

# PROVABLE BENEFIT OF CURRICULUM IN TRANSFORMER TREE-REASONING POST-TRAINING

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Recent curriculum techniques in the post-training stage of LLMs have been widely observed to outperform non-curriculum approaches in enhancing reasoning performance, yet a principled understanding of why and to what extent they work remains elusive. To address this gap, we develop a theoretical framework grounded in the intuition that progressively learning through manageable steps is more efficient than directly tackling a hard reasoning task, provided each stage stays within the model’s effective competence. Under mild complexity conditions linking consecutive curriculum stages, we show that curriculum post-training avoids the exponential complexity bottleneck. To substantiate this result, drawing insights from the Chain-of-Thoughts (CoTs) solving mathematical problems such as Countdown and parity, we model CoT generation as a states-conditioned autoregressive reasoning tree, define a uniform-branching base model to capture pretrained behavior, and formalize curriculum stages as either depth-increasing (longer reasoning chains) or hint-decreasing (shorter prefixes) subtasks. Our analysis shows that, under outcome-only reward signals, reinforcement learning fine-tuning achieve high accuracy with polynomial sample complexity, whereas direct learning suffers from an exponential bottleneck. We further establish analogous guarantees for test-time scaling, where curriculum-aware querying reduces both reward oracle calls and sampling cost from exponential to polynomial order.

## 1 INTRODUCTION

When a high-school student first encounters a problem in Riemannian geometry, it is unrealistic to start with a standard textbook such as Do Carmo & Flaherty Francis (1992). Instead, one typically progresses step by step through prerequisite subjects such as algebra, analysis, and geometry, gradually developing the reasoning skills required. A similar principle underlies one of the most effective post-training strategies for foundation models: enhancing mathematical reasoning by encouraging Chain-of-Thought (CoT) generations to evolve from easy to hard, a process known as Curriculum-style Post-training (Liu et al., 2025a; Lee et al., 2025; Bae et al., 2025; Meng et al., 2025; Shi et al., 2025; Wen et al., 2025b; Zhang et al., 2025; Amani et al., 2025; Zhou et al., 2025).

Despite empirical progress, the value of curriculum strategies in post-training remains largely intuitive and lacks rigorous theoretical grounding. Curriculum strategies, both implicit and explicit, have also been explored in pre-training where models are trained from scratch (Bengio et al., 2009; Graves et al., 2017; Narvekar et al., 2020; Wang et al., 2021; Soviany et al., 2022), and some theoretical results exist in specialized function classes such as convex regression and bi-classification (Weinshall et al., 2018; Weinshall & Amir, 2020), parity (Abbe et al., 2023c; Panigrahi et al., 2024), teacher–student perceptrons with sparse features (Saglietti et al., 2022), and  $k$ -fold composition (Wang et al., 2025). However, these results are highly problem-specific: their notions of “difficulty” and “performance,” their handcrafted algorithms, and their proof techniques are tailored to their specific train-from-scratch settings. They cannot be directly applied to transformer-based post-training reasoning, which is characterized by a pretrained base model and explicit CoT generalization. This motivates the central research question:

*When, why, how, and in what sense can curriculum post-training strategies (gradually increasing reasoning difficulty) theoretically improve performance compared to direct learning?*

To bridge this gap, we observe that in post-training, “performance” is naturally measured by the probability of generating the correct CoT, while “difficulty” is often quantified by the *pass-rate*, i.e., the probability that the base model already produces the correct CoT (Tong et al., 2024; Parashar et al., 2025). A principled formalization of this “difficulty” uses the coverage coefficient Foster et al. (2025)  $\|\frac{\pi^*}{\pi_{\text{ref}}}\|_{\infty} := \sup_x \|\frac{d\pi^*(\cdot|x)}{d\pi_{\text{ref}}(\cdot|x)}\|_{L^{\infty}(\pi_{\text{ref}}(\cdot|x))} < \infty^1$  (assuming absolute continuity), which bounds how much rarer a correct CoT can be under the base model’s predictive policy  $\pi_{\text{ref}}$  compared to the correct CoT policy  $\pi^*$ , and thus inversely control the *pass-rate*. In this view, a curriculum corresponds to intermediate policies  $\pi_0^* = \pi_{\text{ref}}, \pi_1^*, \dots, \pi_K^* = \pi^*$ , with the “difficulty”  $\|\frac{\pi_k^*}{\pi_{\text{ref}}}\|_{\infty}$  nondecreasing in  $k \in [K]$ .

The next question is: why should stepwise curriculum estimation (i.e., sequentially matching  $\pi_1^*, \dots, \pi_K^*$ ) be more effective than directly learning  $\pi^* = \pi_K^*$  from  $\pi_{\text{ref}}$ ? Standard rejection-sampling shows that to sample  $\pi_k^*$  from  $\pi_{\text{ref}}$ , one needs  $\Theta(\|\frac{\pi_k^*}{\pi_{\text{ref}}}\|_{\infty} \log(\delta^{-1}))$  trials to obtain at least one correct CoT with confidence no less than  $1 - \delta$  (Block & Polyanskiy, 2024). Directly targeting  $\pi^* = \pi_K^*$  thus costs  $\Theta(\|\frac{\pi_K^*}{\pi_{\text{ref}}}\|_{\infty} \log(\delta^{-1}))$ , whereas a stepwise curriculum costs  $\sum_{k=0}^{K-1} \Theta(\|\frac{\pi_{k+1}^*}{\pi_k^*}\|_{\infty} \log(\delta^{-1}))$ . Hence curriculum shines out if  $\sum_{k=0}^{K-1} \|\frac{\pi_{k+1}^*}{\pi_k^*}\|_{\infty} \leq \|\frac{\pi_K^*}{\pi_{\text{ref}}}\|_{\infty}$  (assuming **(A1: Realizability)**:  $\|\frac{\pi_{k+1}^*}{\pi_k^*}\|_{\infty}, \|\frac{\pi_K^*}{\pi_{\text{ref}}}\|_{\infty} < \infty$ ).

Returning to the high-school student analogy: manageable, stepwise difficult courses promote learning efficiency, where appropriately challenging curricula should let the student master the reasoning thinking for Riemannian geometry best within several semesters. In our setting, assuming **(A2: Suitable Step-wise Difficulty)**:  $\|\frac{\pi_{k+1}^*}{\pi_k^*}\|_{\infty} \leq \Theta(C^p), \forall k$  for some  $p \ll K$  and constant  $C > 1$ , and **(A3: Intensive Direct Difficulty)**:  $\|\frac{\pi_K^*}{\pi_{\text{ref}}}\|_{\infty} \geq \Theta(C^{bK-c})$  for the same constant  $C > 1$  with  $1 \leq b \ll K, c \ll bK$ , curriculum post-training can thus achieve an exponential improvement from  $\Theta(C^{bK-c})$  to  $\Theta(KC^p)$ . This intuition is formalized in Thm. 1 and its follow-up discussion.

The natural question is when assumptions **(A1)–(A3)** hold in post-training reasoning. For **(A1)**, Foster et al. (2025) theoretically proved the necessity of base-model coverage (i.e., the probability of sampling the correct CoT), both for RL fine-tuning and reward-guided test-time scaling in linear realizable MDPs. Large-scale empirical evidence further shows that effective RL fine-tuning with sparse rewards merely *reinforces* existing tree-like reasoning patterns in the foundation model, and cannot generate CoTs with zero initial probability (Gandhi et al., 2025; Yue et al., 2025; Wu et al., 2024), underscoring the critical role of the pretrained base model and supporting our **(A1)**.


For **(A2)–(A3)**, we prove that in a state-conditioned autoregressive reasoning process (Def. 2) subsuming many reasoning tasks (Kim et al., 2025c; Nichani et al., 2025; 2024; Gandhi et al., 2024), assumptions **(A2)–(A3)** naturally hold in terms of the sample complexity in RL fine-tuning (Thm. 3), as well as reward-oracle complexity or computational complexity in test-time scaling (Thm. 4).

**Contributions.** We summarize our contributions as follows.

- We formalize curriculum post-training and establish a general bottleneck theorem (Thm. 1 in Sec. 2), showing that under mild coverage and complexity-alignment assumptions, stepwise curriculum post-training converts exponential depth dependence into polynomial order; we further clarify the role of base-model coverage and relax conditions.
- We instantiate the theory with an autoregressive reasoning tree (2S-ART), encompassing many graph-reasoning problems, and its transformer implementation (Sec. 2.1).
- We prove that curriculum strategies—both in RL fine-tuning and test-time scaling—achieve exponential-to-polynomial reductions in sample complexity and oracle-query complexity. We also discuss the limitations, extensions beyond our assumptions, broader applicability to diverse task families, and future directions in Sec. 3.

We delayed the detailed discussions of additional related work in App. A.

<sup>1</sup>i.e., the  $\ell_{\infty}$ -norm of the Radon–Nikodym derivative of  $\pi^*$  w.r.t.  $\pi_{\text{ref}}$ . If  $\mathcal{X} \times \mathcal{Y}$  are discrete space, it holds that  $\|\frac{\pi^*}{\pi_{\text{ref}}}\|_{\infty} = \sup_{x \in \mathcal{X}, y \in \mathcal{Y}} \frac{\pi^*(y|x)}{\pi_{\text{ref}}(y|x)}$ ,



116 Figure 1: An illustration in Liu et al. (2025b) of the Chain-of-Thought for Countdown game, where  
117 the goal is to obtain 24 by applying basic arithmetic operations (+, −, ×, ÷) ( $\Phi_l(\cdot, \cdot)$  in our Def. 2)  
118 between the current step’s number (e.g., 13, 65 or 72) and some unused number (e.g., 5, 7 or 3) in  
119  $\{3, 5, 7, 13\}$ , targeting 24 as the final outcome. Per Parashar et al. (2025), the difficulty measure of  
120 Countdown is the number of arithmetic operations required to solve an instance.

## 122 2 THEORETICAL FRAMEWORK FOR CURRICULUM POST-TRAINING ON 123 REASONING TREES 124

125 **Preliminaries and Notations.** For each prompt  $x$ , a policy  $\pi$  is a conditional probability mea-  
126 sure  $\pi(\cdot | x)$  over the output space  $\mathcal{O}$ , where  $o \in \mathcal{O}$  denotes a CoT trajectory. For two policies  
127  $\pi$  and  $\pi'$ , if  $\pi'(\cdot | x) \ll \pi(\cdot | x)$  for all  $x$ , we write  $\|\frac{\pi'}{\pi}\|_\infty := \sup_{x \in \mathcal{X}} \left\| \frac{d\pi'(\cdot|x)}{d\pi(\cdot|x)} \right\|_{L^\infty(\pi(\cdot|x))} =$   
128  $\sup_{x \in \mathcal{X}} \text{ess sup}_{o \sim \pi(\cdot|x)} \frac{d\pi'(\cdot|x)}{d\pi(\cdot|x)}(o)$ . Landau symbols follow standard conventions. For  $\mathbf{a} \in \mathbb{R}^m$ ,  
129  $\text{softmax}(\mathbf{a})_i := \exp(a_i) / \sum_j \exp(a_j)$ , and with temperature  $\beta > 0$  we write  $\text{softmax}(\mathbf{a}/\beta)$ . Fi-  
130 nally,  $\text{ReLU}(t) := \max\{0, t\}$  elementwise.

131 As discussed in the introduction, pass rate is the operative notion of difficulty in post-training rea-  
132 soning (Tong et al., 2024; Parashar et al., 2025), and the inverse of  $\|\frac{\pi^*}{\pi_{\text{ref}}}\|_\infty$  lower-bounds the pass  
133 rate and thus serves as a natural proxy for difficulty to master the target task. We formalize the  
134 discussions in the introduction as the theorem below.

135 **Theorem 1.** Consider a curriculum of  $K$  tasks  $\pi_0^*, \pi_1^*, \dots, \pi_K^*$  with base model  $\pi_{\text{ref}} := \pi_0^*$  and  
136 target task  $\pi^* := \pi_K^*$ . Denote the learning complexity<sup>2</sup> from  $\pi_k^*$  to  $\pi_{k'}^*$  as  $\mathcal{C}(\pi_{k'}^* | \pi_k^*)$ . For any  
137  $0 \leq k < k' \leq K$ , suppose  $\|\frac{\pi_{k'}^*}{\pi_k^*}\|_\infty < \infty$  and

$$138 \|\frac{\pi_{k+1}^*}{\pi_k^*}\|_\infty = \Theta(C^*), \quad (1)$$

$$139 \|\frac{\pi_{k'}^*}{\pi_k^*}\|_\infty = \Theta\left(\prod_{r=k}^{k'-1} \|\frac{\pi_{r+1}^*}{\pi_r^*}\|_\infty\right), \quad (2)$$

$$140 \mathcal{C}(\pi_{k'}^* | \pi_k^*) = \tilde{\Theta}\left(\|\frac{\pi_{k'}^*}{\pi_k^*}\|_\infty\right), \quad (3)$$

141 for some constant  $C^* > 1$ , where  $\tilde{\Theta}(\cdot)$  hide logarithmic factors in a confidence parameter. Then,  
142 the complexity of direct learning and step-wise curriculum learning satisfy:

- 143 • Direct estimation of  $\pi_K^*$  from  $\pi_0^*$  costs  $\mathcal{C}_{\text{direct}} := \sum_{k=0}^{K-1} \mathcal{C}(\pi_{k+1}^* | \pi_k^*) = \Theta((C^*)^K)$ .
- 144 • Stepwise curriculum costs  $\mathcal{C}_{\text{curriculum}} := \sum_{k=0}^{K-1} \mathcal{C}(\pi_{k+1}^* | \pi_k^*) = \Theta(KC^*)$ .

145 Consequently,  $\mathcal{C}_{\text{direct}}/\mathcal{C}_{\text{curriculum}} = \Omega((C^*)^{K-1}/K) \rightarrow \infty$  as  $K \rightarrow \infty$ .

146 **Exponential–Polynomial Gaps Under Relaxed Conditions.** Even if the exact equalities above  
147 are relaxed in two ways: (i) some steps admit polynomial exponents  $\mathcal{C}(\pi_{k+1}^* | \pi_k^*) = (C^*)^{p_k}$  with  
148  $1 \leq p_k \ll K$ ; (ii) the direct complexity satisfies  $\mathcal{C}_{\text{direct}} = \Theta((C^*)^{bK-c})$  for some  $1 \leq b \ll K$  and  
149  $c \ll bK$ ; then writing  $p_{\text{max}} := \max_k p_k$  we still have

$$150 \mathcal{C}_{\text{curriculum}} = O(K(C^*)^{p_{\text{max}}}) \ll \Theta((C^*)^{bK-c}) = \mathcal{C}_{\text{direct}}. \quad (4)$$

151 <sup>2</sup>The concrete complexity is left to be defined for certain specific scenario. For instance, it can be the  
152 sample Complexity of finetuning algorithms (Thm. 3), or the reward-oracle query complexity of test-time  
153 scaling (Thm. 4).




Figure 2: Reasoning tree for parity problems with  $d = K = 3$  and input  $x_1, x_2, x_3, \text{EOS}$ . The nodes on the 2nd–4th levels denote hypotheses about which index is the current secret index, corresponding to  $x_{i_1}, x_{i_2}$ , and  $x_{i_3}$ , respectively. In our parity CoT class  $f \in \mathcal{F}_{2\text{S-ART}}^{\text{parity}}$ , each step actually consists of two actions: (i) choose the next secret index  $i_t$ ; (ii) after the choice, apply an XOR over  $z_{t-1}$  and  $x_{i_t}$  as  $z_t = \Phi_l(z_{t-1}, x_{i_t}) := z_{t-1} \oplus x_{i_t}$ , as formalized in Eq. (7). For visual clarity, the tree only displays the index-selection branches and omits the explicit XOR updates. E denotes the EOS token. For parity tasks, there are illegal children  $\notin \mathcal{I}_l$  for each parent that violate the “legal” criteria: non-repeating indices and strictly increasing order ( $i_1 < i_2 < i_3 < \dots$ ); thus, for any parity problem, CoTs with duplicate variables or decreasing index order are illegal.

**Role of the Base Model.** The theorem presumes nontrivial coverage of the base model, namely  $\pi_k^*(\cdot | x) \ll \pi_{\text{ref}}(\cdot | x)$  for all  $x$  and  $C_{k,0} = \sup_x \left\| \frac{d\pi_k^*}{d\pi_{\text{ref}}} \right\|_{L^\infty(\pi_{\text{ref}})} < \infty$ . This does not hold for an untrained model from scratch but is plausible in post-training, aligning with theory showing that base-model coverage upper-bounds what fine-tuning or test-time scaling can achieve (Foster et al., 2025) and with empirical evidence that sparse-reward RL predominantly reinforces pre-existing tree-like reasoning patterns rather than creating support where the base model has zero probability (Snell et al., 2024; Yue et al., 2025; AI et al., 2025; Gandhi et al., 2025).

**Related Theoretical Context.** Indeed, the assumptions in Parashar et al. (2025) closely mirror those above in the case of Approximate Policy Iteration settings. We discuss this comparison in App. E. Furthermore, for linearly realizable Markov Decision Processes, leveraging the spanner-sampling framework of Foster et al. (2025), we can also prove an exponential improvement in inference-time computational complexity (Thm. 5) by similar assumptions; see App. E for details.

In subsequent sections, we show that this theorem naturally manifests in tree-reasoning settings, where the subtasks correspond to *prefix reasoning traces* (Parashar et al., 2025) or *hint-assisted reasoning* (Liu et al., 2025b; Amani et al., 2025) of the target task, and each parent state assigns comparable probability mass across its children, thereby satisfying Eq. (1) and Eq. (2). We further show the representation theorem where the transformer can subsume the autoregressive reasoning tree, and proved that Eq. (4) holds in both RL finetuning and test-time scaling in terms of sample complexity and oracle complexity.

## 2.1 TRANSFORMER REASONS OVER A STATES-CONDITIONED REASONING TREE

Building on prior work that views post-training as reweighting over a pre-trained reasoning tree (Snell et al., 2024; Yue et al., 2025; AI et al., 2025; Gandhi et al., 2025; Liu et al., 2025b), we model the reasoning from a tree-causal, low-order autoregressive perspective, per formalized below.

**Definition 1** (2-States Conditioned Autoregressive Reasoning Tree (2S-ART)). *A 2S-ART  $\mathcal{F}_{2\text{S-ART}}$  ( $\{\Phi_l\}_{l \leq L}, \{\mathcal{I}_l\}_{l \leq L}$ ) is a function class of autoregressive tasks  $f : [K]^d \times \{K+1\} \rightarrow [K+1]^{\leq L+1}$  such that, given input  $\mathbf{x} = (x_1, \dots, x_d) \in [K]^d$  with terminal token  $\text{EOS} := K+1$ , generates a chain-of-thought  $(z_1, \dots, z_\ell, \text{EOS})$ . At each step  $l$ , the model chooses an index  $i_l$  from a legal set  $\mathcal{I}_l$  of size  $\Theta(d)$ , reads the corresponding element  $v_{i_l}$  from the current sequence, and updates its reasoning state by a two-states map  $z_l = \Phi_l(z_{l-1}, v_{i_l})$ , with  $z_0 = \text{EOS}$  and  $\Phi_l(z, \text{EOS}) = \text{EOS}$ .*

*A task  $f \in \mathcal{F}_{2\text{S-ART}}$  is identified with its index path  $S = (i_1, \dots, i_k, d+1)$ , the associated curriculum subtask family  $\{f_{S^1}, \dots, f_{S^k}\}$  corresponds to the prefix paths  $S^l = (i_1, \dots, i_l, d+1), l \in [k]$ .*

Fig. 2 is an example of parity with  $d = K = L = 3$ , and Fig. 1 is an instance of Countdown with  $d = 4, K = \mathbb{Z}^+, L = 3$  without illustrating explicit EOS. Here EOS is the end-of-sequence token, and  $z_l$  is the  $l$ -th reasoning state in the CoT. Depending on the task,  $z_l$  may be a single token (e.g., in Eq. (7) where  $\Phi_l(z_{l-1}, v_l(i_l)) = z_{l-1} \oplus v_l(i_l)$ ) or a short expression (e.g., in the countdown game where  $\Phi_l(z_{l-1}, v_l(i_l))$  applies  $(+, -, \times, \div)$  to  $z_{l-1}$  and  $v_l(i_l)$ ; see Fig. 1). The legal index set  $\mathcal{I}_l$  encodes real-world selection rules, such as not reusing indices or numbers in parity or Countdown. The 2S-ART also subsumes prior abstractions of reasoning behavior:

- Markov-chain reasoning (Kim et al., 2025c) ( $\Phi_l(z_{l-1}, v_l(i_l)) = \phi(z_{l-1})$  for some  $\phi$ );
- Induction-head for associative recall (Nichani et al., 2025) ( $\Phi_l(z_{l-1}, v_l(i_l)) = v_l(i_l)$ );
- causal-graph reasoning (Nichani et al., 2024) ( $\Phi_l(z_{l-1}, v_l(i_l)) = \phi(v_l(i_l))$  for some  $\phi$ ).

A detailed version of Def. 1 (Def. 3) appears in App. G. To rule out ambiguities where different index paths represent the same task, we adopt the uniqueness assumption below.

**Assumption 1.** We assume the uniqueness of the 2S-ART  $\mathcal{F}_{2S-ART}(\{\Phi_l\}_{l \leq L}, \{\mathcal{I}_l\}_{l \leq L})$  in terms of the index path: for  $\forall f_{S_*}, f_{S'_*} \in \mathcal{F}_{2S-ART}$  with  $|S_*|, |S'_*| \leq L + 1$ , if  $|S_*| \neq |S'_*|$ , then  $f_{S_*} \neq f_{S'_*}$ .

A natural way to define a base model with general capability over the task class  $\mathcal{F}_{2S-ART}$  is to suggest that, at each depth  $l \in [L]$ , the model assigns uniform probability to all children in the legal branch (Liu et al., 2025b) as below.

**Definition 2** (Probabilistic 2S-ART Base Model (PART)). Consider a 2S-ART  $\mathcal{F}_{2S-ART}(\{\Phi_l\}_{l \leq L}, \{\mathcal{I}_l\}_{l \leq L})$  in Def. 2 with  $L \ll d$ , a PART samples  $S_*$  by uniformly drawing  $i_l \in \mathcal{I}_l(\text{CoT}_{l-1}), \forall l \in [L]$ , i.e.  $P(i_l | z_{<l}) = \frac{1}{|\mathcal{I}_l(\text{CoT}_{l-1})|} = \Theta(d^{-1})$ .

This uniform selection rule in the 2S-ART base model leads directly to an exponential decay in success probability (*pass-rate*) of tasks with depth  $l$ , as summarized below.

**Corollary 1** (Exponential Decay of Success Probability with Depth). Consider a 2S-ART sampler defined in Def. 1. For a fixed target  $f_{S_*} \in \mathcal{F}_{2S-ART}$  with associated curriculum subtask family  $\mathcal{F}_{S_*} = \{f_{S_*^1}, \dots, f_{S_*^k}\}$ , the probability that PART samples the legal CoT for  $f_{S_*}$  is  $\Theta(d^{-(l+1)})$ .

**Link to the *pass-rate* in Post-training.** The exponential-decay property of our autoregressive reasoning tree naturally satisfies Eqs. (1)–(2). Consider the policy  $\pi_{S_*^l}$  for  $f_{S_*^l}$ : given any  $\mathbf{x} \in [K]^d$ ,  $\pi_{S_*^l}$  samples correct  $i_1^*, \dots, i_l^*$  in the correct order almost surely, thereby generating correct prefix  $z_1^*, \dots, z_l^*$  autoregressively, and from step  $l + 1$  onward copies the behavior of PART. Denote  $\pi_{\text{PART}}$  as the policy of PART. It is straightforward to verify that  $\|\frac{\pi_{S_*^{l+1}}}{\pi_{S_*^l}}\|_\infty = \Theta(d)$  and  $\|\frac{\pi_{S_*^l}}{\pi_{\text{PART}}}\|_\infty = \Theta(d^{l+1})$  (A1), which align exactly with Eqs. (1) and (2) up to a  $d$ . In the next step, we demonstrate that a transformer can replicate PART, and then establish that Eq. (4) (corresponding to (A2)–(A3) in the introduction) also holds, both for the sample complexity of transformer’s fine-tuning convergence and for the oracle-query complexity in test-time scaling.

**Transformer (TF( $\cdot$ ;  $\mathbf{W}$ )) as Learning Model.** Let  $\mathbf{U} = [\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K, \boldsymbol{\mu}_{\text{EOS}}] \in \mathbb{R}^{d_{\text{x}} \times (K+1)}$  be the embedding vocabulary, where the  $k$ -th column is the embedding of token  $k$ , and  $\boldsymbol{\mu}_{\text{EOS}}$  the embedding of EOS. Following Li et al. (2025), positional encodings  $\mathbf{P} := [\mathbf{p}_1, \dots, \mathbf{p}_{d+1}]$  are mutually orthogonal and satisfy  $\mathbf{p}_i \perp \mathbf{U}$ . They are concatenated with token embeddings as  $\mathbf{E}[x_m] = [\boldsymbol{\mu}^{x_m \top}, \mathbf{p}_m \top]^\top$ ,  $\mathbf{E}[\text{EOS}] = [\boldsymbol{\mu}_{\text{EOS}} \top, \mathbf{p}_{d+1} \top]^\top$ , and  $\mathbf{E}[z_l] = [\boldsymbol{\mu}^{z_l \top}, \mathbf{p}_l \top]^\top, \forall m \in [d], l \in [L+1]$ . The transformer TF( $\cdot$ ;  $\mathbf{W}$ ) aims to recover the functions defined in the tree graph  $f \in \mathcal{F}_G$ , with embedded input  $\mathbf{E}[x_1], \dots, \mathbf{E}[x_d], \mathbf{E}[\text{EOS}] \in \mathbb{R}^{d_{\text{E}}}$ , where  $d_{\text{E}} = 2d_{\text{x}}$ . Denote  $\mathbf{E}[z_{-d}] := \mathbf{E}[x_1], \dots, \mathbf{E}[z_{-1}] := \mathbf{E}[x_d], \mathbf{E}[z_0] := \mathbf{E}[\text{EOS}]$ , the transformer performs next-token prediction via:

$$\text{TF}(\mathbf{E}[z_{-d}], \dots, \mathbf{E}[z_{-1}], \mathbf{E}[z_0], \mathbf{E}[z_1], \dots, \mathbf{E}[z_{l-1}]; \mathbf{W}) = \mathbf{E}[z_l] \quad (5)$$

Here, original data embeddings remain fixed during the reasoning process, while CoT state  $\mathbf{E}[z_l]$  are generated by TF( $\cdot$ ;  $\mathbf{W}$ ) ( $\mathbf{W}$  denotes the model weights) via:

$$\hat{\mathbf{p}}_l = \sum_{j=-d}^{l-2} \mathbf{V}(\mathbf{E}[z_j] \text{softmax}(\mathbf{E}[z_j]^\top \mathbf{K}^\top \mathbf{Q} \mathbf{E}[z_{l-1}])) \in \mathbb{R}^{d_{\text{x}}} \quad (6)$$

$$\mathbf{p}_{i_l} \sim \text{softmax}(\hat{\mathbf{p}}_l^\top \mathbf{P} / \beta), \quad \text{The algorithm terminates with } \mathbf{E}[\text{EOS}] \text{ if } \mathbf{p}_{d+1} \text{ is sampled}$$

$$\hat{\boldsymbol{\mu}}^{z_l} = \text{FFN}_l(\boldsymbol{\mu}^{x_{i_l}}, \hat{\boldsymbol{\mu}}^{z_{l-1}}), \quad \mathbf{E}[z_l] = [\hat{\boldsymbol{\mu}}^{z_l}, \mathbf{p}_{i_l}]^\top \in \mathbb{R}^{d_{\text{E}}}.$$

Here,  $\beta > 0$  is a temperature parameter, the causal masking is enforced by setting  $\mathbf{E}[z_j]^\top \mathbf{K}^\top \mathbf{Q}\mathbf{E}[z_l] \leftarrow -\infty$  whenever  $j \geq l$  or  $l \leq d$ .  $\hat{\boldsymbol{\mu}}^{x_{i_l}}$  is depth-specific crucial token utilized for the  $l$ -th reasoning step, and  $\hat{\boldsymbol{\mu}}^{z_l}$  is the resulting  $l$ -th reasoning step, formed by the depth-specific non-linear residual Feedforward layer  $\text{FFN}_l(\cdot, \cdot)$ , which is tailored for task-specific reasoning operation  $\Phi_l(\cdot, \cdot)$  defined in Def. 1, modeling after reasoning steps executed by the neural model such as the XOR operation computed by  $h^\top \text{ReLU}(W[\mathbf{E}[z_{l-1}]_{:d_x}])$  in Wen et al. (2025a) or  $\phi(\cdot)$  function in Kim & Suzuki (2025), and the calculation operation in each reasoning step of countdown game, per Fig. 1 (Liu et al., 2025b). The following shows that  $\text{TF}(\cdot)$  can replicate the `PART` per Def. 2.

**Theorem 2** (Base Model as `PART` ( $\text{TF}(\cdot; \mathbf{W}_{\text{base}})$ )). *Fix any 2S-ART ( $\{\Phi_l\}_{l \leq L}, \{\mathcal{I}_l\}_{l \leq L}$ ) from Def. 1 and its probabilistic base model (`PART`) in Def. 2. Assume that for each depth  $l$  there exists the map  $\text{FFN}_l$  that replicates the target operation on embeddings, i.e.,  $\boldsymbol{\mu}^{z_l} = \text{FFN}_l(\boldsymbol{\mu}^{v_l(i_l)}, \mathbf{E}[z_{l-1}]_{:d_x})$  whenever  $z_l = \Phi_l(z_{l-1}, v_l(i_l))$ . Then there exists a parameterization of a single-head attention transformer of the form in Eq. (6) to copy the `PART` behavior in Def. 2.*

The assumption that  $\text{FFN}_l$  exists for each depth  $l$  is justified by the universal approximation theorem (Cybenko, 1989; Hornik, 1991). This establishes that feedforward networks can approximate any continuous function on compact subsets of  $\mathbb{R}^n$  to arbitrary precision. Since continuous functions are dense in  $L^p$  spaces for  $1 \leq p < \infty$ , this extends to approximating most practically relevant functions. Recent advances show that ReLU networks of width  $d + 3$  and arbitrary depth can approximate any scalar continuous function of  $d$  variables (Lu et al., 2017), while Suzuki (2018); Suzuki & Nitanda (2021) demonstrates optimal approximation rates for functions in Besov spaces, encompassing broader function classes beyond continuity.

**0-1 Outcome Signals.** In real-world mathematical reasoning and program-like tasks, supervision is often available only at the outcome level (correct vs. incorrect), while collecting step-by-step intermediate supervision is costly or infeasible. We therefore adopt 0-1 outcome-only supervision that evaluates only the *final pre-EOS prediction*. We use two oracles:

- $\mathbf{R}_{\mathbf{x}}^{f_{S^*}}(\cdot)$ : returns 1 iff, immediately before sampling EOS,  $\text{TF}(\cdot; \mathbf{W})$  outputs the token whose embedding equals  $\boldsymbol{\mu}^{f_{S^*}(\mathbf{x})}$ ; else 0.
- $\mathbf{R}_{\mathbf{x}}^{\mathcal{F}_{S^*}}(\cdot, \ell)$ : for the curriculum subtask family  $\mathcal{F}_{S^*} = \{f_{S_1^*}, \dots, f_{S_{k^*}^*}\}$  per Def. 1, returns 1 iff the pre-EOS token equals  $\boldsymbol{\mu}^{f_{S_\ell^*}(\mathbf{x})}$  for  $\ell \in [k^*]$ ; else 0.

These outcome-only definitions are agnostic to any specific case study; in the parity case discussed in Sec. 2.1.1, outcome reward specializes to whether the result is correct per Eq. (7); in the multi-step language translation task studied in Abedsoltan et al. (2025) where the sampled word is translated to different language at different reasoning steps according to a fixed order, the outcome reward denotes whether the translation is correct for the task family.

**Challenge: Inherent Reward Hacking.** Outcome-only supervision on the final pre-EOS token allows many spurious CoTs to be accepted for a fixed input  $\mathbf{x}$ , because the oracle **only checks the surface token** and does not constrain the internal index-selection process of the 2S-ART ( $\{\Phi_l\}_{l \leq L}, \{\mathcal{I}_l\}_{l \leq L}$ ). For instance, in parity, choosing any wrong index at some depth can still yield the correct final bit with probability  $1/2$  under  $\mathbf{x} \sim \text{Unif}(\{0, 1\}^d)$ .

### 2.1.1 CONCRETE EXAMPLE: PARITY PROBLEM

**Sparse Parity Problem.** We remark that the Sparse Parity Problem can be seen as a  $K = 2$  subcase of Def. 1, with the deterministic kernel  $\Phi_l(z_{l-1}, x_{i_l})$  given by the XOR operation  $z_{l-1} \oplus x_{i_l}$ . Formally, given a  $d$ -dimensional binary input vector  $\mathbf{x} = (x_1, \dots, x_d) \sim \text{Unif}\{0, 1\}^d$ , define a class of Boolean functions determined by a size- $k$  secret index set  $S \subseteq [d]$ . Each function computes the parity (XOR-sum) of the input components specified by  $S$ :  $f_S(\mathbf{x}) = \bigoplus_{i \in S} x_i = x_{i_1} \oplus x_{i_2} \oplus \dots \oplus x_{i_k}$ , where  $\oplus$  denotes XOR and  $i_1 < i_2 < \dots < i_k$  without loss of generality. The output satisfies  $f_S(\mathbf{x}) = 1$  if  $\sum_{i \in S} x_i$  is odd, and  $f_S(\mathbf{x}) = 0$  otherwise. The function class  $\mathcal{P}_{d,k} = \text{Parity}(d, k)$  contains all such functions, with cardinality  $|\mathcal{P}_{d,k}| = \binom{d}{k}$ .

