

# Reasoning with Sampling: Cutting at Decision Points

Felix Zhou  
Yale University

Anay Mehrotra  
Stanford University

Quanquan C. Liu  
Yale University

## Abstract

Frontier reasoning models are produced by post-training base language models with reinforcement learning. Recent work has challenged this by showing that sampling from a sharpened version of the base model’s distribution, a so-called *power distribution*, elicits comparable reasoning without additional training, curated datasets, or verifiers. However, making this method practical requires *efficiently sampling* from the power distribution. A sampler needs to “mix” to the power distribution, which necessitates moving between modes of the target distribution; intuitively, *e.g.*, trying different reasoning strategies. The samplers proposed in prior works repeatedly select a “cut” position in the current reasoning trace uniformly at random and resample the suffix from that position onward. However, reasoning traces typically contain a few consequential decisions (*e.g.*, the choice of proof strategy or algorithm), and we observe that a uniformly chosen cut tends to rewrite local details rather than revisit decision points. We introduce an algorithm (Entropy-Cut Metropolis–Hastings) that uses the base model’s next-token entropy as a proxy to identify key decision points and resample from those positions. We empirically verify that entropy jumps are a useful proxy for decision points and, in a stylized model of reasoning, prove that our method’s mixing time scales with the number of decisions in a trace rather than with the number of tokens, which can be much larger. Across MATH500, HumanEval, GPQA Diamond, and AIME26, our method consistently improves over baselines and RL-trained models.

## 1 Introduction

Frontier reasoning models are typically obtained by posttraining base language models with reinforcement learning. This paradigm has produced striking gains on mathematics, coding, and scientific reasoning [OJKL+24; CBSP+25; GYZS+25]. At the same time, it raises a basic question: does posttraining teach models fundamentally new reasoning procedures, or does it mainly elicit abilities already encoded in the base model? One explanation for these gains is the *sharpening* view. Under this view, base models already assign non-negligible probability to many high-quality reasoning traces, and posttraining improves performance by concentrating more mass on these traces. Indeed, one mechanism by which posttraining improves reasoning performance is to convert strong top- $k$  performance into stronger top-1 performance [NLGS+25; YCLZ+25]. A growing body of evidence supports this perspective: RL-trained models often improve pass@1 without expanding the set of reasoning paths available in the base distribution [HFW25; SLXG+25; SKM25; YCLZ+25].

Building on this, Karan and Du [KD26] showed that strong reasoning behavior can be elicited from a base model without additional training by sampling from a sharpened version of the base

Uniform-Cut MH	
we have $x = 0$ and $y = 3$ .	<i>Current Completion</i>
### Step 1: Calculate $r$	
$r = \sqrt{x^3 + y^3} = \sqrt{0^3 + 3^3}$ [CUT];	
$r = \sqrt{27}$	<i>Proposal Suffix</i>

Entropy-Cut MH (Ours)	
we have $x = 0$ and $y = 3$ .	<i>Current Completion</i>
[CUT]### Step 1: Calculate $r$	
$r = \sqrt{x^3 + y^3} = \sqrt{0^3 + 3^3} = \sqrt{27}$	
$r = \sqrt{27}$	<i>Proposal Suffix</i>
### Step 1: Calculate $r$	
$r = \sqrt{x^2 + y^2} = \sqrt{0^2 + 3^2} = 3$	




Figure 1: **Entropy-Cut MH revises reasoning traces at decision points and improves accuracy.** Left: uniform cuts can splice proposals inside a local calculation, producing suffixes that only rewrite nearby tokens. Entropy-Cut instead cuts near high-uncertainty reasoning steps, allowing the proposal to reconsider the underlying continuation. Right: on Qwen2.5-7B, this targeted proposal improves accuracy over standard sampling, low-temperature sampling, SMC, TMC, and uniform-cut MH across MATH500, HumanEval, and GPQA Diamond.

distribution. Concretely, for a base distribution  $p$  over complete traces, they target the power distribution

$$\Pi_T(x) \propto p(x)^\alpha,$$

which upweights traces already assigned high likelihood by the base model. This approach is attractive because it is training-free, dataset-free, and verifier-free, making it a form of test-time scaling that uses only the base model’s own probabilities. Sampling from  $\Pi_T$  is, however, computationally non-trivial (see Section 2). Karan and Du address this with a stagewise Metropolis–Hastings sampler that iteratively revises a candidate trace: it picks a cut position, keeps the prefix before the cut, resamples the suffix from a proposal model, and accepts or rejects the new trace via the Metropolis–Hastings correction. This sampler succeeds in eliciting strong reasoning, but its efficiency is the bottleneck for the entire approach, raising the key question:

*How can we sample from the power distribution efficiently?*

**Prior Work.** Several recent works have sought to address this question by departing from the stagewise Metropolis–Hastings framework. Azizi, Potraghloo, Ahmadi, Kundu, et al. [APAK+26] replace it with a Sequential Monte Carlo method that maintains particles and corrects their weights token by token, while Ji, Tutunov, Zimmer, and Ammar [JTZA26] sample autoregressively by approximating the next-token conditionals of  $\Pi_T$ . Hence, both of these works depart from the approach of prior work and use new and more sophisticated sampling algorithms. We instead retain the Metropolis–Hastings sampler of Karan and Du [KD26] and modify it to better suit the structure of reasoning traces.

**Our Approach.** For an MCMC method to mix efficiently, it must be able to transition between the modes of its target distribution. For  $\Pi_T$ , which is a distribution over sequences of tokens, these

modes correspond, heuristically, to qualitatively different lines of reasoning: a proof by induction and a proof by contradiction may both have high probability under the sharpened distribution, but moving between them requires revisiting the early decision that chose the proof strategy (Figure 2).

**Our Contributions.** Our contributions are as follows.

- We introduce *Entropy-Cut Metropolis–Hastings* (Section 3), a training-free sampler for  $\Pi_T$  that modifies the stagewise sampler of Karan and Du [KD26] by placing cuts at positions of high next-token entropy, a proxy for decision points (see Figure 1), rather than uniformly. The Metropolis–Hastings correction ensures the target distribution is unchanged.
- We formalize the benefit of cutting at decision points (Section 4) in a *reasoning-tree* model, where the possible reasoning traces are encoded as a tree whose root-to-leaf paths are token sequences and whose branch nodes correspond to key decision points. We show that entropy-cut mixing scales with the number of decision points  $k$ , while uniform-cut mixing can scale with the token depth  $T$ , which can be much larger than  $k$  (Theorem 4.1).
- We empirically demonstrate the effectiveness of our algorithm over a range of models (Qwen2.5-7B, Qwen2.5-Math-7B, Qwen3-8B-Base, Phi-3.5-mini-instruct, Phi-4-mini-instruct) and reasoning tasks (MATH500, HumanEval, GPQA Diamond, AIME26). As summarized in Figure 1, our results show that more suitable power sampling techniques can *better* extract the latent reasoning capabilities of language models. Moreover, over multiple samples, we avoid a collapse in diversity despite the increased single-shot performance.

**Organization.** The next sections review the power-distribution sampling framework, introduce Entropy-Cut MH, and analyze it through a reasoning-tree model. We discuss the closest related methods in context and defer broader related work to Appendix A.

## 2 Preliminaries on Sharpening and Metropolis–Hastings Algorithm

In this section, we set up notation, introduce the *power distribution*, and describe the stagewise Metropolis–Hastings sampler of Karan and Du [KD26] that serves as our starting point.

**Notation.** For nonnegative functions  $f, g : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ , we write  $f \lesssim g$  if there exists an absolute constant  $C < \infty$  such that  $f(x) \leq Cg(x)$  for all  $x \in \mathbb{N}$ . Similarly, we write  $f \gtrsim g$  if  $g \lesssim f$  and  $f \simeq g$  if  $f \lesssim g$  and  $f \gtrsim g$ . We work with finite sequences of tokens and denote a sequence of length  $T + 1$  by  $x_{0:T} := (x_0, x_1, \dots, x_T)$ . For any  $0 \leq t \leq T$ , we write  $x_{<t} := (x_0, \dots, x_{t-1})$  and  $x_{>t} := (x_{t+1}, \dots, x_T)$  for the prefix and suffix around position  $t$ . Let  $x_p$  be a prompt and let  $p$  be an autoregressive language model which, given  $x_p$ , generates a continuation  $x_{0:T}$  with probability  $p(x_{0:T} \mid x_p) = \prod_{t=0}^T p(x_t \mid x_p, x_{<t})$ . Throughout the paper, we condition on the prompt  $x_p$  and suppress it from the notation, writing

$$p(x) := \prod_{t=0}^{\ell} p(x_t \mid x_{<t})$$

for any prefix  $x = x_{0:\ell}$ . Sampling from  $p$  is straightforward: drawing  $x_t \sim p(\cdot \mid x_{<t})$  sequentially for  $t = 0, 1, \dots, \ell$  produces a sample from  $p(x)$  in  $\ell + 1$  forward passes. As we discuss, *this convenience does not extend to our target distribution, and sampling efficiently is a central challenge of our paper.*

**Power Distribution and Sharpening.** A growing body of evidence suggests that RL-posttraining does not teach base models fundamentally new reasoning capabilities, but rather concentrates their sampling on high-likelihood reasoning traces that are already present in the base distribution [HFW25; SLXG+25; SKM25; YCLZ+25]. In other words, among candidate continuations of a fixed prompt, those assigned higher probability by the base model tend to be more likely to be correct. This motivates *sharpening*: reweighting sequences by a monotone function of  $p(x)$  that enlarges the ratio between likely and unlikely sequences, so that the resulting distribution concentrates on the sequences the base model is already most confident about. Building on this insight, Karan and Du [KD26] proposed sampling from the *power distribution* as a concrete, training-free mechanism for sharpening. Fix a parameter  $\alpha > 1$ , called the *sharpening power*. For each sequence length  $\ell$ , the power distribution  $\Pi_\ell$  is obtained by raising  $p(x)$  to the  $\alpha$ -th power and renormalizing:

$$\Pi_\ell(x) := \frac{p(x)^\alpha}{Z_{\ell,\alpha}} \quad \text{where} \quad Z_{\ell,\alpha} := \sum_{y_{0:\ell}} p(y_{0:\ell})^\alpha. \quad (\text{Target: } \alpha\text{-Power Distribution})$$

The parameter  $\alpha$  controls the strength of sharpening: the probability ratio between two sequences is amplified from  $p(x)/p(y)$  under  $p$  to  $(p(x)/p(y))^\alpha$  under  $\Pi_\ell$ . As  $\alpha \rightarrow 1$ , the power distribution reduces to  $p$ , and as  $\alpha \rightarrow \infty$ , it concentrates on the most likely completion. The distribution that we ultimately wish to sample from is  $\Pi_T$ , where  $T$  is the total generation length. Intuitively, sampling from  $\Pi_T$  favors tokens whose continuations concentrate on a few high-likelihood futures, rather than tokens that spread probability over many mediocre ones. This implicit bias toward “planning ahead” is particularly valuable for reasoning, where a few early token choices can determine whether the rest of the trace succeeds or fails [AABB+24; LKC25]. Karan and Du [KD26] show in their simulations that being able to sample from  $\Pi_T$  can elicit comparable reasoning capabilities without additional training.

Next, we explain why sampling from  $\Pi_T$  is hard and how prior work overcomes this.

**Sampling from the Power Distribution Is Hard.** At first sight, it might seem that sampling from  $\Pi_T$  reduces to low-temperature sampling at temperature  $1/\alpha$ , where each next-token conditional  $p(\cdot | x_{<t})$  is independently raised to the  $\alpha$ -th power and renormalized. This is not the case: low-temperature sampling sharpens and normalizes each next-token distribution locally, whereas the power distribution  $\Pi_T$  sharpens complete traces and normalizes globally through  $Z_{T,\alpha}$ . In fact, the two distributions can be made arbitrarily far apart in total variation, even for sequences of just two tokens (see Appendix B.3). More fundamentally, computing  $Z_{T,\alpha}$  requires summing over an exponentially large space of token sequences, making exact sampling from  $\Pi_T$  intractable.

**Metropolis–Hastings for Sampling from the Power Distribution.** Crucially, although the normalizing constant  $Z_{T,\alpha}$  is intractable, the *unnormalized* density  $p(x)^\alpha$  is easy to evaluate: a single autoregressive forward pass yields  $p(x)$ . For sampling from targets specified only through an unnormalized density, a central tool is the Metropolis–Hastings (MH) algorithm [MRRT+53; Has70]. Given a proposal distribution  $q(x|x')$ , a single MH step at state  $x$  draws a candidate  $x' \sim q(x, \cdot)$  and accepts it with probability  $\min\{1, \pi(x')q(x'|x)/(\pi(x)q(x|x'))\}$ , otherwise the chain stays at  $x$ . Iterating this rule produces a Markov chain whose stationary distribution is  $\pi$  (under minimal assumptions of aperiodicity and irreducibility, which hold in our setting); see Appendix B.2.

