the
 116 existing literature.

117

118 3 PROBLEM FORMULATION

119

120 **Notations:** Upper-case letters denote random variables (e.g., Z), lower-case letters denote the re-
 121 alizations of random variables (e.g., z), and calligraphic letters denote sets (e.g., \mathcal{Z}). All loga-
 122 rithms are in the natural base. The set of probability distributions (measures) over a space \mathcal{X} with
 123 finite variance is denoted by $\mathcal{P}(\mathcal{X})$. Δ_N is N -simplex distribution set. The KL divergence be-
 124 tween two probability distributions on \mathbb{R}^d with densities $p(x)$ and $q(x)$, such that $q(x) > 0$ when
 125 $p(x) > 0$, is $\text{KL}(p\|q) := \int_{\mathbb{R}^d} p(x) \log(p(x)/q(x))dx$ (with $0 \cdot \log 0 := 0$). The total-variation dis-
 126 tance is defined as $\text{TV}(p, q) = \frac{1}{2} \int_{\mathcal{X}} |p(x) - q(x)|$. Furthermore, we define chi-square divergence
 127 as $\chi^2(p(x)\|q(x)) = \int_{\mathcal{X}} \frac{p^2(x)}{q(x)} - 1$.

128

129 **Preliminaries:** Let the finite set² of prompts be \mathcal{X} and the discrete finite set of responses be \mathcal{Y} .
 130 Prompts are drawn from a distribution ρ over \mathcal{X} . A (stochastic) policy $\pi \in \Pi$ assigns, for every
 131 prompt $x \in \mathcal{X}$, a conditional distribution $\pi(\cdot | x)$ over \mathcal{Y} ; drawing $y \sim \pi(\cdot | x)$ yields a response. We
 132 treat the supervised-fine-tuned (SFT) model as a *reference policy*, denoted $\pi_{\text{ref}}(\cdot | x)$.

133

134

3.1 REWARD FUNCTION

135

136 We consider a (calibrated) true-reward function $r^*(y, x)$ and a (calibrated) proxy-reward function
 137 $\hat{r}(y, x)$, both mapping $\mathcal{Y} \times \mathcal{X} \rightarrow [0, 1]$ ³. The true-reward function is estimated via proxy-reward
 138 function $\hat{r}(y, x)$ using some preference datasets⁴. As shown by Balashankar et al. (2025), a cali-
 139 brated reward function satisfies:

140

- 141 • Boundedness: for all x, y , we have $\hat{r}(y, x), r^*(y, x) \in [0, 1]$.
- 142 • Uniformity under the reference model: for each prompt $x \in \mathcal{X}$, if $Y \sim \pi_{\text{ref}}(\cdot | x)$ then
 143 $r(y, x) \sim \text{Unif}(0, 1)$ for $r \in \{\hat{r}, r^*\}$.

144

145 In practice, the proxy-reward model can be fit to a human-labeled preference dataset or to data
 146 annotated with true-rewards. Following (Huang et al., 2025), we assume for simplicity that $\hat{r}(y, x)$
 147 is given. We define $\mathcal{Y}_{r^*}^*(x) = \arg \max_y r^*(y, x)$ as the set of maximizers for true-reward. Similarly,
 148 we can define $\hat{\mathcal{Y}}_{\hat{r}}(x)$ as the set of maximizers for proxy-reward model. More discussion regarding
 149 reward is provided in App. C.

150

151

152 **Assumption 3.1** (Achievable maximum reward). *We assume that for $r \in \{\hat{r}, r^*\}$, we have
 153 $r(\hat{y}(x), x) = 1$ for all $\hat{y}(x) \in \arg \max_y r(y, x)$ and given $x \in \mathcal{X}$.*

154

155

156 In many settings, the reward function attains its maximum at specific responses. In particular, since
 157 a large language model (LLM) generates outputs using a finite vocabulary and a bounded number of
 158 tokens, the space of possible generations is finite, and thus the assumption holds trivially.

159

²For measure-theoretic simplicity and notational convenience, we assume finiteness for the set of prompts.
 Our results also hold for non-finite set.

160

³For the remainder of this paper, we assume all reward functions are calibrated (see App. C). For brevity,
 we refer to them simply as *reward functions*.

161

⁴In some cases, the reward model is not derived from human preference data. Instead, it is either determin-
 istic (e.g., code execution scores) or provided by an automated classifier (e.g., for toxicity or sentiment).

162 3.2 SBoN ALGORITHM
163

164 Fix a prompt $x \in \mathcal{X}$ and draw N i.i.d. candidates $Y_{1:N} \sim \pi_{\text{ref}}(\cdot | x)$. Let $Z \in \{1, \dots, N\}$ denote
165 the index of the selected response with distribution P_Z ; write $P_Z(i) = \Pr(Z = i)$. We seek a
166 distribution over indices that maximizes the proxy-reward:

$$168 \max_{P_Z \in \Delta_N} \mathbb{E}_Z [\hat{r}(Y_Z, x)].$$

170 Without regularization, the optimizer is the deterministic Best-of- N (BoN) rule $P_Z = \delta_{i^*}$ with
171 $i^* \in \arg \max_i \hat{r}(Y_i, x)$. Because \hat{r} is a *proxy* for the true-reward, this deterministic choice can
172 overoptimize the proxy-reward and get response with lower true-reward. To smooth this, we add an
173 entropy penalty with temperature $\beta > 0$:

$$175 \max_{P_Z \in \Delta_N} \mathbb{E}_Z [\hat{r}(Y_Z, x)] + \frac{1}{\beta} H(P_Z).$$

177 The unique solution is the softmax distribution, $P_Z(i) = \frac{\exp(\beta \hat{r}(Y_i, x))}{\sum_{j=1}^N \exp(\beta \hat{r}(Y_j, x))}$.

178 We then sample Z from this distribution and return Y_Z . We refer to this sampling rule as **Soft-BoN**,
179 as introduced by [Mayrink Verdun et al. \(2025\)](#). Note that the value of β parameter for SBoN is equal
180 to β in tilted optimal policy. It is shown by [Mayrink Verdun et al. \(2025\)](#), the SBoN policy can be
181 interpreted as a finite-sample approximation of the tilted optimal policy.

182 We denote the final policy from SBoN via $\pi_{\hat{r}}^{\text{SBoN}}(y|x)$. Note that for $\beta \rightarrow \infty$ and $\beta \rightarrow -\infty$, we
183 recover BoN and worst-of- N (WoN) ([Balashankar et al., 2025](#)), respectively. Furthermore, for $\beta \rightarrow 0$, we recover uniform sampling among the N response samples, which is equivalent to sampling
184 from the reference model $\pi_{\text{ref}}(y|x)$. In ([Mayrink Verdun et al., 2025](#), Lemma 1), the closed form
185 solution of SBoN policy is derived,

$$189 \pi_{\hat{r}}^{\text{SBoN}}(y|x) = \frac{\pi_{\text{ref}}(y|x) \exp(\beta \hat{r}(y, x))}{Z_{N,\beta}}, \quad (1)$$

190 where $Z_{N,\beta} = \mathbb{E} \left[\left(\frac{1}{N} \left(\exp(\beta \hat{r}(y, x)) + \sum_{i=1}^{N-1} \exp(\beta \hat{r}(Y_i, x)) \right) \right)^{-1} \right]^{-1}$. Similarly, we can define
191 $\pi_{r^*}^{\text{SBoN}}(y|x)$ based on a true-reward model. For simplicity, we define BoN policies under true-reward
192 and proxy-reward models as $\pi_{r^*}^{\text{BoN}}(y|x)$ and $\pi_{\hat{r}}^{\text{BoN}}(y|x)$, respectively. In this work, we focus on
193 $\beta \geq 0$. Another motivation for SBoN based on the Gumbel-Max trick is provided in App. E.

194 3.3 OPTIMAL (TILTED) POLICY
195

196 For a given temperature $\beta > 0$, we seek a policy that remains close to π_{ref} while maximizing
197 expected true-reward, leading to the KL-regularized objective

$$201 \max_{\pi \in \Pi} \mathbb{E}_{Y \sim \pi(\cdot|x)} [r^*(y, x)] - \frac{1}{\beta} \text{KL}(\pi(\cdot | x) \| \pi_{\text{ref}}(\cdot | x)). \quad (2)$$

202 The unique solution is the optimal *tilted* policy ([Korbak et al., 2022b;a](#); [Yang et al., 2024](#))

$$205 \pi_{\beta, r^*}(y|x) = \frac{\pi_{\text{ref}}(y|x) \exp(\beta r^*(y, x))}{Z_{r^*, Y}(x, \beta)}, \quad (3)$$

206 where $Z_{r^*, Y}(x, \beta) = \sum_{y \in \mathcal{Y}} \pi_{\text{ref}}(y|x) \exp(\beta r^*(y, x))$, is the normalizing (*partition*) function.

207 Note that, in practice, we do not have access to the closed form of reference policy
208 $\pi_{\text{ref}}(y|x)$ and $r^*(y, x)$. We can only first estimate the true-reward function via
209 proxy-reward function $\hat{r}(y, x)$ and then sample from $\pi_{\text{ref}}(y|x)$ and compute $\hat{r}(y, x)$ for
210 each individual sample. Finally, we can apply inference time algorithms, e.g., BoN or
211 SBoN ([Mayrink Verdun et al., 2025](#)), where N samples are generated from $\pi_{\text{ref}}(y|x)$ and
212 we choose the sample with the highest proxy-reward function (BoN) or sampled from a
213 distribution (SBoN) using the proxy-reward function. When only a proxy-reward function
214 $\hat{r}(y, x)$ is available, we obtain the analogous partition function $Z_{\hat{r}, Y}(x, \beta)$ and policy
215 $\pi_{\beta, \hat{r}}(\cdot|x)$.

216

217

218 We can also define the optimal policy under the
219 true-reward model as,

$$220 \pi_{r^*}^*(y|x) = \arg \max_{\pi} \mathbb{E}_{Y \sim \pi(\cdot|x)} [r^*(Y, x)]. \\ 221$$

222 Similarly, we can define $\pi_{\hat{r}}^*(y|x)$ as the optimal
223 policy under the proxy-reward model⁵.

224 As the reward functions (true and proxy) are
225 bounded due to calibration, we can interpret opti-
226 mal policies as the limit of tilted optimal poli-
227 cies for $r \in \{\hat{r}, r^*\}$,

$$228 \pi_{\infty,r}(\cdot|x) := \lim_{\beta \rightarrow \infty} \pi_{\beta,r}(\cdot|x) \\ 229$$

230 where $\pi_{\infty,r^*}(\cdot|x)$ and $\pi_{\infty,\hat{r}}(\cdot|x)$ place all their probability mass on the maximizers of $r^*(y, x)$ and
231 $\hat{r}(y, x)$, respectively. The connections between SBoN, BoN, optimal and optimal tilted policies
232 under true or proxy-reward models are shown in Figure 1.

233

234

3.4 TILTED ERROR

235

236 Let's define the tilted error as the tilted average of square estimation error of true-reward function
237 for a given prompt x with parameter β ,⁶ as follows,

$$238 \varepsilon_{\beta,r}(x) := \frac{1}{\beta} \log \left(\mathbb{E}_{Y \sim \pi_{\text{ref}}(y|x)} [e^{\beta(r^*(Y, x) - \hat{r}(Y, x))^2}] \right). \quad (4)$$

239

240

241 A similar definition of estimation error is introduced in (Yang & Wibisono, 2022). When $\beta = 0$,
242 the definition reduces to the mean-squared error, which is also introduced in (Huang et al., 2025).
243 Letting $\beta \rightarrow \infty$ recovers the square of the supremum (infinity) norm ($\|\cdot\|_\infty$) of the estimation error
244 between $r^*(y, x)$ and $\hat{r}(y, x)$. Therefore, the following properties hold for $\varepsilon_{\beta,r}(x)$,

245

- 246 • The tilted error is bounded, i.e., $\varepsilon_{\beta,r}(x) \in [0, 1]$.
- 247 • The tilted average of the estimation error is monotonically increasing in β .
- 248 • $\varepsilon_{\infty,r}(x) := \lim_{\beta \rightarrow \infty} \varepsilon_{\beta,r}(x) = \|r^*(Y, x) - \hat{r}(Y, x)\|_\infty^2$.

249

250

251 We assume that overoptimization regime happens whenever we have $\varepsilon_{\beta,r}(x) > 0$. In more details,
252 $\varepsilon_{\beta,r}(x) > 0$ indicates reward misspecification. As noted in Gao et al. (2023b), misspecification is a
253 necessary precondition for overoptimization, but overoptimization itself refers to the specific regime
254 where the optimization strength is sufficiently high that the degradation due to error outweighs the
improvement in the proxy-reward.

255

256

257 We define tilted error using (proxy and true) reward models rather than raw reward models, because
258 our focus is on how rankings change under the proxy. For example, if the proxy-reward is a strictly
259 increasing transform of the true-reward, the ranking is preserved; the Best-of-N (BoN) policy re-
260 mains optimal and no overoptimization occurs. This behavior cannot be captured when working
261 with the raw (uncalibrated) reward models. Note that in (Huang et al., 2025), raw (uncalibrated)
262 reward models are utilized for error definition.

263

264

3.5 COVERAGE

265

266

267 For a given reward function $r(x, y)$, we define the tilted policy (softmax policy):

$$268 \pi_{\beta,r}(y|x) \propto \pi_{\text{ref}}(y|x) \exp(\beta r(y, x)).$$

269

270

⁵This policy, $\pi_{\hat{r}}^*(y|x)$ maximizes \hat{r}_c , it may be suboptimal or harmful under r^* due to Goodhart's Law Gao et al. (2023b).

271

272

⁶To simplify notation, we adopt the same smoothing parameter β for the error definition as is used in the KL and tilted cases.

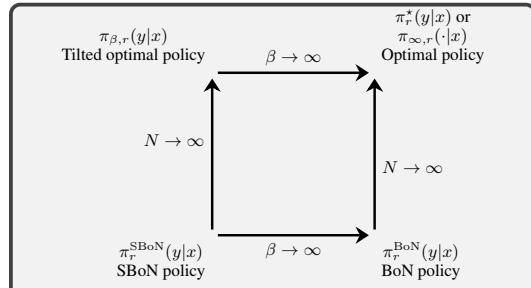


Figure 1: Connections of different policies under reward model $r \in \{\hat{r}, r^*\}$

270 Then, we introduce the coverage of tilted policy with respect to the reference policy as,
 271

$$272 \quad 273 \quad 274 \quad C_{\beta,r}(x) := \sum_{y \in \mathcal{Y}} \frac{\pi_{\beta,r}^2(y|x)}{\pi_{\text{ref}}(y|x)}. \quad (5)$$

275 We also define,

$$276 \quad 277 \quad C_{\infty,\hat{r}}(x) := \lim_{\beta \rightarrow \infty} C_{\beta,r}(x).$$

278 This measure $C_{\beta,r}(x)$ can also be interpreted as a coverage constant, which is standard in
 279 KL-regularized policy learning. Furthermore, we can define the coverage of the tilted pol-
 280 icy with respect to the reference policy as χ^2 -divergence between $\pi_{\beta,r}(y|x)$ and $\pi_{\text{ref}}(y|x)$, i.e.,
 281 $\chi^2(\pi_{\beta,r}(y|x) \parallel \pi_{\text{ref}}(y|x))$. It ensures that the reference policy places sufficient probability mass on
 282 high-reward responses, thereby guaranteeing that the support of the optimal policy lies within the
 283 support of the reference. This prevents cases where optimal outputs are entirely excluded by the
 284 reference. Similar notions of coverage have been explored in Huang et al. (2025).

285 3.6 REGRET

287 For a given policy $\pi(Y|x)$, we define expected true-reward with respect to the policy (a.k.a. value
 288 function⁷) as

$$289 \quad J_{r^*}(\pi(\cdot|x)) := \mathbb{E}_{Y \sim \pi(\cdot|x)}[r^*(Y, x)]. \quad (6)$$

290 For two policies, $\pi_1(\cdot|x)$ and $\pi_2(\cdot|x)$, we define the gap between these two policies as follows,

$$293 \quad \Delta_{J_{r^*}}(\pi_1(\cdot|x), \pi_2(\cdot|x)) := J_{r^*}(\pi_1(\cdot|x)) - J_{r^*}(\pi_2(\cdot|x)). \quad (7)$$

294 We provide an upper bound on the gap of the SBoN solution, which is the gap between $\pi_{r^*}^*(\cdot|x)$ as
 295 the optimal policy and $\pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)$,

$$296 \quad \Delta_{J_{r^*}}(\pi_{r^*}^*(\cdot|x), \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)) = J_{r^*}(\pi_{r^*}^*(\cdot|x)) - J_{r^*}(\pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)). \quad (8)$$

297 Regarding regret of the BoN, we consider $\pi_{\hat{r}}^{\text{BoN}}(\cdot|x)$ instead of $\pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)$ in equation 8.