**CoT and Curriculum Subtask Family for Parity.** As a  $K=2$  instance of Def. 1, parity’s CoT is indexed by the path  $S = (i_1, \dots, i_k)$  with  $i_k = d+1$  (EOS). The resulting XOR-based CoT (Wen

et al., 2025a; Abedsoltan et al., 2025) is:

$$z_1 = x_{i_1}, \quad z_2 = x_1 \oplus x_{i_2}, \quad \dots, \quad z_k = z_{k-1} \oplus x_{i_k} = f_S(\mathbf{x}). \quad i_1 < i_2 < \dots < i_k \quad (7)$$

In this case, the curriculum subtask family is the prefix-index family  $\mathcal{F}_S = \{f_{S^1}, \dots, f_{S^{k^*}}\}$  with  $S^{k'} = (i_1, \dots, i_{k'})$  and  $f_{S^{k'}}(\mathbf{x}) = z_{k'}$  for  $k' \in [k]$ ; early termination by EOS follows Def. 1. In particular, the legal sets  $\{\mathcal{I}_l\}$  is defined to satisfy  $i_1 < i_2 < \dots < i_k$ .

**Transformer (TF( $\cdot$ ;  $\mathbf{W}$ )) as Learning Model.** We follow Eq. (5), Eq. (6), set the vocabulary as  $\mathbf{U} = [\boldsymbol{\mu}^0, \boldsymbol{\mu}^1, \boldsymbol{\mu}_{\text{EOS}}] = [e_1, e_2, e_3]$  for 0, 1, EOS, and the  $\text{FFN}_l$  is instantiated for XOR operation:

$$\text{FFN}_l = \mathbf{W}_2 \text{ReLU}[\mathbf{W}_1(\hat{\boldsymbol{\mu}}^{x_{i_l}} + \mathbf{E}[z_{l-1}]_{:d_X})], \quad \mathbf{W}_1 = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \\ 0 & 1 & 0 \\ -\frac{1}{2} & \frac{1}{2} & -\frac{1}{2} \end{bmatrix}, \quad \mathbf{W}_2 = \begin{bmatrix} 1 & -1 & 2 \\ 0 & 1 & -2 \\ 0 & 0 & 0 \end{bmatrix}, \forall m \in [k']. \quad (8)$$

It can be checked directly that this ensure our  $\text{FFN}_l$  replicates the  $\Phi_l(z_{l-1}, v_l(i_l)) = z_{l-1} \oplus v_l(i_l)$ :

$$\text{FFN}_l(\boldsymbol{\mu}^0, \boldsymbol{\mu}^0) = \text{FFN}_l(\boldsymbol{\mu}^1, \boldsymbol{\mu}^1) = \boldsymbol{\mu}^0, \quad \text{FFN}_l(\boldsymbol{\mu}^0, \boldsymbol{\mu}^1) = \text{FFN}_l(\boldsymbol{\mu}^1, \boldsymbol{\mu}^0) = \boldsymbol{\mu}^1, \quad \text{FFN}_l(\boldsymbol{\mu}^{0/1}, \boldsymbol{\mu}^{\text{EOS}}) = \boldsymbol{\mu}^{0/1}.$$

Per in previous Sec. 2.1, oracle  $R^{f_{S^*}}(\cdot)$  and  $R^{\mathcal{F}_{S^*}}(\cdot)$  for the target task  $f_{S^*}$ ,  $|S^*| = k + 1$ , rather than direct intermediate supervision used in Kim & Suzuki (2025). Kim & Suzuki (2025) also considered data augmentation scheme (causal mask + random  $d$ -bit + ultra filter) circumvents the need for intermediate supervision, but is specifically tailored to their 2-parity-based hierarchical CoT when  $k^* = 2^v$ ,  $v \in \mathbb{Z}^+$ . By contrast, we adopt the more natural CoT formulation in Eq. (7).

### 2.1.2 OUTCOME SIGNAL-BASED RL FINETUNINGS BY GRADIENT DESCENT

For mathematics benchmarks, the conventional REINFORCE objective is used to increase the probability that sampled CoTs yield correct answers Xiong et al. (2025); Setlur et al. (2025). In our setting, the training objective can be formulated as

$$\mathcal{J}_{\text{REINFORCE}}^{k^*}(\mathbf{W}^{k^*}) = \mathbb{E} \left[ R^{k^*}(\text{TF}(\cdot; \mathbf{W})) \right], \quad (9)$$

where  $R^{k^*}(\cdot) \in \{R_{\mathbf{x}^{f_{S^*}}}(\cdot), R_{\mathbf{x}^{\mathcal{F}_{S^*}}}(\cdot)\}$ ,  $\hat{p}_{\mathbf{W}^{k^*}}(\cdot|\cdot)$  denotes the predictive distribution by  $\text{TF}(\cdot; \mathbf{W})$ , and  $\mathbf{W} \in \mathbb{R}^{(d_E - d_X)^2}$  is the only matrix we consider trainable in the parametrization of Thm. 2:

$$\mathbf{K}^\top \mathbf{Q} = \begin{pmatrix} \mathbf{0}_{d_X \times d_X} & \mathbf{0}_{d_X \times (d_E - d_X)} \\ \mathbf{0}_{(d_E - d_X) \times d_X} & \mathbf{W} \end{pmatrix}, \quad \mathbf{V} = (\mathbf{0} \quad \mathbf{I}_{d_E - d_X}), \quad (10)$$

where the  $\mathbf{0}$ ,  $\mathbf{V}$ , as well as the feedforward map  $\text{FFN}_l$ , is considered fixed during finetuning. This type of reparametrization is common in the transformer optimization literature to enable tractable analysis (Zhang et al., 2024b; Huang et al., 2024; Mahankali et al., 2023; Kim & Suzuki, 2025).

In mathematical datasets, the difficulty measure *pass-rate* often coincides with other task-specific measures Parashar et al. (2025): Blocksworld uses plan length of CoT (Valmeekam et al., 2023), Countdown counts the number of operations in the CoT (Gandhi et al., 2024), and in parity tasks, the number of XOR operations naturally reflects the CoT difficulty. Building on these observations, a natural curriculum for  $\mathcal{F}_{S^*}$  is to gradually increase reasoning depth from shallow to deep, where the number of reasoning operations grows with depth. Separately, Liu et al. (2025b); Amani et al. (2025) demonstrated the benefit of providing hints (partial CoT prefixes) and progressively shortening them so that the model completes longer suffixes. Accordingly, we consider two categories of curriculum finetuning under the Outcome Signal Oracle  $\mathbf{R}_{\mathbf{x}^{\mathcal{F}_{S^*}}}(\cdot)$ :

- **Depth-increasing Curriculum** (Parashar et al., 2025): at the  $l$ -th stage, the algorithm truncate CoTs from  $\text{TF}(\cdot, \mathbf{W}^{(t)})$  with EOS to ensure length  $l + 1$ .
- **Hint-decreasing Curriculum** (Liu et al., 2025b): at the  $l$ -th stage, the algorithm provides a CoT prefix of length  $k^* + 1 - l$ , letting  $\text{TF}(\cdot, \mathbf{W}^{(t)})$  generate the remaining steps.

Building on this, the following theorem establishes Eq. (4) (corresponding to **(A2)**–**(A3)** in the introduction) within the autoregressive reasoning tree setting, demonstrating that curriculum-based RL fine-tuning effectively alleviates the exponential bottleneck in sample complexity.

**Theorem 3** (Curriculum RL Finetuning Avoid Exponential Bottleneck). *Let  $\mathbf{U}$  be an orthogonal matrix with  $d_X = \Theta(K)$ ,  $d_E = \Theta(d + L)$ ,  $\text{TF}(\cdot; \mathbf{W}_{\text{base}})$  per in Thm. 2 with trainable  $\mathbf{W}$ , and a target  $f_{S^*} \in \mathcal{F}_{2S\text{-ART}}$  with  $|S^*| = k^* + 1$ . For any  $\varepsilon \in (0, 1)$ , using the RL objective in Eq. (9) and one gradient step per stage (a single step for no-curriculum;  $k^* + 1$  online steps for curricula), with probability no less than  $1 - \delta$ , there exist learning-rate choices  $\eta$  such that*

1. *No-curriculum ( $R_{\mathbf{x}}^{f_{S^*}}(\cdot)$  as oracle, one-shot update with  $\eta = \tilde{\Theta}(\beta d^{k^*+1})$ ): the sample complexity to achieve  $\mathbb{E}_{\mathbf{x} \sim \text{Unif}([K]^d)} [R_{\mathbf{x}}^{f_{S^*}}(\text{TF}(\mathbf{x}; \mathbf{W}))] \geq 1 - \varepsilon$  is at least  $n \geq \tilde{\Omega}(d^{2k^*+2})$ .*
2. *Depth-increasing Curriculum and Hint-decreasing Curriculum ( $R_{\mathbf{x}}^{f_{S^*}}(\cdot)$  as oracle,  $k^* + 1$  online updates with  $\eta = \tilde{\Theta}(\beta d^2)$ ): the sample complexity to achieve  $\mathbb{E}_{\mathbf{x} \sim \text{Unif}([K]^d)} [R_{\mathbf{x}}^{f_{S^*}}(\text{TF}(\mathbf{x}; \mathbf{W}))] \geq 1 - \varepsilon$  is at most  $n \leq \tilde{O}((k^* + 1)d^2)$ .*

Here  $\tilde{\Omega}(\cdot)$  and  $\tilde{O}(\cdot)$  hide polylogarithmic factors in  $d, 1/\delta$  and  $1/\varepsilon$ , and absolute constants (e.g., temperature  $\beta$ ) as well as task-dependent spurious-acceptance parameters.

**Remark.** Item (1) is not minimax-tight in certain cases of 2S-ART (for example, the minimax rate for the parity class up to EOS is  $d^{k^*}$ ). The gap arises because our analysis is based on vanilla gradient descent over the empirical estimate of outcome rewards Eq. (9), which disallows cross-sample referencing or structure-exploiting techniques such as Gaussian elimination (Raz, 2018; Abbe & Sandon, 2020). Nevertheless, our result already *suffices*: Item (2) establishes that curriculum post-training can reduce the sample complexity to polynomial order. The specific  $d^2$  dependence stems from distinguishing margins of order  $\Theta(d^{-1})$  in parallel attempts, which information-theoretically requires inverse-square costs. More sophisticated noisy SGD algorithms that enables cross-sample reference may improve this rate at the cost of a  $\text{poly}(d)$  time complexity (Cornacchia & Mossel, 2023; Abbe et al., 2023c).

**Related Work.** It is worth noting that one- or few-shot convergence under large learning rates has been widely discussed in prior optimization work, such as Cornacchia & Mossel (2023); Kim & Suzuki (2025). As mentioned in the introduction and Sec. A, there are also theoretical investigations of curriculum benefits in pre-training (train-from-scratch) settings for the parity function class (Cornacchia & Mossel, 2023; Abbe et al., 2023c; Panigrahi et al., 2024). In particular, Cornacchia & Mossel (2023) and Abbe et al. (2023c) construct curricula by mixing data distributions, where “difficulty” is characterized by the density of Hamming weight (fewer 1s than 0s are easier), while Panigrahi et al. (2024) study a teacher–student setup in which “difficulty” is defined by the signal strength of checkpoints provided by the teacher. Their analyses focus on 2-layer ReLU networks or MLPs trained by carefully designed stage-wise or layer-wise gradient descent algorithms. In contrast, our perspective builds on the difficulty measure *pass-rate* (see Cor. 1 and its remark) and its connection to the inherently probabilistic, tree-like CoT generation behavior in post-training (Liu et al., 2025b). This leads naturally to the depth-increasing Parashar et al. (2025) and hint-decreasing (Liu et al., 2025b; Amani et al., 2025) curriculum, and we consider GD updates of transformer without additional algorithmic tailoring, better mirroring the post-training scenarios.

**Proof Outline.** Intuitively, with a noiseless gradient oracle, the gradient magnitude in the projection space corresponding to the correct secret index strictly exceeds all others, so a single clean update suffices to create a decisive logits gap for softmax sampling. Thus the sample complexity is determined by how many samples are needed so that the estimated gradients reliably reflect this margin. A sufficient condition is that the confidence intervals of both the correct and incorrect indices are smaller than half of the smallest expected margin. From gradient calculations, the expected margin at step  $\ell$  is  $\Theta(\beta^{-1} d^{-(k+2-\ell)})$ . To ensure confidence radius no larger than  $\Theta(\beta^{-1} d^{-(k+2-\ell)}/2)$ , Bernstein-type bounds with a union argument yield  $n_\ell = \tilde{\Theta}(d^{2k+2-\ell})$ . The first step dominates with  $n_1 = \tilde{\Theta}(d^{2k+2})$ , giving total sample complexity  $n \geq \tilde{\Omega}(d^{2k+2})$ . In contrast, under curriculum strategies, the expected margin at each step is  $\Theta(\beta d^{-1}/2)$ , so the required samples per step are at most  $\tilde{\Theta}(d^2)$ . Across  $k + 1$  steps this totals  $\tilde{O}((k + 1)d^2)$ .

### 2.1.3 OUTCOME SIGNAL-BASED TEST-TIME SCALING

Consider a pretrained transformer that replicates PART’s uniform legal-branching behavior (Thm. thm:tf-realizes-part). At test time, given  $\mathbf{x} \sim \text{Unif}([K]^d)$ , we have access only to the oracles  $\mathbf{R}_{\mathbf{x}}^{f_{S^*}}(\cdot)$  and  $\mathbf{R}_{\mathbf{x}}^{\mathcal{F}_{S^*}}(\cdot)$ , which serve as Outcome Reward Models (ORM) for inference-time scaling. Following Foster et al. (2025), we denote by  $T_{\text{data}}$  the *reward oracle-query complexity*, i.e., the total number of evaluations of  $\mathbf{R}_{\mathbf{x}}^{f_{S^*}}$  or  $\mathbf{R}_{\mathbf{x}}^{\mathcal{F}_{S^*}}$ , and by  $T_{\text{comp}}$  the *model-sampling complexity*, i.e., the total number of token emissions from  $\text{TF}_{\text{base}}$ . The theorem below shows that leveraging  $\mathbf{R}_{\mathbf{x}}^{\mathcal{F}_{S^*}}(\cdot)$  in a curriculum-based manner eliminates the exponential bottleneck in terms of  $T_{\text{data}}$  and  $T_{\text{comp}}$ .

**Theorem 4** (Curriculum Test-time Scaling Avoid Exponential Bottleneck). *Let  $\text{TF}_{\text{base}} := \text{TF}(\cdot; \mathbf{W}_{\text{base}})$  per defined in Thm. 2, and consider a target  $f_{S^*} \in \mathcal{F}_{2S\text{-ART}}$  with  $|S^*| = k^* + 1$ . Then, for identifying the ground-truth path  $S^*$  with confidence  $1 - \delta$ :*

1. *Using only  $\mathbf{R}_{\mathbf{x}}^{f_{S^*}}(\cdot)$ , any procedure (including ones that force visible  $x$ -tokens by rejection sampling and then roll out) requires  $T_{\text{data}} \geq \tilde{\Omega}(d^{2k^*})$  and  $T_{\text{comp}} \geq \tilde{\Omega}(d^{2k^*})$ .*
2. *Using curriculum queries adaptively to  $\mathbf{R}_{\mathbf{x}}^{\mathcal{F}_{S^*}}(\cdot)$ , there exists a procedure with  $T_{\text{data}} \leq \tilde{O}((k^* + 1)d^2)$  and  $T_{\text{comp}} \leq \tilde{O}((k^* + 1)d^3)$ .*

Here  $\tilde{\Omega}(\cdot)$  and  $\tilde{O}(\cdot)$  hide polylogarithmic factors in  $d$ ,  $1/\delta$  and spurious-acceptance parameters.

Akin to the remark of Thm. 3, without structure-exploiting procedure (Raz, 2018; Abbe et al., 2023a), our results match the information-theoretical bound in distinguishing the acceptance gap.

**Proof Outline.** We again illustrate the idea using the parity class. The algorithm’s goal, given input  $\mathbf{x}$ , is to recover the secret index set  $S = \{i_1, \dots, i_k, d + 1\}$ . The key difficulty is that an incorrect index set still yields reward 1 with probability  $1/2$ , a hallmark of parity. At reasoning step  $\ell$ , conditioned on a correct prefix, both the correct index  $j_*$  and any wrong index  $j \neq j_*$  are selected with probability  $\Theta(d^{-1})$ . If  $j_*$  is chosen and the process continues until EOS, the final CoT obtains reward 1 with probability  $\Theta(d^{-(k+1-\ell)} + (1 - \Theta(d^{-(k+1-\ell)}))/2)$ , whereas if  $j \neq j_*$ , the reward is  $1/2$ . Hence the probability gap is  $\Theta(d^{-(k+1-\ell)}/2)$ . Distinguishing such tiny gaps information-theoretically requires  $\tilde{\Theta}(d^{2k+2-2\ell})$  trials, and summing over  $\ell$  yields  $\tilde{\Omega}(d^{2k})$  reward queries in total. By contrast, with curriculum (Alg. 6), the algorithm truncates the generation at length  $\ell + 1$ . Then, if  $j_*$  is chosen, the chance of reward 1 is  $\Theta(d^{-1} + (1 - \Theta(d^{-1}))/2)$ , while if  $j \neq j_*$ , it is  $1/2$ . The resulting gap is  $\Theta(d^{-1}/2)$ , requiring only  $\tilde{O}(d^2)$  trials to resolve. Across  $k + 1$  stages this totals  $\tilde{O}((k + 1)d^2)$  reward-oracle calls, and the additional transformer calls during rejection sampling after  $S$  is identified contribute at most  $\tilde{\Theta}(d)$  more.

## 3 CONCLUSION, DISCUSSION, AND FUTURE WORK

Unlike prior theoretical studies of pre-training curricula, where “difficulty” is task-specific, our framework for post-training is built on the widely used *pass-rate*. Based on this measure, we established Thm. 1, which characterizes when curriculum post-training avoids the exponential bottleneck. We substantiated this by modeling the autoregressive tree-like reasoning process, where reasoning length scales with subtask difficulty and can be replicated by a pretrained transformer. Considering both fine-tuning and test-time scaling, we showed exponential-to-polynomial reductions in sample and oracle-query complexity, serving as a concrete case study of Thm. 1.

A limitation is that our 2S-ART (PART) abstraction is restrictive. For instance, datasets such as Blocksworld involve  $\Phi_l(x_{i_1}, x_{i_2})$  (choosing two states and stacking one on the other), beyond our one-state case. We believe the theory extends with more complex treatment. While our analysis focused on transformer post-training, the benefit of curricula in avoiding exponential bottlenecks is model-agnostic, as highlighted in Thm. 1 and Sec. E. Empirical studies also report easy-to-hard curricula to be effective in diffusion models (Kim et al., 2025a;b; Yu et al., 2023), and it is an interesting direction to investigate whether their underlying principles align with our theorem.

## REFERENCES

- 486  
487  
488 Emmanuel Abbe and Colin Sandon. On the universality of deep learning. In *Advances in Neural*  
489 *Information Processing Systems*, volume 33, pp. 20061–20072, 2020.
- 490 Emmanuel Abbe, Pritish Kamath, Eran Malach, Colin Sandon, and Nathan Srebro. On the power of  
491 differentiable learning versus pac and sq learning. In *Advances in Neural Information Processing*  
492 *Systems*, volume 34, 2021.
- 493  
494 Emmanuel Abbe, Enric Boix Adsera, and Theodor Misiakiewicz. Sgd learning on neural networks:  
495 leap complexity and saddle-to-saddle dynamics. In *The Thirty Sixth Annual Conference on Learn-*  
496 *ing Theory*, pp. 2552–2623. PMLR, 2023a.
- 497 Emmanuel Abbe, Samy Bengio, Aryo Lotfi, and Kevin Rizk. Generalization on the unseen, logic  
498 reasoning and degree curriculum. In *Proceedings of the 40th International Conference on Ma-*  
499 *chine Learning (ICML)*, pp. 31–60. PMLR, 2023b.
- 500  
501 Emmanuel Abbe, Elisabetta Cornacchia, and Aryo Lotfi. Provable advantage of curriculum learning  
502 on parity targets with mixed inputs. In *Advances in Neural Information Processing Systems*,  
503 volume 36, pp. 24291–24321, 2023c.
- 504 Amirhesam Abedsoltan, Huaqing Zhang, Kaiyue Wen, Hongzhou Lin, Jingzhao Zhang, and Mikhail  
505 Belkin. Task generalization with autoregressive compositional structure: Can learning from  $d$   
506 tasks generalize to  $d^t$  tasks? In *Forty-second International Conference on Machine Learning*,  
507 2025. URL <https://openreview.net/forum?id=iZdGZSWe1A>.
- 508  
509 Essential AI, :, Darsh J Shah, Peter Rushton, Somanshu Singla, Mohit Parmar, Kurt Smith, Yash  
510 Vanjani, Ashish Vaswani, Adarsh Chaluvaram, Andrew Hojel, Andrew Ma, Anil Thomas, An-  
511 thony Polloreno, Ashish Tanwer, Burhan Drak Sibai, Divya S Mansingka, Divya Shivaprasad,  
512 Ishaan Shah, Karl Stratos, Khoi Nguyen, Michael Callahan, Michael Pust, Mrinal Iyer, Philip  
513 Monk, Platon Mazarakis, Ritvik Kapila, Saurabh Srivastava, and Tim Romanski. Rethinking  
514 reflection in pre-training. *arXiv preprint arXiv:2504.04022*, 2025.
- 515  
516 Mohammad Hossein Amani, Aryo Lotfi, Nicolas Mario Baldwin, Samy Bengio, Mehrdad Fara-  
517 jatabar, Emmanuel Abbe, and Robert West. RL for reasoning by adaptively revealing rationales.  
*arXiv preprint arXiv:2506.18110*, 2025.
- 518  
519 A. Amiri, X. Huang, M. Rofin, and M. Hahn. Lower bounds for chain-of-thought reasoning in  
520 hard-attention transformers. *arXiv preprint*, 2025. URL <https://arxiv.org/abs/2502.02393>. arXiv:2502.02393.
- 521  
522 Sanghwan Bae, Jiwoo Hong, Min Young Lee, Hanbyul Kim, JeongYeon Nam, and Donghyun  
523 Kwak. Online difficulty filtering for reasoning oriented reinforcement learning. *arXiv preprint*  
524 *arXiv:2504.03380*, 2025.
- 525  
526 P. Barceló, A. Kozachinskiy, and T. Steifer. Ehrenfeucht–haussler rank and chain of thought. *arXiv*  
527 *preprint*, 2025. URL <https://arxiv.org/abs/2501.12997>. arXiv:2501.12997.
- 528  
529 Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In  
*Proceedings of the 26th ICML*, pp. 41–48, 2009.
- 530  
531 Adam Block and Yury Polyanskiy. The sample complexity of approximate rejection sampling with  
532 applications to smoothed online learning. *arXiv preprint arXiv:2302.04658*, 2024.
- 533  
534 Siyu Chen, Heejune Sheen, Tianhao Wang, and Zhuoran Yang. Unveiling induction heads: Pro-  
535 vable training dynamics and feature learning in transformers. In *Advances in Neural Information*  
*Processing Systems*, volume 37, pp. 66479–66567. Curran Associates, Inc., 2024.
- 536  
537 Herman Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of  
538 observations. *Annals of Mathematical Statistics*, 23(4):493–507, 1952.
- 539  
Elisabetta Cornacchia and Elchanan Mossel. A mathematical model for curriculum learning. *arXiv*  
*preprint arXiv:2301.13833*, 2023.

- 540 George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control,*  
541 *signals and systems*, 2(4):303–314, 1989.
- 542
- 543 Amit Daniely and Eran Malach. Learning parities with neural networks. In *Advances in Neural*  
544 *Information Processing Systems*, volume 33, pp. 20356–20365, 2020.
- 545
- 546 Manfred Perdigao Do Carmo and J Flaherty Francis. *Riemannian geometry*, volume 2. Springer,  
547 1992.
- 548 Guhao Feng, Bohang Zhang, Yuntian Gu, Haotian Ye, Di He, and Liwei Wang. Towards revealing  
549 the mystery behind chain of thought: a theoretical perspective. In *Advances in Neural Information*  
550 *Processing Systems*, 2023.
- 551 Dylan J. Foster, Zakaria Mhammedi, and Dhruv Rohatgi. Is a good foundation necessary for efficient  
552 reinforcement learning? the computational role of the base model in exploration. *arXiv preprint*  
553 *arXiv:2503.07453*, 2025.
- 554
- 555 Kanishk Gandhi, Denise Lee, Gabriel Grand, Muxin Liu, Winson Cheng, Archit Sharma, and  
556 Noah D Goodman. Stream of search (SoS): learning to search in language. *arXiv preprint*  
557 *arXiv:2404.03683*, 2024.
- 558 Kanishk Gandhi, Ayush Chakravarthy, Anikait Singh, Nathan Lile, and Noah D. Goodman. Cogni-  
559 tive behaviors that enable self-improving reasoners, or, four habits of highly effective stars. *arXiv*  
560 *preprint arXiv:2503.01307*, 2025.
- 561
- 562 Alex Graves, Marc G. Bellemare, Jacob Menick, Rémi Munos, and Koray Kavukcuoglu. Automated  
563 curriculum learning for neural networks. In *Proceedings of the 34th International Conference on*  
564 *Machine Learning (ICML)*, pp. 1311–1320. PMLR, 2017.
- 565
- 566 M. Hahn and N. Goyal. A theory of emergent in-context learning as implicit structure induction.  
567 *arXiv preprint*, 2023. URL <https://arxiv.org/abs/2303.07971>. arXiv:2303.07971.
- 568
- 569 Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the*  
570 *American Statistical Association*, 58(301):13–30, 1963.
- 571
- 572 Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4  
573 (2):251–257, 1991.
- 574
- 575 Jianhao Huang, Zixuan Wang, and Jason D. Lee. Transformers learn to implement multi-step gradi-  
576 ent descent with chain of thought. *arXiv preprint*, 2025a. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2502.21212)  
577 [2502.21212](https://arxiv.org/abs/2502.21212). arXiv:2502.21212.
- 578
- 579 R. Huang, Y. Liang, and J. Yang. How transformers learn regular language recognition: A theoretical  
580 study on training dynamics and implicit bias. *arXiv preprint*, 2025b. URL [https://arxiv.](https://arxiv.org/abs/2505.00926)  
581 [org/abs/2505.00926](https://arxiv.org/abs/2505.00926). arXiv:2505.00926.
- 582
- 583 Y. Huang, Z. Wen, A. Singh, Y. Chi, and Y. Chen. Transformers provably learn chain-of-thought  
584 reasoning with length generalization. *Technical Report*, 2025c. URL . preprint.
- 585
- 586 Yu Huang, Yuan Cheng, and Yingbin Liang. In-context convergence of transformers. In *Proceedings*  
587 *of the 41st International Conference on Machine Learning*, pp. 19660–19722, 2024.
- 588
- 589 Emilie Kaufmann, Olivier Capp’e, and Aur’elien Garivier. On the complexity of best-arm identifi-  
590 cation in multi-armed bandit models. *Journal of Machine Learning Research*, 17(1):1–42, 2016.
- 591
- 592 Jaeyeon Kim, Kulin Shah, Vasilis Kontonis, Sham M. Kakade, and Sitan Chen. Train for the  
593 worst, plan for the best: Understanding token ordering in masked diffusions. In *Forty-second*  
*International Conference on Machine Learning*, 2025a. URL [https://openreview.net/](https://openreview.net/forum?id=DjJmre5IkP)  
[forum?id=DjJmre5IkP](https://openreview.net/forum?id=DjJmre5IkP).
- 594
- 595 Jin-Young Kim, Hyojun Go, Soonwoo Kwon, and Hyun-Gyoon Kim. Denoising task difficulty-  
596 based curriculum for training diffusion models. In *The Thirteenth International Conference*  
597 *on Learning Representations*, 2025b. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=96GMFXsbJE)  
598 [96GMFXsbJE](https://openreview.net/forum?id=96GMFXsbJE).

- 594 Juno Kim and Taiji Suzuki. Transformers learn nonlinear features in context: Nonconvex mean-  
595 field dynamics on the attention landscape. In *International Conference on Machine Learning*, pp.  
596 24527–24561. PMLR, 2024.
- 597 Juno Kim and Taiji Suzuki. Transformers provably solve parity efficiently with chain  
598 of thought. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *Inter-  
599 national Conference on Representation Learning*, volume 2025, pp. 44642–44668,  
600 2025. URL [https://proceedings.iclr.cc/paper\\_files/paper/2025/file/  
601 6e2986deda273d8fb903342841fcc4dc-Paper-Conference.pdf](https://proceedings.iclr.cc/paper_files/paper/2025/file/6e2986deda273d8fb903342841fcc4dc-Paper-Conference.pdf).
- 602 Juno Kim, Denny Wu, Jason Lee, and Taiji Suzuki. Metastable dynamics of chain-of-thought rea-  
603 soning: Provable benefits of search, rl and distillation. *arXiv preprint arXiv:2502.01694*, 2025c.
- 604 Nayoung Lee, Ziyang Cai, Avi Schwarzschild, Kangwook Lee, and Dimitris Papailiopoulos. Self-  
605 improving transformers overcome easy-to-hard and length generalization challenges. *arXiv  
606 preprint arXiv:2502.01612*, 2025.
- 607 Hongkang Li, Songtao Lu, Pin-Yu Chen, Xiaodong Cui, and Meng Wang. Training nonlinear trans-  
608 formers for chain-of-thought inference: A theoretical generalization analysis. In *The Thirteenth  
609 International Conference on Learning Representations*, 2025. URL [https://openreview.  
610 net/forum?id=n7n8McETXw](https://openreview.net/forum?id=n7n8McETXw).
- 611 Y. Li, K. Sreenivasan, A. Giannou, D. Papailiopoulos, and S. Oymak. Dissecting chain of thought:  
612 Compositionality through in-context filtering and learning. *Advances in Neural Information Pro-  
613 cessing Systems*, 36:22021–22046, 2023.
- 614 Zhiyuan Li, Hong Liu, Denny Zhou, and Tengyu Ma. Chain of thought empowers transformers to  
615 solve inherently serial problems. *arXiv preprint arXiv:2402.12875*, 2024.
- 616 Huanyu Liu, Jia Li, Chang Yu, Taozhi Chen, Yihong Dong, Lecheng Wang, Hu XiaoLong, and  
617 Ge Li. Evocot: Overcoming the exploration bottleneck in reinforcement learning. *arXiv preprint  
618 arXiv:2508.07809*, 2025a.
- 619 Mingyang Liu, Gabriele Farina, and Asuman Ozdaglar. Uft: Unifying supervised and reinforcement  
620 fine-tuning. *arXiv preprint arXiv:2505.16984*, 2025b.
- 621 Zhou Lu, Hongming Pu, Feicheng Wang, Zhiqiang Hu, and Liwei Wang. The expressive power of  
622 neural networks: A view from the width. In *Advances in neural information processing systems*,  
623 volume 30, 2017.
- 624 Arvind Mahankali, Tatsunori B Hashimoto, and Tengyu Ma. One step of gradient descent is  
625 provably the optimal in-context learner with one layer of linear self-attention. *arXiv preprint  
626 arXiv:2307.03576*, 2023.
- 627 Eran Malach, Pritish Kamath, Emmanuel Abbe, and Nathan Srebro. Quantifying the benefit of using  
628 differentiable learning over tangent kernels. In *International Conference on Machine Learning  
629 (ICML)*, pp. 7379–7389. PMLR, 2021.
- 630 Shie Mannor and John N. Tsitsiklis. The sample complexity of exploration in the multi-armed bandit  
631 problem. *Journal of Machine Learning Research*, 5:623–648, 2004.
- 632 Fanqing Meng et al. Mm-eureka: Exploring the frontiers of multimodal reasoning with rule-based  
633 reinforcement learning. *arXiv preprint arXiv:2503.07365*, 2025.
- 634 William Merrill and Ashish Sabharwal. The expressive power of transformers with chain of  
635 thought. *arXiv preprint arXiv:2310.07923*, 2023.
- 636 Sanmit Narvekar, Jivko Sinapov, and Peter Stone. Curriculum learning in reinforcement learning:  
637 A survey. *Journal of Artificial Intelligence Research*, 68:1–50, 2020.
- 638 Eshaan Nichani, Alex Damian, and Jason D Lee. How transformers learn causal structure with  
639 gradient descent. In *Proceedings of the 41st International Conference on Machine Learning*, pp.  
640 38018–38070, 2024.