The key property of MH is that the acceptance rule depends on  $\pi$  only through ratios of densities, so the intractable normalizer  $Z_{T,\alpha}$  cancels in  $\Pi_T(x')/\Pi_T(x)$ , leaving a ratio involving only  $p(x')^\alpha$

and  $p(x)^\alpha$  that is computable from autoregressive forward passes (see [Remark B.1](#)). However, *the efficiency of MH depends critically on the proposal distribution  $q$ : a poor choice can cause the chain to mix exponentially slowly, i.e., require an exponentially large number of iterations before its distribution converges*. This problem is especially acute for  $\Pi_T$ , whose state space  $|V|^T$  grows exponentially in the generation length [[GLP17](#); [BMNW22](#)].

To overcome this, Karan and Du [[KD26](#)] propose a stagewise MH sampler that (partially) addresses this challenge through two design choices (see [Algorithm 3](#)). First, instead of targeting  $\Pi_T$  directly, they grow the generation in blocks of size  $B$ : samples from  $\Pi_{kB}$  are used to initialize an MH chain targeting  $\Pi_{(k+1)B}$ , avoiding pathological cold starts. Second, within each stage, the proposal distribution is defined by a *cut distribution*  $\lambda(\cdot; x)$ , a distribution over positions  $\{1, \dots, \ell\}$ : the chain draws a cut index  $m \sim \lambda(\cdot; x)$ , retains the prefix  $x_{0:m-1}$ , and redraws the suffix  $x_{m:\ell}$  from a proposal model  $p_{\text{prop}}$ . Concretely, the proposal distribution assigns zero mass to any candidate  $x'$  whose prefix disagrees with  $x$  before position  $m$ , and autoregressive probability  $\prod_{t=m, \dots, \ell} p_{\text{prop}}(x'_t | x'_{<t})$  to any candidate that agrees on the prefix (see [Remark B.2](#) for the explicit form). The resulting acceptance probability, after the standard MH correction for the state-dependent cut distribution, is

$$A_m(x, x') = \min \left\{ 1, \frac{p(x')^\alpha}{p(x)^\alpha} \cdot \frac{\lambda(m; x')}{\lambda(m; x)} \cdot \frac{p_{\text{prop}}(x_{m:\ell} | x_{0:m-1})}{p_{\text{prop}}(x'_{m:\ell} | x_{0:m-1})} \right\}. \quad (1)$$

Karan and Du [[KD26](#)] select a *uniform cut distribution*,  $\lambda_{\text{unif}}(m; x) \equiv 1/\ell$ , paired with a proposal  $p_{\text{prop}}$  given by the low-temperature version of  $p$ , which treats every token position as an equally useful place to revise the trace. In [Algorithm 1](#), we present pseudocode for a meta-algorithm that generalizes Karan and Du [[KD26](#)]'s method and allows for a general cut distribution. As we argue in the next section, *the uniform cut distribution from [[KD26](#)] is suboptimal and a key contribution of our work is a carefully designed cut distribution* that improves empirical performance by reducing mixing time.

### 3 Entropy-Cut Metropolis–Hastings

In this section, we introduce our algorithm for sampling from the power distribution  $\Pi_T$ .

Our algorithm is an instantiation of the meta-algorithm from [Algorithm 1](#). It is identical to the algorithm of Karan and Du [[KD26](#)] except in the choice of the *cut distribution*  $\lambda$ —the rule that chooses where each MH proposal is cut. Rather than cutting uniformly at random, our cut distribution concentrates cuts at positions where the model faces a genuine choice. The motivation for this change is that the uniform-cut sampler can mix slowly.

A standard way to think about slow mixing is through *conductance*. Roughly speaking, conductance measures how easily a Markov chain can leave one region of its state space and move to another: if the chain tends to get stuck in a region for many steps, then the conductance is low, and mixing is slow.

In our setting, these regions correspond to different high-level lines of reasoning. For example, a proof by induction and a proof by contradiction may lie in different regions, and moving from one to the other requires the sampler to revisit the early decision that chose the proof strategy. More generally, reasoning traces contain only a few such high-level choices—which proof technique to use, how to split into cases, which algorithm to implement—followed by many tokens that simply work out the consequences of those choices. A sampler that treats every token position as

---

**Algorithm 1:** Meta-Algorithm: Stagewise Metropolis–Hastings

---

**Input** : cut-point sampler  $\lambda$ ; proposal models  $p_{\text{prop}}$ ; power  $\alpha > 1$ ; length  $T$   
**Hyperparameters**: block size  $B$ ; number of MCMC steps per stage  $N_{\text{MCMC}}$   
**Output** : continuation  $x_{0:T}$  approximately distributed according to  $\Pi_T$

- 1 Set  $K \leftarrow \lceil T/B \rceil$ ,  $T_0 \leftarrow 0$ , and  $x^{(0)}$  to the empty continuation
- 2 **for** each stage  $k \leftarrow 1, 2, \dots, K$  **do**
- 3     Set the  $k$ -th stage’s target length  $T_k \leftarrow \min\{kB, T\}$
- 4     Sample a suffix  $s \sim p_{\text{prop}}(\cdot \mid x^{(k-1)})$  of length  $T_k - T_{k-1}$
- 5     Set the current state  $x \leftarrow (x^{(k-1)}, s)$
- 6     *// Metropolis–Hastings chain targeting the power distribution for length- $T_k$*
- 7     **for** each MCMC step  $n \leftarrow 1, 2, \dots, N_{\text{MCMC}}$  **do**
- 8         Sample a cut point  $m \sim \lambda(\cdot; x)$  (from which the current continuation will be revised)
- 9         Sample a (new) suffix  $s \sim p_{\text{prop}}(\cdot \mid x_{0:m-1})$  of length  $T_k - m$
- 10         Set the proposed state as  $x^{\text{prop}} \leftarrow (x_{0:m-1}, s)$
- 11         Compute the acceptance probability  $A_m(x, x^{\text{prop}})$  from Equation (1)
- 12         With probability  $A_m(x, x^{\text{prop}})$ , *accept* the proposal and set  $x \leftarrow x^{\text{prop}}$
- 13         Otherwise, *reject* the proposal and leave  $x$  unchanged
- 14     Set the  $k$ -th stage’s output as  $x^{(k)} \leftarrow x$
- 15 **return**  $x^{(K)}$

---

equally worth revising will usually cut too late, so it mostly rewrites details within the same line of reasoning rather than moving to a different one. The example below makes this concrete in the setting of a math proof.

**Example 3.1** (Proof Strategy as an Early Choice). Consider proving a generalization bound. Early in the trace, the model may choose between a uniform-convergence argument and an algorithmic-stability argument. Once this choice is made, later tokens mostly execute the chosen proof strategy (see Figure 2). If the choice occurs at position  $t^*$  in a trace of length  $\ell$ , a uniform cut revisits it with probability only  $t^*/\ell$ ; when  $t^* \ll \ell$ , the sampler rarely tries a different approach, creating a low-conductance bottleneck that increases mixing time.

A natural way to address this is therefore to revisit such consequential decisions more often. To turn this idea into an algorithm, however, we need a generic way to identify where these decisions occur. In some special cases, one could do this using ad hoc rules—for example, in a proof one might ask the model to write in explicit steps and then cut at the beginning of each step. But such rules do not generalize across the wide variety of reasoning tasks handled by modern LLMs.

Our first **insight is to use the entropy of the base model’s next-token distribution as a proxy for a decision point**. Intuitively, entropy is high when several continuations are plausible, and low when the continuation is largely determined. For each  $0 \leq t \leq \ell$ , we define  $h_t(x) := -\sum_{v \in V} p(v \mid x_{<t}) \log p(v \mid x_{<t})$ , the Shannon entropy of the base model’s next-token distribution after  $x_{<t}$ .

Using  $h_t$  directly has one drawback: after an important decision point, entropy often remains elevated for several subsequent tokens while the consequences of that decision are being worked out. A cut distribution proportional to  $h_t$  would therefore spread mass across an entire high-entropy stretch, even though the most informative place to cut is often where the uncertainty first increases.

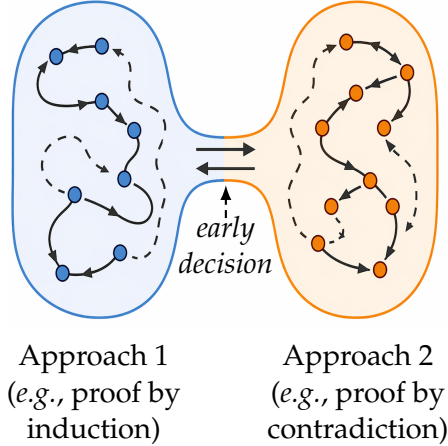


Figure 2: Illustration of a low-conductance bottleneck. Each region represents a family of reasoning traces that share the same high-level approach (e.g., two different proof strategies). Within a region, the sampler can make many local edits. To move from one region to the other, however, the sampler must revisit an *early decision* that determined the approach. When this happens rarely, it creates a low-conductance bottleneck leading to slow mixing.

Motivated by this, we use the *positive entropy jump*

$$\Delta_t(x) := \max \{0, h_t(x) - h_{t-1}(x)\} \quad (1 \leq t \leq \ell), \quad (2)$$

which emphasizes positions where the model’s uncertainty has just increased relative to the preceding token, while suppressing later tokens in the same high-entropy region. This leads to the following cut law.

**Definition 1** (Entropy-Cut Distribution). *Fix a cut power  $\beta \geq 0$ . The entropy-cut distribution*

$$\text{for a state } x \text{ of length } \ell, \quad \lambda_\beta(m; x) \propto \Delta_m(x)^\beta \quad (1 \leq m \leq \ell).$$

*In the edge case where  $\Delta_t(x) = 0$  for every  $t$ , the distribution is uniform over  $\{1, \dots, \ell\}$ .*

Here, the choice  $\beta = 0$  recovers the uniform cut distribution  $\lambda_{\text{unif}}$  of Karan and Du [KD26], while larger  $\beta$  places more mass on larger entropy jumps.

Entropy has recently been used to identify branching points in chain-of-thought reasoning and to allocate additional computation at such points [WYGZ+25; ZWMZ+25; LCGZ26]. Our use of entropy is somewhat different: rather than using it to guide search, training, or evaluation of the output, we use entropy jumps to define the *proposal distribution* of a Metropolis–Hastings sampler targeting the power distribution. In particular, entropy affects the MH sampling process.

**Entropy-Cut MH and its Running Time.** Our algorithm, *Entropy-Cut MH*, is obtained by instantiating the meta-algorithm of Algorithm 1 with the cut distribution  $\lambda_\beta$ ; see Algorithm 2 for the full pseudocode. Relative to the uniform-cut sampler of Karan and Du [KD26], the only change is the choice of cut distribution: the stagewise targets, the proposal model  $p_{\text{prop}}$ , and the Metropolis–Hastings acceptance rule are unchanged. Moreover, the entropies  $h_t(x)$  are already produced by the forward passes used to evaluate  $p(x)^\alpha$ , so *computing  $\lambda_\beta$  adds no cost over the uniform-cut baseline*.

### 3.1 Empirical Validation of the Entropy Proxy

The entropy-cut rule is useful only if large positive entropy jumps tend to occur near genuine decision points in the model’s reasoning. We test this directly by comparing what happens when

we resample from positions with large  $\Delta_t$  to what happens when we resample from positions with very small  $\Delta_t$ .

**Experiment.** We conduct a simple experiment to test how sensitive an LLM’s continuation is to resampling at different points in a chain-of-thought. For each MATH500 problem, we generate a single completion and record the token-level entropies along the way. We then compute the entropy jumps  $\Delta_t$ , select 5 cut positions from the top decile of the  $\Delta_t$  values and 5 from the bottom decile, and resample the suffix 16 times from each selected position while keeping the prefix fixed. We quantify the resulting variation in two ways: *normalized edit distance* between the resampled suffixes, computed as the average pairwise token-level Levenshtein distance divided by the mean suffix length, and *distinct-answer fraction*, computed as the number of unique parsed final answers among the resamples divided by the number of resamples. Together, these give a per-position measure of suffix divergence.

**Findings and Interpretation.** We find that cutting at top-decile  $\Delta_t$  positions produces substantially more divergent continuations than cutting at bottom-decile positions. The normalized edit distance is about  $1.33\times$  higher (0.287 vs. 0.216), and the distinct-answer fraction is about  $1.36\times$  higher (0.63 vs. 0.47). This is direct evidence that large entropy jumps tend to mark branching points in the reasoning trace: resampling at such points leads to more varied downstream text and more varied conclusions, whereas resampling at low- $\Delta_t$  positions typically yields much more consistent suffixes.

