

What Drives the Zero-to-One Jump? Diagnosing and Steering Hard-Prompt Learning in RLVR

Rui Zhu and Yumou Liu and Zhenzhe Zheng* and Fan Wu

Shanghai Jiao Tong University

Shanghai, China

{jameszhu, liuyumou, zhengzhenzhe, wu-fan}@sjtu.edu.cn

Abstract

Reinforcement Learning with Verifiable Rewards (RLVR) has substantially improved mathematical reasoning, yet progress on hard prompts often stalls: a prompt can remain unsolved for many training steps before abruptly transitioning to solved—a discrete “zero-to-one” jump whose underlying dynamics and steerability remain poorly understood. We investigate two questions: *what drives this transition, and can it be steered?* Through event-based analysis measuring the time-to-first-positive (TTFP), we find that negative rollouts preceding successful transitions consistently retain higher early-stage generation entropy, distinguishing near-solvable failures from dead ends. This signal motivates **Negative-Augmented Policy Optimization (NAPO)**, a method integrated with Group Relative Policy Optimization (GRPO) that ranks negatives in unsolved groups via an early-branching driver and assigns graded advantages, injecting informative gradients precisely where standard group-relative updates vanish. Under fixed rollout budgets, NAPO accelerates early solution discovery and improves Pass@256 by 10.0 points on AIME24 and 6.6 points on AIME2025, demonstrating that zero-to-one transitions are diagnosable and steerable through driver-guided negative learning. Our code is available at <https://github.com/jameszhu-pixel-time/Z20>.

1 Introduction

Reinforcement Learning with Verifiable Rewards (RLVR) (Lambert et al., 2025) has rapidly emerged as a dominant paradigm for post-training reasoning models, particularly in mathematical problem solving where correctness can be automatically verified (Yang et al., 2025a; Guo et al., 2025). At each training step, the policy samples multiple rollouts

per prompt and receives binary rewards indicating whether each final answer is correct. When some rollouts in a group succeed, the resulting contrast between positive and negative samples provides a clear gradient that drives steady improvement. However, on sufficiently hard prompts the on-policy pass rate is so low that entire rollout groups contain only incorrect samples for many consecutive steps. In group-relative methods such as GRPO (Shao et al., 2024; Guo et al., 2025), this eliminates reward contrast and yields near-zero advantages on precisely the prompts that most need exploration. As a result, hard prompts often remain stagnant for many steps and then transition abruptly from unsolved to solved—a discrete “zero-to-one” jump rather than smooth refinement. The dynamics that drive this transition, and whether they can be steered, remain largely unexplored.

Answering these questions, however, requires confronting two intertwined obstacles that current RLVR pipelines leave unresolved. The first obstacle is diagnosing zero-to-one events under binary supervision. Standard RLVR diagnostics offer little purchase on what happens inside a hard prompt before it is first solved: pass@k (Chen et al., 2021) aggregates end-of-training capability and cannot localize the moment a previously unsolved prompt becomes solvable, while group-level rewards remain identically zero across long stretches of training and therefore carry no information about which unsolved prompts are close to a breakthrough. What is missing is (i) an event-based view of training that pinpoints when transitions occur and (ii) an internal, rollout-level signal that distinguishes near-solvable failures from true dead ends using only the information already available at sampling time.

The second obstacle is turning heterogeneous negatives into informative gradient on fully-negative groups. A majority of RLVR methods build on grouped objectives such as GRPO and its

*Zhenzhe Zheng is the corresponding author.




Figure 1: Overview of NAPO pipeline

variants (Shao et al., 2024; Yu et al., 2025; Zheng et al., 2025a; Wu et al., 2026), where all-negative groups produce vanishing advantages and provide little learning signal (Ma et al., 2026; Yu et al., 2025). Recent variants restore a gradient through weighted negative learning, virtual rewards, or entropy-aware updates (Zhu et al., 2025; Nan et al., 2025; Le et al., 2026), but they apply roughly uniform pressure across negatives. The open problem is therefore how to assign graded credit among them under only binary rewards, in a way that injects gradient into unsolved groups without destabilizing the rest of training.

We analyze token-level generation entropy during RLVR training and find that negative rollouts from prompts that later undergo a Negative-to-Positive (N→P) transition exhibit front-loaded uncertainty, with early-response entropy patterns that distinguish them from rollouts on prompts that remain negative (N→N). Based on this signal, we propose NAPO (Negative-Augmented Policy Optimization), a GRPO-integrated method that ranks negative samples by front-loaded entropy and injects informative gradients into unsolved groups.

Our contributions are threefold:

- We study zero-to-one learning dynamics in RLVR through event-based metrics, especially N→P transitions and TTFP, to characterize how hard prompts first become solvable.

- We identify a token-entropy pattern associated with N→P transitions and propose early-branching drivers that rank negative rollouts by their latent learning potential.
- We introduce NAPO, a GRPO-integrated algorithm that assigns graded advantages to negative samples according to their driver-based ranking while maintaining stable updates via in-group normalization, improving token-efficient exploration on hard prompts as reflected by better TTFP and higher pass@k.

2 Preliminaries

Related Work. Reinforcement learning (RL) (Sutton and Barto, 1998) has been shown to improve LLM reasoning via both exploration and exploitation (Auer et al., 2002; Cheng et al., 2026). Approaches such as RLHF (Ouyang et al., 2022), Test-Time RL (Zuo et al., 2025), and RLVR (Guo et al., 2025) train policies using verifiable rewards or reward models, often with policy-gradient methods like PPO, GRPO, and REINFORCE (Schulman et al., 2017; Shao et al., 2024; Williams, 1992). These methods can stall on very hard prompts (Ma et al., 2026). Recent work addresses this by re-allocating rollouts toward learnable prompts and by extracting informative signals from negative samples, including token-level entropy to encourage exploration and reduce overfitting (Zhu et al., 2025;

Le et al., 2026; Fang et al., 2026). We further introduce the **negative-to-positive transition** to predict and leverage latent prompt learnability to overcome this challenge.

Group-Relative Policy Optimization for RLVR . Following recent RLVR practice for reasoning models, we sample a group of G rollouts $\{y_i\}_{i=1}^G$ for each prompt and optimize the policy with Group Relative Policy Optimization (GRPO) (Guo et al., 2025; Shao et al., 2024). Let $\pi_{\theta_{\text{old}}}$ be the snapshot policy and define the token-level likelihood ratio $\rho_{i,t}(\theta) = \frac{\pi_{\theta}(a_{i,t}|s_{i,t})}{\pi_{\theta_{\text{old}}}(a_{i,t}|s_{i,t})}$. GRPO then applies the clipped surrogate objective

$$\begin{aligned} \max_{\theta} \mathbb{E}_x \mathbb{E}_{y_i \sim \pi_{\theta_{\text{old}}}} \left[\sum_{i=1}^G \sum_t \ell_{i,t} - \beta \text{KL}(\pi_{\theta} \parallel \pi_{\text{ref}}) \right], \\ \ell_{i,t} = \min \left(\rho_{i,t}(\theta) \hat{A}_{i,t}, \right. \\ \left. \text{clip}(\rho_{i,t}(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_{i,t} \right), \end{aligned} \quad (1)$$

where $\hat{A}_{i,t}$ is the group-relative advantage, π_{ref} is a reference policy, and β controls KL regularization. For brevity, we use π_{θ} and π_{ref} to denote the distributions conditioned on a given prompt x , i.e., $\pi_{\theta}(\cdot | x)$ and $\pi_{\text{ref}}(\cdot | x)$, respectively.