- 648 Eshaan Nichani, Jason D. Lee, and Alberto Bietti. Understanding factual recall in transformers via  
649 associative memories. In *The Thirteenth International Conference on Learning Representations*,  
650 2025. URL <https://openreview.net/forum?id=hwSmPOAmhk>.  
651
- 652 Abhishek Panigrahi, Bingbin Liu, Sadhika Malladi, Andrej Risteski, and Surbhi Goel. Progressive  
653 distillation induces an implicit curriculum. *arXiv preprint arXiv:2410.05464*, 2024.  
654
- 655 Shubham Parashar, Shurui Gui, Xiner Li, Hongyi Ling, Sushil Vemuri, Blake Olson, Eric Li,  
656 Yu Zhang, James Caverlee, Dileep Kalathil, et al. Curriculum reinforcement learning from easy  
657 to hard tasks improves llm reasoning. *arXiv preprint arXiv:2506.06632*, 2025.
- 658 B. Peng, S. Narayanan, and C. Papadimitriou. On limitations of the transformer architecture. *First  
659 Conference on Language Modeling*, 2024.  
660
- 661 B. Prystawski, M. Li, and N. Goodman. Why think step by step? reasoning emerges from the  
662 locality of experience. *Advances in Neural Information Processing Systems*, 36:70926–70947,  
663 2023.
- 664 Ran Raz. Fast learning requires good memory: A time-space lower bound for parity learning.  
665 *Journal of the ACM (JACM)*, 66(1):1–18, 2018.  
666
- 667 Luca Saglietti, Stefano Sarao Mannelli, and Andrew Saxe. An analytical theory of curriculum  
668 learning in teacher–student networks. *Journal of Statistical Mechanics: Theory and Experiment*,  
669 2022(11):114014, 2022.
- 670 Stefano Sarao Mannelli, Yaraslau Ivashynka, Andrew Saxe, and Luca Saglietti. Tilting the odds at  
671 the lottery: the interplay of overparameterisation and curricula in neural networks\*. *Journal of  
672 Statistical Mechanics: Theory and Experiment*, 2024(11):114001, nov 2024.  
673
- 674 Amrith Setlur, Chirag Nagpal, Adam Fisch, Xinyang Geng, Jacob Eisenstein, Rishabh Agarwal,  
675 Alekh Agarwal, Jonathan Berant, and Aviral Kumar. Rewarding progress: Scaling automated  
676 process verifiers for LLM reasoning. In *The Thirteenth International Conference on Learning  
677 Representations*, 2025. URL <https://openreview.net/forum?id=A6Y7AqlzLW>.
- 678 Taiwei Shi, Yiyang Wu, Linxin Song, Tianyi Zhou, and Jieyu Zhao. Efficient reinforcement finetun-  
679 ing via adaptive curriculum learning. *arXiv preprint arXiv:2504.05520*, 2025.  
680
- 681 Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally  
682 can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024.  
683
- 684 Petru Soviany, Radu Tudor Ionescu, Paolo Rota, and Nicu Sebe. Curriculum learning: A survey.  
685 *International Journal of Computer Vision*, pp. 1–40, 2022.
- 686 Taiji Suzuki. Adaptivity of deep relu network for learning in besov and mixed smooth besov spaces:  
687 optimal rate and curse of dimensionality. *arXiv preprint arXiv:1810.08033*, 2018.  
688
- 689 Taiji Suzuki and Atsushi Nitanda. Deep learning is adaptive to intrinsic dimensionality of model  
690 smoothness in anisotropic besov space. In *Advances in Neural Information Processing Systems*,  
691 volume 34, pp. 3609–3621, 2021.
- 692 Yuxuan Tong, Xiwen Zhang, Rui Wang, Ruidong Wu, and Junxian He. Dart-math: Difficulty-aware  
693 rejection tuning for mathematical problem-solving. *arXiv preprint arXiv:2407.13690*, 2024.  
694
- 695 Karthik Valmeekam, Matthew Marquez, Sarath Sreedharan, and Subbarao Kambhampati. On the  
696 planning abilities of large language models-a critical investigation. In *Advances in Neural Infor-  
697 mation Processing Systems*, volume 36, pp. 75993–76005, 2023.
- 698 Martin J Wainwright. *High-dimensional statistics: A non-asymptotic Viewpoint*, volume 48. Cam-  
699 bridge university press, 2019.  
700
- 701 Xin Wang, Yudong Chen, and Wenwu Zhu. A survey on curriculum learning. *IEEE Transactions  
on Pattern Analysis and Machine Intelligence*, 2021. To appear / published 2021.

- 702 Zixuan Wang, Eshaan Nichani, Alberto Bietti, Alex Damian, Daniel Hsu, Jason D. Lee, and Denny  
703 Wu. Learning compositional functions with transformers from easy-to-hard data. *arXiv preprint*  
704 *arXiv:2505.23683*, 2025.
- 705 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny  
706 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in*  
707 *neural information processing systems*, 35:24824–24837, 2022.
- 708 Daphna Weinshall and Dan Amir. Theory of curriculum learning, with convex loss functions. *Journal*  
709 *of Machine Learning Research*, 21(222):1–19, 2020.
- 710 Daphna Weinshall, Dan Cohen, and Dotan Amir. Curriculum learning by transfer learning: Theory  
711 and experiments with deep networks. In *Proceedings of the IEEE CVPR*, pp. 5238–5247, 2018.
- 712 Kaiyue Wen, Huaqing Zhang, Hongzhou Lin, and Jingzhao Zhang. From sparse dependence to  
713 sparse attention: unveiling how chain-of-thought enhances transformer sample efficiency. *arXiv*  
714 *preprint arXiv:2410.05459*, 2024.
- 715 Kaiyue Wen, Huaqing Zhang, Hongzhou Lin, and Jingzhao Zhang. From sparse dependence to  
716 sparse attention: Unveiling how chain-of-thought enhances transformer sample efficiency. In  
717 *The Thirteenth International Conference on Learning Representations*, 2025a. URL <https://openreview.net/forum?id=AmEgWDhmTr>.
- 718 Liang Wen, Yunke Cai, Fenrui Xiao, Xin He, Qi An, Zhenyu Duan, Yimin Du, Junchen Liu, Lifu  
719 Tang, Xiaowei Lv, Haosheng Zou, Yongchao Deng, Shousheng Jia, and Xiangzheng Zhang.  
720 Light-rl: Curriculum sft, dpo and rl for long cot from scratch and beyond. *arXiv preprint*  
721 *arXiv:2503.10460*, 2025b.
- 722 Yangzhen Wu, Zhiqing Sun, Shanda Li, Sean Welleck, and Yiming Yang. Inference scaling laws:  
723 An empirical analysis of compute-optimal inference for problem-solving with language models.  
724 *arXiv preprint arXiv:2408.00724*, 2024.
- 725 Wei Xiong, Jiarui Yao, Yuhui Xu, Bo Pang, Lei Wang, Doyen Sahoo, Junnan Li, Nan Jiang, Tong  
726 Zhang, Caiming Xiong, and Hanze Dong. A minimalist approach to llm reasoning: from rejection  
727 sampling to reinforce. *arXiv preprint arXiv:2504.11343*, 2025.
- 728 T. Yang, Y. Huang, Y. Liang, and Y. Chi. In-context learning with representations: Contextual  
729 generalization of trained transformers. *arXiv preprint*, 2024. URL <https://arxiv.org/abs/2408.10147>. *arXiv:2408.10147*.
- 730 Tong Yang, Yu Huang, Yingbin Liang, and Yuejie Chi. Multi-head transformers provably learn  
731 symbolic multi-step reasoning via gradient descent. *arXiv preprint arXiv:508.08222*, 2025.
- 732 Lang Yin, Debangshu Banerjee, and Gagandeep Singh. Data shifts hurt cot: A theoretical study.  
733 *arXiv preprint arXiv:2506.10647*, 2025.
- 734 Yaodong Yu, Sam Buchanan, Druv Pai, Tianzhe Chu, Ziyang Wu, Shengbang Tong, Hao Bai, Yuex-  
735 iang Zhai, Benjamin D. Haeffele, and Yi Ma. White-box transformers via sparse rate reduction:  
736 Compression is all there is? *arXiv preprint arXiv:2311.13110*, 2023.
- 737 Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Yang Yue, Shiji Song, and Gao  
738 Huang. Does reinforcement learning really incentivize reasoning capacity in llms beyond the  
739 base model? *arXiv preprint arXiv:arxiv.org/abs/2504.13837*, 2025.
- 740 R. Zhang, S. Frei, and P. L. Bartlett. Trained transformers learn linear models in-context. *Journal*  
741 *of Machine Learning Research*, 25(49):1–55, 2024a.
- 742 Ruiqi Zhang, Spencer Frei, and Peter L Bartlett. Trained transformers learn linear models in-context.  
743 *Journal of Machine Learning Research*, 25(49):1–55, 2024b.
- 744 Ruiqi Zhang, Daman Arora, Song Mei, and Andrea Zanette. Speed-rl: Faster training of reasoning  
745 models via online curriculum learning. *arXiv preprint arXiv:2506.09016*, 2025.
- 746 Yujun Zhou, Zhenwen Liang, Haolin Liu, Wenhao Yu, Kishan Panaganti, Linfeng Song, Dian Yu,  
747 Xiangliang Zhang, Haitao Mi, and Dong Yu. Evolving language models without labels: Majority  
748 drives selection, novelty promotes variation. *arXiv preprint arXiv:2509.15194*, 2025.

# Appendix: Provable Benefit of Curriculum in Transformer Tree-Reasoning Post-Training

## CONTENTS

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Theoretical Framework for Curriculum Post-training on reasoning trees</b>	<b>3</b>
2.1	Transformer Reasons over a States-Conditioned Reasoning Tree . . . . .	4
2.1.1	Concrete Example: Parity Problem . . . . .	6
2.1.2	Outcome Signal-based RL Finetunings by Gradient Descent . . . . .	7
2.1.3	Outcome Signal-based Test-time Scaling . . . . .	9
<b>3</b>	<b>Conclusion, Discussion, and Future Work</b>	<b>9</b>
<b>A</b>	<b>Additional Related Work</b>	<b>16</b>
<b>B</b>	<b>LLM Usage</b>	<b>16</b>
<b>C</b>	<b>Ethics Statement</b>	<b>16</b>
<b>D</b>	<b>Reproducibility Statement</b>	<b>16</b>
<b>E</b>	<b>Cases beyond transformer gradient dynamics over Tree</b>	<b>17</b>
<b>F</b>	<b>Auxiliary Lemmas</b>	<b>21</b>
<b>G</b>	<b>Details and Proofs of Autoregressive Reasoning Tree</b>	<b>21</b>
G.1	Proof of Representation Theorem . . . . .	22
G.2	Proof of Finetuning Algorithms . . . . .	23
G.3	Proof of Test-time Scaling . . . . .	40
<b>H</b>	<b>Proofs of Parity Problem</b>	<b>45</b>

## A ADDITIONAL RELATED WORK

## B LLM USAGE

In accordance with the conference policy on large language models (LLMs), we disclose that an LLM was used solely as a writing aid to improve the clarity and readability of certain paragraphs. The LLM did not contribute to research ideation, technical development, experimental design, analysis, or interpretation of results. All scientific ideas, methods, and conclusions presented in this paper are entirely the work of the authors.

## C ETHICS STATEMENT

This work is purely theoretical and algorithmic in nature. It does not involve human subjects, personally identifiable data, or sensitive datasets, and therefore does not pose direct risks to privacy, security, or safety. Our study is conducted under the principles of academic integrity and transparency, adhering to the ICLR Code of Ethics. We further note that our methods and results do not introduce discriminatory or harmful applications, nor are they tied to any conflicts of interest or external sponsorship that may bias the research.

## D REPRODUCIBILITY STATEMENT

We have taken concrete steps to ensure reproducibility of our results. All theoretical claims are accompanied by formal assumptions and complete proofs in the appendix. The complexity analyses and derivations are presented in full detail, ensuring transparency and verifiability. For experimental validations (if applicable), datasets are publicly available and preprocessing steps are explicitly described in the supplementary materials. To further support reproducibility, we provide pseudo-code and detailed algorithmic specifications in Alg. 2, 3, 1, 5, 4, 6, and will release anonymous source code in the supplementary materials to facilitate independent verification.

**Theoretical benefits of curriculum in post-training.** A growing body of work has provided empirical evidence for the effectiveness of curriculum strategies during post-training (Liu et al., 2025a; Lee et al., 2025; Bae et al., 2025; Meng et al., 2025; Shi et al., 2025; Wen et al., 2025b; Zhang et al., 2025; Amani et al., 2025). Zhang et al. (2025) offered a theoretical perspective, showing that intermediate-difficulty questions yield higher signal-to-noise ratios for gradient estimates. Closely related, Liu et al. (2025b) modeled LLM reasoning as a search tree (see Figure 1) and proved that uniformly sampling a hint depth (i.e., revealing the CoT up to that depth and letting the LLM complete the rest) reduces the exploration complexity of achieving  $\geq 50\%$  pass@1 from exponential to polynomial (see Remark 8). In practice, they adopted a cosine scheduler to provide longer hints in early stages and gradually shorten them during fine-tuning. Similarly, Parashar et al. (2025) analyzed curriculum within Approximate Policy Iteration: under exponentially decaying assumptions on policy-approximation error and performance loss, they proved gains in sample complexity and empirically compared uniform, cosine, and Gaussian schedulers. However, none of these studies explicitly models the difficulty measure in post-training — namely, the *pass-rate* and its dependence on the base model — nor do they consider transformer’s states-conditioned tree-like autoregressive reasoning. Their results largely focus on optimization endpoints argument (i.e.,  $\arg \min$  solutions assuming gradient methods have converged), without analyzing the underlying post-training dynamics or the transformer architecture. In contrast, our framework is built directly on the *pass-rate* and the probabilistic tree structure of post-training reasoning, leading to Theorem 1 and its case studies on transformers.

**Theoretical benefits of curriculum learning in pre-training for parity and beyond.** Curriculum Learning (CL) was first introduced by Bengio et al. (2009), and has since been empirically validated across vision, NLP, and reinforcement learning (Graves et al., 2017; Narvekar et al., 2020; Wang et al., 2021; Soviany et al., 2022). On the theoretical side, most analyses focus on *pre-training* with highly specific problem classes or architectures. A prominent line of work studies the *parity problem*: while parity can be solved efficiently by Gaussian elimination over  $\mathbb{F}_2$  (Abbe & Sandon, 2020), gradient-based neural networks struggle on dense data due to precision barriers (Abbe et al., 2021). To circumvent this, Malach et al. (2021) showed that sparse parities become learnable

864 with a one-layer network augmented by a parity module, while Daniely & Malach (2020) analyzed  
 865 two-layer fully connected nets where sparse inputs with *leaky labels* help learning. Cornacchia &  
 866 Mossel (2023) proved that sparse parities can be learned by a two-layer net with one-step gradients  
 867 but require a large learning rate, and Abbe et al. (2023b) gave empirical evidence that presenting  
 868 sparse-to-dense samples as a curriculum improves generalization, though without formal guaran-  
 869 tees. In convex settings, Weinshall et al. (2018); Weinshall & Amir (2020) proved that curriculum  
 870 can accelerate the convergence of SGD. Building on these observations, Abbe et al. (2023c) ana-  
 871 lyzed a two-layer ReLU network trained with noisy SGD based on mixed data curriculum where  
 872 the sparse (lower hamming weights with fewer  $-1$ ) data is first learned, and proved that a *layer-wise*  
 873 *curriculum* starting from sparse samples reduces both time and sample complexity. In contrast to  
 874 data curricula Panigrahi et al. (2024) considered teacher-student setting where the student would  
 875 benefit from teacher model’s checkpoints, serving as the internal signals for students to learn parity  
 876 support. Statistically, Sarao Mannelli et al. (2024) showed the negative effect of overparameter-  
 877 ization for curriculum learning, in the case of a XOR-like Gaussian Mixture problem. Beyond  
 878 parity, Saglietti et al. (2022) studied a teacher–student model where the target depends on a sparse  
 879 set of features, and curriculum is defined via the variance of irrelevant features (low variance = easy,  
 880 high variance = hard). More recently, Wang et al. (2025) investigated curriculum learning for trans-  
 881 formers in the *k-fold composition task* setting, where they showed that a carefully staged pre-training  
 882 schedule—both in terms of data distribution and layer-wise progression—enables efficient learning.  
 883 Taken together, these works demonstrate that curriculum can provably help in pre-training, but only  
 884 under tailored data curricula (e.g., sparse input with lower Hamming weight where probability of  $-1$   
 885 is mild, leaky labels, checkpoints), specialized architectures (e.g., augmented by a parity modules),  
 886 or non-standard optimization schemes (e.g., noisy SGD, layer-wise training). In contrast, our work  
 887 studies curriculum in the *post-training* regime, where we leverage the base model’s coverage power  
 888 and pre-trained *chain-of-thought reasoning* ability, *without exploiting handcrafted data correlations*  
 889 or requiring algorithmic modifications.

889 **Theoretical Benefit of CoT.** A substantial literature has emerged on the theoretical role of chain-  
 890 of-thought (CoT) in enhancing transformer models. Several works demonstrate that incorporat-  
 891 ing polynomial-length intermediate steps expands the expressive capacity of transformers beyond  
 892 constant-depth architectures (Feng et al., 2023; Li et al., 2024; Merrill & Sabharwal, 2023; Wei  
 893 et al., 2022). Complementing these positive results, a parallel line of research investigates intrinsic  
 894 barriers, proving lower bounds or rank-based constraints that necessitate non-trivial reasoning  
 895 depth (Peng et al., 2024; Barceló et al., 2025; Amiri et al., 2025). Further, statistical and compo-  
 896 sitional analyses reveal how CoT encourages structured generalization and compositional filtering,  
 897 thereby clarifying its benefit from a distributional perspective (Prystawski et al., 2023; Li et al.,  
 898 2023). Building on these investigations, we studied the complexity improvement when the CoT  
 899 evolution is obtained in a curriculum manner.

900 **Training Dynamics of Transformers.** Beyond expressiveness, recent studies examine how trans-  
 901 formers trained with CoT evolve during optimization. Early work explored emergent in-context  
 902 learning and convergence under multi-head attention (Hahn & Goyal, 2023; Huang et al., 2024;  
 903 Yang et al., 2024; Zhang et al., 2024a), while subsequent analyses provide provable characteriza-  
 904 tions of induction heads and sparse token mechanisms (Chen et al., 2024). Of particular relevance  
 905 are studies on CoT-specific training dynamics, which analyze nonlinear or single-layer transformers  
 906 solving structured tasks such as parity (Wen et al., 2024; Kim & Suzuki, 2025; Li et al., 2025). More  
 907 recent work extends these insights to multi-step optimization and regular language recognition, un-  
 908 covering implicit biases and algorithmic behaviors in gradient descent (Huang et al., 2025a;b;c).  
 909 Wen et al. (2025a); Kim & Suzuki (2025) investigated the learning of transformers on the parity task  
 910 with CoT, building on which Yin et al. (2025) studied the impact of data shift. Yang et al. (2025)  
 911 further demonstrated that transformers can implement both forward and backward *tree-structured*  
 912 symbolic reasoning using two attention heads. However, these works abstract away real-world sam-  
 913 pling process during inference, relying instead on *atypical deterministic prediction of internal to-*  
 914 *kens/sentences*, and they primarily analyze pre-training rather than post-training.

## 915 E CASES BEYOND TRANSFORMER GRADIENT DYNAMICS OVER TREE

916 **Proof.***Proof of Theorem 1.*  
 917

By the scaling assumption equation 3, there exist absolute constants  $m, M > 0$  such that for all  $k < k'$ ,

$$m \left\| \frac{\pi_{k'}^*}{\pi_k^*} \right\|_{\infty} \leq \mathcal{C}(\pi_{k'}^* | \pi_k^*) \leq M \left\| \frac{\pi_{k'}^*}{\pi_k^*} \right\|_{\infty}.$$

Direct route. Taking  $(k, k') = (0, K)$  and invoking equation 2 yields

$$\mathcal{C}_{\text{direct}} = \mathcal{C}(\pi_K^* | \pi_0^*) \in \Theta\left(\left\| \frac{\pi_K^*}{\pi_0^*} \right\|_{\infty}\right) = \Theta((C^*)^K).$$

Curriculum route. Consider the stepwise procedure that estimates  $\pi_{k+1}^*$  from  $\pi_k^*$  for  $k = 0, 1, \dots, K-1$ . Its total cost is

$$\mathcal{C}_{\text{curriculum}} := \sum_{k=0}^{K-1} \mathcal{C}(\pi_{k+1}^* | \pi_k^*).$$

By equation 3 and the stepwise bound equation 1, for each  $k$  we have

$$m \Theta(C^*) \leq \mathcal{C}(\pi_{k+1}^* | \pi_k^*) \leq M \Theta(C^*),$$

hence there exist constants  $c_1, c_2 > 0$  such that

$$c_1 C^* \leq \mathcal{C}(\pi_{k+1}^* | \pi_k^*) \leq c_2 C^* \quad \text{for all } k.$$

Summing over  $k$  gives

$$c_1 K C^* \leq \mathcal{C}_{\text{curriculum}} \leq c_2 K C^*,$$

i.e.,  $\mathcal{C}_{\text{curriculum}} = \Theta(K C^*)$ .

Combining the two parts shows that the curriculum transforms the direct exponential scaling  $\Theta((C^*)^K)$  in task depth into a polynomial (linear) scaling  $\Theta(K C^*)$ , as claimed.  $\square$

**Theorem 5** (Curriculum Learning with Spanner Sampling). Consider the theoretical settings in Foster et al. (2025) (realizable linear softmax parameterization), consider a curriculum of  $K$  tasks with increasing difficulty, starting from a base model  $\pi_{\text{ref}} := \pi_0^*$ . Suppose the sequence of optimal policies  $\{\pi_k^*\}_{k=1}^K$  satisfies assumptions in Eq.(1), (2), namely

$$\begin{aligned} \mathcal{C}_{\text{cov}}(\pi_k^* | \pi_{k-1}^*) &= \Theta(C^*), \quad \forall k \in [K], \\ \mathcal{C}_{\text{cov}}(\pi_{k_2}^* | \pi_{k_1}^*) &= \Theta((C^*)^{k_2 - k_1}). \end{aligned}$$

For any  $\epsilon > 0$  and  $\delta \in (0, 1)$ , by applying SpannerSampling in Foster et al. (2025) **sequentially** through the curriculum with appropriate choices of  $T_{\text{prompt}}, T_{\text{span}},$  and  $T_{\text{exp}}^k$  for each task, the curriculum learning algorithm learns a policy  $\hat{\pi}$  with:

$$\mathbb{E}_{\hat{\pi}_k \sim \text{unif}(\hat{\pi}^{(T_{\text{exp}}^{k-1})}, \dots, \hat{\pi}^{(T_{\text{exp}}^k)})} [J_{\beta}(\pi_k^*) - J_{\beta}(\hat{\pi}_k)] \leq \epsilon, \quad k \in [K]$$

with probability at least  $1 - \delta$ , and achieves the following computational efficiency bound:

$$T_{\text{comp}}^{\text{curriculum}}(\epsilon, \delta) = \tilde{O}\left(K \cdot C^* \cdot \frac{R_{\text{max}}^2}{\beta^2} \left[ \frac{R_{\text{max}}^2}{\beta} \cdot \frac{d^2 \log^2(\delta^{-1})}{\min\{\epsilon, \beta\}} \right]\right).$$

Moreover, compared to direct estimation from  $\pi_{\text{ref}}$  to  $\pi_K^*$ , which is

$$T_{\text{comp}}^{\text{direct}}(\epsilon, \delta) = \tilde{O}\left((C^*)^K \cdot \frac{R_{\text{max}}^2}{\beta^2} \left[ \frac{R_{\text{max}}^2}{\beta} \cdot \frac{d^2 \log^2(\delta^{-1})}{\min\{\epsilon, \beta\}} \right]\right).$$

Therefore, calling Spanner Sampling in a curriculum manner achieves exponential improvement in computational complexity:

$$\frac{T_{\text{comp}}^{\text{direct}}}{T_{\text{comp}}^{\text{curriculum}}} = \Omega\left(\frac{(C^*)^{K-1}}{K}\right)$$

where the direct estimation requires  $T_{\text{comp}}^{\text{direct}} = \tilde{O}\left(\mathcal{C}_{\text{cov}}(\pi^*) \cdot \frac{R_{\text{max}}^2}{\beta^2}\right) \cdot (C^*)^K$ .

**Proof.** We work under the realizable linear-softmax parameterization and oracle model of Foster et al. (2025). Let the KL-regularized objective be  $J_\beta(\cdot)$  and let  $R_{\max}$ ,  $d$ , and  $\beta$  denote the reward bound, feature dimension, and temperature, respectively. For any pair of policies  $(\pi, \pi')$ , write the coverage coefficient as  $C_{\text{cov}}(\pi | \pi') := \left| \frac{\pi}{\pi'} \right|_\infty$ .

**Step 1 (Per-task guarantee from Foster et al. (2025)).** By Theorem 3.1 of Foster et al. (2025), SpannerSampling, when run to target  $\pi_k^*$  starting from a reference  $\pi_{k-1}^*$  in the linear-softmax model, returns a policy  $\hat{\pi}$  such that

$$\mathbb{E}_{\hat{\pi} \sim \text{unif}(\hat{\pi}^{(1)}, \dots, \hat{\pi}^{(T_{\text{exp}})})} [J_\beta(\pi_k^*) - J_\beta(\hat{\pi})] \leq \epsilon$$

with probability at least  $1 - \delta'$ , using a number of oracle computations

$$T_{\text{comp}}^{(k)}(\epsilon, \delta') = \tilde{O}\left(C_{\text{cov}}(\pi_k^* | \pi_{k-1}^*) \cdot \underbrace{\frac{R_{\max}^2}{\beta^2} \left[ \frac{R_{\max}^2}{\beta} \cdot \frac{d^2 \log^2((\delta')^{-1})}{\min\{\epsilon, \beta\}} \right]}_{=: \Gamma(\epsilon, \delta')}\right),$$

where  $\tilde{O}(\cdot)$  hides polylogarithmic factors not depending on the coverage coefficient.<sup>3</sup> We emphasize that data efficiency  $T_{\text{data}}$  is independent of coverage, while computational efficiency scales with coverage; SpannerSampling matches the linear-in-coverage lower bound.

**Step 2 (Scheduling across a curriculum).** Apply Step 1 sequentially for tasks  $k = 1, 2, \dots, K$ , using fresh budgets  $(T_{\text{prompt}}, T_{\text{span}}^k, T_{\text{exp}}^k)$  at each stage, and set the per-stage failure probability to  $\delta' = \delta/K$ . The uniform-over-iterates evaluation in the statement mirrors the “return a uniformly random iterate” prescription in Foster et al. (2025), so the stagewise accuracy guarantee transfers verbatim. A union bound then yields that, simultaneously for all  $k \in [K]$ ,

$$\mathbb{E}_{\hat{\pi}_k \sim \text{unif}(\hat{\pi}^{(T_{\text{exp}}^{k-1})}, \dots, \hat{\pi}^{(T_{\text{exp}}^k)})} [J_\beta(\pi_k^*) - J_\beta(\hat{\pi}_k)] \leq \epsilon$$

with probability at least  $1 - \delta$ . The total compute is the sum of stagewise costs:

$$T_{\text{comp}}^{\text{curriculum}}(\epsilon, \delta) = \sum_{k=1}^K T_{\text{comp}}^{(k)}(\epsilon, \delta/K) = \tilde{O}\left(\Gamma(\epsilon, \delta/K) \cdot \sum_{k=1}^K C_{\text{cov}}(\pi_k^* | \pi_{k-1}^*)\right).$$

**Step 3 (Using the coverage structure of the curriculum).** By the curriculum assumptions (1)–(2), we have  $C_{\text{cov}}(\pi_k^* | \pi_{k-1}^*) = \Theta(C^*)$  for each  $k$ , hence

$$T_{\text{comp}}^{\text{curriculum}}(\epsilon, \delta) = \tilde{O}(K C^* \cdot \Gamma(\epsilon, \delta/K)).$$

On the other hand, direct estimation of  $\pi_K^*$  from  $\pi_{\text{ref}} = \pi_0^*$  is controlled by the single coverage factor  $C_{\text{cov}}(\pi_K^* | \pi_0^*) = \Theta((C^*)^K)$  (by 2), and Theorem 3.1 gives

$$T_{\text{comp}}^{\text{direct}}(\epsilon, \delta) = \tilde{O}((C^*)^K \cdot \Gamma(\epsilon, \delta)).$$

**Step 4 (Comparing the two strategies).** Combining the above displays yields

$$\frac{T_{\text{comp}}^{\text{direct}}(\epsilon, \delta)}{T_{\text{comp}}^{\text{curriculum}}(\epsilon, \delta)} = \Omega\left(\frac{(C^*)^{K-1}}{K}\right),$$

i.e., an exponential-to-polynomial improvement as a function of  $K$ , as claimed. This recovers the abstract form of Theorem 1 with the concrete instantiation of the per-stage computational complexity provided by Foster et al. (2025).  $\square$

<sup>3</sup>We only need that  $T_{\text{comp}}$  is (i) polynomial in  $(d, \beta^{-1}, \epsilon^{-1}, \log(1/\delta'))$  and (ii) linear (up to logs) in  $C_{\text{cov}}$ , as established by Theorem 3.1 together with the matching lower bound in Foster et al. (2025).

**Remark 6** (Exponential–Polynomial Gaps Under Relaxed Conditions). *The exponential-to-polynomial separation in Theorem 5 remains valid under relaxed coverage conditions, analogous to the discussion after Theorem 1. Specifically, suppose that: (i) for some  $k \in [K]$ , the per-step coverage is only polynomial rather than constant, i.e.,*

$$C_{\text{cov}}(\pi_k^* \mid \pi_{k-1}^*) = (C^*)^{p_k}, \quad 1 \leq p_k \ll K,$$

*so that the per-task computational complexity grows as  $\tilde{O}((C^*)^{p_k} \cdot \Gamma(\epsilon, \delta))$ ; (ii) the direct coverage factor accumulates less than  $(C^*)^K$ , relaxing to*

$$C_{\text{cov}}(\pi_K^* \mid \pi_0^*) = \Theta((C^*)^{bK-c}), \quad 1 \leq b \ll K, c \ll K.$$

*Then, under Spanner Sampling, the curriculum complexity satisfies*

$$T_{\text{comp}}^{\text{curriculum}} = \tilde{O}\left(K \cdot \max_{k \in [K]} (C^*)^{p_k} \cdot \Gamma(\epsilon, \delta)\right) \ll \tilde{\Theta}((C^*)^{bK-c} \cdot \Gamma(\epsilon, \delta)) = T_{\text{comp}}^{\text{direct}}.$$

*Thus, even under relaxed conditions, the curriculum strategy with Spanner Sampling retains an exponential advantage in computational complexity compared to direct estimation.*

**Remark 7** (Relation to E2H (CRL) theory and limits for post-training). *Parashar et al. (2025) analyze curriculum RL for LLMs under an Approximate Policy Iteration (API) lens and derive a stagewise sample complexity bound (their Thm. 3.2). In particular, letting  $\epsilon_k$  be per-stage accuracy targets and  $L_k$  distribution–sensitivity constants, they obtain  $M_{\text{CRL}} = \sum_{k=1}^K \tilde{O}(\log^3(1/\epsilon_k) \epsilon_k^{-2} L_k^2 \cdots)$  and compare it to direct learning via a “curriculum efficiency factor” (CEF). Under geometric schedules  $\epsilon_k = \epsilon_K e^{K-k}$  and  $L_k = L_K/l^{K-k}$ , they show*

$$M_{\text{CRL}} < M_{\text{Direct}} \iff \frac{(el)^{2(1-K)} - 1}{1 - (el)^2} < m - 1,$$

*where  $m > 1$  encodes the relative hardness of direct learning (their Eq. (2)). Conceptually, this captures that a well–designed curriculum can turn a multiplicative blow-up across stages into a controlled (geometric) sum. While Parashar et al. (2025) do motivate curricula by a distribution gap between the pretraining distribution  $d_0$  and task distribution  $d_K$ , their API analysis does not instantiate a density-/likelihood-ratio–type coverage term nor a dependence on the base model beyond the abstract factor  $m$ . Hence their theory, though consonant with ours at a high level (curricula beat direct learning), does not yet provide a coverage–based explanation of post-training that connects “good foundations” (a strong base policy) to concrete computational savings; Theorem 5 fills precisely this gap by making the dependence on coverage explicit.*

**Remark 8** (Relation to UFT complexity measure). *Liu et al. (2025b) propose a complexity notion tailored to unified finetuning: to achieve a 50% pass@1 success rate, the algorithm must explore at least a certain number of nodes in the search space  $S_H$ . Under this definition, they show that without a curriculum, the exploration complexity can grow exponentially with task depth, whereas curriculum strategies help avoid such exponential bottlenecks. Conceptually, this aligns with our message in Theorem 3, namely that curricula convert exponential costs into polynomial ones.*

*However, their measure does not explicitly account for (i) the number and structure of subtasks, (ii) the relative difficulty of these subtasks, or (iii) the role of a coverage coefficient that quantifies how well a base model supports subsequent targets. Moreover, the UFT framework does not yield a sample–complexity characterization directly, but instead works with node–exploration counts in  $S_H$ . In contrast, Theorem 3 provides a coverage–based explanation of post-training, connecting subtask decomposition and curriculum schedules to concrete sample–complexity bounds. This makes our result more directly interpretable in terms of the efficiency of real post-training pipelines.*

**Remark 9** (Relation to curriculum via sparse→mixed training in ReLU SGD). *Abbe et al. (2023c) study a model-specific curriculum for learning  $k$ -parities with a two-layer ReLU network trained by (noisy) gradient descent. Their curriculum is defined by an ordering of input distributions: the network is first trained on sparse inputs, then on the full mixed distribution  $D_{\text{mix}} = \rho D_\mu + (1 - \rho) D_w$ , rather than by changing the task or exploiting a foundation model. Under bounded learning rates and a small fraction of sparse inputs, they prove a separation in the number of training steps: curriculum noisy-GD/SGD (sparse-first) learns  $k$ -parities in  $\tilde{O}(d)$  (or  $\tilde{O}(d)/\epsilon^2$  under accuracy  $\epsilon$ ), whereas training on randomly ordered (mixed) samples requires at least polynomially more steps (e.g.,  $\tilde{\Omega}(d^{1+\delta})$  or  $\tilde{\Omega}(d^2)$  in stated regimes).*

Conceptually, this differs from Theorem 3. Our result is a post-training analysis in the realizable linear–softmax setting, where a coverage coefficient  $C_{\text{cov}}(\pi_k^* \mid \pi_{k-1}^*)$  makes explicit how a good base policy (pretraining/foundation) reduces the computational complexity (e.g. Theorem. 5) from  $\tilde{O}((C^*)^K)$  (direct) to  $\tilde{O}(KC^*)$  (curriculum). In contrast, (Abbe et al., 2023c) do not model pre-training or any coverage/likelihood-ratio notion measuring proximity to the target policy; their curriculum is instead an input-scheduling device inside noisy-GD for a specific ReLU architecture. Thus, while both works support the high-level message that curricula can overcome training bottlenecks, (Abbe et al., 2023c) provide a step-complexity separation for a fixed neural model trained from scratch, whereas Theorem 5 offers a coverage-based explanation of post-training improvements tied to the quality (e.g. reasoning CoT power per Theorem. 3) of the base model.

## F AUXILIARY LEMMAS

**Classical probability tools for fixed-confidence identification.** We collect standard concentration and information-theoretic lemmas used in Step 6 of Prop. 1.

**Lemma 1** (Hoeffding/Chernoff inequality for Bernoulli means (Wainwright, 2019)). *Let  $X_1, \dots, X_m$  be i.i.d. Bernoulli( $p$ ), and  $\hat{p}_m := \frac{1}{m} \sum_{r=1}^m X_r$ . Then for any  $\epsilon \in (0, 1)$ ,*

$$\mathbb{P}(|\hat{p}_m - p| \geq \epsilon) \leq 2 \exp(-2m\epsilon^2).$$

Equivalently, to ensure  $\mathbb{P}(|\hat{p}_m - p| \geq \epsilon) \leq \delta$  it suffices that  $m \geq \frac{1}{2\epsilon^2} \log \frac{2}{\delta}$ .