## 4 Theoretical Results on Mixing Time

In this section, we study the mixing time of our method under the following stylized model:

**Definition 2** (Reasoning Tree). *Fix a prompt and a target length  $T$ . A reasoning tree for  $(\Pi_T, p_{\text{prop}})$  is a rooted, token-labeled tree of depth  $T$  whose root-to-leaf paths are token sequences  $x_{0:T}$ , such that  $\Pi_T$  and  $p_{\text{prop}}$  are both supported on these paths. For a root-to-leaf path  $x = x_{0:T}$ , the node at depth  $t \in \{1, \dots, T\}$  is a branch node on the path if the prefix  $x_{0:t-1}$  has at least two children in the tree, and a chain node otherwise.*

Intuitively, a *chain node* represents positions where the next token is effectively determined—for instance, the next step of a calculation already in progress, or completing the implementation of an already-described function. A *branch node*, in contrast, represents a position where a decision must be made, *e.g.*, which calculation to do or which function to implement. For instance, in [Example 3.1](#) the reasoning tree has a branch at depth  $b_1 = O(1)$  where the path splits by proof technique; subsequent branches encode decisions internal to the chosen technique.

The reasoning tree constrains the support of the power distribution  $\Pi_T$  but does not place any constraint on how it distributes its mass across these paths. How this mass is distributed can substantially affect the time a method takes to converge to  $\Pi_T$ . Indeed, if  $\Pi_T$  places nearly all of its mass on a single path  $p$ , then any instantiation of Meta-Algorithm 1 lands on  $p$  quickly and the choice of the cut distribution is irrelevant. The challenging regime is where  $\Pi_T$  places comparable mass on several paths and the sampling algorithm must “explore” all of them. To isolate the effect of the cut distribution, we analyze our method in this regime, formalized in the condition below.

**Condition 1** (Approximately Symmetric Reasoning Tree). *Let  $U$  be the uniform distribution over root-to-leaf paths. For some  $\eta \geq 0$ , the reasoning tree satisfies the following symmetry conditions:*

1. For each root-to-leaf path  $x$ ,  $e^{-\eta}U(x) \leq \Pi_T(x) \leq e^\eta U(x)$ .
2. For each root-to-leaf path  $x$ , the proposal distribution satisfies  $\frac{U(x_{>t}|x_{0:t})}{p_{\text{prop}}(x_{>t}|x_{0:t})} \leq e^\eta$ .
3. For each root-to-leaf path  $x$ , the entropy jump  $\Delta_t(x)$  is  $\Delta$  at branch nodes and 0 otherwise.
4. For each branch node in the tree, the subtrees rooted at its children are isomorphic.

The first two conditions control how far the power distribution  $\Pi_T$  and the proposal distribution  $p_{\text{prop}}$  can deviate from the uniform distribution  $U$  over the root-to-leaf paths. Here, the parameter  $\eta$  controls the amount of deviation, and  $\eta = 0$  corresponds to perfect symmetry, the hardest regime. Condition 3 formalizes our observation in [Section 3](#) that branch nodes have a much larger entropy jump than chain nodes. To simplify the statement, we assume all branch nodes have the same entropy jump  $\Delta$ . Our result extends to the more general case at the cost of a more complicated statement. Finally, Condition 4 assumes the reasoning tree is symmetric, so all decisions at a branch point appear “equally good.” Again, this isolates the hardest case for the analysis, and the result extends to more general trees. In particular, due to symmetry, this condition implies that each root-to-leaf path has the same number of branch nodes, say  $k$ . We call  $k$  the *semantic depth* of the reasoning tree.

Recall that the *mixing time*  $\tau_{\text{mix}}^\lambda(\epsilon)$  of an MH chain measures how many steps it needs before its distribution is within  $\epsilon$  of the target in TV-distance.<sup>1</sup> Our main result compares the mixing times  $\tau_{\text{mix}}^{\text{ec}}$  and  $\tau_{\text{mix}}^{\text{unif}}$  of the entropy-cut and uniform-cut MH chains, respectively.

**Theorem 4.1** (Mixing Separation Under Approximate Symmetry). *Consider the reasoning-tree model under [Condition 1](#). Fix  $\beta > 0$  and set the number of stages  $K = 1$ . Let the number of branch nodes in a root-to-leaf path be  $k$  (i.e., the semantic depth is  $k$ ).<sup>2</sup> Then, for every  $\epsilon \in (0, 1/2)$ ,*

1. The entropy-cut chain satisfies  $\tau_{\text{mix}}^{\text{ec}}(\epsilon) \lesssim e^{2\eta} k \log 1/\epsilon$ ;
2. Moreover, if the first branch node is at depth  $b_1 \leq T/2$ , then  $\tau_{\text{mix}}^{\text{unif}}(\epsilon) = \Omega_\epsilon(T/b_1)$ .

Thus, for  $\eta = O(1)$  and constant accuracy, entropy-cut mixes in  $O(k)$  steps while uniform-cut requires  $\Omega(T/b_1)$  steps. **Thus, the relevant scale for entropy-cut is the number of semantic decisions  $k$ , whereas for uniform-cut it is the token depth  $T/b_1$  of the earliest decision.** When  $k \ll T/b_1$ , the separation is large: if the first decision occurs after  $O(1)$  tokens and  $k = o(T)$  (as in [Figure 2](#)), entropy-cut mixes in  $o(T)$  steps while uniform-cut requires  $\Omega(T)$ . The proof of [Theorem 4.1](#) appears in [Appendix B.7](#).

## 5 Evaluation of Performance on Reasoning and Instruction Following

In this section, we present our main empirical results.

<sup>1</sup>Formally,  $\tau_{\text{mix}}^\lambda(\epsilon) := \min \{n \geq 0 : \sup_x d_{\text{TV}}(P_\lambda^n(x, \cdot), \Pi_T) \leq \epsilon\}$ , where the supremum is over root-to-leaf paths.

<sup>2</sup>Due to Item 4 in [Condition 1](#), the number of branch nodes in each root-to-leaf path is the same.

## 5.1 Setup

We use a standard suite of reasoning benchmarks across math, coding, and STEM. All of our evaluations are in the *single-shot* setting, *i.e.*, with a single response string. We report the average scores across 8 independent runs for all datasets except AIME26, which we repeat 64 times. Each run is completed on a single H200 GPU. Our implementation leverages the `vllm v0.9.2` inference library [KLZS+23].

**Reasoning and Instruction-Following Benchmarks.** We evaluate on MATH500, HumanEval, GPQA Diamond, and AIME26, covering competition mathematics, code generation, graduate-level science questions, and olympiad-style mathematical reasoning. Dataset details and evaluation protocols are given in [Appendix C.1](#).

**Language Models.** Following Karan and Du [KD26], we use the base models `Qwen2.5-7B`, `Qwen2.5-Math-7B`, `Qwen3-8B-Base`, as well as the instruction-tuned models `Phi-3.5-mini-instruct` and `Phi-4-mini-instruct`.

**Baselines.** To demonstrate the efficacy of our algorithm, we compare against the following:

- **Standard:** Sampling from the conditional next-token distribution provided by the model.
- **Low-Temperature:** Sampling from the exponentiated conditional next-token distribution with exponent (*temperature*)  $\tau = 1/\alpha$ .
- **Sequential Monte Carlo (SMC):** SMC [PS99; DDJ06; Joh09] is a well-studied particle-gibbs method for sampling from distributions over sequences given access to unnormalized probabilities. In the context of power sampling, SMC maintains a weighted population of partial completions, updates their weights as tokens are decoded, and resamples when the weights become too uneven [APAK+26].
- **Twisted Monte Carlo (TMC):** Another well-studied sequential sampling technique is to *modulate* the conditional next-token distribution with some estimate of the future “rewards” (a *twist* function). For power sampling, the twist is taken to be an estimate of the future probabilities of the completion [JTZA26].
- **Uniform-Cut Metropolis–Hastings (MH):** Metropolis–Hastings with uniform cut-point sampling.

As SMC and TMC do not have public implementations, we implement our own following the description in their respective works.

**Hyperparameters.** Unless otherwise stated, we follow the hyperparameter choices of Karan and Du [KD26]. We set the maximum length  $T = 3072$  and block size  $B = 3072/16 = 192$  for MATH500, HumanEval, and double  $T$  and  $B$  for the more challenging GPQA Diamond and AIME26 datasets. We set the power exponent  $\alpha = 4.0$  for all datasets except for HumanEval, for which we found a higher  $\alpha = 5.0$  to perform better. The cut law exponent is set to  $\beta = 4.0$  for all datasets. The number of MH transition steps is set to  $N_{\text{MCMC}} = 10$ . We take the proposal distribution  $p_{\text{prop}}$  to be the low-temperature distribution with temperature  $\tau = 1/\alpha$ . The hyperparameters for SMC and TMC are based on their best-reported settings. See [Appendix C](#) for more details about SMC, TMC, and hyperparameter ablations.

## 5.2 Results

Our main findings are displayed in Table 1. **Across base and instruction-tuned models from different families, our sampling algorithm consistently improves accuracy on reasoning tasks.**

In particular, using Qwen2.5-7B, we observe a **+36.0** gain on MATH500 and a **+35.9** gain on HumanEval relative to standard sampling. The consistent improvement of entropy-cut over uniform-cut MH across model families and tasks supports our theoretical prediction (Theorem 4.1) that cutting at semantic decision points yields faster mixing, while reaffirming that improved power sampling can further extract the latent reasoning capabilities of existing models [KD26].

Table 1: **Entropy-Cut MH (ours) is competitive with and often improves over prior (power) sampling techniques across model families and tasks.** We benchmark our sampler on MATH500, HumanEval, GPQA Diamond, and AIME26. All scores are percentages. We bold the best score in each column within a model family and underline our scores (last row per block).

Model	Method	MATH500	HumanEval	GPQA Diamond	AIME26
Qwen2.5-7B	Standard	35.9	33.0	29.4	2.0
	Low-Temperature	62.3	64.3	29.5	5.7
	SMC	69.8	64.9	28.8	7.4
	TMC	70.4	66.2	27.5	<b>9.4</b>
	Uniform-Cut MH	67.4	66.8	29.4	8.2
	Entropy-Cut MH (ours)	<b>71.9</b>	<b>68.9</b>	<b>30.2</b>	<b>9.4</b>
Qwen2.5-Math-7B	Standard	50.1	31.0	29.0	4.8
	Low-Temperature	68.3	55.8	33.1	6.4
	SMC	78.9	55.2	32.0	11.8
	TMC	78.2	51.2	32.5	10.6
	Uniform-Cut MH	73.1	59.8	32.4	9.5
	Entropy-Cut MH (ours)	<b>79.0</b>	<b>59.9</b>	<b>34.1</b>	<b>13.1</b>
Qwen3-8B-Base	Standard	58.5	47.6	31.7	3.2
	Low-Temperature	77.3	73.7	37.6	6.4
	SMC	79.3	76.7	38.1	10.1
	TMC	79.7	75.2	38.6	9.8
	Uniform-Cut MH	79.5	75.8	38.8	9.5
	Entropy-Cut MH (ours)	<b>80.2</b>	<b>79.3</b>	<b>40.0</b>	<b>10.3</b>
Phi-3.5-mini-instruct	Standard	42.9	40.5	27.8	0.6
	Low-Temperature	45.3	62.2	32.5	1.6
	SMC	<b>49.8</b>	65.2	31.0	1.7
	TMC	49.3	<b>65.6</b>	31.5	3.0
	Uniform-Cut MH	49.1	64.9	31.8	2.9
	Entropy-Cut MH (ours)	<b>49.7</b>	<b>65.6</b>	<b>32.8</b>	<b>3.1</b>
Phi-4-mini-instruct	Standard	55.5	39.9	30.7	3.5
	Low-Temperature	66.4	65.5	32.1	6.8
	SMC	68.1	66.7	31.9	8.6
	TMC	66.1	66.1	32.1	6.9
	Uniform-Cut MH	68.7	67.9	32.0	8.0
	Entropy-Cut MH (ours)	<b>68.8</b>	<b>68.4</b>	<b>33.3</b>	<b>8.7</b>

## 5.3 Analysis

We analyze several important characteristics of various (power) sampling algorithms and defer additional details, such as hyperparameter ablations and running times, to Appendix C.

**Sequence Log-Probabilities.** We plot the distribution of the sequence log-likelihood values computed via the base model ( $Q_{\text{wen2.5-Math-7B}}$ ) on the MATH500 dataset across 8 random seeds in Figure 3. In addition, the 25th, 50th (median), and 75th percentile values are marked with dashed, solid, and dotted vertical lines. We observe that entropy-cut MH samples from higher-likelihood regions than standard, low-temperature, and uniform-cut MH. See Appendix C.3 for average confidence plots.




Figure 3: **Entropy-Cut MH (ours) samples from higher-probability regions than baselines.** We plot  $Q_{\text{wen2.5-Math-7B}}$  log-likelihoods for MATH500 responses across samplers. The 25th ( $Q_1$ ), 50th ( $Q_2$ ; median), and 75th ( $Q_3$ ) percentile values are marked with dashed, solid, and dotted vertical lines, respectively.

**Diversity and Pass@k.** We plot the pass@k accuracy (at least one correct answer among k) in Figure 4 using an unbiased estimator [CTJY+21]. This demonstrates that our method maintains diversity comparable to standard sampling despite its improved pass@1 performance.




Figure 4: **Entropy-Cut MH (ours) maintains diversity across multiple passes despite the improved pass@1 performance.** We plot  $Q_{\text{wen2.5-Math-7B}}$  pass@k accuracy for MATH500 responses across samplers.