PSR/NSR decomposition. For verifiable-reward post-training, each rollout receives a binary outcome. We view the learning signal through the positive-sample reward (PSR) and negative-sample reward (NSR) decomposition (Zhu et al., 2025). Let $r(x, y) \in \{-1, +1\}$ be the verifiable reward for completion y on prompt x . A simplified RLVR objective is

$$\mathcal{L}_{\text{RLVR}}(\theta) := -\mathbb{E}_{x \sim \mathcal{D}} \mathbb{E}_{y \sim \pi_{\theta}(\cdot | x)} [r(x, y)]. \quad (2)$$

Using $r = \mathbb{I}[r = +1] - \mathbb{I}[r = -1]$, we can write

$$\begin{aligned} \mathcal{L}_{\text{RLVR}}(\theta) = & \underbrace{-\mathbb{E}_{x \sim \mathcal{D}} \mathbb{E}_{y \sim \pi_{\theta}(\cdot | x)} [\mathbb{I}[r(x, y) = +1]]}_{\mathcal{L}_{\text{PSR}}(\theta)} \\ & + \underbrace{\mathbb{E}_{x \sim \mathcal{D}} \mathbb{E}_{y \sim \pi_{\theta}(\cdot | x)} [\mathbb{I}[r(x, y) = -1]]}_{\mathcal{L}_{\text{NSR}}(\theta)}. \end{aligned} \quad (3)$$

PSR reinforces verified rollouts, whereas NSR suppresses failed ones. Both sides can provide useful supervision on reasoning tasks (Zhu et al., 2025).

Token-level uncertainty and rollout-level drivers. To analyze how N→P happens, we analyze the patterns of negative rollouts before N→P

happens. We measure token-level uncertainty with the entropy of the next-token distribution (Wang et al., 2025; Yang et al., 2025b). For each rollout, let $p_t^{(k)}$ denote the top- k truncated next-token distribution at position t after renormalization, with invalid and padding positions masked out. We define the token-level entropy sequence

$$H_t^{(k)} = - \sum_v p_t^{(k)}(v) \log p_t^{(k)}(v), \quad (4)$$

and its exponential form

$$\rho_t^{(k)} = \exp(H_t^{(k)}), \quad (5)$$

which we use as a token-level perplexity / effective branching statistic. Our goal is not to measure absolute early-stage uncertainty in isolation, but to summarize the *relative prominence* of early uncertainty within each rollout.¹

N→P Transition. To localize progress on hard prompts, let $q_t(x)$ indicate whether evaluation at training step t contains at least one positive rollout for prompt x . We define the time-to-first-positive (TTFP) as

$$\text{TTFP}(x) = \min\{t : q_t(x) = 1\},$$

with $\text{TTFP}(x) = \infty$ if no positive appears within the training horizon. We define an N→P transition as

$$q_{t-1}(x) = 0 \wedge q_t(x) = 1.$$

This implies that the model solves a problem on which all previous attempts failed.

3 Signals of Near-Solvable Failures

This section answers the question of what drives the N→P transition.

3.1 Experimental Setup

We examine two axes of variation in this section: (i) *within-batch*, comparing rollouts inside all-negative groups, and (ii) *between-replay*, comparing two training continuations that share a base state but differ in sampling stochasticity.

¹In off-policy settings, rollouts may come from another policy π_{ϕ} , but unless stated otherwise we still compute H_t under the training policy π_{θ} . $k = 20$ unless otherwise stated.

	AMC23		AIME2025		Math-Beyond		Training Set		
Replay	N→P	P→N	N→P	P→N	N→P	P→N	N→P	P→N	TTFP
<i>n1</i>	2/3	0.33	2.67/20	2	12/134	5.33	120	116	1.213
<i>n2</i>	1.67/3	2.67	2/20	2.33	8.33/134	5.33	117.67	102.67	1.225

Table 1: **Transition statistics at $t^* = 1$ for the $n1$ – $n2$ replay pair.** A prompt is labeled as P if at least one rollout is correct under the fixed rollout budget, and as N otherwise. In the validation columns, a/b reports a N→P transitions among b prompts that are negative at the base checkpoint. Results are averaged over three runs, and TTFP is computed over the first three steps. Best value is shown in bold.

Fixed-batch protocol. RLVR on hard prompts exhibits strong temporal correlations across training steps (Li et al., 2025; Zheng et al., 2025b). To expose this structure directly, we train Qwen2.5-Math-7B on a fixed set of 1024 DAPO-17k prompts for 10 consecutive gradient steps, sampling $G = 8$ rollouts per prompt per step. Holding the prompt batch fixed across steps lets us align rollout-internal uncertainty patterns with later N→P events using the metrics of §2. Unless otherwise noted, we use temperature $T = 1.0$, maximum context length 4096, learning rate 10^{-6} , and zero entropy/KL penalties.²

Critical-step replay. Across the 10-step window, we identify the step t^* at which N→P events peak and treat it as the *critical checkpoint*. Starting from this shared state, we replay training with the same base policy, optimizer state, data order, and hyperparameters, varying only rollout stochasticity. This yields two representative continuations, denoted n_1 and n_2 . We do not label either replay as “good” or “bad” at this point.

Outcome measures. (i) *Prompt-level transition label.* Among prompts that are negative at the base checkpoint t^* , we call a prompt **N→P** if at least one later checkpoint within the replay window produces a positive rollout for it, and **N→N** otherwise. Each negative rollout inherits its prompt’s label. At training time the binary reward gives both groups identical labels; the partition into N→P and N→N is only observable *after* the replay window, which is precisely what makes the within-batch driver signal of §3.2.2 non-trivial.

(ii) *Replay-level transition counts.* For each replay we report the number of N→P and P→N transitions under 64 rollouts per prompt, together with the mean TTFP over prompts within the replay window. Smaller first-positive indices indicate earlier

solution discovery.

(iii) *Validation sets.* We evaluate replay checkpoints on AMC23, AIME 2025, and Math-Beyond (Mayilvahanan et al., 2025) under the same 64-rollout budget.

3.2 Findings and Implications

We first define a rollout-time driver that compresses the entropy–position pattern into a single scalar (§3.2.1), then validate it at each level in turn (§3.2.2 and §3.2.3).

3.2.1 An early-branching driver

Why a ratio, not an absolute level. A naive metric would be the average entropy of a rollout, but the contrast we care about is not “more entropy” in absolute terms; it is a difference in the *timing* of uncertainty—how much of a rollout’s total uncertainty is allocated to its opening tokens (Fang et al., 2026).

Definition. Let $T_\alpha = \lceil \alpha T \rceil$ denote the number of tokens in the first α fraction of a response of length T . The *early mean-perplexity ratio* is

$$d_{\text{early-ppl}} = \frac{\frac{1}{T_\alpha} \sum_{t=1}^{T_\alpha} \rho_t^{(k)}}{\frac{1}{T} \sum_{t=1}^T \rho_t^{(k)}}, \quad (6)$$


where $\rho_t^{(k)} = \exp(H_t^{(k)})$ is the exponential of the top- k truncated entropy defined in Eq. (5). Values of $d_{\text{early-ppl}} > 1$ indicate that uncertainty is concentrated early in the response³.

Choice of the early fraction. We introduce the early-fraction parameter α and compare 0.1, 0.2, and 0.5 on the n_1 – n_2 replay pair (Fig. 5)⁴. Among


³Since entropy is measured on a logarithmic scale, $\exp(H_t^{(k)})$ can be interpreted as a top- k truncated local perplexity, i.e., the effective branching factor of plausible next-token continuations. This is related to the branching-factor view in (Yang and Holtzman, 2025).

⁴Details are deferred to §5.

²Full hyperparameters appear in Appendix A and Table 4.



(a) Entropy-position curves for the bad replay n_2 .



(b) Comparison of the $N \rightarrow P$ group between n_1 and n_2 .

Figure 2: **Entropy-position plots.** The main separation is localized in the early part of generation. Prompts that later undergo $N \rightarrow P$ transitions preserve more early-stage uncertainty than prompts that remain $N \rightarrow N$, and this pattern is sharper in the good replay n_1 than in the bad replay n_2 .

these candidates, 0.2 yields the clearest and most stable separation between $N \rightarrow P$ and $N \rightarrow N$ rollouts while remaining robust across both replays, and we use it throughout the remainder of the paper.