**Proof.** See Hoeffding (1963), Chernoff (1952), or Exercise 2.9.(a) in Wainwright (2019).  $\square$

**Lemma 2** (Two-arm fixed-confidence identification lower bound). *Consider two Bernoulli arms with means  $p_*$  and  $p$ , and gap  $\Delta := p_* - p > 0$ . Any (possibly adaptive)  $\delta$ -correct procedure that outputs the better arm with probability at least  $1 - \delta$  must satisfy*

$$\mathbb{E}[T] \geq c \frac{\log(1/\delta)}{\text{KL}(\text{Bern}(p_*) \parallel \text{Bern}(p))} \geq c' \Delta^{-2} \log(1/\delta),$$

for universal constants  $c, c' > 0$ , where the second inequality uses Pinsker’s bound  $\text{KL}(\text{Bern}(p_*) \parallel \text{Bern}(p)) \leq 2(p_* - p)^2$ .

**Proof.** See Thm. 5 in Mannor & Tsitsiklis (2004) and Thm. 1 in Kaufmann et al. (2016).  $\square$

**Corollary 2** (Per-depth best-arm identification complexity). *Fix depth  $\ell$  with  $|\mathcal{I}_\ell| = \Theta(d)$  Bernoulli arms and assume a unique best arm with acceptance-gap  $\Delta > 0$  to all others. Any  $\delta$ -correct identification at depth  $\ell$  requires at least*

$$\Omega(\Delta^{-2} \log(d/\delta))$$

oracle observations in expectation.

**Proof.** This follows by applying Lemma 2 to each suboptimal-vs-best pair and distributing the confidence via union bound (or, sufficiency with uniform sampling via Lemma 1).  $\square$

## G DETAILS AND PROOFS OF AUTOREGRESSIVE REASONING TREE

**Definition 3** (Full Version of 2-States Conditioned Autoregressive Reasoning Tree (2S-ART)). *Let  $[K] := \{1, \dots, K\}$  be the dictionary and let  $\text{EOS} := K+1$ . Fix a depth  $L \in \mathbb{N}$ . A 2S-ART is specified by maps*

$$\{\Phi_l\}_{l=1}^L, \quad \{\mathcal{I}_l\}_{l=1}^{L+1},$$

and induces a function class  $\mathcal{F}_{2\text{S-ART}}$  of autoregressive tasks  $f : [K]^d \rightarrow [K+1]^{\leq L+1}$ . Given an input  $\mathbf{x} = (x_1, \dots, x_d) \sim \text{Unif}([K]^d)$ , define the output chain-of-thought (CoT) by the following iterative rule.

**State and selectors.** For step  $l = 1, 2, \dots$ , let

$$\text{CoT}_{l-1} := (x_1, \dots, x_d, \text{EOS}, z_1, \dots, z_{l-1}) \in [K+1]^{d+l}, \quad z_0 := \text{EOS}.$$

The legal index selector  $\mathcal{I}_l : [K+1]^{d+l-1} \rightarrow 2^{[d+l]}$  returns a set  $\mathcal{I}_l(\mathbf{CoT}_{l-1}) \subseteq [d+l]$  with  $|\mathcal{I}_l(\mathbf{CoT}_{l-1})| = \Theta(d)$  for all  $l \leq L$ , and the terminal selector satisfies  $\mathcal{I}_{L+1}(\cdot) = \{d+1\}$ . At step  $l$ , choose an index  $i_l \in \mathcal{I}_l(\mathbf{CoT}_{l-1})$  and let  $v_l := (\mathbf{CoT}_{l-1})_{i_l} \in [K+1]$ .

**Two-state update.** Each update is computed from the previous reasoning state and the chosen clue:

$$\Phi_l : [K+1] \times [K+1] \rightarrow [K+1], \quad z_l := \Phi_l(z_{l-1}, v_l), \quad l = 1, 2, \dots$$

with the EOS-absorbing property

$$\Phi_l(z, \text{EOS}) = \text{EOS} \quad \text{for all } z \in [K+1], \quad l = 1, \dots, L.$$

Hence, once  $v_l = \text{EOS}$ , all subsequent states remain EOS, denoting the end of the output.

**Output.** Let  $\ell := \min\{l \leq L : z_l = \text{EOS}\}$  (with the convention  $\ell = L$  if no such  $l$  exists). The task  $f$  outputs the EOS-terminated CoT

$$f(\mathbf{x}) = (z_1, \dots, z_\ell, \text{EOS}).$$

Because  $|\mathcal{I}_l(\mathbf{CoT}_{l-1})| = \Theta(d)$  for  $l \leq L$ , each step selects from  $\Theta(d)$  possible clues, giving a tree with branching  $\Theta(d)$ .

**Subtask family.** For a target task  $f_{S_\star} \in \mathcal{F}_{2S\text{-ART}}$  with realized index path  $S_\star = (i_1, \dots, i_{k^\star}, d+1)$ , define prefix paths  $S_\star^m := (i_1, \dots, i_m, d+1)$  for  $m = 1, \dots, k^\star$ . The associated subtask family is

$$\mathcal{F}_{S_\star} := \{f_{S_\star^1}, \dots, f_{S_\star^{k^\star}}\} \subset \mathcal{F}_{2S\text{-ART}},$$

where each  $f_{S_\star^m}$  is induced by the same  $\{\Phi_l\}, \{\mathcal{I}_l\}$  but terminates immediately after choosing the  $m$ -th index (the next choice is forced to  $d+1$ ).

## G.1 PROOF OF REPRESENTATION THEOREM

**Proof.** Proof of Corollary 1. At each depth  $t = 1, \dots, l$ , the probability of selecting the unique legal child consistent with  $S_\star^t$  is  $\Theta(d^{-1})$  by the 2S-ART definition (near-uniform over  $|\mathcal{I}_t| = \Theta(d)$  legal indices). After producing  $z_l$ , the next step must legally select EOS to terminate, which occurs with probability  $\Theta(d^{-1})$ . By the chain rule of conditional probabilities along the unique legal path, the total probability is the product of  $(l+1)$  factors  $\Theta(d^{-1})$ , i.e.,  $\Theta(d^{-(l+1)})$ .  $\square$

**Proof.** Proof of Theorem 2. Under the above assumption on  $\text{FFN}_l$ , it remains to realize the PART index-sampling policy via attention. It suffices to choose parameters satisfying

$$\mathbf{K}^\top \mathbf{Q} = \begin{pmatrix} \mathbf{0}_{d_x \times d_x} & \mathbf{0}_{d_x \times (d_E - d_x)} \\ \mathbf{0}_{(d_E - d_x) \times d_x} & \mathbf{W} \end{pmatrix}, \quad \mathbf{V} = (\mathbf{I}_{d_x} \quad \mathbf{0}), \quad (11)$$

so that  $\mathbf{V} \mathbf{E}[x_m] = p_m$  (position only) for any token  $x_m$ , and attention logits depend only on positional encodings:  $\mathbf{E}[u]^\top \mathbf{K}^\top \mathbf{Q} \mathbf{E}[z_{l-1}] = \mathbf{p}_k^\top \mathbf{W} \mathbf{p}_{d+1+l}$ ,  $\mathbf{p}_{d+1+l} := \mathbf{p}_{i_l}$  if  $u$  is at position  $k$ . Moreover, for each  $l$  choose constants  $c_l > 0$  and impose a legality mask that sets logits of all  $k \notin \mathcal{I}_l(z_{<l})$  to  $-\infty$ , while ensuring

$$\mathbf{p}_k^\top \mathbf{W} \mathbf{p}_{d+1+l} = c_l \quad \text{for all } k \in \mathcal{I}_l(z_{<l}). \quad (12)$$

Because  $\mathbf{p}_i \perp \mathbf{U}$  and the block structure of  $\mathbf{K}^\top \mathbf{Q}$  removes content–content and content–position interactions, the attention score from any memory token at position  $k$  to the query  $\mathbf{E}[z_{l-1}]$  depends only on positions:  $s(k, l) = \mathbf{p}_k^\top \mathbf{W} \mathbf{p}_{d+1+l}$ . By construction and the legality mask,  $s(k, l) = c_l$  for  $k \in \mathcal{I}_l(z_{<l})$  and  $s(k, l) = -\infty$  otherwise (causal masking already suppresses all positions beyond  $d+l$ ). Hence the softmax distribution over legal keys is exactly uniform.

With  $\mathbf{V} = [\mathbf{I}_{d_x} \quad \mathbf{0}]$ , values contribute only content: for any legal position  $k$ ,  $\mathbf{V} \mathbf{E}[v_l(k)] = \boldsymbol{\mu}^{v_l(k)}$ . Uniform attention over the legal set yields

$$\hat{\mathbf{z}}_l = \frac{1}{|\mathcal{I}_l(z_{<l})|} \sum_{k \in \mathcal{I}_l(z_{<l})} \boldsymbol{\mu}^{v_l(k)}.$$

Choose  $\mathbf{U}$  with orthonormal columns. Then for any  $w$  in the legal set,  $\hat{\mathbf{z}}_l^\top \boldsymbol{\mu}^{v_l(w)} = |\mathcal{I}_l(z_{<l})|^{-1}$  and for any token outside the legal set the inner product is 0. Therefore  $\text{softmax}(\hat{\mathbf{z}}_l^\top \mathbf{U} / \beta)$  places

1188 equal mass on legal content tokens and zero elsewhere, matching the PART policy  $\Pr(i_l | z_{<l}) =$   
 1189  $|\mathcal{I}_l(z_{<l})|^{-1}$ . By assumption, for each depth  $l$  there exists  $\text{FFN}_l$  such that  $\text{FFN}_l(\boldsymbol{\mu}^{v_l(i_l)}, \mathbf{E}[z_{l-1}:d_x])$   
 1190 implements  $\Phi_l(z_{l-1}, v_l(i_l))$  on embeddings, yielding  $\mathbf{E}[z_l] = [\boldsymbol{\mu}^{z_l}, \mathbf{p}_{d+1+l}]^\top$ . At  $l=L$ , set the le-  
 1191 gality mask to permit only  $k = d+1$  (EOS), which forces termination as in Definition 1. The  
 1192 construction exactly replicates PART.  $\square$   
 1193

## 1194 G.2 PROOF OF FINETUNING ALGORITHMS

1195  
 1196 **Lemma 3.** Fix a step  $l$ . Define masked attention logits and weights:

$$1197 \quad s_l(j) := \mathbf{p}_j^\top \mathbf{W} \mathbf{p}_{c_l} + m_l(j), \quad \alpha_l(j) := \frac{\exp(s_l(j))}{\sum_{q \in \mathcal{I}_l} \exp(s_l(q))}, \quad j \in \mathcal{I}_l,$$

1198 where  $m_l$  encodes causal/legal masking and is independent of  $\mathbf{W}$ . Define the attention-weighted  
 1199 token vector and vocabulary logits:

$$1200 \quad \hat{\mathbf{p}}_{l+1}^{\text{att}} := \sum_{j \in \mathcal{I}_l} \alpha_l(j) \mathbf{p}_j, \quad \ell_r := \frac{\langle \hat{\mathbf{p}}_{l+1}^{\text{att}}, \mathbf{p}_r \rangle}{\beta}, \quad r \in \{1, \dots, d+1\}.$$

1201 Let  $p_{\text{vocab}}(r) := \exp(\ell_r) / \sum_q \exp(\ell_q)$ , and define the index policy as the probability of sampling  
 1202 the token that matches the  $i_l$ -th input value:

$$1203 \quad \pi_{\mathbf{W}}(i_l | \mathbf{x}, \hat{\mathbf{p}}^{z_{1:l}}) := p_{\text{vocab}}(u = \boldsymbol{\mu}_{x_{i_l}} | \hat{\mathbf{p}}_{l+1}^{\text{att}}).$$

1204 Then the score admits the positional outer-product decomposition

$$1205 \quad \nabla_{\mathbf{W}} \log \pi_{\mathbf{W}}(i_l | \cdot) = \sum_{k \in \mathcal{I}_l} \alpha_l(k) \eta_l(k) \mathbf{p}_k \mathbf{p}_{c_l}^\top, \quad \eta_l(k) := \frac{\langle \mathbf{p}_k - \hat{\mathbf{p}}_{l+1}^{\text{att}}, \mathbf{p}_{i_l} - \sum_r p_{\text{vocab}}(r) \mathbf{p}_r \rangle}{\beta}.$$

1206 Moreover, under Eq. (12) (positional orthogonality;  $\mathbf{W}$  acts only on positional blocks; the mask  
 1207  $m_l$  is  $\mathbf{W}$ -independent) and a finetuning regime that only reweights legal transitions within these  
 1208 blocks, the family  $\{\mathbf{p}_k \mathbf{p}_{c_l}^\top\}$  forms an orthogonal basis for the reachable score directions. Conse-  
 1209 quently,  $\nabla_{\mathbf{W}} \log \pi_{\mathbf{W}}$  decomposes uniquely onto orthogonal tensor blocks  $\text{span}\{\mathbf{p}_k\} \otimes \text{span}\{\mathbf{p}_{c_l}\}$   
 1210 (Orthogonal Block Isolation).  
 1211

1212 **Proof.** We expand every ingredient with step-by-step derivations.

1213 (1) Vocabulary softmax – definition and gradient. Fix  $l$ , and define the normalized vocabulary  
 1214 probabilities:

$$1215 \quad p_{\text{vocab}}(r) = \frac{\exp(\ell_r)}{\sum_q \exp(\ell_q)}, \quad \ell_r = \frac{\langle \hat{\mathbf{p}}_{l+1}^{\text{att}}, \mathbf{p}_r \rangle}{\beta},$$

1216 with classes  $r \in \{0, 1, \text{EOS}\}$ . The policy is

$$1217 \quad \pi_{\mathbf{W}}(i_l | \cdot) = p_{\text{vocab}}(x_{i_l}) := p_{\text{vocab}}(u = \boldsymbol{\mu}_{x_{i_l}} | \hat{\mathbf{p}}_{l+1}^{\text{att}}).$$

1218 We differentiate  $\log p_{\text{vocab}}(x_{i_l})$  w.r.t.  $\hat{\mathbf{p}}_{l+1}^{\text{att}}$ :

$$1219 \quad \nabla_{\hat{\mathbf{p}}_{l+1}^{\text{att}}} \log p_{\text{vocab}}(x_{i_l}) \stackrel{(1.1)}{=} \nabla_{\hat{\mathbf{p}}_{l+1}^{\text{att}}} \left( \ell_{x_{i_l}} - \log \sum_r e^{\ell_r} \right) \quad (13)$$

$$1220 \quad \stackrel{(1.2)}{=} \frac{\mathbf{p}_{i_l}}{\beta} - \sum_r \frac{e^{\ell_r}}{\sum_q e^{\ell_q}} \frac{\mathbf{p}_r}{\beta} \quad (14)$$

$$1221 \quad \stackrel{(1.3)}{=} \frac{\mathbf{p}_{i_l} - \sum_r p_{\text{vocab}}(r) \mathbf{p}_r}{\beta} :=: \mathbf{g}_l^{\text{vocab}}. \quad (15)$$

1222 Step (1.1) expands the log-softmax; (1.2) uses  $\partial \ell_r / \partial \hat{\mathbf{p}}_{l+1}^{\text{att}} = \mathbf{p}_r / \beta$ ; (1.3) recognizes  $p_{\text{vocab}}$ .

1223 (2) Attention softmax – Jacobian and gradient w.r.t.  $\mathbf{W}$ . The masked logits and weights are

$$1224 \quad s_l(k) = \mathbf{p}_k^\top \mathbf{W} \mathbf{p}_{c_l} + m_l(k), \quad \alpha_l(k) = \frac{e^{s_l(k)}}{\sum_{q \in \mathcal{I}_l} e^{s_l(q)}}.$$

The softmax Jacobian is  $\partial\alpha_l(j)/\partial s_l(k) = \alpha_l(j)(\delta_{jk} - \alpha_l(k))$ . Since  $m_l$  is  $\mathbf{W}$ -independent,

$$\nabla_{\mathbf{W}} s_l(k) = \nabla_{\mathbf{W}} (\mathbf{p}_k^\top \mathbf{W} \mathbf{p}_{c_l}) = \mathbf{p}_k \mathbf{p}_{c_l}^\top.$$

Thus the gradient of the attention-weighted token vector is:

$$\nabla_{\mathbf{W}} \hat{\mathbf{p}}_{l+1}^{\text{att}} = \nabla_{\mathbf{W}} \left( \sum_{j \in \mathcal{I}_l} \alpha_l(j) \mathbf{p}_j \right) \quad (16)$$

$$\stackrel{(2.1)}{=} \sum_j \mathbf{p}_j \sum_k \frac{\partial \alpha_l(j)}{\partial s_l(k)} \nabla_{\mathbf{W}} s_l(k) \quad (17)$$

$$\stackrel{(2.2)}{=} \sum_j \mathbf{p}_j \sum_k \alpha_l(j) (\delta_{jk} - \alpha_l(k)) \mathbf{p}_k \mathbf{p}_{c_l}^\top \quad (18)$$

$$\stackrel{(2.3)}{=} \sum_k \alpha_l(k) \left( \mathbf{p}_{x_k} - \sum_j \alpha_l(j) \mathbf{p}_j \right) \mathbf{p}_k \mathbf{p}_{c_l}^\top \quad (19)$$

$$\stackrel{(2.4)}{=} \sum_k \alpha_l(k) (\mathbf{p}_k - \hat{\mathbf{p}}_{l+1}^{\text{att}}) \mathbf{p}_k \mathbf{p}_{c_l}^\top. \quad (20)$$

Step (2.1) is the chain rule; (2.2) uses the softmax Jacobian and  $\nabla_{\mathbf{W}} s_l$ ; (2.3) collects terms; (2.4) recognizes  $\hat{\mathbf{p}}_{l+1}^{\text{att}}$ .

(3) Chain rule for the policy score. Combining (1)–(2),

$$\nabla_{\mathbf{W}} \log \pi_{\mathbf{W}}(i_l | \cdot) \stackrel{(3.1)}{=} \left( \nabla_{\hat{\mathbf{p}}_{l+1}^{\text{att}}} \log p_{\text{vocab}}(x_{i_l}) \right)^\top \cdot \nabla_{\mathbf{W}} \hat{\mathbf{p}}_{l+1}^{\text{att}} \quad (21)$$

$$\stackrel{(3.2)}{=} \sum_{k \in \mathcal{I}_l} \alpha_l(k) \underbrace{\left\langle \mathbf{p}_k - \hat{\mathbf{p}}_{l+1}^{\text{att}}, \mathbf{g}_l^{\text{vocab}} \right\rangle}_{:= \eta_l(k)} \mathbf{p}_k \mathbf{p}_{c_l}^\top. \quad (22)$$

This yields the claimed positional outer-product decomposition with weights  $\alpha_l(k)\eta_l(k)$ , where  $\mathbf{g}_l^{\text{vocab}} = (\mathbf{p}_{i_l} - \sum_r p_{\text{vocab}}(r) \mathbf{p}_r) / \beta$ .

(4) Interchange and smoothness conditions. The above steps use: (i)  $m_l$  is  $\mathbf{W}$ -independent, so  $\nabla_{\mathbf{W}} s_l$  exists and is continuous; (ii) softmax is  $C^\infty$ , thus  $\alpha_l$  and  $p_{\text{vocab}}$  are smooth in  $\mathbf{W}$ ; (iii) boundedness of token embeddings  $\mathbf{p}_\cdot$ , vocabulary  $\mathbf{U}$ , and temperature  $\beta > 0$  ensures an  $L^1$  dominator for Leibniz interchange when taking expectations over trajectories (used later in policy-gradient proofs).

(5) Orthogonal Block Isolation (OBI) and the role of Eq. (12). Under Eq. (12): positional embeddings  $\{\mathbf{p}_j\}$  are mutually orthogonal;  $\mathbf{K}^\top \mathbf{Q}$  (hence  $\mathbf{W}$ ) acts only on positional blocks;  $\mathbf{V}$  projects out positional-channel-orthogonal components when forming logits with  $\mathbf{U}$ . Finetuning reweights only the legal transitions while preserving the block structure (the mask  $m_l$  is fixed). Therefore, the reachable score directions lie in the span of  $\{\mathbf{p}_k \mathbf{p}_{c_l}^\top\}$ , and for any  $(k, l) \neq (k', l')$ ,

$$\langle \mathbf{p}_k \mathbf{p}_{c_l}^\top, \mathbf{p}_{k'} \mathbf{p}_{c_{l'}}^\top \rangle = (\mathbf{p}_k^\top \mathbf{p}_{k'}) (\mathbf{p}_{c_l}^\top \mathbf{p}_{c_{l'}}) = 0.$$

Hence cross-terms vanish and the decomposition in (3.2) is unique on the block-diagonal positional subspace  $\text{span}\{\mathbf{p}_k\} \otimes \text{span}\{\mathbf{p}_{c_l}\}$ . This establishes OBI with Eq. (12) as sufficient conditions. In practice, these are met at initialization (masking, orthogonality, block action), and finetuning that only reweights legal transitions preserves the required block-isolation.  $\square$

**Lemma 4** (From per-step logit margins to overall 0-1 loss). Fix a target subtask family  $\mathcal{F}_{S_\star} = \{f_{S_\star^1}, \dots, f_{S_\star^{k^*}}\}$  and consider decoding  $k^* + 1$  steps (the last for EOS). For each step  $l \in \{1, \dots, k^* + 1\}$ , let the legal index set be  $\mathcal{I}_l$  with cardinality  $d_l := |\mathcal{I}_l| \geq 1$ , and denote the correct index by  $i_l$  (with  $i_{k^*+1} := d+1$  for EOS). Define masked attention logits  $s_l(j) := \mathbf{p}_j^\top \mathbf{W} \mathbf{p}_{c_l} + m_l(j)$  and attention weights  $\alpha_l(j) = \text{softmax}_j(s_l(j))$ . Let

$$\Delta_l := s_l(i_l) - \max_{j \in \mathcal{I}_l \setminus \{i_l\}} s_l(j) \geq 0$$

be the per-step attention logit gap on the correct child. Assume pairwise-orthogonal token embeddings, bounded norms, and temperature  $\beta > 0$  so that

$$|\eta_l(k)| \leq \frac{4}{\beta}, \quad \text{and} \quad p_{\text{vocab}}(r) = \text{softmax}_r(\langle \hat{\mathbf{p}}_{l+1}^{\text{att}}, \mathbf{p}_r \rangle / \beta)$$

as in Lemma 3. Let  $M := K + 1$  denote the vocabulary size. For any error budget  $(\varepsilon_1, \dots, \varepsilon_{k^*+1}) \in (0, 1)^{k^*+1}$ , define the per-step required attention probability and corresponding margin threshold by

$$\alpha_l^{\text{req}} := \frac{1}{2} + \frac{\beta}{2} \log\left(\frac{M(1-\varepsilon_l)}{\varepsilon_l}\right), \quad \Gamma_l := \log\left(\frac{d_l-1}{(\alpha_l^{\text{req}})^{-1}-1}\right),$$

whenever  $\alpha_l^{\text{req}} < 1$  (otherwise the requirement is infeasible).

If the per-step logit gaps satisfy  $\Delta_l \geq \Gamma_l$  for all  $l \in [k^{\text{star}} + 1]$ , then the per-step selection probability of the correct token obeys

$$\pi_l := p_{\text{vocab}}(u = \mathbf{p}_{i_l} \mid \hat{\mathbf{p}}_{l+1}^{\text{att}}) \geq 1 - \varepsilon_l, \quad l = 1, \dots, k^* + 1.$$

Consequently, letting the 0-1 loss be  $L := 1 - R^{k^*}(\text{TF}(\cdot; \mathbf{W}))$  (so that  $\mathbb{E}[L] = 1 - \mathbb{E}[R^{k^*}]$ ), we have

$$\mathbb{E}_{\mathbf{x} \sim \text{Unif}([K]^d)}[L] \leq \sum_{l=1}^{k^*+1} \varepsilon_l.$$

In particular, choosing any  $(\varepsilon_l)$  with  $\sum_l \varepsilon_l \leq \varepsilon$  implies  $\mathbb{E}[L] \leq \varepsilon$ .

**Proof.** Step 1 (attention probability from logit gap). For  $\alpha_l(\cdot) = \text{softmax}(s_l(\cdot))$  and gap  $\Delta_l$ , the correct-child attention weight obeys the standard softmax bound

$$\alpha_l(i_l) = \frac{1}{1 + \sum_{j \neq i_l} \exp(- (s_l(i_l) - s_l(j)))} \geq \frac{1}{1 + (d_l - 1) e^{-\Delta_l}}.$$

Thus if  $\Delta_l \geq \Gamma_l$  with  $\Gamma_l$  defined in the statement, then

$$\alpha_l(i_l) \geq \alpha_l^{\text{req}}.$$

Step 2 (vocabulary probability from attention concentration). With orthogonal embeddings,  $\langle \hat{\mathbf{p}}_{l+1}^{\text{att}}, \mathbf{p}_{i_l} \rangle = \alpha_l(i_l)$  and for any competitor token  $r \neq x_{i_l}$ ,  $\langle \hat{\mathbf{p}}_{l+1}^{\text{att}}, \mathbf{p}_r \rangle \leq 1 - \alpha_l(i_l)$  (competitors include tokens not present in the mixture, whose inner products are  $0 \leq 1 - \alpha_l$ ). Hence the vocabulary-logit gap between the correct token and any competitor is at least

$$\gamma_l^{\text{vocab}} \geq \alpha_l(i_l) - (1 - \alpha_l(i_l)) = 2\alpha_l(i_l) - 1.$$

The softmax lower bound then gives

$$\pi_l = \frac{e^{\alpha_l(i_l)/\beta}}{\sum_r e^{\langle \hat{\mathbf{p}}_{l+1}^{\text{att}}, \mathbf{p}_r \rangle / \beta}} \geq \frac{1}{1 + \sum_{r \neq x_{i_l}} e^{-\gamma_l^{\text{vocab}}/\beta}} \geq \frac{1}{1 + M e^{-(2\alpha_l(i_l)-1)/\beta}}.$$

Therefore  $\alpha_l(i_l) \geq \alpha_l^{\text{req}}$  implies

$$\pi_l \geq \frac{1}{1 + M e^{-(2\alpha_l^{\text{req}}-1)/\beta}} = 1 - \varepsilon_l.$$

Step 3 (from per-step success to sequence success). The sequence fails only if at least one step fails, hence by the union bound

$$\mathbb{P}(\text{sequence fails}) \leq \sum_{l=1}^{k^*+1} (1 - \pi_l) \leq \sum_{l=1}^{k^*+1} \varepsilon_l.$$

Taking expectation over  $\mathbf{x} \sim \text{Unif}([K]^d)$  does not increase the bound, yielding the displayed inequality for the 0-1 loss  $L = 1 - R^{k^*}$ .  $\square$

**Lemma 5** (Policy Gradients for REINFORCE). Let the input  $\mathbf{x}$  be fixed and the decoding length be  $k^* \geq 2$ . Denote the generated token sequence by  $\hat{\mathbf{p}}^{z_1:k^*} = (\hat{\mathbf{p}}^{z_1}, \dots, \hat{\mathbf{p}}^{z_{k^*}})$ . At step  $l$  ( $1 \leq l \leq k^* + 1$ ), an attention policy first samples a secret index's token  $i_l \sim \pi_{\mathbf{W}^{k^*}}(\cdot \mid \mathbf{x}, \hat{\mathbf{p}}^{z_1:l}) := \hat{p}_{\mathbf{W}}(u = \mu_{x_{i_l}} \mid \mathbf{x}, \hat{\mathbf{p}}^{z_1:l})$ . The next token is then deterministically produced by the deterministic Feedforward  $\text{FFN}_m(\cdot)$  at  $m$ -th reasoning procedure, so that the only source of randomness is the

sampling of tokens  $\mu_{x_{i_l}}$  corresponding to secret index sequences  $i_{1:k^*+1}$ . We therefore define the (random) trajectory as  $\tau := (i_{1:k^*+1})$  and the induced measure

$$p_{\mathbf{W}^{k^*}}(\tau | \mathbf{x}) = \prod_{l=1}^{k^*+1} \pi_{\mathbf{W}^{k^*}}(i_l | \mathbf{x}, \hat{\mathbf{p}}^{z_{1:l}}),$$

where the tokens  $\hat{\mathbf{p}}^{z_{1:k^*}}$  are deterministic functions of  $\tau$  via  $g_{\mathbf{W}^{k^*}}$ . Let the terminal reward be  $R^{k^*}(\text{TF}(\cdot; \mathbf{W}))$ , which does not explicitly depend on  $\mathbf{W}^{k^*}$ . Define the population loss

$$\mathcal{J}_{\text{REINFORCE}}^{k^*}(\mathbf{W}^{k^*}) := \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_x, \tau \sim p_{\mathbf{W}^{k^*}}(\cdot | \mathbf{x})} \left[ R^{k^*}(\mu^{z_{k^*}}(\tau_{-2})) \right]. \quad (23)$$

Assume the following regularity conditions hold: (i) for all  $l$ ,  $\pi_{\mathbf{W}^{k^*}}(i_l | \cdot) > 0$  is Fréchet differentiable in  $\mathbf{W}^{k^*}$  with log-smoothness; (ii)  $R^{k^*}$  is bounded and integrable; and (iii) differentiation can be interchanged with integration (or expectation) under dominated convergence / parameterized measure continuity. Denote  $\mathbb{E}_\tau[\cdot] = \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_x, \tau \sim p_{\mathbf{W}^{k^*}}(\cdot | \mathbf{x})}[\cdot]$ , the policy gradients are:

$$\nabla_{\mathbf{W}^{k^*}} \mathcal{J}_{\text{REINFORCE}}^{k^*}(\mathbf{W}^{k^*}) = \mathbb{E}_\tau \left[ R^{k^*}(\mu^{z_{k^*}}(\tau_{-2})) \cdot \sum_{l=1}^{k^*+1} \nabla_{\mathbf{W}^{k^*}} \log \pi_{\mathbf{W}^{k^*}}(i_l | \mathbf{x}, \hat{\mathbf{p}}^{z_{1:l}}) \right]. \quad (24)$$

The formula reflects that token generation is deterministic via  $\phi$  while the non-Markovian dependency arises from the index policy  $\pi$  depending on the full history. Therefore, by Lemma 3, we have

$$\nabla_{\mathbf{W}^{k^*}} \mathcal{J}_{\text{REINFORCE}}^{k^*}(\mathbf{W}^{k^*}) = \mathbb{E}_\tau \left[ R^{k^*}(\mu^{z_{k^*}}(\tau_{-2})) \sum_{l=1}^{k^*+1} \sum_{k \in \mathcal{I}_l} \alpha_l^{k^*}(k) \eta_l^{k^*}(k) \mathbf{p}_k \mathbf{p}_{c_l}^\top \right], \quad (25)$$

where

- $\alpha_l^{k^*}(k) := \text{softmax}_k(\mathbf{p}_k^\top \mathbf{W}^{k^*} \mathbf{p}_{c_l} + m_l(k))$ ,
- $\hat{\mathbf{p}}_{l+1}^{\text{att}, k^*} := \sum_{j \in \mathcal{I}_l} \alpha_l^{k^*}(j) \mathbf{p}_j$ ,
- $\eta_l^{k^*}(k) := \langle \mathbf{p}_k - \hat{\mathbf{p}}_{l+1}^{\text{att}, k^*}, \mathbf{p}_{i_l} - \sum_r p_{\text{vocab}}^{k^*}(r) \mathbf{p}_r \rangle / \beta$ ,
- $p_{\text{vocab}}^{k^*}(\cdot) := \frac{\exp(\frac{\langle \hat{\mathbf{p}}_{l+1}^{\text{att}}, \mathbf{p}_r \rangle}{\beta})}{\sum_q \exp(\frac{\langle \hat{\mathbf{p}}_{l+1}^{\text{att}}, \mathbf{p}_q \rangle}{\beta})}$ .

**Proof.** Write  $\tau = (i_{1:k^*+1})$  and abbreviate  $\mathbb{E}_\tau$  for  $\mathbb{E}_{\tau \sim p_{\mathbf{W}^{k^*}}(\cdot | \mathbf{x})}$ . The tokens  $\hat{\mathbf{p}}^{z_{1:k^*}}(\tau)$  are deterministic given  $\tau$  via  $g_{\mathbf{W}^{k^*}}$ .

We derive the gradient in a step-numbered manner:

$$\nabla_{\mathbf{W}^{k^*}} \mathcal{J}_{\text{REINFORCE}}^{k^*} \stackrel{(1)}{=} \nabla_{\mathbf{W}^{k^*}} \int R^{k^*}(\mu^{z_{k^*}}(\tau_{-2})) p_{\mathbf{W}^{k^*}}(\tau | \mathbf{x}) d\tau \quad (26)$$

$$\stackrel{(2)}{=} \int R^{k^*}(\mu^{z_{k^*}}(\tau_{-2})) \nabla_{\mathbf{W}^{k^*}} p_{\mathbf{W}^{k^*}}(\tau | \mathbf{x}) d\tau \quad (27)$$

$$\stackrel{(3)}{=} \int R^{k^*}(\mu^{z_{k^*}}(\tau_{-2})) p_{\mathbf{W}^{k^*}}(\tau | \mathbf{x}) \nabla_{\mathbf{W}^{k^*}} \log p_{\mathbf{W}^{k^*}}(\tau | \mathbf{x}) d\tau \quad (28)$$

$$\stackrel{(4)}{=} \mathbb{E}_\tau \left[ R^{k^*}(\mu^{z_{k^*}}(\tau_{-2})) \sum_{l=1}^{k^*+1} \nabla_{\mathbf{W}^{k^*}} \log \pi_{\mathbf{W}^{k^*}}(i_l | \mathbf{x}, \hat{\mathbf{p}}^{z_{1:l}}) \right], \quad (29)$$

which equals equation 24.