## 6 Conclusion, Limitations, and Future Work

We introduced Entropy-Cut Metropolis–Hastings, a simple modification of the stagewise sampler of Karan and Du [KD26] that places cuts at positions of high next-token entropy rather than

uniformly over tokens, focusing proposals on points where the model appears to face a genuine choice. This preserves the same power-distribution target while better matching the structure of reasoning traces. Our theory shows that, in a reasoning-tree model, entropy-cut mixing scales with the number of semantic decisions rather than the full token length, and empirically, Entropy-Cut MH improves performance across models and reasoning benchmarks while maintaining sample diversity.

Several directions remain open. Our experiments focus on small and medium open models; evaluating entropy-cut sampling at larger scale would clarify how much further reasoning performance can be extracted from stronger base distributions. We also study single-shot completions, but many practical reasoning systems revise partial solutions, call tools, or branch across attempts; extending entropy-cut proposals to these settings could make the sampler a useful component in broader inference-time reasoning. Finally, our entropy proxy is deliberately simple; future work could learn or adapt the cut distribution from richer signals such as model uncertainty, verifier feedback, or cross-sample disagreement, while retaining the Metropolis–Hastings correction that keeps the target distribution unchanged.

## **Acknowledgments**

Felix Zhou acknowledges the support of the Natural Sciences and Engineering Research Council of Canada (NSERC).

## References

- [AABB+24] Marah Abdin, Jyoti Aneja, Harkirat Behl, Sebastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, James R. Lee, Yin Tat Lee, Yuanzhi Li, Weishung Liu, Caio C. T. Mendes, Anh Nguyen, Eric Price, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Xin Wang, Rachel Ward, Yue Wu, Dingli Yu, Cyril Zhang, and Yi Zhang. *Phi-4 Technical Report*. 2024. arXiv: 2412.08905. URL: <https://arxiv.org/abs/2412.08905> (cit. on pp. 4, 19).
- [APAK+26] Seyedarmin Azizi, Erfan Baghaei Potraghloo, Minoos Ahmadi, Souvik Kundu, and Massoud Pedram. “Power-SMC: Low-Latency Sequence-Level Power Sampling for Training-Free LLM Reasoning”. In: *arXiv preprint arXiv:2602.10273* (2026) (cit. on pp. 2, 10, 28).
- [Art26] Art of Problem Solving. 2026 AIME I and AIME II Problems. AIME I: [https://artofproblemsolving.com/wiki/index.php/2026\\_AIME\\_I\\_Problems](https://artofproblemsolving.com/wiki/index.php/2026_AIME_I_Problems) AIME II: [https://artofproblemsolving.com/wiki/index.php/2026\\_AIME\\_II\\_Problems](https://artofproblemsolving.com/wiki/index.php/2026_AIME_II_Problems). Accessed: 2026-04-20. 2026 (cit. on p. 28).
- [BMNW22] Afonso S. Bandeira, Antoine Maillard, Richard Nickl, and Sven Wang. *On Free Energy Barriers in Gaussian Priors and Failure of Cold Start MCMC for High-Dimensional Unimodal Distributions*. 2022. arXiv: 2209.02001. URL: <https://arxiv.org/abs/2209.02001> (cit. on p. 5).
- [CBIL+23] Charlie Chen, Sebastian Borgeaud, Geoffrey Irving, Jean-Baptiste Lespiau, Laurent Sifre, and John Jumper. *Accelerating Large Language Model Decoding with Speculative Sampling*. 2023. arXiv: 2302.01318 [cs.CL]. URL: <https://arxiv.org/abs/2302.01318> (cit. on p. 18).
- [CBSP+25] Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, Luke Marris, Sam Petulla, Colin Gaffney, Asaf Aharoni, Nathan Lintz, Tiago Cardal Pais, Henrik Jacobsson, Idan Szpektor, Nan-Jiang Jiang, Krishna Haridasan, Ahmed Omran, Nikunj Saunshi, Dara Bahri, Gaurav Mishra, Eric Chu, Toby Boyd, Brad Hekman, Aaron Parisi, Chaoyi Zhang, Kornrathop Kawintiranon, Tania Bedrax-Weiss, Oliver Wang, Ya Xu, Ollie Purkiss, Uri Mendlovic, Ilai Deutel, Nam Nguyen, Adam Langley, Flip Korn, Lucia Rossazza, Alexandre Ramé, Sagar Waghmare, Helen Miller, Nathan Byrd, Ashrith Sheshan, Raia Hadsell, Sangnie Bhardwaj, Pawel Janus, Tero Rissa, Dan Horgan, Alvin Abdagic, Lior Belenki, James Allingham, Anima Singh, Theo Guidroz, Srivatsan Srinivasan, Herman Schmit, Kristen Chiafullo, Andre Elisseeff, Nilpa Jha, Prateek Kolhar, Leonard Berrada, Frank Ding, Xiance Si, Shrestha Basu Mallick, Franz Och, Sofia Erell, Eric Ni, Tejasi Latkar, Sherry Yang, Petar Sirkovic, Ziqiang Feng, Robert Leland, Rachel Hornung, Gang Wu, Charles Blundell, Hamidreza Alvari, Po-Sen Huang, Cathy Yip, Sanja Deur, Li Liu, Gabriela Surita, Pablo Duque, Dima Damen, Johnson Jia, Arthur Guez, Markus Mircea, Animesh Sinha, Alberto Magni, Paweł Stradomski, Tal Marian, Vlado Galić, Wenhua Chen, Hisham Husain, Achintya Singhal, Dominik Grewe, François-Xavier Aubet, Shuang Song, Lorenzo Blanco, Leland Rechis, and 3326 others. “Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities”. In: *arXiv preprint arXiv:2507.06261* (2025). arXiv: 2507.06261 [cs.CL]. URL: <https://arxiv.org/abs/2507.06261> (cit. on p. 1).
- [CTJY+21] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. “Evaluating large language models trained on code”. In: *arXiv preprint arXiv:2107.03374* (2021) (cit. on pp. 12, 28).
- [DDJ06] Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. “Sequential monte carlo samplers”. In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* 68.3 (2006), pp. 411–436 (cit. on p. 10).

- [FKKG24] Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. “Detecting Hallucinations in Large Language Models Using Semantic Entropy”. In: *Nature* 630 (2024), pp. 625–630 (cit. on p. 18).
- [FRSP+24] Ekaterina Fadeeva, Aleksandr Rubashevskii, Artem Shelmanov, Sergey Petrakov, Haonan Li, Hamdy Mubarak, Evgenii Tsymbalov, Gleb Kuzmin, Alexander Panchenko, Timothy Baldwin, Preslav Nakov, and Maxim Panov. “Fact-Checking the Output of Large Language Models via Token-Level Uncertainty Quantification”. In: *Findings of the Association for Computational Linguistics: ACL 2024*. Association for Computational Linguistics, 2024, pp. 9367–9385. URL: <https://aclanthology.org/2024.findings-acl.558/> (cit. on p. 18).
- [GLP17] Reza Gheissari, Eyal Lubetzky, and Yuval Peres. *Exponentially Slow Mixing in the Mean-Field Swendsen–Wang Dynamics*. 2017. arXiv: [1702.05797](https://arxiv.org/abs/1702.05797) (cit. on p. 5).
- [GYZS+25] 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, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Hanwei Xu, Honghui Ding, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jingchang Chen, Jingyang Yuan, Jinhao Tu, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaichao You, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingxu Zhou, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. “DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning”. In: *Nature* 645.8081 (2025), pp. 633–638. URL: <http://dx.doi.org/10.1038/s41586-025-09422-z> (cit. on p. 1).
- [Has70] W. K. Hastings. “Monte Carlo Sampling Methods Using Markov Chains and Their Applications”. In: *Biometrika* 57.1 (1970), pp. 97–109 (cit. on p. 4).
- [HFW25] Andre He, Daniel Fried, and Sean Welleck. *Rewarding the Unlikely: Lifting GRPO Beyond Distribution Sharpening*. 2025. arXiv: [2506.02355](https://arxiv.org/abs/2506.02355) [cs.LG] (cit. on pp. 1, 4).
- [Joh09] Adam Johansen. “A tutorial on particle filtering and smoothing: Fifteen years later”. In: (2009) (cit. on p. 10).
- [JTZA26] Xiaotong Ji, Rasul Tutunov, Matthieu Zimmer, and Haitham Bou Ammar. “Scalable Power Sampling: Unlocking Efficient, Training-Free Reasoning for LLMs via Distribution Sharpening”. In: *arXiv preprint arXiv:2601.21590* (2026) (cit. on pp. 2, 10, 28).

- [KD26] Aayush Karan and Yilun Du. “Reasoning with Sampling: Your Base Model is Smarter Than You Think”. In: *The Fourteenth International Conference on Learning Representations (ICLR 2026)*. Oral. OpenReview.net, 2026. URL: <https://openreview.net/forum?id=Vsgq21dr4K> (cit. on pp. 1–5, 7, 10–12, 19, 23, 29).
- [KHRS+24] Jannik Kossen, Jiatong Han, Muhammed Razzak, Lisa Schut, Shreshth Malik, and Yarín Gal. “Semantic Entropy Probes: Robust and Cheap Hallucination Detection in LLMs”. In: *arXiv preprint arXiv:2406.15927* (2024). arXiv: 2406.15927 [cs.CL]. URL: <https://arxiv.org/abs/2406.15927> (cit. on p. 18).
- [KLZS+23] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. “Efficient Memory Management for Large Language Model Serving with PagedAttention”. In: *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*. 2023 (cit. on p. 10).
- [LCGZ26] Xianzhi Li, Ethan Callanan, Abdellah Ghassel, and Xiaodan Zhu. “Entropy-Gated Branching for Efficient Test-Time Reasoning”. In: *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Mar. 2026, pp. 5054–5069. URL: <https://aclanthology.org/2026.eacl-long.235/> (cit. on pp. 7, 18).
- [LKBE+24] Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. “Let’s Verify Step by Step”. In: *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7–11, 2024*. OpenReview.net, 2024. URL: <https://openreview.net/forum?id=v8L0pN6EOi> (cit. on p. 28).
- [LKC25] Marvin Li, Aayush Karan, and Sitan Chen. *Blink of an Eye: A Simple Theory for Feature Localization in Generative Models*. 2025. arXiv: 2502.00921. URL: <https://arxiv.org/abs/2502.00921> (cit. on pp. 4, 19).
- [LKM23] Yaniv Leviathan, Matan Kalman, and Yossi Matias. “Fast Inference from Transformers via Speculative Decoding”. In: *Proceedings of the 40th International Conference on Machine Learning*. Vol. 202. Proceedings of Machine Learning Research. PMLR, 2023, pp. 19274–19286 (cit. on p. 18).
- [LLXL+24] Zicheng Lin, Tian Liang, Jiahao Xu, Qiuzhi Lin, Xing Wang, Ruilin Luo, Chufan Shi, Siheng Li, Yujiu Yang, and Zhaopeng Tu. “Critical Tokens Matter: Token-Level Contrastive Estimation Enhances LLM’s Reasoning Capability”. In: *arXiv preprint arXiv:2411.19943* (2024). arXiv: 2411.19943 [cs.CL]. URL: <https://arxiv.org/abs/2411.19943> (cit. on p. 19).
- [LPW17] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. 2nd ed. American Mathematical Society, 2017 (cit. on p. 20).
- [MG21] Andrey Malinin and Mark Gales. “Uncertainty Estimation in Autoregressive Structured Prediction”. In: *International Conference on Learning Representations*. 2021. URL: <https://openreview.net/forum?id=jN5y-zb5Q7m> (cit. on p. 18).
- [MRRT+53] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. “Equation of State Calculations by Fast Computing Machines”. In: *The Journal of Chemical Physics* 21.6 (1953), pp. 1087–1092 (cit. on p. 4).
- [NLGS+25] Vaskar Nath, Elaine Lau, Anisha Gunjal, Manasi Sharma, Nikhil Baharte, and Sean Hendryx. “Adaptive Guidance Accelerates Reinforcement Learning of Reasoning Models”. In: *arXiv preprint arXiv:2506.13923* (2025). arXiv: 2506.13923 [cs.LG]. URL: <https://arxiv.org/abs/2506.13923> (cit. on p. 1).
- [OJKL+24] OpenAI, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Alex Low, Alice Helyar, Aleksander Madry, Alex Beutel, Alex Carney, Alex Iftimie, Alex Karpenko, Alexandre Passos, Alex Neitz, Alexander Prokofiev, Alexander Wei, Allison Tam, Amanda Bennett, Ananya Kumar, et al. “OpenAI o1 System Card”. In: *arXiv preprint arXiv:2412.16720*