3.2.2 Within-batch evidence: negatives are not homogeneous

Entropy-position pattern. Figure 2a plots mean token entropy as a function of position for negative rollouts in replay n_2 . The $N \rightarrow P$ group maintains noticeably higher uncertainty at the beginning of the response, whereas the $N \rightarrow N$ group collapses earlier to lower-entropy continuations. The two groups converge later in the response, confirming that the discriminative signal is localized in the early stage.

Driver separation. Figure 3a shows the distribution of $d_{\text{early-ppl}}$ for the two groups in replay n_1 . Negatives associated with future $N \rightarrow P$ prompts are systematically shifted toward larger driver values than negatives associated with $N \rightarrow N$ prompts, mirroring the qualitative pattern of Fig. 2a under the single scalar of Eq. (6).

Claim A (within-batch). The driver provides a rollout-level ordering signal that separates near-solvable failures from dead ends inside an otherwise indistinguishable negative pool.

3.2.3 Between-replay evidence: the driver tracks aggregate transition tendency

Entropy-position pattern. Figure 2b restricts attention to $N \rightarrow P$ negatives and compares n_1 against n_2 . The good replay n_1 preserves more early-stage uncertainty than the bad replay n_2 , sharpening the front-loaded pattern of Claim A but now across replays rather than within one. This is consistent with


the relevant contrast being a matter of timing rather than magnitude: the same shape of distribution is shifted earlier in n_1 than in n_2 .

Driver distribution. Figure 3b overlays all four (replay $\times$ label) groups. n_1 allocates more probability mass to high-driver rollouts than n_2 does in both groups.⁵


Realized transitions. Table 1 reports the downstream outcome. Starting from an identical base state, the good replay n_1 realizes more $N \rightarrow P$ transitions on the training set and attains a smaller mean first-positive index than n_2 . The ranking of realized transitions matches the ranking of right-tail driver mass, indicating that the batch-level driver shift is not merely a cosmetic difference but is related to *which* replay will subsequently break more prompts out of the all-negative regime. This is consistent with recent RLVR observations that low-success prompts exhibit structured temporal dynamics rather than noise-like fluctuations throughout training (Zheng et al., 2025b).

Claim B (between-replay). The driver also functions as a batch-level diagnostic: batches whose driver distribution has a heavier right-tail mass go on to realize more $N \rightarrow P$ events.

⁵Batch-level summary statistics quantifying this right-tail dominance are reported in Appendix B.



(a) Driver distribution for the good replay $n1$. See Appendix B for statistics’ demonstration



(b) Driver distribution for the bad replay $n2$ and good replay $n1$.

Figure 3: **Within-batch separation under the practical driver $d_{\text{early-ppl}}$.** Rollouts associated with prompts that later undergo $N \rightarrow P$ transitions are systematically right-shifted relative to rollouts from prompts that remain $N \rightarrow N$. The separation is present in both replays, but is cleaner and more pronounced in the good replay $n1$.

Takeaway

Prompts approaching the $N \rightarrow P$ boundary exhibit a distinctive form of front-loaded uncertainty in their negative rollouts. The early-branching driver $d_{\text{early-ppl}}$ captures this pattern at two levels of aggregation: it separates $N \rightarrow P$ from $N \rightarrow N$ negatives *within* a batch (Claim A), and it separates higher-yield from lower-yield batches *across* replays (Claim B).

4 Algorithm

We now turn the empirical findings above into a training design. Our goal is to translate the driver into a practical update rule for RLVR.

4.1 Theoretical Interpretation

We provide a view of Fisher information to theoretically interpret the driver on the negative rollouts.

Fisher information as a measure of policy plasticity. Larger local Fisher means that a failed continuation is easier to reshape under a small parameter update, so the policy remains more *plastic* around that rollout. Let $\mathbf{p}_t = \pi_\theta(\cdot \mid x, y_{<t})$ be the next-token distribution at step t , and let $\mathbf{F}_t$ denote the corresponding token-level Fisher contribution.

Lemma 4.1 (Token-level Fisher information and entropy). *For token distribution $\mathbf{p}_t$, the token-level Fisher trace is*

$$\text{Tr}(\mathbf{F}_t) = 1 - \sum_{k \in V} p_{t,k}^2.$$

Accordingly, Shannon entropy serves as a useful proxy for local Fisher information: higher-entropy

token distributions typically correspond to more policy-plastic states, although the two quantities are not identical in general.

Implication for negative learning. Lemma 4.1 suggests that uniform penalties are wasteful: low-entropy failures are already rigid, so pushing them down mostly reinforces dead ends. By contrast, failures generated in higher-entropy states are more *plastic* and retain nearby alternatives, making negative updates more likely to redirect the policy toward useful exploration. Empirically, the informative plasticity is concentrated early in the rollout, which is why NAPO uses the early-branching driver rather than generic response-level entropy. In short, we want to penalize failures more strongly where the policy is still movable.

Takeaway

The driver indicates that failures arising in more *plastic* regions of the policy. The $N \rightarrow P$ phenomenon suggests that failures with smaller driver values are not random or hopeless: they come from prompts whose trajectories can cross into correctness under nearby policy shifts.

This leads to a design intuition that larger penalties can be assigned to negative trajectories with larger drivers where the policy still retains meaningful uncertainty and plausible alternatives.

4.2 Negative-Augmented Policy Optimization

Motivated by the heterogeneity of negative rollouts, we propose **NAPO**, a GRPO-integrated algorithm that augments standard policy optimization




Figure 4: Cumulative solved prompts on validation sets. NAPO achieves faster early progress, solving more prompts within the first 50 steps, while maintaining comparable performance at later stages.

Algorithm 1 NAPO: Negative-Augmented Policy Optimization

- 1: Sample prompts x ; for each x , sample rollouts $\{y_i\}$ and rewards $r(x, y_i)$;
 - 2: **if** $\exists i : r(x, y_i) = +1$ **then**
 - 3: Compute standard GRPO advantages $A_{\text{GRPO}}(x, y_i)$;
 - 4: Update the policy with GRPO loss.
 - 5: **else**
 - 6: Compute driver scores for all negative rollouts by (6);
 - 7: Re-normalize scores into $[0, 1]$ by (7);
 - 8: Assign rollout-level advantages by (8);
 - 9: Update the policy with GRPO loss.
 - 10: **end if**
-

with *driver-weighted negative learning* on unsolved groups.

Driver-weighted negative advantage. NAPO uses the driver to modulate negative learning strength inside unsolved groups. Given driver scores $d(x, y_i)$ for rollouts in a fully negative group, we first re-normalize them into $[0, 1]$ to avoid exploding gradient while preserving a relevant advantage magnitude:

$$\tilde{d}(x, y_i) = \frac{d(x, y_i) - \min_j d(x, y_j)}{\max_j d(x, y_j) - \min_j d(x, y_j) + \epsilon}. \quad (7)$$

We then assign rollout-level advantages as

$$A(x, y) = \begin{cases} A_{\text{GRPO}}(x, y), & \text{if } p(x) = 1, \\ -\tilde{d}(x, y), & \text{if } p(x) = 0, \end{cases} \quad (8)$$

where $p(x) = \mathbb{I}[\exists i : r(x, y_i) = +1]$ indicates whether the current group is solved.




Figure 5: Number of solved problems across training steps. Setting the early fraction gate to 0.2 leads to the fastest exploration of correct solutions, consistent with our inference-time findings.