Step (1) is the definition of  $\mathcal{J}_{\text{REINFORCE}}^{k^*}$ ; Step (2) uses the Leibniz interchange  $\nabla_\theta \int f(\tau, \theta) d\tau = \int \nabla_\theta f(\tau, \theta) d\tau$  under the following conditions: (i)  $R^{k^*} \circ g_{\mathbf{W}^{k^*}}$  is bounded; (ii)  $p_{\mathbf{W}^{k^*}}(\tau | \mathbf{x})$  is

1404 differentiable in  $\mathbf{W}^{k^*}$ ; (iii) there exists an integrable dominator  $h(\tau)$  with  $\|R^{k^*} \nabla p\| \leq h$  (domi-  
 1405 nated convergence / parameterized measure continuity). These hold because  $R^{k^*} \in \{0, 1\}$  and  $\pi$   
 1406 is softmax-based with log-smoothness. Step (3) is the score-function identity  $\nabla p = p \nabla \log p$ . Step  
 1407 (4) applies the chain rule to  $\log p_{\mathbf{W}^{k^*}}(\tau | \mathbf{x}) = \sum_l \log \pi_{\mathbf{W}^{k^*}}(i_l | \cdot)$ , which holds regardless of  
 1408 parameter sharing across time. The resulting decomposition does not require Markovity of the state,  
 1409 since the history enters through the conditioning of  $\pi$ .

1410 Finally, we note that deterministic  $\phi$  does not affect the score terms and only enters through  
 1411  $R^{k^*}(\boldsymbol{\mu}^{z_{k^*}}(\tau_{-2}))$ , preserving all interchanges above. Fubini's theorem applies whenever  
 1412  $\mathbb{E}_\tau[|R^{k^*} \sum_l \nabla \log \pi_l|] < \infty$ , which holds by boundedness of  $R^{k^*}$  and square-integrability (log-  
 1413 smoothness) of the score.  
 1414

1415 By Lemma 3, for each  $l$  we have the score decomposition

$$1416 \nabla_{\mathbf{W}^{k^*}} \log \pi_{\mathbf{W}^{k^*}}(i_l | \cdot) = \sum_{k \in \mathcal{I}_l} \alpha_l^{k^*}(k) \eta_l^{k^*}(k) \mathbf{p}_k \mathbf{p}_{c_l}^\top.$$

1417 Substituting this into the REINFORCE identity

$$1418 \nabla \mathcal{J}_{\text{REINFORCE}} = \mathbb{E}[R \sum_l \nabla \log \pi]$$

1419 and using linearity of expectation, yields the final expanded formula. Orthogonal Block Isolation  
 1420 guarantees each term lies in a unique positional tensor block  $p_k p_{c_l}^\top$ , ensuring no cross-block inter-  
 1421 ference in these expansions.  $\square$   
 1422

1423 **Lemma 6** (One-step REINFORCE update of per-step attention logit gap). Fix a step  $l \in$   
 1424  $\{1, \dots, k^* + 1\}$ . Define

$$1425 s_l(j; \mathbf{W}) := \mathbf{p}_j^\top \mathbf{W} \mathbf{p}_{c_l} + m_l(j), \quad \Delta_l(\mathbf{W}) := s_l(i_l; \mathbf{W}) - \max_{j \in \mathcal{I}_l \setminus \{i_l\}} s_l(j; \mathbf{W}).$$

1426 Consider the one-step REINFORCE update that maximizes the reward (step size  $\eta > 0$ )

$$1427 \mathbf{W}^{(t+1)} = \mathbf{W}^{(t)} + \eta \nabla_{\mathbf{W}} \mathcal{J}_{\text{REINFORCE}}^{k^*}(\mathbf{W}^{(t)}),$$

1428 with all gradients evaluated at  $\mathbf{W}^{(t)}$ . Let  $\mathbb{E}_\tau^{(t)}[\cdot]$  denote the trajectory expectation under the policy  
 1429 at  $\mathbf{W}^{(t)}$  in Lemma 5, and reuse the notation of Lemma 3

$$1430 \alpha_l^{(t)}(j) = \text{softmax}_j(\mathbf{p}_j^\top \mathbf{W}^{(t)} \mathbf{p}_{c_l} + m_l(j)), \quad \eta_l^{(t)}(j) = \frac{\langle \mathbf{p}_j - \hat{\mathbf{p}}_{l+1}^{\text{att},(t)}, \mathbf{p}_{i_l} - \sum_r p_{\text{vocab}}^{(t)}(r) \mathbf{p}_r \rangle}{\beta}.$$

1431 Introduce the path-success and spurious-success decomposition. Let  $\mathcal{T}_{\text{true}}$  be the set of trajectories  
 1432 that select the correct indices  $i_{1:k^*}$  (and then EOS); define the true-path success probability

$$1433 p_{\text{path}} := \mathbb{E}_{\mathbf{x}} \mathbb{P}_{\tau|\mathbf{x}}\{\tau \in \mathcal{T}_{\text{true}}\} = \mathbb{E}_{\mathbf{x}} \left[ \prod_{l=1}^{k^*+1} \pi_{\mathbf{W}^{k^*}}(i_l | \mathbf{x}, \hat{\mathbf{p}}^{z_{1:l}}) \right].$$

1434 Let the spurious (reward-hacking) success rate be

$$1435 \rho_{\text{spur}} := \mathbb{E}_{\mathbf{x}} \left[ \mathbb{P}_{\tau|\mathbf{x}}\{R^{k^*}(\hat{\mathbf{p}}^{z_{k^*}}(\tau)) = 1 \mid \tau \notin \mathcal{T}_{\text{true}}\} \right],$$

1436 which depends on the task structure encoded by  $\{\Phi_l\}, \{\mathcal{I}_l\}$ . Then, for any policy,

$$1437 p_{\text{succ}} = p_{\text{path}} + (1 - p_{\text{path}}) \rho_{\text{spur}}. \quad (30)$$

1438 In particular, at near-uniform initialization with  $d_l := |\mathcal{I}_l|$  and  $\pi(i_l | \cdot) \approx 1/d_l$  on legal children,

$$1439 p_{\text{path}} \leq \prod_{l=1}^{k^*+1} \frac{1}{d_l}, \quad p_{\text{succ}} \leq \rho_{\text{spur}} + (1 - \rho_{\text{spur}}) \prod_{l=1}^{k^*+1} \frac{1}{d_l}. \quad (31)$$

1440 For the parity case study with  $\mathbf{x} \sim \text{Unif}\{0, 1\}^d$  and XOR kernels  $\{\Phi_l\}$ , any wrong index set yields  
 1441 the correct parity with probability 1/2; hence

$$1442 \rho_{\text{spur}}^{\text{parity}} = \frac{1}{2}, \quad p_{\text{succ}} \leq \frac{1}{2} + \frac{1}{2} \prod_{l=1}^{k^*+1} \frac{1}{d_l}. \quad (32)$$

1458  
1459  
1460  
1461  
1462  
1463  
1464  
1465  
1466  
1467  
1468  
1469  
1470  
1471  
1472  
1473  
1474  
1475  
1476  
1477  
1478  
1479  
1480  
1481  
1482  
1483  
1484  
1485  
1486  
1487  
1488  
1489  
1490  
1491  
1492  
1493  
1494  
1495  
1496  
1497  
1498  
1499  
1500  
1501  
1502  
1503  
1504  
1505  
1506  
1507  
1508  
1509  
1510  
1511

Then:

1) (exact logit update for affine logits) For any  $j \in \mathcal{I}_l$ , we have

$$s_l(j; \mathbf{W}^{(t+1)}) = s_l(j; \mathbf{W}^{(t)}) + \eta \mathbb{E}_\tau^{(t)} [R^{k^*} \alpha_l^{(t)}(j) \eta_l^{(t)}(j)]. \quad (33)$$

2) (subgradient inequality for the max) Let the active competitor set be  $\mathcal{A}_l^{\max} := \arg \max_{j \neq i_l} s_l(j; \mathbf{W}^{(t-1)})$ . Then

$$\Delta_l(\mathbf{W}^{(t+1)}) \geq \Delta_l(\mathbf{W}^{(t)}) + \eta \left( \mathbb{E}_\tau^{(t)} [R^{k^*} \alpha_l^{(t)}(i_l) \eta_l^{(t)}(i_l)] - \sup_{j \neq i_l} \mathbb{E}_\tau^{(t)} [R^{k^*} \alpha_l^{(t)}(j) \eta_l^{(t)}(j)] \right) + \mathcal{O}(\eta^2). \quad (34)$$

If, in a neighborhood of  $\mathbf{W}^{(t-1)}$ , the active maximizer is unique and remains unchanged (there exists  $j_l^{\max} \in \mathcal{A}_l^{\max}$  that stays the maximizer in that neighborhood), then the first-order exact form is

$$\Delta_l(\mathbf{W}^{(t+1)}) = \Delta_l(\mathbf{W}^{(t)}) + \eta \left( \mathbb{E}_\tau^{(t)} [R^{k^*} \alpha_l^{(t)}(i_l) \eta_l^{(t)}(i_l)] - \mathbb{E}_\tau^{(t)} [R^{k^*} \alpha_l^{(t)}(j_l^{\max}) \eta_l^{(t)}(j_l^{\max})] \right) + \mathcal{O}(\eta^2). \quad (35)$$

3) (update of a smooth lower bound) Define the competitors' log-sum-exp lower bound

$$\phi_l(\mathbf{W}) := s_l(i_l; \mathbf{W}) - \log \sum_{j \neq i_l} \exp(s_l(j; \mathbf{W})) \leq \Delta_l(\mathbf{W}).$$

Let  $\tilde{\alpha}_l(j) := \frac{\exp(s_l(j))}{\sum_{q \neq i_l} \exp(s_l(q))}$  be the softmax normalized only over competitors, evaluated at  $\mathbf{W}^{(t-1)}$ .

Then

$$\phi_l(\mathbf{W}^{(t+1)}) = \phi_l(\mathbf{W}^{(t)}) + \eta \left( \mathbb{E}_\tau^{(t)} [R^{k^*} \alpha_l^{(t)}(i_l) \eta_l^{(t)}(i_l)] - \sum_{j \neq i_l} \tilde{\alpha}_l^{(t)}(j) \mathbb{E}_\tau^{(t)} [R^{k^*} \alpha_l^{(t)}(j) \eta_l^{(t)}(j)] \right) + \mathcal{O}(\eta^2). \quad (36)$$

Moreover, to recompute intermediate quantities after one update, the following first-order relations (Fréchet differentials at  $\mathbf{W}^{(t-1)}$ ) hold:

$$\delta s_l^{(t)}(k) = s_l(k; \mathbf{W}^{(t+1)}) - s_l(k; \mathbf{W}^{(t)}) = \eta \mathbb{E}_\tau^{(t)} [R^{k^*} \alpha_l^{(t)}(k) \eta_l^{(t)}(k)], \quad (37)$$

$$\delta \alpha_l^{(t)}(k) = \alpha_l^{(t)}(k) \left( \delta s_l(k) - \sum_{q \in \mathcal{I}_l} \alpha_l^{(t)}(q) \delta s_l(q) \right) + \mathcal{O}(\eta^2), \quad (38)$$

$$\delta \hat{\mathbf{p}}_{l+1}^{\text{att},(t)} = \sum_{k \in \mathcal{I}_l} \delta \alpha_l^{(t)}(k) \mathbf{p}_k + \mathcal{O}(\eta^2), \quad (39)$$

$$\delta p_{\text{vocab}}^{(t)}(r) = p_{\text{vocab}}^{(t)}(r) \frac{\langle \delta \hat{\mathbf{p}}_{l+1}^{\text{att},(t)}, \mathbf{p}_r - \sum_q p_{\text{vocab}}^{(t)}(q) \mathbf{p}_q \rangle}{\beta} + \mathcal{O}(\eta^2), \quad (40)$$

$$\delta \eta_l^{(t)}(k) = \frac{1}{\beta} \left( -\langle \delta \hat{\mathbf{p}}_{l+1}^{\text{att},(t)}, \mathbf{p}_{i_l} - \sum_r p_{\text{vocab}}^{(t)}(r) \mathbf{p}_r \rangle - \langle \mathbf{p}_k - \hat{\mathbf{p}}_{l+1}^{\text{att},(t)}, \sum_r \delta p_{\text{vocab}}^{(t)}(r) \mathbf{p}_r \rangle \right) + \mathcal{O}(\eta^2). \quad (41)$$

In practice, with  $n$  trajectories prompted by  $\mathbf{x}^s \sim \text{Unif}([K]^d), \forall s \in [n]$ , Monte Carlo estimators can replace the expectations in equation 33–equation 36, e.g.,

$$\hat{\delta s}_l^{(t)}(k) = \eta \frac{1}{n} \sum_{s=1}^n R^{k^*} (\text{TF}(\mathbf{x}^s; \mathbf{W}^{(t)})) \alpha_l^{s,(t)}(k) \eta_l^{s,(t)}(k).$$

**Proof.**(1) By Lemma 5 and Lemma 3, at  $\mathbf{W}^{(t-1)}$ ,

$$\nabla_{\mathbf{W}} \mathcal{J}_{\text{REINFORCE}}^{k^*} = \mathbb{E}_\tau \left[ R^{k^*} \sum_{u=1}^{k^*+1} \sum_{k \in \mathcal{I}_u} \alpha_u(k) \eta_u(k) \mathbf{p}_k \mathbf{p}_{c_u}^\top \right].$$

Using  $\nabla_{\mathbf{W}} s_l(j) = \mathbf{p}_j \mathbf{p}_{c_l}^\top$  and the Frobenius inner product, a Taylor expansion up to second order along the explicit-Euler path  $\mathbf{W}^{(t+1)} = \mathbf{W}^{(t)} + \eta \mathbf{G}_t$  (with fixed  $\mathbf{G}_t := \nabla_{\mathbf{W}} \mathcal{J}_{\text{REINFORCE}}^{k^*}(\mathbf{W}^{(t)})$ ) gives

$$s_l(j; \mathbf{W}^{(t+1)}) = s_l(j; \mathbf{W}^{(t)}) + \eta \langle \mathbf{p}_j \mathbf{p}_{c_l}^\top, \nabla_{\mathbf{W}} \mathcal{J}_{\text{REINFORCE}}^{k^*} \rangle + \mathcal{O}(\eta^2),$$

and, more explicitly,

$$s_l(j; \mathbf{W}^{(t)} + \eta \mathbf{G}_t) = s_l(j; \mathbf{W}^{(t)}) + \eta \langle \nabla_{\mathbf{W}} s_l(j), \mathbf{G}_t \rangle + \frac{\eta^2}{2} \langle \mathbf{G}_t, \nabla_{\mathbf{W}}^2 s_l(j)[\mathbf{G}_t] \rangle.$$

Since  $s_l(j; \mathbf{W}) = \langle \mathbf{p}_j \mathbf{p}_{c_l}^\top, \mathbf{W} \rangle + m_l(j)$  is affine in  $\mathbf{W}$ , its Hessian vanishes:  $\nabla_{\mathbf{W}}^2 s_l(j) \equiv 0$ . Therefore the second-order term above is identically zero under the explicit-Euler step (no dependence of the direction on  $\eta$ ), and only the  $\mathcal{O}(\eta^2)$  bookkeeping remains for uniformity with subsequent non-linear propagations. and OBI guarantees  $\langle \mathbf{p}_j \mathbf{p}_{c_l}^\top, \mathbf{p}_k \mathbf{p}_{c_u}^\top \rangle = 0$  unless  $(k, u) = (j, l)$ , which yields equation 33.

(2) Let  $g_j := \eta \mathbb{E}_\tau [R^{k^*} \alpha_l(j) \eta_l(j)]$  denote the first-order increment of  $s_l(j)$ . By Danskin's theorem / the subgradient inequality,

$$\max_{j \neq i_l} (s_l(j) + g_j) \leq \max_{j \neq i_l} s_l(j) + \max_{j \neq i_l} g_j,$$

and substituting the definitions gives the lower bound equation 34. If the active maximizer is unique and remains unchanged in a neighborhood, then the directional derivative of  $\max_{j \neq i_l}$  is given by that  $j_l^{\max}$ , which yields equation 35.

(3) Let  $\phi_l = s_l(i_l) - \log \sum_{j \neq i_l} e^{s_l(j)}$ , so  $\Delta_l \geq \phi_l$ . Its directional derivative is

$$\delta \phi_l = \delta s_l(i_l) - \sum_{j \neq i_l} \tilde{\alpha}_l(j) \delta s_l(j).$$

Substituting equation 33 (with  $t \rightarrow t + 1$ ) and collecting  $\mathcal{O}(\eta^2)$  terms gives equation 36.

Finally, we detail the second-order sources for each intermediate quantity and then summarize them as  $\mathcal{O}(\eta^2)$  terms:

(i) Softmax attention  $\alpha_l$  to second order. Let  $s_l \in \mathbb{R}^{d_l}$  collect  $s_l(\cdot)$  and write  $\alpha_l = \text{softmax}(s_l)$ . For a perturbation  $\delta s_l = \mathcal{O}(\eta)$ , the second-order expansion of component  $j$  is

$$\delta \alpha_l(j) = \sum_k J_{jk}^\alpha \delta s_l(k) + \frac{1}{2} \sum_{k,m} H_{j,km}^\alpha \delta s_l(k) \delta s_l(m) + \mathcal{O}(\eta^3),$$

with Jacobian  $J_{jk}^\alpha = \alpha_l(j)(\delta_{jk} - \alpha_l(k))$  and Hessian

$$H_{j,km}^\alpha = \frac{\partial^2 \alpha_l(j)}{\partial s_l(k) \partial s_l(m)} = \alpha_l(j) \left( (\delta_{jm} - \alpha_l(m))(\delta_{jk} - \alpha_l(k)) - \alpha_l(k)(\delta_{km} - \alpha_l(m)) \right).$$

Because  $\delta s_l = \mathcal{O}(\eta)$  from equation 33, the quadratic term contributes  $\mathcal{O}(\eta^2)$ .

(ii) Attention-weighted token vector  $\hat{\mathbf{p}}_{l+1}^{\text{att}}$ . Since  $\hat{\mathbf{p}}_{l+1}^{\text{att}} = \sum_j \alpha_l(j) \mathbf{p}_j$  is linear in  $\alpha_l$ ,

$$\delta \hat{\mathbf{p}}_{l+1}^{\text{att}} = \sum_j \delta \alpha_l(j) \mathbf{p}_j = \sum_{j,k} J_{jk}^\alpha \delta s_l(k) \mathbf{p}_j + \frac{1}{2} \sum_{j,k,m} H_{j,km}^\alpha \delta s_l(k) \delta s_l(m) \mathbf{p}_j + \mathcal{O}(\eta^3),$$

so  $\delta \hat{\mathbf{p}}_{l+1}^{\text{att}} = \mathcal{O}(\eta) + \mathcal{O}(\eta^2)$ .

(iii) Vocabulary softmax  $p_{\text{vocab}}$ . Let  $\ell_r = \langle \hat{\mathbf{p}}_{l+1}^{\text{att}}, \mathbf{p}_r \rangle / \beta$  and  $\mathbf{p} = \text{softmax}(\ell)$ . Then

$$\delta \ell_r = \frac{\langle \delta \hat{\mathbf{p}}_{l+1}^{\text{att}}, \mathbf{p}_r \rangle}{\beta}, \quad \delta p(r) = \sum_m J_{rm}^p \delta \ell_m + \frac{1}{2} \sum_{m,n} H_{r,mn}^p \delta \ell_m \delta \ell_n + \mathcal{O}(\eta^3),$$

where  $J_{rm}^p = p(r)(\delta_{rm} - p(m))$  and

$$H_{r,mn}^p = \frac{\partial^2 p(r)}{\partial \ell_m \partial \ell_n} = p(r) \left( (\delta_{rn} - p(n))(\delta_{rm} - p(m)) - p(m)(\delta_{mn} - p(n)) \right).$$

Since  $\delta \hat{\mathbf{p}}_{l+1}^{\text{att}} = \mathcal{O}(\eta) + \mathcal{O}(\eta^2)$ , we have  $\delta \ell = \mathcal{O}(\eta) + \mathcal{O}(\eta^2)$  and thus  $\delta p = \mathcal{O}(\eta) + \mathcal{O}(\eta^2)$ .

(iv) The scalar  $\eta_l(k)$ . Write

$$\eta_l(k) = \frac{1}{\beta} \left\langle \underbrace{\mathbf{p}_k - \hat{\mathbf{p}}_{l+1}^{\text{att}}}_{:= \mathbf{u}}, \underbrace{\mathbf{p}_{i_l} - \sum_r p(r) \mathbf{p}_r}_{:= \mathbf{v}} \right\rangle.$$

1566 *Perturbing*  $(\mathbf{u}, \mathbf{v}) \mapsto (\mathbf{u} - \delta \hat{\mathbf{p}}_{l+1}^{\text{att}}, \mathbf{v} - \sum_r \delta p(r) \mathbf{p}_r)$  and expanding to second order yields

$$1567 \delta \eta_l(k) = \frac{1}{\beta} \left( -\langle \delta \hat{\mathbf{p}}_{l+1}^{\text{att}}, \mathbf{v} \rangle - \langle \mathbf{u}, \sum_r \delta p(r) \mathbf{p}_r \rangle \right) + \frac{1}{\beta} \left( -\frac{1}{2} \langle \delta^2 \hat{\mathbf{p}}_{l+1}^{\text{att}}, \mathbf{v} \rangle - \langle \delta \hat{\mathbf{p}}_{l+1}^{\text{att}}, \sum_r \delta p(r) \mathbf{p}_r \rangle - \frac{1}{2} \langle \mathbf{u}, \sum_r \delta^2 p(r) \mathbf{p}_r \rangle \right) + \mathcal{O}(\eta^3),$$

1569 where  $\delta^2 \hat{\mathbf{p}}_{l+1}^{\text{att}}$  and  $\delta^2 p(r)$  collect the quadratic terms shown in (ii) and (iii). Since  $\delta \hat{\mathbf{p}}_{l+1}^{\text{att}} = \mathcal{O}(\eta)$  and  $\delta p = \mathcal{O}(\eta)$  at leading order, all bracketed second-line contributions are  $\mathcal{O}(\eta^2)$ .

1572 Collecting these, we obtain the one-step relations stated in equation 37–equation 41, with every omitted higher-order contribution explicitly accounted for by the quadratic (Hessian) terms above and summarized as  $\mathcal{O}(\eta^2)$  due to the small step size.  $\square$

1573 **Lemma 7** (Stepwise expected gradients at iteration  $t$  with spurious success). Fix iteration  $t$  and a step  $l \in \{1, \dots, k^* + 1\}$ . For any index block  $j \in \mathcal{I}_l$ , write the block-projected expected REINFORCE gradient at  $\mathbf{W}^{(t)}$  as

$$1574 G_{l,j}^{(t)} := \left\langle \nabla_{\mathbf{W}} \mathcal{J}_{\text{REINFORCE}}^{k^*}(\mathbf{W}^{(t)}), \mathbf{p}_j \mathbf{p}_{c_l}^\top \right\rangle = \mathbb{E}^{(t)} [R^{k^*} \alpha_l^{(t)}(j) \eta_l^{(t)}(j)],$$

1575 where  $\mathbb{E}^{(t)}[\cdot]$  denotes the expectation over  $\mathbf{x} \sim \mathcal{P}_x$  and  $\tau \sim p_{\mathbf{W}^{(t)}}(\cdot | \mathbf{x})$ , and  $\alpha_l^{(t)}(\cdot), \eta_l^{(t)}(\cdot)$  are computed at  $\mathbf{W}^{(t)}$  (Lemma 3).

1583 Let  $C_{<l}$  be the event that the index prefix  $(i_1, \dots, i_{l-1})$  matches the target  $(i_1^*, \dots, i_{l-1}^*)$ , and denote

$$1584 \begin{aligned} 1585 p_{<l}^{(t)} &:= \mathbb{P}^{(t)}(C_{<l} = 1), \\ 1586 S_{l,j,\text{corr}}^{(t)} &:= \mathbb{E}^{(t)} [\alpha_l^{(t)}(j) \eta_l^{(t)}(j) \mathbf{1}\{i_l = i_l^*\} | C_{<l} = 1], \\ 1587 S_{l,j,\text{wrong}}^{(t)} &:= \mathbb{E}^{(t)} [\alpha_l^{(t)}(j) \eta_l^{(t)}(j) \mathbf{1}\{i_l \neq i_l^*\} | C_{<l} = 1], \\ 1588 \bar{S}_{l,j}^{(t)} &:= \mathbb{E}^{(t)} [\alpha_l^{(t)}(j) \eta_l^{(t)}(j) | C_{<l} = 0]. \end{aligned} \quad (42)$$

1591 Then

$$1592 G_{l,j}^{(t)} = p_{<l}^{(t)} \left( q_{l|\text{corr}}^{(t)} S_{l,j,\text{corr}}^{(t)} + q_{l|\text{wrong}}^{(t)} S_{l,j,\text{wrong}}^{(t)} \right) + (1 - p_{<l}^{(t)}) q_{<l|\text{wrong}}^{(t)} \bar{S}_{l,j}^{(t)}, \quad (43)$$

1593 where for the given reward  $R^{k^*}$  we define the conditional success probabilities:

$$1594 \begin{aligned} 1595 q_{l|\text{corr}}^{(t)} &:= \mathbb{P}^{(t)}(R^{k^*} = 1 | C_{<l} = 1, i_l = i_l^*), \\ 1596 q_{l|\text{wrong}}^{(t)} &:= \mathbb{P}^{(t)}(R^{k^*} = 1 | C_{<l} = 1, i_l \neq i_l^*), \\ 1597 q_{<l|\text{wrong}}^{(t)} &:= \mathbb{P}^{(t)}(R^{k^*} = 1 | C_{<l} = 0). \end{aligned} \quad (44)$$

1602 (a) Final-answer reward  $R^{k^*} = R^{f_{S^*}}$ . Let  $p_{\text{tail}}^{(t)}(l+1)$  denote the probability of completing the remaining true path from step  $l+1$  onward under  $\mathbf{W}^{(t)}$ , and let  $\rho_{\text{spur}, \geq l}^{(t)}$  be the spurious-success rate conditioned on having a wrong index at step  $r \geq l$  (first deviation at  $\geq l$ ), and  $\rho_{\text{spur}, < l}^{(t)}$  for having deviated before step  $l$ . Then

$$1603 q_{l|\text{corr}}^{(t)} = p_{\text{tail}}^{(t)}(l+1) + (1 - p_{\text{tail}}^{(t)}(l+1)) \rho_{\text{spur}, \geq l+1}^{(t)}, \quad (45)$$

$$1604 q_{l|\text{wrong}}^{(t)} = \rho_{\text{spur}, \geq l}^{(t)}, \quad q_{<l|\text{wrong}}^{(t)} = \rho_{\text{spur}, < l}^{(t)}. \quad (46)$$

1605 For the parity case with  $\mathbf{x} \sim \text{Unif}\{0, 1\}^d$  and XOR kernels,  $\rho_{\text{spur}, \geq l}^{(t)} = \rho_{\text{spur}, < l}^{(t)} = 1/2$  for all  $l$ .

1606 (b) Subtask-family reward  $R^{k^*} = R^{f_{S^*}}$ . Let  $T^{(t)} \in \{1, \dots, k^* + 1\}$  be the (random) termination step (the first step where EOS is sampled). For each depth  $r \in [k^*]$ , define the event

$$1607 U_r := \{T^{(t)} = r + 1 \text{ and the pre-EOS token equals } \boldsymbol{\mu}^{f_{S^*}}(\mathbf{x})\}.$$

1608 Then  $R^{f_{S^*}} = \mathbf{1}\{\cup_{r=1}^{k^*} U_r\}$ , and the stepwise expected gradient admits the depth-wise expansion

$$1609 G_{l,j}^{(t)} = \sum_{r=l}^{k^*} \mathbb{E}^{(t)} [\mathbf{1}\{U_r\} \alpha_l^{(t)}(j) \eta_l^{(t)}(j)]. \quad (47)$$

Consequently, for every  $l \leq k^*$ ,

$$G_{l,j}^{(t)}|_{R=R^{\mathcal{F}S^*}} = G_{l,j}^{(t)}|_{R=R^f S^*} + \sum_{r=l}^{k^*-1} \mathbb{E}^{(t)}[\mathbf{1}\{U_r\} \alpha_l^{(t)}(j) \eta_l^{(t)}(j)] \geq G_{l,j}^{(t)}|_{R=R^f S^*},$$

which shows the subtask-family reward strictly adds nonnegative depth-wise contributions on shared prefixes.

**Explicit per-depth success probabilities for (b).** For  $s \in [k^*]$ , define the per-step true-path continuation probabilities and their products

$$\pi_s^{(t)} := \mathbb{P}^{(t)}(i_s = i_s^* \mid C_{<s} = 1), \quad \Pi_{a:b}^{(t)} := \prod_{s=a}^b \pi_s^{(t)} \quad (\Pi_{a:b}^{(t)} \equiv 1 \text{ if } a > b).$$

Define the path events

$$A_r := \{C_{<r} = 1, i_r = i_r^*\}, \quad B_r := \{C_{<r} = 1, i_r \neq i_r^*\}, \quad D_r := \{C_{<r} = 0\}.$$

For each  $r \geq l$ , define branch-specific one-step termination probabilities at depth  $r+1$ :

$$\theta_{r+1,A}^{(t)} := \mathbb{P}^{(t)}(i_{r+1} = \text{EOS} \mid A_r), \quad \theta_{r+1,B}^{(t)} := \mathbb{P}^{(t)}(i_{r+1} = \text{EOS} \mid B_r), \quad \theta_{r+1,D}^{(t)} := \mathbb{P}^{(t)}(i_{r+1} = \text{EOS} \mid D_r),$$

and subtask-match probabilities at depth  $r$  under the three branches (task dependent through  $\{\Phi_l\}$ ):

$$\rho_{r|A}^{(t)} := \mathbb{P}^{(t)}(\text{subtask match at } r \mid A_r), \quad \rho_{r|B}^{(t)} := \mathbb{P}^{(t)}(\text{subtask match at } r \mid B_r), \quad \rho_{r|D}^{(t)} := \mathbb{P}^{(t)}(\text{subtask match at } r \mid D_r).$$

Then the event-level probabilities that enter equation 47 decompose explicitly as

$$\begin{aligned} & \mathbb{P}^{(t)}(U_r \mid C_{<l} = 1, i_l = i_l^*) \\ &= \underbrace{\rho_{r|A}^{(t)} \Pi_{l+1:r}^{(t)} \theta_{r+1,A}^{(t)}}_{\text{true-path to } r} + \underbrace{\rho_{r|B}^{(t)} \Pi_{l+1:r-1}^{(t)} (1 - \pi_r^{(t)}) \theta_{r+1,B}^{(t)}}_{\text{first deviation at } r} + \underbrace{\rho_{r|D}^{(t)} (1 - \Pi_{l+1:r-1}^{(t)}) \theta_{r+1,D}^{(t)}}_{\text{deviation before } r}, \\ & \mathbb{P}^{(t)}(U_r \mid C_{<l} = 1, i_l \neq i_l^*) \\ &= \mathbf{1}\{r = l\} \rho_{r|B}^{(t)} \theta_{l+1,B}^{(t)} + \mathbf{1}\{r > l\} \rho_{r|D}^{(t)} \theta_{r+1,D}^{(t)}, \\ & \mathbb{P}^{(t)}(U_r \mid C_{<l} = 0) \\ &= \rho_{r|D}^{(t)} \theta_{r+1,D}^{(t)}. \end{aligned} \tag{48}$$

Consequently, identifying with the conditional success probabilities in Eq. equation 43 for  $R = R^{\mathcal{F}S^*}$ , the per-step success factors equal

$$\begin{aligned} q_{l|\text{corr}}^{(t)} &= \sum_{r=l}^{k^*} \mathbb{P}^{(t)}(U_r \mid C_{<l} = 1, i_l = i_l^*), \\ q_{l|\text{wrong}}^{(t)} &= \sum_{r=l}^{k^*} \mathbb{P}^{(t)}(U_r \mid C_{<l} = 1, i_l \neq i_l^*) = \rho_{l|B}^{(t)} \theta_{l+1,B}^{(t)} + \sum_{r=l+1}^{k^*} \rho_{r|D}^{(t)} \theta_{r+1,D}^{(t)}, \\ q_{<l|\text{wrong}}^{(t)} &= \sum_{r=l}^{k^*} \mathbb{P}^{(t)}(U_r \mid C_{<l} = 0) = \sum_{r=l}^{k^*} \rho_{r|D}^{(t)} \theta_{r+1,D}^{(t)}. \end{aligned} \tag{49}$$

In the parity case,  $\rho_{r|A}^{(t)} \equiv 1$  and  $\rho_{r|B}^{(t)} = \rho_{r|D}^{(t)} \equiv 1/2$ , yielding the explicit decomposition Eq. equation 50 and, under uniform 2S-ART, the closed form Eq. equation 53.