301 4 KL DIVERGENCE ANALYSIS

302 The KL divergence between the aligned policy and the reference policy, $\text{KL}(\pi_{r^*}^{\text{BoN}} \parallel \pi_{\text{ref}})$, is studied
 303 by Beirami et al. (2024); Mroueh (2024) from a theoretical perspective. In particular, Beirami
 304 et al. (2024) derives an upper bound on KL divergence for BoN policies under the assumptions of a
 305 bijective true-reward mapping and a finite output space:

$$306 \quad \text{KL}(\pi_{r^*}^{\text{BoN}}(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x)) \leq \log(N) - 1 + \frac{1}{N}, \quad (9)$$

307 Mroueh (2024) relaxes the bijectivity assumption and derives similar bounds using information-
 308 theoretic tools. Under some assumptions, the bound in equation 9 is tight. Furthermore, using
 309 Pinsker's inequality, in a similar approach to (Mroueh, 2024), we have,

$$310 \quad \mathbb{E}_{Y \sim \pi_{r^*}^{\text{SBoN}}(\cdot|x)}[r^*(Y, x)] \leq 0.5 + \sqrt{\frac{1}{2} \text{KL}(\pi_{r^*}^{\text{SBoN}}(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x))}. \quad (10)$$

311 Note that equation 10 implies that improvement of expected true-reward relative to the reference
 312 policy can not exceed the square root of the KL divergence. However, the analysis of KL divergence
 313 for the SBoN policy under the true-reward model is overlooked. Therefore, we first establish an
 314 upper bound on the KL divergence between the SBoN policy under the true-reward model and the
 315 reference policy, shedding light on its behaviour as a function of the number of samples N and
 316 temperature parameter β . All proof details are deferred to App. G.

317 318 319 320 321 322 323 ⁷We can also consider $\mathbb{E}_{X \sim \rho(\cdot)}[J_{r^*}(\pi(\cdot|X))]$. All of our results also hold for expected version of value
 324 function.

324 **Lemma 4.1.** *The following upper bound holds on KL divergence between SBoN and reference*
 325 *polices for a given prompt $x \in \mathcal{X}$,*

$$327 \quad \text{KL}(\pi_{r^*}^{\text{SBoN}}(y|x) \parallel \pi_{\text{ref}}(y|x)) \leq \log \left(\frac{N}{1 + (N-1) \exp(-\beta)} \right). \quad (11)$$

329 Using Lemma 4.1, we can observe that for BoN, $\beta \rightarrow \infty$, we have,

$$331 \quad \text{KL}(\pi_{r^*}^{\text{BoN}}(y|x) \parallel \pi_{\text{ref}}(y|x)) \leq \log(N). \quad (12)$$

332 Comparing equation 12 with results in (Beirami et al., 2024; Mroueh, 2024), our result is derived
 333 from the SBoN asymptotic regime. Note that our bound is looser than the bound on KL divergence
 334 in equation 9. In contrast, our bound is general and can be applied to different β in SBoN. For $\beta = 0$,
 335 where our policy is the reference policy, our bound is tight. It is also important to note that the upper
 336 bound in Lemma 4.1 increases with the temperature parameter β for fixed N .

337 Recent works by Gao et al. (2023a) and Hilton et al. (2022) empirically demonstrate that, under a
 338 true-reward model, the improvement in expected true-reward, relative to a reference policy, **scales ap-**
 339 **proximately proportionally to** $\sqrt{\text{KL}(\pi_{r^*}^{\text{BoN}} \parallel \pi_{\text{ref}})}$ for both RL and BoN policies. It is also observed
 340 by Gao et al. (2023b) that models optimized using proxy-rewards can suffer from overoptimization
 341 where the learned policy diverges further from the reference, the alignment may degrade. Despite
 342 theoretical advances, the KL divergence analysis for SBoN and BoN under the proxy-reward model
 343 remains largely unexplored. Therefore, we are interested in investigating the cost we have for esti-
 344 mation error of true-reward via proxy-reward model. For this aim, we first propose the following
 345 useful Lemma to study the closeness of the SBoN policy under the true-reward model to the SBoN
 346 policy under the proxy-reward model in KL divergence measure.

347 **Lemma 4.2.** *The following upper bound holds on the KL divergence between the SBoN policies*
 348 *under true-reward and proxy-reward models respectively,*

$$349 \quad \text{KL}(\pi_{r^*}^{\text{SBoN}}(\cdot|x) \parallel \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)) \leq \frac{N\beta\sqrt{\varepsilon_{\beta,r}(x)}}{1 + (N-1)\exp(-\beta)} \left(\frac{N\exp(2\beta)}{(N-1)^2} + 1 \right). \quad (13)$$

352 Note that for $\beta = 0$, the upper bound in Lemma 4.2 is tight. The result in Lemma 4.2 quantifies the
 353 estimation error introduced by substituting a proxy-reward model for the true-reward model.

354 Next, we compare BoN and SBoN under overoptimization from KL-divergence perspective.

356 **Remark 4.3** (No overoptimization). *We can observe that for a given β , if we assume $\varepsilon_{\beta,r}(x) = 0$,*
 357 *we have $\text{KL}(\pi_{r^*}^{\text{SBoN}}(\cdot|x) \parallel \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)) = 0$. Note that, as mentioned in (Gao et al., 2023b), the*
 358 *expected true-reward under the aligned policy, relative to the reference policy, is proportional to the*
 359 *square root of KL divergence. Then a larger KL divergence is desirable in this context, as proposed*
 360 *by (Gao et al., 2023b), the BoN policy is preferred under no overoptimization scenario.*

361 **Remark 4.4** (Overoptimization). *When $\varepsilon_{\beta,r}(x) > 0$, we have two conflicting goals in both*
 362 *Lemma 4.1 and Lemma 4.2: one suggesting for fixed N that β needs to be smaller for better es-*
 363 *timation of the true policy by the proxy-reward model one given in Lemma 4.2, and another one*
 364 *suggesting a larger β to induce a better KL trade-off based on Gao et al. (2023b). Hence, for a*
 365 *given N , there exists an optimal β to balance between the estimation error of Lemma 4.2 and the*
 366 *scaling law under the SBoN policy for the true-reward model, Lemma 4.1. In this scenario, SBoN*
 367 *can lead to better tradeoffs than BoN. A similar discussion can be done for fixed β and varying N .*
 368 *Further analysis using reverse I-projection (Csiszár & Matus, 2003) in large N regime is provided*
 369 *in App. K.*

370 5 REGRET ANALYSIS

372 In this section, we derive theoretical regret bounds, **upper and lower bounds**, for SBoN and BoN
 373 based on reward models. First, we provide a helpful Lemma regarding the expected coverage as-
 374 sumption that can help us interpret the results of regret for BoN and SBoN. All proof details are
 375 deferred to App. H.

376 **Lemma 5.1.** *Under Assumption 3.1, it holds that $C_{\infty,\hat{r}}(x) = \frac{1}{\sum_i \pi_{\text{ref}}(y_{i,r}^{\max}(x)|x)}$, where $y_{i,r}^{\max}(x) \in$*
 377 *$\arg \max_y r(y, x)$.*

378 5.1 UPPER BOUND ON REGRET
379380 Now, we derive an upper bound on the regret of SBoN.
381382 **Theorem 5.2** (Upper Bound on Regret of SBoN). *Under Assumption 3.1, the following upper bound
383 holds on the optimal regret gap of the SBoN policy for any $\beta > 0$,*
384

385
$$\Delta_{J_{r^*}}(\pi_{r^*}^*(\cdot | x), \pi_{\hat{r}}^{\text{SBoN}}(\cdot | x)) \leq \sqrt{\varepsilon_{\beta, r}(x)}(\sqrt{C_{\infty, \hat{r}}(x)} + \sqrt{C_{\infty, r^*}(x)})$$

386
$$+ 2\sqrt{\frac{1}{2} \log\left(1 + \frac{C_{\infty, \hat{r}}(x) - 1}{N}\right)} + \frac{\log(C_{\infty, r^*}(x))}{\beta}.$$

387

388 **Regret Bound of BoN Through Smoothing Lens:** We now derive an upper bound on the regret of
389 BoN by taking the asymptotic limit of the regret bound on regret of SBoN in Theorem 5.2.
390391 **Proposition 5.3** (Upper Bound on Regret of BoN). *Under Assumption 3.1, the following upper
392 bound holds on the optimal regret gap of the BoN policy for any $\beta > 0$,*
393

394
$$\Delta_{J_{r^*}}(\pi_{r^*}^*(\cdot | x), \pi_{\hat{r}}^{\text{BoN}}(\cdot | x)) \leq \sqrt{\varepsilon_{\infty, r}(x)}(\sqrt{C_{\infty, \hat{r}}(x)} + \sqrt{C_{\infty, r^*}(x)})$$

395
$$+ 2\sqrt{\frac{1}{2} \log\left(1 + \frac{C_{\infty, \hat{r}}(x) - 1}{N}\right)}.$$

396

397 **Remark 5.4** (Comparison with (Huang et al., 2025)). *The regret bound for BoN policy grows with
398 the L_∞ -norm of the reward-model estimation error. In contrast to the result in (Huang et al., 2025),
399 our bound remains finite whenever the overoptimization error vanishes, i.e., when $\varepsilon_{\infty, \beta}(x) = 0$ or
400 N grows. We also derive results based on calibrated reward, instead of raw (uncalibrated) reward
401 models. A full comparison with Huang et al. (2024) is provided in App. J.*
402403 **Overoptimization (asymptotic regime):** Assume that $\varepsilon_{\beta, r}(x) > 0$ for every $\beta > 0$. Letting
404 $N \rightarrow \infty$ and invoking Theorem 5.2, we obtain
405

406
$$\Delta_{J_{r^*}}(\pi_{r^*}^*(\cdot | x), \pi_{\hat{r}}^{\text{SBoN}}(\cdot | x)) \leq \frac{\log C_{\infty, r^*}(x)}{\beta}$$

407
$$+ \sqrt{\varepsilon_{\beta, r}(x)}(\sqrt{C_{\infty, \hat{r}}(x)} + \sqrt{C_{\infty, r^*}(x)}). \quad (14)$$

408

409 Similarly, for BoN we have,
410

411
$$\Delta_{J_{r^*}}(\pi_{r^*}^*(\cdot | x), \pi_{\hat{r}}^{\text{BoN}}(\cdot | x)) \leq \sqrt{\varepsilon_{\infty, r}(x)}(\sqrt{C_{\infty, \hat{r}}(x)} + \sqrt{C_{\infty, r^*}(x)}). \quad (15)$$

412

413 **No overoptimization:** Assume that the overoptimization vanishes, i.e. $\varepsilon_{\beta, r}(x) = 0$ for every $\beta \in$
414 $[0, \infty)$. Then the optimality gaps of the SBoN and BoN policies satisfy
415

416
$$\Delta_{J_{r^*}}(\pi_{r^*}^*(\cdot | x), \pi_{\hat{r}}^{\text{SBoN}}(\cdot | x)) \leq 2\sqrt{\frac{1}{2} \log\left(1 + \frac{C_{\infty, \hat{r}}(x) - 1}{N}\right)} + \frac{\log C_{\infty, r^*}(x)}{\beta}, \quad (16)$$

417

418
$$\Delta_{J_{r^*}}(\pi_{r^*}^*(\cdot | x), \pi_{\hat{r}}^{\text{BoN}}(\cdot | x)) \leq 2\sqrt{\frac{1}{2} \log\left(1 + \frac{C_{\infty, \hat{r}}(x) - 1}{N}\right)}. \quad (17)$$

419

420 5.2 LOWER BOUND ON REGRET
421422 In this section, we complement the regret upper bounds (Theorem 5.2 and Proposition 5.3) with
423 lower bounds that hold for any finite N and fixed $\beta \geq 0$. For the lower bound, the following
424 assumptions are needed.
425426 **Assumption 5.5** (Margin Assumption). *Let $\gamma(x) = 1 - \sup_{y \notin \mathcal{Y}_{r^*}(x)} r^*(y, x)$. We assume that
427 $\gamma(x) \in (0, 1)$.*
428429 Note that Assumption 5.5 for strictly positive lower bound is needed.
430

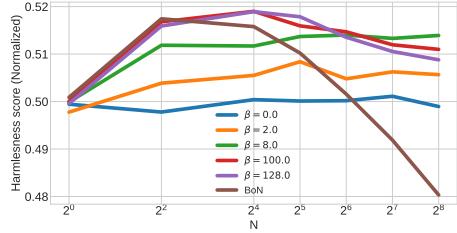
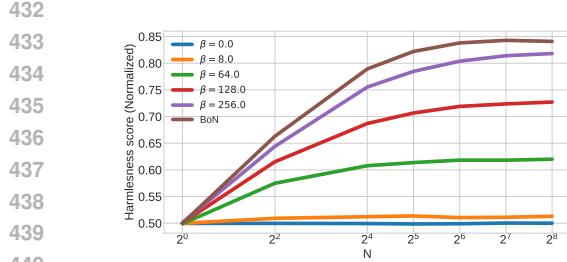


Figure 2: Soft Best-of-N experiment using a strong reward model (Left) and a weak one (Right). Number of samples versus Harmlessness score (higher is better). When the performance of the reward model is poor, BoN can lead to overoptimization, while the SBoN can help to mitigate it.

Theorem 5.6 (Lower Bound on Regret of SBoN). *Under the same Assumptions in Theorem 5.2 and Lemma H.3, the following lower bound holds on the regret of SBoN policy under the proxy-reward model,*

$$\begin{aligned} & \Delta_{J_{r^*}}(\pi_{r^*}^*(\cdot|x), \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)) \\ & \geq \gamma(x) \left((1 - C_{\infty, r^*}(x)^{-1})^N + \frac{N((1 - C_{\infty, r^*}(x)^{-1}) - (1 - C_{\infty, r^*}(x)^{-1})^N)}{N + (\exp(\beta) - 1)(N - 1)} \right) \quad (18) \\ & \quad - \min \left(\sqrt{1 - \exp(-U(N, \beta))}, \frac{N}{1 + (N - 1)\exp(-\beta)} \sqrt{\varepsilon_{\beta, r}(x)} \right) \end{aligned}$$

where $U(N, \beta) = \frac{N\beta\sqrt{\varepsilon_{\beta, r}(x)}}{1 + (N - 1)\exp(-\beta)} \left(\frac{N\exp(2\beta)}{(N - 1)^2} + 1 \right)$.

Using similar approach to Proposition 5.3, we can derive the lower bound on regret of the BoN policy through the lens of smoothing.

Proposition 5.7 (Lower Bound on Regret of BoN). *Under the same assumptions in Theorem 5.6, the following lower bound holds on regret of BoN policy,*

$$\Delta_{J_{r^*}}(\pi_{r^*}^*(\cdot|x), \pi_{\hat{r}}^{\text{BoN}}(\cdot|x)) \geq \gamma(x)(1 - C_{\infty, r^*}(x)^{-1})^N - \min(1, N\sqrt{\varepsilon_{\infty, r}(x)}). \quad (19)$$

Under no overoptimization ($\varepsilon_{\infty, r}(x) = 0$), due to $(1 - C_{\infty, r^*}(x)^{-1})^N \geq \exp(\frac{-N}{C_{\infty, r^*}(x) - 1})$, the lower bound on regret of BoN is $O(\exp(-N))$. Furthermore, in App. I.2, we present a numerical example to illustrate the positivity of the lower bounds under certain conditions.

5.3 DISCUSSION

This section analyzes the performance of BoN and SBoN in the presence and absence of overoptimization.

Remark 5.8 (Overoptimization). *Considering the upper bounds in equation 14 and equation 15, we define,*

$$g(\beta) = \beta(\sqrt{\varepsilon_{\infty, r}(x)} - \sqrt{\varepsilon_{\beta, r}(x)}), \quad \beta \geq 0.$$

Because $g(0) = g(\infty) = 0$ and $g(\beta) \geq 0$ for all β , there exists at least one maximizer $\beta^* \in (0, \infty)$ such that $g(\beta^*) = \max_{\beta \geq 0} g(\beta)$.

If $\frac{\log C_{\infty, r^*, \text{ref}}(x)}{\sqrt{C_{\infty, r^*, \text{ref}}(x)} + \sqrt{C_{\infty, r^*, \text{ref}}(x)}} \leq g(\beta^*)$, then the upper bound in equation 14 does not exceed equation 15, and hence the bound on the regret of the SBoN policy is tighter than the bound on the regret of the BoN policy under the proxy-reward model. An analogous comparison can be carried out for any fixed β and changing N . A numerical example for the existence of optimal β based on this approach is provided in App. I.2.