Objective: Grouped advantage on solved groups and weighted NSR on unsolved groups. Let $p(x) = \mathbb{I}[\exists i : r(x, y_i) = +1]$ denote whether the current group is solved. Equivalently, the rollout-level advantage in Eq. (8) can be written as

$$A(x, y_i) = p(x) A_{\text{GRPO}}(x, y_i) - (1 - p(x)) \tilde{d}(x, y_i). \quad (9)$$

When $p(x) = 1$, Eq. (9) reduces to the standard GRPO clipped update. When $p(x) = 0$, we have $A(x, y_i) = -\tilde{d}(x, y_i)$, so the same clipped surrogate becomes a weighted NSR loss in which larger-driver negatives receive larger-magnitude penalties while clipping preserves the usual PPO/GRPO trust-region effect.

5 Evaluation

We empirically evaluate the effectiveness of NAPO on hard prompts.

Setup. We evaluate NAPO under a fixed rollout budget and compare against GRPO and negative-sample-based baselines, including NSR’s weighted variant (w-REINFORCE). We also evaluate the impact of early fraction ratio, where NAO02 refers to taking the first 20% of the response in the driver, while NAO05 takes 50% and NAO01 takes 10%. We report (i) Pass@k on validation sets and (ii) TTFP. Experiments are conducted on multiple validation sets (AMC23, AIME2025, AIME24, and Math-Beyond), covering a wide range of difficulty levels.⁶

⁶We additionally report training-batch transition statistics, with extended ablations deferred to the appendix.

Method k	Pass@ k								
	1	2	4	8	16	32	64	128	256
AMC23									
Base	0.2847	0.4512	0.6237	0.7546	0.8400	0.9010	0.9495	0.9838	1.0000
GRPO	0.7352	0.7910	<u>0.8211</u>	<u>0.8443</u>	<u>0.8681</u>	0.8972	0.9306	0.9579	<u>0.9750</u>
W-REINFORCE	0.4514	0.5700	0.6570	0.7245	0.7900	0.8555	0.9086	0.9450	<u>0.9750</u>
NAPO02	0.7141	0.7718	0.8051	0.8328	0.8628	0.8940	0.9222	0.9467	<u>0.9750</u>
NAPO05	<u>0.7191</u>	<u>0.7828</u>	0.8247	0.8556	0.8845	0.9147	<u>0.9438</u>	<u>0.9654</u>	<u>0.9750</u>
AIME2025									
Base	0.0305	0.0561	0.0964	0.1496	0.2074	0.2657	0.3216	0.3684	0.4000
GRPO	0.1346	0.1753	0.2165	<u>0.2563</u>	0.2910	0.3271	0.3740	0.4294	<u>0.4667</u>
W-REINFORCE	0.0254	0.0439	0.0689	0.0970	0.1300	0.1767	0.2425	0.3278	0.4333
NAPO02	<u>0.1220</u>	<u>0.1648</u>	<u>0.2113</u>	0.2631	0.3175	0.3695	0.4195	0.4745	0.5333
NAPO05	0.1138	0.1576	0.2008	0.2475	<u>0.3006</u>	<u>0.3555</u>	<u>0.4130</u>	<u>0.4741</u>	0.5333
AIME24									
Base	0.0857	0.1508	0.2415	0.3425	0.4343	0.5112	0.5819	0.6575	0.7333
GRPO	0.2832	0.3574	0.4242	0.4860	0.5391	0.5869	0.6330	0.6719	0.7000
W-REINFORCE	0.1516	0.2322	0.3105	0.3706	0.4152	0.4500	0.4821	0.5209	0.5667
NAPO02	0.2988	0.3878	0.4648	0.5314	0.5917	0.6504	0.7048	0.7564	0.8000
NAPO05	<u>0.2898</u>	<u>0.3762</u>	<u>0.4490</u>	<u>0.5133</u>	<u>0.5734</u>	<u>0.6301</u>	<u>0.6810</u>	<u>0.7248</u>	<u>0.7667</u>
Math-Beyond									
Base	0.0426	0.0702	0.1051	0.1446	0.1879	0.2357	0.2866	0.3391	0.3978
GRPO	0.1337	0.1574	0.1780	0.1984	0.2205	0.2454	0.2723	0.3014	0.3370
W-REINFORCE	0.0587	0.0866	0.1147	0.1400	0.1639	0.1899	0.2197	0.2564	0.3094
NAPO02	<u>0.1348</u>	<u>0.1612</u>	<u>0.1882</u>	<u>0.2130</u>	<u>0.2334</u>	0.2527	<u>0.2763</u>	<u>0.3105</u>	<u>0.3591</u>
NAPO05	0.1383	0.1680	0.1948	0.2164	0.2337	<u>0.2508</u>	0.2704	0.2926	0.3204

Table 2: Pass@ k comparison across datasets. The row corresponding to our method is highlighted in gray, the best value within each dataset block and column is shown in bold, and the second-best value is underlined. For W-REINFORCE method, we set the positive advantage weight to 0.5 to handle the much harder prompts.

Training protocol. We use the verl framework (Sheng et al., 2025) and train on DAPO-17k (Yu et al., 2025), a challenging dataset for Qwen2.5-Math-7B (Yang et al., 2024), where the initial pass rate is below 0.2. For each prompt, we sample $G = 8$ rollouts per step. Unless otherwise specified, we use temperature $T = 1.0$ during training, a maximum context length of 4096, and train for one epoch.⁷ We report both large-scale validation results and controlled experiments under identical training batches (as in § 3.2).

Results. (1) Improved sample efficiency on hard prompts. As shown in Figure 4, NAPO discovers first-positive rollouts earlier across datasets, solving more prompts during early training under the same rollout budget. This indicates improved sample efficiency, particularly in low-success regimes where standard GRPO suffers from weak learning signals. These gains require

little additional computation: NAPO uses no extra rollout, verifier call, auxiliary model, or policy forward pass. We compute the uncertainty statistic from logits already produced during the old-policy forward pass and retain one scalar per valid response token rather than a vocabulary-sized tensor; the remaining operations are masking, sequence-wise reductions, and within-group min-max normalization. Across common logged steps, the advantage calculation time is 0.935 ± 0.030 s/step for NAPO and 0.525 ± 0.016 s/step for GRPO (mean $\pm$ standard deviation). The resulting increment is only **0.410 s/step**, approximately **0.03%** of the observed end-to-end step time. Thus, the trace shows no material training-time overhead.

(2) Effect of negative scaling. Comparing NAPO and NAPO05 reveals a clear trade-off rather than a uniform winner. The milder NAPO05 variant performs best on AMC23, on the highest- k region of AIME2025, and on low- k Math-Beyond, whereas full NAPO is clearly strongest on AIME24 and in the mid- k regime of AIME2025. This pat-

⁷All other hyperparameters follow Section 3.1, with full details provided in Appendix A and Table 4.

Dataset	$\Delta\text{Pass}@k$ (%)		
	$k = 1$	$k = 8$	$k = 256$
AMC23	+14.40	+7.95	+1.88
AIME2025	+10.48	+15.31	+8.33
AIME24	+20.15	+27.28	+17.50
Math-Beyond	+7.43	+8.73	+0.41

Table 3: Pass@ k gains of NAPO over Uniform(-0.1) across datasets at $T = 1.0$, averaged over four rollout reevaluations with 256 generations per question. Values are percentage points; positive values favor NAPO.

tern suggests that negative updates should be tuned to benchmark difficulty: stronger shaping helps when there is substantial room for zero-to-one improvement, while milder scaling is more robust once exploration is less constrained.

(3) Within-group differentiation versus uniform negative feedback. To isolate the effect of within-group differentiation, we add a Uniform(-0.1) control that assigns the same small negative advantage to every rollout in an all-negative group. All training and evaluation settings are otherwise matched. Table 3 reports the mean paired improvement of NAPO over this control across four independent evaluations.

Among the evaluation results above, the separation is strongest on the AIME benchmarks—for example, +27.28 points at Pass@8 on AIME24—and the mean remains positive at Pass@256 on all four datasets. In a separate single-evaluation control with Uniform(-1), NAPO improves Pass@256 by 17.50, 23.33, 30.00, and 17.13 percentage points on AMC23, AIME2025, AIME24, and Math-Beyond, respectively. Thus, simply activating constant negative feedback does not reproduce NAPO in these runs; these controls isolate the benefit of driver-based within-group differentiation over a constant update. Because the four reevaluations use fixed checkpoints rather than independent training runs, they establish evaluation-level consistency but not training-seed significance.