- (2024). arXiv: [2412.16720](https://arxiv.org/abs/2412.16720) [cs.AI]. URL: <https://arxiv.org/abs/2412.16720> (cit. on p. 1).
- [PCIF+25] Mihir Prabhudesai, Lili Chen, Alex Ippoliti, Katerina Fragkiadaki, Hao Liu, and Deepak Pathak. “Maximizing confidence alone improves reasoning”. In: *arXiv preprint arXiv:2505.22660* (2025) (cit. on p. 28).
- [PS99] Michael K Pitt and Neil Shephard. “Filtering via simulation: Auxiliary particle filters”. In: *Journal of the American statistical association* 94.446 (1999), pp. 590–599 (cit. on p. 10).
- [RHSP+24] David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. “Gpqa: A graduate-level google-proof q&a benchmark”. In: *First conference on language modeling*. 2024 (cit. on p. 28).
- [SKM25] Yuda Song, Julia Kempe, and Remi Munos. *Outcome-based Exploration for LLM Reasoning*. 2025. arXiv: [2509.06941](https://arxiv.org/abs/2509.06941) [cs.LG] (cit. on pp. 1, 4).
- [SLXG+25] Rulin Shao, Shuyue Stella Li, Rui Xin, Scott Geng, Yiping Wang, Sewoong Oh, Simon Shaolei Du, Nathan Lambert, Sewon Min, Ranjay Krishna, Yulia Tsvetkov, Hannaneh Hajishirzi, Pang Wei Koh, and Luke Zettlemoyer. *Spurious Rewards: Rethinking Training Signals in RLVR*. 2025. arXiv: [2506.10947](https://arxiv.org/abs/2506.10947) [cs.AI] (cit. on pp. 1, 4).
- [WYGZ+25] Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xionghui Chen, Jianxin Yang, Zhenru Zhang, Yuqiong Liu, An Yang, Andrew Zhao, Yang Yue, Shiji Song, Bowen Yu, Gao Huang, and Junyang Lin. “Beyond the 80/20 Rule: High-Entropy Minority Tokens Drive Effective Reinforcement Learning for LLM Reasoning”. In: *arXiv preprint arXiv:2506.01939* (2025). Accepted to NeurIPS 2025. arXiv: [2506.01939](https://arxiv.org/abs/2506.01939) [cs.CL] (cit. on pp. 7, 18).
- [YCLZ+25] Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Yang Yue, Shiji Song, and Gao Huang. *Does Reinforcement Learning Really Incentivize Reasoning Capacity in LLMs Beyond the Base Model?* 2025. arXiv: [2504.13837](https://arxiv.org/abs/2504.13837) [cs.AI] (cit. on pp. 1, 4).
- [ZWMZ+25] Jinghan Zhang, Xiting Wang, Fengran Mo, Yeyang Zhou, Wanfu Gao, and Kunpeng Liu. “Entropy-based Exploration Conduction for Multi-step Reasoning”. In: *Findings of the Association for Computational Linguistics: ACL 2025*. Association for Computational Linguistics, 2025, pp. 3895–3906. URL: <https://aclanthology.org/2025.findings-acl.201/> (cit. on pp. 7, 18).

## A Additional Related Work

In this section, we discuss additional related lines of work and how they compare to our work.

**Speculative Decoding.** The structural idea of drawing a candidate from a cheap proposal and accepting it with a ratio-based rule resembles *speculative decoding* [CBIL+23; LKM23], in which a small draft model proposes continuations that are then accepted or rejected according to ratios of probabilities under a larger target model. The two methods share the accept–reject template but serve different purposes: speculative decoding accelerates exact sampling from a fixed autoregressive target, whereas MH produces samples from a target, here  $\Pi_T$ , that is not itself autoregressive.

**Entropy as an Uncertainty Signal.** Several works use entropy or related uncertainty estimates from a language model’s predictive distribution to identify unreliable generations. In autoregressive sequence prediction, token- and sequence-level uncertainty have been studied for error detection and out-of-domain detection [MG21]. For factual generation, Fadeeva, Rubashevskii, Shelmanov, Petrakov, et al. [FRSP+24] use token-level uncertainty to identify claims that are likely to be unreliable. These works use uncertainty primarily as a diagnostic for the generated output. We use the same broad signal for a different purpose: rather than deciding whether a completed answer should be trusted, we use local changes in uncertainty to decide where an MH proposal should cut and resample a reasoning trajectory.

**Entropy and Reasoning Decisions.** Closest to our setting are recent works that treat high-entropy tokens as places where a reasoning trajectory branches. Li, Callanan, Ghassel, and Zhu [LCGZ26] propose entropy-gated branching, which allocates additional test-time search at high-entropy moments where multiple continuations are plausible. Wang, Yu, Gao, Zheng, et al. [WYGZ+25] study entropy patterns in chain-of-thought reasoning and find that a small fraction of high-entropy tokens can act as “forks” that steer the model toward different reasoning paths. Some work also considers changes in uncertainty over time: Zhang, Wang, Mo, Zhou, et al. [ZWMZ+25] use changes in entropy and variance-entropy across consecutive reasoning steps to decide whether to deepen, branch, or stop exploration. Taking inspiration from these observations, we use entropy differences as a lightweight proxy for decision points in the model’s reasoning. Unlike beam- or tree-expansion methods, however, we do not allocate additional samples at every such point. Instead, we use the signal to choose cut locations inside an MH sampler, with the goal of revisiting uncertain reasoning decisions while avoiding cuts inside long low-entropy stretches of local computation.

**Semantic Entropy.** A related line of work argues that token-level entropy can be misleading in open-ended generation, because many different strings may express the same meaning. Farquhar, Kossen, Kuhn, and Gal [FKKG24] therefore propose *semantic entropy*, which measures uncertainty over semantic equivalence classes of sampled responses rather than over surface forms, and use it to detect confabulations in LLM outputs. Follow-up work such as Kossen, Han, Razzak, Schüt, et al. [KHRS+24] aims to approximate semantic entropy more efficiently from model representations. This line of work supports the general idea that uncertainty in the model distribution can reveal meaningful properties of the generation, but it also highlights a limitation of purely token-level entropy. In this paper we do not use semantic entropy: our cut rule is based on token-level entropy changes because they are cheap to compute during generation and naturally aligned with token-level cut positions. Understanding whether semantic entropy, or another meaning-aware uncertainty measure, can further improve the choice of cut locations is an interesting direction for

future work.

**Critical Tokens and Decision Points.** Several recent works observe that reasoning trajectories often contain a small number of high-leverage tokens. The *Phi-4* technical report introduces *pivotal token search*, which finds tokens whose generation substantially changes the probability of final correctness and uses them to construct targeted DPO pairs [AABB+24]. Lin, Liang, Xu, Lin, et al. [LLXL+24] identify *critical tokens* in incorrect mathematical reasoning traces and show that modifying these tokens can substantially improve downstream correctness. Li, Karan, and Chen [LKC25] study the related phenomenon of *critical windows*, where model behavior changes sharply over a narrow region of the generation process. Karan and Du [KD26] connect this idea to power sampling, showing that the power distribution can favor tokens whose futures concentrate on high-likelihood completions rather than tokens that merely look locally likely. Our method is in the same spirit, but uses entropy jumps as a verifier-free inference-time proxy for these decision points, allowing the Metropolis–Hastings sampler to revisit consequential choices without changing its target distribution.

## B Additional Discussion and Remarks

### B.1 Remarks on Metropolis–Hastings

**Remark B.1** (MH Acceptance Ratio Is Tractable for  $\Pi_T$ ). Because  $p$  is autoregressive,  $p(x) = \prod_{t=0}^{\ell} p(x_t | x_{<t})$  is computable from a single forward pass, and the MH ratio depends on  $p$  only through  $p(x')^\alpha / p(x)^\alpha$ . The intractable sum defining  $Z_{T,\alpha}$  therefore never needs to be computed. The cut-based proposal kernel below can likewise be sampled by a single autoregressive pass, so both the acceptance ratio and the proposal run in polynomial time despite the exponentially large state space.

**Remark B.2** (Explicit Form of the Cut-Based Proposal Kernel). For a cut index  $m$ , the proposal kernel that retains  $x_{0:m-1}$  and redraws  $x_{m:\ell}$  from  $p_{\text{prop}}$  is

$$Q_m(x, x') = \begin{cases} p_{\text{prop}}(x'_{m:\ell} | x_{0:m-1}), & \text{if } x'_{0:m-1} = x_{0:m-1}, \\ 0, & \text{otherwise,} \end{cases}$$

where  $p_{\text{prop}}(x_{m:\ell} | x_{0:m-1}) := \prod_{t=m}^{\ell} p_{\text{prop}}(x_t | x_{<t})$ . The factor  $\lambda(m; x') / \lambda(m; x)$  in the acceptance probability corrects for the state-dependence of the cut law and is what ensures that  $\Pi_\ell$  is the stationary distribution of the resulting chain.

### B.2 Irreducibility and Aperiodicity of the Stagewise MH Chain

In this section, we recall the standard conditions under which the stagewise Metropolis–Hastings chain converges to its target distribution, and explain why they hold in our setting.

Fix a stage length  $\ell$ , let

$$\Omega_\ell := \{x : \Pi_\ell(x) > 0\}$$

be the support of the target distribution at that stage, and let  $P_\ell$  denote one step of the corresponding MH chain. We adopt the convention that the earliest admissible cut redraws the entire continuation; with the indexing  $x_{0:\ell}$  used in the main text, this corresponds to allowing  $m = 0$ .

**Definition 3** (Irreducible). An MH chain  $P_\ell$  is said to be irreducible on  $\Omega_\ell$  if for every  $x, y \in \Omega_\ell$ , there exists  $r \geq 1$  such that  $P_\ell^r(x, y) > 0$ .

**Definition 4** (Aperiodic). An MH chain  $P_\ell$  is said to be aperiodic if for every  $x \in \Omega_\ell$ ,  $\gcd \{r \geq 1 : P_\ell^r(x, x) > 0\}$  is 1.

Since  $\Omega_\ell$  is finite and the Metropolis–Hastings construction makes  $P_\ell$  reversible with stationary distribution  $\Pi_\ell$ , irreducibility and aperiodicity imply that  $\|P_\ell^n(x, \cdot) - \Pi_\ell\|_{\text{TV}} \rightarrow 0$  for every  $x \in \Omega_\ell$  as  $n \rightarrow \infty$ ; see, e.g., Levin, Peres, and Wilmer [LPW17].

It remains to explain why these conditions hold in our setting. Assume first that the proposal model  $p_{\text{prop}}$  has full support on  $\Omega_\ell$ , in the sense that whenever  $y \in \Omega_\ell$ , every suffix of  $y$  has positive probability under  $p_{\text{prop}}$  conditional on its preceding prefix. Assume also that the cut law has full support on the admissible cut locations, i.e.,

$$\lambda(m; x) > 0 \quad \text{for every } x \in \Omega_\ell \text{ and every admissible } m.$$

This is immediate for the uniform-cut baseline, and for entropy-cut, it can be enforced by adding an arbitrarily small uniform floor, e.g.,  $\tilde{\lambda}_{\beta, \varepsilon}(m; x) \propto \Delta_m(x)^\beta + \varepsilon$ , ( $\varepsilon > 0$ ). Under these two support assumptions, irreducibility and aperiodicity are immediate. For irreducibility, fix any  $x, y \in \Omega_\ell$ . With positive probability, the chain chooses the earliest cut and then proposes  $y$  exactly by redrawing the entire continuation from  $p_{\text{prop}}$ ; by the support assumption, this proposal probability is strictly positive. The reverse proposal from  $y$  back to  $x$  is also positive, and both  $\Pi_\ell(x)$  and  $\Pi_\ell(y)$  are positive on  $\Omega_\ell$ , so the MH acceptance probability is strictly positive as well. Hence  $P_\ell(x, y) > 0$ , so the chain is irreducible. For aperiodicity, from any state  $x \in \Omega_\ell$ , there is positive probability of choosing any admissible cut and then redrawing exactly the current suffix, thereby proposing  $x$  itself. Such a proposal is accepted with probability one, so  $P_\ell(x, x) > 0$  for every  $x$ . Thus the chain has a self-loop at every state and therefore period one.

### B.3 Low-Temperature Sampling Is Not Power Sampling

In this section, we show that sampling from the power distribution  $\Pi_T(x_{0:T}) \propto p(x_{0:T})^\alpha$  is not equivalent to low-temperature sampling at temperature  $1/\alpha$ .

Concretely, let  $\Phi_T$  denote the low-temperature distribution where each next-token conditional  $p(\cdot \mid x_{<t})$  is independently raised to the  $\alpha$ -th power and renormalized. Both power sampling and low-temperature sampling alter the base model, but they do so in different ways:  $\Phi_T$  alters and normalizes each next-token distribution locally, whereas  $\Pi_T$  raises complete traces to the  $\alpha$ -th power and normalizes *globally*. We show that the two distributions can be almost disjoint, even for sequences of just two tokens.

**Proposition B.3.** Fix any  $\alpha > 1$  and any  $\varepsilon > 0$ . There exists an autoregressive distribution  $p$  over two-token strings such that the power distribution  $\Pi_1$  and the low-temperature distribution  $\Phi_1$  satisfy

$$\|\Pi_1 - \Phi_1\|_{\text{TV}} \geq 1 - \varepsilon.$$

The construction is simple. One first-token choice has smaller marginal probability under the base model but leads to a single high-probability completion. The other has larger marginal probability

but spreads its mass over many low-probability completions. Low-temperature sampling favors the second choice because it sharpens marginal next-token probabilities locally; the power distribution favors the first because it sharpens complete traces, rewarding the concentrated completion.