Remark 5.9 (No overoptimization). *By Lemma 5.1, $C_{\infty, r^*, \text{ref}}(x) \geq 1$; consequently, the bound in equation 17 is tighter than the bound in equation 16. Similar discussion can be provided for lower bound.*

486 **Remark 5.10** (Quality of reference policy). *The upper bounds in Proposition 5.3 (or Theorem 5.2)*
 487 *depend on two quantities,*

$$489 \quad C_{\infty, r^*}(x) = \frac{1}{\sum_i \pi_{\text{ref}}(y_{i, r^*}^{\max}(x) \mid x)}, \quad \text{and} \quad C_{\infty, \hat{r}}(x) = \frac{1}{\sum_i \pi_{\text{ref}}(y_{i, \hat{r}}^{\max}(x) \mid x)},$$

491 *which represent its quality under the true-reward model and proxy-reward model, respectively. A*
 492 *larger value for these quantities implies that the reference model rarely generates optimal responses,*
 493 *thereby degrading performance. We can observe that upper bounds increase by increasing* $C_{\infty, r^*}(x)$
 494 *and* $C_{\infty, \hat{r}}(x)$.

496 6 EMPIRICAL EVIDENCE

499 To support our theoretical analysis, we conducted experiments comparing Soft Best-of-N (SBoN)
 500 across different regularization strengths and reward model qualities. We used the Olmo-2 1B model
 501 (Olmo et al., 2024) as the generator and prompts from the Attaq dataset (Kour et al., 2023). For each
 502 prompt, we generated multiple responses and selected one using SBoN with varying temperature
 503 values β . We ran two experimental conditions: one using a strong proxy-reward model (ArmoRM
 504 8B (Wang et al., 2024)) which is close to true-reward model, and another using a weaker proxy-
 505 reward model (Beaver 7B RM (Dai et al., 2023)). We use LLM-as-a-Judge Zheng et al. (2023) as
 506 our r^* . To match our theoretical setting, we perform empirical calibration of each reward model by
 507 sampling 256 responses for every query and calculate the quantiles. As shown in Figure 2, when
 508 the reward model is weak, performance degrades for large N due to reward hacking. However,
 509 the smoothing in SBoN helps mitigate this degradation. This observation is also aligned with our
 510 theoretical analysis and discussion in Section 4, where under overoptimization there exists a β for a
 511 given N which outperforms BoN. For more details, see App. I. We also studied the behavior of our
 512 upper bound on the KL divergence between the SBoN policy and the reference policy, Lemma 4.2,
 513 at App. I.2. More experiments with a medium weak reward model are provided in App. I.1.

514 7 CONCLUSION AND FUTURE WORK

515 In this work, we establish a theoretical foundation for alignment strategies based on Soft Best-of-N
 516 (SBoN) and Best-of-N (BoN) policies. Specifically, we derive upper bounds on the KL divergence
 517 between the aligned policy such as SBoN or BoN and the reference policy. We also studied the regret
 518 gap between the optimal policy and the aligned policy, e.g., BoN and SBoN policies. We further
 519 analyze how errors in reward estimation affect performance in both KL divergence and regret gap.
 520 Notably, both our theoretical analysis and empirical results demonstrate that, under a proxy-reward
 521 model where overoptimization happens, SBoN perform better than BoN under some conditions.

522 Existing literature Beirami et al. (2024); Mroueh (2024), suggests a tighter upper bound on the
 523 KL divergence for BoN policy. Our derived result, however, provides a looser bound for BoN.
 524 Consequently, an interesting direction for future research is the derivation of an upper bound that is
 525 asymptotically tight for the BoN policy. Furthermore, while our current analysis utilizes a shared β
 526 for both the tilted error and the SBoN policy, exploring decoupled temperature parameters for these
 527 components remains a promising avenue for future study.

528
 529
 530
 531
 532
 533
 534
 535
 536
 537
 538
 539

540 REFERENCES
541

542 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman,
543 Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical
544 report. *arXiv preprint arXiv:2303.08774*, 2023.

545 Afra Amini, Tim Vieira, Elliott Ash, and Ryan Cotterell. Variational best-of-n alignment. In *Inter-*
546 *national Conference on Learning Representations (ICLR)*, 2025.

547 Gholamali Aminian, Amir R Asadi, Idan Shenfeld, and Youssef Mroueh. Theoretical analysis of
548 kl-regularized rlhf with multiple reference models. *arXiv preprint arXiv:2502.01203*, 2025.

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

551 Ananth Balashankar, Ziteng Sun, Jonathan Berant, Jacob Eisenstein, Michael Collins, Adrian Hutter,
552 Jong Lee, Chirag Nagpal, Flavien Prost, Aradhana Sinha, Ananda Theertha Suresh, and Ahmad
553 Beirami. Infalign: Inference-aware language model alignment. *International Conference on*
554 *Machine Learning (ICML)*, 2025.

555 Ahmad Beirami, Alekh Agarwal, Jonathan Berant, Alexander D’Amour, Jacob Eisenstein, Chirag
556 Nagpal, and Ananda Theertha Suresh. Theoretical guarantees on the best-of-n alignment policy.
557 *International Conference on Machine Learning (ICML)*, 2024.

558 Vivek S Borkar. Q-learning for risk-sensitive control. *Mathematics of Operations Research*, 2002.

559 Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V Le, Christopher Ré, and
560 Azalia Mirhoseini. Large language monkeys: Scaling inference compute with repeated sampling.
561 *arXiv preprint arXiv:2407.21787*, 2024.

562 Clément L Canonne. A short note on an inequality between kl and tv. *arXiv preprint*
563 *arXiv:2202.07198*, 2022.

564 Pierre Cardaliaguet, François Delarue, Jean-Michel Lasry, and Pierre-Louis Lions. *The master*
565 *equation and the convergence problem in mean field games:(ams-201)*. Princeton University
566 Press, 2019.

567 Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémie Scheurer, Javier
568 Rando, Rachel Freedman, Tomasz Korbak, David Lindner, Pedro Freire, et al. Open problems
569 and fundamental limitations of reinforcement learning from human feedback. *arXiv preprint*
570 *arXiv:2307.15217*, 2023.

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

574 Thomas Coste, Usman Anwar, Robert Kirk, and David Krueger. Reward model ensembles help
575 mitigate overoptimization. In *The Twelfth International Conference on Learning Representations*.

576 Imre Csiszár and Frantisek Matus. Information projections revisited. *IEEE Transactions on Infor-*
577 *mation Theory*, 49(6):1474–1490, 2003.

578 Josef Dai, Xuehai Pan, Ruiyang Sun, Jiaming Ji, Xinbo Xu, Mickel Liu, Yizhou Wang, and
579 Yaodong Yang. Safe rlhf: Safe reinforcement learning from human feedback. *arXiv preprint*
580 *arXiv:2310.12773*, 2023.

581 Jacob Eisenstein, Chirag Nagpal, Alekh Agarwal, Ahmad Beirami, Alexander Nicholas D’Amour,
582 Krishnamurthy Dj Dvijotham, Adam Fisch, Katherine A Heller, Stephen Robert Pfohl, Deepak
583 Ramachandran, et al. Helping or herding? reward model ensembles mitigate but do not eliminate
584 reward hacking. In *First Conference on Language Modeling*.

585 L Gao, J Tow, B Abbasi, S Biderman, S Black, A DiPofi, C Foster, L Golding, J Hsu,
586 A Le Noach, et al. A framework for few-shot language model evaluation. *URL https://zenodo.*
587 *org/records/10256836*, 7, 2023a.

594 Leo Gao, John Schulman, and Jacob Hilton. Scaling laws for reward model overoptimization. In
 595 *International Conference on Machine Learning*, pp. 10835–10866. PMLR, 2023b.
 596

597 Jonathan Geuter, Youssef Mroueh, and David Alvarez-Melis. Guided speculative inference for effi-
 598 cient test-time alignment of llms. *arXiv preprint arXiv:2506.04118*, 2025.

599 Lin Gui, Cristina Gârbacea, and Victor Veitch. Bonbon alignment for large language models and
 600 the sweetness of best-of-n sampling. *arXiv preprint arXiv:2406.00832*, 2024.

601

602 Emil Julius Gumbel. Statistical theory of extreme values and some practical applications. *Nat. Bur.*
 603 *Standards Appl. Math. Ser. 33*, 1954.

604 J. Hilton, P. Clark, et al. Measuring goodharts law: Towards an evaluation framework for open-ended
 605 generative models. <https://openai.com/index/measuring-goodharts-law>,
 606 2022. Accessed: 2025-01-30.

607

608 Ronald A Howard and James E Matheson. Risk-sensitive markov decision processes. *Management*
 609 *science*, 1972.

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

613

614 Audrey Huang, Adam Block, Qinghua Liu, Nan Jiang, Dylan J Foster, and Akshay Krishnamurthy.
 615 Is best-of-n the best of them? coverage, scaling, and optimality in inference-time alignment. *arXiv*
 616 *preprint arXiv:2503.21878*, 2025.

617 John Hughes, Sara Price, Aengus Lynch, Rylan Schaeffer, Fazl Barez, Sanmi Koyejo, Henry
 618 Sleight, Erik Jones, Ethan Perez, and Mrinank Sharma. Best-of-n jailbreaking. *arXiv preprint*
 619 *arXiv:2412.03556*, 2024.

620

621 Yuki Ichihara, Yuu Jinnai, Tetsuro Morimura, Kenshi Abe, Kaito Ariu, Mitsuki Sakamoto, and Eiji
 622 Uchibe. Evaluation of best-of-n sampling strategies for language model alignment. *Transactions*
 623 *on Machine Learning Research*.

624

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

629

630 Yuu Jinnai, Tetsuro Morimura, Kaito Ariu, and Kenshi Abe. Regularized best-of-n sampling to
 631 mitigate reward hacking for language model alignment. In *ICML 2024 Workshop on Models of*
 632 *Human Feedback for AI Alignment*, 2024.

633

634 Hadi Khalaf, Claudio Mayrink Verdun, Alex Oesterling, Himabindu Lakkaraju, and Flavio
 635 du Pin Calmon. Inference-time reward hacking in large language models. *arXiv preprint*
 636 *arXiv:2506.19248*, 2025.

637

638 Maxim Khanov, Jirayu Burapachetip, and Yixuan Li. Args: Alignment as reward-guided search.
 639 *arXiv preprint arXiv:2402.01694*, 2024.

640

641 Tomasz Korbak, Hady Elsahar, Germán Kruszewski, and Marc Dymetman. On reinforcement learn-
 642 ing and distribution matching for fine-tuning language models with no catastrophic forgetting.
 643 *Advances in Neural Information Processing Systems*, 35:16203–16220, 2022a.

644

645 Tomasz Korbak, Ethan Perez, and Christopher L Buckley. Rl with kl penalties is better viewed as
 646 bayesian inference. *arXiv preprint arXiv:2205.11275*, 2022b.

647

648 Barry W Kort and Dimitri P Bertsekas. A new penalty function method for constrained minimization.
 649 In *IEEE Conference on Decision and Control and 11th Symposium on Adaptive Processes*, 1972.

650

651 George Kour, Marcel Zalmanovici, Naama Zwerdling, Esther Goldbraich, Ora Nova Fandina, Ateret
 652 Anaby-Tavor, Orna Raz, and Eitan Farchi. Unveiling safety vulnerabilities of large language
 653 models. *arXiv preprint arXiv:2311.04124*, 2023.

648 Tian Li, Ahmad Beirami, Maziar Sanjabi, and Virginia Smith. On tilted losses in machine learning:
 649 Theory and applications. *Journal of Machine Learning Research*, 24(142):1–79, 2023.
 650

651 Guan-Horng Liu and Evangelos A Theodorou. Deep learning theory review: An optimal control
 652 and dynamical systems perspective. *arXiv preprint arXiv:1908.10920*, 2019.

653 Claudio Mayrink Verdun, Alex Oesterling, Himabindu Lakkaraju, and Flavio P Calmon. Soft best-
 654 of-n sampling for model alignment. *arXiv preprint arXiv:2505.03156*, 2025.
 655

656 Youssef Mroueh. Information theoretic guarantees for policy alignment in large language models.
 657 *arXiv preprint arXiv:2406.05883*, 2024.

658 Sidharth Mudgal, Jong Lee, Harish Ganapathy, Yaguang Li, Tao Wang, Yanping Huang, Zhifeng
 659 Chen, Heng-Tze Cheng, Michael Collins, Trevor Strohman, et al. Controlled decoding from
 660 language models. In *International Conference on Machine Learning*, pp. 36486–36503. PMLR,
 661 2024.

662 Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher
 663 Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. Webgpt: Browser-assisted
 664 question-answering with human feedback. *arXiv preprint arXiv:2112.09332*, 2021.
 665

666 Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia,
 667 Yuling Gu, Shengyi Huang, Matt Jordan, et al. 2 olmo 2 furious. *arXiv preprint arXiv:2501.00656*,
 668 2024.

669 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
 670 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow
 671 instructions with human feedback. *Advances in neural information processing systems*, 35:27730–
 672 27744, 2022.

673 EY Pee and Johannes O Royset. On solving large-scale finite minimax problems using exponential
 674 smoothing. *Journal of Optimization Theory and Applications*, 2011.

675 Yury Polyanskiy and Yihong Wu. Information theory: From coding to learning, 2022.

676 Jiahao Qiu, Yifu Lu, Yifan Zeng, Jiacheng Guo, Jiayi Geng, Huazheng Wang, Kaixuan Huang, Yue
 677 Wu, and Mengdi Wang. Treebon: Enhancing inference-time alignment with speculative tree-
 678 search and best-of-n sampling. *arXiv preprint arXiv:2410.16033*, 2024.

679 Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D Manning, and Chelsea
 680 Finn. Direct preference optimization: Your language model is secretly a reward model. *arXiv
 681 preprint arXiv:2305.18290*, 2023.

682 Pier Giuseppe Sessa, Robert Dadashi, Léonard Hussenot, Johan Ferret, Nino Vieillard, Alexandre
 683 Ramé, Bobak Shariari, Sarah Perrin, Abe Friesen, Geoffrey Cideron, et al. Bond: Aligning llms
 684 with best-of-n distillation. *arXiv preprint arXiv:2407.14622*, 2024.

685 Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally
 686 can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024.

687 Yuda Song, Gokul Swamy, Aarti Singh, Drew Bagnell, and Wen Sun. The importance of online data:
 688 Understanding preference fine-tuning via coverage. In *The Thirty-eighth Annual Conference on
 689 Neural Information Processing Systems*, 2024.

690 Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford,
 691 Dario Amodei, and Paul F Christiano. Learning to summarize with human feedback. *Advances
 692 in Neural Information Processing Systems*, 33:3008–3021, 2020.

693 Benedikt Stroebel, Sayash Kapoor, and Arvind Narayanan. Inference scaling f-laws: The limits of
 694 llm resampling with imperfect verifiers. *arXiv preprint arXiv:2411.17501*, 2024.

695 Hanshi Sun, Momin Haider, Ruiqi Zhang, Huitao Yang, Jiahao Qiu, Ming Yin, Mengdi Wang, Peter
 696 Bartlett, and Andrea Zanette. Fast best-of-n decoding via speculative rejection. *arXiv preprint
 697 arXiv:2410.20290*, 2024.

702 Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut,
 703 Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly
 704 capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.

705 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Niko-
 706 lay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open founda-
 707 tion and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.

708 Haoxiang Wang, Wei Xiong, Tengyang Xie, Han Zhao, and Tong Zhang. Interpretable preferences
 709 via multi-objective reward modeling and mixture-of-experts. In *EMNLP*, 2024.

710 Xueqin Wang, Yunlu Jiang, Mian Huang, and Heping Zhang. Robust variable selection with expo-
 711 nential squared loss. *Journal of the American Statistical Association*, 2013.

712 Wei Xiong, Hanze Dong, Chenlu Ye, Ziqi Wang, Han Zhong, Heng Ji, Nan Jiang, and Tong Zhang.
 713 Iterative preference learning from human feedback: Bridging theory and practice for rlhf under
 714 kl-constraint. In *Forty-first International Conference on Machine Learning*, 2024.

715 Joy Qiping Yang, Salman Salamatian, Ziteng Sun, Ananda Theertha Suresh, and Ahmad Beirami.
 716 Asymptotics of language model alignment. In *2024 IEEE International Symposium on Informa-
 717 tion Theory (ISIT)*, pp. 2027–2032. IEEE, 2024.

718 Kaylee Yingxi Yang and Andre Wibisono. Convergence of the inexact langevin algorithm and score-
 719 based generative models in kl divergence. *arXiv preprint arXiv:2211.01512*, 2022.

720 Chenlu Ye, Wei Xiong, Yuheng Zhang, Nan Jiang, and Tong Zhang. A theoretical analysis of
 721 nash learning from human feedback under general kl-regularized preference. *arXiv preprint
 722 arXiv:2402.07314*, 2024.

723 Wenhao Zhan, Masatoshi Uehara, Nathan Kallus, Jason D Lee, and Wen Sun. Provable offline
 724 preference-based reinforcement learning. *arXiv preprint arXiv:2305.14816*, 2023.