6 Conclusion

We studied the zero-to-one problem in RLVR, motivated by two obstacles: diagnosing $N \rightarrow P$ transitions under binary supervision, and extracting informative gradients from heterogeneous negatives in fully-unsolved groups. Using event-based metrics (TTFP and $N \rightarrow P$ events), we showed that failures preceding successful transitions re-

tain higher early-stage uncertainty, providing a practical signal to separate near-solvable rollouts from dead ends. Building on this, we proposed NAPO, a GRPO-integrated method that assigns graded negative advantages according to driver scores under a Fisher-information view, thereby injecting structured gradients into unsolved groups while preserving stable updates on solved ones. NAPO accelerates early solution discovery and improves Pass@256 by 10.0 points on AIME24 and 6.6 points on AIME2025 compared with GRPO, confirming that driver-guided negative learning is an effective way to broaden reasoning boundaries in low-success RLVR.

Limitations

Our analysis and method are currently focused on math-oriented RLVR settings with binary verifiers. Although these tasks provide a clean testbed for studying zero-to-one transitions, the same entropy-based driver may behave differently in broader domains, longer-horizon reasoning tasks, or environments with partial-credit, noisy, or subjective feedback. Extending NAPO to such settings may require more robust driver definitions and calibration procedures.

NAPO also relies on rollout-level statistics collected during training. Its benefits may therefore depend on the rollout budget, group size, decoding temperature, and the quality of the base policy. When the model rarely approaches a correct solution, early uncertainty patterns may be insufficient to identify useful negative samples; conversely, when success rates are already high, the additional weighting of negatives may provide limited gains. Future work should evaluate the method under more diverse model scales, verifier designs, and compute budgets.

Acknowledgments

This work was supported in part by China NSF grant No.U2268204, 62322206, 62432007, 62441236, 62132018, 62025204, 62272307, 62372296, U25A6024. The opinions, findings, conclusions, and recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agencies or the government.

References

- Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. 2002. [Finite-time analysis of the multiarmed bandit problem](#). *Machine Learning*, 47(2–3):235–256.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Pondé de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, and 39 others. 2021. [Evaluating large language models trained on code](#). *CoRR*, abs/2107.03374.
- Daixuan Cheng, Shaohan Huang, Xuekai Zhu, Bo Dai, Xin Zhao, Zhenliang Zhang, and Furu Wei. 2026. [Reasoning with exploration: An entropy perspective](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(36):30377–30385.
- Yangyi Fang, Jiaye Lin, Xiaoliang Fu, Cong Qin, Haolin Shi, Chaowen Hu, Lu Pan, Ke Zeng, and Xunliang Cai. 2026. [How to allocate, how to learn? dynamic rollout allocation and advantage modulation for policy optimization](#). *Preprint*, arXiv:2602.19208.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, and 175 others. 2025. [Deepseek-r1 incentivizes reasoning in llms through reinforcement learning](#). *Nature*, 645(8081):633–638.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2024. [Live-codebench: Holistic and contamination free evaluation of large language models for code](#). *Preprint*, arXiv:2403.07974.
- Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James Validad Miranda, Alisa Liu, Nouha Dziri, Xinxin Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Christopher Wilhelm, Luca Soldaini, and 4 others. 2025. [Tulu 3: Pushing frontiers in open language model post-training](#). In *Second Conference on Language Modeling*.
- Thanh-Long V. Le, Myeongho Jeon, Kim Vu, Viet Dac Lai, and Eunho Yang. 2026. [No prompt left behind: Exploiting zero-variance prompts in LLM reinforcement learning via entropy-guided advantage shaping](#). In *The Fourteenth International Conference on Learning Representations*.
- Xuefeng Li, Haoyang Zou, and Pengfei Liu. 2025. [Limr: Less is more for rl scaling](#). *Preprint*, arXiv:2502.11886.
- Shichao Ma, Zhiyuan Ma, Ming Yang, Xiaofan Li, Xing Wu, Jintao Du, Yu Cheng, Weiqiang Wang, Qiliang Liu, Zhengyang Zhou, and Yang Wang. 2026. [Tspo: Breaking the double homogenization dilemma in multi-turn search policy optimization](#). *Preprint*, arXiv:2601.22776.
- Prasanna Mayilvahanan, Ricardo Dominguez-Olmedo, Thaddäus Wiedemer, and Wieland Brendel. 2025. [Math-beyond: A benchmark for rl to expand beyond the base model](#). *Preprint*, arXiv:2510.11653.
- Gongrui Nan, Siye Chen, Jing Huang, Mengyu Lu, Dexun Wang, Chunmei Xie, Weiqi Xiong, Xianzhou Zeng, Qixuan Zhou, Yadong Li, and Xingzhong Xu. 2025. [Ngrpo: Negative-enhanced group relative policy optimization](#). *Preprint*, arXiv:2509.18851.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, Inc.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms](#). *Preprint*, arXiv:1707.06347.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#). *Preprint*, arXiv:2402.03300.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2025. [Hybridflow: A flexible and efficient rlhf framework](#). In *EuroSys*, pages 1279–1297.
- Richard S Sutton and Andrew G Barto. 1998. *Reinforcement learning: An introduction*, volume 1. MIT press.
- Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xiong-Hui Chen, Jianxin Yang, Zhenru Zhang, Yuqiong Liu, An Yang, Andrew Zhao, Yang Yue, Shiji Song, Bowen Yu, Gao Huang, and Junyang Lin. 2025. [Beyond the 80/20 rule: High-entropy minority tokens drive effective reinforcement learning for llm reasoning](#). In *Advances in Neural Information Processing Systems*, volume 38, pages 115452–115486. Curran Associates, Inc.
- Ronald J. Williams. 1992. [Simple statistical gradient-following algorithms for connectionist reinforcement learning](#). *Machine Learning*, 8(3–4):229–256.
- Junkang Wu, Kexin Huang, Jiancan Wu, An Zhang, Xiang Wang, and Xiangnan He. 2026. [Quantile advantage estimation: Stabilizing RLVR for LLM reasoning](#). In *The Fourteenth International Conference on Learning Representations*.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025a. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.

An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, Keming Lu, Mingfeng Xue, Runji Lin, Tianyu Liu, Xingzhang Ren, and Zhenru Zhang. 2024. [Qwen2.5-math technical report: Toward mathematical expert model via self-improvement](#). *CoRR*, abs/2409.12122.

Chenghao Yang, Lin Gui, Chenxiao Yang, Victor Veitch, Lizhu Zhang, and Zhuokai Zhao. 2025b. [Let it calm: Exploratory annealed decoding for verifiable reinforcement learning](#). *Preprint*, arXiv:2510.05251.

Chenghao Yang and Ari Holtzman. 2025. [Llm probability concentration: How alignment shrinks the generative horizon](#). *Preprint*, arXiv:2506.17871.

Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gao-hong Liu, juncai liu, LingJun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, and 17 others. 2025. [Dapo: An open-source llm reinforcement learning system at scale](#). In *Advances in Neural Information Processing Systems*, volume 38, pages 113222–113244. Curran Associates, Inc.

Chujie Zheng, Shixuan Liu, Mingze Li, Xiong-Hui Chen, Bowen Yu, Chang Gao, Kai Dang, Yuqiong Liu, Rui Men, An Yang, Jingren Zhou, and Junyang Lin. 2025a. [Group sequence policy optimization](#). *Preprint*, arXiv:2507.18071.