*Proof.* Let  $\delta = \varepsilon/2$ . Choose  $R > 1$  large enough that  $\frac{1}{1+R^\alpha} \leq \delta$ , and then choose  $N$  large enough that  $\frac{R^\alpha}{N^{\alpha-1}} \leq \delta$ . Consider two-token strings whose first token is either  $a$  or  $b$ . Let the second-token alphabet contain a special token  $\star$  and tokens  $1, \dots, N$ . Define the base distribution  $p$  by

$$p(a, \star) = \frac{1}{1+R} \quad \text{and} \quad p(b, i) = \frac{R}{N(1+R)} \quad \text{for } i = 1, \dots, N,$$

and assign probability zero to all other two-token strings. Equivalently, the base model chooses  $a$  with probability  $1/(1+R)$ , chooses  $b$  with probability  $R/(1+R)$ , completes  $a$  deterministically with  $\star$ , and completes  $b$  uniformly among  $N$  choices.

Let  $A$  be the event that the first token is  $a$ . Under low-temperature sampling, the first-token conditional is sharpened locally, so

$$\Phi_1(A) = \frac{(1/1+R)^\alpha}{(1/1+R)^\alpha + (R/1+R)^\alpha} = \frac{1}{1+R^\alpha} \leq \delta.$$

Under the power distribution, however, the probability of  $A$  is obtained by sharpening the probabilities of complete two-token strings:

$$\Pi_1(A) = \frac{(1/1+R)^\alpha}{(1/1+R)^\alpha + N(R/(N+NR))^\alpha} = \frac{1}{1+R^\alpha/N^{\alpha-1}} \geq \frac{1}{1+\delta} \geq 1-\delta.$$

Therefore,

$$\|\Pi_1 - \Phi_1\|_{\text{TV}} \geq |\Pi_1(A) - \Phi_1(A)| \geq 1 - 2\delta = 1 - \varepsilon. \quad \square$$

## B.4 Illustration of a Reasoning Tree

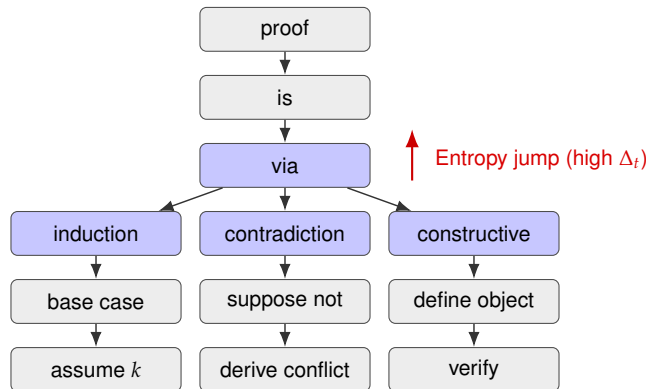


Figure 5: Illustration of a reasoning tree. Chain nodes (gray) are positions where the next token is effectively determined; branch nodes (blue) are positions where a semantic decision must be made. Entropy jumps occur at branch nodes. For example, the branch after the token “via” corresponds to the choice of proof strategy.

## B.5 Algorithm: Entropy-Cut Stagewise Metropolis–Hastings

In this section, we present the pseudocode for our algorithm.

Our algorithm is the special case of the meta-algorithm in [Algorithm 1](#) obtained by taking the cut distribution to be the entropy-cut distribution  $\lambda_\beta$  from [Definition 1](#). Thus, instead of treating every token position as equally likely to be revised, the sampler places more mass on positions where the model’s uncertainty has just increased, as measured by the positive entropy jump  $\Delta_t$ . In this way, cut points are biased toward likely decision points in the reasoning trace.

Unlike the uniform-cut case, the cut distribution now depends on the current state, so the cut-law ratio does not cancel from the Metropolis–Hastings acceptance probability. This correction is precisely what ensures that the resulting chain still targets the same power distribution  $\Pi_T$ . The resulting sampler is shown in [Algorithm 2](#). As before, we keep the proposal model  $p_{\text{prop}}$  explicit below; in our experiments, it is instantiated as a low-temperature version of the base model.

---

### Algorithm 2: Entropy-Cut Stagewise Metropolis–Hastings

---

**Input** : base model  $p$ ; proposal model  $p_{\text{prop}}$ ; power  $\alpha > 1$ ; cut power  $\beta \geq 0$ ; length  $T$   
**Hyperparameters**: block size  $B$ ; number of MCMC steps per stage  $N_{\text{MCMC}}$   
**Output** : continuation  $x_{0:T}$  approximately distributed according to  $\Pi_T$

- 1 Set  $K \leftarrow \lceil T/B \rceil$ ,  $T_0 \leftarrow 0$ , and  $x^{(0)}$  to the empty continuation
- 2 **for** each stage  $k \leftarrow 1, 2, \dots, K$  **do**
- 3     Set the  $k$ -th stage’s target length  $T_k \leftarrow \min \{kB, T\}$
- 4     Sample a suffix  $s \sim p_{\text{prop}}(\cdot \mid x^{(k-1)})$  of length  $T_k - T_{k-1}$
- 5     Set the current state  $x \leftarrow (x^{(k-1)}, s)$
- 6     *// Metropolis–Hastings chain targeting the power distribution for length- $T_k$*
- 7     **for** each MCMC step  $n \leftarrow 1, 2, \dots, N_{\text{MCMC}}$  **do**
- 8         Compute entropy profile  $h_t(x) \leftarrow -\sum_{v \in V} p(v \mid x_{<t}) \log p(v \mid x_{<t})$  for  $t = 0, 1, \dots, T_k$
- 9         Compute entropy jumps  $\Delta_t(x) \leftarrow \max\{0, h_t(x) - h_{t-1}(x)\}$  for  $t = 1, 2, \dots, T_k$
- 10         **if**  $\sum_{r=1}^{T_k} \Delta_r(x)^\beta = 0$  **then**
- 11             | Sample a cut point  $m$  uniformly from  $\{1, \dots, T_k\}$
- 12         **else**
- 13             | Sample a cut point  $m \sim \lambda_\beta(\cdot; x)$
- 14         Sample a (new) suffix  $s \sim p_{\text{prop}}(\cdot \mid x_{0:m-1})$  of length  $T_k - m$
- 15         Set the proposed state as  $x^{\text{prop}} \leftarrow (x_{0:m-1}, s)$
- 16         Compute the acceptance probability
 
$$A_m(x, x^{\text{prop}}) \leftarrow \min \left\{ 1, \frac{\Pi_{T_k}(x^{\text{prop}})}{\Pi_{T_k}(x)} \cdot \frac{\lambda_\beta(m; x^{\text{prop}})}{\lambda_\beta(m; x)} \cdot \frac{p_{\text{prop}}(x_{m:T_k} \mid x_{0:m-1}^{\text{prop}})}{p_{\text{prop}}(x_{m:T_k}^{\text{prop}} \mid x_{0:m-1})} \right\}$$
- 17         With probability  $A_m(x, x^{\text{prop}})$ , *accept* the proposal and set  $x \leftarrow x^{\text{prop}}$
- 18         Otherwise, *reject* the proposal and leave  $x$  unchanged
- 19     Set the  $k$ -th stage’s output as  $x^{(k)} \leftarrow x$
- 20 **return**  $x^{(K)}$

---

## B.6 Algorithm: Uniform-Cut Stagewise Metropolis–Hastings

In this section, we present the pseudocode for the sampler of Karan and Du [KD26].

The sampler of Karan and Du [KD26] is the special case of the meta-algorithm in [Algorithm 1](#) obtained by taking the cut distribution to be uniform:

$$\lambda_{\text{unif}}(m; x) \equiv \frac{1}{\ell} \quad \text{for each } m \in \{1, \dots, \ell\}.$$

Thus every token position is treated as equally likely to be revised. Under this choice, the cut-law ratio in the Metropolis–Hastings acceptance probability is identically one, so the general meta-algorithm simplifies to the uniform-cut sampler shown in [Algorithm 3](#). In the implementation of Karan and Du [KD26], the proposal model  $p_{\text{prop}}$  is taken to be a low-temperature version of the base model, but we keep  $p_{\text{prop}}$  explicit below.

---

### Algorithm 3: Uniform-Cut Stagewise Metropolis–Hastings

---

**Input** : base model  $p$ ; proposal model  $p_{\text{prop}}$ ; power  $\alpha > 1$ ; length  $T$   
**Hyperparameters**: block size  $B$ ; number of MCMC steps per stage  $N_{\text{MCMC}}$   
**Output** : continuation  $x_{0:T}$  approximately distributed according to  $\Pi_T$

- 1 Set  $K \leftarrow \lceil T/B \rceil$ ,  $T_0 \leftarrow 0$ , and  $x^{(0)}$  to the empty continuation
- 2 **for** each stage  $k \leftarrow 1, 2, \dots, K$  **do**
- 3     Set the  $k$ -th stage’s target length  $T_k \leftarrow \min \{kB, T\}$
- 4     Sample a suffix  $s \sim p_{\text{prop}}(\cdot \mid x^{(k-1)})$  of length  $T_k - T_{k-1}$
- 5     Set the current state  $x \leftarrow (x^{(k-1)}, s)$
- 6     *// Metropolis–Hastings chain targeting the power distribution for length- $T_k$*
- 7     **for** each MCMC step  $n \leftarrow 1, 2, \dots, N_{\text{MCMC}}$  **do**
- 8         Sample a cut point  $m$  uniformly from  $\{1, \dots, T_k\}$
- 9         Sample a (new) suffix  $s \sim p_{\text{prop}}(\cdot \mid x_{0:m-1})$  of length  $T_k - m$
- 10         Set the proposed state as  $x^{\text{prop}} \leftarrow (x_{0:m-1}, s)$
- 11         Compute the acceptance probability
 
$$A_m(x, x^{\text{prop}}) \leftarrow \min \left\{ 1, \frac{\Pi_{T_k}(x^{\text{prop}})}{\Pi_{T_k}(x)} \cdot \frac{p_{\text{prop}}(x_{m:T_k} \mid x_{0:m-1}^{\text{prop}})}{p_{\text{prop}}(x_{m:T_k}^{\text{prop}} \mid x_{0:m-1})} \right\}$$
- 12         With probability  $A_m(x, x^{\text{prop}})$ , *accept* the proposal and set  $x \leftarrow x^{\text{prop}}$   
           Otherwise, *reject* the proposal and leave  $x$  unchanged
- 13     Set the  $k$ -th stage’s output as  $x^{(k)} \leftarrow x$
- 14 **return**  $x^{(K)}$

---

## B.7 Proof of [Theorem 4.1](#)

In this section, we prove [Theorem 4.1](#).

*Proof of [Theorem 4.1](#).* We divide the proof into three steps. First, we reduce the chain to the space of branch choices and define the relevant conductance quantities. Second, we show that entropy-cut

has large conductance, while uniform-cut has a low-conductance bottleneck at the first branch. Third, we convert these conductance estimates into the claimed mixing-time bounds.

**Step A (Reduction to Branch Choices and Conductance Notation).** By Item 4 of [Condition 1](#), all subtrees below the children of any branch node are isomorphic. This implies that every root-to-leaf path has branch nodes at the same depths

$$1 \leq b_1 < b_2 < \dots < b_k \leq T,$$

and that the  $j$ -th branch has the same number of children, say  $d_j \geq 2$ , on every path. Indeed, before the first branch node there is only one possible path, so the first branch depth  $b_1$  is common to all leaves. At that branch, the child subtrees are isomorphic, so the pattern of future branch depths and the number of children at each future branch are identical in every child subtree. Repeating this argument inductively gives common branch depths  $b_1, \dots, b_k$  and common branching factors  $d_1, \dots, d_k$ .

Since every non-branch node has only one child, a leaf is uniquely determined by its sequence of branch choices. Thus, we identify the state space with

$$\Omega := [d_1] \times [d_2] \times \dots \times [d_k].$$

We write  $r = (r_1, \dots, r_k) \in \Omega$  for a branch-choice sequence, and we write  $\Pi_T(r)$  for the target mass of the corresponding leaf. Let  $x(r)$  denote the root-to-leaf token sequence corresponding to  $r$ . Since all tokens before the first branch are deterministic, the prefix  $x(r)_{0:b_1-1}$  is the same for every  $r$ ; denote this common prefix by  $x_{0:b_1-1}^\circ$ .