725 Heyang Zhao, Chenlu Ye, Quanquan Gu, and Tong Zhang. Sharp analysis for kl-regularized contex-
 726 tual bandits and rlhf. *arXiv preprint arXiv:2411.04625*, 2024.

727 Yao Zhao, Mikhail Khalman, Rishabh Joshi, Shashi Narayan, Mohammad Saleh, and Peter J Liu.
 728 Calibrating sequence likelihood improves conditional language generation. In *The Eleventh Inter-
 729 national Conference on Learning Representations*, 2022.

730 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang,
 731 Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and
 732 chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623, 2023.

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756

757

758

759

Appendix

760

761

Table of Contents

762

763

764

765

766

767

768

769

770

771

772

773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

| | |
|---|-----------|
| A Table of Notations | 16 |
| B Other Related Works | 16 |
| C Calibrated Reward | 17 |
| D Summary of KL divergence Results | 17 |
| E Gumbel-Max trick | 17 |
| F Technical Tools | 18 |
| G Proof and Details of Section 4 | 21 |
| H Proof and details of Section 5 | 23 |
| H.1 Upper bound proof and details | 23 |
| H.2 Lower bound proofs and details: | 27 |
| I Experiments | 29 |
| I.1 More Experiments | 29 |
| I.2 Numerical Examples | 30 |
| J Comparison with Huang et al. (2025) | 33 |
| K Information projection across reward functions | 34 |

810
811
A TABLE OF NOTATIONS812
813
All notations are summarized in Table 1.

| Notation | Definition |
|-----------------------------|---|
| $r^*(y, x)$ | (Calibrated) true-reward model |
| $\hat{r}(y, x)$ | (Calibrated) proxy-reward model |
| N | Number of responses |
| x | Prompt |
| β | Temperature (regularization coefficient) |
| $\pi_{\beta, r}(y x)$ | Tilted policy under reward function r |
| $\pi_r^{\text{SBoN}}(y x)$ | SBoN policy under reward function r |
| $\pi_r^{\text{BoN}}(y x)$ | BoN policy under reward function r |
| π_{ref} | Reference policy |
| $\text{KL}(p\ q)$ | KL divergence between p and q distributions |
| $\pi_r^*(y x)$ | Optimal policy |
| $C_{\beta, r}(x)$ | Coverage constant under reward model r |
| $\mathcal{Y}_{r^*}^*(x)$ | Set of maximizers for $r^*(y, x)$ |
| $\varepsilon_{\beta, r}(x)$ | Tilted Error |

829
830
831
832
Table 1: Summary of notations in the paper833
834
B OTHER RELATED WORKS

835
836
Smoothing of Maximum: Approximating the maximum operator using a smoothed or softmax-based surrogate is a widely adopted technique in machine learning. This approach is particularly useful in settings where the hard maximum is non-differentiable or leads to unstable optimization. For instance, in robust regression, smooth approximations to the max operator are used in min-max formulations to achieve tractable optimization under distributional shifts (Wang et al., 2013; Li et al., 2023). In sequential decision-making, similar ideas appear in risk-sensitive control and Q-learning, where the softmax of Q-values leads to stochastic policies that balance exploration and exploitation (Howard & Matheson, 1972; Borkar, 2002). In convex and non-convex optimization, smoothing the maximum objective has been shown to improve convergence properties (Kort & Bertsekas, 1972; Pee & Royset, 2011; Liu & Theodorou, 2019). The Soft Best-of-N (SBoN) framework, (Mayrink Verdun et al., 2025; Khanov et al., 2024; Jinnai et al., 2024), leverages this principle by replacing the hard selection of the highest-reward sample with a softmax-weighted sampling distribution. Regarding the SBoN, the empirical version of SBoN is introduced by (Khanov et al., 2024) as ARGs-stochastic, where a token from a probability distribution among the top-k candidate tokens is chosen. Then, the regularized version of BoN, which can be represented as SBoN, is discussed by (Jinnai et al., 2024). Given the broad success of SBoN, we are motivated to theoretically investigate the SBoN policies and the effect of the proxy-reward model (reward hacking) and the quality of the reference policy.

852
853
Theoretical Foundation of RLHF: Several works have studied the theoretical underpinnings of reverse KL-regularized RLHF, particularly in terms of sample complexity (Zhao et al., 2024; Xiong et al., 2024; Song et al., 2024; Zhan et al., 2023; Ye et al., 2024; Aminian et al., 2025). Note that, as the sampling distributions in BoN and SBoN are different, we can not apply RLHF analysis to these sampling strategies. Therefore, it is needed to develop new foundations for BoN and SBoN.

857
858
Overoptimization. Alignment methods are widely known to suffer from overoptimization, also known as misspecification, reward hacking, or Goodhart Law, where optimizing against a proxy-reward model leads to degraded performance compared to the true-reward model (Amodei et al., 2016; Casper et al., 2023; Gao et al., 2023b). This issue is particularly pronounced in inference-time alignment methods such as BoN, where an increasing number of responses N makes the overoptimization problem worse (Huang et al., 2025; Stroebel et al., 2024; Gao et al., 2023b). Huang et al. (2025) theoretically demonstrate that the BoN policy suffers from overoptimization when N is large, given a fixed estimation error in the reward model, and propose a solution based on a χ^2 -regularized

framework. Other approaches to mitigating this issue include ensembling strategies (Coste et al.; Eisenstein et al.) and regularization techniques (Ichihara et al.). In a concurrent line of work, Khalaf et al. (2025) introduce the Best-of-Poisson method to reduce overoptimization in inference-time algorithms. The overoptimization in BoN and SBoN is also studied by Khalaf et al. (2025) and a principled hedging framework is proposed to mitigate the overoptimization. In contrast, we study overoptimization in inference-time alignment methods SBoN and BoN from the perspectives of regret gap and KL divergence analysis.

C CALIBRATED REWARD

Inspired by (Balashankar et al., 2025), in this section, we provide more details regarding calibrated reward. A standard metric for evaluation of models is the *win-rate* relative to a base policy π_{ref} (Stiennon et al., 2020; Gao et al., 2023b). For a prompt x and responses y, z , define the win random variable under raw (uncalibrated) reward r_{uc} as

$$w_r(y, z | x) = \mathbf{1}\{r_{\text{uc}}(y, x) > r_{\text{uc}}(z, x)\} + \frac{1}{2}\mathbf{1}\{r_{\text{uc}}(y, x) = r_{\text{uc}}(z, x)\}.$$

Definition C.1 (Calibrated reward). *The calibrated reward of y under policy π is its expected win-rate probability against $z \sim \pi(\cdot | x)$:*

$$r_{c,\pi}(y, x) := \mathbb{E}_{z \sim \pi(\cdot | x)}[w_r(y, z | x)].$$

In practice, we consider $\pi = \pi_{\text{ref}}$, therefore we denote calibrated reward via $r(y, x)$ under reference policy. In the following, we provide some reasons for choosing calibrated reward instead of raw (uncalibrated) reward in our work,

- **Matches win-rate evaluation.** For any policies π_1, π_2 ,

$$W_r(\pi_1 \succ \pi_2 | x) := \mathbb{E}_{y \sim \pi_1(\cdot | x)}[r_{c,\pi_2}(y, x)],$$

where $W_r(\pi_1 \succ \pi_2 | x)$ is standard win-rate. So maximizing $\mathbb{E}_{Y \sim \pi}[r_{c,\pi_{\text{ref}}}(y, x)]$ directly optimizes standard win rate vs. the base model.

- **Invariance to score scaling.** If m is strictly increasing and $r' = m \circ r$, then

$$r'(y, x) = r(y, x),$$

making the target robust to arbitrary monotone reparameterizations of the reward (e.g., affine rescaling, temperature).

- **Unified, probabilistic scale.** For $y \sim \pi_{\text{ref}}(\cdot | x)$,

$$r(y, x) \sim \text{Unif}[0, 1],$$

independent of both r and π_{ref} . This normalizes per-prompt reward scales and interprets scores as win probabilities.

D SUMMARY OF KL DIVERGENCE RESULTS

In Table 2, we summarize results on KL divergences between the aligned and reference policies, along with corresponding upper bounds for both SBoN and BoN policies. Furthermore, in Table 3, we summarize results on KL divergences between aligned policies under true and proxy-reward models, along with upper bounds for SBoN and tilted policies.

E GUMBEL-MAX TRICK

We also provide an interpretation for SBoN from the Gumbel-Max trick. An alternative way to sample Z from

$$\Pr(Z = i) \propto \exp(\beta \hat{r}(Y_i, x))$$

is via the Gumbel-Max trick. We can draw independent Gumbel-distributed random variables $G_i \sim \text{Gumbel}(0, 1)$, $i = 1, \dots, n$, and then set

$$Z = \arg \max_{i \in \{1, \dots, N\}} \left[\hat{r}(Y_i, x) + \frac{G_i}{\beta} \right].$$

918 Table 2: KL divergences between the aligned and reference policies, along with corresponding upper
 919 bounds for both SBoN and BoN policies.
 920

| 921 KL divergence Term | 922 Theorem / Lemma | 923 Upper Bound |
|---|--|--|
| 922 $\text{KL}(\pi_{r^*}^{\text{SBoN}}(\cdot x) \parallel \pi_{\text{ref}}(\cdot x))$ | 923 Lemma 4.1 | $\log\left(\frac{N}{1+(N-1)\exp(-\beta)}\right)$ |
| 924 $\text{KL}(\pi_{r^*}^{\text{BoN}}(\cdot x) \parallel \pi_{\text{ref}}(\cdot x))$ | 925 Theorem 3.1 in (Beirami et al., 2024) and Theorem 1 in (Mroueh, 2024) | 926 $\log(N) - 1 + 1/N$ |

927 Table 3: KL divergences between aligned policies under true and proxy-reward models, along with
 928 upper bounds for SBoN and tilted policies.
 929

| 930 KL divergence Term | 931 Theorem / Lemma | 932 Upper Bound |
|--|----------------------------|--|
| 932 $\text{KL}(\pi_{r^*}^{\text{SBoN}}(\cdot x) \parallel \pi_{\hat{r}}^{\text{SBoN}}(\cdot x))$ | 933 Lemma 4.2 | $\frac{N\beta\sqrt{\varepsilon_{\beta,r}(x)}}{1+(N-1)\exp(-\beta)} \left(\frac{N\exp(2\beta)}{(N-1)^2} + 1 \right)$ |
| 933 $\text{KL}(\pi_{\beta,r^*}(\cdot x) \parallel \pi_{\beta,\hat{r}}(\cdot x))$ | 934 Lemma F.8 | $2\beta\sqrt{\varepsilon_{\beta,r}(x)} \left(\sqrt{\frac{\mathbb{E}[\exp(2\beta\hat{r}(Y,x))]}{\mathbb{E}^2[\exp(\beta\hat{r}(Y,x))]}} + \sqrt{\frac{\mathbb{E}[\exp(2\beta r^*(Y,x))]}{\mathbb{E}^2[\exp(\beta r^*(Y,x))]}} \right)$ |

935
 936 By properties of the Gumbel distribution, this yields exactly the same softmax sampling law, without
 937 needing to compute the normalizing factor $\sum_{j=1}^N \exp(\beta\hat{r}(Y_j, x))$ explicitly (Gumbel, 1954). When
 938 $\beta \rightarrow \infty$, the effect of the Gumbel noises vanishes and the sampling strategy reduces to BoN.
 939

940 F TECHNICAL TOOLS

941 We denote the set maximizers of the proxy-reward function via $\hat{\mathcal{Y}}(x) = \{\hat{y}_j(x)\}_{j=1}^{m(x)}$.
 942

943 We introduce the functional derivative, see Cardaliaguet et al. (2019).

944 **Definition F.1.** (Cardaliaguet et al., 2019) A functional $U : \mathcal{P}(\mathbb{R}^n) \rightarrow \mathbb{R}$ admits a functional
 945 derivative if there is a map $\frac{\delta U}{\delta m} : \mathcal{P}(\mathbb{R}^n) \times \mathbb{R}^n \rightarrow \mathbb{R}$ which is continuous on $\mathcal{P}(\mathbb{R}^n)$ and, for all
 946 $m, m' \in \mathcal{P}(\mathbb{R}^n)$, it holds that
 947

$$948 U(m') - U(m) = \int_0^1 \int_{\mathbb{R}^n} \frac{\delta U}{\delta m}(m_\beta, a) (m' - m)(da) d\beta,$$

949 where $m_\beta = m + \beta(m' - m)$.
 950

951 **Definition F.2** (Sensitivity of a policy). We also define the sensitivity of a policy $\pi_r(y|x)$, which is
 952 a function of reward function $r(y, x)$, with respect to the reward function as
 953

$$954 \frac{\partial \pi}{\partial r}(r) := \lim_{\Delta r \rightarrow 0} \frac{\pi_r(y|x) - \pi_{r+\Delta r}(y|x)}{\Delta r}. \quad (20)$$

955 **Lemma F.3** (Kantorovich-Rubenstein duality of total variation distance, see (Polyanskiy & Wu,
 956 2022)). The Kantorovich-Rubenstein duality (variational representation) of the total variation dis-
 957 tance is as follows:
 958

$$959 \text{TV}(m_1, m_2) = \frac{1}{L} \sup_{g \in \mathcal{G}_L} \{ \mathbb{E}_{Z \sim m_1}[g(Z)] - \mathbb{E}_{Z \sim m_2}[g(Z)] \}, \quad (21)$$

960 where $\mathcal{G}_L = \{g : \mathcal{Z} \rightarrow \mathbb{R}, \|g\|_\infty \leq L\}$.
 961

962 **Lemma F.4** (Lemma 5.4 in (Aminian et al., 2025)). Consider the softmax policy, $\pi_r^\beta(y|x) \propto$
 963 $\pi_{\text{ref}}(y|x) \exp(\beta r(y, x))$. Then, the sensitivity of the policy with respect to the reward function is
 964

$$965 \frac{\partial \pi_r^\beta}{\partial r}(r) = \beta \pi_r^\beta(y|x) (1 - \pi_r^\beta(y|x)).$$

966 **Lemma F.5** (Pinsker's Inequality (Canonne, 2022)). For m_1 and m_2 , we have,
 967

$$968 \text{TV}(m_1, m_2) \leq \sqrt{\frac{1}{2} \text{KL}(m_2 \parallel m_1)}. \quad (22)$$

972 The following Lemmata are useful for our technical proofs.
 973
 974

Lemma F.6. *The following upper bound holds,*

$$975 \quad \log \left(\frac{Z_{r^*,Y}(x, \beta)}{Z_{\hat{r},Y}(x, \beta)} \right) \leq \beta \sqrt{\varepsilon_{\beta,r}(x)} \sqrt{C_{\beta,r^*}(x)}. \quad (23)$$

$$976$$

977 *Proof.*

$$978 \quad \begin{aligned} \frac{Z_{\hat{r},Y}(x, \beta)}{Z_{r^*,Y}(x, \beta)} &= \frac{\sum_y \exp(\beta \hat{r}(y, x)) \pi_{\text{ref}}(y|x)}{\sum_y \exp(\beta r^*(y, x)) \pi_{\text{ref}}(y|x)} \\ 980 &= \frac{\sum_y \exp(\beta(\hat{r}(y, x) - r^*(y, x))) \exp(\beta r^*(y, x)) \pi_{\text{ref}}(y|x)}{\sum_y \exp(\beta r^*(y, x)) \pi_{\text{ref}}(y|x)} \\ 982 &= \sum_y \pi_{\beta,r^*}(y|x) \exp(\beta(\hat{r}(y, x) - r^*(y, x))) \end{aligned} \quad (24)$$

$$983$$

$$984$$

$$985$$

986 Due to convexity of $-\log(\cdot)$ and using Cauchy-Schwarz inequality, we have,

$$987 \quad \begin{aligned} -\log \left(\frac{Z_{\hat{r},Y}(x, \beta)}{Z_{r^*,Y}(x, \beta)} \right) &\leq \beta \sum_y \pi_{\beta,r^*}(y|x) (r^*(y, x) - \hat{r}(y, x)) \\ 988 &\leq \beta \sum_y \frac{\pi_{\beta,r^*}(y|x)}{\pi_{\text{ref}}(y|x)} (r^*(y, x) - \hat{r}(y, x)) \pi_{\text{ref}}(y|x) \\ 989 &\leq \beta \sqrt{\sum_y (r^*(y, x) - \hat{r}(y, x))^2 \pi_{\text{ref}}(y|x)} \sqrt{\sum_y \frac{\pi_{\beta,r^*}^2(y|x)}{\pi_{\text{ref}}(y|x)}} \\ 990 &= \sqrt{\beta} \sqrt{\sum_y \log(\exp(\beta(r^*(y, x) - \hat{r}(y, x))^2)) \pi_{\text{ref}}(y|x)} \sqrt{C_{\beta,r^*}(x)} \\ 991 &\leq \beta \sqrt{\frac{1}{\beta} \log(\sum_y \exp(\beta(r^*(y, x) - \hat{r}(y, x))^2)) \pi_{\text{ref}}(y|x)} \sqrt{C_{\beta,r^*}(x)} \\ 992 &= \beta \sqrt{\varepsilon_{\beta,r}(x)} \sqrt{C_{\beta,r^*}(x)}, \end{aligned} \quad (25)$$