Haizhong Zheng, Yang Zhou, Brian Bartoldson, Bhavya Kailkhura, Fan Lai, Jiawei Zhao, and Beidi Chen. 2025b. [Act only when it pays: Efficient reinforcement learning for llm reasoning via selective roll-outs](#). In *Advances in Neural Information Processing Systems*, volume 38, pages 124321–124346. Curran Associates, Inc.

Xinyu Zhu, Mengzhou Xia, Zhepei Wei, Wei-Lin Chen, Danqi Chen, and Yu Meng. 2025. [The surprising effectiveness of negative reinforcement in llm reasoning](#). In *Advances in Neural Information Processing Systems*, volume 38, pages 126546–126573. Curran Associates, Inc.

Yuxin Zuo, Kaiyan Zhang, Li Sheng, Shang Qu, Ganqu Cui, Xuekai Zhu, Haozhan Li, yuchen zhang, Xinwei Long, Ermo Hua, Bqing Qi, Youbang Sun, Zhiyuan Ma, Lifan Yuan, Ning Ding, and Bowen Zhou. 2025. [Ttrl: Test-time reinforcement learning](#). In *Advances in Neural Information Processing Systems*, volume 38, pages 131459–131483. Curran Associates, Inc.

A Experiment Settings

A.1 Chat template

Chat template for Qwen

System: You are a helpful assistant.

User: {input} Please reason step by step, and put your final answer within `\boxed{}`.

Assistant: Model generates a step-by-step solution; the final answer is enclosed in `\boxed{}`.

A.2 Hyperparameter Settings

Category	Setting
<i>Model</i>	
Model	Qwen2.5-Math-7B
<i>Data</i>	
Train batch size	1024
Max prompt length	1024
Max response length	3072
<i>RL Algorithm</i>	
KL coefficient	0.0
Entropy coefficient	0.0
Critic warmup	0
<i>Optimization</i>	
Learning rate	1×10^{-6}
PPO minibatch size	256
<i>Sampling</i>	
top_k	-1
top_p	1
temperature	1.0

Table 4: Detailed training hyperparameters.

All experiments are conducted on a single node equipped with eight NVIDIA H100 GPUs. We report final evaluation results using the checkpoint at step 200 for Qwen2.5-Math-7B and the checkpoint at step 100 for Llama-3.1-8B-Instruct.

B Definitions of Batch-Level Driver Statistics

In this appendix, we define the batch-level summary statistics used in Table 5 and Figure 3 to characterize the distribution of rollout-time drivers within the $N \rightarrow P$ group. For Figure 3 We use kernel density estimation (KDE) only for visualization, to provide a smoothed view of the empirical distribution of the driver scores. KDE is a standard non-parametric density estimation method that places

a smooth kernel around each observed sample and aggregates these kernels to estimate a continuous density. It does not affect training, evaluation, or the computation of the proposed driver; all quantitative analyses are performed on the original empirical scores.

Let $\{d_i\}_{i=1}^n$ denote the driver values computed for the rollouts in a given subset (e.g., the N→P rollout group of a batch), where n is the number of rollouts in that subset.

Mean-shift statistic (msk). The statistic msk summarizes the average magnitude of the driver values within the selected rollout group:

$$\text{msk} = \frac{1}{n} \sum_{i=1}^n d_i. \quad (10)$$

In our analysis, larger msk indicates that the corresponding batch tends to allocate relatively more uncertainty to the early stage of generation.

Right-tail mass statistic (rt). The statistic rt summarizes the heaviness of the right tail of the driver distribution. Let τ be a fixed upper-tail threshold (we set $\tau = 1.95$ in practice). We define

$$\text{rt} = \frac{1}{n} \sum_{i=1}^n \mathbf{1}[d_i > \tau], \quad (11)$$

where $\mathbf{1}[\cdot]$ is the indicator function. Larger rt indicates that a larger fraction of rollouts fall into the high-driver region, i.e., the distribution has a heavier right tail.

Interpretation. These statistics are not intended to serve as calibrated predictors of future success or as exact estimates of future reward. In the main text, we use them only to support a weaker but more practically relevant claim: larger driver values, especially a heavier right tail, are enriched in rollout groups associated with stronger N→P tendency. This is the only property required by our method design, where the driver is used as a surrogate ordering signal rather than as a direct estimator of future pass rate.

C Additional results

C.1 Generalization beyond Qwen2.5-Math-7B

Transition-level evidence. Starting from a shared Llama-3.1-8B-Instruct checkpoint, we run five stochastic GRPO continuations. Across the 696 prompts that are negative at the shared checkpoint, the mean N→P rate is $(14.86 \pm 0.44)\%$

Replay	Training N→P	msk	rt
<i>n1</i>	120	0.924	0.039
<i>n2</i>	117.67	0.812	0.026

Table 5: Representative batch-level driver statistics for a good(*n1*)/bad(*n2*) replay comparison. Larger msk and rt indicate more mass in the high-driver region.

(mean $\pm$ standard error across continuations). Thus, delayed first-solution discovery is not confined to Qwen2.5-Math-7B.

Training-level evidence. We apply NAPO using the same early-fraction window $\alpha = 0.2$ selected on Qwen2.5-Math-7B (Figure 5), without Llama-specific tuning. Table 6 compares the resulting step-100 checkpoint against GRPO.

The results are mixed, where NAPO performs best on Math-Beyond and high-budget AIME24, while degrading on AIME2025. Thus, the evaluation indicates that the driver window and resulting advantage magnitudes may require calibration across model families, while remaining robust to a certain degree.

C.2 Beyond mathematics: binary-verifiable code

Setup. We examine whether the transition-level entropy signal extends beyond mathematical reasoning. We train Qwen3-4B-Instruct (Yang et al., 2025a) on 512 distinct coding problems from LiveCodeBench v5 (Jain et al., 2024) using execution-based binary rewards, and evaluate on 100 held-out problems from LiveCodeBench v6. Five consecutive GRPO steps each visit all 512 training prompts and sample $G = 8$ rollouts per prompt, yielding four adjacent transitions for the analysis below. All other settings follow Section 3.1.

Code-domain entropy statistic. On the fixed sensitivity grid $\alpha \in \{0.1, 0.2, 0.5\}$ evaluated on the same transitions, $\alpha = 0.5$ gives the clearest observed separation, consistent with the different output structure and prompt formatting of coding tasks. For this experiment, msk denotes the prompt-level mean early-to-global plain-entropy ratio across the eight source-step rollouts. Unlike the NAPO training driver $d_{\text{early-ppl}}$, this analysis-only statistic uses entropy directly rather than its exponential.

Pooling the four transitions gives 83 N→P and 550 N→N prompt–transition observations, whose mean msk values are 1.0985 and 1.0709, respec-




Figure 6: TTFP distribution




Figure 7: Cumulative solved number

Method (Llama)	Pass@k (%)			
k	1	2	8	256
AMC23				
NAPO02	40.50	47.68	57.82	80.83
GRPO	37.55	47.29	59.89	80.00
AIME2025				
NAPO02	0.06	0.11	0.45	13.33
GRPO	0.17	0.33	1.25	20.00
AIME24				
NAPO02	4.75	6.67	11.96	37.78
GRPO	7.11	10.41	16.15	33.33
Math-Beyond				
NAPO02	5.90	7.32	9.68	17.50
GRPO	4.93	6.39	8.92	16.21

Table 6: Pass@k comparison across datasets for **Llama-3.1-8B-Instruct** at temperature $T = 1.0$. Values are percentages. The row corresponding to NAPO02 is highlighted in gray, and the best value between GRPO and NAPO02 within each dataset block and column is shown in bold. Due to limited computational resources, we evaluate the checkpoint at step 100.