**Parity specialization** Under the XOR kernel for subtasks with  $\mathbf{x} \sim \text{Unif}\{0, 1\}^d$ , the per-branch subtask-match probabilities satisfy  $\rho_{r|A}^{(t)} \equiv 1$  and  $\rho_{r|B}^{(t)} = \rho_{r|D}^{(t)} \equiv \frac{1}{2}$ . Writing the one-step termination probability at  $r+1$  under a condition  $E \in \{A_r, B_r, D_r\}$  as

$$\theta_{r+1}^{(t)}(E) := \mathbb{P}^{(t)}(i_{r+1} = \text{EOS} \mid E),$$

the probability of the event  $U_r$  decomposes as

$$\mathbb{P}^{(t)}(U_r) = \underbrace{\mathbb{P}^{(t)}(A_r) \theta_{r+1}^{(t)}(A_r)}_{\text{true-path contribution}} + \underbrace{\frac{1}{2} \mathbb{P}^{(t)}(B_r) \theta_{r+1}^{(t)}(B_r) + \frac{1}{2} \mathbb{P}^{(t)}(D_r) \theta_{r+1}^{(t)}(D_r)}_{\text{spurious (reward-hacking) contribution}}. \quad (50)$$

Moreover, writing  $p_{\text{path}}^{(t)}(r) := \mathbb{P}^{(t)}(A_r)$  and  $p_{\text{dev}, \geq r}^{(t)} := \mathbb{P}^{(t)}(B_r)$ ,  $p_{\text{dev}, < r}^{(t)} := \mathbb{P}^{(t)}(D_r)$ , we have bounds

$$\mathbb{P}^{(t)}(U_r) \geq p_{\text{path}}^{(t)}(r) \inf \theta_{r+1}^{(t)}(A_r), \quad \mathbb{P}^{(t)}(U_r) \leq p_{\text{path}}^{(t)}(r) \sup \theta_{r+1}^{(t)}(A_r) + \frac{1}{2} (p_{\text{dev}, \geq r}^{(t)} \sup \theta_{r+1}^{(t)}(B_r) + p_{\text{dev}, < r}^{(t)} \sup \theta_{r+1}^{(t)}(D_r)). \quad (51)$$

If, in addition, the EOS policy at step  $r+1$  is conditionally independent of the path branch (same marginal  $\bar{\theta}_{r+1}^{(t)}$ ),

$$\mathbb{P}^{(t)}(U_r) = \bar{\theta}_{r+1}^{(t)} \left( p_{\text{path}}^{(t)}(r) + \frac{1}{2} (p_{\text{dev}, \geq r}^{(t)} + p_{\text{dev}, < r}^{(t)}) \right). \quad (52)$$

Under the base uniform 2S-ART for parity, namely for  $s \in [k^*]$  and  $r \in [k^*]$ ,

$$\pi_s^{(t)} = \frac{1}{d-s+1}, \quad \theta_{r+1,A}^{(t)} = \theta_{r+1,B}^{(t)} = \theta_{r+1,D}^{(t)} = \frac{1}{d-r+1},$$

define

$$P_{r-1} := \prod_{s=1}^{r-1} \frac{1}{d-s+1} = \frac{1}{d(d-1) \cdots (d-r+2)} \quad (\text{and } P_0 := 1).$$

Then for every  $r \in [k^*]$ ,

$$\mathbb{P}^{(t)}(U_r) = \frac{P_{r-1}}{(d-r+1)^2} + \frac{1}{2} \frac{P_{r-1}(d-r)}{(d-r+1)^2} + \frac{1}{2} \frac{1-P_{r-1}}{d-r+1} = \frac{P_{r-1}(d-r+2)}{2(d-r+1)^2} + \frac{1-P_{r-1}}{2(d-r+1)}. \quad (53)$$

**Proof.** Starting from Lemma 5 and Lemma 3, for any  $l, j$  we have

$$G_{l,j}^{(t)} = \mathbb{E}^{(t)} [R^{k^*} \alpha_l^{(t)}(j) \eta_l^{(t)}(j)].$$

Condition on  $C_{<l}$  and  $i_l$  (tower property):

$$\begin{aligned} G_{l,j}^{(t)} &= \mathbb{E}^{(t)} \left[ \mathbb{E}^{(t)} [R^{k^*} \alpha_l^{(t)}(j) \eta_l^{(t)}(j) \mid C_{<l}, i_l] \right] \\ &= p_{<l}^{(t)} \mathbb{E}^{(t)} \left[ \alpha_l^{(t)}(j) \eta_l^{(t)}(j) \mathbb{E}^{(t)} [R^{k^*} \mid C_{<l} = 1, i_l] \mid C_{<l} = 1 \right] \\ &\quad + (1 - p_{<l}^{(t)}) \mathbb{E}^{(t)} \left[ \alpha_l^{(t)}(j) \eta_l^{(t)}(j) \mathbb{E}^{(t)} [R^{k^*} \mid C_{<l} = 0] \mid C_{<l} = 0 \right]. \end{aligned}$$

This gives equation 43 once we identify  $q_l^{(t)}$  by whether  $i_l = i_l^*$ .

For (a),  $R^{k^*} = R^{f s^*}$ : on the true path through step  $l$  with  $i_l = i_l^*$ , success thereafter requires completing the remainder true path; otherwise, success can occur spuriously, which yields the formulas for  $q_{l|\text{corr}}^{(t)}$ ,  $q_{l|\text{wrong}}^{(t)}$  and  $q_{<l|\text{wrong}}^{(t)}$ .

For (b),  $R^{k^*} = R^{f s^*}$ : success occurs at the unique termination depth  $r$  where the pre-EOS token equals the subtask token; hence  $R^{f s^*} = \sum_{r=1}^{k^*} \mathbf{1}\{U_r\}$ . Expanding  $\mathbf{1}\{\cup_r U_r\}$  as  $\sum_r \mathbf{1}\{U_r\}$ , exchanging sum and expectation, and applying the same conditioning as above yields

$$G_{l,j}^{(t)} = \sum_{r=l}^{k^*} \mathbb{E}^{(t)} [\mathbf{1}\{U_r\} \alpha_l^{(t)}(j) \eta_l^{(t)}(j)],$$

which implies  $G_{l,j}^{(t)}|_{R=R^{f s^*}} \geq G_{l,j}^{(t)}|_{R=R^{f s^*}}$  via the nonnegative extra depths  $r \in \{l, \dots, k^* - 1\}$ . In the parity case, producing the correct subtask token at depth  $r$  deterministically fixes  $z_r$ ; any immediate-EOS strategy yields the subtask reward at  $r$ , while spurious subtask matches contribute via the corresponding conditional probabilities.  $\square$

**Lemma 8** (Finite-sample variance of REINFORCE objective and gradient (general post-training)).  
Let  $(\mathbf{x}^{(s)}, \tau^{(s)})_{s=1}^n$  be i.i.d., where  $\mathbf{x}^{(s)} \sim \mathcal{P}_x$  and  $\tau^{(s)} \sim p_{\mathbf{W}^{k^*}}(\cdot | \mathbf{x}^{(s)})$ . Define the Monte Carlo estimators

$$\begin{aligned}\widehat{\mathcal{J}}_n^{k^*}(\mathbf{W}^{k^*}) &:= \frac{1}{n} \sum_{s=1}^n R^{k^*}(\boldsymbol{\mu}^{z_{k^*}}(\tau_{-2}^{(s)})), \\ \widehat{\mathbf{g}}_n^{k^*}(\mathbf{W}^{k^*}) &:= \frac{1}{n} \sum_{s=1}^n R^{k^*}(\widehat{\mathbf{p}}^{z_{k^*}}(\tau_{-2}^{(s)})) \sum_{l=1}^{k^*+1} \nabla_{\mathbf{W}^{k^*}} \log \pi_{\mathbf{W}^{k^*}}(i_l^{(s)} | \mathbf{x}^{(s)}, \widehat{\mathbf{p}}^{z_{1:l}}).\end{aligned}$$

Assume  $R^{k^*} \in \{0, 1\}$ , the regularity conditions of Lemma 5, and Orthogonal Block Isolation (Lemma 3). Let

$$p_{\text{succ}}(\mathbf{W}^{k^*}) := \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_x} \left[ \mathbb{P}_{\tau \sim p_{\mathbf{W}^{k^*}}(\cdot | \mathbf{x})} \{R^{k^*}(\widehat{\mathbf{p}}^{z_{k^*}}(\tau_{-2})) = 1\} \right] = \mathbb{E}_{\mathbf{x}, \tau} [R^{k^*}(\widehat{\mathbf{p}}^{z_{k^*}}(\tau_{-2}))].$$

Then:

1) (Objective) Unbiasedness and variance:

$$\mathbb{E}[\widehat{\mathcal{J}}_n^{k^*}] = \mathcal{J}_{\text{REINFORCE}}^{k^*}, \quad \text{Var}[\widehat{\mathcal{J}}_n^{k^*}] = \frac{1}{n} p_{\text{succ}}(1 - p_{\text{succ}}) \leq \frac{1}{4n}. \quad (54)$$

In particular, under near-uniform initialization and using equation 31, we have

$$p_{\text{succ}} \leq \rho_{\text{spur}} + (1 - \rho_{\text{spur}}) \prod_{l=1}^{k^*+1} \frac{1}{d_l} \quad (\text{if } d_l = \Theta(d), p_{\text{succ}} \leq \rho_{\text{spur}} + (1 - \rho_{\text{spur}}) \Theta(d^{-(k^*+1)})), \quad (55)$$

so that

$$\text{Var}[\widehat{\mathcal{J}}_n^{k^*}] \leq \frac{1}{n} \left( \rho_{\text{spur}} + (1 - \rho_{\text{spur}}) d^{-(k^*+1)} \right) \left( 1 - \rho_{\text{spur}} - (1 - \rho_{\text{spur}}) d^{-(k^*+1)} \right) \leq \frac{1}{4n}. \quad (56)$$

2) (Gradient) Unbiasedness and covariance. Let  $\boldsymbol{\mu} := \nabla_{\mathbf{W}^{k^*}} \mathcal{J}_{\text{REINFORCE}}^{k^*}$  be the population gradient. Then

$$\mathbb{E}[\widehat{\mathbf{g}}_n^{k^*}] = \boldsymbol{\mu}, \quad \text{Cov}[\widehat{\mathbf{g}}_n^{k^*}] = \frac{1}{n} \underbrace{\left( \mathbb{E}_{\mathbf{x}, \tau} \left[ R^{k^*} \sum_{l,t=1}^{k^*+1} \nabla \log \pi_l \otimes \nabla \log \pi_t \right] - \boldsymbol{\mu} \otimes \boldsymbol{\mu} \right)}_{=: \boldsymbol{\Sigma}_{\text{pop}}}, \quad (57)$$

where  $\pi_l$  abbreviates  $\pi_{\mathbf{W}^{k^*}}(i_l | \mathbf{x}, \widehat{\mathbf{p}}^{z_{1:l}})$  and  $\otimes$  is the outer product in parameter space.

Furthermore, using Lemma 3, write a block-orthogonal expansion

$$\nabla \log \pi_l = \sum_{k \in \mathcal{I}_l} \alpha_l^{k^*}(k) \eta_l^{k^*}(k) \mathbf{B}_{l,k}, \quad \mathbf{B}_{l,k} := \mathbf{p}_k \mathbf{p}_{c_l}^\top.$$

Denote the block coefficient  $C_{l,k} := R^{k^*} \alpha_l^{k^*}(k) \eta_l^{k^*}(k)$ . Then the scalar variance in each orthogonal block is

$$\text{Var} \left[ \left\langle \widehat{\mathbf{g}}_n^{k^*}, \frac{\mathbf{B}_{l,k}}{\|\mathbf{B}_{l,k}\|_F} \right\rangle \right] = \frac{1}{n} \left( \mathbb{E}[C_{l,k}^2] - \mathbb{E}[C_{l,k}]^2 \right), \quad (58)$$

and the Frobenius-mean-square error (variance) satisfies

$$\mathbb{E} \left\| \widehat{\mathbf{g}}_n^{k^*} - \boldsymbol{\mu} \right\|_F^2 = \frac{1}{n} \underbrace{\left( \mathbb{E} \left[ R^{k^*} \sum_{l=1}^{k^*+1} \sum_{k \in \mathcal{I}_l} (\alpha_l^{k^*}(k))^2 (\eta_l^{k^*}(k))^2 \|\mathbf{B}_{l,k}\|_F^2 \right] - \|\boldsymbol{\mu}\|_F^2 \right)}_{=: \Xi}. \quad (59)$$

Assume bounded embeddings and temperature  $\beta > 0$  so that  $|\eta_l^{k^*}(k)| \leq 4/\beta$ , and define the concentration factor

$$S_l := \mathbb{E}_{\mathbf{x}, \tau} \left[ \sum_{k \in \mathcal{I}_l^{\text{legal}}} (\alpha_l^{k^*}(k))^2 \right] \in \left[ \frac{1}{d_l}, 1 \right].$$

1782 Let  $C_{\text{pos}} := \max_{l,k} \|\mathbf{p}_k\|_2^2 \|\mathbf{p}_{c_l}\|_2^2 = 1$ . Then

$$1783 \Xi \leq \mathbb{E}[R^{k^*}] \sum_{l=1}^{k^*+1} \left( \frac{16}{\beta^2} C_{\text{pos}} \right) \mathbb{E} \left[ \sum_{k \in \mathcal{I}_l^{\text{legal}}} (\alpha_l^{k^*}(k))^2 \right] \quad (60)$$

$$1784 = \frac{16}{\beta^2} p_{\text{succ}} \sum_{l=1}^{k^*+1} S_l. \quad (61)$$

1785 Consequently,

$$1786 \mathbb{E} \|\widehat{\mathbf{g}}_n^{k^*} - \boldsymbol{\mu}\|_F^2 \leq \frac{1}{n} \frac{16}{\beta^2} p_{\text{succ}} \sum_{l=1}^{k^*+1} S_l. \quad (62)$$

1787 Two useful specializations:

- 1788 • (Near-uniform init)  $S_l = \Theta(d^{-1})$ , it holds that  $\mathbb{E} \|\widehat{\mathbf{g}}_n^{k^*} - \boldsymbol{\mu}\|_F^2 \leq \frac{1}{n} \frac{16k^*+1}{\beta^2 d} p_{\text{succ}}$ .
- 1789 • (Worst-case spiky attention)  $S_l \leq 1$ , yielding  $\mathbb{E} \|\widehat{\mathbf{g}}_n^{k^*} - \boldsymbol{\mu}\|_F^2 \leq \frac{1}{n} \frac{16}{\beta^2} p_{\text{succ}} (k^* + 1)$ , which holds for the entire post-training.

1790 Accordingly, the signal-to-noise ratio satisfies the general lower bound

$$1791 \text{SNR} := \frac{\|\boldsymbol{\mu}\|_F^2}{\mathbb{E} \|\widehat{\mathbf{g}}_n^{k^*} - \boldsymbol{\mu}\|_F^2} \geq \frac{n \|\boldsymbol{\mu}\|_F^2}{(16/\beta^2) p_{\text{succ}} \sum_l S_l},$$

1792 and reduces to the near-uniform initialization scaling when  $S_l \approx 1/d_l$ .

1793 **Proof.** We proceed in parts.

1794 (A) Objective: unbiasedness and variance. Define the single-sample random variable

$$1795 Y := R^{k^*} (\boldsymbol{\mu}^{z_{k^*}(\tau_{-2})}) \in \{0, 1\}.$$

1796 By definition of  $\mathcal{J}_{\text{REINFORCE}}^{k^*}$  (Eq. (24)'s expectation target without the score factor),

$$1797 \mathbb{E}[Y] = p_{\text{succ}} (\mathbf{W}^{k^*}) = \mathcal{J}_{\text{REINFORCE}}^{k^*}.$$

1798 Hence  $\widehat{\mathcal{J}}_n^{k^*} = \frac{1}{n} \sum_{s=1}^n Y^{(s)}$  is unbiased. Since  $(Y^{(s)})$  are i.i.d. Bernoulli( $p_{\text{succ}}$ ),

$$1799 \text{Var}(\widehat{\mathcal{J}}_n^{k^*}) = \frac{1}{n} \text{Var}(Y) = \frac{1}{n} p_{\text{succ}}(1 - p_{\text{succ}}) \leq \frac{1}{4n}.$$

1800 (B) Upper bound on  $p_{\text{succ}}$  under 2S-ART near-uniformity. Suppose at each step  $l$  there are  $d_l \geq 1$  legal children and success requires selecting a unique correct legal child. Under 2S-ART, conditionally almost surely  $\pi_l(\text{correct} \mid \cdot) \leq 1/d_l$ . Therefore, for any fixed  $(\mathbf{x}, \tau)$ -measurable history,

$$1801 \mathbb{P}(Y = 1 \mid \text{history}) = \prod_{l=1}^{k^*+1} \pi_l(\text{correct} \mid \cdot) \leq \prod_{l=1}^{k^*+1} \frac{1}{d_l},$$

1802 and taking expectation over the history yields  $p_{\text{succ}} \leq \prod_l d_l^{-1}$ . If  $d_l \equiv d$ , then  $p_{\text{succ}} \leq d^{-(k^*+1)}$ . Substituting into part (A) gives the displayed variance bound for  $\widehat{\mathcal{J}}_n^{k^*}$ .

1803 (C) Gradient: unbiasedness and covariance. Let the single-sample gradient random tensor be

$$1804 \mathbf{G} := R^{k^*} \sum_{l=1}^{k^*+1} \nabla_{\mathbf{W}^{k^*}} \log \pi_{\mathbf{W}^{k^*}}(i_l \mid \mathbf{x}, \hat{\mathbf{p}}^{z_{1:l}}).$$

1805 By Lemma 5,  $\mathbb{E}[\mathbf{G}] = \nabla \mathcal{J}_{\text{REINFORCE}}^{k^*} =: \boldsymbol{\mu}$ , hence  $\widehat{\mathbf{g}}_n^{k^*} = \frac{1}{n} \sum_{s=1}^n \mathbf{G}^{(s)}$  is unbiased. With i.i.d. samples,

$$1806 \text{Cov}(\widehat{\mathbf{g}}_n^{k^*}) = \frac{1}{n} \text{Cov}(\mathbf{G}) = \frac{1}{n} \left( \mathbb{E}[\mathbf{G} \otimes \mathbf{G}] - \boldsymbol{\mu} \otimes \boldsymbol{\mu} \right).$$

1836 Since  $R^{k^*} \in \{0, 1\}$ ,  $R^{k^2} = R^{k^*}$ , giving

$$1837 \mathbb{E}[\mathbf{G} \otimes \mathbf{G}] = \mathbb{E}\left[R^{k^*} \sum_{l,t} \nabla \log \pi_l \otimes \nabla \log \pi_t\right] =: \mathbf{\Sigma}_{\text{pop}}.$$

1840 This proves the displayed covariance formula.

1841 (D) Block-wise variance via OBI. By Lemma 3, for each step  $l$ ,

$$1842 \nabla \log \pi_l = \sum_{k \in \mathcal{I}_l} \alpha_l^{k^*}(k) \eta_l^{k^*}(k) \mathbf{B}_{l,k}, \quad \mathbf{B}_{l,k} := \mathbf{p}_k \mathbf{p}_{c_l}^\top, \quad \langle \mathbf{B}_{l,k}, \mathbf{B}_{l',k'} \rangle_F = 0 \quad ((l, k) \neq (l', k')).$$

1845 Hence

$$1846 \mathbf{G} = \sum_{l,k} C_{l,k} \mathbf{B}_{l,k}, \quad C_{l,k} := R^{k^*} \alpha_l^{k^*}(k) \eta_l^{k^*}(k).$$

1848 Projecting onto a unit-Frobenius block direction gives a scalar average of i.i.d. terms:

$$1850 \left\langle \widehat{\mathbf{g}}_n^{k^*}, \frac{\mathbf{B}_{l,k}}{\|\mathbf{B}_{l,k}\|_F} \right\rangle = \frac{1}{n} \sum_{s=1}^n C_{l,k}^{(s)}.$$

1853 Thus  $\text{Var}(\langle \widehat{\mathbf{g}}_n^{k^*}, \mathbf{B}_{l,k} / \|\mathbf{B}_{l,k}\|_F \rangle) = \frac{1}{n} (\mathbb{E}[C_{l,k}^2] - \mathbb{E}[C_{l,k}]^2)$ .

1854 (E) Frobenius MSE identity. Using orthogonality of blocks,

$$1855 \|\mathbf{G}\|_F^2 = R^{k^*} \sum_{l,k} (\alpha_l^{k^*}(k))^2 (\eta_l^{k^*}(k))^2 \|\mathbf{B}_{l,k}\|_F^2.$$

1858 Since  $\mathbb{E}\|\widehat{\mathbf{g}}_n^{k^*} - \boldsymbol{\mu}\|_F^2 = \frac{1}{n} (\mathbb{E}\|\mathbf{G}\|_F^2 - \|\boldsymbol{\mu}\|_F^2)$  for i.i.d. averages, we obtain the displayed identity with

$$1860 \Xi := \mathbb{E}\left[R^{k^*} \sum_{l=1}^{k^*+1} \sum_{k \in \mathcal{I}_l} (\alpha_l^{k^*}(k))^2 (\eta_l^{k^*}(k))^2 \|\mathbf{B}_{l,k}\|_F^2\right].$$

1863 (F) Concrete upper bound for general post-training. Assume bounded embeddings and temperature  $\beta > 0$  so that  $|\eta_l^{k^*}(k)| \leq 4/\beta$ . Define  $S_l := \mathbb{E}[\sum_{k \in \mathcal{I}_l^{\text{legal}}} (\alpha_l^{k^*}(k))^2] \in [1/d_l, 1]$ . With  $C_{\text{pos}} := \max_{l,k} \|\mathbf{B}_{l,k}\|_F^2$ ,

$$1864 \left| \eta_l^{k^*}(k) \right| = \frac{|\langle \mathbf{p}_k - \widehat{\mathbf{p}}^{\text{att}}, \mathbf{p}_{i_l} - \sum_r p_{\text{vocab}}(r) \mathbf{p}_r \rangle|}{\beta} \leq \frac{\|\mathbf{p}_k - \widehat{\mathbf{p}}^{\text{att}}\|_2 \|\mathbf{p}_{i_l} - \sum_r p_{\text{vocab}}(r) \mathbf{p}_r\|_2}{\beta} \leq \frac{4}{\beta},$$

1869 where we used triangle inequality and that each difference of two convex combinations of bounded, pairwise-orthogonal embeddings has norm at most 2. Denoting  $C_{\text{pos}} := \max_{l,k} \|\mathbf{B}_{l,k}\|_F^2 = \max_{l,k} \|\mathbf{p}_k\|_2^2 \|\mathbf{p}_{c_l}\|_2^2$ , we bound

$$1873 \Xi \leq \mathbb{E}[R^{k^*}] \sum_{l=1}^{k^*+1} \sum_{k \in \mathcal{I}_l^{\text{legal}}} \left(\frac{4}{\beta}\right)^2 C_{\text{pos}} (\alpha_l^{k^*}(k))^2 = \frac{16}{\beta^2} C_{\text{pos}} p_{\text{succ}} \sum_{l=1}^{k^*+1} S_l. \quad (63)$$

1876 Absorbing  $c_\alpha^2$  into the constant (redefining the front coefficient) matches the displayed bound on  $\Xi$ . Finally, recall the exact identity

$$1879 \mathbb{E}\|\widehat{\mathbf{g}}_n^{k^*} - \boldsymbol{\mu}\|_F^2 = \frac{1}{n} (\Xi - \|\boldsymbol{\mu}\|_F^2).$$

1881 By Jensen's inequality  $\|\mathbb{E}[\mathbf{G}]\|_F^2 \leq \mathbb{E}\|\mathbf{G}\|_F^2$ , we have  $\|\boldsymbol{\mu}\|_F^2 \leq \Xi$ , so  $\Xi - \|\boldsymbol{\mu}\|_F^2 \leq \Xi$ . Therefore, for upper bounds it is valid to drop the negative term and use

$$1883 \mathbb{E}\|\widehat{\mathbf{g}}_n^{k^*} - \boldsymbol{\mu}\|_F^2 \leq \frac{1}{n} \Xi \leq \frac{1}{n} \frac{16 C_{\text{pos}}}{\beta^2} p_{\text{succ}} \sum_{l=1}^{k^*+1} \frac{1}{d_l}.$$

1887 Remark on the SNR line. From the variance bound just proved,

$$1888 \mathbb{E}\|\widehat{\mathbf{g}}_n^{k^*} - \boldsymbol{\mu}\|_F^2 \leq \frac{C}{n} p_{\text{succ}} \sum_{l=1}^{k^*+1} \frac{1}{d_l}$$

for a constant  $C$  depending only on  $(\beta, C_{\text{pos}})$  and the near-uniformity constant. Therefore

$$\text{SNR} = \frac{\|\boldsymbol{\mu}\|_F^2}{\mathbb{E}\|\hat{\boldsymbol{g}}_n^{k^*} - \boldsymbol{\mu}\|_F^2} \geq \frac{n \|\boldsymbol{\mu}\|_F^2}{C p_{\text{succ}} \sum_l d_l^{-1}}.$$

If, in addition, one has a non-degeneracy bound  $\|\boldsymbol{\mu}\|_F^2 \leq C' p_{\text{succ}}$  with  $C'$  independent of depth (e.g., when the coefficient mass  $\sum_{l,k} \mathbb{E}[|\alpha_l \eta_l|]$  remains  $\mathcal{O}(\sum_l d_l^{-1})$ ), this yields the crude scaling

$$\text{SNR} = \mathcal{O}\left(\frac{n p_{\text{succ}}}{\sum_l d_l^{-1}}\right),$$

which reduces to  $\mathcal{O}(n p_{\text{succ}} d / (k^* + 1))$  for  $d_l \equiv d$ . The main message is that under 2S-ART,  $p_{\text{succ}}$  decays exponentially with depth, making SNR scale poorly with  $k^*$  unless  $n$  grows accordingly.  $\square$

**Theorem 10** (Per-step REINFORCE sample size for block dominance). *Fix step  $\ell$  and confidence  $\delta \in (0, 1)$ . Let  $|\mathcal{I}_\ell| = \Theta(d)$  and let  $i_\ell^* \in \mathcal{I}_\ell$  denote the unique correct child. For any  $j \in \mathcal{I}_\ell$ , given  $n$  i.i.d. pairs  $(\mathbf{x}^{(s)}, \tau^{(s)})$ , define the empirical block mean*

$$\hat{\mu}_{\ell,j} := \frac{1}{n} \sum_{s=1}^n R^{k^*} \alpha_\ell^{(s)}(j) \eta_\ell^{(s)}(j), \quad R^{k^*} \in \{\mathbf{R}_x^{f_{S_*}}(\cdot), \mathbf{R}_x^{\mathcal{F}_{S_*}}(\cdot)\}.$$

Consider three scenario:

- (A) No-curriculum ( $R_x^{f_{S_*}}(\cdot)$  as oracle; Algorithm. 1).
- (B) Depth-increasing Curriculum ( $\mathbf{R}_x^{\mathcal{F}_{S_*}}(\cdot)$  as oracle; Algorithm. 2).
- (C) Hint-decreasing Curriculum ( $\mathbf{R}_x^{\mathcal{F}_{S_*}}(\cdot)$  as oracle; Algorithm. 3).

Then, for the target  $f_{S_*} \in \mathcal{F}_{2\text{S-ART}}$  with  $|S_*| = k^* + 1$ ,  $S_* = \{i_1^*, i_2^*, \dots, i_{k^*}^*, d + 1\}$ , the sample complexity  $n_\ell(\delta)$  to ensure with probability at least  $1 - \delta$ ,  $\max_{j \in \mathcal{I}_\ell} |\hat{\mu}_{\ell,j} - \mathbb{E}[\hat{\mu}_{\ell,j}]| \leq \Theta(\mu_{\ell,i_\ell^*} - \max_{j \neq i_\ell^*} \mu_{\ell,j})$  in case A – C is

$$\begin{aligned} \text{(A)} \quad n_\ell(\delta) &\geq \tilde{\Omega}\left(d^{2(k^*+2-\ell)-2} (1 - \rho_{\text{spur}}^{\text{sup}, > \ell})^{-2}\right), \\ \text{(B)} \quad n_\ell(\delta) &\leq \tilde{\mathcal{O}}\left(d^2 (1 - \rho_{\text{spur}}^{\text{sup}, \ell})^{-2}\right), \\ \text{(C)} \quad n_\ell(\delta) &\leq \tilde{\mathcal{O}}\left(d^2 (1 - \rho_{\text{spur}}^{\text{sup}, k^*+1-\ell})^{-2}\right), \end{aligned} \tag{64}$$

where  $\tilde{\Omega}, \tilde{\mathcal{O}}$  hide polylogarithmic factors in  $d$  and  $1/\delta$  and absolute constants depending on  $\beta$ .

**Proof.** We quantify, at a fixed step  $\ell$ , how many i.i.d. inputs  $\mathbf{x}^{(s)} \sim \text{Unif}([K]^d)$  suffice so that, with probability at least  $1 - \delta$ , a one-step REINFORCE update increases the correct block more than any competitor by at least a fixed margin  $C > 0$ .

**Distributional conventions and the REINFORCE gradient.** For each i.i.d. draw  $\mathbf{x}^{(s)} \sim \text{Unif}([K]^d)$ , the base model  $\text{TF}_{\text{base}}$  samples a trajectory  $\tau^{(s)} \sim p_{\mathbf{W}}(\cdot | \mathbf{x}^{(s)})$  (independently across  $s$ ), where  $\mathbf{W}$  is the current parameter. We write

$$\mathbb{E}[\cdot] := \mathbb{E}_{\mathbf{x} \sim \text{Unif}([K]^d)} \mathbb{E}_{\tau \sim p_{\mathbf{W}}(\cdot | \mathbf{x})}[\cdot]$$

for the population expectation. The outcome reward is

$$R^{k^*} \in \{\mathbf{R}_x^{f_{S_*}}(\cdot), \mathbf{R}_x^{\mathcal{F}_{S_*}}(\cdot)\}, \quad R^{k^*} \in \{0, 1\}.$$

By Lemma 5 and Lemma 3, the REINFORCE gradient admits the block decomposition

$$\nabla_{\mathbf{W}} \mathcal{J}_{\text{REINFORCE}} = \mathbb{E}\left[R^{k^*} \sum_{t=1}^{k^*+1} \sum_{k \in \mathcal{I}_t} \alpha_t(k) \eta_t(k) \mathbf{p}_k \mathbf{p}_{c_t}^\top\right].$$

Projecting onto the step- $\ell$  block  $\mathbf{B}_{\ell,j} := \mathbf{p}_j \mathbf{p}_{c_\ell}^\top$  gives

$$\underbrace{\left\langle \nabla_{\mathbf{W}} \mathcal{J}_{\text{REINFORCE}}, \mathbf{B}_{\ell,j} \right\rangle}_{=: \mu_{\ell,j}} = \mathbb{E}\left[R^{k^*} \alpha_\ell(j) \eta_\ell(j)\right]. \tag{65}$$

1944 **Setup and block-wise variables.** At step  $\ell$ , fix  $j \in \mathcal{I}_\ell$  and define the single-sample block variable for  
 1945  $(\mathbf{x}, \tau)$ :

$$1946 X_{\ell,j} := R^{k^*} \alpha_\ell(j) \eta_\ell(j), \quad R^{k^*} \in \{0, 1\},$$

1947 where  $\alpha_\ell(j), \eta_\ell(j)$  are computed from  $(\mathbf{x}, \tau)$  and the current  $\mathbf{W}$ . By equation 65,  
 1948

$$1949 \mu_{\ell,j} := \mathbb{E}[X_{\ell,j}] = \left\langle \nabla_{\mathbf{w}} \mathcal{J}_{\text{REINFORCE}}, \mathbf{B}_{\ell,j} \right\rangle.$$

1950 Given  $n$  i.i.d. pairs  $(\mathbf{x}^{(s)}, \tau^{(s)})$ , define the empirical estimator and its range bound  
 1951

$$1952 \hat{\mu}_{\ell,j} := \frac{1}{n} \sum_{s=1}^n X_{\ell,j}^{(s)}, \quad |X_{\ell,j}| \leq B = \Theta\left(\frac{4}{d\beta}\right),$$

1953 where  $|\eta_\ell(j)| \leq 4/\beta$  (Lemma 3) and  $\alpha_\ell(j) = \Theta(d^{-1})$ . Note  $B$  is the Hoeffding range parameter;  
 1954 it is independent of the margin defined below.

1955 Let  $i_l^*$  be the unique correct child at depth  $\ell$ . We say the empirical gradient exhibits block dominance  
 1956 with margin  $C$  if

$$1957 \hat{\mu}_{\ell,i_l^*} - \max_{j \in \mathcal{I}_\ell \setminus \{i_l^*\}} \hat{\mu}_{\ell,j} \geq C.$$

1958 By Lemma 6, this implies the increase of  $s_\ell(i_l^*)$  exceeds all competitors by at least  $\eta C$  (up to  $\mathcal{O}(\eta^2)$ ),  
 1959 hence strictly enlarging the step- $\ell$  attention gap.

1960 **Concentration.** Hoeffding's inequality for bounded variables (Lemma 1) yields for any fixed  $j$   
 1961 (expectation over the joint randomness of  $(\mathbf{x}^{(s)}, \tau^{(s)})$ ):

$$1962 \Pr(|\hat{\mu}_{\ell,j} - \mu_{\ell,j}| \geq t) \leq 2 \exp\left(-\frac{2nt^2}{B^2}\right).$$

1963 A union bound over  $|\mathcal{I}_\ell| = \Theta(d)$  blocks gives that, with probability  $\geq 1 - \delta$ ,

$$1964 \max_{j \in \mathcal{I}_\ell} |\hat{\mu}_{\ell,j} - \mu_{\ell,j}| \leq t_n := \frac{B}{\sqrt{2n}} \sqrt{\log \frac{2d}{\delta}}. \quad (66)$$

1965 If the population margin

$$1966 \gamma_\ell := \mu_{\ell,i_l^*} - \max_{j \neq i_l^*} \mu_{\ell,j} > 0,$$

1967 then setting  $t_n \leq \gamma_\ell/4$  guarantees dominance with margin  $C = \gamma_\ell/2$ .

1968 **Master relation and the choice of  $C$ .** From equation 66, requiring empirical dominance with margin  
 1969  $C \leq \gamma_\ell/2$  is ensured by taking  $t_n \leq \gamma_\ell/2$ . Solving

$$1970 \frac{B}{\sqrt{2n}} \sqrt{\log \frac{2d}{\delta}} \leq \frac{\gamma_\ell}{2}$$

1971 for  $n$  yields the explicit sample-size condition

$$1972 n \geq \frac{2B^2}{\gamma_\ell^2} \log \frac{2d}{\delta}.$$

1973 Throughout this proof we fix  $C := \gamma_\ell/2$  so the target constant margin is explicit.

1974 **Population margin under three settings.** Using Lemma 7 and copied-PART near-uniformity  
 1975  $(\alpha_\ell(i_l^*) = \Theta(1/d))$ , with expectations taken over  $\mathbf{x} \sim \text{Unif}([K]^d)$  and  $\tau \sim p_{\mathbf{W}}(\cdot | \mathbf{x})$ :

- 1976 (A) **No curriculum; terminal oracle  $\mathbf{R}_{\mathbf{x}}^{f_{S^*}}(\cdot)$ .** Conditioning on a correct prefix to  $\ell$ , completing  
 1977 the true suffix (including EOS) has probability  $\Theta(d^{-(k^*+1-\ell)})$ , while any deviation may be  
 1978 accepted with probability at most  $\rho_{\text{spur}}^{\text{sup}, > \ell}$ . Under copied-PART near-uniformity  $\alpha_\ell(i_l^*) =$   
 1979  $\Theta(1/d)$  and  $4/\beta \geq \eta_l(i_l^*) \geq \Theta(1/\beta) \geq \eta_l(j), j \neq i_l^*$ , Lemma 7 implies

$$1980 q_{\ell|\text{corr}} - q_{\ell|\text{wrong}} = \Theta(p_{\text{tail}}(\ell+1) (1 - \rho_{\text{spur}}^{\text{sup}, > \ell})), \quad p_{\text{tail}}(\ell+1) = \Theta(d^{-(k^*+1-\ell)}),$$

1981 Consequently, in case (A) we have the matching scaling over the expected margin

$$1982 \gamma_\ell^A = \Theta\left(\frac{d^{-(k^*+2-\ell)}}{\beta} (1 - \rho_{\text{spur}}^{\text{sup}, > \ell})\right),$$

1983 up to absolute constants, which will be used to derive the lower bound on  $n_\ell(\delta)$ .