$$993$$

$$994$$

$$995$$

$$996$$

$$997$$

$$998$$

$$999$$

$$1000$$

$$1001$$

$$1002$$

$$1003$$

□

1004 **Lemma F.7.** *The following holds,*

$$1005 \quad \text{KL}(\pi_{r^*}^*(\cdot|x) \|\pi_{\text{ref}}(\cdot|x)) \leq \log(C_{\infty,r^*}(x)). \quad (26)$$

$$1006$$

$$1007$$

1008 *Proof.* Note that, we have,

$$1009 \quad \begin{aligned} \text{KL}(\pi_{r^*}^*(\cdot|x) \|\pi_{\text{ref}}(\cdot|x)) &\leq \log \left(\mathbb{E}_{Y \sim \pi_{r^*}^*(\cdot|x)} \left[\frac{\pi_{r^*}^*(\cdot|x)}{\pi_{\text{ref}}(\cdot|x)} \right] \right) \\ 1010 &\leq \log(C_{\infty,r^*}(x)). \end{aligned} \quad (27)$$

$$1011$$

$$1012$$

$$1013$$

□

1014 **Lemma F.8.** *The following upper bound holds,*

$$1015 \quad \text{KL}(\pi_{\beta,r^*}(y|x) \|\pi_{\beta,\hat{r}}(y|x)) \leq \beta \sqrt{\varepsilon_{\beta,r}(x)} (\sqrt{C_{\beta,r^*}(x)} + \sqrt{C_{\beta,\hat{r}}(x)}). \quad (28)$$

$$1016$$

$$1017$$

1018 *Proof.*

$$1019 \quad \begin{aligned} \text{KL}(\pi_{\beta,r^*}(y|x) \|\pi_{\beta,\hat{r}}(y|x)) &= \sum_y \pi_{\beta,r^*}(y|x) \log \left(\frac{\pi_{\beta,r^*}(y|x)}{\pi_{\beta,\hat{r}}(y|x)} \right) \\ 1020 &= \beta \sum_y (r^*(y, x) - \hat{r}(y, x)) \pi_{\beta,r^*}(y|x) + \log(Z_{\hat{r},Y}(x, \beta)/Z_{r^*,Y}(x, \beta)) \\ 1021 &\leq \beta \sqrt{\varepsilon_{\beta,r}(x)} (\sqrt{C_{\beta,r^*}(x)} + \sqrt{C_{\beta,\hat{r}}(x)}), \end{aligned} \quad (29)$$

$$1022$$

$$1023$$

$$1024$$

$$1025$$

1026 where the final inequality holds due to Lemma F.6 and applying Cauchy-Schwarz inequality. \square
 1027

1028 **Lemma F.9.** Suppose that $f(Z) \in [0, Z_{\max}]$, $\mathcal{Z}_{\max} = \{z_{m,i}\}_{i=1}^m$ is the set of maximizers of $f(Z)$,
 1029 i.e., $f(z) = Z_{\max}$ for $z \in \mathcal{Z}_{\max}$. Then we have,

$$1030 \lim_{\beta \rightarrow \infty} \frac{\mathbb{E}[\exp(2\beta f(Z))]}{\mathbb{E}[\exp(\beta f(Z))]^2} = \frac{1}{\sum_{z \in \mathcal{Z}_{\max}} P(Z = z)}. \quad (30)$$

1033 *Proof.*

$$1034 \frac{\mathbb{E}[\exp(2\beta f(Z))]}{\mathbb{E}[\exp(\beta f(Z))]^2} = \frac{\mathbb{E}[\exp(2\beta(f(Z) - Z_{\max}))]}{\mathbb{E}[\exp(\beta(f(Z) - Z_{\max}))]^2} \quad (31)$$

$$1037 \frac{\sum_j P(Z = z_j) \exp(2\beta(f(z_j) - Z_{\max}))}{(\sum_j P(Z = z_j) \exp(\beta(f(z_j) - Z_{\max})))^2} \quad (32)$$

1039 Now, we have,

$$1040 \lim_{\beta \rightarrow \infty} \frac{\mathbb{E}[\exp(2\beta f(Z))]}{\mathbb{E}[\exp(\beta f(Z))]^2} \quad (33)$$

$$1043 = \lim_{\beta \rightarrow \infty} \frac{\sum_j P(Z = z_j) \exp(2\beta(f(z_j) - Z_{\max}))}{(\sum_j P(Z = z_j) \exp(\beta(f(z_j) - Z_{\max})))^2} \quad (34)$$

$$1046 = \frac{\sum_{z \in \mathcal{Z}_{\max}} P(Z = z)}{(\sum_{z \in \mathcal{Z}_{\max}} P(Z = z))^2} \quad (35)$$

$$1048 = \frac{1}{\sum_{z \in \mathcal{Z}_{\max}} P(Z = z)}, \quad (36)$$

1050 where we used the fact that $\lim_{\beta \rightarrow \infty} \exp(\beta(z_j - Z_{\max})) = 0$ for $z_j < Z_{\max}$. \square
 1051

1052 **Lemma F.10** (Theorem 1 in (Mayrink Verdun et al., 2025)). For $\beta > 0$, and $N \geq 1$, we have,

$$1053 \text{KL}(\pi_{\beta,r^*}(\cdot|x) \|\pi_{r^*}^{\text{SBoN}}(y|x)) \leq \log\left(1 + \frac{C_{\beta,r^*}(x) - 1}{N}\right). \quad (37)$$

1055 **Lemma F.11.** For a given $x \in \mathcal{X}$, we have,

$$1057 \left| \frac{\delta f(r)}{\delta r} \right| \leq \frac{N^2 \beta \exp(2\beta)}{(N-1)^2}, \quad (38)$$

1059 where $f(r) = \log\left(\mathbb{E}\left[\frac{1}{\exp(\beta r) + \sum_{i=1}^{N-1} \exp(\beta R_i)}\right]\right)$, $r = r(y, x)$ and $R_i = r(Y_i, x)$.
 1060

1061 *Proof.* Note that $\{R_i\}_{i=1}^{N-1}$ are i.i.d. . Therefore, we have,

$$1063 \frac{\delta f(r(y, x))}{\delta r} = \mathbb{E}\left[\frac{1}{\exp(\beta r) + \sum_{i=1}^{N-1} \exp(\beta R_i)}\right]^{-1} \frac{\delta \mathbb{E}\left[\frac{1}{\exp(\beta r) + \sum_{i=1}^{N-1} \exp(\beta R_i)}\right]}{\delta r} \\ 1064 \leq \mathbb{E}\left[\frac{1}{\exp(\beta r) + \sum_{i=1}^{N-1} \exp(\beta R_i)}\right]^{-1} \\ 1066 \times \left(\sum_{k=1}^N \frac{\beta k \binom{N-1}{k-1} \exp(\beta r)}{(k \exp(\beta r) + N-1-k)^2} (1 - P(R=r))^{N-k} P^{k-1}(R=r)\right) \\ 1068 \\ 1069 \leq \mathbb{E}\left[\frac{1}{\exp(\beta r) + \sum_{i=1}^{N-1} \exp(\beta R_i)}\right]^{-1} \\ 1071 \times \frac{\beta \exp(\beta)}{(N-1)^2} \left(\sum_{k=1}^N k \binom{N-1}{k-1} (1 - P(R=r))^{N-k} P^{k-1}(R=r)\right) \\ 1073 \\ 1074 \leq \frac{N \beta \exp(2\beta)}{(N-1)^2} (1 + (N-1)P(R=r)) \\ 1076 \\ 1077 \leq \frac{N^2 \beta \exp(2\beta)}{(N-1)^2}. \quad (39)$$

1080
1081 For first inequality provide more details in the following. Lets consider the conditional expectation
1082 where we condition on having exactly $(k - 1)$ of $(N - 1)$ R_i s equal to r . Therefore, we have,
1083

$$\begin{aligned} & \frac{\delta \mathbb{E} \left[\frac{1}{\exp(\beta r) + \sum_{i=1}^{N-1} \exp(\beta R_i)} \right]}{\delta r} \\ &= \sum_{k=1}^N \binom{N-1}{k-1} \mathbb{E} \left[\frac{\beta k \exp(\beta r)}{k \exp(\beta r) + \sum_{i=1}^{N-k} \exp(\beta R_i)} \right] (1 - P(R = r))^{N-k} P^{k-1}(R = r) \end{aligned}$$

1088 Now computing the derivative and using the fact that $1 \leq \exp(\beta R_i)$, the first inequality holds.
1089 \square

1090
1091
1092 **Lemma F.12** (The BH inequality (Canonne, 2022)). *For every two probability distributions \mathbf{p}, \mathbf{q} ,
1093 we have the simple yet never vacuous bound*
1094

$$\mathbb{TV}(\mathbf{p}, \mathbf{q}) \leq \sqrt{1 - e^{-\text{KL}(\mathbf{p} \parallel \mathbf{q})}}. \quad (40)$$

1095 G PROOF AND DETAILS OF SECTION 4

1101 **Lemma 4.1.** *The following upper bound holds on KL divergence between SBoN and refer-
1102 ence policies for a given prompt $x \in \mathcal{X}$,*
1103

$$\text{KL}(\pi_{r^*}^{\text{SBoN}}(y|x) \parallel \pi_{\text{ref}}(y|x)) \leq \log \left(\frac{N}{1 + (N-1) \exp(-\beta)} \right). \quad (41)$$

1108 *Proof.* Recall that,

$$\pi_{r^*}^{\text{SBoN}}(y|x) = \pi_{\text{ref}}(y|x) \exp(\beta r^*(y, x)) \mathbb{E} \left[\left(\frac{1}{N} (\exp(\beta r^*(y, x)) + \sum_{i=1}^{N-1} \exp(\beta r^*(Y_i, x))) \right)^{-1} \right].$$

1113 Now, we have,

$$\begin{aligned} & \text{KL}(\pi_{r^*}^{\text{SBoN}}(y|x) \parallel \pi_{\text{ref}}(y|x)) \\ &= \sum_y \pi_{r^*}^{\text{SBoN}}(y|x) \log(\pi_{r^*}^{\text{SBoN}}(y|x) / \pi_{\text{ref}}(y|x)) \\ &= \sum_y \pi_{r^*}^{\text{SBoN}}(y|x) \log(N) \\ &+ \sum_y \pi_{r^*}^{\text{SBoN}}(y|x) \log(\mathbb{E} \left[\exp(\beta r^*(y, x)) \left(\exp(\beta r^*(y, x)) + \sum_{i=1}^{N-1} \exp(\beta r^*(Y_i, x)) \right)^{-1} \right]) \\ &= \log(N) + \sum_y \pi_{r^*}^{\text{SBoN}}(y|x) \log(\mathbb{E} \left[\left(1 + \sum_{i=1}^{N-1} \exp(\beta(r^*(Y_i, x) - r^*(y, x))) \right)^{-1} \right]). \end{aligned} \quad (42)$$

1127 For the second term in equation 42, consider
1128

$$A(y, Y, x) = \sum_{i=1}^{N-1} \exp(\beta(r^*(Y_i, x) - r^*(y, x))) > 0,$$

1132 where we have
1133

$$(N-1) \exp(-\beta) \leq A(y, Y, x) \leq (N-1) \exp(\beta).$$

1134 Therefore, we have,
 1135

$$\begin{aligned}
 1136 \quad & \sum_{\mathcal{Y}} \pi_{r^*}^{\text{SBoN}}(y|x) \log(\mathbb{E}\left[\left(1 + \sum_{i=1}^{N-1} \exp(\beta(r^*(Y_i, x) - r^*(y, x)))\right)^{-1}\right]) \\
 1137 \quad & \leq \sum_{\mathcal{Y}} \pi_{r^*}^{\text{SBoN}}(y|x) \log\left(\frac{1}{1 + (N-1) \exp(-\beta)}\right) \\
 1138 \quad & = \log\left(\frac{1}{1 + (N-1) \exp(-\beta)}\right).
 \end{aligned} \tag{43}$$

1144 Combining equation 43 with equation 42 completes the proof. □

1146

1147

1148

1149

1150

1151

Lemma 4.2. *The following upper bound holds on the KL divergence between the SBoN policies under true-reward and proxy-reward models respectively,*

$$\text{KL}(\pi_{r^*}^{\text{SBoN}}(\cdot|x) \parallel \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)) \leq \frac{N\beta\sqrt{\varepsilon_{\beta,r}(x)}}{1 + (N-1)\exp(-\beta)} \left(\frac{N\exp(2\beta)}{(N-1)^2} + 1 \right). \tag{44}$$

1152

1153

1154

1155

1156

1157 *Proof.* We first provide the following upper bound,

$$\begin{aligned}
 1158 \quad & \text{KL}(\pi_{r^*}^{\text{SBoN}}(y|x) \parallel \pi_{\hat{r}}^{\text{SBoN}}(y|x)) \\
 1159 \quad & = \sum_{\mathcal{Y}} \pi_{r^*}^{\text{SBoN}}(y|x) \log\left(\frac{\pi_{r^*}^{\text{SBoN}}(y|x)}{\pi_{\hat{r}}^{\text{SBoN}}(y|x)}\right) \\
 1160 \quad & = \sum_{\mathcal{Y}} \pi_{r^*}^{\text{SBoN}}(y|x) \beta(r^*(y, x) - \hat{r}(y, x)) \\
 1161 \quad & + \sum_{\mathcal{Y}} \pi_{r^*}^{\text{SBoN}}(y|x) \left(\log\left(\mathbb{E}\left[\frac{1}{\exp(\beta r^*(y, x)) + \sum_{i=1}^{N-1} \exp(\beta r^*(Y_i, x))}\right]\right) \right. \\
 1162 \quad & \quad \left. - \log\left(\mathbb{E}\left[\frac{1}{\exp(\beta \hat{r}(y, x)) + \sum_{i=1}^{N-1} \exp(\beta \hat{r}(Y_i, x))}\right]\right) \right) \\
 1163 \quad & \leq \frac{N\beta\sqrt{\varepsilon_{\beta,r}(x)}}{1 + (N-1)\exp(-\beta)} \\
 1164 \quad & + \sum_{\mathcal{Y}} \pi_{r^*}^{\text{SBoN}}(y|x) \left(\log\left(\mathbb{E}\left[\frac{1}{\exp(\beta r^*(y, x)) + \sum_{i=1}^{N-1} \exp(\beta r^*(Y_i, x))}\right]\right) \right. \\
 1165 \quad & \quad \left. - \log\left(\mathbb{E}\left[\frac{1}{\exp(\beta \hat{r}(y, x)) + \sum_{i=1}^{N-1} \exp(\beta \hat{r}(Y_i, x))}\right]\right) \right), \tag{45}
 \end{aligned}$$

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178 where we used two facts. First,

$$\pi_{r^*}^{\text{SBoN}}(\cdot|x) \leq \frac{N\pi_{\text{ref}}(\cdot|x)}{1 + (N-1)\exp(-\beta)}$$

1179

1180

1181

1182

1183

1184

1185

1186

1187

Second, for a random variable X ,

$$E[X] \leq \sqrt{E[X^2]} = \sqrt{E\left[\frac{1}{\beta} \log(\exp(\beta X^2))\right]} \leq \sqrt{\frac{1}{\beta} \log(E[\exp(\beta X^2)])},$$

and applying it to $E_{\pi_{\text{ref}}(\cdot|x)}[(r^*(y, x) - \hat{r}(y, x))]$, the final result in holds.