For a cut distribution  $\lambda$ , let  $P_\lambda$  denote the one-step MH transition kernel on  $\Omega$ . For any set  $A \subseteq \Omega$  with  $0 < \Pi_T(A) \leq 1/2$ , define its conductance under  $\lambda$  by

$$\Phi_\lambda(A) := \Pr_\lambda(X_1 \in A^c \mid X_0 \sim \Pi_T(\cdot \mid A)) = \frac{1}{\Pi_T(A)} \sum_{r \in A} \Pi_T(r) P_\lambda(r, A^c). \quad (3)$$

The global conductance is

$$\Phi_\lambda := \min_{0 < \Pi_T(A) \leq 1/2} \Phi_\lambda(A).$$

We will use the following proposal-imbalance quantity for the proposal obtained by cutting at the first branch and resampling the suffix from  $p_{\text{prop}}$ :

$$M_1 := \sup_{r \in \Omega} \frac{\Pi_T(r)}{p_{\text{prop}}(x(r)_{b_1:T} \mid x_{0:b_1-1}^\circ)}. \quad (4)$$

This definition uses  $p_{\text{prop}}$  directly, rather than introducing a separate full-suffix proposal distribution.

We next bound  $M_1$  using the first two items of [Condition 1](#). For any  $r \in \Omega$ , Item 1 gives  $\Pi_T(r) \leq e^l U(r)$ . Since the prefix before the first branch is deterministic,

$$U(r) = U(x(r)_{b_1:T} \mid x_{0:b_1-1}^\circ).$$

Applying Item 2 with the cut prefix  $x_{0:b_1-1}^\circ$  gives

$$U(x(r)_{b_1:T} \mid x_{0:b_1-1}^\circ) \leq e^\eta p_{\text{prop}}(x(r)_{b_1:T} \mid x_{0:b_1-1}^\circ).$$

Combining the above two yields  $\Pi_T(r) \leq e^{2\eta} p_{\text{prop}}(x(r)_{b_1:T} \mid x_{0:b_1-1}^\circ)$ , and therefore

$$M_1 \leq e^{2\eta}. \quad (5)$$

**Step B (Conductance Bounds).** We now prove the conductance estimates that drive the separation. The statement below isolates the two relevant facts: entropy-cut has uniformly large conductance because it can reopen the first semantic decision with probability  $1/k$ , whereas uniform-cut has a low-conductance bottleneck at the first branch because it reopens that decision only with probability  $b_1/T$ .

**Lemma B.4** (Conductance Bounds). *Under the assumptions of [Theorem 4.1](#), the entropy-cut chain satisfies  $\Phi_{\text{ec}}(A) \geq 1/2M_1k$  for every  $A \subseteq \Omega$  with  $0 < \Pi_T(A) \leq 1/2$ . In particular,*

$$\Phi_{\text{ec}} \geq \frac{1}{2M_1k} \geq \frac{e^{-2\eta}}{2k}.$$

Moreover, there exists a first-branch set  $S \subseteq \Omega$  with  $0 < \Pi_T(S) \leq 1/2$  such that  $\Phi_{\text{unif}}(S) \leq b_1/T$ . Consequently, for this same set  $S$ ,

$$\Phi_{\text{ec}}(S) \geq \frac{T}{2M_1kb_1} \Phi_{\text{unif}}(S) \geq \frac{e^{-2\eta}T}{2kb_1} \Phi_{\text{unif}}(S). \quad (6)$$

*Proof of [Lemma B.4](#).* We first show that entropy-cut chooses uniformly among the semantic branch positions. By Item 3 of [Condition 1](#), the positive entropy jumps occur exactly at branch nodes and have the same value  $\Delta$ . Thus, for every root-to-leaf path  $x$ ,

$$\Delta_t(x) = \begin{cases} \Delta, & t \in \{b_1, \dots, b_k\}, \\ 0, & \text{otherwise.} \end{cases}$$

Therefore, for any  $\beta > 0$ ,

$$\lambda_\beta(b_j; x) = \frac{\Delta^\beta}{\sum_{i=1}^k \Delta^\beta} = \frac{1}{k} \quad \text{for each } j \in \{1, \dots, k\}.$$

In particular, entropy-cut chooses the first branch depth  $b_1$  with probability  $1/k$ .

Now condition on the event that entropy-cut chooses  $m = b_1$ . Because all tokens before the first branch are deterministic in the branch-choice representation, this cut redraws the entire branch-choice sequence by sampling

$$x(r')_{b_1:T} \sim p_{\text{prop}}(\cdot \mid x_{0:b_1-1}^\circ).$$

Moreover, since  $\lambda_\beta(b_1; x) = 1/k$  for every state  $x$ , the cut-distribution correction in the MH ratio cancels. Thus, conditional on choosing  $m = b_1$ , the transition is the independence Metropolis–

Hastings kernel with target  $\Pi_T$  and proposal mass  $p_{\text{prop}}(x(r')_{b_1:T} \mid x_{0:b_1-1}^\circ)$ . We claim that this conditional kernel satisfies the minorization

$$K_{b_1}(r, \cdot) \geq \frac{1}{M_1} \Pi_T(\cdot) \quad \text{for every } r \in \Omega, \quad (7)$$

where  $K_{b_1}$  denotes the MH kernel conditional on cutting at  $b_1$ . To prove this, define

$$\omega(r) := \frac{\Pi_T(r)}{p_{\text{prop}}(x(r)_{b_1:T} \mid x_{0:b_1-1}^\circ)}.$$

By definition,  $\omega(r) \leq M_1$  for every  $r$ . For two states  $r, r' \in \Omega$ , the accepted proposal mass that  $K_{b_1}$  assigns to  $r'$  is

$$p_{\text{prop}}(x(r')_{b_1:T} \mid x_{0:b_1-1}^\circ) \min \left\{ 1, \frac{\Pi_T(r') p_{\text{prop}}(x(r)_{b_1:T} \mid x_{0:b_1-1}^\circ)}{\Pi_T(r) p_{\text{prop}}(x(r')_{b_1:T} \mid x_{0:b_1-1}^\circ)} \right\}.$$

Equivalently, this is

$$p_{\text{prop}}(x(r')_{b_1:T} \mid x_{0:b_1-1}^\circ) \min \left\{ 1, \frac{\omega(r')}{\omega(r)} \right\}.$$

If  $\omega(r') \leq \omega(r)$ , then the accepted proposal mass equals

$$p_{\text{prop}}(x(r')_{b_1:T} \mid x_{0:b_1-1}^\circ) \frac{\omega(r')}{\omega(r)} = \frac{\Pi_T(r')}{\omega(r)} \geq \frac{1}{M_1} \Pi_T(r').$$

If  $\omega(r') > \omega(r)$ , then the accepted proposal mass equals

$$p_{\text{prop}}(x(r')_{b_1:T} \mid x_{0:b_1-1}^\circ) = \frac{\Pi_T(r')}{\omega(r')} \geq \frac{1}{M_1} \Pi_T(r').$$

Thus the accepted proposal mass dominates  $M_1^{-1} \Pi_T(r')$  for every  $r'$ . Rejections only add mass to the current state, so (7) follows. Since the full entropy-cut kernel uses  $K_{b_1}$  with probability  $1/k$ , we obtain

$$P_{\text{ec}}(r, \cdot) \geq \frac{1}{M_1 k} \Pi_T(\cdot) \quad \text{for every } r \in \Omega. \quad (8)$$

Consequently, for every  $A \subseteq \Omega$  with  $0 < \Pi_T(A) \leq 1/2$ , and every  $r \in A$ ,

$$P_{\text{ec}}(r, A^c) \geq \frac{1}{M_1 k} \Pi_T(A^c) \geq \frac{1}{2M_1 k}.$$

Averaging over  $r \sim \Pi_T(\cdot \mid A)$  gives  $\Phi_{\text{ec}}(A) \geq \frac{1}{2M_1 k}$ . Since this holds for every such  $A$ ,  $\Phi_{\text{ec}} \geq \frac{1}{2M_1 k}$ . Using (5), we further get  $\Phi_{\text{ec}} \geq \frac{e^{-2\eta}}{2k}$ .

We next exhibit a low-conductance set for the uniform-cut chain. For  $i \in [d_1]$ , let  $S_i \subseteq \Omega$  be the set of paths whose first branch choice is  $i$ . These sets partition  $\Omega$ , so there exists  $i_* \in [d_1]$  such that  $\Pi_T(S_{i_*}) \leq \frac{1}{d_1} \leq \frac{1}{2}$ . Set  $S := S_{i_*}$ . Then  $0 < \Pi_T(S) \leq 1/2$ .

Under the uniform-cut chain, the first branch can change only if the cut occurs at or before

depth  $b_1$ . This happens with probability  $b_1/T$ . Any later cut preserves the first branch choice, and any rejection also leaves the state unchanged. Therefore, for every  $r \in S$ ,  $P_{\text{unif}}(r, S^c) \leq \frac{b_1}{T}$ . Averaging over  $r \sim \Pi_T(\cdot | S)$  gives  $\Phi_{\text{unif}}(S) \leq \frac{b_1}{T}$ . Finally, applying the entropy-cut lower bound to this same set  $S$ , and using the uniform-cut upper bound, gives

$$\Phi_{\text{ec}}(S) \geq \frac{1}{2M_1k} \geq \frac{T}{2M_1kb_1} \Phi_{\text{unif}}(S) \geq \frac{e^{-2\eta}T}{2kb_1} \Phi_{\text{unif}}(S).$$

This proves the lemma.  $\square$

**Step C (From Conductance to Mixing Time).** We now convert the above bounds into mixing-time estimates. For entropy-cut, we use the stronger minorization (8). Let  $\rho := \frac{1}{M_1k}$ . Then (8) implies that for every  $r \in \Omega$ ,  $P_{\text{ec}}(r, \cdot) \geq \rho \Pi_T(\cdot)$ . Equivalently, each step of the entropy-cut chain has probability at least  $\rho$  of refreshing exactly from stationarity. Hence, after  $n$  steps,

$$\sup_{r \in \Omega} d_{\text{TV}}(P_{\text{ec}}^n(r, \cdot), \Pi_T) \leq (1 - \rho)^n \leq \exp(-\rho n).$$

Therefore, if  $n \geq M_1k \log 1/\varepsilon$ , then the total-variation distance is at most  $\varepsilon$ . Using  $M_1 \leq e^{2\eta}$  from (5), we obtain

$$\tau_{\text{mix}}^{\text{ec}}(\varepsilon) \leq e^{2\eta}k \log 1/\varepsilon. \quad (9)$$

It remains to lower bound the mixing time of the uniform-cut chain. Start the uniform-cut chain from any state  $r^{(0)} \in S$ , where  $S$  is the first-branch set constructed in the proof of Lemma B.4. Until the first branch is reopened, the chain must remain in  $S$ . Under uniform cuts, the first branch is reopened at each step with probability  $b_1/T$ . Hence,

$$P_{\text{unif}}^n(r^{(0)}, S) \geq \left(1 - \frac{b_1}{T}\right)^n. \quad (10)$$

Total variation distance is at least the discrepancy on the event  $S$ , so

$$d_{\text{TV}}\left(P_{\text{unif}}^n(r^{(0)}, \cdot), \Pi_T\right) \geq P_{\text{unif}}^n(r^{(0)}, S) - \Pi_T(S) \geq \left(1 - \frac{b_1}{T}\right)^n - \frac{1}{2}. \quad (11)$$

Fix  $\varepsilon \in (0, 1/2)$  and define

$$c_\varepsilon := \frac{1}{4} \log \frac{1}{\varepsilon + 1/2} > 0.$$

Since  $b_1 \leq T/2$ , we have  $-\log\left(1 - \frac{b_1}{T}\right) \leq 2\frac{b_1}{T}$ . Therefore, for every  $n \leq c_\varepsilon \frac{T}{b_1}$ , we have

$$\left(1 - \frac{b_1}{T}\right)^n \geq \exp(-2c_\varepsilon) = \sqrt{\varepsilon + 1/2}.$$

Since  $\varepsilon < 1/2$ , we have  $\sqrt{\varepsilon + 1/2} - 1/2 > \varepsilon$ . Plugging this into (11) gives  $d_{\text{TV}}\left(P_{\text{unif}}^n(r^{(0)}, \cdot), \Pi_T\right) > \varepsilon$  for all  $n \leq c_\varepsilon T/b_1$ . Thus,  $\tau_{\text{mix}}^{\text{unif}}(\varepsilon) = \Omega_\varepsilon\left(\frac{T}{b_1}\right)$ . Together with (9), this proves the theorem, with the entropy-cut upper bound carrying the factor  $e^{2\eta}$  under the symmetry condition as stated.  $\square$

## C Additional Experiments and Details

### C.1 Benchmark Details

In this section, we describe the datasets we use.

- **MATH500:** The MATH dataset [LKBE+24] is a set of competition math problems spanning seven categories, including geometry, number theory, and precalculus. There are 12500 total problems, of which 5000 are test problems. MATH500 is a particular randomly chosen subset of the test set standardized by OpenAI.
- **HumanEval:** HumanEval [CTJY+21] consists of 164 handwritten programming problems covering algorithms, reasoning, mathematics, and language comprehension. Each problem comes with associated unit tests (7.7 on average) and a solution must pass all tests to be considered correct.
- **GPQA Diamond:** GPQA [RHSP+24] is a dataset of graduate-level multiple-choice science questions. GPQA Diamond is a subset of 198 questions of the highest quality.
- **AIME26:** The AIME 2026 [Art26] dataset comprises 30 high-school olympiad-level mathematical problems sourced from the 2026 American Invitational Mathematics Examination (AIME I and AIME II). The dataset spans algebra, geometry, number theory, and combinatorics, with each problem requiring multi-step logical deduction.