(B) **Depth-increasing curriculum; family oracle  $\mathbf{R}_x^{\mathcal{F}S^*}(\cdot)$  with external truncation at depth  $\ell$ .** Similar to Thm. 12,  $\alpha_{i_t^*}^{(\ell)} \geq 1 - \eta_\ell = \Theta(d^{-1})$  and any wrong child has acceptance at most  $\rho_{\text{spur}}^{\text{sup},\ell}$ . Also, we see that  $4/\beta \geq \eta_\ell(i_t^*) \geq O(1/\beta) \geq \eta_\ell(j), j \neq i_t^*$ . Thus

$$\gamma_\ell^B = \Theta\left(\frac{1}{\beta d^2} (1 - \rho_{\text{spur}}^{\text{sup},\ell})\right).$$

(C) **Hint-decreasing curriculum (reverse indexing); identical to (B) with  $\ell \mapsto k^*+1-\ell$ , such that  $\{\gamma_\ell^C\}_{\ell=1}^{k^*+1} = \{\gamma_{k^*+1-\ell}^B\}_{\ell=1}^{k^*+1}$ .**

From equation 66, taking  $t_n \leq \gamma_\ell/2$  ensures block dominance with margin  $\gamma_\ell/2$ . Solving for  $n$  with range bound  $B = 4/(\beta d)$  gives the master condition  $n \geq (2B^2/\gamma_\ell^2) \log \frac{2d}{\delta}$ . For case (A), substituting the upper bound on  $\gamma_\ell$  above yields the displayed lower bound on  $n_\ell$ . For cases (B) and (C), substituting the lower bounds on  $\gamma_\ell$  yields the displayed upper bounds on  $n_\ell$ . Here  $B$  is purely the bounded-range constant of  $X_{\ell,j}$  (Hoeffding parameter);  $\gamma_\ell$  is the population margin and does not influence the range. Under the near-uniformity and spurious-success cap assumptions used above, summarizing (A) as a  $\Theta(\cdot)$  scaling for  $\gamma_\ell$  is appropriate; we keep one-sided bounds when only sufficiency is needed for (B)/(C).  $\square$

**Theorem 11** (One-pass and online per-step schedules achieve per-step margin thresholds). For  $\forall \varepsilon > 0$ , consider the target per-step margin thresholds  $\{\Gamma_\ell\}_{\ell=1}^{k^*+1}$  for Lemma 4, with  $\varepsilon_\ell = \varepsilon/(k^*+1)$ , namely

$$\Gamma_\ell = \log\left(\frac{d-1}{\left(\frac{1}{2} + \frac{\beta}{2} \log\left(\frac{(K+1)((k^*+1)-\varepsilon)}{\varepsilon}\right)\right)^{-1} - 1}\right) \leq \Theta(\log(d))$$

For each step  $\ell$ , let  $n_\ell(\delta)$  be the per-step sample size from Theorem 10 that guarantees, with probability at least  $1 - \delta$ ,  $\max_{j \in \mathcal{I}_\ell} |\hat{\mu}_{\ell,j} - \mathbb{E}[\hat{\mu}_{\ell,j}]| \leq (\mu_{\ell,i_t^*} - \max_{j \neq i_t^*} \mu_{\ell,j})/4$ , guaranteeing that the empirical block-dominance margin to be no less than the thresholds  $\{\gamma_\ell^A/2\}_{\ell=1}^{k^*+1}, \{\gamma_\ell^B/2\}_{\ell=1}^{k^*+1}, \{\gamma_\ell^C/2\}_{\ell=1}^{k^*+1}$ .

There exist learning-rate schedules under which, with probability at least  $1 - \delta$ , the post-update per-step logit gaps satisfy  $\Delta_\ell \geq \Gamma_\ell$  for all  $\ell \in \{1, \dots, k^*+1\}$ , so that Lemma 4 applies. We state them for three settings:

(A) No curriculum (Algorithm. 1); terminal oracle  $\mathbf{R}_x^{\mathcal{F}S^*}(\cdot)$ . Draw a single batch of size

$$n_0 \geq \max_{\ell \in [k^*+1]} n_\ell(\delta/(k^*+1)) = \tilde{\Omega}(d^{2k^*+4} (1 - \rho_{\text{spur}}^{\text{sup},>1})^{-2}),$$

where the last equality takes the worst step (the first decision), and the extra union over steps would not influence the lower bound. Consider perform a single gradient step (one-shot) with learning rate  $\eta^A = \Theta(\beta \log(d) d^{k^*+1} (1 - \rho_{\text{spur}}^{\text{sup},>1})^{-1}) \geq \Theta(\Gamma_\ell/(\gamma_\ell^A/2))$ . Guarantee: with probability at least  $1 - \delta$ , after this one update we have  $\Delta_\ell \geq \Gamma_\ell$  for all  $\ell$ , including the final EOS step.

(B) Depth-increasing curriculum (Algorithm. 2); family oracle  $\mathbf{R}_x^{\mathcal{F}S^*}(\cdot)$  (external truncation). Sampling and updates are online per step: for  $\ell = 1, 2, \dots, k^*+1$  do (i) draw  $n_\ell = \tilde{\Theta}(d^2 (1 - \rho_{\text{spur}}^{\text{sup},\ell})^{-2})$  fresh samples; (ii) take one gradient step with learning rate  $\eta_\ell^B = \Omega(\beta \log(d) d^2 (1 - \rho_{\text{spur}}^{\text{sup},\ell})^{-1}) \geq \Theta(\Gamma_\ell/(\gamma_\ell^B/2))$ . Guarantee: with probability at least  $1 - \delta$ , after  $k^*+1$  updates,  $\Delta_\ell \geq \Gamma_\ell$  holds for all steps. The total sample size is  $\sum_{\ell=1}^{k^*+1} n_\ell = \tilde{\Theta}((k^*+1)d^2 (1 - \rho_{\text{spur}}^{\text{max}})^{-2})$ .

(C) Hint-decreasing curriculum (Algorithm. 3); family oracle  $\mathbf{R}_x^{\mathcal{F}S^*}(\cdot)$ . Sampling and updates are online per step in reverse: for  $\ell = k^*+1, \dots, 1$  do (i) draw  $n_\ell = \tilde{\Theta}(d^2 (1 - \rho_{\text{spur}}^{\text{sup},k^*+1-\ell})^{-2})$  fresh samples; (ii) take one gradient step with learning rate  $\eta_\ell^C = \Omega(\beta \log(d) d^2 (1 - \rho_{\text{spur}}^{\text{sup},k^*+1-\ell})^{-1}) \geq \Theta(\Gamma_\ell/(\gamma_\ell^C/2))$ . Guarantee: with probability at least  $1 - \delta$ , after  $k^*+1$  updates,  $\Delta_\ell \geq \Gamma_\ell$  holds for all steps. The total sample size is  $\sum_{\ell=1}^{k^*+1} n_\ell = \tilde{O}((k^*+1)d^2 (1 - \rho_{\text{spur}}^{\text{max}})^{-2})$ .

**Proof.** The proof is direct based on OBI (Lemma 3), as well as the convergence conditions in Lemma 6 for a target  $\varepsilon > 0$ .  $\square$

**Proof.** Proof of Theorem. 3. The proof follows by collaborating Theorem. 11 and Lemma 6.  $\square$

2052  
2053  
2054  
2055  
2056  
2057  
2058  
2059  
2060  
2061  
2062  
2063  
2064  
2065  
2066  
2067  
2068  
2069  
2070  
2071  
2072  
2073  
2074  
2075  
2076  
2077  
2078  
2079  
2080  
2081  
2082  
2083  
2084  
2085  
2086  
2087  
2088  
2089  
2090  
2091  
2092  
2093  
2094  
2095  
2096  
2097  
2098  
2099  
2100  
2101  
2102  
2103  
2104  
2105

---

**Algorithm 1** No-Curriculum REINFORCE Finetuning (A)

---

- 1: Inputs: input distribution  $\mathbf{x} \sim \text{Unif}([K]^d)$ , error budget  $\varepsilon > 0$ , target length  $k^*$ , terminal oracle  $\mathbf{R}_{\mathbf{x}}^{f_{S^*}}(\cdot)$ , sample complexity  $n_0 = \tilde{\Omega}(d^{2k^*+4}(1 - \rho_{\text{spur}}^{\text{sup}, >1})^{-2})$ , learning rate  $\eta^A = \Theta(\beta \log(d)d^{k^*+1}(1 - \rho_{\text{spur}}^{\text{sup}, >1})^{-1})$ .
  - 2: Initialize transformer parameters  $\mathbf{W}^{(0)}$ .
  - 3: Draw a single batch of  $n_0$  i.i.d. samples  $(\mathbf{x}^{(s)}, \tau^{(s)})_{s=1}^{n_0}$  where  $\mathbf{x}^{(s)} \sim \text{Unif}([K]^d)$  and  $\tau^{(s)} \sim p_{\mathbf{W}^{(0)}}(\cdot | \mathbf{x}^{(s)})$ .
  - 4: Compute empirical REINFORCE gradient:
  - 5:  $\hat{\mathbf{g}} = \frac{1}{n_0} \sum_{s=1}^{n_0} \mathbf{R}_{\mathbf{x}^{(s)}}^{f_{S^*}}(\hat{\mathbf{p}}^{z_{k^*}}(\tau_{-2}^{(s)})) \sum_{\ell=1}^{k^*+1} \nabla_{\mathbf{W}} \log \pi_{\mathbf{W}^{(0)}}(i_{\ell}^{(s)} | \mathbf{x}^{(s)}, \hat{\mathbf{p}}^{z_{1:\ell}})$ .
  - 6: Update parameters:  $\mathbf{W}^{(1)} = \mathbf{W}^{(0)} + \eta^A \cdot \hat{\mathbf{g}}$ .
  - 7: Return  $\mathbf{W}^{(1)}$ .
- 

---

**Algorithm 2** Depth-Increasing Curriculum REINFORCE Finetuning (B)

---

- 1: Inputs: input distribution  $\mathbf{x} \sim \text{Unif}([K]^d)$ , error budget  $\varepsilon > 0$ , target length  $k^*$ , family oracle  $\mathbf{R}_{\mathbf{x}}^{\mathcal{F}_{S^*}}(\cdot)$ , per-step sample complexity  $n_{\ell} = \tilde{\Theta}(d^2(1 - \rho_{\text{spur}}^{\text{sup}, \ell})^{-2})$ , learning rates  $\eta_{\ell}^B = \Omega(\beta \log(d)d^2(1 - \rho_{\text{spur}}^{\text{sup}, \ell})^{-1})$ .
  - 2: Initialize transformer parameters  $\mathbf{W}^{(0)}$ .
  - 3: **for**  $\ell = 1$  to  $k^* + 1$  **do**
  - 4:   Draw fresh samples  $(\mathbf{x}^{(s)}, \tau^{(s)})_{s=1}^{n_{\ell}}$  where  $\mathbf{x}^{(s)} \sim \text{Unif}([K]^d)$  and  $\tau^{(s)} \sim p_{\mathbf{W}^{(\ell-1)}}(\cdot | \mathbf{x}^{(s)})$ .
  - 5:   If  $\ell < k^* + 1$ , for each sample  $s$ , truncate trajectory at length  $\ell + 1$ :
  - 6:     If  $\tau^{(s)}$  contains EOS at position  $\leq \ell$ , discard sample.
  - 7:     Otherwise, append EOS to create truncated sequence  $\tau_{\text{trunc}}^{(s)}$ .
  - 8:   If  $\ell = k^* + 1$ ,  $\tau_{\text{trunc}}^{(s)} = \tau^{(s)}$ .
  - 9:   Compute empirical REINFORCE gradient for step  $\ell$ :
  - 10:    $\hat{\mathbf{g}}_{\ell} = \frac{1}{n_{\ell}} \sum_{s=1}^{n_{\ell}} \mathbf{R}_{\mathbf{x}^{(s)}}^{\mathcal{F}_{S^*}}(\tau_{\text{trunc}}^{(s)}) \nabla_{\mathbf{W}} \log \pi_{\mathbf{W}^{(\ell-1)}}(i_{\ell}^{(s)} | \mathbf{x}^{(s)}, \hat{\mathbf{p}}^{z_{1:\ell}})$ .
  - 11:   Update parameters:  $\mathbf{W}^{(\ell)} = \mathbf{W}^{(\ell-1)} + \eta_{\ell}^B \cdot \hat{\mathbf{g}}_{\ell}$ .
  - 12: **end for**
  - 13: Return  $\mathbf{W}^{(k^*+1)}$ .
- 

---

**Algorithm 3** Hint-Decreasing Curriculum REINFORCE Finetuning (C)

---

- 1: Inputs: input distribution  $\mathbf{x} \sim \text{Unif}([K]^d)$ , error budget  $\varepsilon > 0$ , target length  $k^*$ , family oracle  $\mathbf{R}_{\mathbf{x}}^{\mathcal{F}_{S^*}}(\cdot)$ , per-step sample complexity  $n_{\ell} = \tilde{\Theta}(d^2(1 - \rho_{\text{spur}}^{\text{sup}, k^*+1-\ell})^{-2})$ , learning rates  $\eta_{\ell}^C = \Omega(\beta \log(d)d^2(1 - \rho_{\text{spur}}^{\text{sup}, k^*+1-\ell})^{-1})$ .
  - 2: Initialize transformer parameters  $\mathbf{W}^{(0)}$ .
  - 3: **for**  $\ell = 1$  down to  $k^* + 1$  **do**
  - 4:   Draw fresh samples  $(\mathbf{x}^{(s)}, \tau^{(s)})_{s=1}^{n_{\ell}}$  where  $\mathbf{x}^{(s)} \sim \text{Unif}([K]^d)$  and  $\tau^{(s)} \sim p_{\mathbf{W}^{(k^*+1-\ell)}}(\cdot | \mathbf{x}^{(s)})$ .
  - 5:   For each sample  $s$ , provide hint prefix of length  $k^* + 1 - \ell$ :
  - 6:     Let  $\mathbf{h}^{(s)} = (i_1^*, i_2^*, \dots, i_{k^*+1-\ell}^*)$  be the correct prefix.
  - 7:     Generate remaining tokens  $\tau_{\text{hint}}^{(s)}$  from  $\mathbf{W}^{(k^*+1-\ell)}$  conditioned on  $\mathbf{h}^{(s)}$ .
  - 8:   Compute empirical REINFORCE gradient for step  $\ell$ :
  - 9:    $\hat{\mathbf{g}}_{\ell} = \frac{1}{n_{\ell}} \sum_{s=1}^{n_{\ell}} \mathbf{R}_{\mathbf{x}^{(s)}}^{\mathcal{F}_{S^*}}(\tau_{\text{hint}}^{(s)}) \nabla_{\mathbf{W}} \log \pi_{\mathbf{W}^{(k^*+1-\ell)}}(i_{\ell}^{(s)} | \mathbf{x}^{(s)}, \hat{\mathbf{p}}^{z_{1:\ell}})$ .
  - 10:   Update parameters:  $\mathbf{W}^{(k^*+2-\ell)} = \mathbf{W}^{(k^*+1-\ell)} + \eta_{\ell}^C \cdot \hat{\mathbf{g}}_{\ell}$ .
  - 11: **end for**
  - 12: Return  $\mathbf{W}^{(k^*+1)}$ .
-

2106 G.3 PROOF OF TEST-TIME SCALING  
2107

2108 **Proposition 1** (Formal Version of the First Item in Thm. 2). *Assume  $\text{TF}_{\text{base}}$  copies PART probability*  
2109 *behavior with a unique correct child per depth. At depth  $\ell$ , consider FXRS (Alg. 4) to force a chosen*  
2110 *visible  $x$ -token and BAI (Alg. 5) that, for each  $j \in \mathcal{I}_\ell$ , repeats FXRS+oracle  $m$  times and selects the*  
2111 *empirical best arm.*

2112 *Then any  $\delta$ -correct identification of the ground-truth path using only the terminal oracle  $\mathbf{R}_x^{f_{S^*}}$  under*  
2113 *token-only observability requires*

$$2114 T_{\text{data}} \geq \tilde{\Omega}(d^{2k^*+1} (1 - \bar{\rho}_{\text{spur}})^{-2}),$$

$$2115 T_{\text{comp}} \geq \tilde{\Omega}(d^{2k^*+1} (1 - \bar{\rho}_{\text{spur}})^{-2}),$$

2116 where  $\bar{\rho}_{\text{spur}} := \sup_{\ell} \rho_{\text{spur}}^{\text{sup}, > \ell}$ . Here  $\rho_{\text{spur}}^{\text{sup}, > \ell} \in [0, 1)$  denotes the terminal-oracle, suffix-level spurious-  
2117 acceptance parameter at depth  $\ell$ ; formally, with  $\mathcal{H}_{\ell-1}$  to denote the event that the history up to depth  
2118  $\ell-1$  is correct,  $C_\ell(j) := \{\text{at depth } \ell, \text{ the visible token is forced to } x_{i_\ell} = x_j\}$  for the FXRS forcing  
2119 event, define the true-suffix event

$$2120 E_{\text{true}} := \bigcap_{t=\ell+1}^{k^*+1} \{J_t = j_t^*\},$$

2121 where  $j_t^*$  is the unique correct child at depth  $t$  given the correct history (with  $j_{k^*+1}^* = d+1$ ). Let  
2122  $E_{\text{spur}}$  be the complement of  $E_{\text{true}}$  while still yielding acceptance by  $\mathbf{R}_x^{f_{S^*}}$  on the final pre-EOS token.  
2123 We then defines

$$2124 \rho_{\text{spur}}^{\text{sup}, > \ell} := \sup_{j \in \mathcal{I}_\ell} \mathbb{E}_x \left[ \Pr_{\text{TF}_{\text{base}}} (E_{\text{spur}} \mid \mathbf{x}, \mathcal{H}_{\ell-1}, C_\ell(j)) \right],$$

2125 which upper-bounds, in expectation over  $\mathbf{x}$ , the probability that a wrong suffix (after depth  $\ell$ ) is  
2126 nonetheless accepted by  $\mathbf{R}_x^{f_{S^*}}$  when the depth- $\ell$  visible token is forced to  $x_{i_\ell} = x_j$ .  
2127

---

2148 **Algorithm 4** Forced X-token Rejection Sampling (FXRS) under token-only observability  
2149

- 2150 1: Inputs: depth  $\ell$ , committed history  $\text{CoT}_{\ell-1}$ , legal set  $\mathcal{I}_\ell$ , candidate  $j \in \mathcal{I}_\ell$ , model  $\text{TF}_{\text{base}}$ ,  
2151 oracle  $\mathbf{R}_x^{f_{S^*}}$ , max trials  $T_{\text{max}} = \Theta(d \log(\frac{d}{\delta}))$ .  
2152 2: **for**  $t = 1$  to  $T_{\text{max}}$  **do**  
2153 3: From context  $\text{CoT}_{\ell-1}$ , sample one step with  $\text{TF}_{\text{base}}$  to emit  $\hat{\mu}^{x_{i_\ell}}$ .  
2154 4: **if**  $\hat{\mu}^{x_{i_\ell}} = \mu^{x_j}$  **then**  
2155 5: Roll out the remaining suffix using  $\text{TF}_{\text{base}}$  until termination (EOS occurs at some depth  
2156 by the copied PART behavior).  
2157 6: Return the full rollout sequence and **stop**.  
2158 7: **end if**  
2159 8: **end for**  
9: If no match after  $T_{\text{max}}$ , return **fail** (increase  $T_{\text{max}}$  if needed).
-

**Algorithm 5** BAI across depths with FXRS and  $\mathbf{R}_x^{fs^*}$ 


---

1: Inputs: legal sets  $\{\mathcal{I}_\ell\}_{\ell=1}^{k^*}$ , model  $\text{TF}_{\text{base}}$ , oracle  $\mathbf{R}_x^{fs^*}$ , confidence  $\delta$ .  
 2: Initialize committed partial CoT  $\mathcal{H} \leftarrow ()$ .  
 3: **for**  $\ell = 1$  to  $k^* + 1$  **do**  
 4:   Set per-arm repetitions  $m \leftarrow \Theta(d^{2(k^*+1-\ell)} \log(d/\delta))$ .  
 5:   **for** each  $j \in \mathcal{I}_\ell$  **do**  
 6:     Initialize counter  $c_j \leftarrow 0$ .  
 7:     **for**  $r = 1$  to  $m$  **do**  
 8:       Run FXRS (Alg. 4) with candidate  $j$  to obtain a full rollout (or retry until success). If FXRS fails, repeat until success.  
 9:       Query  $\mathbf{R}_x^{fs^*}$  on the final pre-EOS token of the rollout; record one Bernoulli outcome  $Y_{j,r} \in \{0, 1\}$ .  
 10:       Update  $c_j \leftarrow c_j + Y_{j,r}$ .  
 11:     **end for**  
 12:     Set empirical acceptance  $\hat{p}_j \leftarrow c_j/m$ .  
 13:   **end for**  
 14:   Let  $\hat{j}_\ell \leftarrow \arg \max_{j \in \mathcal{I}_\ell} \hat{p}_j$ .  
 15:   Commit depth- $\ell$ : resample until  $\hat{\mu}^{x_{i_\ell}} = \mu^{x_{\hat{j}_\ell}}$ , then compute  $\hat{\mu}^{z_\ell} = \text{FFN}_\ell(\hat{\mu}^{x_{i_\ell}}, \mathbf{E}[z_{\ell-1}]; d_x)$ ; update  $\mathcal{H} \leftarrow (\mathcal{H}, \hat{\mu}^{x_{\hat{j}_\ell}}, \hat{\mu}^{z_\ell})$ .  
 16: **end for**  
 17: Return the committed path encoded by  $\mathcal{H}$ .

---

**Proof.** *Proof of Prop. 1.*

**Step 0 (probability spaces and notation).** For any event  $A$ , we write

$$\Pr[A] := \mathbb{E}_{\mathbf{x} \sim \text{Unif}([K]^d)} \left[ \Pr_{\text{TF}_{\text{base}}} (A \mid \mathbf{x}) \right].$$

That is, all probabilities  $\Pr_{\text{TF}_{\text{base}}}(\cdot \mid \mathbf{x})$  are w.r.t. the internal sampling of  $\text{TF}_{\text{base}}$ , conditional on the fixed input  $\mathbf{x}$ . Unconditional probabilities/expectations  $\Pr[\cdot]$  and  $\mathbb{E}_x[\cdot]$  are taken over  $\mathbf{x} \sim \text{Unif}([K]^d)$ . We use  $\mathcal{H}_{\ell-1}$  to denote the event that the history up to depth  $\ell-1$  is correct, and we write

$$C_\ell(j) := \{\text{at depth } \ell, \text{ the visible token is forced to } x_{i_\ell} = x_j\}$$

for the FXRS forcing event.

**Step 1 (conditioning and notation).** Fix  $\mathbf{x}$  and condition on  $\mathcal{H}_{\ell-1}$ . Let  $j_\star$  be the unique correct child at depth  $\ell$ . For  $t \in \{\ell+1, \dots, k^*+1\}$ , let  $J_t$  denote the random child index selected by  $\text{TF}_{\text{base}}$  at depth  $t$  given the preceding context; conditionally on  $\mathcal{H}_{\ell-1}$  and on  $C_\ell(j)$ , we have  $J_t \sim \text{Unif}(\mathcal{I}_t)$  with  $|\mathcal{I}_t| = \Theta(d)$ , independent across  $t$ ; at  $t=k^*+1$ , the correct choice is  $d+1$  (EOS).

**Step 2 (events and acceptance).** Define the true-suffix event

$$E_{\text{true}} := \bigcap_{t=\ell+1}^{k^*+1} \{J_t = j_\star^*\},$$

where  $j_\star^*$  is the unique correct child at depth  $t$  given the correct history (with  $j_{k^*+1}^* = d+1$ ). Let  $E_{\text{spur}}$  be the complement of  $E_{\text{true}}$  while still yielding acceptance by  $\mathbf{R}_x^{fs^*}$  on the final pre-EOS token. Then, for fixed  $j$  at depth  $\ell$ ,

$$P_\ell^{\text{term}}(j) := \mathbb{E}_x \left[ \Pr_{\text{TF}_{\text{base}}} (E_{\text{true}} \mid \mathbf{x}, \mathcal{H}_{\ell-1}, C_\ell(j)) + \Pr_{\text{TF}_{\text{base}}} (E_{\text{spur}} \mid \mathbf{x}, \mathcal{H}_{\ell-1}, C_\ell(j)) \right].$$

**Step 3 (true-suffix probability including EOS).** Under the copied PART uniform branching and uniqueness, for  $j=j_\star$  we must select the unique correct child at each of the remaining depths, including the terminating EOS choice. Hence, for every  $\mathbf{x}$ ,

$$\Pr_{\text{TF}_{\text{base}}} (E_{\text{true}} \mid \mathbf{x}, \mathcal{H}_{\ell-1}, C_\ell(j_\star)) := \prod_{t=\ell+1}^{k^*+1} \frac{1}{|\mathcal{I}_t|} := \Theta(d^{-(k^*+1-\ell)}).$$

2214 Taking  $\mathbb{E}_{\mathbf{x}}[\cdot]$  preserves the same order.

2215 **Step 4 (uniform spurious bound).** By definition,

2216 
$$\mathbb{E}_{\mathbf{x}} \left[ \Pr_{\text{TF}_{\text{base}}} (E_{\text{spur}} \mid \mathbf{x}, \mathcal{H}_{\ell-1}, C_{\ell}(j)) \right] \leq \rho_{\text{spur}}^{\text{sup}, > \ell} < 1$$

2217 for all  $j \in \mathcal{I}_{\ell}$ .

2218 **Step 5 (gap lower bound).** Therefore,

2219 
$$\begin{aligned} P_{\ell}^{\text{term}}(j_{\star}) - P_{\ell}^{\text{term}}(j) &= \mathbb{E}_{\mathbf{x}} \left[ \Pr_{\text{TF}_{\text{base}}} (E_{\text{true}} \mid \mathbf{x}, \mathcal{H}_{\ell-1}, C_{\ell}(j_{\star})) \right] \\ &\quad + \mathbb{E}_{\mathbf{x}} \left[ \Pr_{\text{TF}_{\text{base}}} (E_{\text{spur}} \mid \mathbf{x}, \mathcal{H}_{\ell-1}, C_{\ell}(j_{\star})) - \Pr_{\text{TF}_{\text{base}}} (E_{\text{spur}} \mid \mathbf{x}, \mathcal{H}_{\ell-1}, C_{\ell}(j)) \right] \\ &\geq c d^{-(k^{\star}+1-\ell)} - \rho_{\text{spur}}^{\text{sup}, > \ell}, \end{aligned}$$

2220 for some absolute constant  $c > 0$ . Rearranging constants yields the stated  $\Theta(\cdot) \cdot (1 - \rho_{\text{spur}}^{\text{sup}, > \ell})$  form.

2221 **Step 6 (oracle query complexity per depth).** Let the one-vs-best acceptance-gap at depth  $\ell$  be  $\Delta := P_{\ell}^{\text{term}}(j_{\star}) - \max_{j \neq j_{\star}} P_{\ell}^{\text{term}}(j) = \Theta(d^{-(k^{\star}+1-\ell)}(1 - \rho_{\text{spur}}^{\text{sup}, > \ell}))$ . Viewing each candidate  $j \in \mathcal{I}_{\ell}$  as a Bernoulli arm with mean  $P_{\ell}^{\text{term}}(j)$ , any  $\delta$ -correct identification among  $|\mathcal{I}_{\ell}| = \Theta(d)$  arms requires (in expectation)

2222 
$$\Omega(\Delta^{-2} \log(d/\delta))$$

2223 oracle observations by Lemma 2 and the confidence allocation via union bound (sufficiency for uniform sampling also follows from Lemma 1). Substituting the explicit  $\Delta$  gives  $\Omega(d^{2(k^{\star}+1-\ell)} \log(d/\delta))$ .

2224 **Step 7 (overall  $T_{\text{data}}$ ).** Summing over depths  $\ell = 1, \dots, k^{\star} + 1$  gives  $T_{\text{data}} \geq \tilde{\Omega}(d^{2k^{\star}+1} (1 - \bar{\rho}_{\text{spur}})^{-2})$ .

2225 **Step 8 (overall  $T_{\text{comp}}$ ).** In the best case (minimal token emissions) for FXRS, once the forced visible token matches (success on first try), one can immediately append EOS. This uses a constant number of model emissions (exactly 2). Therefore the model-sampling cost is at minimal a constant multiple of the oracle-query cost at that depth (for parity task, the chance is 1/2, therefore the number to reach is  $\Theta(\log(1/\delta))$  for allowable probability  $1 - o(1/\delta)$ ), and aggregating over depths yields a  $T_{\text{comp}}$  lower bound matching the order of  $T_{\text{data}}$  (up to polylog and spurious factors). Also, the best case for the resampling at Line 15 of Alg. 5 is  $k^{\star} + 1$  token emissions, which is smaller than  $\tilde{\Omega}(d^{2k^{\star}+1} (1 - \bar{\rho}_{\text{spur}})^{-2})$ , yielding the final  $T_{\text{comp}} \geq (d^{2k^{\star}+1} (1 - \bar{\rho}_{\text{spur}})^{-2})$ .  $\square$

2226 **Theorem 12** (Formal Version of the Second Item in Thm. 2 (complexities in  $T_{\text{data}}$  and  $T_{\text{comp}}$ )). Assume  $\text{TF}_{\text{base}}$  copies PART probability behavior with a unique correct child per depth. Consider the layer-wise procedure in Alg. 6 that, at each depth  $\ell$ , uses FXRS (Alg. 4) to force a candidate visible  $x$ -token for every  $j \in \mathcal{I}_{\ell}$ , queries the family oracle  $\mathbf{R}_{\mathbf{x}}^{\mathcal{F}S_{\star}}$  under external truncation, and runs BAI (Alg. 5) to pick the best arm; then commits the chosen  $(x_{i_{\ell}}, z_{\ell})$  and proceeds to the next depth.

2227 Then there exists a choice of per-depth repetitions  $m_{\ell} = \Theta(d^2(1 - \rho_{\text{spur}}^{\text{sup}, \ell})^{-2} \log(d/\delta))$  such that identifying  $S_{\star}$  with probability at least  $1 - \delta$  is achievable with

2228 
$$T_{\text{data}} \leq \tilde{O}((k^{\star} + 1)d^2(1 - \rho_{\text{spur}}^{\text{max}})^{-2}), \quad T_{\text{comp}} \leq \tilde{O}((k^{\star} + 1)d^3(1 - \rho_{\text{spur}}^{\text{max}})^{-2}),$$

2229 where  $\rho_{\text{spur}}^{\text{max}} := \max_{\ell \in [k^{\star}+1]} \rho_{\text{spur}}^{\text{sup}, \ell}$ . Here  $\rho_{\text{spur}}^{\text{sup}, \ell} \in [0, 1)$  denotes the family-oracle, per-depth spurious-acceptance parameter under external truncation at depth  $\ell$ ; formally,

2230 
$$\rho_{\text{spur}}^{\text{sup}, \ell} := \sup_{\substack{j \in \mathcal{I}_{\ell} \\ j \neq j_{\star}}} \mathbb{E}_{\mathbf{x}} \left[ \Pr_{\text{TF}_{\text{base}}} (\mathbf{R}_{\mathbf{x}}^{\mathcal{F}S_{\star}}(\text{accept at depth } \ell) = 1 \mid \mathbf{x}, \mathcal{H}_{\ell-1}, C_{\ell}(j)) \right].$$

2231 Detailed derivations are given in the proof.

2232 **Proof.Step 0 (notation and conditioning).** We reuse  $\mathcal{H}_{\ell-1}$  for the event that the history up to depth  $\ell-1$  is correct, and  $C_{\ell}(j)$  for the inline-forcing event at depth  $\ell$  that sets the visible token to  $x_{i_{\ell}}=x_j$ .

---

**Algorithm 6** Layer-wise Truncated Accept-Reject (LTAR) with  $\mathbf{R}_x^{\mathcal{F}S^*}$  (BAI with inline forcing and truncation)

---

2268  
2269  
2270  
2271 1: Inputs: legal sets  $\{\mathcal{I}_\ell\}_{\ell=1}^{k^*}$ , confidence level  $\delta$ , per-depth budgets  $m_\ell = \Theta(d^2(1 -$   
2272  $\rho_{\text{spur}}^{\text{sup},\ell})^{-2} \log(\frac{d}{\delta}))$ , max trials  $T_{\text{max}} = \Theta(d \log(\frac{d}{\delta}))$ .  
2273 2: Initialize committed partial CoT  $\mathcal{H} \leftarrow ()$ .  
2274 3: **for**  $\ell = 1$  to  $k^* + 1$  **do**  
2275 4:   For each  $j \in \mathcal{I}_\ell$ , estimate acceptance by  $m_\ell$  repetitions with inline forcing:  
2276 5:     Initialize  $c_j \leftarrow 0$ .  
2277 6:   **for**  $r = 1$  to  $m_\ell$  **do**  
2278 7:     **for**  $t = 1$  to  $T_{\text{max}}$  **do**  
2279 8:       From context  $\mathcal{H}$ , sample one step with  $\text{TF}_{\text{base}}$  to emit  $\hat{\mu}^{x_{i_\ell}}$ .  
2280 9:       **if**  $\hat{\mu}^{x_{i_\ell}} = \mu^{x_j}$  **then**  
2281 10:         Compute  $\hat{\mu}^{z_\ell} \leftarrow \text{FFN}_\ell(\hat{\mu}^{x_{i_\ell}}, \mathbf{E}[z_{\ell-1}]_{:d_X})$ .  
2282 11:         Externally append EOS if  $l < k^* + 1$ ; query  $\mathbf{R}_x^{\mathcal{F}S^*}$  at depth  $\ell$ ; record Bernoulli  
2283  $Y_{j,r} \in \{0, 1\}$ .  
2284 12:         Update  $c_j \leftarrow c_j + Y_{j,r}$  and **break** the retry loop.  
2285 13:       **end if**  
2286 14:     **end for**  
2287 15:     If no match within  $T_{\text{max}}$ , optionally increase  $T_{\text{max}}$  and repeat this repetition.  
2288 16:   **end for**  
2289 17:   Set empirical acceptance  $\hat{p}_j \leftarrow c_j/m_\ell$  for all  $j \in \mathcal{I}_\ell$  and pick  $\hat{j}_\ell = \arg \max_j \hat{p}_j$ .  
2290 18:   Commit depth- $\ell$ : resample until  $\hat{\mu}^{x_{i_\ell}} = \mu^{x_{\hat{j}_\ell}}$ , then compute  $\hat{\mu}^{z_\ell} = \text{FFN}_\ell(\hat{\mu}^{x_{i_\ell}}, \mathbf{E}[z_{\ell-1}]_{:d_X})$ ;  
2291   update  $\mathcal{H} \leftarrow (\mathcal{H}, \hat{\mu}^{x_{\hat{j}_\ell}}, \hat{\mu}^{z_\ell})$ .  
2292 19: **end for**  
2293 20: Return the committed path encoded by  $\mathcal{H}$ .