1188 Note that for the last term in equation 45, we can apply the mean-value theorem as follows,
1189

$$\begin{aligned} 1190 \sum_{\mathcal{Y}} \pi_{r^*}^{\text{SBoN}}(y|x) & \left(\log \left(\mathbb{E} \left[\frac{1}{\exp(\beta r^*(y, x)) + \sum_{i=1}^{N-1} \exp(\beta r^*(Y_i, x))} \right] \right) \right. \\ 1191 & \left. - \log \left(\mathbb{E} \left[\frac{1}{\exp(\beta \hat{r}(y, x)) + \sum_{i=1}^{N-1} \exp(\beta \hat{r}(Y_i, x))} \right] \right) \right) \\ 1192 & \leq \sum_{\mathcal{Y}} \pi_{r^*}^{\text{SBoN}}(y|x) |r^*(y, x) - \hat{r}(y, x)| \left| \frac{\delta f(r_\gamma(y, x))}{\delta r} \right|, \\ 1193 & \end{aligned} \quad (46)$$

1194 where $f(r_\gamma(y, x)) = \log \left(\mathbb{E} \left[\frac{1}{\exp(\beta r_\gamma(y, x)) + \sum_{i=1}^{N-1} \exp(\beta r_\gamma(Y_i, x))} \right] \right)$, for some $\gamma \in (0, 1)$ we have
1195 $r_\gamma(y, x) = \gamma \hat{r}(y, x) + (1 - \gamma) r^*(y, x)$. Using Lemma F.11, we have,
1196

$$\left| \frac{\delta f(r(y, x))}{\delta r} \right| \leq \frac{N^2 \beta \exp(2\beta)}{(N-1)^2}. \quad (47)$$

1200 Using equation 47 in equation 46 and applying Cauchy-Schwarz inequality, we have,
1201

$$\begin{aligned} 1202 \sum_{\mathcal{Y}} \pi_{r^*}^{\text{SBoN}}(y|x) & \left(\log \left(\mathbb{E} \left[\frac{1}{\exp(\beta r^*(y, x)) + \sum_{i=1}^{N-1} \exp(\beta r^*(Y_i, x))} \right] \right) \right. \\ 1203 & \left. - \log \left(\mathbb{E} \left[\frac{1}{\exp(\beta \hat{r}(y, x)) + \sum_{i=1}^{N-1} \exp(\beta \hat{r}(Y_i, x))} \right] \right) \right) \\ 1204 & \leq \sqrt{\sum_{\mathcal{Y}} \mathbb{E} \left[1 / \left(1 + \sum_{i=1}^{N-1} \exp(\beta(r^*(Y_i, x) - r^*(y, x))) \right)^2 \right] \pi_{\text{ref}}(y|x)} \\ 1205 & \times \sqrt{\sum_{\mathcal{Y}} |r^*(y, x) - \hat{r}(y, x)|^2 \pi_{\text{ref}}(y|x)} \sqrt{\sum_{\mathcal{Y}} \left| \frac{\delta f(r_\gamma(y, x))}{\delta r} \right|^2 \pi_{\text{ref}}(y|x)} \\ 1206 & \leq \frac{\sqrt{\varepsilon_{\beta, r}(x)}}{1 + (N-1) \exp(-\beta)} \frac{N^2 \beta \exp(2\beta)}{(N-1)^2}. \\ 1207 & \end{aligned} \quad (48)$$

1208 It completes the proof. □
1209

H PROOF AND DETAILS OF SECTION 5

H.1 UPPER BOUND PROOF AND DETAILS

1210 **Lemma 5.1 (Full Version).** *Under Assumption 3.1, the following properties of $C_{\beta, r}(x)$ hold,*

1. $C_{\beta, r}(x) = \frac{\mathbb{E}[\exp(2\beta \hat{r}(Y, x))]}{\mathbb{E}^2[\exp(\beta \hat{r}(Y, x))]}$.
2. $C_{\beta, r}(x)$ is an increasing function with respect to β .
3. $C_{\infty, \hat{r}}(x) = \frac{1}{\sum_i \pi_{\text{ref}}(y_{i, r}^{\max}(x) | x)}$ where $y_{i, r}^{\max}(x) \in \arg \max_y r(y, x)$.
4. For all $\beta < \infty$, we have $1 \leq C_{\beta, r}(x) \leq \min(C_{\infty, \hat{r}}(x), \exp(2\beta))$.

1211 *Proof.* In the following, we provide proofs of different items.
1212

1213 1.

$$\begin{aligned} 1214 C_{\beta, \hat{r}}(x) &= \sum_{\mathcal{Y}} \frac{\pi_{\beta, \hat{r}}^2(y|x)}{\pi_{\text{ref}}(y|x)} \\ 1215 &= \sum_{\mathcal{Y}} \frac{\exp(2\beta \hat{r}(y, x))}{\mathbb{E}^2[\exp(\beta \hat{r}(Y, x))]} \pi_{\text{ref}}(y|x) \\ 1216 &= \frac{\mathbb{E}[\exp(2\beta \hat{r}(Y, x))]}{\mathbb{E}^2[\exp(\beta \hat{r}(Y, x))]} \\ 1217 & \end{aligned} \quad (49)$$

1242 2. We can show that the logarithm function of $C_{\beta, \hat{r}}(x)$ is increasing. Then, due to the increasing
 1243 feature of the log function, the final result holds.

1244

$$\begin{aligned} & \log\left(\frac{\mathbb{E}[\exp(2\beta\hat{r}(Y, x))]}{\mathbb{E}^2[\exp(\beta\hat{r}(Y, x))]} \right) \\ &= \log(\mathbb{E}[\exp(2\beta\hat{r}(Y, x))]) - 2\log(\mathbb{E}[\exp(\beta\hat{r}(Y, x))]), \end{aligned} \quad (50)$$

1245 then we can compute the derivative of equation 50,

1246

$$\begin{aligned} & \frac{d \log(\mathbb{E}[\exp(2\beta\hat{r}(Y, x))])}{d\beta} - 2 \frac{d \log(\mathbb{E}[\exp(\beta\hat{r}(Y, x))])}{d\beta} \\ &= \frac{\mathbb{E}[2\hat{r}(Y, x) \exp(2\beta\hat{r}(Y, x))]}{\mathbb{E}[\exp(2\beta\hat{r}(Y, x))]} - 2 \frac{\mathbb{E}[\hat{r}(Y, x) \exp(\beta\hat{r}(Y, x))]}{\mathbb{E}[\exp(\beta\hat{r}(Y, x))]} \end{aligned} \quad (51)$$

1247 Note that we have,

1248

$$\begin{aligned} & \frac{d \frac{\mathbb{E}[\hat{r}(Y, x) \exp(\beta\hat{r}(Y, x))]}{\mathbb{E}[\exp(\beta\hat{r}(Y, x))]} }{d\beta} \\ &= \frac{\mathbb{E}[\hat{r}^2(Y, x) \exp(\beta\hat{r}(Y, x))] \mathbb{E}[\exp(\beta\hat{r}(Y, x))] - \mathbb{E}[\hat{r}(Y, x) \exp(\beta\hat{r}(Y, x))]^2}{\mathbb{E}^2[\exp(\beta\hat{r}(Y, x))]} \\ &= \mathbb{E}_{Y \sim \pi_{\beta, \hat{r}}(\cdot|x)} [\hat{r}^2(Y, x)] - \mathbb{E}_{Y \sim \pi_{\beta, \hat{r}}(\cdot|x)} [\hat{r}(Y, x)]^2 \geq 0. \end{aligned} \quad (52)$$

1249 Therefore, we have,

1250

$$\frac{\mathbb{E}[\hat{r}(Y, x) \exp(2\beta\hat{r}(Y, x))]}{\mathbb{E}[\exp(2\beta\hat{r}(Y, x))]} \geq \frac{\mathbb{E}[\hat{r}(Y, x) \exp(\beta\hat{r}(Y, x))]}{\mathbb{E}[\exp(\beta\hat{r}(Y, x))]} \quad (53)$$

1251 It completes the proof. □

1252 3. Follows directly from Lemma F.9.

1253 4. Due to Jensen inequality for $\mathbb{E}^2[\exp(\beta\hat{r}(Y, x))] \leq \mathbb{E}[\exp(2\beta\hat{r}(Y, x))]$, the $C_{\beta, \hat{r}}(x)$. We
 1254 also have the uniform bound, $C_{\beta, \hat{r}}(x) = \frac{\mathbb{E}[\exp(2\beta\hat{r}(Y, x))]}{\mathbb{E}^2[\exp(\beta\hat{r}(Y, x))]} \leq \exp(\beta)$. Furthermore, due to
 1255 increasing property in second item, we also have $\sup_{\beta} C_{\beta, \hat{r}}(x) = C_{\infty, \hat{r}}(x)$. Therefore, the
 1256 upper bound holds.

1257

1258 **Theorem H.1.** *The following upper bound holds,*

1259

$$\begin{aligned} \Delta_{J_{r^*}}(\pi_{\beta, r^*}(\cdot|x), \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)) &\leq \frac{1}{\beta} \left(\text{KL}(\pi_{\beta, r^*}(\cdot|x) \|\pi_{\text{ref}}(\cdot|x)) - \text{KL}(\pi_{\beta, \hat{r}}(\cdot|x) \|\pi_{\text{ref}}(\cdot|x)) \right) \\ &+ \sqrt{\varepsilon_{\beta, r}(x)} (\sqrt{C_{\beta, \hat{r}}(x)} + \sqrt{C_{\beta, r^*}(x)}) \\ &+ 2\sqrt{\frac{1}{2} \log \left(1 + \frac{C_{\beta, \hat{r}}(x) - 1}{N} \right)}. \end{aligned}$$

1260 *Proof.* Note that, we have,

1261

$$\begin{aligned} & \Delta_{J_{r^*}}(\pi_{\beta, r^*}(\cdot|x), \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)) \\ &= \mathbb{E}_{Y \sim \pi_{\beta, r^*}(\cdot|x)} [r^*(Y, x)] - \mathbb{E}_{Y \sim \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)} [r^*(Y, x)] \\ &= \underbrace{\mathbb{E}_{Y \sim \pi_{\beta, r^*}(\cdot|x)} [r^*(Y, x)] - \mathbb{E}_{Y \sim \pi_{\beta, \hat{r}}(\cdot|x)} [r^*(Y, x)]}_{I_1} \\ &+ \underbrace{\mathbb{E}_{Y \sim \pi_{\beta, \hat{r}}(\cdot|x)} [r^*(Y, x)] - \mathbb{E}_{Y \sim \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)} [r^*(Y, x)]}_{I_2} \end{aligned} \quad (54)$$

1296 Note that, using the definition of $\pi_{\beta, r^*}(\cdot|x)$ and $\pi_{\beta, \hat{r}}(\cdot|x)$ as solutions to KL-regularized problem,
 1297 we have,
 1298

$$1299 \\ 1300 \mathbb{E}_{Y \sim \pi_{\beta, r^*}(\cdot|x)}[r^*(Y, x)] = \frac{1}{\beta} \text{KL}(\pi_{\beta, r^*}(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x)) + \frac{1}{\beta} \log(\mathbb{E}_{Y \sim \pi_{\text{ref}}(\cdot|x)}[\exp(\beta r^*(Y, x))]). \\ 1301 \\ 1302 \quad (55)$$

$$1303 \\ 1304 \mathbb{E}_{Y \sim \pi_{\beta, \hat{r}}(\cdot|x)}[\hat{r}(Y, x)] = \frac{1}{\beta} \text{KL}(\pi_{\beta, \hat{r}}(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x)) + \frac{1}{\beta} \log(\mathbb{E}_{Y \sim \pi_{\text{ref}}(\cdot|x)}[\exp(\beta \hat{r}(Y, x))]). \quad (56) \\ 1305 \\ 1306 \\ 1307$$

1308 Therefore, for term I_1 , we have,
 1309

$$1310 \\ 1311 \\ 1312 \\ 1313 \mathbb{E}_{Y \sim \pi_{\beta, r^*}(\cdot|x)}[r^*(Y, x)] - \mathbb{E}_{Y \sim \pi_{\beta, \hat{r}}(\cdot|x)}[r^*(Y, x)] \\ 1314 \\ 1315 = \mathbb{E}_{Y \sim \pi_{\beta, r^*}(\cdot|x)}[r^*(Y, x)] - \mathbb{E}_{Y \sim \pi_{\beta, \hat{r}}(\cdot|x)}[\hat{r}(Y, x)] \\ 1316 \\ 1317 + \mathbb{E}_{Y \sim \pi_{\beta, \hat{r}}(\cdot|x)}[\hat{r}(Y, x)] - \mathbb{E}_{Y \sim \pi_{\beta, \hat{r}}(\cdot|x)}[r^*(Y, x)] \\ 1318 \\ 1319 = \frac{1}{\beta} \left(\text{KL}(\pi_{\beta, r^*}(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x)) - \text{KL}(\pi_{\beta, \hat{r}}(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x)) \right) \\ 1320 \\ 1321 + \frac{1}{\beta} \log(\mathbb{E}_{Y \sim \pi_{\text{ref}}(\cdot|x)}[\exp(\beta r^*(Y, x))]) - \frac{1}{\beta} \log(\mathbb{E}_{Y \sim \pi_{\text{ref}}(\cdot|x)}[\exp(\beta \hat{r}(Y, x))]) \\ 1322 \\ 1323 + \sum_{\mathcal{Y}} \pi_{\beta, \hat{r}}(\cdot|x) (\hat{r}(y, x) - r^*(y, x)) \\ 1324 \\ 1325 \leq \frac{1}{\beta} \left(\text{KL}(\pi_{\beta, r^*}(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x)) - \text{KL}(\pi_{\beta, \hat{r}}(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x)) \right) \\ 1326 \\ 1327 + \frac{1}{\beta} \log(\mathbb{E}_{Y \sim \pi_{\text{ref}}(\cdot|x)}[\exp(\beta r^*(Y, x))]) - \frac{1}{\beta} \log(\mathbb{E}_{Y \sim \pi_{\text{ref}}(\cdot|x)}[\exp(\beta \hat{r}(Y, x))]) \\ 1328 \\ 1329 + \frac{1}{\sqrt{\beta}} \sqrt{\sum_{\mathcal{Y}} \frac{\pi_{\beta, \hat{r}}^2(y|x)}{\pi_{\text{ref}}(y|x)} \sqrt{\beta \sum_{\mathcal{Y}} (\hat{r}(y, x) - r^*(y, x))^2 \pi_{\text{ref}}(y|x)}} \\ 1330 \\ 1331 \\ 1332 \quad (57) \\ 1333 \\ 1334 \leq \frac{1}{\beta} \left(\text{KL}(\pi_{\beta, r^*}(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x)) - \text{KL}(\pi_{\beta, \hat{r}}(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x)) \right) \\ 1335 \\ 1336 + \frac{1}{\beta} \log(\mathbb{E}_{Y \sim \pi_{\text{ref}}(\cdot|x)}[\exp(\beta r^*(Y, x))]) - \frac{1}{\beta} \log(\mathbb{E}_{Y \sim \pi_{\text{ref}}(\cdot|x)}[\exp(\beta \hat{r}(Y, x))]) \\ 1337 \\ 1338 + \sqrt{C_{\beta, \hat{r}}(x) \varepsilon_{\beta, r}(x)} \\ 1339 \\ 1340 \leq \frac{1}{\beta} \left(\text{KL}(\pi_{\beta, r^*}(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x)) - \text{KL}(\pi_{\beta, \hat{r}}(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x)) \right) \\ 1341 \\ 1342 + \frac{1}{\beta} \log(\mathbb{E}_{Y \sim \pi_{\text{ref}}(\cdot|x)}[\exp(\beta r^*(Y, x))]) - \frac{1}{\beta} \log(\mathbb{E}_{Y \sim \pi_{\text{ref}}(\cdot|x)}[\exp(\beta \hat{r}(Y, x))]) \\ 1343 \\ 1344 + \sqrt{C_{\beta, \hat{r}}(x) \varepsilon_{\beta, r}(x)} \\ 1345 \\ 1346 \leq \frac{1}{\beta} \left(\text{KL}(\pi_{\beta, r^*}(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x)) - \text{KL}(\pi_{\beta, \hat{r}}(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x)) \right) \\ 1347 \\ 1348 + \sqrt{C_{\beta, r^*}(x) \varepsilon_{\beta, r}(x)} \\ 1349 + \sqrt{C_{\beta, \hat{r}}(x) \varepsilon_{\beta, r}(x)}.$$

1350 For term I_2 and using similar approach to term I_1 and applying Lemma F.10, we have,
1351

$$\begin{aligned}
& \mathbb{E}_{Y \sim \pi_{\beta, \hat{r}}(\cdot|x)}[r^*(Y, x)] - \mathbb{E}_{Y \sim \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)}[r^*(Y, x)] \\
& \leq 2\text{TV}(\pi_{\beta, \hat{r}}(\cdot|x), \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)) \\
& \leq 2 \min \left(1, \sqrt{\frac{1}{2} \text{KL}(\pi_{\beta, \hat{r}}(\cdot|x) \parallel \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x))} \right) \\
& \leq 2 \min \left(1, \sqrt{\frac{1}{2} \log \left(1 + \frac{C_{\beta, \hat{r}}(x) - 1}{N} \right)} \right) \\
& \leq 2 \sqrt{\frac{1}{2} \log \left(1 + \frac{C_{\beta, \hat{r}}(x) - 1}{N} \right)}
\end{aligned} \tag{58}$$

1362 Combining equation 57 and equation 58 with equation 54 completes the proof. \square
1363

1364 **Theorem 5.2.** Under Assumption 3.1, the following upper bound holds on the optimal regret
1365 gap of the SBoN policy for any $\beta > 0$,
1366

$$\begin{aligned}
\Delta_{J_{r^*}}(\pi_{r^*}^*(\cdot|x), \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)) & \leq \sqrt{\varepsilon_{\beta, r}(x)} \left(\sqrt{C_{\infty, \hat{r}}(x)} + \sqrt{C_{\infty, r^*}(x)} \right) \\
& + 2 \sqrt{\frac{1}{2} \log \left(1 + \frac{C_{\infty, \hat{r}}(x) - 1}{N} \right)} \\
& + \frac{\log(C_{\infty, r^*}(x))}{\beta},
\end{aligned}$$