### C.2 Experimental Details

Here we detail the hyperparameters for the baselines.

**SMC Hyperparameters.** The most important hyperparameter for SMC is the number of particles  $N$ , which Azizi, Potraghloo, Ahmadi, Kundu, et al. [APAK+26] sets to be  $N = 64$ . All other hyperparameters, such as the generation length  $T$ , power exponent  $\alpha$ , etc., are not specific to SMC, and we take them to be the same as entropy-cut MH (Section 5.1).

**TMC Hyperparameters.** TMC has a few more hyperparameters, which we set following Ji, Tutunov, Zimmer, and Ammar [JTZA26]. Specifically, we take the block-length (look-ahead length) to be  $B = 192$ , the number of candidate sequences  $K = 8$ , and the number of look-ahead completions per candidate to be  $M = 8$ . Similar to SMC, all other hyperparameters are not specific to TMC, and we take them to be the same as entropy-cut MH (Section 5.1).

### C.3 Likelihood and Confidence

The sequence log-likelihood and average *confidence* [PCIF+25] of a sequence with respect to a model are respectively defined to be

$$\text{LL}(x_{0:T}) = \sum_{t=0}^T \log p(x_t | x_{<t}), \quad \text{Conf}(x_{0:T}) = \frac{1}{T+1} \sum_{t=0}^T \sum_{v \in V} p(v | x_{<t}) \log p(v | x_{<t}).$$

We plot the distribution of the sequence log-likelihood values computed via the base model (Qwen2.5-Math-7B) on the MATH500 dataset across 8 random seeds in Figure 6. Similarly, we plot the average confidence in Figure 7. In addition, the 25th, 50th (median), and 75th percentile

values are marked with dashed, solid, and dotted vertical lines. We observe that entropy-cut MH samples from higher-likelihood regions than standard, low-temperature, and uniform-cut MH, and from regions similar to those sampled by SMC and TMC.

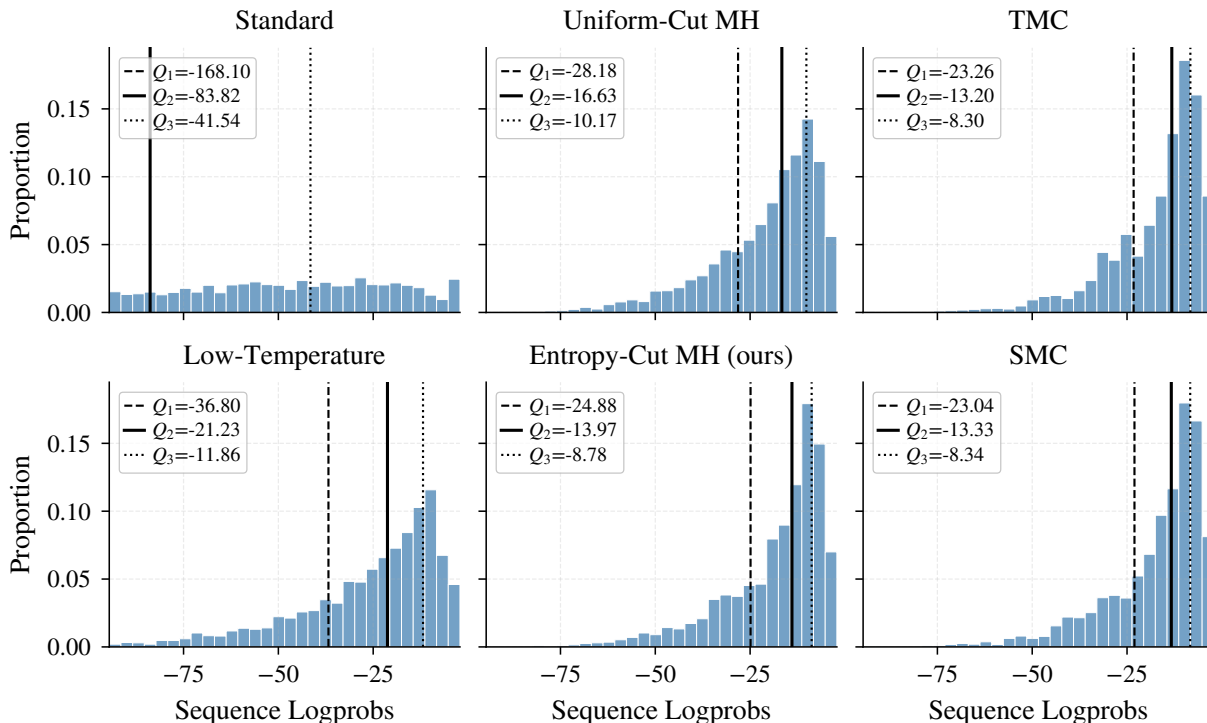


Figure 6: **Entropy-Cut MH (ours) samples from high probability regions.** We plot  $Q_{\text{wen2.5-Math-7B}}$  log-likelihoods for MATH500 responses across samplers.

#### C.4 Ablations

We demonstrate the stability of our algorithm against the specific choices of hyperparameters. Throughout these ablations, we fix the model to be  $Q_{\text{wen2.5-Math-7B}}$ , the MATH500 dataset, and all default hyperparameter settings described in Section 5.1, while varying the parameter in question. We report the average score across 8 repetitions.

**Power Distribution Exponent ( $\alpha$ ).** We plot the MATH500 scores across varying values of  $\alpha$  in Figure 8. At the extreme  $\alpha = 1$ , this is equivalent to the standard sampling algorithm. As noted by Karan and Du [KD26], a higher  $\alpha$  is correlated with better reasoning, but directly optimizing for likelihood is not necessarily optimal for sampling. Nonetheless, we see that the accuracy of entropy-cut MH remains relatively stable beyond  $\alpha \geq 2.0$ .

**Cut Law Exponent ( $\beta$ ).** Similarly, Figure 9 plots the MATH500 scores across various values of  $\beta$ . The accuracy seems similarly stable for values of  $\beta \geq 2.0$ .

**Metropolis-Hastings Steps ( $N_{\text{MCMC}}$ ).** In Figure 10, we plot the MATH500 scores across different numbers of  $N_{\text{MCMC}}$ . While increasing the number of steps increases the performance as expected, the increase past 10 steps seems to be relatively minor.

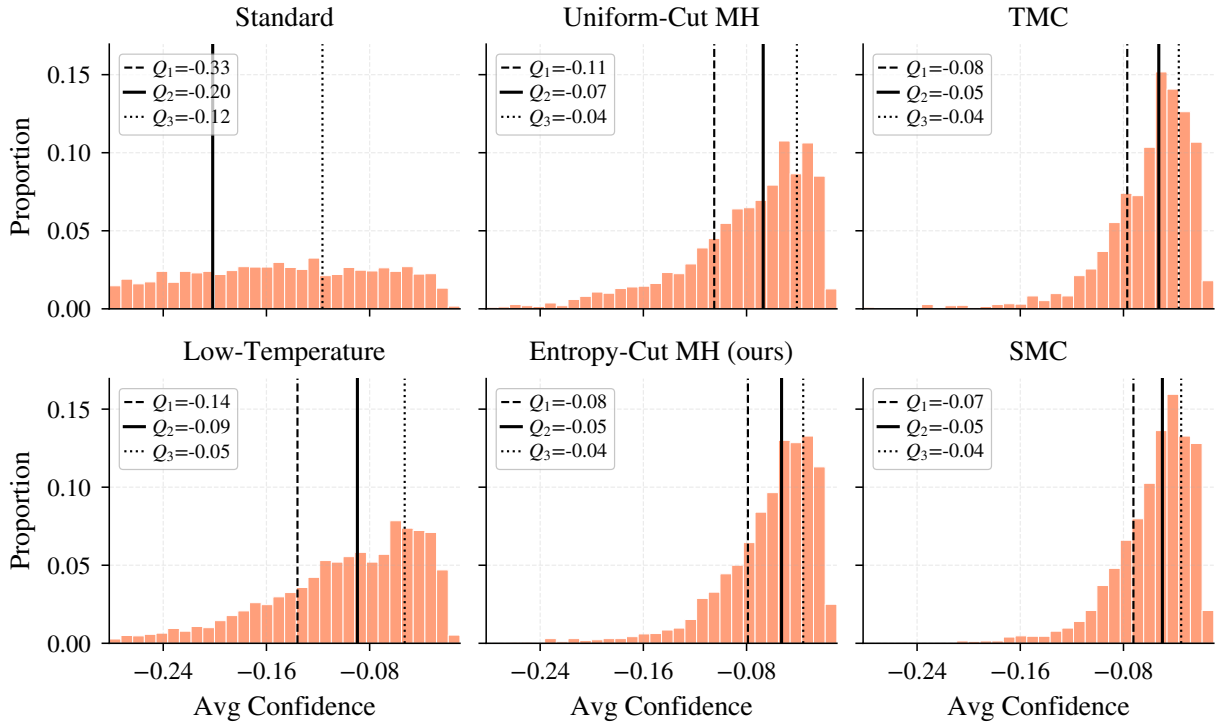


Figure 7: **Entropy-Cut MH (ours) samples from high confidence regions.** We plot Qwen2.5-Math-7B average confidence values for MATH500 responses across samplers.




Figure 8: **Entropy-Cut MH is stable to the choice of power exponent  $\alpha$ .** We plot Qwen2.5-Math-7B scores for MATH500 across various values of  $\alpha$  generated by entropy-cut MH.




Figure 9: **Entropy-Cut MH is stable to the choice of cut law exponent  $\beta$ .** We plot Qwen2.5-Math-7B scores for MATH500 across various values of  $\beta$  generated by entropy-cut MH.




Figure 10: **Entropy-Cut MH improves with additional transition steps  $N_{\text{MCMC}}$ .** We plot Qwen2.5-Math-7B scores for MATH500 across various values of  $N_{\text{MCMC}}$  generated by entropy-cut MH.

## C.5 Running Times

To demonstrate the scalability of our method, we report the average running time to generate the solution to a MATH500 question using `Qwen2.5-Math-7B` for the various (power) samplers in Figure 11. As expected, standard and low-temperature sampling are an order of magnitude faster; uniform-cut MH, entropy-cut MH, and SMC have comparable costs. TMC is the most expensive algorithm, requiring nearly double the cost of other samplers.




Figure 11: **Entropy-Cut MH has running times comparable to other power sampling methods.** We plot `Qwen2.5-Math-7B` average time to generate a solution per MATH500 question across various (power) sampling algorithms.

## C.6 Prompts

For each dataset except HumanEval, we have two sets of prompts. We use one for the Qwen family of base models (`Qwen2.5-7B`, `Qwen2.5-Math-7B`, `Qwen3-8B-Base`), which does not use a chat template, and another for the Phi family of instruction-tuned models (`Phi-3.5-mini-instruct`, `Phi-4-mini-instruct`), which uses the respective chat templates.

### C.6.1 MATH500

**Qwen Family.** We use the following prompt template where `{Question}` is replaced with the specific MATH500 question.

Can you solve the following math problem? Please reason step by step,  $\rightarrow$  and put your final answer within `\boxed{{}}`.

```
{Question}
```

```
Remember to present your final answer within \boxed{{}}!"
```

**Phi Family.** We format the following messages using the respective model chat templates where `{Qwen Prompt}` is replaced with the prompt above.

```
[  
  {"role": "system", "content": "You are an AI math expert."},  
  {"role": "user", "content": {Qwen Prompt}},  
]
```

## C.6.2 HumanEval

For HumanEval, we directly provide the model with the code stub and enforce the following stop words.

```
[
    "\nclass",
    "\ndef",
    "\n#",
    "\nif",
    "\nprint",
    "\nassert",
    "\nimport",
    "\nfrom",
    "\n`",
    "if __name__",
]
```

## C.6.3 GPQA Diamond

**Qwen Family.** For fairness, we randomly permute each GPQA Diamond multiple-choice answer among the choices and use the following prompt template. Here `{Question}` is replaced with the specific GPQA Diamond question and `{A}`, `{B}`, `{C}`, and `{D}` are replaced with the actual multiple-choice responses.

Answer the following multiple-choice question. The last line of your  
↪ response should be of the following format: `'\boxed{{$LETTER}}'` (  
↪ without quotes) where LETTER is one of ABCD. Think step by step  
↪ before answering.

`{Question}`

- A) `{A}`
- B) `{B}`
- C) `{C}`
- D) `{D}`

**Phi Family.** We format the following messages using the respective model chat templates where `{Qwen Prompt}` is replaced with the prompt above.

```
[
    {"role": "user", "content": {Qwen Prompt}},
]
```

## C.6.4 AIME26

For AIME26, we use the same prompt templates as MATH500 for both the Qwen and Phi families.