Transition	N→P / N→N	msk N→P	msk N→N	Δ msk
1 → 2	20/133	1.1044	1.0740	+0.0304
2 → 3	20/135	1.0997	1.0710	+0.0287
3 → 4	18/144	1.0968	1.0679	+0.0289
4 → 5	25/138	1.0941	1.0710	+0.0231

Table 7: Transition-level entropy statistics for **Qwen3-4B-Instruct** on LiveCodeBench at $\alpha = 0.5$. The second column reports the N→P and N→N prompt counts, and the remaining columns report their mean msk values and difference; positive differences indicate higher early relative entropy for N→P.

tively. The difference is +0.0276, with a 95% prompt-clustered bootstrap confidence interval of [0.0066, 0.0484]. The step-stratified AUC is 0.608, with a 95% confidence interval of [0.532, 0.681]. The difference has the same sign at all four transitions rather than being driven by a single favorable step. These results provide preliminary evidence for an entropy-exploration signal in the coding tasks.

D Theoretical Derivation: Fisher Information and Negative-Augmented Policy Optimization

In this section, we provide a theoretical justification for the proposed driver-based weighting scheme. We demonstrate that penalizing negative samples with uniform advantage (standard GRPO) wastes optimization budget on samples that provide negligible Fisher Information. We then derive Negative-Augmented Policy Optimization (NAPO), where the advantage is proportional to the Fisher Informa-

tion of the trajectory, effectively maximizing the information gain of the policy update.

D.1 Fisher Information of Negative Samples

Consider a prompt x with a rollout group $\mathcal{G}^- = \{y_i\}_{i=1}^G$ where all rollouts are negative, i.e., $r(x, y_i) = -1$. In policy optimization, the Fisher Information Matrix (FIM) $\mathbf{F}(\theta)$ measures the curvature of the KL divergence and the informativeness of a sample gradient. For a specific trajectory y , the contribution to the FIM is given by the outer product of the score function:

$$\mathbf{F}_y(\theta) = \nabla_{\theta} \log \pi_{\theta}(y|x) \nabla_{\theta} \log \pi_{\theta}(y|x)^{\top}. \quad (32)$$

A sample y is considered informative for optimization if its gradient vector has a large magnitude in the parameter space, indicating a significant potential to reshape the policy distribution.

We analyze the Fisher Information at the token level. Let $\mathbf{p}_t = \pi_{\theta}(\cdot|x, y_{<t})$ be the probability vector over the vocabulary V at step t . Let $\mathbf{F}_t$

denote the Fisher Information contribution from position t .

Lemma D.1 (Fisher Information and Entropy). *Let $H_t = -\sum_k p_k \log p_k$ be the entropy of the distribution at step t . The trace of the token-level Fisher Information matrix satisfies:*

$$\text{Tr}(\mathbf{F}_t) = 1 - \sum_{k \in V} p_k^2. \quad (33)$$

Furthermore, for a fixed vocabulary size V , the quantity $1 - \sum p_k^2$ and the entropy H_t are positively associated: both are maximized by the uniform distribution (at $\frac{V-1}{V}$ and $\log V$, respectively) and both equal zero only for deterministic distributions. In particular, $H_t \rightarrow 0$ implies $\text{Tr}(\mathbf{F}_t) \rightarrow 0$.

Proof. The score function for a sampled token y_t is $\nabla_\theta \log \pi_\theta(y_t|x, y_{<t})$. Under the common softmax parameterization $p_k = \frac{e^{z_k}}{\sum e^{z_k}}$, the gradient of the log-probability with respect to the logits $\mathbf{z}$ is:

$$\nabla_{\mathbf{z}} \log p(y_t) = \mathbf{e}_{y_t} - \mathbf{p}_t, \quad (34)$$

where $\mathbf{e}_{y_t}$ is the one-hot vector for the sampled token. The Fisher Information matrix is the covariance of this gradient:

$$\mathbf{F}_t = \mathbb{E}_{y_t} [(\mathbf{e}_{y_t} - \mathbf{p}_t)(\mathbf{e}_{y_t} - \mathbf{p}_t)^\top] \quad (35)$$

$$= \mathbb{E}[\mathbf{e}_{y_t} \mathbf{e}_{y_t}^\top] - \mathbf{p}_t \mathbf{p}_t^\top. \quad (36)$$

Since $\mathbb{E}[\mathbf{e}_{y_t} \mathbf{e}_{y_t}^\top] = \text{diag}(\mathbf{p}_t)$, the trace of the FIM is:

$$\text{Tr}(\mathbf{F}_t) = \text{Tr}(\text{diag}(\mathbf{p}_t)) - \text{Tr}(\mathbf{p}_t \mathbf{p}_t^\top) \quad (37)$$

$$= \sum_k p_k - \sum_k p_k^2 = 1 - \sum_k p_k^2. \quad (38)$$

The quantity $G := 1 - \sum p_k^2$ is the Gini impurity (equivalently, the Tsallis entropy of order 2). Both G and the Shannon entropy $H = -\sum p_k \log p_k$ are *Schur-concave* functions of the probability vector: they are both uniquely maximized by the uniform distribution and both equal zero if and only if the distribution is deterministic. While G and H are not in general strict monotonic functions of each other for $V \geq 3$, they are tightly coupled via the Rényi entropy ordering $H_1 \geq H_2$, which gives:

$$1 - \sum_k p_k^2 \leq 1 - e^{-H}.$$

In particular, $H \rightarrow 0$ forces $\text{Tr}(\mathbf{F}_t) \rightarrow 0$, and conversely the uniform distribution uniquely maximizes both quantities. Hence, higher-entropy distributions correspond, in a strong qualitative sense, to larger local Fisher information. $\square$

D.2 Analysis of Standard NSR on Unsolved Groups

Standard NSR assigns uniform negative advantages $A(y) \approx -1$ to all samples in an unsolved group. From the perspective of Fisher Information, this implies that the optimization update treats samples with zero information (confident errors) and high information (uncertain errors) equally.

Proposition D.2 (Gradient Saturation and Fisher Information). *Consider a negative sample y generated from a low-entropy (confident) policy state.*

1. **Zero Information:** *If $H_t \rightarrow 0$, then $\text{Tr}(\mathbf{F}_t) \rightarrow 0$. The gradient $\nabla_\theta \log \pi_\theta(y_t)$ vanishes.*
2. **Saddle Points:** *Negative samples generated with high confidence correspond to “saddle points” or “dead ends” in the likelihood landscape. Uniform penalization ($A = -1$) attempts to push probability mass away from these points, but the vanishing Fisher Information implies that $\Delta\theta \approx 0$ regardless of the advantage magnitude.*

Thus, uniform advantage assignment wastes gradient budget on samples that lie in low-information regions of the parameter space.

Proof. The gradient vector for the sampled token y_t is $\mathbf{g}_t = \mathbf{e}_{y_t} - \mathbf{p}_t$, where $\mathbf{e}_{y_t}$ is the one-hot encoding of the sampled token and $\mathbf{p}_t$ is the probability vector $\pi_\theta(\cdot|x, y_{<t})$. The squared norm of this gradient is:

$$\|\mathbf{g}_t\|_2^2 = \|\mathbf{e}_{y_t} - \mathbf{p}_t\|_2^2 \quad (39)$$

$$= (1 - \pi(y_t))^2 + \sum_{v \neq y_t} \pi(v)^2. \quad (40)$$

As $\pi(y_t) \rightarrow 1$ (confidence approaches certainty, entropy $H_t \rightarrow 0$), the first term approaches 0. Simultaneously, the mass of the probability vector concentrates on y_t , forcing $\pi(v) \rightarrow 0$ for all other v . Thus, $\|\mathbf{g}_t\|_2 \rightarrow 0$.

This implies that confidently generated tokens (which characterize the “dead-end” negative roll-outs in Section 3) contribute negligible gradient signal. Uniformly weighting such samples wastes optimization budget on saturated gradients that cannot effectively redirect the policy. $\square$

D.3 Derivation of NAPO

To maximize the efficiency of the update, we should allocate the “negative learning budget” to the samples whose penalization best reduces the total likelihood of failed trajectories. We formalize this as a constrained optimization over per-rollout weights and show that the solution naturally connects to Fisher Information and, through it, to trajectory entropy.