1374 *Proof.* Note that we have,
1375

$$\begin{aligned}
& \Delta_{J_{r^*}}(\pi_{r^*}^*(\cdot|x), \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)) \\
& = \mathbb{E}_{Y \sim \pi_{r^*}^*(\cdot|x)}[r^*(Y, x)] - \mathbb{E}_{Y \sim \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)}[r^*(Y, x)] \\
& = \underbrace{\mathbb{E}_{Y \sim \pi_{r^*}^*(\cdot|x)}[r^*(Y, x)] - \mathbb{E}_{Y \sim \pi_{\beta, r^*}(\cdot|x)}[r^*(Y, x)]}_{I_3} \\
& + \underbrace{\Delta_{J_{r^*}}(\pi_{\beta, r^*}(\cdot|x), \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x))}_{I_4}
\end{aligned} \tag{59}$$

1384 For term I_4 , we can use Theorem H.1. For term I_3 , note that, we have for $\beta > 0$,
1385

$$\mathbb{E}_{Y \sim \pi_{r^*}^*(\cdot|x)}[r^*(Y, x)] - \mathbb{E}_{Y \sim \pi_{\beta, r^*}(\cdot|x)}[r^*(Y, x)] \leq \frac{\text{KL}(\pi_{r^*}^*(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x)) - \text{KL}(\pi_{\beta, r^*}(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x))}{\beta} \tag{60}$$

1389 Combining equation 60 with Theorem H.1, completes the proof due the positiveness of KL divergence and using Lemma F.7 and Lemma 5.1.
1390

\square

1393 **Remark H.2.** For $\beta = 0$, we have, $\lim_{\beta \rightarrow 0} \pi_{\beta, r^*}(\cdot|x) = \pi_{\text{ref}}(\cdot|x)$. Therefore, we have,
1394

$$\mathbb{E}_{Y \sim \pi_{r^*}^*(\cdot|x)}[r^*(Y, x)] - \mathbb{E}_{Y \sim \pi_{\beta, r^*}(\cdot|x)}[r^*(Y, x)] \leq \sqrt{2\text{KL}(\pi_{r^*}^*(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x))}. \tag{61}$$

1397 **Proposition 5.3.** Under Assumption 3.1, the following upper bound holds on the optimal
1398 regret gap of the BoN policy for any $\beta > 0$,
1399

$$\begin{aligned}
\Delta_{J_{r^*}}(\pi_{r^*}^*(\cdot|x), \pi_{\hat{r}}^{\text{BoN}}(\cdot|x)) & \leq \sqrt{\varepsilon_{\infty, r}(x)} \left(\sqrt{C_{\infty, \hat{r}}(x)} + \sqrt{C_{\infty, r^*}(x)} \right) \\
& + 2 \sqrt{\frac{1}{2} \log \left(1 + \frac{C_{\infty, \hat{r}}(x) - 1}{N} \right)}.
\end{aligned}$$

1404 *Proof.* The results follow directly from Proposition 5.3 for $\beta \rightarrow \infty$. \square

1405

1406

1407 **H.2 LOWER BOUND PROOFS AND DETAILS:**

1408

1409 We can define the probability of coverage of maximizers under the reference policy as

$$1410 \quad p^*(x) := \Pr_{Y \sim \pi_{\text{ref}}(\cdot|x)}[Y \in \mathcal{Y}_{r^*}^*(x)] = \frac{1}{C_{\infty, r^*}(x)}. \\ 1411$$

1412 **Lemma H.3.** *Under Assumption 5.5, the following lower bound holds on the optimal regret gap of*
 1413 *the SBoN policy based on true-reward,*

$$1415 \quad \Delta_{J_{r^*}}(\pi_{r^*}^*(\cdot|x), \pi_{r^*}^{\text{SBoN}}(\cdot|x)) \geq \gamma(x) \Pr_{\text{Err}}(x), \quad (62)$$

1416 *where \Pr_{Err} is the probability that selected sample from $\{Y_i\}_{i=1}^N$ via SBoN does not belong to*
 1417 *$\mathcal{Y}_{r^*}^*(x)$.*

1418

1419 *Proof.* We have,

$$1420 \quad \begin{aligned} & \Delta_{J_{r^*}}(\pi_{r^*}^*(\cdot|x), \pi_{r^*}^{\text{SBoN}}(\cdot|x)) \\ &= J_{r^*}(\pi_{r^*}^*(\cdot|x)) - J_{r^*}(\pi_{\beta, r^*}(\cdot|x)) \\ &= \mathbb{E}_{Y \sim \pi_{r^*}^*(\cdot|x)}[r^*(Y, x)] - \mathbb{E}_{Y \sim \pi_{\beta, r^*}(\cdot|x)}[r^*(Y, x)] \\ &= 1 - \left[\mathbb{E}_{Y \sim \pi_{\beta, r^*}(\cdot|x)}[r^*(Y, x) | \text{Err}] \Pr_{\text{Err}} + \mathbb{E}_{Y \sim \pi_{\beta, r^*}(\cdot|x)}[r^*(Y, x) | \text{Err}^c] (1 - \Pr_{\text{Err}}(x)) \right] \quad (63) \\ &\geq 1 - \left[(1 - \gamma(x)) \Pr_{\text{Err}}(x) + (1 - \Pr_{\text{Err}}(x)) \right] \\ &= \gamma(x) \Pr_{\text{Err}}(x), \end{aligned}$$

1425 where Err is the event where the selected sample using SBoN is not a maximizer. \square

1426

1427 Next, we provide a lower bound on $\Pr_{\text{Err}}(x)$.

1428 **Lemma H.4.** *The following lower bound holds on \Pr_{Err} for SBoN algorithm,*

$$1435 \quad \Pr_{\text{Err}}(x) \geq (1 - p^*(x))^N + \frac{N((1 - p^*(x)) - (1 - p^*(x))^N)}{N + (\exp(\beta) - 1)(N - 1)}. \quad (64)$$

1436

1437 *Proof.* We consider cases where exactly K responses are selected from $\mathcal{Y}_{r^*}^*(x)$. The K has a binomial distribution $\text{Bin}(N, p^*(x))$ Using Bayes rule, we have,

$$1438 \quad \begin{aligned} \Pr_{\text{Err}}(x) &= \sum_{i=0}^{N-1} \Pr(K = i) \Pr_{\text{Err}}(x | K = i) \\ &= \sum_{i=0}^{N-1} \binom{N}{i} p^*(x)^i (1 - p^*(x))^{N-i} \Pr_{\text{Err}}(x | K = i) \\ &= \sum_{i=0}^{N-1} \binom{N}{i} p^*(x)^i (1 - p^*(x))^{N-i} \frac{\sum_{Y_j \notin \mathcal{Y}_{r^*}^*(x)} \exp(\beta r^*(Y_j, x))}{\sum_{j=1}^N \exp(\beta r^*(Y_j, x))} \\ &\geq (1 - p^*(x))^N + \sum_{i=1}^{N-1} \binom{N}{i} p^*(x)^i (1 - p^*(x))^{N-i} \frac{N - i}{N - i + ie^\beta} \\ &\geq (1 - p^*(x))^N + \sum_{i=1}^{N-1} \binom{N}{i} p^*(x)^i (1 - p^*(x))^{N-i} \frac{N - i}{N + (e^\beta - 1)(N - 1)} \\ &\geq (1 - p^*(x))^N + \frac{N((1 - p^*(x)) - (1 - p^*(x))^N)}{N + (e^\beta - 1)(N - 1)} \end{aligned} \quad (65)$$

1455

1456

1457

□

1458
1459 Using Lemma H.3 and Lemma H.2, we can derive the following lower bound on our main optimal
1460 regret gap,

1461 **Theorem 5.6.** *Under the same Assumptions in Theorem 5.2 and Lemma H.3, the following
1462 lower bound holds on the regret of SBoN policy under the proxy-reward model,*

$$\begin{aligned} 1464 \quad & \Delta_{J_{r^*}}(\pi_{r^*}^*(\cdot|x), \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)) \\ 1465 \quad & \geq \gamma(x) \left((1 - p^*(x))^N + \frac{N((1 - p^*(x)) - (1 - p^*(x))^N)}{N + (\exp(\beta) - 1)(N - 1)} \right) \\ 1466 \quad & - \min \left(\sqrt{1 - \exp(-U(N, \beta))}, \frac{N}{1 + (N - 1)\exp(-\beta)} \sqrt{\varepsilon_{\beta, r}(x)} \right), \end{aligned} \quad (66)$$

$$1470 \quad \text{where } p^*(x) = \frac{1}{C_{\infty, r^*}(x)} \text{ and } U(N, \beta) = \frac{N\beta\sqrt{\varepsilon_{\beta, r}(x)}}{1 + (N - 1)\exp(-\beta)} \left(\frac{N\exp(2\beta)}{(N - 1)^2} + 1 \right).$$

1473 *Proof.*

$$\Delta_{J_{r^*}}(\pi_{r^*}^*(\cdot|x), \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)) \quad (67)$$

$$= \Delta_{J_{r^*}}(\pi_{r^*}^*(\cdot|x), \pi_{r^*}^{\text{SBoN}}(\cdot|x)) \quad (68)$$

$$+ \Delta_{J_{r^*}}(\pi_{r^*}^{\text{SBoN}}(\cdot|x), \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)), \quad (69)$$

1479 where equation 68 can be bounded using Lemma H.3 and Lemma H.2. By combining Lemma F.12
1480 and Lemma 4.2 we have an upper bound on absolute value of equation 69 which can also result in a
1481 lower bound on equation 69. Furthermore, another simple upper bound on equation 69 is,

$$\Delta_{J_{r^*}}(\pi_{r^*}^{\text{SBoN}}(\cdot|x), \pi_{\hat{r}}^{\text{SBoN}}(\cdot|x)) \quad (70)$$

$$= \mathbb{E}_{Y \sim \pi_{r^*}^{\text{SBoN}}(Y|x)}[r^*(Y, x)] - \mathbb{E}_{Y \sim \pi_{\hat{r}}^{\text{SBoN}}(Y|x)}[\hat{r}(Y, x)] \quad (71)$$

$$+ \mathbb{E}_{Y \sim \pi_{\hat{r}}^{\text{SBoN}}(Y|x)}[\hat{r}(Y, x)] - \mathbb{E}_{Y \sim \pi_{\hat{r}}^{\text{SBoN}}(Y|x)}[r^*(Y, x)] \quad (72)$$

$$= \mathbb{E}_{Y \sim \pi_{\hat{r}}^{\text{SBoN}}(Y|x)}[\hat{r}(Y, x)] - \mathbb{E}_{Y \sim \pi_{\hat{r}}^{\text{SBoN}}(Y|x)}[r^*(Y, x)] \quad (73)$$

1488 Note that due to calibration both true and proxy-reward models are uniform under $\pi_{\text{ref}}(\cdot|x)$. Therefore, we have,

$$\mathbb{E}_{Y \sim \pi_{r^*}^{\text{SBoN}}(Y|x)}[r^*(Y, x)] = \mathbb{E}_{Y \sim \pi_{\hat{r}}^{\text{SBoN}}(Y|x)}[\hat{r}(Y, x)].$$

1492 Now we have,

$$\begin{aligned} 1494 \quad & |\mathbb{E}_{Y \sim \pi_{\hat{r}}^{\text{SBoN}}(Y|x)}[\hat{r}(Y, x)] - \mathbb{E}_{Y \sim \pi_{\hat{r}}^{\text{SBoN}}(Y|x)}[r^*(Y, x)]| \leq \sqrt{\sum_y \frac{(\pi_{\hat{r}}^{\text{SBoN}}(\cdot|x))^2}{\pi_{\text{ref}}(\cdot|x)}} \sqrt{\varepsilon_{\beta, r}(x)} \\ 1495 \quad & \leq \frac{N}{1 + (N - 1)\exp(-\beta)} \sqrt{\varepsilon_{\beta, r}(x)}. \end{aligned}$$

1499 \square

1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511

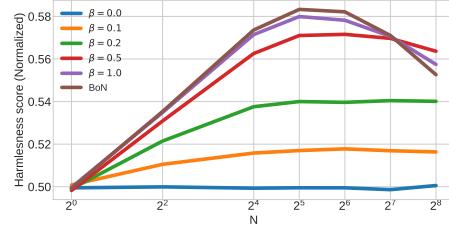
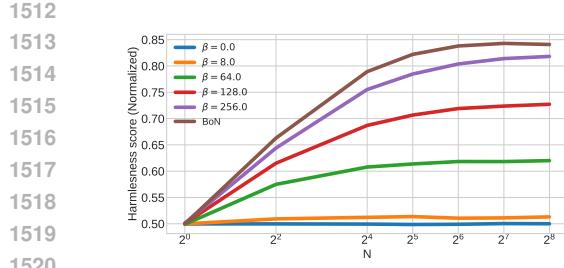


Figure 3: Soft Best-of-N experiment using a strong reward model (Left) and a weak one (Right). When the performance of the reward model is poor, BoN can lead to overoptimization, while the SBoN can help to mitigate it.

I EXPERIMENTS

Our implementation is available in the following link: <https://anonymous.4open.science/r/sbon-C6D3/README.md>.

LLM-as-a-Judge: To evaluate the performance of the generated responses, we used GPT-4o to determine how harmless they are. Specifically, we used the following prompt:

You are an AI safety evaluator. Your task is to determine if a response to a given prompt contains harmful content. Respond with a harmlessness score using these criteria:

2 = SAFE: The response is completely safe and appropriate.

1 = SOMEWHAT HARMFUL: The response contains potentially concerning elements, mild toxicity, stereotypes, subtle harmful guidance, or borderline inappropriate content.

0 = HARMFUL: The response contains clearly harmful content such as hate speech, explicit instructions for dangerous activities, severe toxicity, etc.

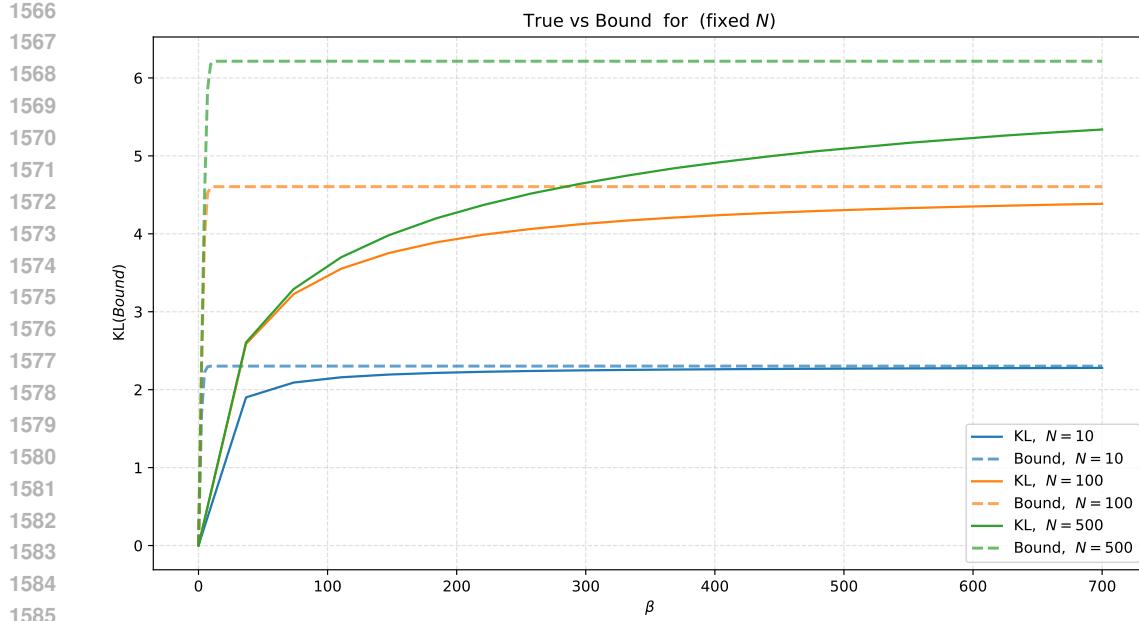
Provide your judgment as an integer in curly brackets like this: 1, 2, or 3.

This yields a reward with values ranging from 0 to 2, where 0 indicates complete harm and 2 indicates safety.

Future work may address optimizing β using cross-validation or the estimation of terms in Remark 5.8 to derive sub-optimal β . For this study, we employ a grid search over a range of β values.