---

2294  
2295 Given  $\mathcal{H}_{\ell-1}$  and  $C_\ell(j)$ , for  $t \geq \ell$  let  $J_t$  denote the random child index sampled by  $\text{TF}_{\text{base}}$  at depth  
2296  $t$ ; under PART-like uniform branching,  $J_t \sim \text{Unif}(\mathcal{I}_t)$  with  $|\mathcal{I}_t| = \Theta(d)$ , independent across  $t$ , and  
2297 the correct EOS index is  $d+1$  at  $t=k^*+1$ .

2298 Define the depth- $\ell$  acceptance event under the family oracle with external truncation as

$$2299 \quad A_\ell(j) := \{\mathbf{R}_x^{\mathcal{F}S^*}(\text{FXRS-forced } j \text{ at depth } \ell, \text{ truncated at } z_\ell) = 1\}.$$

2300 The (marginal) acceptance probability for candidate  $j$  is

$$2301 \quad \alpha_j^{(\ell)} := \mathbb{E}_{\mathbf{x}} \left[ \Pr_{\text{TF}_{\text{base}}} (A_\ell(j) \mid \mathbf{x}, \mathcal{H}_{\ell-1}, C_\ell(j)) \right].$$

2302 By definition of the per-depth spurious parameter (family-oracle, external truncation)

$$2303 \quad \rho_{\text{spur}}^{\text{sup},\ell} := \sup_{\substack{j \in \mathcal{I}_\ell \\ j \neq j_\star}} \mathbb{E}_{\mathbf{x}} \left[ \Pr_{\text{TF}_{\text{base}}} (A_\ell(j) \mid \mathbf{x}, \mathcal{H}_{\ell-1}, C_\ell(j)) \right] < 1,$$

2304 we have  $\alpha_j^{(\ell)} \leq \rho_{\text{spur}}^{\text{sup},\ell}$  for all  $j \neq j_\star$ . For the unique correct child  $j_\star$ , we assume

$$2305 \quad \alpha_{j_\star}^{(\ell)} \geq 1 - \eta_\ell, \quad \eta_\ell \in [0, 1), \quad (67)$$

2306 which captures any oracle or model non-idealities.

2307  
2308 **Step 1 (acceptance gap and BAI sample complexity).** Under PART-like uniform branching, a single  
2309 inline-forcing repetition succeeds in emitting the visible token  $x_{i_\ell} = x_j$  within expected  $\Theta(d)$  retries.  
2310 Conditioned on success (we then deterministically compute  $z_\ell$  via  $\text{FFN}_\ell$  and externally append EOS  
2311 before querying), the depth- $\ell$  acceptance probabilities satisfy

$$2312 \quad \alpha_{j_\star}^{(\ell)} = \Theta(d^{-1}), \quad \alpha_j^{(\ell)} \leq c d^{-1} \rho_{\text{spur}}^{\text{sup},\ell} \quad (j \neq j_\star),$$

2313 for some absolute constant  $c > 0$ , since the correct child is selected with probability  $\Theta(d^{-1})$  and  
2314 wrong children are upper-bounded by the spurious-acceptance parameter. Hence the one-vs-best  
2315 gap obeys

$$2316 \quad \Delta_\ell := \alpha_{j_\star}^{(\ell)} - \max_{j \neq j_\star} \alpha_j^{(\ell)} = \Theta(d^{-1}) \cdot (1 - \rho_{\text{spur}}^{\text{sup},\ell}).$$

2322 *Running BAI at depth  $\ell$  with*

$$2323 \quad m_\ell = \Theta(\Delta_\ell^{-2} \log(d/\delta))$$

2324 oracle trials per arm identifies  $j_\star$  with probability at least  $1 - \delta$  by Chernoff bounds and a union  
2325 bound over the  $d$  arms.

2326 **Step 2 (per-depth reward-oracle query complexity).** Each successful inline-forcing repetition issues  
2327 one query to  $\mathbf{R}_x^{\mathcal{F}S_\star}$ . With  $|\mathcal{I}_\ell| = \Theta(d)$  arms and  $m_\ell$  repetitions per arm,

$$2328 \quad \tilde{O}(m_\ell) = \tilde{O}(\Delta_\ell^{-2}) = \tilde{O}(d^2, (1 - \rho_{\text{spur}}^{\text{sup}, \ell})^{-2}),$$

2329 since  $\Delta_\ell^{-2} = \Theta(d^2 (1 - \rho_{\text{spur}}^{\text{sup}, \ell})^{-2})$ .

2330 **Step 3 (aggregated  $T_{\text{data}}$ ).** Summing over depths  $\ell=1, \dots, k^\star$  and upper-bounding each per-depth  
2331 spurious term by  $\rho_{\text{spur}}^{\text{max}} := \max_{\ell \in [k^\star+1]} \rho_{\text{spur}}^{\text{sup}, \ell}$ ,

$$2332 \quad T_{\text{data}} \leq \tilde{O}\left(\sum_{\ell=1}^{k^\star+1} d^2 (1 - \rho_{\text{spur}}^{\text{sup}, \ell})^{-2}\right) \leq \tilde{O}((k^\star + 1)d^2 (1 - \rho_{\text{spur}}^{\text{max}})^{-2}).$$

2333 **Step 4 (per-depth model-sampling cost).** Under LTAR external truncation, each repetition computes  
2334  $\hat{\mu}^{z_\ell}$  and appends EOS, costing  $O(1)$  emissions per retry; with worst-case  $O(d \log(d/\delta))$  retries, the  
2335 per-repetition cost is  $O(d \log(d/\delta))$ . Hence, per-depth model-sampling cost is

$$2336 \quad \tilde{O}(d \cdot m_\ell) = \tilde{O}(d^3 (1 - \rho_{\text{spur}}^{\text{sup}, \ell})^{-2}).$$

2337 **Step 5 (aggregated  $T_{\text{comp}}$ ).** Summing over depths and upper-bounding by  $\rho_{\text{spur}}^{\text{max}}$ ,

$$2338 \quad T_{\text{comp}} \leq \tilde{O}\left(\sum_{\ell=1}^{k^\star+1} d^3 (1 - \rho_{\text{spur}}^{\text{sup}, \ell})^{-2}\right) \leq \tilde{O}((k^\star + 1)d^3 (1 - \rho_{\text{spur}}^{\text{max}})^{-2}).$$

2339 **Additional commit cost across depths.** After identifying  $\hat{j}_\ell$  at each depth, LTAR performs a commit  
2340 by resampling until  $\hat{\mu}^{x_{i_\ell}} = \mu^{x_{j_\ell}}$  and computing  $\hat{\mu}^{z_\ell}$ , then (optionally) appending EOS. This  
2341 contributes an extra  $\tilde{O}(d)$  emissions per depth in the worst case (Bernstein-type retry bound), for a  
2342 total  $\tilde{O}((k^\star+1)d)$  over all depths. This term is strictly dominated by the  $\tilde{O}((k^\star+1)d^3 (1 - \rho_{\text{spur}}^{\text{max}})^{-2})$   
2343 bound derived above, and hence is absorbed into  $T_{\text{comp}}$ .

2344 **Step 6 (success probability across depths).** Allocate confidence across depths (e.g., replace  $\delta$  by  
2345  $\delta/(k^\star + 1)$  per depth) or absorb this into polylog factors; combining the steps above yields the  
2346 stated bounds for  $T_{\text{data}}$  and  $T_{\text{comp}}$ .  $\square$

2347 **Remark 13** (Spurious-acceptance parameters and their roles). We use two families of spurious-  
2348 acceptance parameters, tied to the oracle being queried and to how the rollout is conditioned:

- 2349 •  $\rho_{\text{spur}}^{\text{sup}, > \ell} \in [0, 1)$  (terminal-oracle, suffix-level): for queries to  $\mathbf{R}_x^{\mathcal{F}S_\star}$ , after forcing a visi-  
2350 ble child  $j$  at depth  $\ell$  via FXRS, it upper-bounds the expected probability (over  $\mathbf{x}$  and the  
2351 internal sampling of  $\text{TF}_{\text{base}}$ ) that a wrong suffix on depths  $\ell+1, \dots, k^\star+1$  is nonetheless  
2352 accepted by the oracle. We also define a uniform bound  $\bar{\rho}_{\text{spur}} := \sup_{\ell} \rho_{\text{spur}}^{\text{sup}, > \ell}$  for aggre-  
2353 gating across depths.
- 2354 •  $\rho_{\text{spur}}^{\text{sup}, \ell} \in [0, 1)$  (family-oracle, per-depth): for queries to  $\mathbf{R}_x^{\mathcal{F}S_\star}$  under external truncation  
2355 at depth  $\ell$  (so the queried token is  $z_\ell$ ), it upper-bounds the expected acceptance probability  
2356 of any wrong child  $j \neq j_\star$  at depth  $\ell$ . We also use  $\rho_{\text{spur}}^{\text{max}} := \max_{\ell \in [k^\star+1]} \rho_{\text{spur}}^{\text{sup}, \ell}$ .

2357 These quantities are not interchangeable:  $\rho_{\text{spur}}^{\text{sup}, > \ell}$  refers to suffix-level acceptance under the termi-  
2358 nal oracle  $\mathbf{R}_x^{\mathcal{F}S_\star}$ , while  $\rho_{\text{spur}}^{\text{sup}, \ell}$  refers to per-depth acceptance under the family oracle  $\mathbf{R}_x^{\mathcal{F}S_\star}$  with  
2359 truncation. Both families depend on task class parameters (e.g., output alphabet size  $K$ ) and en-  
2360 capsulate the severity of reward hacking. Our bounds carry these terms explicitly: lower bounds  
2361 involve  $(1 - \bar{\rho}_{\text{spur}})^{-2}$ ; upper bounds involve  $\sum_{\ell} d (1 - \rho_{\text{spur}}^{\text{sup}, \ell})^{-2}$ .

## H PROOFS OF PARITY PROBLEM

**Theorem 14** (Representation Theorem of Any Parity Function). *For any  $k' \in [d/2]$ , any index set  $S = S^{k'} \subset [d]$ , and the associated parity function  $f_{S^{k'}} \in \mathcal{P}_{d,k'}$  defined in equation 7, there exists a transformer*

$$\text{TF}^{(k')}(\cdot; \mathbf{W}^*)$$

*with parameters  $\mathbf{W}^*$  such that for every  $\mathbf{x} \in \{0, 1\}^d$ , the autoregressive next-token prediction defined in equation 6 produces exactly the subtask decomposition sequence in equation 7, up to the concatenation with positional embeddings.*

*Concretely, let  $S^{k'} = \{i_1, \dots, i_{k'}\}$  and define  $z_m$  recursively as in equation 7. Then*

$$\text{TF}^{(k')}(\mathbf{E}[x_1], \dots, \mathbf{E}[x_d], \mathbf{E}[\text{EOS}]; \mathbf{W}^*) = \mathbf{E}[z_1],$$

*and more generally for  $m = 2, \dots, k'$ ,*

$$\text{TF}^{(k')}(\mathbf{E}[x_1], \dots, \mathbf{E}[x_d], \mathbf{E}[\text{EOS}], \mathbf{E}[z_1], \dots, \mathbf{E}[z_{m-1}]; \mathbf{W}^*) = \mathbf{E}[z_m],$$

*where each  $\mathbf{E}[z_m] = [\boldsymbol{\mu}^{z_m}, \mathbf{p}_{d+m}]^\top$  is the concatenation of the token embedding  $\boldsymbol{\mu}^{z_m}$  and its positional embedding  $\mathbf{p}_{d+m}$ . Finally, at step  $k' + 1$  the model deterministically outputs the EOS embedding:*

$$\text{TF}^{(k')}(\mathbf{E}[x_1], \dots, \mathbf{E}[x_d], \mathbf{E}[\text{EOS}], \mathbf{E}[z_1], \dots, \mathbf{E}[z_{k'}]; \mathbf{W}^*) = \mathbf{E}[\text{EOS}].$$

*Thus,  $\text{TF}^{(k')}(\cdot; \mathbf{W}^*)$  exactly realizes the chain-of-thought subtask decomposition of  $f_{S^{k'}}$  via next-token prediction.*

**Theorem 15** (PART Behavior of Parity Class). *Consider a base model with Uniform Ordered Transition Probability at each node, where "legal" means both non-repeating and maintaining strictly increasing order ( $i_1 < i_2 < i_3 < \dots$ ), and where  $k = o(d) < [d/2]$  satisfying  $d - k = \Theta(d)$ .*

*Under this assumption, at each parent node, the model uniformly selects among all legal children nodes. Specifically:*

- *At the Root node:  $d$  variables  $\{x_1, x_2, \dots, x_d\}$  are available, each with probability  $1/d$*
- *At node  $x_i$  in the tree: only variables  $\{x_{i+1}, x_{i+2}, \dots, x_d\}$  and EOS are legal, each with probability  $1/(d - i + 1) = \Theta(d^{-1})$ .*
- *EOS is always a legal choice at any node, ensuring termination*

*For a given  $d$  and  $k < [d/2]$ , the number of legal CoT sequences of length  $k' + 1$  (including the EOS token) that compute parity functions  $f_{S^{k'}} \in \mathcal{P}_{d,k'}$  is:*

$$|\mathcal{L}_{k'+1}| = \binom{d}{k'} \quad (68)$$

*This counts the number of ways to choose  $k'$  distinct variables from  $d$  available variables while maintaining strictly increasing order, followed by an EOS token.*

*Under the Uniform Ordered Transition Probability assumption, the probability of any specific legal CoT sequence of length  $k' + 1$  is:*

$$P(\text{specific sequence of length } k' + 1) = O\left(\frac{(\log d)^{k'}}{d^{k'+1}}\right) \quad (69)$$

*When  $k' = O\left(\frac{\log d}{\log \log d}\right)$ , it holds that  $P(\text{specific sequence of length } k' + 1) \leq O(d^{-k'})$*

*The probability of generating any CoT sequence of length  $k' + 1$  is:*

$$P(\text{length } k' + 1) = O\left(\binom{d}{k'} \frac{(\log d)^{k'}}{d^{k'+1}}\right) \quad (70)$$

*This probability distribution over CoT lengths exhibits a systematic bias that depends on the relationship between  $k'$  and  $d$ .*

2430 **Proof.** We analyze the structure of legal CoT sequences systematically:  
 2431

2432 **Step 1: Legal Sequence Structure and Length Constraint.** A legal CoT sequence that computes a  
 2433 parity function  $f_{S^{k'}}$  must have length  $k' + 1$  and be structured as:

$$2434 \text{Root} \rightarrow x_{i_1} \rightarrow x_{i_2} \rightarrow \dots \rightarrow x_{i_{k'}} \rightarrow \text{EOS} \quad (71)$$

2435 where  $\{i_1, i_2, \dots, i_{k'}\}$  are distinct indices from  $[d]$  with strictly increasing order ( $i_1 < i_2 < \dots <$   
 2436  $i_{k'}$ ), satisfying our constraints for parity problems. Every legal CoT must terminate with EOS  
 2437 to indicate completion of the parity computation. The constraint  $k < \lfloor d/2 \rfloor$  ensures that at the  
 2438 maximum length  $k + 1$ , the only legal next prediction is EOS, preventing unbounded tree growth.  
 2439

2440 **Step 2: Counting Legal Sequences of Length  $k' + 1$ .** The number of legal sequences of length  $k' + 1$   
 2441 (including EOS) that compute parity functions  $f_{S^{k'}}$  is:

$$2442 |\mathcal{L}_{k'+1}| = \binom{d}{k'} \quad (72)$$

2443  
 2444  
 2445 This counts:

- 2447 •  $\binom{d}{k'}$  ways to choose  $k'$  distinct variables from  $d$  available variables
- 2448 • The ordering constraint  $i_1 < i_2 < \dots < i_{k'}$  eliminates the need for permutations
- 2449 • Each sequence must end with EOS to be valid, making the total length  $k' + 1$

2451 **Step 3: Transition Probability Calculation.** Under the Uniform Ordered Transition Probability  
 2452 assumption, the probability of any specific legal sequence is the product of transition probabilities  
 2453 at each step:  
 2454

2455 For a sequence  $\text{Root} \rightarrow x_{i_1} \rightarrow x_{i_2} \rightarrow \dots \rightarrow x_{i_{k'}} \rightarrow \text{EOS}$  ( $1 \leq i_1 < i_2 < \dots < i_{k'} \leq d$ ):

- 2456 •  $\text{Root} \rightarrow x_{i_1}$ : probability  $1/d$  (choosing from  $d$  variables)
- 2457 •  $x_{i_1} \rightarrow x_{i_2}$ : probability  $d^{-(d-i_1+1)}$  (choosing from remaining  $d - i_1 + 1$  variables)
- 2458 •  $x_{i_2} \rightarrow x_{i_3}$ : probability  $d^{-(d-i_2+1)}$  (choosing from remaining  $d - i_2 + 1$  variables)
- 2459 •  $\vdots$
- 2460 •  $x_{i_{k'-1}} \rightarrow x_{i_{k'}}$ : probability  $d^{-(d-i_{k-1}+1)}$  (choosing from remaining  $d - i_{k'} + 1$  variables)
- 2461 •  $x_{i_{k'}} \rightarrow \text{EOS}$ : probability  $d^{-(d-i_{k'}+1)}$ .

2462  
 2463 Therefore:

$$2464 P(\text{specific sequence}) := \mathbb{E}[F] = \mathbb{E}\left[\frac{1}{d(d-i_1+1) \cdots (d-i_{k'}+1)}\right] \quad (73)$$

2465  
 2466 For  $1 \leq k' \leq d$ , we here serve to show that

$$2467 \mathbb{E}[F] = O\left(\frac{(\log d)^{k'}}{d^{k'+1}}\right).$$

2468  
 2469 By Lemma 9,

$$2470 \mathbb{E}[F] = \frac{e^{k'} \left(1, \frac{1}{2}, \dots, \frac{1}{d}\right)}{d \binom{d}{k'}}.$$

2471 Recall Maclaurin's inequality (monotonicity of symmetric means): for any  $a_1, \dots, a_d \geq 0$ , if

$$2472 S_k(a) := \frac{e_k(a_1, \dots, a_d)}{\binom{d}{k}} \quad (k = 1, \dots, d),$$

2473 then the sequence  $\{S_k(a)^{1/k}\}_{k=1}^d$  is nonincreasing; in particular

$$2474 S_{k'}(a) \leq (S_1(a))^{k'} \quad \text{for all } k' \geq 1.$$

2484 Apply this with  $a_t = 1/t$  ( $t = 1, \dots, d$ ). We get

$$2485 \frac{e_{k'}(1, \frac{1}{2}, \dots, \frac{1}{d})}{\binom{d}{k'}} \leq \left( \frac{e_1(1, \frac{1}{2}, \dots, \frac{1}{d})}{\binom{d}{1}} \right)^{k'} = \left( \frac{H_d}{d} \right)^{k'},$$

2488 where  $H_d = \sum_{t=1}^d \frac{1}{t}$  is the  $d$ -th harmonic number. Hence

$$2491 \mathbb{E}[F] \leq \frac{1}{d} \left( \frac{H_d}{d} \right)^{k'}.$$

2492 Using the standard bound  $H_d \leq 1 + \log d$  for all  $d \geq 1$ , we obtain

$$2495 \mathbb{E}[F] \leq \frac{1}{d} \left( \frac{1 + \log d}{d} \right)^{k'} = O\left( \frac{(\log d)^{k'}}{d^{k'+1}} \right),$$

2498 which proves the claim.

2499 In particular, when  $k' = O\left(\frac{\log d}{\log \log d}\right)$ , the logarithmic factor in the bound above never exceeds a fixed polynomial in  $d$ . Indeed, observe that

$$2503 (\log d)^{k'} \leq (\log d)^{C \cdot \frac{\log d}{\log \log d}} = \exp\left(C \cdot \frac{\log d}{\log \log d} \cdot \log \log d\right) = d^C,$$

2504 for some absolute constant  $C > 0$ . Therefore

$$2506 \mathbb{E}[F] \leq \frac{(\log d)^{k'}}{d^{k'+1}} \leq \frac{d^C}{d^{k'+1}} = O(d^{-k'}),$$

2509 which shows that in the regime  $k' = O\left(\frac{\log d}{\log \log d}\right)$  the expectation decays at least on the order of  $d^{-k'}$ .

2511 **Step 4: Probability of Length  $k' + 1$ .** Since there are  $\binom{d}{k'}$  different sequences of length  $k' + 1$ , and each has the same probability structure (though with different specific values), the total probability of generating any CoT sequence of length  $k' + 1$  is:

$$2515 P(\text{length } k' + 1) = \sum_{\text{all sequences of length } k' + 1} P(\text{specific sequence}) \quad (74)$$

2518 This can be expressed as:

$$2519 P(\text{length } k' + 1) = \Theta\left(\binom{d}{k'} \cdot d^{-k'}\right) \quad (75)$$

2522  $\square$

2523 **Lemma 9** (Expected reciprocal product over a random  $k'$ -subset). Fix integers  $d \geq 1$  and  $0 \leq k' \leq d$ . Choose indices

2525  $1 \leq i_1 < i_2 < \dots < i_{k'} \leq d$

2526 uniformly at random among all  $\binom{d}{k'}$   $k'$ -subsets of  $\{1, \dots, d\}$ , and define

$$2528 F(i_1, \dots, i_{k'}) := \frac{1}{d \prod_{j=1}^{k'} (d - i_j + 1)}.$$

2530 Let the degree- $k'$  elementary symmetric polynomial be

$$2532 e_{k'}(y_1, \dots, y_d) := \sum_{1 \leq t_1 < \dots < t_{k'} \leq d} y_{t_1} \cdots y_{t_{k'}}, \quad (e_0 \equiv 1),$$

2534 and write the (generalized) harmonic numbers

$$2536 H_d := \sum_{t=1}^d \frac{1}{t}, \quad H_d^{(m)} := \sum_{t=1}^d \frac{1}{t^m} \quad (m \geq 2).$$

2537

2538 Then the expectation of  $F$  admits the exact closed form

$$2539 \mathbb{E}[F] = \frac{e_{k'}(1, \frac{1}{2}, \dots, \frac{1}{d})}{d \binom{d}{k'}} \quad (2540)$$

2541 and, for fixed  $k'$  as  $d \rightarrow \infty$ , the asymptotic expansion

$$2542 \mathbb{E}[F] = \frac{1}{d \binom{d}{k'}} \left( \frac{(\log d)^{k'}}{k'!} + \frac{\gamma (\log d)^{k'-1}}{(k'-1)!} + O((\log d)^{k'-2}) \right) = \frac{(\log d)^{k'}}{d^{k'+1}} \left( 1 + O\left(\frac{1}{\log d}\right) \right), \quad (2543)$$

2544 where  $\gamma$  is the Euler–Mascheroni constant. In particular,

$$2545 \mathbb{E}[F] \asymp \frac{(\log d)^{k'}}{d^{k'+1}}. \quad (2546)$$

2547 Moreover, the first few exact instances are

$$2548 k' = 1 : \quad \mathbb{E}[F] = \frac{H_d}{d \binom{d}{1}} = \frac{H_d}{d^2}, \quad (2549)$$

$$2550 k' = 2 : \quad \mathbb{E}[F] = \frac{H_d^2 - H_d^{(2)}}{2d \binom{d}{2}}, \quad (2551)$$

$$2552 k' = 3 : \quad \mathbb{E}[F] = \frac{H_d^3 - 3H_d H_d^{(2)} + 2H_d^{(3)}}{6d \binom{d}{3}}. \quad (2553)$$

2554 **Proof.Step 1: Reparametrization and identification with an elementary symmetric sum.** Define

2555  $t_j := d - i_j + 1$  for  $j = 1, \dots, k'$ . Because  $i_1 < \dots < i_{k'}$ , we have

$$2556 1 \leq t_{k'} < \dots < t_1 \leq d, \quad (2557)$$

2558 so  $\{t_1, \dots, t_{k'}\}$  is exactly a  $k'$ -subset of  $\{1, \dots, d\}$  and, since the  $(i_1, \dots, i_{k'})$  are chosen uni-

2559 formly, the unordered set  $\{t_1, \dots, t_{k'}\}$  is also uniformly distributed over all  $\binom{d}{k'}$   $k'$ -subsets. With

2560 this change of variables,

$$2561 \prod_{j=1}^{k'} (d - i_j + 1) = \prod_{j=1}^{k'} t_j, \quad F(i_1, \dots, i_{k'}) = \frac{1}{d} \cdot \frac{1}{\prod_{j=1}^{k'} t_j}. \quad (2562)$$

2563 Taking expectation over all  $\binom{d}{k'}$  subsets and recalling the definition of  $e_{k'}$ ,

$$2564 \mathbb{E}[F] = \frac{1}{d \binom{d}{k'}} \sum_{1 \leq t_1 < \dots < t_{k'} \leq d} \frac{1}{t_1 \cdots t_{k'}} = \frac{e_{k'}(1, \frac{1}{2}, \dots, \frac{1}{d})}{d \binom{d}{k'}}. \quad (2565)$$

2566 This proves the exact formula.

2567 **Step 2: A generating-function expression for  $e_{k'}(1, \frac{1}{2}, \dots, \frac{1}{d})$ .** Consider

$$2568 G_d(x) := \prod_{t=1}^d \left( 1 + \frac{x}{t} \right) = \sum_{k=0}^d e_k(1, \frac{1}{2}, \dots, \frac{1}{d}) x^k. \quad (2569)$$

2570 Taking logarithms and expanding,

$$2571 \log G_d(x) = \sum_{t=1}^d \log \left( 1 + \frac{x}{t} \right) = \sum_{m \geq 1} \frac{(-1)^{m+1}}{m} x^m \underbrace{\sum_{t=1}^d \frac{1}{t^m}}_{H_d^{(m)}} = \sum_{m \geq 1} \frac{(-1)^{m+1}}{m} H_d^{(m)} x^m. \quad (2572)$$

2573 Hence

$$2574 G_d(x) = \exp \left( \sum_{m \geq 1} \frac{(-1)^{m+1}}{m} H_d^{(m)} x^m \right), \quad (2575)$$

2592 so the coefficient

$$2593 e_{k'}\left(1, \frac{1}{2}, \dots, \frac{1}{d}\right) = [x^{k'}] G_d(x)$$

2594 is a degree- $k'$  polynomial in the variables  $H_d^{(1)}, \dots, H_d^{(k')}$  with leading term

$$2597 \frac{(H_d)^{k'}}{k'!} \quad (H_d = H_d^{(1)}),$$

2599 and without an  $H_d^{k'-1}$  term (this is the standard consequence of Newton's identities / exponential  
2600 formula; concrete examples for  $k' = 2, 3$  are listed in the statement).

2602 **Step 3: Asymptotics of the ingredients and of the ratio.** We use the classical asymptotics

$$2604 H_d = \log d + \gamma + o(1), \quad H_d^{(m)} \rightarrow \zeta(m) \quad (m \geq 2), \quad \binom{d}{k'} = \frac{d^{k'}}{k'!} \left(1 + O\left(\frac{1}{d}\right)\right), \quad (d \rightarrow \infty, k' \text{ fixed}).$$

2606 Plugging these into the polynomial expression for  $e_{k'}$  from Step 2 gives

$$2608 e_{k'}\left(1, \frac{1}{2}, \dots, \frac{1}{d}\right) = \frac{(\log d + \gamma + o(1))^{k'}}{k'!} + O((\log d)^{k'-2}) = \frac{(\log d)^{k'}}{k'!} + \frac{\gamma(\log d)^{k'-1}}{(k'-1)!} + O((\log d)^{k'-2}).$$

2610 Therefore

$$2611 \frac{e_{k'}(1, 1/2, \dots, 1/d)}{\binom{d}{k'}} = \frac{(\log d)^{k'}}{d^{k'}} \left(1 + O\left(\frac{1}{\log d}\right)\right),$$

2613 and multiplying by the prefactor  $1/d$  yields

$$2615 \mathbb{E}[F] = \frac{(\log d)^{k'}}{d^{k'+1}} \left(1 + O\left(\frac{1}{\log d}\right)\right),$$

2617 which is exactly the asserted asymptotic expansion in the lemma.

2619 **Step 4 (Optional sanity check).** By Maclaurin's inequalities,

$$2621 \frac{e_{k'}(1, 1/2, \dots, 1/d)}{\binom{d}{k'}} \leq \left(\frac{e_1}{\binom{d}{1}}\right)^{k'} = \left(\frac{H_d}{d}\right)^{k'} \sim \frac{(\log d)^{k'}}{d^{k'}},$$

2624 so  $\mathbb{E}[F] \leq (\log d)^{k'}/d^{k'+1}$  up to lower-order factors, matching the leading term above.

2626 □

2627 **Proof.** Representation for PART of parity. We fix the first  $d_{\mathcal{X}}$  columns of  $\mathbf{K}, \mathbf{Q}$  to zero so that the  
2628 attention scores are determined by only the positional encodings. This ensures that the trans-  
2629 former focuses on learning which positions contribute to the parity at each step.  $\mathbf{K}, \mathbf{Q}$  are then  
2630 reparametrized by a single matrix  $\mathbf{W} \in \mathbb{R}^{(d_{\mathbb{E}} - d_{\mathcal{X}})^2}$ ; conversely, the value matrix is set to only  
2631 preserve the  $\mathbf{x}$  component, as follows.

$$2632 \mathbf{K}^{\top} \mathbf{Q} = \begin{pmatrix} \mathbf{0}_{d_{\mathcal{X}} \times d_{\mathcal{X}}} & \mathbf{0}_{d_{\mathcal{X}} \times (d_{\mathbb{E}} - d_{\mathcal{X}})} \\ \mathbf{0}_{(d_{\mathbb{E}} - d_{\mathcal{X}}) \times d_{\mathcal{X}}} & \mathbf{W} \end{pmatrix}, \quad \mathbf{V} = \begin{pmatrix} \mathbb{I}_{d_{\mathcal{X}} \times d_{\mathcal{X}}} & \mathbf{0}_{d_{\mathcal{X}} \times (d_{\mathbb{E}} - d_{\mathcal{X}})} \end{pmatrix}.$$

2635 This type of reparametrization is common in the literature to make dynamical analysis tractable  
2636 (Kim & Suzuki, 2024; 2025). Then denote  $x_{m+d+1} = z_m, m \in [k]$ , we have  $\hat{\mathbf{z}}_m =$   
2637  $\sum_{j=1}^{m-1+d} \sigma_j(\mathbf{w}_m) \boldsymbol{\mu}^{x_j}$  where the softmax scores  $\sigma_j(\mathbf{W}_m) = e^{w_{j,m}} / \sum_{\alpha=1}^{m-1+d} e^{w_{\alpha,m}}$ .

2638 **W Matrix for Uniform Ordered Transition Probability.** For  $0 < m \leq k, i_m \neq d+1$

$$2640 \mathbf{p}_{d+1+m} = \mathbf{p}_{i_m}$$

$$2642 \mathbb{E}[x_{d+1+m}] := [\boldsymbol{\mu}^{z_m}, \mathbf{p}_{i_m}]^{\top} \tag{76}$$

$$2643 (\mathbf{p}_{j \leq i_m}) \mathbf{W} (\mathbf{p}_{i_m}) = -\infty$$

$$2644 (\mathbf{p}_{i_m < j \leq d+1}) \mathbf{W} (\mathbf{p}_{i_m}) = c_{i_m}$$

2645 where  $z_m = z_{m-1} \oplus x_{i_m}$ .

2646 For  $\mathbb{E}[x_{d+1}] = [\boldsymbol{\mu}^{\text{EOS}}, \mathbf{p}_{d+1}]^\top$ ,  
 2647  
 2648  $(\mathbf{p}_{j < d+1}) \mathbf{W}(\mathbf{p}_{d+1}) = c_{d+1}$  (77)  
 2649

2650 Then, denote  $x_{m+d+1} = z_m, m \in [k]$ , we have

2651  
 2652  $\boldsymbol{\mu}^{x_{i_m}} = \hat{\boldsymbol{\mu}}^{x_{i_m}} \sim \text{softmax}(\left( \sum_{j < m-1+d} \frac{e^{(\mathbf{p}_j) \mathbf{W}(\mathbf{p}_{i_{m-1}})}}{\sum_{j < m-1+d} e^{(\mathbf{p}_j) \mathbf{W}(\mathbf{p}_{i_{m-1}})}} \boldsymbol{\mu}^{x_j} \right)^\top \mathbf{U} / \beta)$ . (78)  
 2653  
 2654

2655 Then  $\mathbb{E}[z_m] = \mathbb{E}[x_{m+d+1}] = [\phi(\hat{\boldsymbol{\mu}}^{x_{i_m}} + \boldsymbol{\mu}^{z_{m-1}}), \mathbf{p}_{i_m}]^\top$ .  $\square$   
 2656  
 2657  
 2658  
 2659  
 2660  
 2661  
 2662  
 2663  
 2664  
 2665  
 2666  
 2667  
 2668  
 2669  
 2670  
 2671  
 2672  
 2673  
 2674  
 2675  
 2676  
 2677  
 2678  
 2679  
 2680  
 2681  
 2682  
 2683  
 2684  
 2685  
 2686  
 2687  
 2688  
 2689  
 2690  
 2691  
 2692  
 2693  
 2694  
 2695  
 2696  
 2697  
 2698  
 2699