Exact formulation via the Gram matrix. Let $\mathcal{G}^- = \{y_i\}_{i=1}^G$ be a fully negative rollout group for prompt x . Write $g_y = \nabla_{\theta} \log \pi_{\theta}(y \mid x)$ for the trajectory-level score vector, and define the *Gram matrix* $\mathbf{G} \in \mathbb{R}^{G \times G}$ with entries

$$G_{ij} = g_{y_i}^{\top} g_{y_j}. \quad (43)$$

Given a weight vector $w = (w_1, \dots, w_G)^{\top}$, the parameter update is $\delta = -\eta \sum_i w_i g_{y_i}$. The total first-order reduction in the log-likelihood of every negative rollout in the group is

$$R(w) = -\sum_{j=1}^G g_{y_j}^{\top} \delta = \eta \mathbf{1}^{\top} \mathbf{G} w. \quad (44)$$

We seek the weight vector that maximizes $R(w)$ subject to a variance budget $\|w\|^2 \leq B$ (which controls the variance of the stochastic gradient and ensures stable updates):

$$\max_w \mathbf{1}^{\top} \mathbf{G} w \quad \text{s.t.} \quad \|w\|^2 \leq B. \quad (45)$$

By the Cauchy–Schwarz inequality, $\mathbf{1}^{\top} \mathbf{G} w \leq \|\mathbf{G}^{\top} \mathbf{1}\| \|w\|$, with equality when $w \propto \mathbf{G}^{\top} \mathbf{1}$. Since $\mathbf{G}$ is symmetric, the exact solution is

$$w^* \propto \mathbf{G} \mathbf{1}, \quad (46)$$

$$w^*(y_i) \propto [\mathbf{G} \mathbf{1}]_i = \sum_{j=1}^G g_{y_i}^{\top} g_{y_j}.$$

Computing $\mathbf{G} \mathbf{1}$ exactly is impractical (it requires full parameter-space inner products between all pairs of rollout gradients in a multibillion-parameter model). We therefore seek a *tractable diagonal proxy* and state a structural condition under which the ranking induced by $\mathbf{G} \mathbf{1}$ coincides with the ranking induced by the diagonal alone.

Uniform-correlation condition. We decompose each entry of $\mathbf{G} \mathbf{1}$:

$$[\mathbf{G} \mathbf{1}]_i = \underbrace{\|g_{y_i}\|^2}_{=\text{Tr}(\mathbf{F}_{y_i})} + \sum_{j \neq i} g_{y_i}^{\top} g_{y_j}.$$

The diagonal term is exactly the trajectory-level Fisher Information; the off-diagonal sum captures cross-sample interference. Directly computing the latter is as expensive as the full Gram matrix. We therefore introduce the following structural assumption.

Assumption D.3 (Uniform pairwise correlation). There exists a constant $c \geq 0$ such that, for all distinct rollouts $y_i, y_j \in \mathcal{G}^-$,

$$g_{y_i}^{\top} g_{y_j} \approx c \|g_{y_i}\| \|g_{y_j}\|.$$

It states that *all same-prompt rollout gradients are roughly equally correlated with each other*—a plausible condition because they share the same prompt, model architecture, and sampling distribution, so no single pair of rollouts is systematically more aligned than any other. The discriminating factor across rollouts is then their *gradient magnitude*, not their mutual orientation.

Theorem D.4 (Optimality of Fisher-Weighted Advantage). *Under Assumption D.3, the optimal weight $w^*(y)$ from Eq. (46) is a monotonically increasing function of $\text{Tr}(\mathbf{F}_y) = \|g_y\|^2$. By Lemma D.1 and the token-independence approximation, this quantity is positively associated with the trajectory entropy $\sum_t H_t$. Therefore, we define the Entropy-Weighted Advantage:*

$$A_{\text{NAPO}}(y) = -\frac{d(y)}{\frac{1}{|\mathcal{G}|} \sum_{y' \in \mathcal{G}} d(y')}, \quad (47)$$

where $d(y)$ is the driver score (normalized entropy) defined in Eq. (6) of the main text.

Proof. Step 1: Exact solution (Gram matrix). The constrained optimum of Eq. (45) is $w^* \propto \mathbf{G} \mathbf{1}$ as derived in Eq. (46).

Step 2: Reduction under uniform correlation. Let $S = \sum_{j=1}^G \|g_{y_j}\|$ denote the total gradient norm across the group. Under Assumption D.3:

$$\begin{aligned} [\mathbf{G} \mathbf{1}]_i &= \|g_{y_i}\|^2 + c \|g_{y_i}\| \sum_{j \neq i} \|g_{y_j}\| \\ &= \|g_{y_i}\|^2 + c \|g_{y_i}\| (S - \|g_{y_i}\|) \\ &= (1 - c) \|g_{y_i}\|^2 + c S \|g_{y_i}\|. \end{aligned} \quad (48)$$

Define $f(u) = (1 - c) u^2 + c S u$ for $u \geq 0$. Its derivative is $f'(u) = 2(1 - c) u + c S$. Since $0 \leq c \leq 1$ (cosine similarity is bounded by one) and $S \geq 0$, we have $f'(u) \geq 0$ for all $u \geq 0$. Hence f is monotonically increasing in $u = \|g_{y_i}\|$, and therefore also in $\|g_{y_i}\|^2 = \text{Tr}(\mathbf{F}_{y_i})$.

The ranking of $w^*(y)$ across the group thus coincides with the ranking of $\text{Tr}(\mathbf{F}_y)$.

Step 3: Connection to trajectory entropy. Since $g_y = \sum_{t=1}^T \nabla_{\theta} \log \pi_{\theta}(y_t|x, y_{<t})$, under the approximation that token-level score vectors are approximately uncorrelated across positions we have:

$$\|g_y\|^2 \approx \sum_{t=1}^T \|\nabla_{\theta} \log \pi_{\theta}(y_t|x, y_{<t})\|^2.$$

Taking the expectation over the sampled token at each position and applying Lemma D.1:

$$\mathbb{E}[\|g_y\|^2] \approx \sum_{t=1}^T \text{Tr}(\mathbf{F}_t) = \sum_{t=1}^T \left(1 - \sum_k p_{t,k}^2\right),$$

which is positively associated with $\sum_t H_t$ by Lemma D.1.

The driver $d(y)$ in Eq. (6) of the main text is a normalized measure of early-stage entropy and thus serves as a tractable proxy for $\text{Tr}(\mathbf{F}_y)$. The advantage assignment A_{NAPPO} therefore implements an approximation to the optimal Fisher-weighted update, prioritizing samples where the policy is “plastic” (high entropy) over samples where it is rigid (low entropy). $\square$

Remark D.5 (Interpretation as a diagonal preconditioner). The uniform-correlation reduction plays the role of a *diagonal preconditioner* for the Gram matrix. Just as diagonal preconditioning of a full Newton step retains the correct coordinate-wise scaling while ignoring off-diagonal curvature, ranking by $\text{Tr}(\mathbf{F}_y)$ retains the correct per-rollout scaling while treating cross-rollout interactions as uniform. When the uniform-correlation assumption is violated (e.g., if two rollouts share unusually long common prefixes), the proxy may over- or under-weight individual rollouts, but the overall ranking remains a principled first-order approximation.

Remark D.6 (Relation to orthogonality). Setting $c = 0$ recovers the special case of orthogonal score vectors, for which $w^*(y) \propto \|g_y\|^2$ exactly. However, same-prompt rollouts typically share early tokens and pass through shared model parameters, making $c > 0$ the realistic regime. The uniform-correlation condition accommodates this shared structure while still yielding a tractable ranking criterion.

Remark D.7 (Practical implementation). In practice, we approximate the theoretically optimal weighting via monotonic min–max rescaling,

which preserves the sample ordering, and retain the PPO/GRPO clipping mechanism for stable policy updates.