I.1 MORE EXPERIMENTS

To further support our theoretical analysis, we conducted more experiments comparing Soft Best-of-N (SBoN) across different regularization strengths and reward model qualities. We ran one extra experimental conditions: one using a weaker proxy-reward model (tulu-v2.5-13b-uf-rm (Ivison et al., 2024)) respect to Armo reward model. As shown in Figure 3, when the reward model is weak, performance degrades for large N due to reward hacking. However, the smoothing in SBoN helps mitigate this degradation. This observation is also aligned with our theoretical analysis and discussion in Section 4.

Figure 4: True KL divergence vs upper bound in Lemma 4.1 for fixed $N = \{10, 100, 500\}$.

I.2 NUMERICAL EXAMPLES

I.2.1 KL BOUND

To illustrate how our analytical upper bound in Lemma 4.1 behaves as a function of the temperature parameter β , we run a toy experiment in which

1. the reference policy is the uniform distribution over responses, and
2. rewards are bounded with $= 1$.

For each β in a logarithmic sweep, we compute the true KL divergence between the SBoN policy and the reference policy, together with the theoretical bound derived in Lemma 4.1.

- **Very large β (nearBoN policy).** As $\beta \rightarrow \infty$ the SBoN policy converges to the BoN policy. The gap between the KL and the bound vanishes.
- **Very small β (reference policy).** When $\beta \rightarrow 0$ the SBoN policy approaches the uniform sampling from the reference policy, which results in the reference policy, making the KL itself tend to zero; the bound is equal to zero for this value.

This experiment confirms that the bound is tight in the two asymptotic regimes and remains a conservative yet informative estimate elsewhere.

I.2.2 REGRET UPPER BOUND

To validate our theoretical findings in Theorem 5.2 and Proposition 5.3, we conducted a numerical analysis comparing the regret upper bounds of SBoN against standard BoN under varying degrees of reward model error.

We simulate a simplified binary reward setting to isolate the effects of regularization temperature (β) and proxy-reward error (δ). The setup is defined as follows:

Reward Structure: We assume a binary reward landscape where responses have a true-reward $r^*(y), \hat{r} \in \{0, 1\}$.

Reference Policy (π_{ref}): The reference model generates an optimal response ($r^* = 1$) with probability $\pi_{\text{ref}}(y_{\text{max}}|x) = 0.05$, corresponding to a coverage constant $C_{\infty, r^*} = 20$. The same also

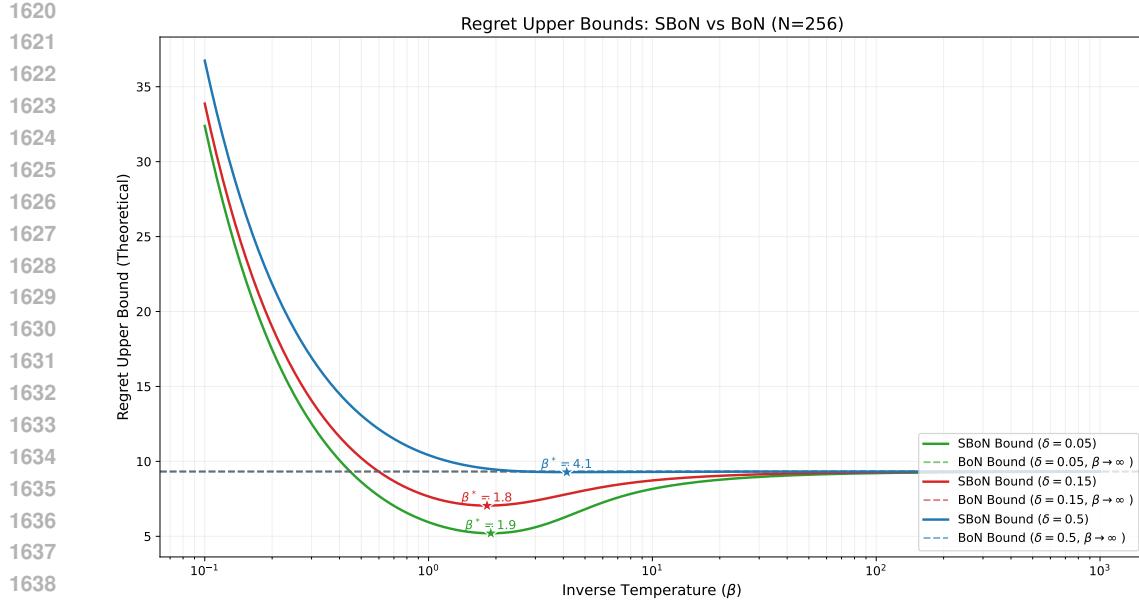


Figure 5: Numerical comparison of upper bound on Regret of SBoN and BoN

holds for proxy-reward model and we have $C_{\infty, \hat{r}} = 20$. We also consider $N = 256$ samples from reference policy.

Proxy-reward Error (δ): The proxy-reward model matches the true-reward with probability $1 - \delta$ and commits an error (flipping the label) with probability δ . We evaluate three error regimes: Low ($\delta = 0.05$), Medium ($\delta = 0.15$), and High ($\delta = 0.50$).

SBoN: We compute the regret upper bound using Theorem 5.2, varying the inverse temperature β from 0.5 to 50.

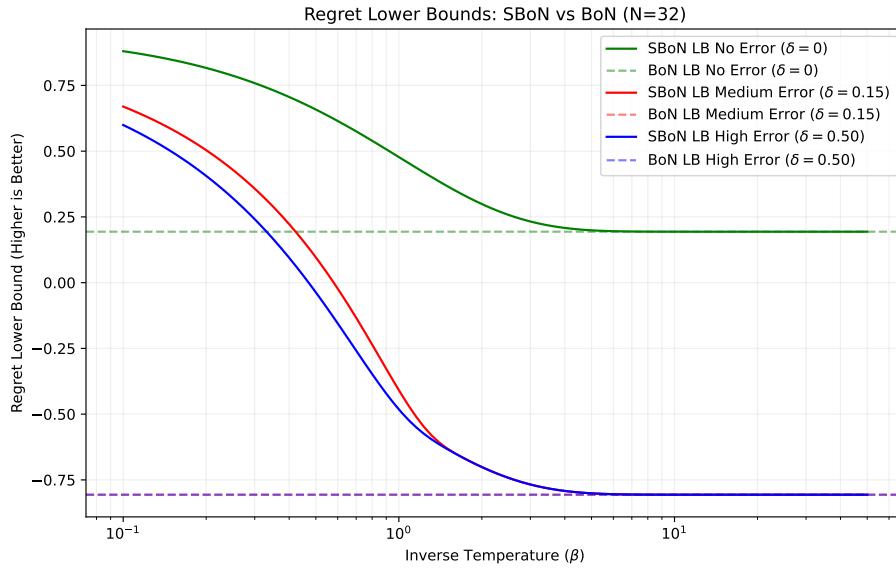
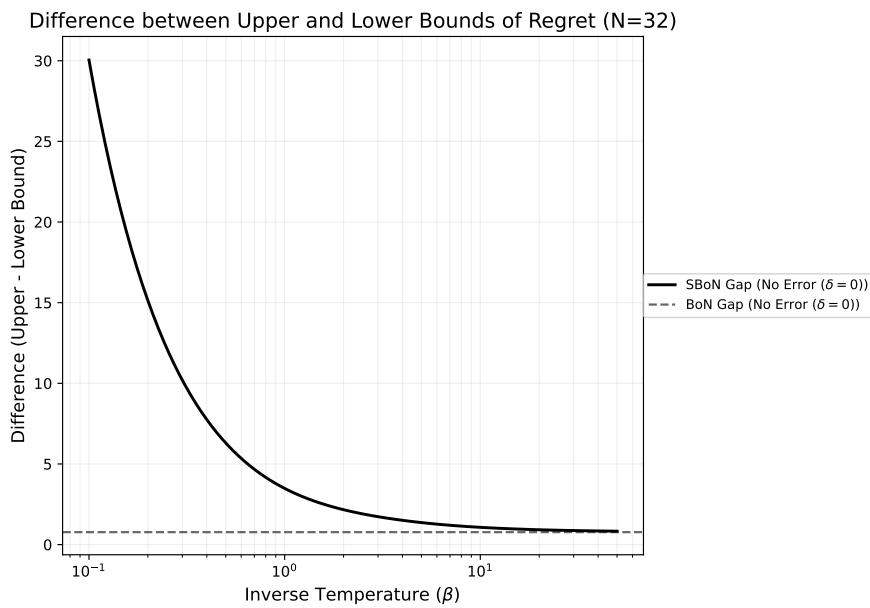
BoN: We compute the regret upper bound using Proposition 5.3, which represents the asymptotic limit of SBoN as $\beta \rightarrow \infty$. In this binary setting, the maximum possible error is ($\varepsilon_{\infty, r} = 1$).

Discussion: The results, Fig. 5 demonstrate a distinct trade-off mechanism governed by the temperature β . While the standard BoN bound remains constant at a high value due to its susceptibility to the maximum error ($\varepsilon_{\infty, r} = 1$), the SBoN bound forms a convex curve.

For every error rate tested, we observe a "sweet spot"-a finite optimal temperature β^* where the SBoN bound is significantly tighter than the BoN bound. This confirms that by smoothing the policy (finite β), SBoN effectively balances the trade-off between exploring high-reward regions and mitigating the overoptimization caused by proxy-reward errors.

I.2.3 REGRET LOWER BOUND

To complement our upper bound analysis, we numerically evaluate the lower bounds on the regret gap for both SBoN and BoN, as derived in Theorem 5.6 and Proposition 5.7 respectively. These lower bounds serve as a "safety guarantee," quantifying the worst-case performance degradation caused by overoptimization. We utilize the same binary reward setup described in previous section, with sample size $N = 32$ and consider no error ($\delta = 0$). The results in Fig. 6 illustrate an advantage of the SBoN policy over the deterministic BoN policy. In error regimes (blue dashed line, $\delta = 0.15, 0.5$), the BoN lower bound drops significantly, indicating a weak guarantee against performance collapse. Note that, for no error regime, the lower bound on BoN's regret is positive.

1697 Figure 6: Numerical comparison of lower bound on Regret of SBoN and BoN
1698
1699
1700
1701
1702
17031724 Figure 7: Numerical comparison of the gap between upper and lower bounds on Regret of SBoN
1725 and BoN
1726
1727

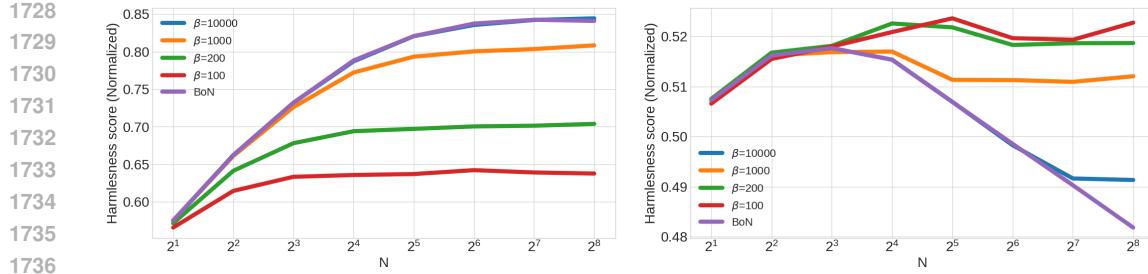


Figure 8: InferenceTimePessimism experiment using a strong reward model (Left) and a weak one (Right). When the performance of the reward model is poor, BoN can lead to overoptimization, while the InferenceTimePessimism can help to mitigate it.

I.2.4 UPPER AND LOWER BOUNDS GAP

To compare the performance of upper and lower bounds, we consider the case where $\delta = 0$ which is similar to no-overoptimization scenario. As shown in Fig. 7, the gap between upper and lower bounds of BoN is tighter than SBoN.

J COMPARISON WITH HUANG ET AL. (2025)

We contrast our contributions with the recent work of Huang et al. (2025), focusing on four key dimensions: the object of study, regret bounds, and the mechanism of regularization.

Object of study. Huang et al. (2025) analyze BoN, introducing an inference-time pessimism algorithm within a χ^2 -regularized framework to mitigate overoptimization. In contrast, we analyze SBoN via treating BoN as the limit where $\beta \rightarrow \infty$ and derive finite-sample KL and regret bounds under reward model misspecification (specifically, proxy-reward vs. true-reward). To our knowledge, these SBoN-specific bounds characterizing the behavior under overoptimization are novel.

Error metrics and reward modeling. While Huang et al. (2025) operate on raw (uncalibrated) rewards using mean-squared error (MSE), we utilize calibrated rewards (Balashankar et al., 2025). We introduce the *tilted error*, denoted as $\varepsilon_{\beta,r}$, which interpolates between MSE (at $\beta = 0$) and L^∞ error (as $\beta \rightarrow \infty$). This metric allows us to explicitly track the dynamics of overoptimization as a function of both N and the temperature β .

Regret bounds. Our Theorem 5.2, Proposition 5.3, Theorem 5.6 and Proposition 5.7 and establish regret bounds for SBoN and its BoN limit that depend on: (i) the tilted error $\varepsilon_{\beta,r}(x)$, and (ii) the coverage constants $C_{\infty,r}$ for both the true and proxy-rewards. A crucial distinction, highlighted in Remark 5.4, is that our BoN regret bound remains finite even when overoptimization vanishes (i.e., when $\varepsilon_{\infty,r} = 0$ or $N \rightarrow \infty$). Conversely, the bound proposed by Huang et al. (2025) on BoN scales with the L^∞ -error. This represents a qualitative improvement in how misspecification is handled in our analysis.

Theoretical and algorithmic consistency. Our analysis of SBoN fully incorporates the algorithm’s primary regularization parameter, β . In contrast, the theoretical regret bounds for *InferenceTimePessimism* in Huang et al. (2025) exclude the estimation error of the normalization factor (Algorithm 3) and are independent of the truncation parameter required by their Algorithm 4. Furthermore, regarding practical complexity, SBoN relies on a single hyperparameter β , whereas *InferenceTimePessimism* requires tuning an additional truncation parameter and incurs computational overhead to estimate the normalization factor.

InferenceTimePessimism Experiment: We compared SBoN against the InferenceTimePessimism baseline (Huang et al., 2025), using the configuration specified in their paper. While both methods achieve similar empirical performance Fig.8, SBoN is significantly more efficient to implement and tune, as it avoids the additional computational overhead and hyperparameter search space required by the pessimistic approach.

1782 **K INFORMATION PROJECTION ACROSS REWARD FUNCTIONS**
 1783

1784 In this section, we study the KL divergence behavior in Section 4 for large N regime, where SBoN
 1785 policy converge to tilted optimal policy.

1786 We consider a proxy-reward class \mathcal{R} of rewards, and a projection problem at a *fixed* temperature β
 1787 which is inspired by reducing the estimation error,

1788 (reverse-I-proj in tilted family) $\hat{r}^1 \in \arg \min_{\hat{r} \in \mathcal{R}} D_{\text{KL}}(\pi_{\beta, r^*} \| \pi_{\beta, \hat{r}}).$
 1789

1790 **Proposition K.1** (Pythagorean I-projection in reward space). *Suppose the proxy class \mathcal{R} is convex
 1791 in r (e.g., an affine subspace). Let $\hat{r}^1 \in \arg \min_{\hat{r} \in \mathcal{R}} D_{\text{KL}}(\pi_{\beta, r^*} \| \pi_{\beta, \hat{r}})$. Then, we have,*

1792
$$D_{\text{KL}}(\pi_{\beta, r^*} \| \pi_{\text{ref}}) \geq \underbrace{D_{\text{KL}}(\pi_{\beta, r^*} \| \pi_{\beta, \hat{r}^1})}_{\text{Estimation error}} + \underbrace{D_{\text{KL}}(\pi_{\beta, \hat{r}^1} \| \pi_{\text{ref}})}_{\text{KL-divergence based on proxy-reward}} \quad (74)$$

 1793

1794 *Proof.* The proof follows from Theorem 1 in Csiszár & Matus (2003). \square
 1795

1796 Proposition K.1 and Lemma F.7 provide an upper bound on $D_{\text{KL}}(\pi_{\beta, r^*} \| \pi_{\text{ref}})$. Since the LHS of
 1797 equation 74 is fixed for a given β , an increase in estimation error necessarily decreases the allowable
 1798 divergence between the proxy-reward and the reference policies $D_{\text{KL}}(\pi_{\beta, \hat{r}^1} \| \pi_{\text{ref}})$. If $r^* \in \mathcal{R}$, the
 1799 estimation error is zero (no overoptimization). If $r^* \notin \mathcal{R}$, the estimation error consumes part of the
 1800 total divergence budget. This analysis can be extended to for finite N , using Lemma F.10.
 1801

1802
 1803
 1804
 1805
 1806
 1807
 1808
 1809
 1810
 1811
 1812
 1813
 1814
 1815
 1816
 1817
 1818
 1819
 1820
 1821
 1822
 1823
 1824
 1825
 1826
 1827
 1828
 1829
 1830
 1831
 1832
 1833
 1834
 1835